The Journal of

Economic Perspectives

A journal of the
American Economic Association

Summer 2025

The Journal of
Economic Perspectives

A journal of the American Economic Association

The Journal of
Economic Perspectives

Contents   *Volume 39 • Number 3 • Summer 2025*

Statement of Purpose

The *Journal of Economic Perspectives* aims to bridge the gap between the general interest business and financial press and standard academic journals of economics. The journal aims to publish articles that will serve several goals: to synthesize and integrate lessons learned from active lines of economic research; to provide economic analysis of public policy issues; to encourage cross-fertilization of ideas among the fields of economics; to offer readers an accessible source for state-of-the-art economic thinking; to suggest directions for future research; to provide insights and readings for classroom use; and to address issues relating to the economics profession. Articles appearing in the journal are normally solicited by the editors and associate editors. Proposals for topics and authors should be directed to the journal office, at the address inside the front cover.

Policy on Data Availability

It is the policy of the *Journal of Economic Perspectives* to publish papers only if the data used in the analysis are clearly and precisely documented and are readily available to any researcher for purposes of replication. Details of the computations sufficient to permit replication must be provided. The Editor should be notified at the time of submission if the data used in a paper are proprietary or if, for some other reason, the above requirements cannot be met.

Policy on Disclosure

Authors of articles appearing in the *Journal of Economic Perspectives* are expected to disclose any potential conflicts of interest that may arise from their consulting activities, financial interests, or other nonacademic activities.

How Regional Inequality and Migration Drive Housing Prices and Rents

Greg Howard and Jack Liebersohn

For several decades after World War II, regional incomes converged within the United States, but in recent decades, this trend has stopped and even reversed. Instead, more recent income growth—and the economic opportunities that come with it—has been increasingly concentrated in a small number of cities. From 2000 to 2018, median household incomes increased by about 15 percent more in cities that were already in the top 10 percent as compared to cities that were initially in the bottom half.

During this time, the house price index calculated by the Federal Housing Finance Agency (FHFA) rose by 31 percentage points more than the overall inflation rate as measured by the Consumer Price Index (CPI). This housing price increase is unevenly distributed and driven in large part by the higher rate of housing price increases in major cities. Relative to CPI, house prices rose by 49 percentage points in New York, 96 percentage points in Seattle, 129 percentage points in Los Angeles, and 113 percentage points in Miami (US Bureau of Labor Statistics 2024b). Our essay argues that this is no coincidence: the workhorse model of urban economics predicts that an increase in regional inequality leads to higher average national housing rents and prices, and the data support this theory.

To understand the link between regional inequality and average housing prices, we start with the relationship between regional inequality and cross-sectional differences in housing prices. Classic models in spatial economics predict that if a location experiences relatively higher productivity growth, then the associated rise in income

■ *Greg Howard is Associate Professor of Economics, University of Illinois, Urbana-Champaign, Illinois. Jack Liebersohn is Assistant Professor of Economics, University of California-Irvine, Irvine, California. Their email addresses are glhoward@illinois.edu and cjlieber@uci.edu.*

will increase demand to migrate there. If there is a mismatch between the places where people increasingly want to live and places with the space and zoning policies to absorb population, then housing costs can increase a lot in the in-demand cities, while decreasing only a little elsewhere. Since 2000, rising spatial inequality has increased relative demand to live in high-income cities. Because these cities face more geographic and regulatory obstacles to new housing, this has raised the prices and rents of housing in aggregate. In the rest of this paper, we discuss the empirical evidence behind the essential ingredients of our theory; we present the theory graphically in a two-region model; and we conclude by discussing implications of the theory.

Three Building Blocks

Three facts about regional inequality are crucial building blocks to our thesis in this paper. First, the impetus for our story is that high-income places have seen faster income growth from 2000 to 2018—that is, regional inequality has increased. Higher inequality has many possible causes, but a sufficient reason is that productivity growth has been higher in cities that had high incomes to begin with—in large part because they had an initially high concentration of industries that experienced growing productivity. Second, these high-income places have also seen higher growth in the cost of housing during the same time period. Together with the first fact, we interpret the second fact as evidence that people have increasingly wanted to move to the highest-income places because of the increased regional income inequality. Third, high-income places have less capacity to build new housing than lower-income regions because they both are denser and have more housing-related regulation. Taken together, these cross-sectional facts are building blocks that form a bridge between regional inequality and aggregate housing costs.

Regional Inequality Has Increased

US metropolitan areas that had initially higher incomes in 2000 also experienced faster income growth by 2018. In Figure 1, panel A, each point in the figure is a "core-based statistical area" (CBSA). A CBSA is defined by the Office of Management and Budget as including both "metropolitan" statistical areas with an urban cluster of at least 50,000 population and "micropolitan" areas with a cluster of at least 10,000. CBSAs are defined at the county level, with multiple counties included when there are sufficiently extensive commuting patterns between them. In 2000, there were 362 metropolitan statistical areas and 560 micropolitan statistical areas in the United States. Each datapoint in Figure 1 is a CBSA, and the size of each point represents the initial population. Each additional $10,000 of median household income is associated with an additional two percentage points of income growth over the 2000–2018 period. In the highest-income cities, there was a particularly large increase. The relationship is statistically significant; the figure notes include regression coefficients, standard errors, and *t*-statistics from a population-weighted linear regression.

Figure 1

The Growth in Income, House Prices, Rents, and Population 2000–2018 (Relative to Levels in 2000)

Panel A. Income growth

Panel B. House price growth

Panel C. Rent growth

Panel D. Population growth
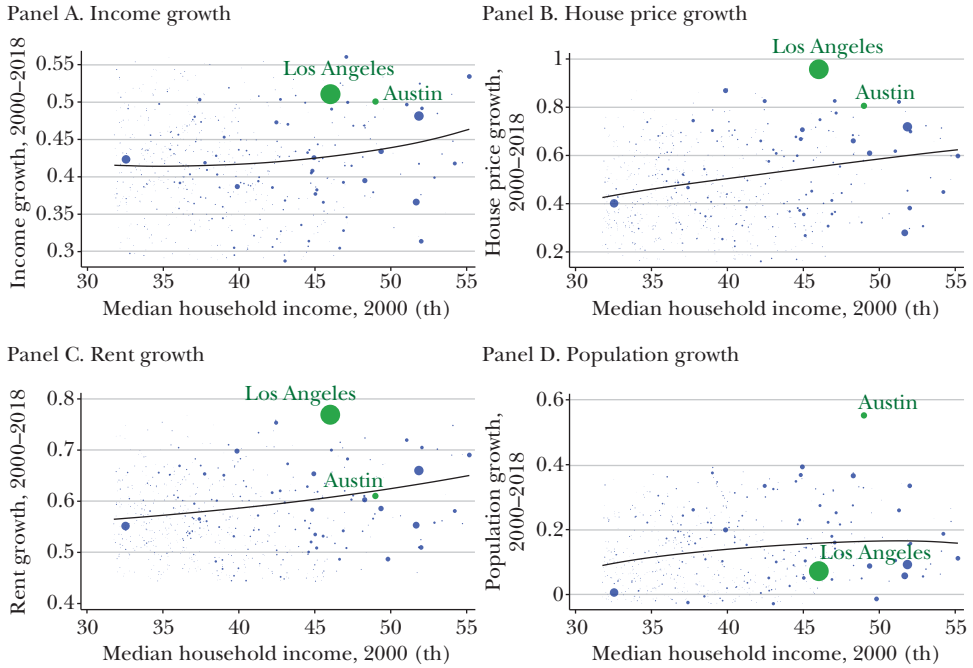
Source: Population data are from the 2000 decennial census and 2018 inter-censal estimates as compiled by NHGIS (Manson et al. 2025) and SEER (2025) using correspondence files from the Missouri Census Data Center (2025) and Fowler (2020). Population, rent, and income are provided via NHGIS (Manson et al. 2025). House price growth is from the Federal Housing Finance Agency (2025).
Note: The panels in this figure show the relationship between nominal 2000 median household income and nominal median income growth, house price growth, rent growth, and population growth. Each point is a core-based statistical area weighted by 2000 population. Growth is measured in log-changes from 2000 to 2018. We use the 2000 decennial census to measure median income and rent in 2000, and we use the 2016–2018 five-year ACS to measure median income and rent in 2018. All data are CBSA aggregates of county-level data, collecting non-CBSA counties into one unit. Figure omits top and bottom 5 percent of CBSAs. Fit line is from a local linear polynomial with smoothing parameter 10. Slope (robust SE; *t*) of population-weighted OLS fit lines: Household income = 0.002 (0.001; 2.51); House price growth = 0.008 (0.002; 3.99); Rent growth = 0.004 (0.001; 4.70); Population growth = 0.003 (0.001; 1.96).

One plausible hypothesis is that lower-income areas, with lower wages and lower costs of housing and land, are likely to attract more economic growth. However, agglomeration economies can ensure that industries remain concentrated in high-income metropolitan areas, even though their higher cost gives some firms the incentive to relocate (for a review of agglomeration economies, see Glaeser and Gottlieb 2009; in this journal, Duranton and Puga 2020). The higher-income urban areas have long specialized in certain industries, especially related to technology,

health care, finance, and highly-skilled services. These industries have grown quickly since the 1980s, causing faster growth in cities that were already high-income (Giannone 2017; Rubinton 2020; Eckert, Ganapati, and Walsh 2022). Much of regional income growth is caused by the relative productivity improvements across industries. For example, Barro and Sala-i-Martin (1991) argues that there is a steady 2 percent rate of convergence across US states, but also that deviations from this rate can be explained by differential industry growth rates. Such models also helped motivate the shift-share instruments used in Blanchard and Katz (1992) and Bartik (1991), and the many subsequent papers that have tried to predict regional growth.[1] In short, we do not need a new economic theory to understand the increase in regional inequality; it can be explained by the fact that the industries that did well in recent years were characterized by agglomeration economies and were already concentrated in high-income regions.

Of course, we are not the first to document this widening regional disparity. From 1940 to 1980, incomes converged, but since then the relative growth rates have reversed. Moretti (2012) describes the widening disparities as "the Great Divergence" and documents growing disparities in education, health, and quality of life, as well as income. Ganong and Shoag (2017) and Giannone (2017) focus on the reversal from previous convergence. Economists have hypothesized that the relevant agglomeration economies may be the result of market access, thick labor markets, or innovation spillovers, each of which is the subject of a large literature. Whatever the underlying cause, these forces make the high-productivity firms choose to stay in the high-productivity cities, despite facing higher wages.

Migration Links Income Changes and Housing Cost Changes

Migration in response to real income differences across cities creates a strong correlation between local incomes and housing costs. Economists have long emphasized a relationship between incomes and housing costs in the cross-section, starting with classical economists who recognized a link between ground rents and agricultural productivity.[2] The foundational spatial economics models of Rosen (1979) and Roback (1982) make the link between housing costs, wages, and migration explicit. Because people can move across space, differences in wages between cities cause net migration to high-wage locations. Consistent with these models, survey evidence agrees that many people move for jobs or because of housing

---

[1] We are not trying to claim that these factors are the only ones that determine regional growth. Besides changes in industry productivity, we would expect local economic policy (for example, Suárez Serrato and Zidar 2016), agglomeration spillovers (for example, Greenstone, Hornbeck, and Moretti 2010; Howard 2020; Duranton and Puga 2004), or changes in amenities (for example, Rappaport 2009) to also have effects on regional growth. Our only point is that differential industry productivity growth is sufficient to explain the increase in regional inequality.

[2] For example, Ricardo (1817) argued that land rents are determined by the value of agricultural output, which in turn is determined by the demand for the land's output, writing that "rent is always the difference between the produce obtained by the employment of two equal quantities of capital and labor." Mill (1848) extends the argument to nonagricultural land, whose rent is at least as high as its agricultural use, but with a premium for "beauty or convenience."

Figure 2
One-City Model
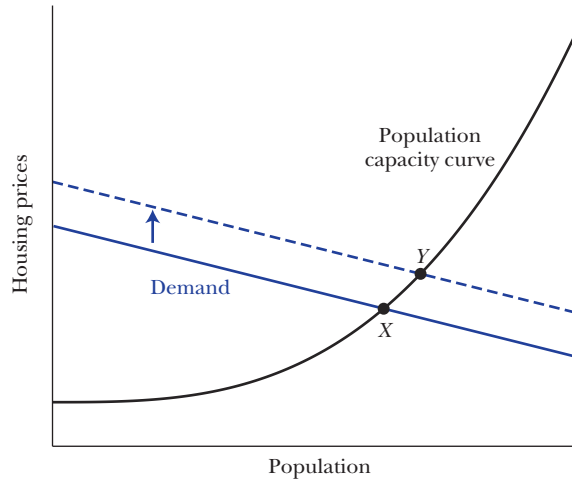
Source: Authors' illustration.

costs (Molloy, Smith, and Wozniak 2011). Modern urban models allow for heterogeneity across people, stronger attachments to particular cities, migration costs, and a variety of other refinements, but still retain the essential concept (for example, Redding and Rossi-Hansberg 2017; Monte, Redding, and Rossi-Hansberg 2018).

In Figure 1, panels B and C, we show how house price growth and rent growth vary with the initial income of core-based statistical areas in 2000. The same high-income cities where median incomes increased the most also had the greatest increase in house prices and rents, consistent with evidence from Van Nieuwerburgh and Weill (2010). As with incomes, the rise in housing costs is most prominent on the right-tail of the figure, which contains high-income cities like San Francisco, California; San Jose, California; and Stamford, Connecticut. Figure 1, panel B, shows that the relationship between incomes and house price growth is, if anything, even stronger than the relationship between incomes and rents. In Howard and Liebersohn (2023), we argue that this is because trends in regional inequality are persistent, so rent growth today predicts further increases in the future. Because house prices are partially determined by future rents, this can lead to a larger increase in house prices than rents.

A likely reason that housing costs reflect incomes is that people will migrate in response to changes in wages. We can think of this using a supply and demand relationship, as shown in Figure 2. The demand curve traces out the number of people who would like to live in a city for a given cost of housing. In this case, the relevant supply curve might be called a "population capacity curve"; that is, we are thinking more broadly of the number of people the city will accommodate at a given price. This measure incorporates two incentives as prices rise: more housing will be built,

and people will consume less housing per capita, whether through smaller units or more occupants per unit. Because it is hard to adjust the size of housing units, one can think of this curve as primarily determined by the housing supply curve. When incomes rise, the demand curve for living in a city increases, raising both the population and the housing cost.[3] The degree to which housing prices will rise depends on how much the demand curve moves in response to income and the steepness of the population capacity curve.

For how much the demand curve moves, we can use the fact that people should be (to first-order) indifferent about living in a particular city if their incomes and their housing costs increase by the same dollar amount. Thus, a rise in a city's income will shift the housing demand curve up by the same dollar amount of the income increase.[4] If the population capacity curve is vertical or the demand curve is horizontal, then all the change in incomes will be reflected in housing costs. If the capacity curve is horizontal and the demand curve is not, then there will be no change at all in housing costs. Here, we consider the more plausible case of an upward-sloping supply curve and downward-sloping demand curve.

In the next section, we discuss heterogeneity in the steepness of the population capacity curve across urban areas with different per capita income level. Here, we focus on the slope of the demand curve. Many economists have tried to estimate the elasticity of this demand curve with many different methods and have gotten a wide range of answers (for a list of elasticities used in the literature, see Fajgelbaum et al. 2019, Table A.17; Howard and Liebersohn 2021, Table 1).[5] A typical estimate is that a 1 percent increase in wages would cause a 2 percent change or more in the population. This is flatter than the best recent estimates of housing supply—for example, Baum-Snow and Han (2024) estimate average housing supply elasticities around 0.3–0.5. (In the short-run, it is hard to redivide housing, making it hard to change housing quantities per capita; therefore, we expect the elasticity of housing demand to be even smaller.) Thus, when we draw Figure 2, we would expect the demand curve to be relatively flat compared to the population capacity curve, and when the demand curve shifts up, we would expect to see an increase in housing costs. Taking this idea to the data, we expect that income differences are reflected in housing price differences across cities. Indeed, Figure 1, panel B, suggests that the same places that experience income gains were also experiencing house price gains.

---

[3] We ignore the income effect on housing demand: that is, the fact that people will want more housing or fewer roommates in response to increased income. Given that we are expressing our quantities as population, such a force would actually shift the housing supply curve to the left, meaning we underestimate the change in housing costs. However, because our focus is on relative changes in income across space, not absolute changes, we wish to primarily focus on the increase in demand to live in the city.

[4] To translate this to elasticities, the elasticity with respect to housing costs is equal to the elasticity with respect to incomes times the share of housing in consumption.

[5] To translate that into an elasticity of population with respect to rents, we first need to calculate the size of a 1 percent increase in rents, as a percent of wages. For that, we need to multiply by the housing share of consumption, which is somewhere near one-third. Taking these together, the elasticity of population with respect to rents should be about one-third as large as the elasticity of population with respect to income.

Some economists are skeptical of a relatively flat demand curve for housing in a certain area because of widely reported findings that Americans have become less mobile in recent decades. There is some controversy over the extent to which the internal migration rate has fallen during this time. The fall in migration is primarily found in the Current Population Survey (as discussed in Molloy, Smith, and Wozniak 2017; Kaplan and Schulhofer-Wohl 2017), while the decline appears much smaller in other surveys, such as the American Community Survey or in administrative data from the Internal Revenue Service (Foster, Ellis, and Fiorio 2023).[6] Here, we emphasize that we are primarily focused on long-run migration rates, which are still fairly high in the United States. More than 30 percent of Americans live in a different state than the one in which they were born (Molloy, Smith, and Wozniak 2011). More importantly, there is significant heterogeneity in the growth rate of US urban areas. At the ninetieth percentile, urban areas grew by 42 percent from 2000 to 2018, while cities at the tenth percentile neither gained nor lost population. These facts emphasize that in the long-run, there is a lot of migration, and so we might expect the demand curve for housing to be fairly flat.

So far, our evidence has relied upon simple and transparent measures of housing costs and incomes. The implicit assumption behind these measures is that someone on the margin of moving would be able to get a higher income in a high-income city than in a low-income city. But differences in average income reflect not only differences in wages being offered, but also differences in the types of workers who live in higher-income cities as compared to workers elsewhere. For example, they may be more educated or more experienced and receive higher wages as a result. The incentive to migrate to high-income places depends on counterfactual incomes, so the ideal measure of incomes would adjust for differences in average ability, education, job quality, taxes/transfers, and so on.

How would our conclusions change if we adjusted for measures of worker quality? Not much. To approximate this type of measure for incomes, Card, Rothstein, and Yi (2025) use a "mover design"—comparing movers and nonmovers while adjusting for differences between these groups—to hold worker characteristics constant. They find that differences in income across space are compensated by differences in housing costs. Diamond and Moretti (2021) show that the cost of living also varies across space, with high-income places also being the most expensive places to live. Both papers show a strong cross-sectional relationship between incomes and the cost of living, confirming the basic assumption of our model. The exact extent to which wages are imputed into housing costs depends on a variety of features of the local economy; for example, the extent of pass-through of rents to other prices can temper the house-price response (Couture et al. 2024).

---

[6] A related concern might be that if migration costs are as high as they are sometimes estimated—for example, Kennan and Walker (2011) estimate costs of over $300,000 per interstate move—that might limit the degree to which housing costs would move in response to income changes. Howard (2024) offers a critique of interpreting this number as a literal moving cost. In particular, these moving cost estimates depend on the modeler's choice of timing; over long time periods, the estimated moving costs will be mechanically lower.

A limitation of sophisticated data based on movers, of the kind these papers use, is that it is not available very far back historically, so we cannot use the data to study changes in economic opportunity over long time periods. Making use of fuller historical data, we adjust for worker characteristics using linear regressions in Howard and Liebersohn (2021). This adjustment does not change the basic fact that regional inequality has increased. Even further back in time, data from the mid- and early-twentieth century are less likely to have worker characteristics, and researchers studying that time period often just use average local wages.

An ideal measure of housing costs would adjust for differences in the type of housing that is available and would measure housing costs on a per-period basis, similar to incomes. Rents are often a better measure of housing costs than the price of housing, because house prices are related to both current rents and expectations of future rents.

Greater availability of rent data in recent years, much of them from private sources such as Zillow, has made it increasingly feasible to study rents directly rather than house prices.

Land prices depend less on housing quality than house prices do, so one potential way of controlling for housing quality is to look at land prices. A disadvantage is that it is difficult to know how much of the cost of housing is due to land alone. However, researchers have used different methods and time periods to back out land prices using data on structure costs or transactions involving vacant lots (for example, Ahlfeldt and McMillen 2018; Davis and Palumbo 2008; Albouy, Ehrlich, and Shin 2018; Davis et al. 2021). All find that land values are highest in the highest-income urban areas. Also, Davis and Palumbo (2008) and Davis et al. (2021) find that land prices have increased the most in expensive coastal metros, similar to rents and house prices. In other words, looking at land prices instead of rents or house prices does not change our basic story.

Spatial equilibrium models also predict that housing costs depend on differences in urban amenities, which again are hard to measure. Some natural amenities may be fixed or slow to change and can be differenced-out by studying changes across time. This assumption may be reasonable over short time periods, but changes in the urban environment or climate are likely to matter more over long time periods. Using an index of many amenities, Diamond (2016) shows that some amenities change endogenously in response to population changes. In particular, this theory emphasizes the role of college-educated people who are more mobile and also have an equilibrium effect on the local amenities. In another approach, Almagro and Domínguez-Iino (2024) allow for preference heterogeneity over multiple amenities instead of using a single quality index. Amenities improve the most in areas with the greatest growth, in this way amplifying the effects of income shocks.

Summing up, we think that migration towards cities with higher real incomes is a key driver of the relationship between income growth and housing cost growth. Looking at the data, Figure 1, panel D, shows how population growth relates to median incomes. All else equal, one would expect population growth to be the greatest in the highest-income core-based statistical areas, because economic and

income growth has been the highest there. But instead, the relationship between population growth and median incomes is inverse-U shaped, with the highest population growth in the middle of the income distribution rather than the top. Average population growth in locations at the top quintile of income was about 15 percent, equal to population growth in the second-lowest quintile. Population growth in the middle quintile of the income distribution was higher than both, at 18 percent. Why did people move more to middle-income places instead of the very highest-income places? This brings us to our final building block.

Higher-Income Regions Have Not Produced Additional Housing

The urban areas with the highest per capita income levels also tend to have the greatest constraints on new housing construction, which means that changes in housing demand have larger effects on equilibrium housing costs but smaller effects on equilibrium population. The key quantitative measure is the elasticity of housing supply with respect to prices—that is, by what percent does the quantity of housing increase for every 1 percent increase in house prices? In areas where the housing supply is elastic, small price increases lead to a large percent increase in housing construction, whereas in areas where it is inelastic, even large price increases do not lead to new homes being built.

Why is the housing supply more elastic in some areas than others? Two factors seem particularly important. First, it is harder to build in areas with more people because higher density makes it harder to assemble plots and fit new structures (Baum-Snow 2023). Given that people migrate towards higher-income areas, it is not surprising that higher-income areas are denser on average. Figure 3, panel A, shows how median incomes relate to population density: here, each data point represents a Metropolitan Statistical Area—that is, the larger urban areas that were part of the core-based statistical areas defined earlier. For every $10,000 higher median income in the year 2000, the population density was 8 percent higher. Again, we include coefficients, standard errors, and *t*-statistics from corresponding population-weighted linear regressions in the figure notes.

Not only are higher-income places denser, they also regulate new construction more heavily. Glaeser and Gyourko (2018) show how stringent land-use restriction in high-income places leads to spatial misallocation (see also Hsieh and Moretti 2019); indeed, such regulations tend to redistribute wealth from lower-income house buyers to higher-income house sellers. Figure 3, panel B, shows the relationship between median incomes and housing supply regulation. To measure regulation, we use data from the Wharton Residential Land Use Regulatory Index, a summary index of land use regulation based on a survey of local land use authorities conducted by Gyourko, Saiz, and Summers (2008a). The relationship between incomes and housing restrictions is positive and strong. Every $10,000 increase in income is associated with about a 0.4-point higher value of the index, a bit less than half a standard deviation.

The strong relationship is notable in part because, unlike population density, land use regulation is under the direct control of local governments. One might think that local governments would respond to house price increases by loosening

Figure 3
Supply Constraints and Median Incomes

Panel A. Population density

Panel B. Wharton residential land use restriction index

Note: The panels in this figure show the relationship between nominal 2000 median household income and land use regulation and log population per square mile. Each point is a Metropolitan Statistical Area weighted by 2000 population. Figure omits top and bottom 5 percent of core-based statistical areas. Fit line is from a local linear polynomial with smoothing parameter 10. Slope (robust SE; *t*) of population-weighted ordinary least squares fit lines: log pop./mi^2 = 0.080 (0.027; 2.99); Wharton Reg. Index = 0.043 (0.012; 3.55).

regulation. But the opposite is true: in an update to the Wharton index, Gyourko, Hartley, and Krimmel (2021) show that land use regulation has increased over time, particularly in large coastal markets. Unless there is a dramatic change of course in land use regulation in high-income areas, we expect the correlation with income to continue.

Why do the high-income urban areas make it harder to build new housing? One reason has to do with fears that new housing will erode home values, which is described as the "Home-Voter Hypothesis" by Fischel (2005). These fears are exacerbated by the fact that moving is costly and that homes have a high value, representing a source of wealth that it is difficult to diversify. A related argument is that zoning can also be a way for homeowners to raise the tax base (Epple, Romer, and Filimon 1988) or to manage congestion costs in high-density areas (Hilber and Robert-Nicoud 2013). A robust literature discusses the political origins of zoning restrictions, but without firm conclusions as to the reasons (Gyourko and Molloy 2015). The link between housing costs and zoning restrictions holds both across and within cities, and has been found in papers using a variety of methods and time periods (Glaeser and Ward 2009; Glaeser, Gyourko, and Saks 2005). In a study of the aggregate implications of the home-voter hypothesis in a spatial equilibrium model, Parkhomenko (2023) finds that these regulations reduce aggregate productivity, because they keep people away from the highest-productivity locations.

Combining data on regulation, population, and other features of the local environment, several recent papers have estimated measures of the housing supply elasticity. The most widely used measure of housing supply elasticity is the one in Saiz (2010a), which estimates an elasticity for each Metropolitan Statistical Area. Baum-Snow and Han (2024) make a similar calculation, but at the neighborhood level and for a later and shorter time period, and estimate a much smaller average elasticity. But the same comparative static still holds in both papers, with the housing supply more inelastic in higher-income areas.[7]

If housing supply becomes more inelastic as incomes and housing costs rise, that should have implications for changes in population. Going back to the supply and demand graph in Figure 2, we have drawn the population capacity curve to be convex, reflecting that as income rises and demand increases, it will be harder to build new housing. Thus, for a high-income city, an increase in demand will lead to a big increase in prices, but a smaller increase in population. For a city that starts out low income, an increase in demand will have only a small change in prices, but a large change in population.

We see evidence consistent with such elasticities in our data on populations and housing costs. In the previous section, we documented that the highest-income parts of the country saw the largest housing cost increases. Without differences in housing supply elasticities, that would mean that we would expect those places to grow the most too. Yet in Figure 1, panel D, the highest-income places did not expand more than the middle-income places. This is consistent with lower housing supply elasticities in the highest-income places, where it is relatively hard to build new housing due to density and regulation.

---

[7] There is some debate about the strength of this heterogeneity. Louie, Mondragon, and Wieland (2025) find that existing measures of housing supply elasticity are not an important mediator of the relationship between income growth and price growth. In response, Furth (2025) argues that their measure of income growth is incorrect.

Of course, one way to put more people into a city without changing housing supply is for people to consume less housing—for example, by having roommates or dividing up existing houses. However, we hypothesize that this margin of adjustment is also more elastic in less-expensive areas. When housing is cheap and people are consuming a lot of it, they will adjust consumption in response to price changes, but when housing is expensive and people are closer to consuming what they view as the bare necessities, then they are unwilling to continue to shrink their housing consumption even if it continues to get more expensive. This is consistent with, although not necessarily implied by, a well-document fact about housing demand: people's consumption rises with their income, but less than one-for-one (Mayo 1981; Quigley and Raphael 2004; Albouy, Ehrlich, and Liu 2016; Cardullo and Sechi 2023). This heterogeneity works in the same way as the heterogeneity in housing supply: people will consume less housing in richer cities than they would in poorer cities, where housing is relatively cheap compared to other goods. Thus, as house prices change, the amount of adjustment they are willing or able to do is smaller.

Finally, it is worth mentioning again that at least some of the relationship between housing supply elasticity and income is a policy choice and therefore could be changed. Several cities with a high level of income and robust income growth sustained large population increases while avoiding substantial increases in housing costs. Many of these cities are located in Texas and the southwestern United States, where there are few natural barriers to construction and new housing is relatively uninhibited by regulation. To illustrate, compare Austin, Texas, and Los Angeles, California. In 2000, our data show that median incomes were similar in both cities—around $49,000 in Austin and $46,000 in Los Angeles. Real income growth from 2000 to 2018 was about 15 percent in Austin and 16 percent in Los Angeles. However, population growth was much greater in Austin, where the total population increased by 74 percent, than in Los Angeles, where it only increased by 7.5 percent. Real house prices increased by 74 percent in Austin—a substantial increase—but much less than the 110 percent increase in Los Angeles. While Austin is an extreme example, similar patterns of robust income growth, large population increases, and relatively moderate housing costs occurred in Houston, Dallas, and Las Vegas.

Aggregate Implications of Regional Inequality

The cross-sectional data themselves are not sufficient to tell us how regional inequality affects national housing prices. We might be tempted to take the model from Figure 2 and see if we can apply it to the country as a whole. But when we think about the population capacity curve of housing at the national level, it ought to be much more elastic than the curve for individual cities. In particular, if we wanted to induce housing construction that occurs anywhere in the United States, it might not require a very large rise in house prices at all. Thus, we need a model that incorporates differences across space to put all the facts together and draw conclusions about the aggregate.

Here, we begin by considering a model that builds on Figure 2 by having multiple places that interact in equilibrium. For graphical simplicity, we show City A and City B. The top half of Figure 4 shows population capacity curves—what share of the total population a city can hold for a given price of housing. The key additional assumption in this figure is that total population is fixed and that people choose between living in City A and City B. Thus, housing costs in both locations are determined by one variable: the share that choose to live in City B. As the population in B increases to the right, housing costs in B rise. At the same time, the population in A necessarily decreases, so housing costs in A fall. We also plot the average housing cost, which will come in handy later. We plot the unweighted average in Figure 4 because we want to focus on an increase in average housing costs that is the result of housing cost changes within cities. When population becomes more unequal, the housing cost for the average person will rise by even more because more people will be choosing to live in the more expensive city. However, this increase does not reflect a change in the person's budget set in the same way that a change in the unweighted average does.

The population capacity curve becomes less elastic as the population increases: as a result, the average will be higher when the population is unevenly distributed. As in the one-city model in Figure 2, we will assume any shifts in the capacity curves due to changes in regional inequality to be quantitatively negligible.

Because there are two places to live, people decide where to live based on not just the housing cost of one city, as before, but rather the difference in housing costs between the two cities, which is shown in the bottom panel of Figure 4. Thus, the relative demand curve is a function of how many people will choose to live in City B given the difference in housing costs between the two cities. The slope of this curve is related to the slope of the demand curve in the one-city model, but the one-city model did not consider any changes in the utility of the outside option, while here we are explicit that the only alternative is to live in the other city. The average level of the demand curve depends on people's common preference for City A or City B. For example, if amenities in City B improve, the relative demand curve will move up and City B will become relatively more expensive. The slope comes from the heterogeneity of people's preferences for each city and how important that heterogeneity is compared to housing prices. If house prices go up in City B but not City A, people on net will move from City B to City A. If many people are close to indifferent across the cities, that will induce many people to move, and the relative demand curve will be fairly flat. If most people have a strong attachment to one city or the other, then not many people will be induced to move, and the relative demand curve will be steep. The spillover at the center of this assumption—people moving in response to house price changes causing house price changes in the place they move to—has been well-documented in Schubert (2021). The bottom panel plots the relative population capacity curve, which is the difference between the two curves in the top panel.

The equilibrium of this model will be where the relative demand curve meets the relative population capacity curve, at point "*X*" in the bottom panel. This point

Figure 4
Two-City Model

Housing prices
in B

Housing prices
in A

A's population
capacity curve

B's population
capacity curve

P_Y

P_X

Average
price

Y   *X*

Share of population in B
(= 1 – share of population in A)

Relative
housing prices
(B – A)

Relative population
capacity (B – A)

0

X

Relative demand

Y

Share of population in B
(= 1 – share of population in A)
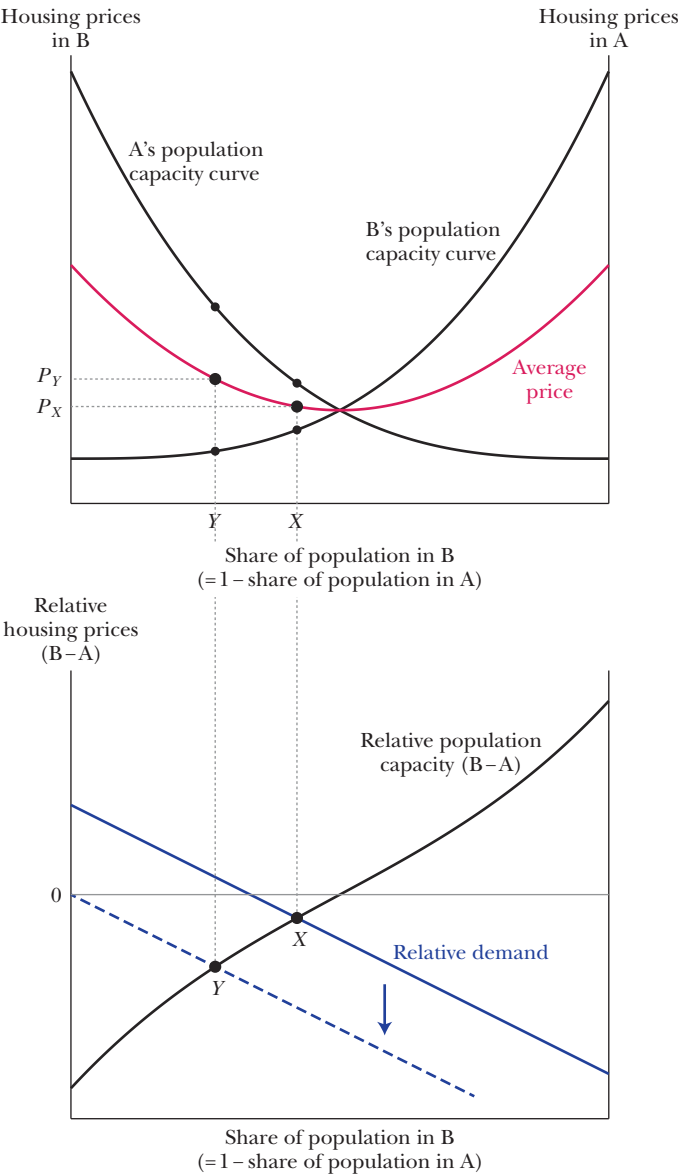
tells us the populations in both cities, as well as the relative housing costs of the two places. Of course, we are primarily interested not in the relative housing costs, but in the individual housing costs and the average housing cost. For those we can trace the equilibrium up to the top panel, where we can see the average housing costs as

a function of the population distribution. The average price corresponding to equilibrium X is P_X in the top panel of Figure 4. The prices in each city can be found by tracing up from the X to the two population capacity curves (to the smaller dots) and then over to the y-axis. As we would expect given the relative housing costs in the lower figure, housing prices in A are higher than housing prices in B.

Now consider a shift in the relative demand curve, representing an increase in regional inequality. Initially, City A was slightly preferred to City B, which we will assume was due to City A having a slightly higher income. Now imagine that City A's income rises compared to City B. That will lower the relative demand curve, so that fewer people wish to live in City B at any given relative price for housing. In the bottom panel, we can see that equilibrium shifts to the left and down to point Y.

The effect of rising relative income in City A for the overall housing costs can be seen by tracing the new equilibrium Y up to the top panel. In City B, housing costs have fallen slightly due to their having a relatively elastic housing supply. In City A, housing costs have risen steeply due to a less elastic housing supply, as they were already the place with more population. Overall, the average house price between the two cities rises because the increase in City A is larger than the decrease in City B. Note that the dot above Y on A's Population Capacity Curve is much higher than the corresponding dot above X, while the two dots on B's Population Capacity Curve are of more similar heights. In addition, house prices for the average person rise by even more because people are moving to the more expensive city. The average housing cost has increased to P_Y.

We can see the importance of our third building block (that higher-income cities have less elastic population capacity) by imagining that instead of the population capacity curves being concave, they were linear and symmetric. In this case, the average housing supply is no longer a function of population—and a change in regional inequality would have no effect on average housing costs nationally.

Generalization of a Model to Many Cities

The model in Figure 4 has only two regions, but the world is a big place, and the regional inequality that we have documented is not universal: there were plenty of lower-income cities that outperformed and plenty of higher-income cities that underperformed. Should we expect our result to hold generally? Is a positive correlation between income growth and initial incomes enough to expect housing costs to rise?

Yes. The key assumption is the equilibrium assumption that everyone must live somewhere. If one person moves from a lower-income city with more housing supply elasticity to a higher-income city with less housing supply elasticity, they raise house prices in the less elastic city more than they lower house prices in the old region. Their net effect is to raise house prices on average. So when regional inequality increases and people—on average—move to higher-income places, average prices will rise. Indeed, in Howard and Liebersohn (2021), we show that the covariance of relative demand changes and housing supply elasticities is the important statistic for the increase in housing costs.

Regional Inequality and Housing Costs in the Data

To this point, we have established a rationale for why increased regional inequality would lead to higher housing costs through migration. Moreover, we have argued that the data are consistent with many of the necessary assumptions for our theory to work. But whether our theory explains a lot of the increase in housing costs is an inherently empirical question. To answer it, we look at the time series of regional inequality and rents.

Figure 5 shows several measures of rents, plotted in the left-hand axis. The top line uses the measure of rents in the Consumer Price Index divided by the overall Consumer Price Index, indexed so that 1969 is equal to 1 (US Bureau of Labor Statistics 2024a). While this index focuses on urban areas, it is the only rent index that is even close to nationally representative and has data for a long period of time. However, one of the key criticisms of this rent index is that it reflects existing rents, which may not be the current price being charged in the market for starting a rental. As a result, it can lag market-priced rents by several years, so we should focus on medium- and long-term trends rather than year-to-year variation.

Crone, Nakamura, and Voith (2010a) argue that prior to 1985, the CPI rent index is mismeasured because it omitted units that had a change of tenants, thus leading to nonresponse bias. We include their rent measure as well, which includes a much smaller decline over the 1970s. We also show the real rent index from Lyons et al. (2024a). It is a hedonic rent index based on newspaper real estate listings in 30 major cities. While it is less nationally representative, the timing is likely to align more closely to market rents, and it is unlikely to suffer from the bias that Crone, Nakamura, and Voith (2010a) discuss.

To measure regional inequality, we use the standard deviation of the log per capita personal income, available through the Bureau of Economic Analysis (BEA) Regional Accounts (US Bureau of Economic Analysis 2023), as measured on the right-hand axis. The BEA has published data at the county level yearly since 1969. We aggregate counties to core-based statistical areas, following our earlier data analysis. We also take a three-year moving average so that the standard deviation is less noisy. Of course, this measure is not ideal, as per capita personal income will include income that is likely not available to someone on the margin of moving between cities, for example business income. Nonetheless, better measures do not have the same geographic and temporal coverage.

We would not use the time series shown in Figure 5 to make causal claims about regional inequality and rents, but there does appear to be a correlation in the values. In addition to the general upward trend in both rents and regional inequality since 1980, many of the fluctuations show up in both series with similar timing. In both series, there is a decline in the first few years, followed by a rise in the 1980s, another fall into the mid-1990s, and then a large increase since then.

Although our theory predicts many of the long-run changes in the price of housing, it does not get all of the smaller fluctuations right. The decline in regional inequality before 1980 and the subsequent rise in inequality are both reflected in falling and then rising rents—although the fall before 1980 is not present in the

Figure 5
Regional Inequality and Real Rents
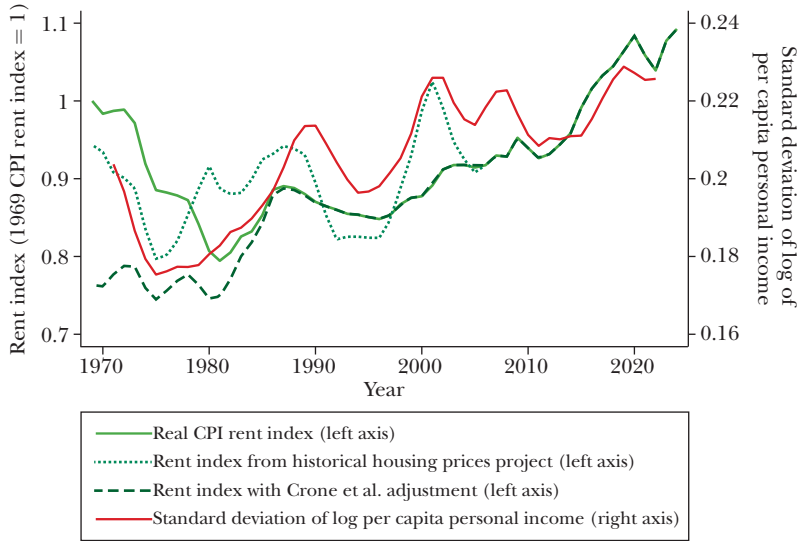
Source: Real CPI rent index is from the US Bureau of Labor Statistics (2024a). Rent index from historical housing prices project is from Lyons et al. (2024b). Rent index with Crone et al. adjustment comes from Crone, Nakamura, and Voith (2010b). Standard deviation of log per capita personal income is from the US Bureau of Economics Analysis (2023).
Note: Real CPI rent index is calculated by using the Consumer Price Index: Rent of Primary Residence in US City Average and dividing by the Consumer Price Index: All Items in US City Average. It is indexed to be 1 in 1969. The dashed line shows the rent index from Crone, Nakamura, and Voith (2010b), similarly adjusted for overall inflation. This series ends in 2000 and is indexed to end at the same level as the Real CPI Rent Index in 2000. The dotted line shows the rent index from Lyons et al. (2024b). This series ends in 2006 and is indexed to have the same level as the Real CPI rent index in 2006. For the standard deviation of log per capita personal income, we use county-level data on per capita personal income. We aggregate the data to core-based statistical areas based on the BEA's population estimates (for Virginia and Alaska, the BEA reports per capita personal income for county groups rather than individual counties. We treat each county group as a city rather than try to inexactly assign them to CBSAs). We take a backward-looking three-year moving average for each city, and then we take the standard deviation in each year of the log of that moving average.

Crone, Nakamura, and Voith (2010a) rent series. The period of flat rents in the 1990s corresponds to a period of falling regional inequality, and the same thing is true immediately in the early 2010s. At the same time, there are a few periods (like after 2000) when inequality falls but rents in the CPI index do not—although they do fall in the Lyons et al. (2024a) rent series. Also, as we discuss in Howard, Liebersohn, and Ozimek (2023), the COVID-19 pandemic increased housing prices while raising the demand to live in places with an elastic housing supply.

We argue that the rise in housing costs since 2019 is due to an increase in demand for housing, rather than demand to live in more expensive areas. Although we would be hesitant to draw too strong a conclusion from this graph, the data do

seem to be consistent with regional inequality playing a role in determining rents. Perhaps another way to put it is that given the cross-sectional evidence in favor of the assumptions of our theory, a comparison of the macro trends in rents and regional inequality does not provide any evidence against it.

For a more quantitative assessment of a related theory, we refer interested readers to Howard and Liebersohn (2021). In that paper, we use a quantitative spatial model to study the importance of the growth for demand to live in housing-supply-inelastic areas. Using elasticities from the literature and similar data to those presented in Figure 1, we argue that the majority of the national rent increase from 2000 to 2018 is due to the change in relative demand for high-income, low-elasticity cities.

Discussion and Implications

The link between regional inequality and housing affordability has applications for a variety of important economic questions and policies. We close by mentioning a few of them.

First, our essay has focused on housing costs in the United States, but there are good reasons to think that the same theory holds elsewhere. Housing costs have trended up in most countries of the Organization for Economic Cooperation and Development (OECD) in recent decades, as has regional inequality (Knoll, Schularick, and Steger 2017; International Monetary Fund 2019). In countries where researchers have studied the elasticity of housing supply, they have often found similar regional differences as in the United States (Beze 2023; Hilber and Vermeulen 2016). For an example from the United Kingdom, Hilber and Mense (2021) find that the interaction of supply elasticities and demand shocks can explain fluctuations in London house prices.

Second, while our discussion has largely avoided differentiating between rents and house prices, regional inequality matters for both. In particular, economists have had success with matching major swings in house prices by considering changes in regional inequality combined with assumptions about the way people extrapolate recent changes in the cost of housing to the future (Chodorow-Reich, Guren, and McQuade 2024; Howard and Liebersohn 2023).

Third, the emerging trend of remote work since the COVID-19 pandemic has changed the geography of income and is likely to continue to do so. Remote workers tend to be high income, raising the average incomes in high-amenity, low-cost regions. If this has a big impact on regional inequality, it could lead to long-term changes in housing affordability as well (Davis, Ghent, and Gregory 2024; Delventhal and Parkhomenko 2024; Howard, Liebersohn, and Ozimek 2023; Brueckner, Kahn, and Lin 2023). Of course, in general, the rise of remote work also led to demand for larger living spaces (for example, for home offices), which caused house prices to rise during the pandemic.

Fourth, cities that are stagnant or shrinking in population have a particularly inelastic housing supply due to the durability of buildings—a notable departure

from the typical negative relationship between median income and housing supply elasticity (Glaeser and Gyourko 2005; Notowidigdo 2020). Over the last few decades, this has not been particularly relevant for aggregate housing costs because very few cities are shrinking (weighted by population, only about 10 percent of cities from 2000 to 2018 shrunk, and most of those that shrunk did so by less than 3 percent). Nonetheless, should it become relevant, it could undo or even reverse the economic forces we highlighted here. If, for example, housing supply in the biggest cities were to increase, migration out of the lower-income regions would lead to lower housing supply elasticities there. In that case, further increases in regional inequality might not raise aggregate housing costs as much—and might even lower them.

Finally, while housing supply policy is primarily determined by local governments, there are significant spillovers across cities in the demand for housing. Therefore, policies that may affect relative demand to live across cities also affect national affordability. For example, place-based policies often target lower-income regions (Neumark and Simpson 2015; Glaeser and Gottlieb 2008), with an explicit goal of reducing regional inequality. Increasing housing affordability nationally can be an important unintended consequence. On the other hand, a variety of policies implicitly subsidize living in the highest-cost areas, including housing vouchers, the state and local tax deduction, and tax credits for housing costs. These policies will have the unintended consequence of reducing affordability nationally.

References

Ahlfeldt, Gabriel M., and Daniel P. McMillen. 2018. "Tall Buildings and Land Values: Height and Construction Cost Elasticities in Chicago, 1870–2010." *Review of Economics and Statistics* 100 (5): 861–75.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu. 2016. "Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis." NBER Working Paper 22816.

Albouy, David, Gabriel Ehrlich, and Minchul Shin. 2018. "Metropolitan Land Values." *Review of Economics and Statistics* 100 (3): 454–66.

Almagro, Milena, and Tomás Domínguez-Iino. 2024. "Location Sorting and Endogenous Amenities: Evidence from Amsterdam." NBER Working Paper 32304.

Barro, Robert J., and Xavier Sala-i-Martin. 1991. "Convergence across States and Regions." *Brookings Papers on Economic Activity* 22 (1): 107–82.

Bartik, Timothy J. 1991. "Boon or Boondoggle? The Debate over State and Local Economic Development Policies." In *Who Benefits from State and Local Economic Development Policies?*, 1–16. Upjohn Institute Press.

Baum-Snow, Nathaniel. 2023. "Constraints on City and Neighborhood Growth: The Central Role of Housing Supply." *Journal of Economic Perspectives* 37 (2): 53–74.

Baum-Snow, Nathaniel, and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–946.

Beze, Eyayaw. 2023. "Geographic Constraints and the Housing Supply Elasticity in Germany." Ruhr Economic Paper 1003.

Blanchard, Olivier Jean, and Lawrence F. Katz. 1992. "Regional Evolutions." *Brookings Papers on Economic Activity* 23 (1): 1–75.

Brueckner, Jan K., Matthew E. Kahn, and Gary C. Lin. 2023. "A New Spatial Hedonic Equilibrium in the Emerging Work-from-Home Economy?" *American Economic Journal: Applied Economics* 15 (2): 285–319.

Card, David, Jesse Rothstein, and Moises Yi. 2025. "Location, Location, Location." *American Economic Journal: Applied Economics* 17 (1): 297–336.

Cardullo, Gabriele, and Agnese Sechi. 2023. "Local Labor Markets with Non-homothetic Preferences." IZA Discussion Paper 16533.

Chodorow-Reich, Gabriel, Adam M. Guren, and Timothy J. McQuade. 2024. "The 2000s Housing Cycle with 2020 Hindsight: A Neo-Kindlebergerian View." *Review of Economic Studies* 91 (2): 785–816.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst. 2024. "Income Growth and the Distributional Effects of Urban Spatial Sorting." *Review of Economic Studies* 91 (2): 858–98.

Crone, Theodore M., Leonard I. Nakamura, and Richard Voith. 2010a. "Rents Have Been Rising, Not Falling, in the Postwar Period." *Review of Economics and Statistics* 92 (3): 628–42.

Crone, Theodore M., Leonard I. Nakamura, and Richard Voith. 2010b. *Data for: "Rents Have Been Rising, Not Falling, in the Postwar Period." Review of Economics and Statistics* 92 (3): 628–42.

Davis, Morris A., Andra C. Ghent, and Jesse Gregory. 2024. "The Work-from-Home Technology Boon and Its Consequences." *Review of Economic Studies* 91 (6): 3362–401.

Davis, Morris A., William D. Larson, Stephen D. Oliner, and Jessica Shui. 2021. "The Price of Residential Land for Counties, ZIP Codes, and Census Tracts in the United States." *Journal of Monetary Economics* 118: 413–31.

Davis, Morris A., and Michael G. Palumbo. 2008. "The Price of Residential Land in Large US Cities." *Journal of Urban Economics* 63 (1): 352–84.

Delventhal, Matt, and Andrii Parkhomenko. 2024. "Spatial Implications of Telecommuting." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.3746555.

Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000." *American Economic Review* 106 (3): 479–524.

Diamond, Rebecca, and Enrico Moretti. 2021. "Where Is Standard of Living the Highest? Local Prices and the Geography of Consumption." NBER Working Paper 29533.

Duranton, Gilles, and Diego Puga. 2004. "Micro-foundations of Urban Agglomeration Economies." In *Handbook of Regional and Urban Economics*, Vol. 4, edited by J. Vernon Henderson and Jacques-François Thisse, 2063–117. Elsevier.

Duranton, Gilles, and Diego Puga. 2020. "The Economics of Urban Density." *Journal of Economic Perspectives* 34 (3): 3–26.

Eckert, Fabian, Sharat Ganapati, and Conor Walsh. 2022. "Urban-Biased Growth: A Macroeconomic Analysis." NBER Working Paper 30515.

Epple, Dennis, Thomas Romer, and Radu Filimon. 1988. "Community Development with Endogenous Land Use Controls." *Journal of Public Economics* 35 (2): 133–62.

Fajgelbaum, Pablo D., Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar. 2019. "State Taxes and Spatial Misallocation." *Review of Economic Studies* 86 (1): 333–76.

Federal Housing Finance Agency. 2025. *County House Price Index (HPI).* US Federal Housing. https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index-Datasets.aspx (accessed May 2025).

Fischel, William A. 2005. *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies.* Harvard University Press.

Foster, Brad, Mark Ellis, and Lee Fiorio. 2023. "Agree to Disagree? Comparing IRS, NCOA, and Census Bureau Survey Migration Measures." Presented at the Federal Committee on Statistical Methodology, Hyattsville, MD, October 25.

Fowler, C. S., & Jensen, L. 2020. *Data for: "Bridging the gap between geographic concept and the data we have: The case of labor markets in the USA." Environment and Planning A: Economy and Space* 52(7): 1395–1414. https://doi.org 10.1177/0308518X20906154.

Furth, Salim. 2025, "Response to 'Supply Constraints Do Not Explain House Price and Quantity Growth Across U.S. Cities' by Louie, Mondragon, and Wieland." Preprint, SSRN. https://ssrn.com/abstract=5227968.

Ganong, Peter, and Daniel Shoag. 2017. "Why Has Regional Income Convergence in the US Declined?" *Journal of Urban Economics* 102: 76–90.

Giannone, Elisa. 2017. "Skilled-Biased Technical Change and Regional Convergence." Unpublished.

Glaeser, Edward L., and Joshua D. Gottlieb. 2008. "The Economics of Place-Making Policies." *Brookings Papers on Economic Activity* 39 (1): 155–239.

Glaeser, Edward L., and Joshua D. Gottlieb. 2009. "The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States." *Journal of Economic Literature* 47 (4): 983–1028.

Glaeser, Edward L., and Joseph Gyourko. 2005. "Urban Decline and Durable Housing." *Journal of Political Economy* 113 (2): 345–75.

Glaeser, Edward, and Joseph Gyourko. 2018. "The Economic Implications of Housing Supply." *Journal of Economic Perspectives* 32 (1): 3–30.

Glaeser, Edward L., Joseph Gyourko, and Raven E. Saks. 2005. "Why Have Housing Prices Gone Up?" *American Economic Review* 95 (2): 329–33.

Glaeser, Edward L., and Bryce A. Ward. 2009. "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston." *Journal of Urban Economics* 65 (3): 265–78.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." *Journal of Political Economy* 118 (3): 536–98.

Gyourko, Joseph, Jonathan S. Hartley, and Jacob Krimmel. 2021. "The Local Residential Land Use Regulatory Environment across US Housing Markets: Evidence from a New Wharton Index." *Journal of Urban Economics* 124: 103337.

Gyourko, Joseph, and Raven Molloy. 2015. "Regulation and Housing Supply." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 1289–337. North-Holland.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008a. "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45 (3): 693–729.

Gyourko, Joseph, Albert Saiz, and Anita Summers. 2008b. *Data for: "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index."* Urban Studies 45 (3): 693–729.

Hilber, Christian A. L., and Andreas Mense. 2021. "Why Have House Prices Risen So Much More than Rents in Superstar Cities?" Centre for Economic Performance Discussion Paper 1743.

Hilber, Christian A. L., and Frédéric Robert-Nicoud. 2013. "On the Origins of Land Use Regulations: Theory and Evidence from US Metro Areas." *Journal of Urban Economics* 75: 29–43.

Hilber, Christian A. L., and Wouter Vermeulen. 2016. "The Impact of Supply Constraints on House Prices in England." *Economic Journal* 126 (591): 358–405.

Howard, Greg. 2020. "The Migration Accelerator: Labor Mobility, Housing, and Demand." *American Economic Journal: Macroeconomics* 12 (4): 147–79.

Howard, Greg. 2024. "Moving Cost Magnitudes in Moving Cost Models." Unpublished.

Howard, Greg, and Jack Liebersohn. 2021. "Why Is the Rent So Darn High? The Role of Growing Demand to Live in Housing-Supply-Inelastic Cities." *Journal of Urban Economics* 124: 103369.

Howard, Greg, and Jack Liebersohn. 2023. "Regional Divergence and House Prices." *Review of Economic Dynamics* 49: 312–50.

Howard, Greg, and Jack Liebersohn. 2025. *Data and Code for: "How Regional Inequality and Migration Drive Housing Prices and Rents."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E230944V1.

Howard, Greg, Jack Liebersohn, and Adam Ozimek. 2023. "The Short- and Long-Run Effects of Remote Work on US Housing Markets." *Journal of Financial Economics* 150 (1): 166–84.

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* 11 (2): 1–39.

International Monetary Fund. 2019. *Global Economic Outlook.* International Monetary Fund.

Kaplan, Greg, and Sam Schulhofer-Wohl. 2017. "Understanding the Long-Run Decline in Interstate Migration." *International Economic Review* 58 (1): 57–94.

Kennan, John, and James R. Walker. 2011. "The Effect of Expected Income on Individual Migration Decisions." *Econometrica* 79 (1): 211–51.

Knoll, Katharina, Moritz Schularick, and Thomas Steger. 2017. "No Price Like Home: Global House Prices, 1870–2012." *American Economic Review* 107 (2): 331–53.

Louie, Schuyler, John A. Mondragon, and Johannes Wieland. 2025. "Supply Constraints Do Not Explain House Price and Quantity Growth across US Cities." NBER Working Paper 33576.

Lyons, Ronan C., Allison Shertzer, Rowena Gray, and David N. Agorastos. 2024a. "The Price of Housing in the United States, 1890–2006." NBER Working Paper 32593.

Lyons, Ronan C., Allison Shertzer, Rowena Gray, and David N. Agorastos. 2024b. *Data for: "The Price of Housing in the United States, 1890–2006."* NBER Working Paper 32593.

Manson, Steven, Jonathan Schroeder, David Van Riper, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles. 2025. *IPUMS National Historical Geographic Information System: Version 19.0.* IPUMS. http://doi.org/10.18128/D050.V19.0.

Mayo, Stephen K. 1981. "Theory and Estimation in the Economics of Housing Demand." *Journal of Urban Economics* 10: 95–116.

Mill, John Stuart. 1848. *Principles of Political Economy.* John W. Parker.

Missouri Census Data Center. 2025. *Geographic Correspondence Engine and Files.* University of Missouri Center for Health Policy. https://mcdc.missouri.edu/applications/geocorr.html (accessed May 2025).

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2011. "Internal Migration in the United States." *Journal of Economic Perspectives* 25 (3): 173–96.

Molloy, Raven, Christopher L. Smith, and Abigail Wozniak. 2017. "Job Changing and the Decline in Long-Distance Migration in the United States." *Demography* 54 (2): 631–53.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. "Commuting, Migration, and Local Employment Elasticities." *American Economic Review* 108 (12): 3855–90.

Moretti, Enrico. 2012. *The New Geography of Jobs.* Houghton Mifflin Harcourt.

Neumark, David, and Helen Simpson. 2015. "Place-Based Policies." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 1197–287. North-Holland.

Notowidigdo, Matthew J. 2020. "The Incidence of Local Labor Demand Shocks." *Journal of Labor Economics* 38 (3): 687–725.

Parkhomenko, Andrii. 2023. "Local Causes and Aggregate Implications of Land Use Regulation." *Journal of Urban Economics* 138 : 103605.

Quigley, John M., and Steven Raphael. 2004. "Is Housing Unaffordable? Why Isn't It More Affordable?" *Journal of Economic Perspectives* 18 (1): 191–214.

Rappaport, Jordan. 2009. "The Increasing Importance of Quality of Life." *Journal of Economic Geography* 9 (6): 779–804.

Redding, Stephen J., and Esteban Rossi-Hansberg. 2017. "Quantitative Spatial Economics." *Annual Review of Economics* 9: 21–58.

Ricardo, David. 1817. *Principles of Political Economy and Taxation.* John Murray.

Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life." *Journal of Political Economy* 90 (6): 1257–78.

Rosen, Sherwin. 1979. "Wage-Based Indexes of Urban Quality of Life." In *Current Issues in Urban Economics*, edited by Peter Mieszkowski and Mahlon Straszheim, 74–104. Johns Hopkins University Press.

Rubinton, Hannah. 2020. "The Geography of Business Dynamism and Skill Biased Technical Change." Federal Reserve Bank of St. Louis Working Paper 2020-20.

Saiz, Albert. 2010a. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.

Saiz, Albert. 2010b. *Data for: "The Geographic Determinants of Housing Supply." Quarterly Journal of Economics* 125 (3): 1253–96.

Schubert, Gregor. 2021. "House Price Contagion and US City Migration Networks." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4984580.

Suárez Serrato, Juan Carlos, and Owen Zidar. 2016. "Who Benefits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review* 106 (9): 2582–624.

SEER (National Cancer Institute Surveillance, Epidemiology, and End Results Program). 1969–2022.

U.S. County Population Estimates: 19 Age Groups, White | Black | Other (Adjusted)—All States Combined. Bethesda, MD: National Cancer Institute, Division of Cancer Control and Population Sciences. https://seer.cancer.gov/popdata.thru.2022/yr1969_2022.19ages/us.1969_2022.19ages.adjusted. txt.gz (accessed July 8, 2025).

US Bureau of Economic Analysis. 2023. *Regional GDP and Personal Income.* https://www.bea.gov/itable/ regional-gdp-and-personal-income (accessed October 31, 2023).

US Bureau of Labor Statistics. 2024a. *Consumer Price Index for All Urban Consumers: Rent of Primary Residence (CUUR0000SEHA).* Distributed by Federal Reserve Bank of St. Louis. https://fred.stlouisfed. org/series/CUUR0000SEHA (accessed June 3, 2024).

US Bureau of Labor Statistics. 2024b. *FRED Economic Data.* Distributed by Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/graph/?g=1IZBh (accessed May 2025).

Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill. 2010. "Why Has House Price Dispersion Gone Up?" *Review of Economic Studies* 77 (4): 1567–606.

Curbing Rising Housing Costs:
A Model-Based Policy Comparison

Boaz Abramson and Tim Landvoigt

Housing is becoming less affordable in the US economy, as housing prices have grown strongly relative to incomes. In 1980, the median house price was \$47,500 and the median household income was \$18,200. By 2020, the median house price rose to \$250,000, reflecting an annual growth rate of approximately 4.2 percent. At the same time, median income grew on average by only 3.5 percent annually, reaching \$73,000 in 2020. The divergence between house prices and income has been particularly pronounced since the Great Recession of 2007–2009. Between 2012 and 2022, while income has continued to grow at a steady annual pace of 3.5 percent, house prices appreciated at a much faster pace of 5.8 percent annually.

The ratio of house price to income, a common measure of (lack of) affordability, is higher for lower-income homeowners. This is illustrated in the left panel of Figure 1. In 2020, the price-to-income ratio was above ten for owners in the bottom income decile, compared to only three for the highest-income owners. In addition, the rise in the price-to-income ratio has been more dramatic for lower-income homeowners since 2000. A similar pattern is observed in the rental market, as illustrated by the right panel of Figure 1. Since 2000, rents have grown faster relative to renters' income, and particularly so for the lowest income renters.

■ *Boaz Abramson is Assistant Professor of Business, Columbia Business School, New York City, New York. Tim Landvoigt is Associate Professor of Finance, Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania. Their email addresses are ba2695@columbia.edu and timland@wharton.upenn.edu.*

Figure 1
Housing Affordability
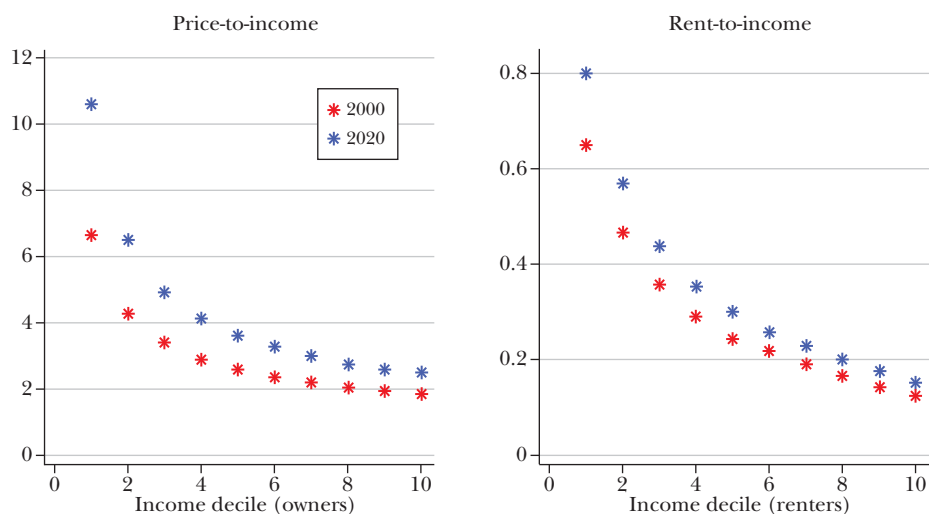
Source: Tabulations are based on American Community Survey data.
Note: The left (right) panel plots a bin-scatter of the price-to-income (rent-to-income) ratio as a function of homeowners (renters) household income in 2000 (in red) and in 2020 (in blue).

Numerous policies have been proposed to make housing more affordable. Most of these policies fall in one of two categories. For economists, at least, perhaps the most obvious remedy for high prices is to increase the supply of housing. Examples of supply-side policies include easing of housing supply regulations, tax credits for developers, and public housing. An alternative set of policies focus on making mortgages more available and less expensive. Examples include down-payment assistance for first-time buyers, mortgage interest deductions, subsidies for mortgage interest payments, or government guarantees of mortgage credit.

At a high level, both types of policies have intuitive appeal. Increasing supply should lower house prices, and cheaper mortgage credit should reduce housing costs for homeowners at a given level of house prices. However, as is often the case, the effects of these policies on housing affordability are less obvious. While there is widespread agreement that more housing construction is needed, it is not clear what type of housing should be constructed and in which locations. Cheap and easily available mortgage credit might cause upward pressure on prices through increased demand for housing and hence exacerbate the affordability problem instead of solving it. Should policymakers ease housing supply regulations, subsidize construction of small apartments, or subsidize credit for first-time homebuyers? How effective are these policies and how can they be compared?

To answer these questions, we develop a model of a city's housing market that allows us (1) to understand the sources of rising price-to-income ratios and (2) to

evaluate different policies in a unified framework. Using the model, we can evaluate policies, even if these policies have yet to be implemented or were implemented in a way that does not lend itself to directly observing causal effects in the data.[1] In the model, households optimally choose consumption of housing and other goods given their budget. These households differ by their level of wealth. The supply of housing is differentiated by its level of quality. In equilibrium, the wealthiest households end up owning the highest quality houses and the poorest households end up renting the lowest quality houses. This is known as an "assignment equilibrium" in housing markets (Määttänen and Terviö 2014; Landvoigt, Piazzesi, and Schneider 2015).[2]

To apply the model to data, we calibrate the key parameters—most prominently the distributions of wealth and housing quality—so that the model matches the observed house prices and rents in San Francisco in 2013.[3] The calibrated model shows that rising wealth dispersion, together with stagnating housing supply, can explain the increase in price-to-income and rent-to-income ratios observed in Figure 1. Why did this confluence of trends likely contribute to the rise in prices and rents? Rising wealth inequality combined with stagnating supply mean that poor households end up competing with richer households for the same housing units.

Having shown that the model performs well in matching the status quo observed in the data, we then use it to evaluate how supply-side and demand-side policies would affect price-to-income and rent-to-income ratios. Here, we want to emphasize a few conclusions from our analysis upfront. First, measures that target the demand side of the housing market are mostly ineffective: in our model, a down-payment subsidy for potential home buyers increases upwards pressure on house prices, especially in those segments of the market in which buyers benefit most from the subsidies. The model strongly suggests that any solution to the affordability crisis involves more supply. Indeed, we find that additional housing supply is always most beneficial for the poorest residents of the city. One perhaps surprising conclusion is that new construction in the high-end segments of the market decreases price-to-income ratios by more in *all segments* of the housing market compared to new construction in bottom-end segments. This is because new construction in high-end segments reduces competition for units in medium- and low-quality housing segments, improving affordability throughout the wealth distribution.

---

[1] A large literature in economics studies how housing supply or mortgage subsidies affect house prices. For example, several papers analyze how changes in construction activity for different types of housing at the city level have affected housing prices (Baum-Snow and Marion 2009; Glaeser and Ward 2009; Mense 2025; Pennington 2021; Asquith, Mast, and Reed 2023).

[2] An early model of segmented housing markets is Ortalo-Magné and Rady (2006). More recent examples include Nathanson (2023), Nikolakoudis (2024), Mense (2025), and Fonseca, Liu, and Mabille (2025).

[3] The San Francisco metro area serves as the main example in our analysis. In Supplemental Appendix D, we perform the identical analysis for Chicago and Dallas, and highlight differences and commonalities.

Construction in low-tier segments is similarly beneficial for the poorest residents, but it improves affordability much less for middle-income and wealthy households.[4]

A Model of Segmented Housing Markets

The main building blocks of our model are households, who are the residents of the city we analyze, and a housing market structure, which specifies the type of residences available for purchase or rent and how these properties are traded. We describe all model elements in greater detail, using graphical tools.[5]

Households

Households in the model maximize utility over two life stages, referred to as "today" and "tomorrow," but which can be considered to be roughly a decade apart. Households are identical except for initial wealth, which fully determines their choices—households with the same wealth behave identically. With only two periods, we do not distinguish between income and financial assets; the sum of both is referred to as "wealth."

Utility depends on two goods: housing services and a composite of all other consumption. Utility increases with either good, but at a decreasing rate. The household's objective is a weighted average of utility across both periods, with full weight on today and a smaller weight on tomorrow. Households simultaneously choose total consumption in each period, the share allocated to housing, and savings.

Their decisions are constrained by today's total wealth, which must cover today's consumption, housing costs (rent or purchase), and savings. Tomorrow, households receive savings with interest, sell their house if owned, and gain additional wealth. These sources fund tomorrow's consumption. Households can save a positive amount, interpreted as a savings account. If purchasing a home, they can have negative savings, which we interpret as borrowing via a mortgage, subject to a loan-to-value limit. Renters cannot borrow.

The solution to the household's problem depends on the utility function, model parameters (like the interest rate), and house prices and rents. Model parameters are ultimately chosen to match relevant aspects of the data in a process called calibration, which we describe below.

---

[4] Other studies have reached similar conclusions. For example, Nathanson (2023) argues, using a different type of model, that low-end construction may not be the best policy to improve affordability in the Boston-Cambridge-Newton, MA-NH, metropolitan area. Recent structural papers focusing on various aspects of housing affordability include Diamond and McQuade (2019), Hsieh and Moretti (2019), Anenberg and Kung (2020), Abramson (2021), Favilukis and Van Nieuwerburgh (2021), Favilukis, Mabille, and Van Nieuwerburgh (2023), İmrohoroğlu and Zhao (2022), Couture et al. (2023), Corbae, Glover, and Nattinger (2024), Abramson and Van Nieuwerburgh (2024), and Gupta, Hansman, and Mabille (2024).

[5] A full and formal mathematical description of the model is presented in Supplemental Appendix A.

Housing Markets

The other key component of the model is the residential housing market, where households either purchase or rent homes. This market is segmented by quality: lower-quality segments yield less housing services and are available for rent, while higher-quality segments offer more housing services and are only available for purchase. The market comprises N competitive submarkets, each representing a distinct quality level, with quality defined by the housing services entering a household's utility if the home is acquired or rented. Thus, each household's housing decision reduces to choosing a quality segment.

The model assumes market equilibrium, with prices determined by supply and demand in each segment. Supply is provided by real estate developers and construction firms outside the model. While it would be straightforward to include a construction sector explicitly, this analysis assumes the city government can direct supply distribution across segments. Demand in each segment results from the household optimization problem described above. Each household buys or rents exactly one property, so market clearing requires that the number of households in each segment matches the supply.

Equilibrium Structure

To illustrate how these market clearing conditions pin down the prices, consider a hypothetical example with three quality segments, illustrated in Figure 2. There is a total mass of households of 1, distributed uniformly between a lower bound of wealth W and an upper bound W. Supply of houses within each quality segment is given by H_0, H_1, and H_2, respectively. The two top housing segments correspond to the owner-occupied market, meaning that houses in these segments are offered for purchase, while the bottom segment corresponds to the rental market, meaning that houses in this segment are offered for rent. We normalize the price of living in the bottom segment so that it is less than the lower bound of the wealth distribution.

While solving for equilibrium outcomes in this setting may seem complex at first glance, it turns out that the ordering of households by wealth greatly simplifies the analysis. In particular, it is straightforward to show that any equilibrium in this model has a so-called "monotone assignment structure"; that is, households are assigned to houses by wealth, with the poorest households living in the bottom segment and the richest households in the top segment. After all, the willingness to pay of households is increasing in their wealth, and wealthier households are thus always willing to outbid poorer households competing for the same house.

To understand this equilibrium, focus on the households with threshold wealth levels W_1 and W_2. Households with wealth above W_2 strictly prefer buying a house in the top segment with quality h_2 and corresponding price p_2. Households in the interval $[W_1, W_2]$ strictly prefer buying a house with quality h_1 at price p_1. Households with wealth below W_1 are priced out of the owner market and forced to rent in segment 0. Households that are exactly at the threshold levels are exactly indifferent between the segments below or above.

Figure 2
Assignment Equilibrium Structure
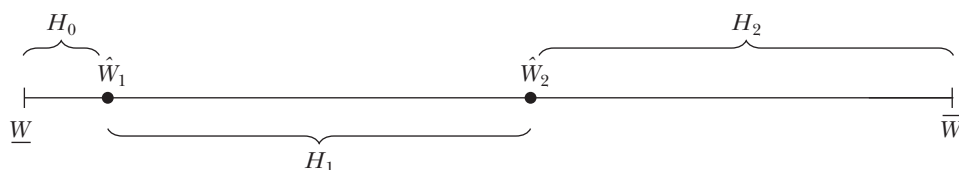
Note: This figure illustrates the assignment equilibrium with a uniform wealth distribution and three housing segments. The equilibrium is characterized by threshold levels of wealth W_1 and W_2.

How are the threshold levels determined? Because we know that each household buys exactly one house, and how many houses exist in each segment, the threshold levels can be directly calculated. The prices p_1 for housing level H_1 and p_2 for housing quality level H_2 are then determined such that the households at the thresholds are indifferent between neighboring segments. Solving the model comes down to computing the prices that solve the equilibrium conditions using numerical methods.

Applying the Model to Data

We now bring this model to the data. Because the model only has two periods, we focus our analysis on the long-run equilibrium in the housing market: one period of time in the model corresponds to ten years. Thus, our model-based projections apply to long-term effects. We believe this focus is justified, because residential real estate is long-lived and changes to the housing stock are very persistent. We first establish a baseline by matching the model to a specific city in a specific year: the San Francisco-Oakland-Hayward Metropolitan Statistical Area in 2013. This urban area has had one of the most severe affordable housing problems in the country; the San Francisco-Oakland-Hayward MSA ranked among the top ten in terms of median housing price-to-income ratio in every year between 2012 and 2022.[6] We choose 2013 as baseline year, because it marks the end of the recovery period from the 2008 housing bust and allows us to analyze the impact of housing policy choices in the subsequent decade. We measure realized changes in the distributions of (1) housing supply and (2) the wealth of residents over the 2013–2022 period. By feeding these measured changes as inputs into the model, we investigate whether realized shifts in wealth and housing supply can explain the decline in affordability

---

[6] Our calculations here are based on data from the American Community Survey. We perform the same analysis for Dallas and Chicago in Supplemental Appendix D, thus demonstrating the broader applicability of this approach.

ratios over the 2013–2022 period. The result is that we have a representation of San Francisco in 2022 in the calibrated model. In the next section, we study several experiments within this 2022 version of the model.

Matching the Model to Baseline Data

As a starting point to establish a baseline for analysis, we calibrate the model to the San Francisco-Oakland-Hayward Metropolitan Statistical Area in 2013. The goal of the calibration is to ensure that the model matches the economic environment that is being studied. Because we are studying housing markets, the key empirical moments to be matched are house prices, rents, and the allocation of households into housing units.[7]

In the model, prices and housing allocation are determined by the intersection of demand and supply. Demand is determined by households' wealth and preferences. We calibrate the wealth distribution of households to match the empirical wealth distribution of households in the data. Using the American Community Survey (ACS) and the Survey of Consumer Finances (SCF), we estimate the wealth distribution of households in San Francisco in 2013. The resulting wealth distribution is illustrated in blue in the top left panel of Figure 4. We calibrate household preferences using standard values from the literature on housing economics. Supply is governed by the distribution of the housing stock across quality segments. We assume that there are 13 quality segments in the city. Houses in the top ten segments are houses that households can own. They account for 54.7 percent of the total housing stock, which is the observed home-ownership rate in San Francisco in 2013. Each of the top ten segments is assumed to be of equal size.

Housing in the bottom three quality segments are residences that households can rent. The top two segments of the rental market correspond to the formal rental market and are assumed to be of equal size. The lowest segment corresponds to informal renting arrangements and accounts for 1.5 percent of the housing stock, which is the observed homelessness rate in the data (Abramson 2021; Abramson and Van Nieuwerburgh 2024). The distribution of the housing stock across quality segments is illustrated by the blue bars in the top right panel of Figure 4. The calibration of housing supply ensures that the model matches the homeownership rate, the rentership rate, and the share of households experiencing housing insecurity—all important statistical moments for studying housing affordability.

Finally, we ensure the model matches observed house prices and rents. To do so, we estimate the house quality in each of the housing segments so that the model implied prices are in line with the data. We estimate the median house price in each segment of the owner-occupied market using Corelogic, a private vendor that compiles data on the universe of US housing transactions and property tax records. Median rents are measured using American Community Survey data. The intuition for the estimation is straightforward. Given household preferences, wealth

---

[7] We discuss the calibration in detail in Supplemental Appendix B.

distribution, the distribution of the housing stock across segments, and the house quality in each segment, we can solve for the prices that equilibrate all housing markets. All else equal, different vectors of house qualities translate to different vectors of house prices.

How Good Is the Model?

At this point in the analysis, we confront the question of whether the model provides an adequate account of the economic environment that is being studied. We will answer this question in two steps.

The first step is evaluating the model's fit to targeted data moments. These are moments that the model was specifically calibrated to match—in our context, house prices and rents. As explained above, the vector of house qualities is estimated to target the median house price within each segment of the owner-occupied market as well as the median rent within each segment of the rental market. The left panel of Figure 3 compares the model-implied prices and rents (in blue circles) to the targeted prices and rents in the data (in red dots). Clearly, the model fits the data well. This result is not especially surprising, as the model was fitted to the 2013 San Francisco data. But of course, if the model did not fit the data well, it would indicate that it is missing something crucial.

A second step to establish model credibility is to evaluate the model's fit to data that the calibration of the model did not specifically target, but that are important for the context under consideration. In our context, two important affordability measures are the price-to-income ratio and the rent-to-income ratio. The middle panel of Figure 3 plots a bin-scatter of the price-to-income ratio as a function of homeowners' income in the model (in blue) and in the data (in red, computed based on the 2013 American Community Survey).[8] The right panel of Figure 3 similarly compares the rent-to-income ratio as a function of renters' income in the model and in the data. Both panels illustrate that the model closely aligns with the data. In addition, the downward-sloping pattern illustrates that concerns about housing affordability are especially prevalent for the bottom of the income distribution of owners and renters.[9]

---

[8] In the model, income is a component of household wealth and does not play an independent role. When solving the model, we draw household wealth from the calibrated (empirical) wealth distribution but do not explicitly track how much of each household's wealth is due to income. To compute the price-to-income ratio in the model, we do explicitly track household incomes. We do so by jointly drawing household wealth and income from the empirical joint distribution of (imputed) wealth and income calculated from the 2013 ACS.

[9] The rent-to-income in the seventh decile of the renter distribution of the model is higher than in the sixth decile because the renter market is separated into two discrete segments, while the income distribution is continuous. Because renter households sort into the two segments of the rental market based on their income, there is an income cutoff such that renter households with an income above the threshold rent in the upper segment, while renter households with an income below the threshold rent in the bottom segment. The rent-to-income ratio of the renter household just above the cutoff is therefore higher than that of the renter just below the cutoff.

Figure 3
Quantitative Model
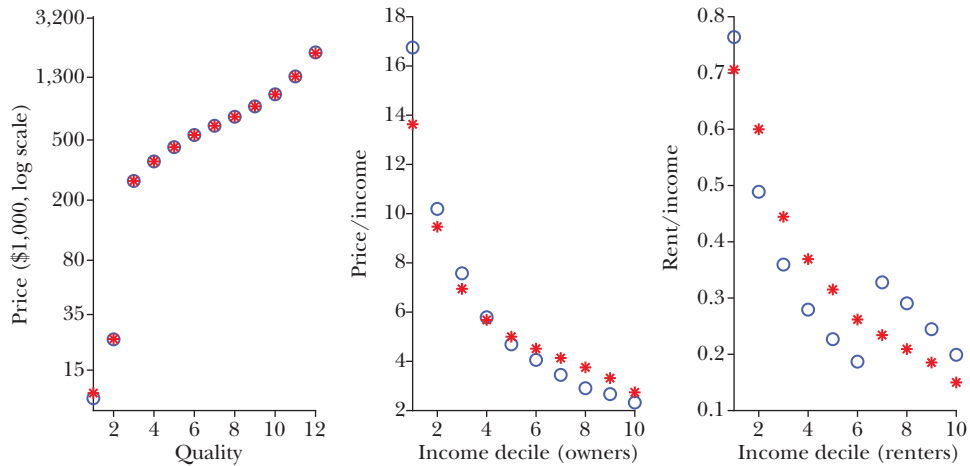
Source: Data moments are computed based on the American Community Survey and Corelogic data.
Note: This figure illustrates equilibrium in the quantified model. The left panel shows equilibrium prices for each segment in the model (blue) and in the data (red). The mid (right) panel plots a bin-scatter of the price-to-income (rent-to-income) ratio as a function of homeowners (renters) household income in the model (in blue) and in the data (in red). The estimation is described in more detail in Supplemental Appendix B.

Drivers of the Affordability Crisis

We have calibrated the model to San Francisco in 2013. Next, we use the calibrated model to quantify how changes in the wealth and housing supply distributions between 2013 and 2022 affect affordability.

In 2022, the wealth distribution shifts upward: existing residents become wealthier, along with in-migration of high-wealth households. This change is illustrated in the top-left panel of Figure 4 by the shift in the empirical wealth distribution from 2013 (in blue) to 2022 (in red). To give a sense of magnitude, the median real wealth in San Francisco more than doubled from $270,000 in 2013 to $644,000 in 2022. The wealth distribution in 2022 is estimated by combining data from the American Community Survey and the Survey of Consumer Finances, following the same steps described for the baseline calibration.

The distribution of houses across the 13 housing segments in San Francisco also changed. First, the share of houses supplied to the owner-occupied market (the top ten housing quality segments) increased from 54.7 percent (the ownership rate in 2013) to 57.1 percent (the ownership rate in 2022). For the two rental market segments, we maintain the baseline assumption that they are of equal size. Further, we estimate that the size of the lowest segment has increased from 1.5 percent to 2.4 percent. The increase in the size of the bottom segment is estimated so that the increase in rent burden in the bottom income decile between 2013 and 2022

implied by the model matches the increase observed in the American Community Survey data, which is 11 percent. The increase in the size of the first segment implies that between 2013 and 2022 the share of the San Francisco population who are either homeless or live in informal renting arrangements, such as doubling-up, has increased by 0.9 percentage points. In other words, on net, housing supply has not kept up with population growth.

Finally, within the top ten owner-occupied segments, we estimate that supply has shifted from the middle segments to the bottom and top segments, as illustrated by the top-right panel of Figure 4. Our estimate of this change in the quality distribution within the owner market requires a method to separate supply and demand factors in house price changes from 2013 to 2022. This task is notoriously difficult, and our approach establishes an upper bound on the shifts in relative supply by effectively attributing all relative price changes (that are not due to the aggregate house price trend in San Francisco) to supply changes.[10] Even though we compute an upper bound, Figure 4 shows that the changes in relative supply are small in an absolute sense, with the share of supply coming from the lowest and highest segments gaining a few percentage points from the middle.

The bottom panels of Figure 4 illustrate how the observed changes in the supply of housing and in the wealth of households affected housing prices, price-to-income, and rent-to-income ratios in San Francisco. Relative to 2013 (in blue), house prices and rents increase in 2022 (in red), and housing becomes less affordable for both owners and renters, especially among the lowest-income households. These patterns mimic the national patterns presented in Figure 1. We have further verified that the model's predictions regarding prices, price-to-income and rent-to-income ratios in San Francisco in 2022 are closely in line with the data, providing further model validation.[11]

Through the lens of the model of utility-maximizing households, a housing market divided into quality segments, and an assignment equilibrium, we can now explain this deterioration in price-to-income and rent-to-income measures. All else equal, the increase in households' wealth drives up house prices. Intuitively, when households have more funds at their disposal, but the distribution of house qualities is largely unchanged, prices must increase across the quality distribution. The rise in wealth can in theory be mitigated by supplying more housing, but in fact, the overall supply of housing decreased (as illustrated in Figure 4 by the increase in the size of the bottom segment), further deepening the affordability problem in the rental market. The modest increase in supply in the top segments does little to improve affordability because it comes at the expense of supply in other segments. As illustrated in the next section, a more substantial increase in supply in the top segments that increases overall housing supply can alleviate the affordability crisis.

---

[10] We provide details on this procedure in Supplemental Appendix B.1.
[11] This validation is spelled out in Supplemental Appendix C.1.

Figure 4
Drivers of the Affordability Crisis
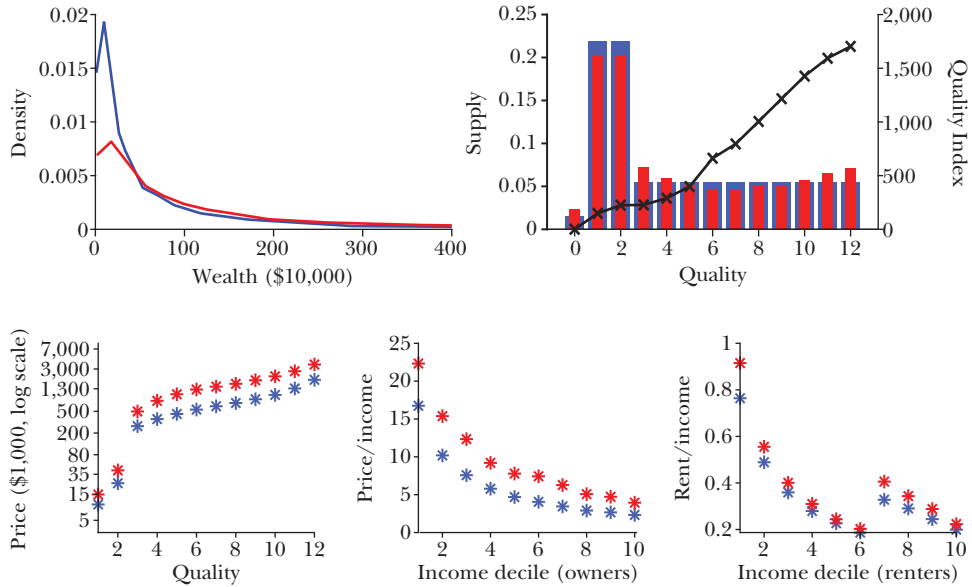
Note: This figure illustrates the effect of the observed change in housing supply and in the wealth distribution between 2013 and 2022. The top-left graph plots the density of the calibrated wealth distributions. The top-right panel shows the distributions of housing supply across segments (bars, left axis) and the quality index associated with housing segments (line, right axis). The bottom-left panel shows equilibrium prices for each segment. The bottom-middle (right) panel plots a bin-scatter of the price-to-income (rent-to-income) ratio as a function of homeowners' (renters') household income. Blue corresponds to the baseline 2013 economy and red corresponds to the 2022 economy.

Evaluating Proposed Solutions to High Housing Prices

How can policymakers lower the cost of housing? In this section, we conduct a sequence of numerical experiments within our quantitative model to evaluate the implications of three commonly proposed policies. More specifically, we identify different parameters of the model that represent policy choices. We vary these parameters, one at a time, while keeping all other parameters fixed at their 2022 values. We then analyze the effects of these changes on house prices, rents, and the allocation of housing.

Adding Housing Supply at the Bottom

Increasing the supply of housing should lead to lower house prices. But what type of housing that should be constructed? To answer this question, we compare the effects of increasing housing supply in the bottom segment to the effects of increasing housing supply in the top segment. In both experiments, we

Figure 5
Results of Three Experiments
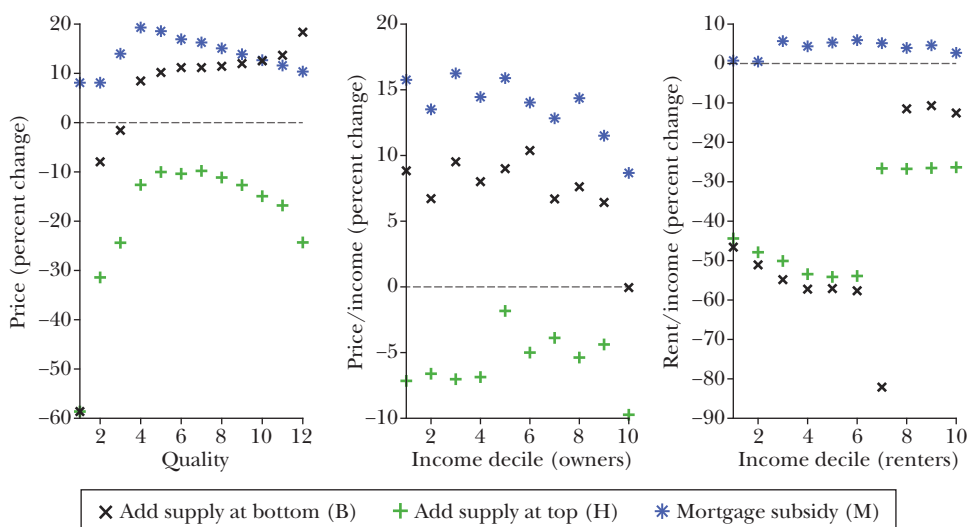
Note: This figure illustrates the effect of three experiments: the black crosses show the effect of an increase in supply at the bottom of the rental market; the green plus signs show the effect of an increase in supply at the top quality segment of the housing market; and the blue asterisks show the effect of a mortgage subsidy. The left panel shows the percent change in equilibrium prices. The middle (right) panel plots the percent change in the price-to-income (rent-to-income) ratio as a function of homeowners' (or renters') income.

increase the total supply of housing in the economy by 2 percent, which in the San Francisco-Oakland-Hayward area accounts for approximately 33,000 additional housing units. We envisage this change happening exogenously—say, via local government decisions on zoning, permitting, and tax credits that facilitate construction of housing.

In our first experiment, the newly built housing units are in the bottom of the two segments of the rental market. In this experiment, we further increase supply in the bottom segment by converting 10 percent of the housing units in the top housing segment to units in the bottom rental segment. Overall, the changes resulting from this policy, which we refer to as policy B, are illustrated by the black crosses in the three panels of Figure 5. The left panel shows that the policy is highly effective in reducing rents in the bottom rental segment. Indeed, the expanded supply of these more affordable units allows a substantial share of households previously forced to reside in the "homelessness" segment to move to higher-quality rental housing. There are moderate spillovers of the supply expansion to the top rental segment and the lowest owner-occupied segment, lowering prices in those segments slightly. However, by shrinking the supply of high-quality owner-occupied

units, this policy causes prices in the upper quality segments to rise. This effect may seem counterintuitive at first, but it is a robust outcome of the economic logic embedded in the model. When we shrink housing supply at the very top end of the housing market, the marginal household that is indifferent between the top-most and the second-best segment is now wealthier. By reducing the quantity of luxury housing, wealthier households are forced into the neighboring lower-quality segments, where they drive up prices. The black dots in the other two panels of Figure 5 show how these price changes reduce the rent/income ratio of renters, but raise price/income ratios of homeowners.

Experiment 2: Adding Housing Supply at the Top

In our second experiment, the newly built housing units are in the top segment of the owner-occupied market. In this experiment, we further increase supply in the top segment by converting 10 percent of the rental units in the bottom segment of the rental market to units in the top housing segment. At an intuitive level, this conversion can be thought of as investing in the quality of the existing housing stock. The changes caused by this policy, which we refer to as policy H, are illustrated by the green plus signs in Figure 5.

The left panel shows, perhaps surprisingly, that rents and prices drop in all segments due to this policy. In particular, expanding high-quality units is as effective at making rental units more affordable as expanding the supply of low-cost rental units. However, this policy also leads to large reductions in house prices in the topmost segments, and moderate declines in middle segments. By raising supply in the highest segment, we prevent the wealthiest households from competing for units in lower segments, which in turn reduces demand in all lower segments, causing a "trickle-down" effect. The middle and right panel show that price/income ratios and rent/income ratios decline substantially. This result contrasts with the previous experiment, where price/income ratios increased. In short, increasing the supply of housing in the top segment is more effective at reducing house-price-to-income ratios than adding supply in the bottom rental segment.[12]

Experiment 3: Mortgage Availability

Another policy that is often proposed by policymakers is subsidizing homeownership, for example by providing down-payment assistance for first-time buyers or by lowering property taxes. In this section, we evaluate the effects of these policies by simulating an economy where owners are granted $100,000 towards the purchase of their house. The main result is that subsidizing ownership without building more houses unintentionally worsens housing affordability in the city. The

---

[12] Increasing housing supply might lead to an influx of migrants from other cities (French and Gilbert 2023). In the model, because the mass of residents is normalized to 1, population growth corresponds to a drop in housing supply. Such in-migration would therefore dampen the effect of new construction on prices and incumbents' welfare.

changes caused by this policy, which we refer to as policy M, are illustrated by the blue asterisks in Figure 5.

The left panel plots the price increases resulting from this pure demand-side policy. The economic intuition is straightforward: transferring wealth to home-buyers increases demand without increasing supply. The consequence are higher prices across all owner-occupied segments. Perhaps surprisingly, subsidizing homeownership through down-payment assistance also makes rental housing more expensive. The intuition is that increasing demand in upper segments of the housing market trickles down to lower segments. By effectively making the marginal home buyer wealthier, the down-payment subsidy also affects the wealth of the marginal household who is indifferent between renting in the top rental segment and buying in the lowest owner segment, which translates to higher rents.

As we can see in the middle and right panels of Figure 5, price-to-income and rent-to-income ratios rise unanimously as result of a subsidized down-payment policy. The lesson is that only policies that raise supply (or decrease demand) will make housing more affordable. Boosting demand for owner-occupied housing through subsidies will boost demand and raise prices, unless the policy also increases supply at the same time.

Computing Welfare

The model allows for a direct welfare comparison of the three potential policies. As illustrated in Figure 5, each of the proposed policies affects equilibrium house prices and rents. This means that household utility, which increases with the quality of the house a household lives in as well as with the amount of other consumption goods, is also affected by these policies. For example, all else equal, a policy that drives households to worse quality housing lowers utility and dampens welfare. To compare the welfare effects of the different policies, we compute a measure called "compensating variation," which is the amount of wealth a household would need to be paid under each policy so that it has equal utility relative to the base-line economy. We display the measures of compensating variation implied by the proposed policies in Figure 6, with compensating variation in absolute dollars on the left and as a share of wealth on the right.

The down-payment subsidy performs worst, because it makes housing uniformly more expensive and therefore lowers households' utility. This translates to a negative compensating variation for all households (blue asterisks). Adding supply to the bottom segment (black crosses) leads to positive welfare gains for relatively poor households but to welfare losses for relatively rich households. For example, the compensating variation is 35 percent of wealth for households in the second percentile of wealth but is negative 3 percent for the richest households. In contrast, adding supply to the top (green plus signs) leads to similarly large welfare gains for poor households, but also improves welfare of rich households. Thus, within our model, this policy is a rare case of a so-called "Pareto superiority," implying that one policy change makes everyone involved better off relative to an alternative policy.

Figure 6
Measuring Welfare
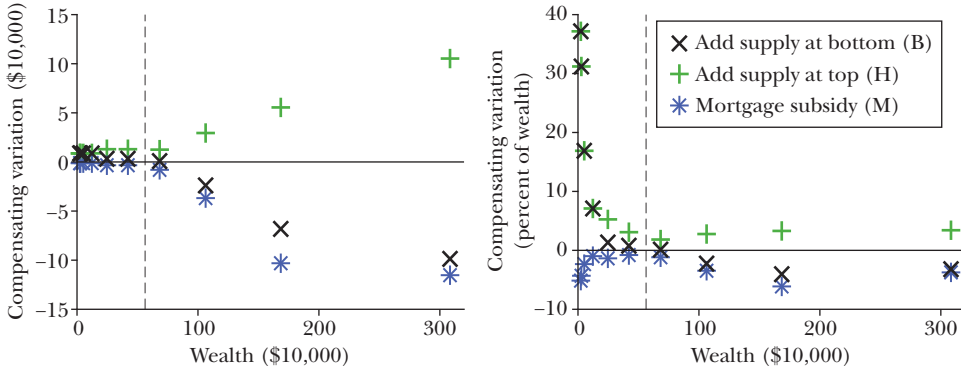
Note: This figure shows the average compensating variation that households would need to receive in the baseline 2022 model in order to be as well off as under each of the proposed policies. A positive compensating variation means that the household is better off under the policy relative to the baseline. The estimates are binned by wealth decile. The left panel shows the compensating variation in 2013 dollars, and the right panel as a percent of household wealth.

Underlying Assumptions

Our calibrated model is complex enough to capture key heterogeneity in household wealth and in housing qualities, yet simple enough to yield clear insights. It reflects how interventions in one segment of the housing market affect the entire market. However, as is the case with most models, the model has limitations due to simplifying assumptions.

Notably, it lacks realistic transition dynamics, operating over just two periods—today and tomorrow. This makes it best suited for analyzing the long-run effects of persistent housing policies, which are slow to implement but have lasting impacts. The model assumes instantaneous housing reallocation as higher-wealth households move into newly supplied higher-quality segments, freeing up units for others. In reality, this process may unfold over years, especially until it reaches the lower end of the market, which is crucial for improving rental affordability. A mitigating factor is that prices typically adjust quickly in anticipation of future supply shifts, potentially reducing housing costs before the long-run equilibrium is reached.

Second, the model omits capital gains from home sales. It assumes lower house prices are always beneficial, as they reduce the cost of housing services. In reality, major shifts in household-unit assignments involve widespread buying and selling. For existing homeowners, falling prices due to increased supply can mean financial losses, helping explain resistance to new construction. Incorporating dynamics into the model would allow us to capture these capital gains and losses. Some studies already explore how policy reforms affect capital gains for heterogeneous

households: useful starting points are Floetotto, Kirker, and Stroebel (2016) and
İmrohoroğlu, Matoba, and Tüzel (2018). Taking these effects into account could
change the compensating variation number presented in Figure 6.

Third, our model does not include a construction sector. In reality, increasing
housing supply requires developers and builders to be compensated at market
rates. By abstracting from construction costs, the model assumes uniform costs
across housing segments. However, evidence from affordable housing mandates
and rent control shows that developers tend to favor high-end construction.
Without regulation, new building would likely occur in the middle- or upper-market
segments—consistent with our policy simulations. In contrast, mandates to build
lower-quality housing at below-market prices may distort incentives and reduce
overall supply over time (Diamond, McQuade, and Qian 2019; Krimmel and Wang
2023).

A further implication of excluding the construction sector is our assumption
that supply remains fixed when analyzing mortgage subsidies. In practice, rising
prices often prompt developers to build more, but this response varies by city due
to natural and regulatory constraints. Supply elasticity captures this variation, with
places like San Francisco exhibiting very low responsiveness. Saiz (2010) estimates a
medium- to long-run elasticity of 0.66 for San Francisco, while Baum-Snow and Han
(2024) estimate a lower long run elasticity of 0.27,[13] supporting our fixed-supply
assumption. Additionally, our model does not consider how mortgage subsidies are
funded. Because they likely draw from limited tax revenues, accounting for funding
costs would further reduce the policy's net welfare benefits.

A final caveat is that our model analyzes the housing market of a city in isola-
tion. Our approach accounts for realized changes in population size over time, but
not endogenous in- or out-migration in response to housing becoming cheaper or
more expensive. Any policy that lowers housing costs in San Francisco will attract
new residents, which will increase housing pressures again.

Conclusion

We use an equilibrium model of a local housing market to examine the causes of
and potential solutions to housing unaffordability. Our analysis indicates that effec-
tive policies must increase housing supply. While supply interventions may target
specific market segments, their effects extend across the entire housing market.
Notably, expanding high-end housing supply reduces competition for mid- and
lower-tier units, improving affordability more broadly. A straightforward long-term
strategy is for local governments to permit and even support developers in building
the housing they find most profitable—often at the higher end of the market.

---

[13] Baum-Snow and Han (2024) estimate supply elasticities at the level of census tracts. We obtain the
aggregate elasticity by averaging across their tract-level estimates. The elasticities measure the percentage
change in supply for a 1 percent increase in housing prices.

As for any structural analysis, our conclusions depend on the validity of our modeling assumptions. In reality, household decisions are shaped by more than wealth, and housing quality is not one-dimensional. A more complex model could incorporate these nuances and address some of the omissions we discussed above, potentially altering some findings. Still, our framework likely captures the key economic mechanisms that would persist even in a richer model.

References

Abramson, Boaz. 2021. "The Equilibrium Effects of Eviction Policies." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4112426.

Abramson, Boaz, and Stijn Van Nieuwerburgh. 2024. "Rent Guarantee Insurance." NBER Working Paper 32582.

Abramson, and Tim Landvoigt. 2025. *Data and Code for: "Curbing Rising Housing Costs: A Model-Based Policy Comparison."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E2334017V1.

Anenberg, Elliot, and Edward Kung. 2020. "Can More Housing Supply Solve the Affordability Crisis? Evidence from a Neighborhood Choice Model." *Regional Science and Urban Economics* 80: 103363.

Asquith, Brian J., Evan Mast, and Davin Reed. 2023. "Local Effects of Large New Apartment Buildings in Low-Income Areas." *Review of Economics and Statistics* 105 (2): 359–75.

Baum-Snow, Nathaniel, and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–946.

Baum-Snow, Nathaniel, and Justin Marion. 2009. "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods." *Journal of Public Economics* 93 (5–6): 654–66.

Board of Governors of the Federal Reserve System. 2013-2022. *Survey of Consumer Finances (SCF).* https://www.federalreserve.gov/econres/scfindex.htm (accessed June 18, 2025).

Corbae, Dean, Andrew Glover, and Michael Nattinger. 2024. "Equilibrium Evictions." NBER Working Paper 32898.

Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst. 2023. "Income Growth and the Distributional Effects of Urban Spatial Sorting." *Review of Economic Studies* 91 (2): 858–98.

Diamond, Rebecca, and Tim McQuade. 2019. "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development." *Journal of Political Economy* 127 (3): 1063–117.

Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." *American Economic Review* 109 (9): 3365–94.

Díaz, Antonia, and Belén Jerez. 2013. "House Prices, Sales, and Time on the Market: A Search-Theoretic Framework." *International Economic Review* 54 (3): 837–72.

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh. 2023. "Affordable Housing and City Welfare." *Review of Economic Studies* 90 (1): 293–330.

Favilukis, Jack, and Stijn Van Nieuwerburgh. 2021. "Out-of-Town Home Buyers and City Welfare." *Journal of Finance* 76 (5): 2577–638.

Floetotto, Max, Michael Kirker, and Johannes Stroebel. 2016. "Government Intervention in the Housing Market: Who Wins, Who Loses?" *Journal of Monetary Economics* 80: 106–23.

Fonseca, Julia, Lu Liu, and Pierre Mabille. 2025. "Unlocking Mortgage Lock-In: Evidence from a Spatial Housing Ladder Model." Working Paper.

French, Robert, and Valentine Gilbert. 2023. "Suburban Housing and Urban Affordability: Evidence from Residential Vacancy Chains." Unpublished.

Garriga, Carlos, and Aaron Hedlund. 2020. "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession." *American Economic Review* 110 (6): 1603–34.

Glaeser, Edward L., and Bryce A. Ward. 2009. "The Causes and Consequences of Land Use Regulation: Evidence from Greater Boston." *Journal of Urban Economics* 65 (3): 265–78.

Gupta, Arpit, Christopher Hansman, and Pierre Mabille. 2024. "Financial Constraints and the Racial Housing Gap." INSEAD Working Paper 2022/58/FIN.

Hsieh, Chang-Tai, and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* 11 (2): 1–39.

İmrohoroğlu, Ayşe, Kyle Matoba, and Şelale Tüzel. 2018. "Proposition 13: An Equilibrium Analysis." *American Economic Journal: Macroeconomics* 10 (2): 24–51.

İmrohoroğlu, Ayşe, and Kai Zhao. 2022. "Homelessness." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4308222.

Krimmel, Jacob, and Betty Wang. 2023. "Upzoning with Strings Attached: Evidence from Seattle's Affordable Housing Mandate." *Cityscape* 25 (2): 257–78.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. "The Housing Market(s) of San Diego." *American Economic Review* 105 (4): 1371–407.

Määttänen, Niku, and Marko Terviö. 2014. "Income Distribution and Housing Prices: An Assignment Model Approach." *Journal of Economic Theory* 151: 381–410.

Mehra, Rajnish, and Edward C. Prescott. 1985. "The Equity Premium: A Puzzle." *Journal of Monetary Economics* 15 (2): 145–61.

Mense, Andreas. 2025. "The Impact of New Housing Supply on the Distribution of Rents." *Journal of Political Economy Macroeconomics* 3 (1): 1–42.

Nathanson, Charles G. 2023. "Trickle-Down Housing Economics." Unpublished.

Nikolakoudis, George. 2024. "The Economics of Segmented Housing Markets." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4251925.

Ortalo-Magné, François, and Sven Rady. 2006. "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." *Review of Economic Studies* 73 (2): 459–85.

Pennington, Kate. 2021. "Does Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.3867764.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. 2020. "Segmented Housing Search." *American Economic Review* 110 (3): 720–59.

Rubin, Donald B. 1986. "Statistical Matching Using File Concatenation with Adjusted Weights and Multiple Imputations." *Journal of Business and Economic Statistics* 4 (1): 87–94.

S&P Dow Jones Indices LLC. 2013-2022. *S&P CoreLogic Case-Shiller CA-San Francisco Home Price Index [SFXRSA].* https://fred.stlouisfed.org/series/SFXRSA (accessed June 18, 2025).

Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3): 1253–96.

US Census Bureau. 1980–2022. *American Community Survey (ACS).* https://usa.ipums.org/usa-action/variables/group (accessed June 18, 2025).

The Folk Economics of Housing

Christopher S. Elmendorf, Clayton Nall, and Stan Oklobdzija

Restrictions on housing development are kaleidoscopic in their variety. Local governments set minimum lot sizes. They cap dwelling-unit densities and building heights. They control project designs with arcane and subjective review procedures. They designate zones as historic or environmentally sensitive. They condition project approvals on the developer paying for parks, infrastructure, and affordable housing. Intentionally or inadvertently, state and local governments may also reduce the financial feasibility of development or increase investors' uncertainty with real-estate transfer taxes, landlord-tenant regulations, building-code requirements, and construction-defect liability rules. Cumulatively, the regulatory morass both limits the amount of housing that may be built within a given distance of city centers and increases the cost of building whatever housing is nominally allowed.

About two decades ago, prominent economists began raising alarms about the effect of severe land-use restrictions on housing affordability (Glaeser and Gyourko 2002; Glaeser, Gyourko, and Saks 2005). Subsequent work finds that housing supply elasticities have been declining over the last several decades and are near-zero in many suburban communities (for example, Baum-Snow 2023; Baum-Snow and

■ *Christopher S. Elmendorf is Martin Luther King Jr. Professor of Law, University of California—Davis, Davis, California, and Nonresident Senior Fellow, Institute for Progress, Washington, DC. Clayton Nall is Associate Professor of Political Science and Affiliate, Department of Geography, University of California—Santa Barbara, Santa Barbara, California. Stan Oklobdzija is Assistant Professor, School of Public Policy, University of California—Riverside, Riverside, California. Their email addresses are cselmendorf@ucdavis.edu, nall@ucsb.edu, and stano@ucr.edu.*

Han 2024; Orlando and Redfearn 2024). Unsurprisingly, localities with more severe land-use regulations have less elastic housing supply (Bartik, Gupta, and Milo 2025).

A recent study comparing multifamily housing projects in heavily regulated California cities with similar projects in more lightly regulated Texas finds that "production costs per net rentable square foot for market-rate housing in California are 2.3 times the average cost in Texas" (Ward and Schlake 2025). In the median of California's urban and suburban jurisdictions, 84 percent of residential land is zoned exclusively for single-family homes (Menendian, Nawaz, and Gambhir 2024). A national study of municipal land-use controls found that between 2006 and 2018, no highly restrictive jurisdiction relaxed its controls, whereas numerous jurisdictions became more restrictive (Gyourko, Hartley, and Krimmel 2021).

Land-use regulations can mitigate negative externalities from development, but the available evidence suggests that the social costs of extant restrictions substantially outweigh their benefits (for example, Albouy and Ehrlich 2018). So why do the regulations persist? The standard political-economy explanations focus on homeowner self-interest and jurisdictional fragmentation (Fischel 2001). Homeowners are thought to oppose development because of fears that it may reduce the value of their property. Small local governments tend to regulate too stringently because the costs of foregone housing development are externalized onto nonresidents, who suffer higher housing costs and reduced access to productive economic agglomerations.

This paper suggests another, complementary explanation for the over-regulation of metropolitan land use: the failure of ordinary voters to appreciate that large, exogenous increases in regional housing stocks would put downward pressure on prices, coupled with voters' tendency to blame housing providers (developers and landlords) for high prices. Using original data from several preregistered, nationally representative surveys of residents of urbanized zip codes, we find that the mass public's views about the price effects of housing supply shocks are weak and unstable—significantly more so than their views about the price effects of supply shocks in other markets. Moreover, while ordinary people have little feel for whether a big, positive shock to their metro region's housing supply would make housing more affordable, they do have a clear set of "folk economics" beliefs about who is to blame for the high cost of housing—and some strongly-held preferences concerning what should be done about it. The survey data suggest that developers and landlords bear the brunt of the blame. Price controls, restrictions on big investors, and mandates that housing developers provide below-market-rate units are the preferred fixes. Such policies are already well established in major cities with high housing demand (Bartik, Gupta, and Milo 2025), and their spread to other places would probably make housing supply even less elastic.

We begin this essay by outlining the "homevoter" and jurisdictional fragmentation theories of the housing shortage, along with certain facts that these theories struggle to explain. Next, we discuss the theoretical considerations that led us to expect that voters would have difficulty with supply-and-demand reasoning in the housing context. Then we present the main findings from our surveys. We close by discussing the implications of our findings for policymakers and researchers.

The Homevoter and the Jurisdictional Fragmentation Hypotheses

Why do local governments regulate housing development too stringently? The *homevoter hypothesis* emphasizes asymmetries of interest and organizational capacity between homeowners and renters (Fischel 2001; Ortalo-Magné and Prat 2014). The theory was originally developed to explain the actions of small local governments with respect to proposed developments. Risk-averse homeowners, who have a substantial share of their wealth tied up in their home, were said to lobby against nearby development because of fears that associated disamenities such as noise, traffic, and loss of views would be capitalized into the value of their home. Renters, who in principle would benefit from new supply, have little incentive to mobilize for any one development because their utility is determined by housing prices across the full set of neighborhoods where they might live, which no single project stands to affect (Muth 1969; Alonso 1964; Rosen 1979; Roback 1982). Homeowners are also better positioned than renters to overcome barriers to collective action, owing to their generally longer tenure in a community (Acolin 2022) and membership in homeowners' associations (Clarke and Freedman 2019).

The implications of the homevoter hypothesis for regulation by regional or statewide actors are less clear. Some scholars have posited that because demand for housing is less elastic at larger geographic scales, homeowners given the opportunity to regulate regionally would act as a cartel, restricting supply in order to drive up prices (for example, Ellickson 1977). Other scholars have developed models in which a regional regulator internalizes costs of municipal land-use controls that would be borne by homeowners in other municipalities, with the result that regional regulation is more accommodating of housing development (Favilukis and Song 2023).

The latter argument exemplifies the *jurisdictional fragmentation hypothesis*, which holds that over-regulation of land use results from geographic mismatches between the costs and benefits of housing development and the governmental units responsible for land-use policy. When a metropolitan region has been fragmented into numerous local governments, a locality can severely restrict housing or commercial development and thereby avoid local disamenities from new construction (traffic, noise, loss of open space or views) without materially limiting the access of its incumbent residents to labor markets and cultural amenities created by the region's economic agglomeration (Marantz and Lewis 2022; Favilukis and Song 2023). A small local government has no incentive to consider how its own restrictions limit the value of the regional agglomeration for residents of other municipalities. Locally optimal land-use regulations are thus even more restrictive than the metro area's median homeowner would prefer. Other political-fragmentation accounts reveal a similar dynamic among metropolitan regions (for example, Duranton and Puga 2023; Parkhomenko 2023).

The homevoter and jurisdictional fragmentation hypotheses offer valuable insights. An index of jurisdictional fragmentation can explain nearly 20 percent of the variance across US metro regions in minimum lot size requirements (Favilukis

and Song 2023), although the causal arrow could also run from popular demand for lot-size controls to greater fragmentation. In France, forced consolidations of small municipalities into larger planning units increased housing production (Tricaud 2025). Formal models of endogenous land-use regulation by local governments have been shown to embody key facts about urban housing markets, such as the persistent wedge, at the perimeter of high-demand metro areas, between housing prices and the cost of construction (Duranton and Puga 2023; Parkhomenko 2023).

However, other facts about housing regulation and public opinion have proven harder to explain within these frameworks. A literature review concluded a decade ago that there is "little empirical evidence that areas with more homeowners adopt stricter housing supply regulations" (Gyourko and Molloy 2015, p. 1308). Results since then have been mixed, with some papers finding little or no correlation between homeownership and enactment of restrictive land-use regulations or inelastic housing supply (Orlando and Redfearn 2024; Tricaud 2025), and others finding evidence more in line with the homevoter hypothesis (Khan 2021; Fang, Stewart, and Tyndall 2023).

Some survey-based work by scholars in political science finds that owners and renters have similar views about development and land-use issues once respondents are matched on key demographics, suggesting that homeownership itself is not the driving factor behind opposition to development (Larsen and Nyholt 2023); other surveys find modest differences (Marble and Nall 2021; Broockman, Elmendorf, and Kalla 2024). Even without adjusting for demographics, homeowners and renters are broadly and about equally supportive of dense housing development in commercial areas and transit corridors (Kansal, Aurand, and Saadian 2024). Homeownership probably does make people somewhat more politically attuned to localized amenity effects associated with land-use changes in their neighborhood (Ahlfeldt and Maennig 2015; Hall and Yoder 2022), but there is little if any evidence that homeowners are less supportive than demographically similar renters of measures to increase the supply of housing citywide or regionally. Indeed, in a study of state-level policy preferences, Elmendorf, Nall, and Oklobdzija (2024b) find that homeowners are *more* supportive than renters of their state allowing more market-rate homes to be built on open land, and about equally supportive of allowing more market-rate apartments in existing neighborhoods.

Another puzzle is why the land-use policies of fragmented municipalities remain uncoordinated, notwithstanding the existence of national and state governments through which the municipal collective-action problem could be solved. If homevoter-controlled local governments have adopted more restrictive regulations than the median homeowner in a metro region prefers, we should expect homeowners and their agents in local government to push state legislatures to loosen the fetters. Yet when YIMBY-aligned ("Yes In My Backyard") legislators introduce bills to liberalize housing supply, local governments regularly line up against them (Stahl 2018). In conversations with housing advocates, we have never heard of municipal lobbies or homeowner associations leading the charge for expanded housing supply.

Finally, neither the homevoter nor the jurisdictional fragmentation hypothesis explains why government actors have so commonly responded to concerns about housing affordability with policies that regulate prices. Substantial federal tax expenditures fund the construction of price-restricted affordable housing. Local governments and an increasing number of states also demand or incentivize the inclusion of price-restricted units in new developments. Several of the parent states of large, high-cost US cities make local governments plan for deed-restricted affordable housing (Elmendorf 2019). In big cities, a cadre of "supply skeptics" argue that new market-rate housing development actually increases rents for existing housing, and that potentially developable sites should therefore be taken off the market and reserved for future social housing with price caps and income-based occupancy restrictions (Been, Ellen, and O'Regan 2019). To be sure, such narratives may just be window dressing for the material interests of neighborhood groups that seek side payments from developers or greater control over developable sites. But that explanation raises the question of why the groups justify their demands in this way.

When Economic Research Meets Rationally Ignorant Voters

One plausible explanation for the similarity of land-use preferences between homeowners and renters, for the absence of a bottom-up groundswell for states to solve local governments' collective-action problem with respect to housing supply, and for the prevalence of price controls in affordable-housing policy, is that ordinary voters simply do not perceive any connection between local constraints on housing supply and the equilibrium price of housing.

The thesis that ordinary voters are rationally ignorant about many questions of public policy and adopt political views that bear little relation to their economic self-interest has an illustrious lineage (for example, Downs 1957; Citrin and Green 1990; Caplan 2011). And there are good reasons to think that questions about the effects of regional or statewide land-use liberalization on housing prices are especially tough for voters to handle.

Laypeople may roughly intuit how local amenities or disamenities from specific development projects would affect nearby property values, but it is much more challenging to envision how tens of thousands of land-use decisions by numerous local governments cumulate into regional shortages, or how changes in the size of a metro region's housing stock in turn affect housing prices throughout the region and beyond. The underlying economics are not obvious. In simple models with homogeneous agents and costless mobility between places, demand for housing in any given place is perfectly elastic (Rosen 1979; Roback 1982). Local regulatory reforms that increase housing productivity in one jurisdiction will be capitalized into the price of land there but will not affect housing prices (Albouy and Ehrlich 2018). Relaxing the assumption of costless mobility or identical preferences for places generates the familiar downward-sloping demand curve, such that a positive supply shock in a jurisdiction will reduce housing rents locally (Arnott and Stiglitz 1979; Hilber 2017).

Recent empirical studies demonstrate that plausibly exogenous positive shocks to metro-area housing supply do in fact substantially reduce prices within the area. Exploiting weather-related shocks to housing supply in Germany, Mense (2023) finds that "[i]ncreasing the flow of new supply [in a city] by 1 percent lowers average rents by 0.2 percent, with similar effects across all market segments." A major citywide upzoning in Auckland, New Zealand, induced a development boom and lowered rents by 20–30 percent relative to a synthetic control city (Greenaway-McGrevy 2023). A broad upzoning of São Paulo, Brazil, led to a 1.4 percent increase in new housing construction and a 0.4 to 0.9 percent decline in rents (Anagol, Ferreira, and Rexer 2023).

Studies of interregional migration in response to negative employment shocks have tended to find that mobility is lower than economists had expected (Autor, Dorn, and Hanson 2013; Greenland, Lopresti, and McHenry 2019; Yagan 2019; Choi et al. 2024). Similarly, US metro areas with greater regulatory barriers to new supply experienced less housing-stock growth and less in-migration, but no more out-migration, than less constrained metros from 1990 to 2000 (Molloy, Nathanson, and Paciorek 2022). Such frictions mean that regional housing supply shocks will have larger effects on regional than on national prices.

But again, laypeople have no reason to know of these research studies. Nor do ordinary people have first-hand experience with large, sudden shocks to their metro-region's housing supply. Regional housing stocks have changed slowly, especially in recent years. Between 2008 and 2023, the total housing stock in the United States increased at an annualized rate of only 0.7 percent, with much of the growth concentrated in only a few active markets (US Census Bureau and US Department of Housing and Urban Development 2023a). Meanwhile, national median home prices have risen consistently except during recessions (US Census Bureau and US Department of Housing and Urban Development 2023b). It would be no great leap for the layperson to infer that only macroeconomic factors and financial markets, and not local supply-side factors, such as land-use restrictions, determine housing prices.

To the extent that ordinary people form loose mental associations between "housing development" and "housing affordability," they may well associate more development with higher prices rather than greater affordability. New housing, being new, tends to be more expensive than existing, depreciated housing. Ordinary people may also observe neighborhood-level correlations between new development and prices for existing housing going up. This will occur if developers strategically choose to build where they anticipate rising demand, or if new construction has local amenity effects that raise the price of nearby housing. Such amenity effects have been documented in poor neighborhoods with deteriorated housing stocks (Diamond and McQuade 2019; González-Pampillón 2022), although most studies find that new development in places favored by market-rate developers does not increase rents nearby (Pennington 2021; Li 2022; Asquith, Mast, and Reed 2023; Büchler and Lutz 2024).

The urbanist orientation of the nascent Yes-In-My-Backyard pro-housing movement, which advocates for allowing apartment buildings in more places, may

also contribute to a sense that YIMBY policies would be futile. Tall buildings are more costly to construct on a per square foot basis and thus only observed where rents are high. If ordinary voters are loosely aware that places like San Francisco, New York, and Los Angeles have tall buildings and high housing prices, they may balk at YIMBY claims that allowing more "urban" forms of housing in the suburbs would make the suburbs more affordable. As a recent ethnography concludes, "densification policies feel disconnected from suburban residents' . . . images of the housing crisis" (Ternullo 2024, p. 25).

Finally, people may resort to what Rubin (2003) calls "folk economic" explanations for high prices. Ordinary people tend to attribute economic outcomes they dislike to conniving bad actors, rather than impersonal market forces (Boyer and Petersen 2018). Real estate developers are an easy target, and landlords make good villains. too. If housing providers are bad and rising prices are bad, then, the folk economist reasons, price gouging and cartelization must be a big part of the reason why prices are going up.

We are certainly not the first to suggest that economic misperceptions may be particularly important to housing policy. For example, Hilber and Schöni (2022, p. 25) speculate that the tendency of advanced democracies to pursue housing affordability though counterproductive policies such as price controls and home-ownership subsidies might be a consequence of voters' "limited knowledge of . . . the housing market" and "poor anticipation of . . . indirect, and often unintended, policy effects." But hard evidence on laypersons' understanding of housing markets has become available only recently. In the next section, we report results from several original surveys of urban and suburban residents of the United States in which we elicited predictions about the effects of supply shocks in housing and other markets.

Survey Evidence of Wants and Beliefs about Housing Prices

We conducted three preregistered online surveys with the goal of understanding (1) whether there is potential majority support for reducing housing prices by increasing supply, notwithstanding that homeowners outnumber renters and are overrepresented among likely voters; (2) whether most people believe that a large, positive shock to their metro region's housing stock would result in lower housing prices; and (3) what factors are correlated with housing "supply skepticism," which we operationalize as the belief that a sizeable regional supply shock would not lower prices.[1]

The surveys were fielded on online samples provided by Forthright, a leading vendor, targeting adult residents of US urban and suburban zip codes, identified by using block groups to calculate population-weighted density (Manson et al. 2024). We directed the survey vendor to maintain equal proportions of homeowners and

---

[1] The preanalysis plans are available at osf.io/96crq (Survey 1), osf.io/e4wr6 (Survey 2), and osf.io/zujyr (Survey 3).

renters in the sample and to balance on age, race, and gender using the vendor's nationally representative population quotas.

All three surveys featured an initial question asking whether the respondent would prefer future home prices and rents in their city to be higher, lower, or the same as today's. Each survey also elicited predictions of the effect of a large, positive regional supply shock on future home prices and rents in the respondent's city. Across the surveys, we then asked different questions about hypothesized correlates of housing supply skepticism. In Surveys 2 and 3, we also asked about the effects of supply shocks in various non-housing markets and sought to assess the public's mental models of housing markets. Finally, on Survey 3, we retested respondents within the survey in order to evaluate the stability of housing-market beliefs relative to analogous beliefs about other markets. The surveys were fielded in March 2022, August 2022 and May 2023. Details of the survey designs and additional, more detailed findings are provided in the Supplemental Appendix.

People Say They Want Lower Housing Prices

The homevoter hypothesis supposes that homeowners want higher housing prices in their local area, whereas renters prefer the opposite. We asked respondents: "Think about possible futures for [name of the respondent's city or town inserted here]. Would you prefer home prices and rents to be higher, lower, or the same as today? Assume that the economy and quality of life would stay the same in other respects. Only home prices and rents would change." We included the other-things-equal condition because we wanted to distinguish preferences with respect to home prices *as such* from preferences with respect to local amenities.

As Figure 1 shows, more than 55 percent of homeowners say they would prefer lower housing prices; unsurprisingly, 85 percent of renters say the same.[2] To be sure, homeowners' stated support for lower prices could be inflated due to "social desirability bias"—in this case, wishing not to appear greedy. Homeowners might also have conflicting desires regarding future housing prices and rents, such as preferring lower rents but higher home prices, a nuance not captured by our question. Or they may have an abstract preference for lower prices that evaporates once they observe a specific threat to their own property's market value, or are told that additional development would have local disamenity effects (contrary to the *ceteris paribus* premise of our question). Even allowing for these caveats, that well over half of urban and suburban American homeowners endorse the goal of lower home prices and rents in their local area suggests that they are not predestined to block state or national supply-side housing policies.

---

[2] This and analogous statements in the text are slight oversimplifications, as the figures in this paper pool data from the surveys and some respondents took more than one survey. The technically precise but cumbersome way to describe the result in Figure 1 would be to say, "more than 55 percent of responses from homeowners state a preference for lower housing prices; unsurprisingly, 85 percent of responses from renters say the same." In all figures, standard errors are clustered on the respondent identifier.

Figure 1
Desired Future Home Prices and Rents in Respondent's City
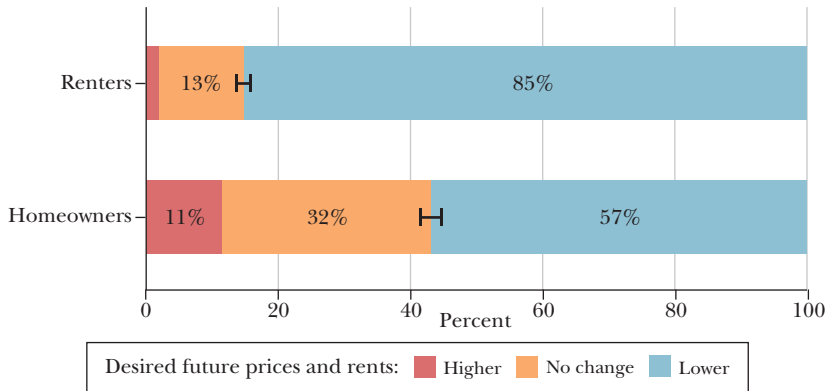
Source: Elmendorf, Nall, and Oklobdzija (2025).
Note: The question asks the respondent to assume no change in the economy or quality of life. This figure pools responses from Surveys 1, 2, and 3. Error bars are 95 percent confidence intervals on the proportion of responses stating a preference for lower prices, with standard errors clustered on the respondent identifier (some individuals took more than one of the surveys). Results by survey are provided in Supplemental Appendix 3.

Beliefs about Housing and Other Markets

While most respondents say they would like lower housing prices in their community, they do not believe that an exogenous positive shock to regional housing supply would reduce prices. Each survey presented scenarios in which land-use deregulation or better home-building technology significantly increases the housing stock of the respondent's metro region (usually specified as a 10 percent increase over five years). We asked respondents to provide their best guess about the effect of the shock on future prices and rents for existing homes in their city.[3]

Figure 2 reports the distribution of respondents' directional predictions about the price effects of supply shocks in markets for houses, cars, agricultural commodities, labor, and consumer goods. Only one-quarter to one-third of respondents predicted that the housing supply shock would reduce prices or rents for existing homes, whereas one-third to one-half said the shock would lead to higher prices. Homeowners and renters made similar predictions; housing supply skepticism does not appear to be a rationalization of homeowner self-interest.[4]

---

[3] As a design-based robustness check, Surveys 2 and 3 randomized numerous aspects of the price-prediction question, including the cause of the supply shock, counterfactual future prices, question complexity, and response format. These variations proved largely inconsequential. The prices used in experimental treatments were derived using data from the Zillow home-value index and observed-rent indexes and Apartment List (Apartment List Research Team 2022; Missouri Census Data Center 2022; Pareto Software 2022; Zillow Group 2022). See Supplemental Appendix 5.

[4] For details, see Supplemental Appendix 6.2.

Figure 2
Beliefs about Price Effects of Housing versus Non-housing Supply Shocks
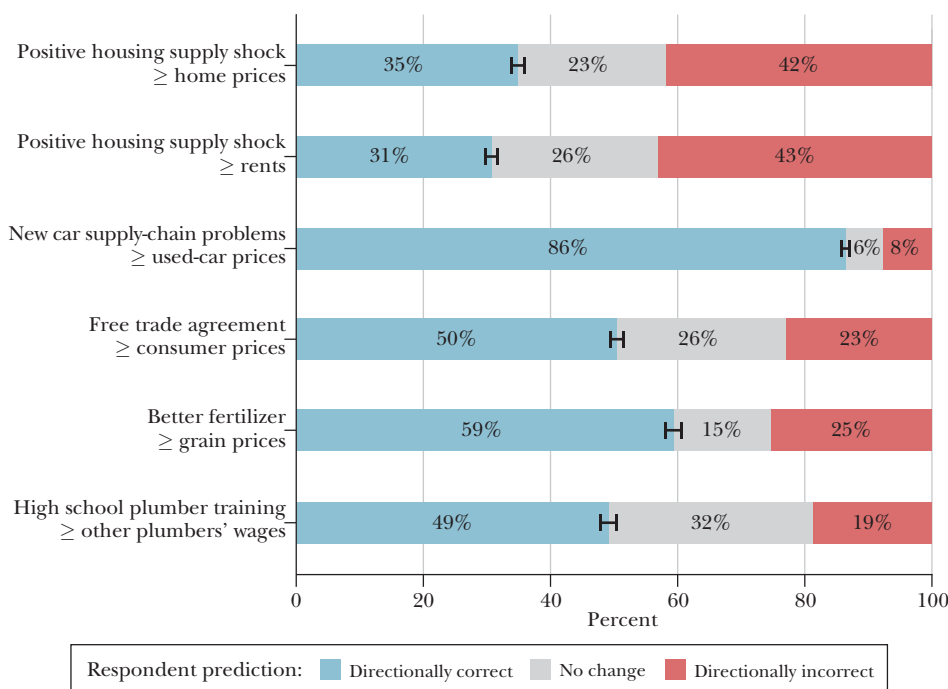
Source: Elmendorf, Nall, and Oklobdzija (2025).
Note: This figure pools responses from Surveys 1–3. Error bars are 95 percent confidence intervals on proportion of responses that are not directionally correct, with standard errors clustered on the respondent. Results disaggregated by survey are provided in Supplemental Appendix 3. Some of the survey questions elicited predictions about the magnitude of the price effect, but for purposes of this figure, we standardize responses to simple directional predictions (increase, decrease, or have no effect), aligning with our preanalysis plans.

Supply skepticism is, however, distinctive to housing. Our respondents largely manifested conventional economic views about the effect of free-trade agreements on consumer-good prices, the effect of yield-improving fertilizers on grain prices, and the effect of a citywide high school plumber-training program on the wages of existing plumbers. Fifty to sixty percent of respondents ventured that these positive supply shocks would reduce prices in the market in question; fewer than 25 percent predicted higher prices.

Even more striking, when asked how supply-chain problems in the market for new cars affect the price of used cars, 85 percent of respondents predicted higher used-car prices. The connections between new-car supply and used-car prices appear easy for people to grasp, in contrast to the connections between new-house supply and used-house prices. Our surveys were fielded in the wake of widely-reported, COVID-era supply-chain problems in the new-car market, which led to a sharp

jump of 45 percent in used-car prices (Boudette 2021). Perhaps recent exposure to supply-and-demand logic in a specific market leads people to adopt the "Econ 101" mentality with respect to that market, if not to generalize beyond it (Cosmides and Tooby 1992).

What Explains Housing Supply Skepticism?

Across the surveys, we asked numerous questions probing for beliefs, abilities, interests, and experiences that may correlate with housing supply skepticism and explain its prevalence. Might supply skepticism be an outgrowth, for example, of motivated reasoning, zero-sum thinking, lack of numeracy or education, exposure to development in places where prices and rents are going up, or beliefs about the neighborhood-level amenity effects of housing development and how those effects manifest in prices locally?

Several of the correlations run in the expected direction, but they are weak, typically in the range of $r = 0.1$ to 0.2.[5] The weakness of these correlations led us to investigate whether housing supply skepticism might be what survey researchers call a "non-attitude"—an opinion made up on the spot by someone who does not actually have a firm opinion on the topic—rather than a settled belief (Campbell et al. 1960). Two pieces of evidence bear out the non-attitude conjecture.

First, respondents who took more than one of our surveys gave inconsistent answers, across surveys, to questions about the price effect of a housing supply shock. Within-respondent, between-survey consistency on the housing items was barely better than would be observed if respondents were blindly guessing. By contrast, the same respondents gave highly consistent responses when asked to recall personal facts, like their partisanship and whether they had voted in the last local election. Preferences for future home prices and rents were also relatively stable between surveys. Non-housing price predictions on topics like cars, free trade, and agricultural and labor shocks were somewhat less consistent, but not nearly as unstable as answers to the housing supply–shock questions. The second piece of evidence comes from Survey 3, on which respondents were tested and then retested on the price effects of supply shocks in various markets. We preregistered a test of the hypothesis that, within the survey, respondents would give less consistent answers to housing than to non-housing-supply-shock questions. Respondents were in fact significantly less consistent on the housing questions.[6]

---

[5] For a presentation of the correlations, see Supplemental Appendix 6.

[6] For details of these results, see Supplemental Appendix 4, which also shows that within-survey retest consistency on economic-shock questions (both housing and non-housing) is about 10 percentage points lower among respondents who predicted that the regional housing shock would increase rents than among respondents who said that the shock would lower rents. That is, the people who answered the regional housing supply–shock question incorrectly were also more prone to contradicting themselves on other economic-knowledge items, which reinforces our view that housing supply skepticism is more a non-attitude than a firmly held belief. However, people do not *say* they are less knowledgeable on the housing items than the other items. Some respondents on Survey 3 were randomly assigned to a version of the supply-shock questions that included a "don't know" response option; this option was used less frequently on housing than non-housing-supply-shock items. See Supplemental Appendix 4.2.2.

Blame Attribution and Policy Preferences in the Cognitive Supply-Side Void

If most people have non-attitudes about the effect of regional housing supply shocks on home prices or rents, then one should expect the preferences of ordinary voters about land-use policies to be poorly explained by their price predictions. On Surveys 1 and 2, we asked about support for regional or statewide measures to allow more housing development. Conditional on respondents having an objective (tenure-based) or subjective (stated-preference) interest in lower housing prices, we observe only weak correlations between price predictions and policy preferences (from $r \approx 0$ to $r \approx 0.2$, depending on the policy question).[7] In another survey of California voters, Elmendorf and Nall (2024) find that support for a proposed state constitutional amendment giving municipalities exclusive control over land use is essentially uncorrelated with housing tenure or stated preferences for future rents.

In principle, an electorate composed of people with little understanding of housing markets but a desire for lower housing prices might be able to bring about pro-supply policy by rewarding incumbents when prices and rents have fallen, and voting them out of office when prices have been going up. However, such an accountability mechanism depends on (1) quick feedback between the enactment of land-use policies and the realization of lower (or higher) housing prices fast enough to influence vote choice at the next election, and (2) voters not having strong prospective views of what policymakers should be doing to make housing more affordable. That is, we would need people to behave as ideal-type retrospective voters, not as ideologues (Healy and Malhotra 2013).

The first condition is unlikely to be met; even effective pro-supply housing policies are unlikely to yield noticeably more affordable rents over electorally relevant timeframes. Regional housing stocks change slowly, delaying any effects of policy change on rents: as one example, Büchler and Lutz (2024) find that upzonings in Zurich, Switzerland increased housing stock in the targeted areas by 9 percent over five- to ten-year periods. To be sure, home prices might respond more quickly to increased supply, as they should reflect expectations about the stream of future rents. But the effect of supply liberalization on house prices is likely to be heterogeneous across parcels. The value of existing improvements should fall, but the redevelopment-option value of well-located parcels may rise (Lange and Teulings 2021). Adjustment of house prices to information about future rents may also be slow insofar as the marginal buyers are unsophisticated owner-users rather than professional investors.

The second condition—purely retrospective evaluation of candidates' housing performance—is vitiated by voters' willingness to blame specific groups for high housing costs and to endorse punitive policies against them. On Surveys 2 and 3, we invited respondents to nominate up to three of eleven actors as "responsible for high housing prices and rents in your area." Housing providers (developers and landlords) were most frequently named, as shown in Figure 3. People who blamed

---

[7] See Supplemental Appendix 7.1.

Figure 3
Attribution of Blame for "High Housing Prices in Your Area"
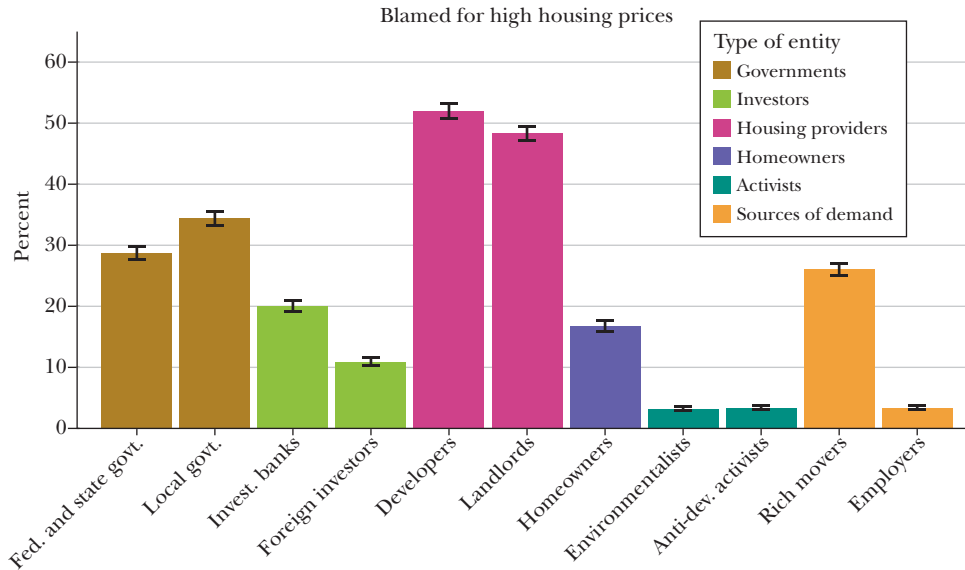
Note: This figure pools responses from Surveys 2 and 3. Respondents were allowed to select up to three actors. Standard errors are clustered on the respondent identifier. Results by survey are provided in Supplemental Appendix 3.

landlords or developers on Survey 2 were likely to do so again on Survey 3, suggesting that this is a real conviction rather than a top-of-mind response.[8] Groups called out in the economics and political science literature, such as homeowners, anti-development activists, and environmentalists, were rarely selected. Homeowners and renters largely blamed the same groups, although renters were somewhat more likely to fault landlords and homeowners more likely to fault developers.[9] While one might expect blame of developers and landlords to be highly correlated, manifesting a folk-economic hostility to real-estate "profiteers" (Shiller 2017), the correlation is weak.

We expect widespread animus towards landlords and developers to translate into punitive affordable-housing policies that saddle providers with regulations and price controls, rather than incentives to produce more housing. We explored this idea in another study in which we elicited preferences over a range of policies that, we told respondents, were being considered by "state lawmakers who think housing has gotten too expensive" (Elmendorf, Nall, and Oklobdzija 2024b). In one block

[8] See Supplemental Appendix 4.3.
[9] See Supplemental Appendix 7.4.

Figure 4
Perceived Relative Efficacy and Overall Support for Proposed Policies
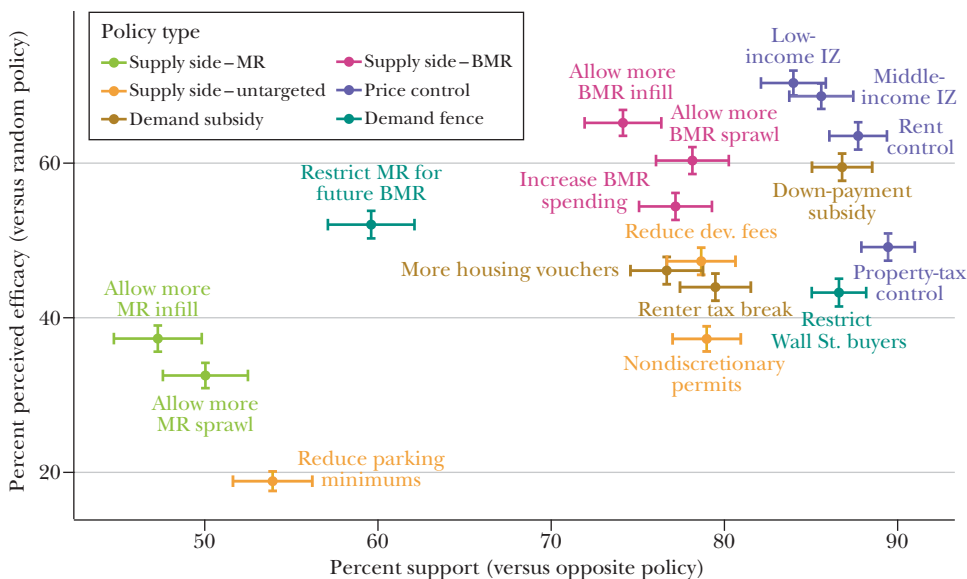
Note: "MR" refers to market-rate housing and "BMR" to price-regulated, below-market-rate housing. See Elmendorf, Nall, and Oklobdzija (2024b, Table 1) for the wording of each policy item. The sampling frame consists of residents of US urban and suburban zip codes, with approximately equal numbers of owners and renters.

of the survey, respondents were given random pairs of policies and asked which one would "be more effective for helping people in [name of the respondent's state inserted here] get housing they can afford."

From that study, Figure 4 locates each housing policy on "support" and "perceived efficacy" dimensions, with support (the horizontal axis) defined as the proportion of respondents who favor the policy, and perceived efficacy (the vertical axis) defined as the proportion of observations in which the policy was identified as more effective than another policy drawn at random. The policies are color-coded using a preregistered typology: (1) supply-side policies targeting market-rate housing, (2) supply-side policies targeting subsidized affordable housing, (3) untargeted supply-side policies, (4) demand subsidies, (5) price controls, and (6) demand fences (policies to keep disfavored actors from buying housing).

The public's take on effective and desirable state housing policies deviates significantly from the views of most economists. In a recent survey, only 2 percent of economists agreed that a rent cap on large apartment buildings would make middle-income Americans better off, versus 74 percent who disagreed (Weldon 2024). By contrast, in our survey of laypeople, price controls and demand subsidies received much more support than any supply-side measure. More than

85 percent of respondents backed rent control, property-tax limits, down-payment subsidies, and restrictions on Wall Street ownership of housing. Large supermajorities also want to require developers to provide price-restricted middle-income or lower-income housing. Policies to reduce development costs and zone for more market-rate housing drew the least support and were also regarded as least effective.[10] Replacing the term "market-rate housing" with "subsidized affordable housing" in questions about allowing more development increased support for the policy by roughly 25 percentage points. (Earlier in the survey, we gave respondents definitions of these housing types and tested their recollection to ensure that they understood.)

Our study of housing-policy preferences also assessed housing's importance relative to other major policy issues. Respondents were asked whether they support or oppose various housing and non-housing policies, with the non-housing policies chosen from items such as marijuana legalization and abortion rights that are prominent on state policy agendas. We then had respondents choose between pairs of randomized platforms, each composed of three policy planks. The platforms were composed so that they had contrasting position on each issue. For example, if one platform included, "allow more market-rate homes on open land," the other platform would include "restrict building of market-rate homes on open land." We measure the importance of an issue as the average effect, on the probability of picking a randomized platform, of switching the respondent's preferred position on the issue from the other platform to the target platform (Sides, Tausanovitch, and Vavreck 2023).

Price-control and investor-restriction housing policies had large effects on respondents' choice of platforms; pro-supply policies did not (Elmendorf, Nall, and Oklobdzija 2024b, Fig. 7). Rent control ranks near the top of the 39 policies we asked about, alongside hot-button issues such as abortion, the minimum wage, prescription drug pricing, gun control, taxing the rich, and marijuana legalization. Property-tax controls, low-income inclusionary zoning, and restrictions on Wall Street homebuyers come in just below it. Moreover, the supporters of rent control and restrictions on Wall Street investors care a lot more about these issues than do opponents, while opponents of market-rate housing development care somewhat more about it than do supporters (Elmendorf, Nall, and Oklobdzija 2024b, Fig. 8).

These findings help to illuminate why state and federal responses to the housing-affordability problem typically take the form of demand subsidies and incentives or requirements for deed-restricted affordable housing, rather than land-use deregulation.

---

[10] Much as people consistently fault landlords and developers yet contradict themselves on the price effects of housing supply shocks in the earlier survey results, we also found on the policy-preferences survey that people have stable preferences for regulating prices but seem less sure about whether to support market-rate housing development (Elmendorf, Nall, and Oklobdzija 2024b, Table 3).

Implications for Policy and Research

In recent years, there has been a concerted effort by Yes-In-My-Backyard advocacy groups, opinion leaders, and allied legislators to shift the locus of land-use policymaking from local governments to the states, or even to the federal government (for example, California YIMBY 2024; Up for Growth 2024). The jurisdictional fragmentation hypothesis motivates this work. Our results offer both a touch of hope and some reasons to worry about efforts to dislodge local control over land use.

The hopeful finding is that even though homeowners are overrepresented among likely voters, most homeowners (as well as nearly all renters) say they would like home prices and rents to be lower in the future. These responses may reflect preference falsification to some extent, but it appears that homeowners would not be inclined to vote en bloc as a regional housing cartel even if they knew how to do it. The worrisome implication is that a supply-skeptical public is unlikely to constitute an electoral constituency for pro-supply policies. While the problem of high housing costs has moved onto state legislative agendas (Kahn and Furth 2023), and even arose in the 2024 presidential election campaign (Green, Tankersley, and Nehamas 2024), our data provide little basis for expecting pro-supply policies becoming front-and-center in state or national legislatures. Some politicians may be persuaded of the merits of "abundance" reforms (Klein and Thompson 2025), but the mass public is primed for attacks on developers and landlords instead.

Our finding that housing supply skepticism in the mass public is closer to a non-attitude than an ideology might also be taken as good news, at least by policymakers accustomed to hearing from activists who demand curtailment of market-rate production (Been, Ellen, and O'Regan 2019). The weakness of the mass public's views suggests that there is room for persuasion. Perhaps publicizing the incredible success of zoning reform in Auckland, New Zealand, would start to change hearts and minds. Or perhaps exposing people to short-form videos explaining the operation of housing markets or even summaries of economics papers would make a difference. We have begun to address this in other work, using survey experiments that treat respondents with economics information or advocacy videos (Elmendorf, Nall, and Oklobdzija 2024a). Some of these treatments have large effects on economic beliefs and policy preferences—roughly two to three times larger than the typical effect sizes in prior studies of economics-information and political-advocacy treatments.

Another implication of our work is that policymakers hoping to build mass support for pro-housing policies will need to anticipate and defang the charge that they are doing the bidding of developers or landlords (see also Monkkonen and Manville 2019). It is probably no coincidence that in the birthplace of YIMBYism, California, the one clear-cut success is the state's accessory dwelling unit (ADU) law (Elmendorf and Nall 2024). The ADU law, which entitles homeowners to add up to two accessory units to their property, launched a new industry and a building boom. Though developers and landlords can also profit from the ADU law, there are plenty of stories of ordinary homeowners using it to build an in-law suite, guest house, home office, or rental unit.

Israel and South Korea have also had notable success with laws that empower a supermajority of condo owners to, in effect, upzone and redevelop their building (Alster 2023; Ji-hyoung 2024). Condo owners whose buildings have more redevelopment value are more strongly supportive of Israel's redevelopment law (Alster 2023). This suggests that market forces may help economically naïve homeowners learn how they can benefit personally from a highly localized upzoning of their neighborhood, even if the owners remain in the dark about the price effects of regional supply shocks. For the market to serve this educational function, the regulatory environment needs to facilitate such upzonings, so that someone (a developer, say, or a lawyer) has an incentive to contact condo owners and explain how they would benefit from redevelopment—or to buy up property of those who refuse to learn.

Ultimately, however, public opinion about who or what is to blame for high housing costs is likely to shape policy too. The journalist Jerusalem Demsas (2023) wrote, "By restricting the [housing] debate to the hyperlocal level, we've blocked out our big-picture values" and "magnifie[d] our selfish concerns . . . ." Our results suggest that Americans' "big-picture values" favor less-expensive housing, but that the devil is in the details of their folk-economic beliefs. Misconceptions about the economic effects of additional housing supply, and related enthusiasm for price controls, demand subsidies, and restrictions on investment, are lingering obstacles to converting big-picture values into effective policy.

References

Acolin, Arthur. 2022. "Owning vs. Renting: The Benefits of Residential Stability?" *Housing Studies* 37 (4): 644–67.

Ahlfeldt, Gabriel M., and Wolfgang Maennig. 2015. "Homevoters vs. Leasevoters: A Spatial Analysis of Airport Effects." *Journal of Urban Economics* 87: 85–99.

Albouy, David, and Gabriel Ehrlich. 2018. "Housing Productivity and the Social Cost of Land-Use Restrictions." *Journal of Urban Economics* 107: 101–20.

Alonso, William. 1964. *Location and Land Use: Toward a General Theory of Land Rent.* Harvard University Press.

Alster, Tal. 2023. "Homeowners Saying 'Yes, in My Back Yard': Evidence from Israel." *Urban Affairs Review* 59 (5): 1408–40.

Anagol, Santosh, Fernando V. Ferreira, and Jonah M. Rexer. 2023. "Estimating the Economic Value of Zoning Reform." NBER Working Paper 29440.

Apartment List Research Team. *Data & Rent Estimates [dataset].* Apartment List. https://www.apartmentlist.com/research/category/data-rent-estimates (accessed October 1, 2022).

Arnott, Richard J., and Joseph E. Stiglitz. 1979. "Aggregate Land Rents, Expenditure on Public Goods, and Optimal City Size." *Quarterly Journal of Economics* 93 (4): 471–500.

Asquith, Brian J., Evan Mast, and Davin Reed. 2023. "Local Effects of Large New Apartment Buildings in Low-Income Areas." *Review of Economics and Statistics* 105 (2): 359–75.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–68.

Bartik, Alexander, Arpit Gupta, and Daniel Milo. 2025. "The Costs of Housing Regulation: Evidence from Generative Regulatory Measurement." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4627587.

Baum-Snow, Nathaniel. 2023. "Constraints on City and Neighborhood Growth: The Central Role of Housing Supply." *Journal of Economic Perspectives* 37 (2): 53–74.

Baum-Snow, Nathaniel, and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–946.

Been, Vicki, Ingrid Gould Ellen, and Katherine O'Regan. 2019. "Supply Skepticism: Housing Supply and Affordability." *Housing Policy Debate* 29 (1): 25–40.

Boudette, Neal E. 2021. "'The Market Is Insane': Cars Are Sold Even Before They Hit the Lot." *New York Times*, November 2. https://www.nytimes.com/2021/07/15/business/car-sales-chip-shortage.html.

Boyer, Pascal, and Michael Bang Petersen. 2018. "Folk-Economic Beliefs: An Evolutionary Cognitive Model." *Behavioral and Brain Sciences* 41: e158.

Broockman, David, Christopher S. Elmendorf, and Joshua Kalla. 2024. "The Symbolic Politics of Housing." Unpublished.

Büchler, Simon, and Elena Lutz. 2024. "Making Housing Affordable? The Local Effects of Relaxing Land-Use Regulation." *Journal of Urban Economics* 143: 103689.

California YIMBY. 2024. "Restoring the California Dream: The Framework for California." https://cayimby.org/resources/policy-framework/.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter.* University of Chicago Press.

Caplan, Bryan. 2011. *The Myth of the Rational Voter: Why Democracies Choose Bad Policies.* Princeton University Press.

Choi, Jiwon, Ilyana Kuziemko, Ebonya Washington, and Gavin Wright. 2024. "Local Economic and Political Effects of Trade Deals: Evidence from NAFTA." *American Economic Review* 114 (6): 1540–75.

Citrin, Jack, and Donald Philip Green. 1990. "The Self-Interest Motive in American Public Opinion." *Research in Micropolitics* 3 (1): 1–28.

Clarke, Wyatt, and Matthew Freedman. 2019. "The Rise and Effects of Homeowners Associations." *Journal of Urban Economics* 112: 1–15.

Cosmides, Leda, and John Tooby. 1992. "Cognitive Adaptations for Social Exchange." In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, edited by Jerome H. Barkow, Leda Cosmides, and John Tooby, 163–228. Oxford University Press.

Demsas, Jerusalem. 2023. "Colorado's Ingenious Idea for Solving the Housing Crisis: And Why Local Governments Hate It." *Atlantic*, May 25. https://www.theatlantic.com/magazine/archive/2023/07/local-government-power-nimby-denver/674164/.

Diamond, Rebecca, and Tim McQuade. 2019. "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development." *Journal of Political Economy* 127 (3): 1063–117.

Downs, Anthony. 1957. "An Economic Theory of Political Action in a Democracy." *Journal of Political Economy* 65 (2): 135–50.

Duranton, Gilles, and Diego Puga. 2023. "Urban Growth and Its Aggregate Implications." *Econometrica* 91 (6): 2219–59.

Ellickson, Robert C. 1977. "Suburban Growth Controls: An Economic and Legal Analysis." *Yale Law Journal* 86 (3): 385–511.

Elmendorf, Christopher S. 2019. "Beyond the Double Veto: Housing Plans as Preemptive Intergovernmental Compacts." *Hastings Law Journal* 71: 79–150.

Elmendorf, Christopher S., and Clayton Nall. 2024. "Plain-Bagel Streamlining? Notes from the California Housing Wars." Preprint, SSRN. https://ssrn.com/abstract=4811580.

Elmendorf, Christopher S., Clayton Nall, and Stan Oklobdzija. 2024a. "Do Housing Supply Skeptics Learn? Evidence from Economics and Advocacy Treatments." Preprint, SSRN. https://doi.org/10.2139/ssrn.4955033.

Elmendorf, Christopher S., Clayton Nall, and Stan Oklobdzija. 2024b. "What State Housing Policies Do Voters Want? Evidence from a Platform-Choice Experiment." *Journal of Political Institutions and Political Economy* 5 (1): 117–52.

Elmendorf, Christopher S., Clayton Nall, and Stan Oklobdzija. 2025. *Data and Code for: "The Folk Economics of Housing."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E233932V1.

Fang, Limin, Nathan Stewart, and Justin Tyndall. 2023. "Homeowner Politics and Housing Supply." *Journal of Urban Economics* 138: 103608.

Favilukis, Jack Y., and Jaehee Song. 2023. "Why Zoning Is Too Restrictive." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4535417.

Fischel, William A. 2001. *The Homevoter Hypothesis: How Home Values Influence Local Government Taxation, School Finance, and Land-Use Policies.* Harvard University Press.

Glaeser, Edward, and Joseph Gyourko. 2002. "Zoning's Steep Price." *Regulation* 25: 24–30.

Glaeser, Edward L., Joseph Gyourko, and Raven E. Saks. 2005. "Why Have Housing Prices Gone Up?" *American Economic Review* 95 (2): 329–33.

González-Pampillón, Nicolás. 2022. "Spillover Effects from New Housing Supply." *Regional Science and Urban Economics* 92: 103759.

Green, Erica L., Jim Tankersley, and Nicholas Nehamas. 2024. "Harris Lays Out Her Economic Vision, Casting Trump's as Backward-Looking." *New York Times*, August 16. https://www.nytimes.com/2024/08/16/us/politics/kamala-harris-raleigh-nc-economy.html.

Greenaway-McGrevy, Ryan. 2023. "Can Zoning Reform Reduce Housing Costs? Evidence from Rents in Auckland." University of Auckland Economic Policy Centre Working Paper 016.

Greenland, Andrew, John Lopresti, and Peter McHenry. 2019. "Import Competition and Internal Migration." *Review of Economics and Statistics* 101 (1): 44–59.

Gyourko, Joseph, Jonathan S. Hartley, and Jacob Krimmel. 2021. "The Local Residential Land Use Regulatory Environment across US Housing Markets: Evidence from a New Wharton Index." *Journal of Urban Economics* 124: 103337.

Gyourko, Joseph, and Raven Molloy. 2015. "Regulation and Housing Supply." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 1289–337. North-Holland.

Hall, Andrew B., and Jesse Yoder. 2022. "Does Homeownership Influence Political Behavior? Evidence from Administrative Data." *Journal of Politics* 84 (1): 351–66.

Healy, Andrew, and Neil Malhotra. 2013. "Retrospective Voting Reconsidered." *Annual Review of Political Science* 16: 285–306.

Hilber, Christian A. L. 2017. "The Economic Implications of House Price Capitalization: A Synthesis." *Real Estate Economics* 45 (2): 301–39.

Hilber, Christian A. L., and Olivier Schöni. 2022. "Housing Policy and Affordable Housing." In *Oxford Research Encyclopedia of Economics and Finance*. https://doi.org/10.1093/acrefore/9780190625979.013.829.

Ji-hyoung, Son. 2024. "Yoon Vows to Remove Hurdle for Housing Redevelopment Projects." *Korea Herald*, January 10. https://www.koreaherald.com/view.php?ud=20240110000619.

Kahn, Eli, and Salim Furth. 2023. "Breaking Ground: An Examination of Effective State Housing Reforms in 2023." Mercatus Center, August 1. https://www.mercatus.org/research/policy-briefs/breaking-ground-examination-effective-state-housing-reforms-2023.

Kansal, Tushar, Andrew Aurand, and Sarah Saadian. 2024. "Homeowners, Renters, and All Income Groups Back Housing Reforms." *Pew*, May 23. https://www.pew.org/en/research-and-analysis/articles/2024/05/23/homeowners-renters-and-all-income-groups-back-housing-reforms.

Khan, Asad R. 2021. "Decentralized Land-Use Regulation with Agglomeration Spillovers: Evidence from Aldermanic Privilege in Chicago." PhD diss., University of Illinois at Urbana-Champaign.

Klein, Ezra, and Derek Thompson. 2025. *Abundance.* Avid Reader Press.

Lange, Rutger-Jan, and Coen N. Teulings. 2021. "The Option Value of Vacant Land: Don't Build When Demand for Housing Is Booming." Tinbergen Institute Discussion Paper 2021-022/IV.

Larsen, Martin Vinaes, and Neils Nyholt. 2024. "Understanding Nimbyism as Local Preservationism." OSF Preprints. October 14. doi:10.31219/osf.io/n2cme.

Li, Xiaodi. 2022. "Do New Housing Units in Your Backyard Raise Your Rents?" *Journal of Economic Geography* 22 (6): 1309–52.

Manson, Steven, Jonathan Schroeder, David Van Riper, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles. 2024. *IPUMS National Historical Geographic Information System: Version 19.0 [dataset].* Minneapolis, MN: IPUMS, 2024. http://doi.org/10.18128/D050.V19.0.

Marantz, Nicholas J., and Paul G. Lewis. 2022. "Jurisdictional Size and Residential Development: Are Large-Scale Local Governments More Receptive to Multifamily Housing?" *Urban Affairs Review* 58 (3): 732–66.

Marble, William, and Clayton Nall. 2021. "Where Self-Interest Trumps Ideology: Liberal Homeowners and Local Opposition to Housing Development." *Journal of Politics* 83 (4): 1747–63.

Menendian, Stephen, Shahan Shahid Nawaz, and Samir Gambhir. 2024. *Single-Family Zoning in California: A Statewide Analysis.* Othering and Belonging Institute, University of California, Berkeley.

Mense, Andreas. 2023. "Secondary Housing Supply." IAB-Discussion Paper 6/2023.

Missouri Census Data Center. 2022. *Geocorr 2022: ZIP Code to County Correspondence File [dataset].* University of Missouri. https://mcdc.missouri.edu/applications/geocorr2022.html (accessed March 15, 2022).

Molloy, Raven, Charles G. Nathanson, and Andrew Paciorek. 2022. "Housing Supply and Affordability: Evidence from Rents, Housing Consumption and Household Location." *Journal of Urban Economics* 129: 103427.

Monkkonen, Paavo, and Michael Manville. 2019. "Opposition to Development or Opposition to Developers? Experimental Evidence on Attitudes toward New Housing." *Journal of Urban Affairs* 41 (8): 1123–41.

Muth, Richard F. 1969. *Cities and Housing: The Spatial Pattern of Urban Residential Land Use.* University of Chicago Press.

Orlando, Anthony W., and Christian L. Redfearn. 2024. "Housing Supply Elasticities: A Structural Vector Autoregression Approach." Unpublished.

Ortalo-Magné, François, and Andrea Prat. 2014. "On the Political Economy of Urban Growth: Homeownership versus Affordability." *American Economic Journal: Microeconomics* 6 (1): 154–81.

Pareto Software, LLC. 2022. *US Zip Codes Database [dataset].* Updated June 11, 2025. https://simplemaps.com/data/us-zips (accessed March 15, 2022).

Parkhomenko, Andrii. 2023. "Local Causes and Aggregate Implications of Land Use Regulation." *Journal of Urban Economics* 138: 103605.

Pennington, Kate. 2021. "Does Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.3867764.

Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life." *Journal of Political Economy* 90 (6): 1257–78.

Rosen, Sherwin. 1979. "Wage-Based Indexes of Urban Quality of Life." In *Current Issues in Urban Economics*, edited by Peter Mieszkowski and Mahlon Straszheim, 74–104. Johns Hopkins University Press.

Rubin, Paul H. 2003. "Folk Economics." *Southern Economic Journal* 70 (1): 157–71.

Shiller, Robert J. 2017. "Narrative Economics." *American Economic Review* 107 (4): 967–1004.

Sides, John, Chris Tausanovitch, and Lynn Vavreck. 2023. *The Bitter End: The 2020 Presidential Campaign and the Challenge to American Democracy.* Princeton University Press.

Stahl, Kenneth. 2018. "'Yes in My Backyard': Can a New Pro-Housing Movement Overcome the Power of NIMBYs?" *Zoning and Planning Law Report* 41 (3): 1–16.

Ternullo, Stephanie. 2024. "The Politics of Suburban Restrictivism." Preprint, SSRN. https://doi.org/10.2139/ssrn.4766660.

Tricaud, Clemence. 2025. "Better Alone? Evidence on the Costs of Intermunicipal Cooperation." *American Economic Journal: Applied Economics* 17 (1): 160–207.

Up for Growth. 2024. "Legislative Agenda." https://upforgrowth.org/up-for-growth-action/legislative-agenda/.

US Census Bureau and US Department of Housing and Urban Development. 2023a. New Privately-Owned Housing Units Started: Total Units. Distributed by Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/HOUST.

US Census Bureau and US Department of Housing and Urban Development. 2023b. *Median Sales Price of Houses Sold for the United States.* Distributed by Federal Reserve Bank of St. Louis. https://fred.stlouisfed.org/series/MSPUS.

Ward, Jason M., and Luke Schlake. 2025. *The High Cost of Producing Multifamily Housing in California: Evidence and Policy Recommendations.* RAND Corporation.

Weldon, Duncan. 2024. "Rent Controls and One-Armed Economists." Kent A. Clark Center for Global Markets, Chicago Booth, August 2. https://kentclarkcenter.org/filterable-categories/on-global-markets/rent-controls-and-one-armed-economists/.

Yagan, Danny. 2019. "Employment Hysteresis from the Great Recession." *Journal of Political Economy* 127 (5): 2505–58.

Zillow Group. 2022. *Zillow Research: Housing Data.* https://www.zillow.com/research/data/ (accessed March 15, 2022).

Building Costs and House Prices

Brian Potter and Chad Syverson

High and increasing housing prices are an active and sometimes heated topic of policy debate (for instance, US GAO 2023; Bundrick 2024; Casselman 2024). In the United States, what were once a collection of metro-specific housing price outliers have recently morphed into widespread price surges. A broadly similar pattern holds in many other countries as well. For example, implied housing price-to-rent ratios in some other countries like Canada are multiples of their US counterparts, even with the latter at historical highs (*Economist* 2024).

Figure 1 compares housing prices and building costs across a selected group of US cities since 1950. Across-city variations in the gap between price and cost growth are visually obvious. Some cities see price growth adhering closely to building cost growth throughout the period. Others experience modest divergences for five- or ten-year periods (often recently), and still others see early and enormous separation of local house prices from the estimated cost of constructing them. Most notable is the extreme rise in housing prices compared to costs of large coastal metros, including the major cities in California (Los Angeles and San Francisco), Seattle, New York, and Boston. These cities have historically been perceived as failing to build sufficient housing to keep up with demand (Badger and Washington 2022), reflected here as a large cost-price disconnect. Cities that have historically built housing in large volumes, such as Dallas, Houston, Atlanta, and Phoenix (Potter 2021), show a much smaller difference between price and cost. However, over the last ten years many of these cities have seen house prices rise dramatically faster than costs.

■ *Brian Potter is Senior Infrastructure Fellow, Institute for Progress, Washington, DC. Chad Syverson is Professor of Economics, University of Chicago Booth School of Business, Chicago, Illinois. Their email addresses are briancpotter@gmail.com and chad.syverson@chicagobooth.edu.*

Figure 1
Selected City-Level RSMeans Housing Cost and FHFA Housing Price Indexes
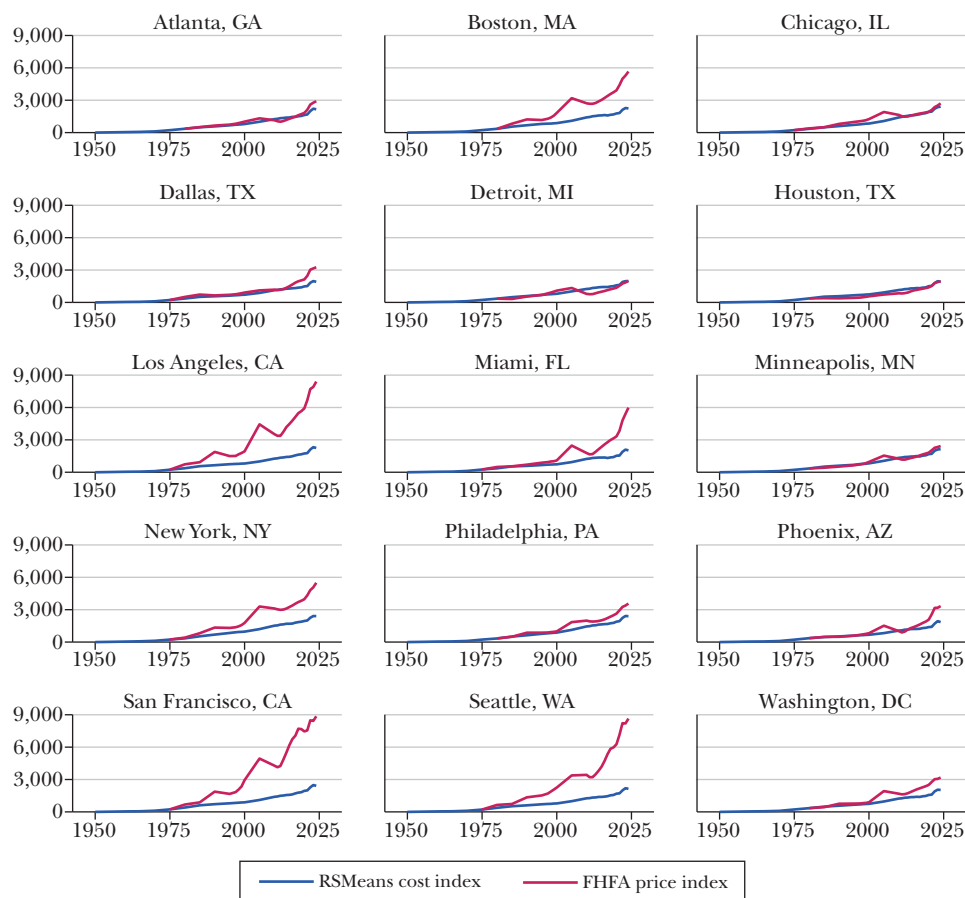
Source: Authors' creation.
Note: The housing cost data come from RSMeans and are shown by the blue lines. Housing price data are from the Federal Housing Finance Agency (FHFA) and are shown by the red lines. These data sources are discussed further later in the paper. The city's FHFA index is normalized to the level of the RSMeans index in the year the FHFA index becomes available, typically 1975 or 1980.

The sources of high house prices are of obvious policy and intellectual interest. A natural candidate explanation is that they reflect high building costs, by which we mean the costs of actually constructing the building itself apart from land acquisition costs. After all, building costs on average account for an estimated 60–70 percent of the full cost of bringing a new house to market in a land-unconstrained area (Glaeser and Gyourko 2018; Lynch 2023). Moreover, the construction sector's poor productivity growth over the past several decades has raised the salience of high and rising building costs in the housing market (D'Amico et al. 2024; Goolsbee and

Syverson 2024; Garcia and Molloy 2025). Yet the divergences in Figure 1 suggest that other factors may often be in play.

In this article we take a long, broad, and theoretically agnostic view toward the connection between building costs and house prices in the US housing market. Our intent is to lay out some facts, drawing connections and tentative conclusions when warranted. We hope this essay will serve as an input into future research on cost-price differences in this highly relevant and interesting market.

Perhaps the clearest conclusion of our analysis is that building costs have limited explanatory power over US housing prices, and even the imperfect correlations of the past have weakened further in recent decades along multiple dimensions. This thesis has a clear conceptual antecedent in Glaeser and Gyourko (2018), published in this journal. They used straightforward but powerful economic theory to organize discussion and interpretation of the connection between housing prices and costs in recent decades. We also explore this connection, but we differ by taking an empirical view that is longer (going back three-fourths of a century) and on some dimensions broader in the cross section. Zooming out in this way highlights a pattern that threads itself through our results in multiple ways; namely, that to the extent construction costs explain house price patterns—already tenuous in certain markets as documented by Glaeser and Gyourko (2018)—they are explaining less than they used to.

How to Measure Housing Construction Costs

The core of our empirical analysis compares metrics of house construction costs to house prices. For construction costs, we rely on the Housing Cost database from RSMeans, a company that has been supplying cost estimate data to the construction industry since the mid-twentieth century. RSMeans describes its cost estimates—which capture the costs and margins associated with actually constructing a building, but exclude ancillary and land acquisition costs—as being compiled "through a combination of field research, industry surveys, cost modeling, and partnerships with construction professionals and suppliers." The validity of the RSMeans data are reflected in part in their extensive use in the construction and building materials industries. Researchers have used them too. For example, Glaeser and Gyourko (2018), Kakar, Daniels, and Grossman (2018), and Duranton and Puga (2023) use RSMeans data to measure construction costs.[1] Our version of the data contains detailed building cost estimates and indexes for various US cities.

To better understand how well RSMeans data track actual construction costs, we compare them to two other widely known indexes of housing construction costs.

First, the Census Single-Family Houses under Construction price index tracks the change in cost in single-family homes using home price data from the Census

---

[1] RSMeans data are available for general purchase from the Gordian company. See https://www.rsmeans.com/ for more details.

Figure 2

Census Single-Family Houses under Construction Price Index and RSMeans Cost Index, 1940–2024
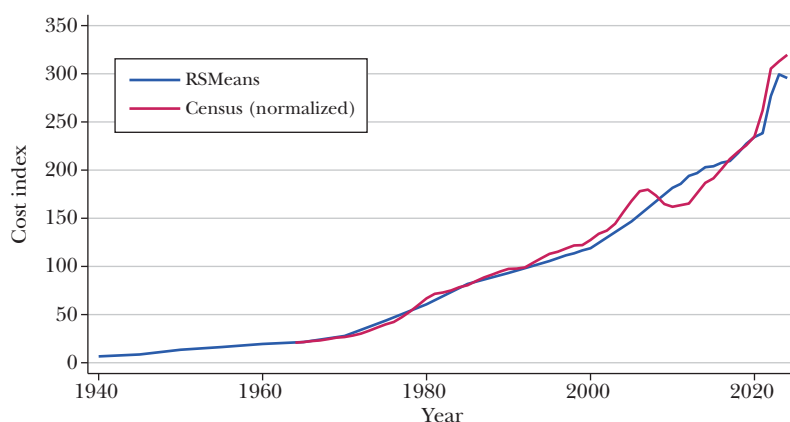
Source: Authors' creation.
Note: This figure shows two housing cost indexes. The RSMeans index is in blue. The Census Single-Family Houses under Construction index (computed from the Census Bureau's Survey of Construction) is in orange. We normalize the 1965 start of the Census index to the contemporaneous level of the RSMeans index.

Bureau's Survey of Construction. The Survey measures housing construction activity by sampling from several hundred local permit-issuing offices, as well as canvassing land areas that do not issue permits. The values under construction exclude the cost of land, and thus track actual house construction costs while avoiding potentially confounding land scarcity effects.

Figure 2 plots the two series together (we normalize the Census index's 1965 start to the contemporaneous level of the RSMeans index). They track each other closely, with the price volatility around the housing boom and bust in the late 2000s being the times of the greatest departures.

The Census comparison is a useful verification that the RSMeans data track nationwide changes in construction costs, but it would also be worthwhile to compare the RSMeans data for individual cities. For this comparison, we use the Turner and Townsend International Construction Market Survey, which tracks construction costs across a range of building types in major cities around the world.[2]

We compare the RSMeans and Turner and Townsend data for seven major US cities in 2019 and five major US cities in 2024. Turner and Townsend construct their data through a multi-input process similar to that used by RSMeans. Here is how we build this comparison: We use each dataset to produce cost estimates

---

[2] This data are available from the Turner and Townsend website at https://www.turnerandtownsend. com/insights/international-construction-market-survey-2024/.

Table 1

Comparison of RSMeans and Turner and Townsend Construction Cost Indexes across Cities

	Townhouse	*High-rise apartment*	*RSMeans*	
Panel A. 2024 (normalized to Houston)				
Houston	1.00	1.00	1.00	
Chicago	1.06	1.42	1.38	
Los Angeles	1.80	1.60	1.32	
San Francisco	1.85	2.25	1.48	
New York	1.87	2.29	1.47	
Correlation	0.70	0.88		
	Townhouse	*Single-family home (medium quality)*	*Single-family home (prestige quality)*	*RSMeans*
Panel B. 2019 (normalized to Houston)				
Houston	1.00	1.00	1.00	1.00
Phoenix	1.02	1.02	1.02	1.04
Atlanta	1.04	1.04	1.04	1.04
Chicago	1.19	1.20	1.40	1.41
Seattle	1.68	1.73	1.33	1.25
New York	1.93	2.15	1.43	1.57
San Francisco	2.45	2.09	1.52	1.52
Correlation	0.82	0.86	0.97	

Note: This table compares RSMeans data and the Turner and Townsend International Construction Market Survey data across cities. We use each dataset to produce cost estimates expressed in dollars, which we then normalize into an index by setting it equal to 1 for Houston, the lowest-cost city in both indexes. Further details can be found in the Supplemental Appendix.

expressed in dollars, which we then normalize into an index by setting it equal to 1 for Houston, the lowest-cost city in both indexes. The data are in Table 1. There are high cross-city correlations between the two measures, ranging from 0.70 to 0.97, though RSMeans indicates somewhat lower costs than Turner and Townsend for the very high-cost cities of New York and San Francisco. While tempered by the fact that the number of comparable cities is small, we take this correspondence as indicating RSMeans data usefully reflect building cost differences across geographies as well.

We begin exploring the connection between homebuilding costs and house prices with a look at long-run trends. Figure 3 shows multiple price indexes over the last three-fourths of a century. In addition to the RSMeans house building cost index already discussed and the Consumer Price Index from the Bureau of Labor Statistics, it also includes several other indexes. The "construction sector gross output price index" is from the Bureau of Economic Analysis (BEA) industry accounts, which the BEA builds to capture the evolution of prices for all of the sector's outputs (including houses, other buildings, and public works). Goolsbee and Syverson (2024) employ this index in their study of construction sector productivity, for example. The housing price index produced by the Federal Housing Finance Agency (FHFA) is constructed from transaction data from Fannie Mae and

Figure 3
Comparison of Price Indexes since 1950 (relative to the GDP deflator)
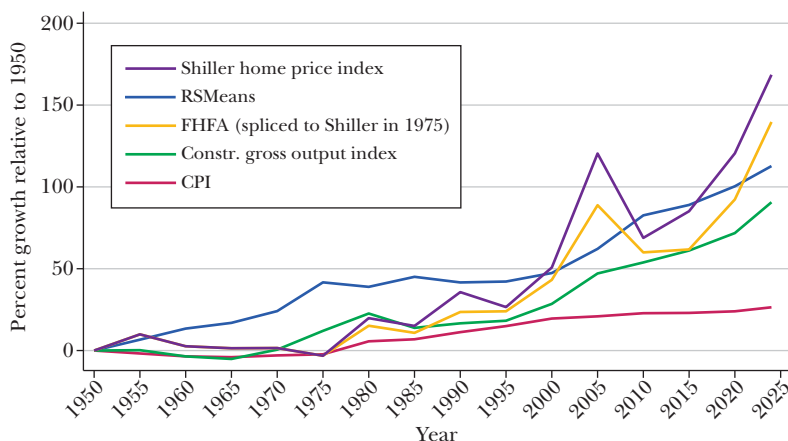
Source: Authors' creation.
Note: The figure compares multiple price indexes: The RSMeans Housing Cost Index, the Consumer Price Index (CPI), the construction sector gross output price index from the Bureau of Economic Analysis industry accounts, the Shiller Home Price Index (calculated based on repeat sales), and the Federal Housing Finance Agency index (which also uses repeat sales, or refinancings). Each of these indexes has been normalized by the GDP deflator over the entire period. The FHFA index only starts in 1975, so we splice it to the level of the Shiller index that year. We use the series values in every fifth year (those ending in "0" or "5").

Freddie Mac using repeat sales or refinancings of the same house. In so doing, it adjusts more thoroughly than cross-sectional samples for quality changes in the housing stock. For examples of research using this index, useful starting points are Guerrieri, Hartley, and Hurst (2013), Charles, Hurst, and Notowidigdo (2018), and Glaeser and Gyourko (2025). We further include the Shiller home price index, which is a composite of Census-division-level single-family home price indices. Like the FHFA index, this is also calculated based on previous sales of the same home. Examples of its use in research include Case and Shiller (2003), Mian and Sufi (2009), and Glaeser, Gottlieb, and Gyourko (2013).

Each of these indexes has already been normalized by the GDP deflator over the entire period. As such, they reflect trends relative to that deflator. Clearly, each grew faster than the GDP deflator over the period.

Several patterns are apparent here. First, construction-related prices of all sorts—whether captured by building costs, construction output deflators, or house prices—grew faster than the other consumer-related prices in the Consumer Price Index. While the CPI saw cumulative growth of about 20 percent relative to the GDP deflator, construction sector gross output prices almost doubled, and housing construction costs and prices both more than doubled. This outsized growth, at least in the latter two-thirds of the sample period, is consistent with the aforementioned

productivity underperformance of the sector relative to the rest of the economy post-1970 (D'Amico et al. 2024; Goolsbee and Syverson 2024; Garcia and Molloy 2025).

Second, for much of the early period, cumulative growth in construction costs, as reflected in the RSMeans data, outpaced housing prices. In fact, cumulative growth in the Shiller house price index over 1950–1975 was 30 percent *less* than comparable change in the RSMeans building cost index. In the late 1970s, however, house prices started rising faster than construction costs. Within the time horizon of the figure, cumulative price growth finally caught up to and surpassed cost growth in the late 1990s.

Third, there were two separate periods where post-1950 cumulative growth in housing prices was notably above the same for building costs: the mid-2000s, during the run-up to the subprime mortgage lending–driven financial crisis; and most recently, from the mid-2010s to the present. While the mid-2000s house price spike did not coincide with a similar acceleration in building costs, this is less true for the contemporaneous bout of price growth. House prices and building costs alike have recently risen substantially faster than their prior growth rates, though housing prices have risen more quickly.

Ultimately, over the entire sweep of the data, house prices cumulatively rose only modestly more than the increase in building costs: about 25 percent more using the Shiller housing price index, and only 10 percent more if we splice the Federal Housing Finance Agency index to the Shiller index in 1975. If we focus on the trends since 1975, however, the differences are starker. Cumulative house price growth using the FHFA index outpaced building costs by over 60 percent; the difference using the Shiller housing index is 80 percent.[3] This pattern of increasingly divergent building costs and house prices in recent decades will repeat itself in our discussion below. Note, interestingly, that this divergence occurs throughout a period where construction productivity growth was very poor. That is, while building cost inflation may well have contributed to high house price growth in recent years, other factors must have made an additional and substantial contribution.

Relationships between Cost and Price Growth in a City Panel

We next drill below the aggregates to look at patterns of house price and building cost growth at the city level. We match city-specific building cost indexes, using the RSMeans data, with housing price indexes, using the Federal Housing Finance Agency indexes at the level of metropolitan statistical areas.

RSMeans data are not always reported at the level of an entire Census-defined metropolitan statistical area. For instance, there are separate indexes for Boston and Fall River in Massachusetts, despite both being part of the same metropolitan statistical area. Another example includes the separate indexes for New York City,

---

[3] For graphical illustration of these points, see Supplemental Appendix Figures A1a and A1b.

Table 2
Summary Statistics for City-Level Building Cost and House Price (in Percent, Showing Compound Annual Growth Rates by Time Period)

Period	N	RSMeans mean	RSMeans standard deviation	FHFA mean	FHFA standard deviation
1975–1980	13	6.84%	1.03	13.26%	4.38
1980–1985	128	5.83	0.78	3.90	3.12
1985–1990	155	2.31	0.48	4.76	4.17
1990–1995	174	2.24	0.53	3.00	2.81
1995–2000	178	2.37	0.61	4.04	1.77
2000–2005	178	4.21	0.47	7.15	4.20
2005–2010	178	4.55	0.51	−0.61	3.49
2010–2015	178	2.68	0.59	1.45	2.06
2015–2020	178	2.68	0.58	4.78	1.89
2020–2024	178	6.27	0.52	10.11	2.01

Note: N is the number of cities in the data for each time period.

New York, and Jersey City, New Jersey, again in the same metropolitan statistical area. In these cases, we use the index corresponding to the center city. The number of cities covered in the Federal Housing Finance Agency data on housing prices is quite small early in the panel. Only 13 cities are covered from 1975 to 1980. (Earlier in the paper, Figure 1 offered illustrations of the cost and price comparison for some of those cities plus others whose series began in 1980.) Coverage increases considerably to 128 cities from 1980 to 1985 and to 178 cities by the 1995–2000 period, where the number of covered cities stabilizes.[4]

As a basis for a broader examination, Table 2 reports summary statistics of building cost and house price growth (showing compound annual growth rates over the respective periods) for our entire panel. We noted above that the aggregates indicate house prices began to systematically outgrow building costs after 1975. This is observed at the city level as well. For seven of the ten five-year periods (the last period is shortened to 2020–2024 given data availability), average city-level house price growth rates exceeded average city-level building cost growth rates. The exceptional periods where building costs outstripped price growth were 1980–1985, 2005–2010, and 2010–2015, with low price growth in the latter two periods driven by home price declines following the global financial crisis. The only time in which any of the annual growth rates were negative was for house prices during 2005–2010.

---

[4] We previously noted in the comparison to the Turner and Townsend data that RSMeans might underestimate construction costs in very expensive metros, but in most cases the large price-cost gap in coastal cities remains even if we adjust for this. Specifically, even if we increase the cost index by 50 percent in large coastal cities shown earlier in Figure 1, the normalized price-cost ratio is still above 2.0 in San Francisco, Los Angeles, and Seattle, and it is roughly 1.5 in Atlanta, Phoenix, and Houston.

Table 3
Cross-Sectional Regressions of Compound Average Growth Rates for House Price on Construction Costs

Period	N	R^2	Coefficient
1975–1980	13	0.641	3.420
			(0.374)
1980–1985	128	0.182	1.705
			(0.386)
1985–1990	155	0.329	4.990
			(0.629)
1990–1995	174	0.110	−1.750
			(0.344)
1995–2000	178	0.000	0.060
			(0.203)
2000–2005	178	0.009	0.851
			(0.569)
2005–2010	178	0.015	0.831
			(0.511)
2010–2015	178	0.035	−0.658
			(0.241)
2015–2020	178	0.011	0.347
			(0.197)
2020–2024	178	0.020	0.551
			(0.298)

Note: House price data are the FHFA metro-specific indexes, and building costs are from RSMeans. Each row is a separate calculation, using data across cities for that time period. Details are available the Supplemental Appendix.

Just as aggregate house price growth has been more volatile over time than aggregate building cost growth, price growth is substantially more variable in the cross section as well. The typical standard deviation across cities of house price growth rates ranges from 2 to 4 percentage points per year, while for building costs it ranges from 0.5 to 1 percentage points. There is very little skewness across cities in the building cost distribution, but a little more among house prices.[5]

We dig into the cross-sectional correlation of growth rates by running separate cross-sectional regressions for each of these five-year periods. The results are in Table 3. Looking down along the R^2 of regressions, we observe—at the city level— an increasing divergence of building costs and house prices over time. The ability of building cost growth to predict price growth in the cross section falls throughout our panel. In fact, there is a notable step down in 2000 from modest explanatory power to very little, if any, explanatory power (single-digit R^2 values). The coefficients vary

---

[5] The patterns here are not solely a function of a few highly divergent cities; the medians (not shown here) show similar patterns. Another way of making this point is that the median for compound aggregate growth rates in housing costs and prices do not differ substantially from the means. For details, see the Supplemental Appendix.

a fair amount from year to year. In fact, from 1990 to 1995—a period where cost growth still explains over 10 percent of price growth—the relationship is negative, hardly the direction implied by simple economic theory.[6]

Which Cities See the Biggest Divergence between Housing Prices and Building Costs?

The cities that saw the largest gaps between house price and building cost growth have changed over time. Table 4 lists the five cities that saw the largest and smallest price growth in excess of building costs during each five-year time period of our sample.[7]

During 1980–1985, Massachusetts cities and those in the New York City region dominated the cities whose house prices most exceeded cost growth. Boston was the highest, with annual price growth 7.7 percent higher than local building cost growth. Excess price growth moved west for the late 1980s, with coastal California cities and Honolulu on top (the last experiencing annual price growth 14 percent higher than costs). The most divergent prices shifted to Utah and the Pacific Northwest for 1990–1995. A mishmash of cities experienced fast relative price growth from 1995 to 2000. During 2000–2005, California cities, including some inland, see excess growth to the tune of 14 percent per year. As housing prices collapsed in the latter half of the 2000s, only one city in the country (!) saw price growth above building cost growth: Odessa, Texas, in the midst of a fracking-driven oil and gas boom. Other cities in Texas, plus Huntsville, Alabama, rounded out the top five from 2005 to 2010, though these all exhibited house price declines as a result of the global financial crisis. California cities lead the house price rebound from 2010 to 2015, and the Pacific Northwest (now including Boise, Idaho) saw the fastest relative price growth from 2015 to 2020. The post-Covid price run-up was led by Miami and saw cities from Tennessee, Georgia, and Maine enter the mix for the first time.

As for the locations seeing the slowest relative price growth (which in all cases is negative, meaning that house prices grew more slowly than building costs did) during 1980–1985, they were in deindustrializing Midwestern cities and Eugene, Oregon. Texas, Oklahoma, and Alaska bore the brunt of the oil price collapse from 1985 to 1990. Smaller New England cities and Greater Los Angeles, following the social unrest there in 1992, saw big relative price declines over 1990–1995. In all these cases, the fastest-declining relative price cities saw prices fall some 6 or 7 percent per year relative to building costs. The 1995–2000 and 2000–2005 periods

---

[6] Supplemental Appendix Table A1 shows a version of Table 3 where we also control for city population growth over the period. This matters little for most of the sample; we cannot reject equality of the coefficients on the construction cost growth rate with and without population growth included for all periods through 2010–2015. However, we can reject equality for the 2015–2020 and 2020–2024 periods. For those periods, controlling for population growth roughly doubles the partial correlation between building cost and house price growth across cities. However, even in those cases, costs have little additional explanatory power over prices. The increase in R^2 when adding cost growth to a regression of price growth on population growth is 0.001 in both 2015–2020 and 2020–2024.
[7] We do not report values for 1975–1980 due to the small number of cities in the panel.

Table 4

Top and Bottom Five Cities' House-Price-to-Building-Cost Growth, by Period

City	1980–1985 %ΔPrice − %ΔCost	City	1985–1990 %ΔPrice − %ΔCost	City	1990–1995 %ΔPrice − %ΔCost
Boston, MA	7.7	Honolulu, HI	14.1	Salt Lake City, UT	9.0
Worcester, MA	6.9	San Francisco, CA	12.2	Ogden, UT	6.9
Springfield, MA	5.3	Los Angeles, CA	11.4	Spokane, WA	6.3
New York, NY	4.9	Oxnard, CA	10.9	Portland, OR	6.2
Newark, NJ	3.3	Anaheim, CA	9.8	Eugene, OR	5.9
Eugene, OR	−6.7	San Antonio, TX	−4.9	Manchester, NH	−5.8
Detroit, MI	−6.8	Oklahoma City, OK	−5.5	Oxnard, CA	−5.8
Davenport, IA	−6.9	Odessa, TX	−6.3	New Haven, CT	−6.6
Waterloo, IA	−7.9	Anchorage, AK	−7.0	Los Angeles, CA	−6.6
Peoria, IL	−8.7	Austin, TX	−7.1	Hartford, CT	−7.0

City	1995–2000 %ΔPrice − %ΔCost	City	2000–2005 %ΔPrice − %ΔCost	City	2005–2010 %ΔPrice − %ΔCost
San Francisco, CA	9.6	Santa Barbara, CA	14.4	Odessa, TX	2.7
Boston, MA	7.2	Riverside, CA	14.3	Austin, TX	−0.5
Charleston, SC	6.6	Fresno, CA	14.1	Beaumont, TX	−0.6
Denver, CO	6.0	Los Angeles, CA	13.5	El Paso, TX	−0.8
Seattle, WA	5.8	Bakersfield, CA	13.4	Huntsville, AL	−0.8
Buffalo, NY	−1.6	Decatur, IL	−1.2	Riverside, CA	−13.8
Camden, NJ	−1.7	Charleston, WV	−1.4	Reno, NV	−14.5
Springfield, IL	−1.7	Springfield, IL	−1.5	Vallejo, CA	−16.2
Utica, NY	−2.0	Memphis, TN	−1.8	Las Vegas, NV	−17.9
Honolulu, HI	−5.0	Sioux City, IA	−3.0	Stockton, CA	−17.9

City	2010–2015 %ΔPrice − %ΔCost	City	2015–2020 %ΔPrice − %ΔCost	City	2020–2024 %ΔPrice − %ΔCost
Stockton, CA	5.3	Boise, ID	8.9	Miami, FL	8.6
San Francisco, CA	5.1	Spokane, WA	7.1	Knoxville, TN	8.1
Vallejo, CA	4.9	Tacoma, WA	7.0	Camden, NJ	8.0
Denver, CO	4.4	Tampa, FL	6.5	Savannah, GA	7.5
Reno, NV	4.1	Las Vegas, NV	6.1	Lewiston, ME	7.3
Columbia, SC	−4.6	Decatur, IL	−1.1	Baton Rouge, LA	−1.2
Tuscaloosa, AL	−4.7	Shreveport, LA	−1.1	Shreveport, LA	−1.8
Mobile, AL	−5.2	Lawton, OK	−1.3	New Orleans, LA	−1.9
Rockford, IL	−5.4	Springfield, IL	−1.5	Odessa, TX	−2.5
Montgomery, AL	−5.7	Peoria, IL	−1.5	Lake Charles, LA	−2.8

Note: House price data are the FHFA metro-specific indexes, and building costs are from RSMeans. Details are available in the Supplemental Appendix.

saw an assortment of cities from various regions experiencing the largest relative price declines, though they were smaller drops in magnitude than those of earlier

periods. Of course 2005–2010 was a banner period for falling prices, with cities in Nevada and inland California seeing annual price drops of 14 to 18 percent relative to costs, reversing gains of the prior half decade. The 2010–2015 and 2015–2020 periods saw a variety of cities in Illinois and Alabama experience the largest relative price drops, and the COVID-influenced period from 2020–2024 was dominated by Louisiana cities along with Odessa, Texas.

While some cities appear in the top or bottom five multiple times, there is a considerable turnover across periods. This suggests relative price-cost growth in a city is not highly persistent. Indeed, formal tests do indeed indicate mean reversion in cities' housing price-cost ratios; that is, faster relative price growth in one period predicts slower relative price growth in the following period, tempering cumulative differences between house prices and building costs. Though as we see below, these differences can still become quite large. If we break this growth rate difference into its price and building cost components, both exhibit mean reversion.[8]

We can of course add up these period-specific growth differences to obtain a measure of the cumulative gap between building costs and house prices in a city. We compute the ratio of cumulative house price growth to cumulative building cost growth in the city since 1980, and then plot the density of this ratio across cities for the 126 cities that we observe from 1980 to the present in Figure 4. Values above 1 indicate more cumulative growth since 1980 in house prices than in building costs.

Some interesting patterns emerge from inspecting these distributions. Despite the average and median growth rate patterns discussed above, several cities experienced lower cumulative house price growth than building cost growth—as is shown by any values below 1. This is true even at the sample's end in 2024: from 1980 to 2024, 33 of the 128 cities saw faster housing cost growth than price growth. Near the start of the time period, the distributions are relatively condensed; for example, the two curves with the highest density peaks are those for 1985 and 1990. Over time, the distributions became quite skewed. The distribution with the biggest rightward skews is for the average price spike of 2005, but the next three curves most skewed to the right are 2024, 2020, and 2015. For instance, in 2024, the 90th percentile cumulative price-cost ratio was 1.90. It was even larger during the price spike of 2005, at 2.21.

Material versus Installation Costs for Cities

Since 2007, the RSMeans Index provides separate indexes for "material" and "installation" indexes, which together comprise the overall building cost index. The material index summarizes the prices of building materials: concrete, lumber, drywall, plumbing fixtures, and so on. The installation index reflects costs tied to the assembly of materials into finished buildings, including some kinds of building-related

---

[8] Supplemental Appendix Table A2 shows that regressing a city's five-year difference between housing price and building cost growth on the same variable computed five years earlier while including city and period fixed effects yields a coefficient of −0.329 (standard error = 0.027). Looking at housing price growth and cost growth separately, five-year price growth regressed on its lag yields a coefficient of −0.288 (standard error = 0.030). The same specification using building cost growth results in a coefficient on lagged cost growth of −0.157 (standard error = 0.028). Details are available in the Supplemental Appendix.

Figure 4
Across-City Density of Cumulative Housing Prices to Costs Ratio Level Ratio
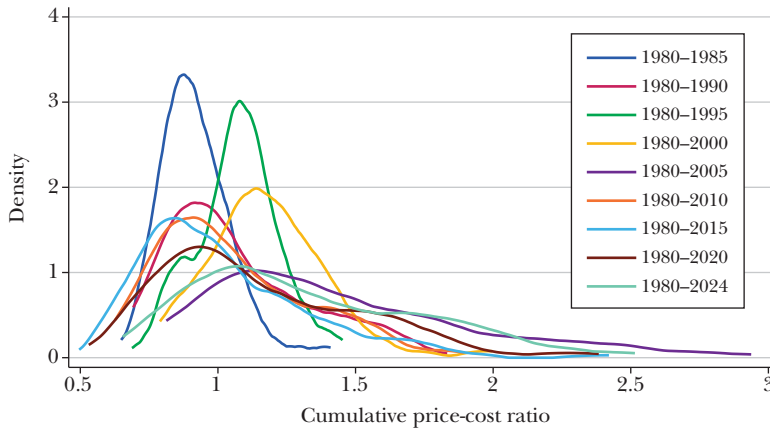
Source: Housing prices are from the FHFA data. Housing costs are from the RSMeans data. *N* is 128 cities for each time period.

overhead, but not "soft costs" like legal and engineering fees, financing costs, and land acquisition. The index therefore mostly captures costs tied to labor-centered services dominated by the wages of contractors and construction crews.

We can compare the growth rates of these two cost components to see if there have been systematic differences in the sources of building cost inflation since 2010. The top panel of Table 5 reports the compound annual growth rates of the total RSMeans index along with its materials and installation components.[9]

In the early 2010s, installation cost growth outpaced growth in material costs by roughly one percentage point per year. This pattern flipped in the latter half of the decade, with materials cost growth outpacing installation cost growth by about the same amount, and an overall acceleration in cost growth of about half a percentage point per year. The overall cost inflation rate doubled in the Covid and post-Covid periods. While both components saw price acceleration, inflation in materials far outpaced installation costs, consistent with the well-documented

---

[9] Unlike the US-wide index used in Figure 2, these components do not report a US overall average, but rather separate indexes for a large number of local areas. We construct US overall versions of each as follows. We first restrict the sample to the 30 cities that RSMeans uses to calculate the US average. (We verified that the across-city mean of the total cost index does in fact equal the reported US aggregate.) Then, for each of the 30 cities, we calculate the material and installation shares of total cost. We take the mean of the shares over these 30 cities for each year. (Over 2010–2024, installation accounts for 43 percent and materials for 57 percent of total building costs.) Finally, we multiply the RSMeans historical cost index for the US average by the material and installation shares for each year to create material and installation specific cost indexes.

Table 5
RSMeans Material and Installation Cost Components
Compound Annual Growth Rates, 2010–2024

	Total Cost	*Material*	*Installation*	
Panel A. US average				
2010–2015	2.35%	1.99%	2.83%	
2015–2020	2.81	3.30	2.16	
2020–2024	5.98	8.05	2.95	
	Material mean	*Material standard deviation*	*Installation mean*	*Installation standard deviation*
Panel B. City-level				
2010–2015	2.51%	0.38	3.67%	2.18
2015–2020	2.53	0.33	3.45	1.92
2020–2024	5.98	0.53	5.32	1.21

supply chain frictions and the accompanying inflation in goods relative to services (for example, di Giovanni et al. 2022; Comin, Johnson, and Jones 2024).

We can also compare these separate indexes for cities during this period since 2010. The bottom panel of Table 5 shows the mean and standard deviation of each index across the cities in our RSMeans data. The means of the city-level distribution look much like the aggregate growth rates, with a few differences.[10] During 2010–2015, the city-level means see the same pattern as in the aggregates of installation cost growth outpacing material cost growth by about one percentage point. From 2015 to 2020, rather than reversing this pattern as the aggregates do, city-level (unweighted) average installation costs continue to grow faster than material costs. As with the aggregates, inflation in both components accelerates in the post-Covid period, and more so for materials. However, the difference between average city-level materials and installation cost growth is smaller than for the aggregates.

An interesting pattern in the city-level numbers that was hidden in the aggregates is the difference in the standard deviations of the cost component growth rates. Installation cost growth varies much more across cities than material cost growth. Before Covid, the ratio of the two components' standard deviations is almost 6. This falls quite a bit after Covid, to around 2 (perhaps because supply chain disruption effects were heterogeneous across geographies), but a clear difference remains.

The greater variance of installation costs may reflect the relative goods intensity of the material index, as opposed to the labor intensity of the installation cost index. Goods are more tradeable, which creates an arbitrage incentive that reduces cross-sectional dispersion in prices. Labor for construction, on the other hand, is a much more local market with higher spatial adjustment costs, and as such it may

---

[10] Again, the median growth rates, not shown in Table 5, are quite close to the means, showing that the patterns are not driven by outliers.

be easier to for greater price dispersion to arise. It also suggests the primary source of any cross-sectional correlation between building costs and house prices—which again we show above is already rather weak, especially in this latter period—is likely to be because of labor costs.

Across-City Housing Cost and Price Level Comparisons

Having used the RSMeans housing cost index to look at cross-sectional relationships between building cost and house price *growth*, we now use data to explore cross-sectional relationships in building cost and price *levels*. In this case, of course, we cannot compare indexes, so we need different data. For building costs, we stick with RSMeans data products and use their construction cost estimator, which reports the expected building cost in dollars of a home with a particular set of attributes in a particular city. Specifically, we use a two-story home of average building quality, with stucco on wood frame walls/frame and floorspace of 2,000 square feet. Because the prices of different house types are highly correlated across cities, using a different set of attributes would likely yield similar results. For house price data, we use the Zillow median house price reported by city. We use data for 2024 and obtain building cost and house price measures for the 100 largest metro areas.[11]
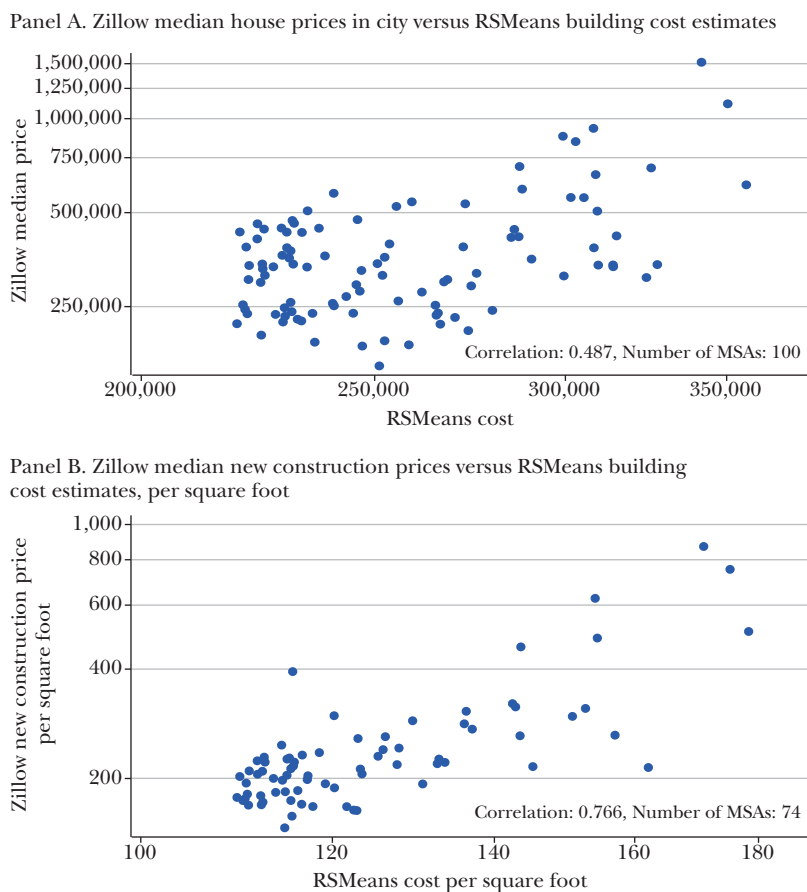
Figure 5, panel A, plots the log of the Zillow median house price against the log of the RSMeans building cost estimate. Costs and prices are positively correlated, with the slope from regressing log prices on log costs yielding a coefficient of 1.57 (standard error $= 0.33$) and an R^2 of 0.24. As with the correlations in growth rates across cities throughout our panel, costs have explanatory power, but the majority of variation is explained by other factors.[12]

As one further check on the housing cost-price relationship, we note our RSMeans dollar cost levels for every city are for a house with standardized attributes,

---

[11] Zillow data can be downloaded from https://www.zillow.com/research/data/. We use *Median Sale Price (Nowcast, All Homes, Monthly), New Construction Median Sale Price (Raw, All Homes, Monthly)*, and *New Construction Median Sale Price per Square Foot (Raw, All Homes, Monthly)*, all at the Metro and US-geography level. We then take the mean over all months in 2024 for each city to get the 2024 values. To further check data quality, we also gathered house price data from realtor.com (https://www.realtor.com/research/data) on metro-level median list prices and list prices per square foot. We download their monthly historical metro data and take the mean over all months in 2024 for each city. These are highly correlated with the Zillow data. The log of prices in the realtor.com and Zillow have a correlation coefficient of 0.97. Their per-square-foot counterparts have a correlation of 0.93.

[12] For details of the calculation, see Supplemental Appendix Table A3. This correlation reflects the relationships between RSMeans cost levels for a city and the median price of all houses in the city. If there is a systematic difference between new construction and the existing housing stock, RSMeans costs for a standardized house might not strongly correlate with overall prices, even if they predict the prices of new construction well. We can check for this possibility, as Zillow also reports median prices for new construction housing for 78 of the metros in our sample. Regressing logged prices on logged building costs in this case yields a coefficient of 2.00 (SE $= 0.31$) and an R^2 of 0.53. The corresponding scatterplot is shown in Supplemental Appendix Figure A2. Building costs thus do a better job explaining new construction prices than prices of existing housing. Flipped on its head, this relative inability to explain the prices of existing housing is yet another facet of the generally weak price-cost relationship.

Figure 5
Across-City Comparisons of Housing Prices and Building Costs

Panel A. Zillow median house prices in city versus RSMeans building cost estimates

Panel B. Zillow median new construction prices versus RSMeans building cost estimates, per square foot

Note: The Zillow data series are available at https://www.zillow.com/research/data/. The building cost data are from RSMeans.

including size. If the size of median homes differs systematically across cities, this could be another reason for a wedge between prices and building costs. To check this, we compare both Zillow prices (though just for new construction) and RSMeans cost estimates in terms of dollars per square foot for 74 of the 100 cities in our original sample. The regression coefficient of prices on costs using this per-unit-area data is 2.16 (SE = 0.33), with an R^2 of 0.59. The corresponding scatter plot is shown in Figure 5, panel B. Adjusting for the fact that RSMeans cost estimates focus on new construction and for differences in house sizes across metros substantially raises the ability of building costs to explain house prices.

That said, even in this best case, almost half of the variation in housing prices remains unexplained. Table 6 lists the top and bottom ten cities (of 74 in the

Table 6

Top Ten and Bottom Ten Residuals from Regressing City's Logged Median House Price per Square Foot on Logged RSMeans Building Cost per Square Foot Estimates, and Ranks of Price and Cost

City	House price residual (%)	Price Rank (of 74)	Cost Rank (of 74)
Top ten			
Miami, FL	98.2	7	46
San Jose, CA	88.5	1	3
Los Angeles, CA	69.5	3	7
San Francisco, CA	54.5	2	2
Seattle, WA	45.2	6	11
Denver, CO	38.7	12	35
Boston, MA	31.2	5	6
Colorado Springs, CO	27.3	21	56
Sarasota, FL	22.2	27	60
Durham, NC	21.1	31	66
Bottom ten			
El Paso, TX	−19.1	69	38
Augusta, GA	−20.6	73	47
Indianapolis, IN	−24.6	70	33
Columbia, SC	−25.1	74	54
Omaha, NE	−25.5	54	21
Baton Rouge, LA	−27.3	71	32
Greenville, SC	−27.9	72	31
Philadelphia, PA	−31.6	17	5
Minneapolis, MN	−33.7	38	10
Chicago, IL	−48.0	39	4

Note: The residuals are calculated from regressing the median per-square-foot price of newly constructed houses, using the Zillow data, on the estimated building costs per square foot, from the RSMeans data. The sample has 74 cities. For details, see Potter and Syverson (2025). The residuals were estimated in log points and then converted to percentages for this table.

sample) in terms of the residual of regressing the median per-square-foot price of newly constructed houses on the estimated building costs per square foot. Cities with positive residuals have high prices relative to those predicted by building costs, while cities with negative residuals have low prices relative to those predicted by building costs.

The highest relative house price cites are in coastal California, the Front Range of Colorado, and Florida, as well as Seattle and Boston. Miami is the highest, with prices per square foot are some 98 percent above what would be predicted by local construction costs.

It is interesting to compare the price and cost ranks of these high-residual cities among the 74 cities in our data. Most of the cities with high price-cost residuals exhibit high house prices overall. The seven highest-residual cities are among the twelve highest-price cities. However, with the exception of Miami and Denver, none of these cities have particularly low building costs. San Jose and

San Francisco (second- and fourth-highest house price residuals) actually have the second- and third-highest estimated building costs among all cities, and Boston, Los Angeles, and Seattle are not far behind. These cities therefore have a large price residual despite their building costs, not because of them. The remainder of the top ten—Colorado Springs, Colorado, Sarasota, Florida, and Durham, North Carolina, have prices out of the top ten but above the median. Their high price-cost residual reflects a large gap between their price and cost ranks rather than astronomical prices per se.

The lowest relative price cities are various locations in the Midwest and South as well as Philadelphia. Chicago has the smallest residual of all cities in our data, indicating its house prices are 48 percent less than predicted by building costs, notably larger than Minneapolis's second-lowest residual pointing to prices 34 percent less than that implied by costs.

Looking at these lowest-residual cities' price and cost ranks, we see a range of patterns. Some cities, like Augusta, Georgia, and Columbia, South Carolina, have very low house prices but also below-median costs—just not as low as their prices, obviously. Others have a negative residual not so much because of low prices, but rather quite high costs: Philadelphia, Minneapolis, and Chicago all have cost ranks in the top ten. Chicago's large negative residual is much more due to high costs than low prices.

Conclusion

Looking at nearly 75 years of RSMeans data has allowed us to examine housing construction cost growth over a long period and across many different metro areas. Comparing this with various measures of housing prices offers insights into the relationship between construction costs and housing prices. Despite construction costs being a very large fraction of the costs of new construction, overall construction costs are not a particularly strong predictor of housing prices overall, especially in the very highest-cost metros and in more recent periods. We observe this across a variety of measures, from the increasing divergence in construction costs and housing prices across major metros, to the larger dispersion of installation costs than material costs in building costs, to contemporary Zillow home price data on both new construction and existing houses.

Patterns in both aggregates and in the cross section suggest the cost-price relationship is weakening over time, particularly post-2000. Moreover, some of the relationship that does exist may be due to the circular effects of high housing costs driving up the costs of local labor, a major input to construction costs.

But while the stagnation (or worse, recession) of construction productivity and its concomitant high building costs in general can only add to high housing prices, these factors are mostly and increasingly overshadowed by other effects on house price. The decoupling of house price and replacement cost may have several drivers,

including increasingly binding restrictions that prevent housing supply expansion in high-demand areas.

References

Badger, Emily, and Eve Washington. 2022. "The Housing Shortage Isn't Just a Coastal Crisis Anymore." *New York Times*, July 14. https://www.nytimes.com/2022/07/14/upshot/housing-shortage-us.html.

Bundrick, Hal. 2024. "Why are home prices so high?" *Yahoo! Finance*, November 19. https://finance.yahoo.com/personal-finance/mortgages/article/why-are-house-prices-so-high-184935574.html.

Case, Karl E., and Robert J. Shiller. 2003. "Is There a Bubble in the Housing Market?" *Brookings Papers on Economic Activity* 34 (2): 299–362.

Casselman, Ben. 2024. "The Housing Market Is Weird and Ugly." *New York Times*, June 20. https://www.nytimes.com/2024/06/20/business/economy/housing-market-explained.html.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2018. "Housing Booms and Busts, Labor Market Opportunities, and College Attendance." *American Economic Review* 108 (10): 2947–94.

Comin, Diego A., Robert C. Johnson, and Callum J. Jones. 2024. "Supply Chain Constraints and Inflation." NBER Working Paper 31179.

D'Amico, Leonardo, Edward L. Glaeser, Joseph Gyourko, William R. Kerr, and Giacomo A. M. Ponzetto. 2024. "Why Has Construction Productivity Stagnated? The Role of Land-Use Regulation." NBER Working Paper 33188.

di Giovanni, Julian, Şebnem Kalemli-Özcan, Alvaro Silva, and Muhammed A. Yildirim. 2022. "Global Supply Chain Pressures, International Trade, and Inflation." NBER Working Paper 30240.

Duranton, Gilles, and Diego Puga. 2023. "Urban Growth and Its Aggregate Implications." *Econometrica* 91 (6): 2219–59.

Economist. 2024. "The House-Price Supercycle Is Just Getting Going." October 1. https://www.economist.com/finance-and-economics/2024/1%1/the-house-price-supercycle-is-just-getting-going.

Garcia, Daniel, and Raven Molloy. 2025. "Reexamining Lackluster Productivity Growth in Construction" *Regional Science and Urban Economics* 113: 104107.

Glaeser, Edward L., Joshua D. Gottlieb, and Joseph Gyourko. 2013. "Can Cheap Credit Explain the Housing Boom?" In *Housing and the Financial Crisis*, edited by Edward L. Glaeser and Todd Sinai, 301–16. University of Chicago Press.

Glaeser, Edward, and Jospeh Gyourko. 2018. "The Economic Implications of Housing Supply." *Journal of Economic Perspectives* 32 (1): 3–30.

Glaeser, Edward L., and Joseph Gyourko. 2025. "America's Housing Supply Problem: The Closing of the Suburban Frontier?" NBER Working Paper 33876.

Goolsbee, Austan, and Chad Syverson. 2024. "The Strange and Awful Path of Productivity in the US Construction Sector." In *Technology, Productivity, and Economic Growth*, edited by Susanto Basu, Lucy Eldridge, John Haltiwanger, and Erich Strassner. University of Chicago Press.

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2013. "Endogenous Gentrification and Housing Price Dynamics." *Journal of Public Economics* 100: 45–60.

Kakar, Venoo, Gerald Daniels, and Aditi Grossman. 2018. "Jobs and Housing (Im)balance in the San

Francisco Bay Area." Center for Applied Housing Research Working Paper 2018-3.

Lynch, Eric. 2023. "Cost of Constructing a Home—2022." National Association of Home Builders, February 1. https://www.nahb.org/-/media/NAHB/news-and-economics/docs/housing-economics-plus/special-studies/2023/special-study-cost-of-constructing-a-home-2022-february-2023.pdf?rev=f771d8b924b14a079010cd59da395406.

Mian, Atif, and Amir Sufi. 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis." *Quarterly Journal of Economics* 124 (4): 1449–96.

Potter, Brian. 2021. "60 Years of Homebuilding." Construction Physics, May 3. https://www.construction-physics.com/p/60-years-of-homebuilding.

Potter, Brian, and Chad Syverson. 2025. *Data and Code for: "Building Costs and House Prices."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E234046V1.

US GAO (Government Accountability Office). 2023. "The Affordable Housing Crisis Grows While Efforts to Increase Supply Fall Short." WatchBlog: Following the Federal Dollar, October 12. https://www.gao.gov/blog/affordable-housing-crisis-grows-while-efforts-increase-supply-fall-short.

International Dimensions of Housing Markets

Cristian Badarinza and Tarun Ramadorai

Discussions of housing markets have traditionally focused on domestic economic factors and policy decisions—such as government choices about zoning or construction permits, rules about what rents or prices can be charged, the types of mortgage contracts that can be offered, or the extent to which regional or national governments have provided or should provide incentives to increase the supply of new construction. Increasingly, however, prices and transactions seem to vary in ways that cannot easily be attributed to local economic factors. In this essay, we make the case that an international perspective on housing markets offers important insights that can help us understand some of these puzzling variations.

To set the stage, we first document substantial cross-country and cross-time variation in house price-to-income ratios. In a broad sense, if house prices are high enough relative to income levels that many people must compromise excessively on physical space, desired amenities, easy access to productive jobs, opportunities for human capital development, or their capacity to meet other essential living expenses, a housing market is considered out of reach for these buyers (Glaeser and Gyourko 2013; Molloy, Nathanson, and Paciorek 2022). We discuss evidence on the extent to which within-country factors can affect cross-country house

■ *Cristian Badarinza is an Associate Professor of Real Estate Finance at the Frankfurt School of Finance and Management, Frankfurt, Germany. Tarun Ramadorai is a Professor of Financial Economics at Imperial College London, and a Research Fellow at the Centre for Economic Policy Research, both in London, United Kingdom. Their email addresses are c.badarinza@fs.de and t.ramadorai@imperial.ac.uk.*

price-to-income ratios, such as construction, household incomes, the popularity of homeownership, and mortgage conditions. While these domestic economic factors clearly matter to housing markets, they also leave unexplained a significant fraction of the cross-country variation in house price-to-income ratios.

International influences on housing markets are widely discussed, often on the front pages of prominent media outlets. The housing market has become a central battleground in debates around globalization, with a number of cities and countries raising concerns—often stridently expressed—that international buyers and flows of capital reduce the affordability of housing for local residents. After all, if the supply of residential real estate is inelastic in response to demand increases (Glaeser and Gyourko 2018; Baum-Snow and Han 2024), foreign demand shocks could raise housing prices and make home-ownership inaccessible to domestic residents. We discuss a growing body of academic evidence that lends credibility to this point of view.

To further explore these issues, we turn to a search and matching framework of the housing market, which accounts for heterogeneity in buyer preferences as well as differences across sellers in their price-setting behavior. In a globalized economy, the search for higher-end real estate crosses international borders. The framework helps us to understand how, especially in higher-profile "target" cities around the world, the global distributions of buyer and seller circumstances can affect local house prices and transactions volumes. We argue that this perspective can be very helpful when studying the global market for housing. Put differently, the global distributions of wealth, risk, and productivity are important and understudied factors that can help to explain why in a growing number of cities and countries, housing affordability is not just a single-market phenomenon.

International Data on House Price-to-Income Ratios

A simple but meaningful housing market metric is the ratio of a property's price to the annual income of the household. One way to think of this ratio is that it approximates the number of years that the household would need to save all its income to be able to purchase the property outright; similar intuition applies to the duration of mortgage repayments required, were the household to finance the purchase using debt.

To compute this price-to-income ratio, we start with two nationally representative datasets that are routinely used to assess household balance sheets, namely, the Survey of Consumer Finance (SCF 2013–2022) in the United States, and the Household Finance and Consumption Survey (HFCS 2010–2021) in the Euro area. For each owner-occupier, the survey records annual gross income as well as an estimate of the current market value of their property. In 2001, when the median US household income was $52,423, and the median price of a residential unit housing a typical family was $123,000, the ratio of these two numbers gives a median-price-to-median-income ratio of roughly 2.5. Two decades later in 2022,

Figure 1
Evidence from Microlevel Survey Data
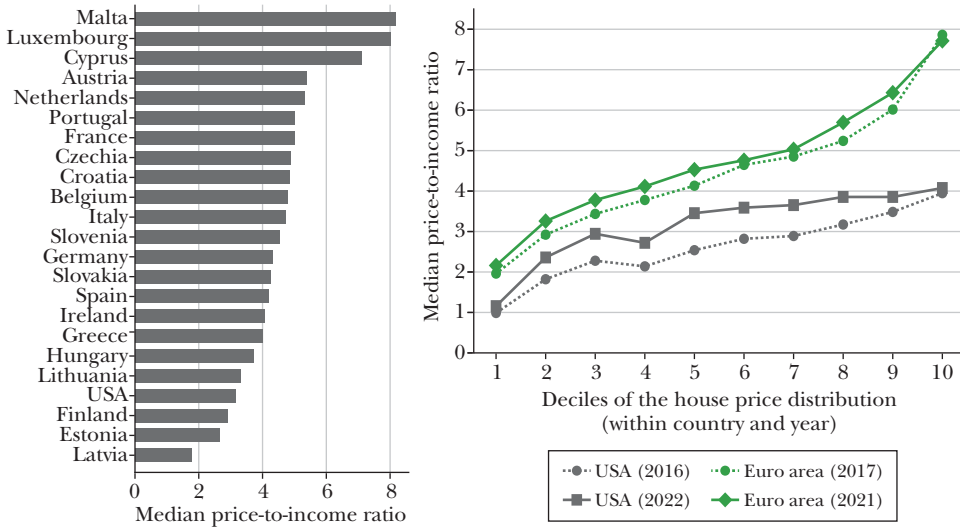
Source: Data are from the SCF (2001–2022) and HFCS (2010–2021) for the period post-2010.
Note: We calculate the price-to-income ratio for households which own their main residence as the ratio between the total value of the home and the total value of annual household income.

median household income increased to $95,120, but the median property price increased to $325,000, that is, the median-price-to-median-income ratio went up to 3.5.[1] Interpreted loosely, a typical US family in the 2020s would need to save or repay a mortgage for approximately 30 percent longer to afford a similar purchase, relative to their situation in the early 2000s.

Figure 1 repeats this calculation in the euro area, where we generally see higher median price-to-income ratios than in the United States; the left-hand plot shows the highest levels in "tax haven" locations such as Malta, Luxembourg, and Cyprus. To gain some understanding about the types of properties particularly subject to high price-to-income ratios, we calculate price-to-income ratios for properties ranked (in deciles) by their price levels in each country and each year. In both the United States and the euro area, the price-to-income ratio has primarily risen for properties that are not at the extremes of the price distribution. That said, price-to-income ratios are on average higher at the top of the price distribution,

---

[1] The corresponding median value of the ratio between the value of a household's main residence and their gross annual income is equal to 2.32 in 2001 and 3.16 in 2022.

consistent with intense competition for housing services in the high-price segment of these housing markets.

The micro data plotted in Figure 1 is only available for a small set of countries. For a broader analysis, we recruit aggregate data reported by the OECD (1960–2024), which cover a larger cross-section of countries over a longer time period. In Figure 2, the OECD price-to-income ratio is reported as a normalized quality-adjusted index relative to the year 2000, because absolute values of these quantities are not directly comparable across countries. This lack of comparability arises from significant differences in the composition of the housing stock between countries, including differences in house types—such as single-family homes versus apartments—and house sizes.[2] Panel A of Figure 2 plots data from this source and shows that in purely nominal terms house prices have generally seen an upward trend, though the steepness of this trend is very different across countries.

More importantly, the large variation in house price trends is also visible in cross-country differences in house price-to-income ratios, shown in panel B of Figure 2. In both panels, the colored lines show a representative set of six countries from different continents where price movements have been most extreme. In Canada, New Zealand, Australia, and Sweden, the price-to-income ratio has roughly doubled since the year 2000. In a smaller subset of countries (such as South Korea and Japan) price-to-income ratios have fallen between then and the most recent period, meaning that the rise in household incomes has more than compensated for the rise of property prices over the past two decades. [3]

Cross-Country Differences and Domestic Factors

Research on the determinants and consequences of house price movements has largely focused on domestic factors such as the rules governing new construction (Glaeser and Gyourko 2018; Baum-Snow 2023; Baum-Snow and Han 2024; Bartik, Gupta, and Milo 2025), political economy considerations such as local control over land use policies (Metcalf 2018), local economic factors including wages, the rate of residential construction, and pandemic-induced changes in demand (Van Nieuwerburgh and Weill 2010; Gupta et al. 2022; Molloy, Nathanson, and Paciorek 2022), the characteristics of mortgages (Green and Wachter 2005; Davis and Van Nieuwerburgh 2015), and preferences for homeownership (Sodini et al. 2023). While this

---

[2] Unlike the micro-level survey data that capture gross income, the OECD captures disposable income—that is, income net of taxes. For the purpose of our analysis, the distinction between per-capita and household income, as well as the difference between gross and net values, is not material. During the two decades that we consider, both family composition and the overall fiscal regime are fairly stable.

[3] Bricongne, Turrini, and Pontuch (2019) offer an alternative approach to compute comparable house price levels across countries, in absolute terms. Their methodology aims to account for cross-national differences in housing-market structures and other relevant factors, potentially enabling more direct comparisons of price levels. However, given the persistent and substantial discrepancies in absolute price-to-income levels between countries, our goal here is to discuss the evolution of this statistic over more recent years, rather than the specific initial levels observed in the year 2000. This allows for a more direct assessment of the factors driving the variation in pricing across different environments, without getting bogged down by inherent structural differences between country-specific housing stocks.

Figure 2
Cross-Country Aggregate Data

Panel A. Nominal house prices
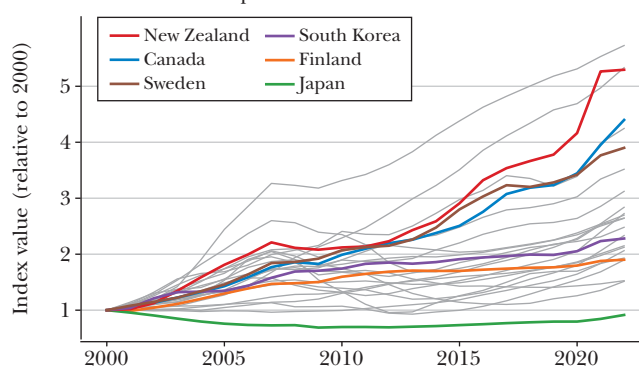
Panel B. Median price-to-income ratios
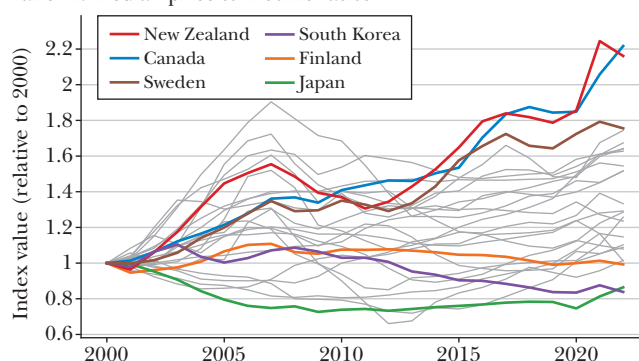
Source: Data for 52 countries obtained from OECD (1960–2024).
Note: For a consistent representation we restrict the view to the period 2000–2022 and to the set of countries for which values for all variables are consistently recorded between 2000–2022. We exclude Romania, Russia, Israel, and Türkiye due to outlier values recorded. We show five selected countries in color, and the rest using gray lines. The left plot shows quality-adjusted nominal house price indexes, and the right plot shows the evolution of the ratio between nominal house prices and disposable income per capita.

large body of work has greatly increased our understanding of the factors driving house price movements, explaining cross-country differences in housing prices in a systematic way remains a difficult analytical task.

Cross-country regressions including an array of possible explanatory variables present well-known difficulties, and in this situation do not provide a high level of explanatory power.[4] Indeed, it can be unclear even whether to expect that the

---

[4] In the Supplemental Appendix, we show for illustrative purposes a regression with the annual level of the (indexed) house price-to-income ratio as the dependent variable. The explanatory variables include the annual level of the (indexed) rent-to-income ratio, short- and long-term interest rates (measured

coefficients in such regressions should be positive or negative. For example, the share of housing financed by mortgages can shed light on the degree of financialization of a housing market. However, it is a priori unclear whether a higher mortgage share will make housing more or less affordable: in equilibrium, a higher mortgage share could also be associated with higher housing prices if credit access expands the number of prospective home buyers.

Yet another puzzling observation seen in the cross-country data relates to the wide dispersion in homeownership rates. Southern European countries such as Spain and Italy have relatively high homeownership rates, while Germany has a low homeownership rate, despite these markets having broadly similar price-to-income ratios. In Belgium, the price-to-income ratio is relatively low (controlling for other domestic factors), even though homeownership and mortgage indebtedness rates are comparable to those in Canada and the United States. Overall, despite the importance of this task, explaining cross-national dispersion in the levels of and trends in house price-to-income ratios is clearly not straightforward. We argue that studying cross-border flows can help us get closer.

The Effect of Foreign Capital

In recent years, a growing literature has shown important effects of international capital flows on housing markets. In the short run, the supply of residential real estate does not quickly respond to increases in demand, meaning that foreign demand shocks could reasonably be expected to raise housing prices relative to income levels.

Accurately estimating the price impact of foreign capital on the local housing market is generally difficult, however, because identification is hard with short time-series data that are afflicted by multiple confounding shocks. For this reason, this literature typically uses granular data and modern causal identification techniques to identify the effects of capital flows from abroad on domestic house prices. For example, Gyourko, Mayer, and Sinai (2013) consider US "superstar cities," which they define as having both an inelastic supply of housing and excess demand for housing—in part because these cities are an attractive destination for international capital. Badarinza and Ramadorai (2018) show that periods of elevated political risk abroad are associated with 2.5 percent higher prices in locations that are most likely to receive foreign capital inflows. Gorback and Keys (2020) look at how US housing markets with a high population of foreign-born Chinese were affected when a number of countries passed laws making it harder for Chinese investors to buy real estate in those countries, thus diverting this demand towards US residential real estate. Looking at overall property prices in Paris, Cvijanović and

---

using one- and ten-year government bond yields, respectively), the level of the population, the level of new construction in each country in each year, homeownership rates, and share of population with a mortgage. We use cross-country comparable variables reported by the OECD, calculated for each year from 2000 to 2022. Obviously, many questions can be raised about such regressions, and we do not relitigate those questions here. Taken at face value, this exercise shows that these domestic factors explain only around 40 percent of the cross-country, cross-time variation in house price-to-income ratios.

Spaenjers (2021) find that out-of-country demand shocks have a small but discernible effect. In California housing markets, Li, Shen, and Zhang (2024) focus on housing purchases by out-of-country Chinese buyers and document an increase in local housing prices and employment, accompanied by modest displacement of lower-income households. In Singapore, Liao et al. (2015) estimate a significant price impact of property purchases associated with foreign buyers in areas where such buyers are more prevalent, also documenting significant spillover effect in neighboring regions.

Faced with these developments, governments have not been passive bystanders; many countries have acted forcefully to limit the foreign ownership of domestic residential real estate in certain segments of the market. For example, in April 2023, Singapore imposed a buyer "stamp duty"—a transactions tax on buyers of residential real estate—of 60 percent for foreign residents. This eye-watering percentage essentially restricted foreign purchase activity to a very minor (and very well-heeled) segment of the market. Australia, Canada, and the United Kingdom have similar policies in place, allowing foreign capital to access their domestic housing markets, but imposing high costs on foreign purchases that accrue as tax revenues to the domestic fiscal authority, presumably to provide a disincentive for foreign purchases and to soften the blow of higher housing prices for domestic residents. In response to one such policy, Pavlov, Somerville, and Wetzel (2024) find that house prices declined by 6 percent in neighborhoods of Vancouver with above median concentrations of foreign buyers, relative to prices in neighborhoods with below median concentrations of foreign buyers. Similarly, in the United Kingdom, Johannesen, Miethe, and Weishaar (2022) find a strong price and volume response to the Brexit referendum, and Collin, Hollenbach, and Szakonyi (2025) document a further decrease in foreign investment activity, but only with a modest price impact, after the imposition of beneficial ownership transparency rules.

Direct Effects of Overseas Housing Demand

Purchases of real estate units by foreigners can have (at least) two underlying motives: an investment motive, which may involve a capital flow linked to diversification, capital preservation, or return chasing; or a physical migration of people accompanied by a capital flow. Either way, the economic mechanism through which the local market is directly affected by such purchases is best represented as a positive wealth shock.

Indeed, countries that choose to open their housing markets to foreign investment tend to experience inflows of financial wealth accompanying high-value house purchases. These can have positive effects on the destination market, potentially outweighing the negative effects of the house price and rental increases these transactions can generate. For example, Favilukis and Van Nieuwerburgh (2021) calibrate a structural model of housing markets using data from the 75 largest US cities. They focus on the effects of out-of-town buyers, both foreign and domestic. They find benefits including increased tax revenues and income growth, and the rise of high-value amenity provision which can affect local neighborhoods. However, as the

higher house prices and rents result in welfare gains for owners and losses for renters, the overall impact can turn out to be a welfare loss for the economy as a whole.

On the other hand, there may be beneficial effects of migration if foreigners contribute to the local labor market, perhaps by improving domestic labor productivity or by bringing in fresh entrepreneurial talent. Again, these impacts of migration can differ for people at different points on the income distribution (Combes et al. 2012; Dustmann, Frattini, and Preston 2013; Ghent and Steiner 2024; Howard, Wang, and Zhang 2024), and benefits arise strictly when the foreign (or out-of-town) buyers work and consume in their new location, but not if they are absentee owners. In that sense, these findings lend some credence to policies aimed at limiting buyer and seller speculation, such as limits on second-home purchases.

Are the absolute magnitudes of purchases by foreign buyers sufficient to rationalize substantial cross-region price variation? At the level of the aggregate national market, the answer is generally "no." In England and Wales, 0.7 percent of all registered property deeds are registered to individuals with an overseas correspondence address (Powell-Smith 2021). This corresponds to a significant increase from the 0.2 percent level in 2010, but taken together with the properties registered by foreign corporations (valued at a total of £70–90 billion), total foreign ownership still accounts for just 1.1 percent, out of which a significant part is attributed to individuals and corporations residing in the United Kingdom's own offshore islands. That said, this low aggregate figure masks significant heterogeneity, with the foreign investor share reaching 10 percent and above in some parts of London (Bomare and Le Guern Herry 2025; Johannesen, Miethe, and Weishaar 2022). Similarly, in the case of Vancouver, another prime target of international real estate investors, the foreign investor share is generally above 1 percent (Pavlov, Somerville, and Wetzel 2024), with variation through time and across locations up to a level of 14 percent. In Paris, nonresident foreigners accounted for 2.8 percent of purchases in 2022 (Paris Property Group 2023). In the US market, 1.3 percent of the existing-home sales went to a foreign buyer in 2024, down from an average level of 4.5 percent between 2010 and the COVID period (National Association of Realtors 2024), and heavily concentrated in the states of Florida, California, and Texas. Given the generally modest overall transaction volumes associated with foreign direct investment in the housing market, the question is whether these capital flows can generate large effects in the destination market; if not, it is difficult to rationalize the significant policy focus on these flows. To understand this better and to guide further empirical analysis, we structure our discussion around a theoretical framework.

Insights from the Search and Matching Framework

Search and matching models have become increasingly popular in the analysis of housing markets (Han and Strange 2015; Landvoigt, Piazzesi, and Scheider 2015; Guren and McQuade 2020). For present purposes, an advantage of the search and matching model is that it provides a structured way of thinking about international

homebuyers and suggests additional factors to explain differences across housing markets.

For any potential buyer, the road to homeownership runs through a search problem. Even if a household has the necessary funds to purchase a house, the household will need to find a set of properties that fit its preferences, inspect these properties to understand their unique amenities, contact sellers, react quickly and strategically to potentially competing offers by other buyers, and finalize negotiations. Different households value a given property differently, as their preferences for physical characteristics and local amenities likely differ. For example, some buyers may prefer a larger home outside the central city, while others prefer a home in the middle of the city. Even within the same neighborhood, a property that is closer to a primary school or a park, or has more space between neighbors or a larger garden, may appeal to different homebuyers.

How do sellers figure in this search and matching process? Initially, the literature assumed that sellers randomly list properties for sale and passively wait for buyers. If a buyer expresses interest, a brief negotiation ensues, a transaction price emerges, and the property changes hands. In more recent models, the likelihood that a given property is listed for sale is linked to the payoff that the seller expects to obtain from moving, say, for example, because a more rosy professional opportunity becomes available in a different location. Starting from a baseline calculation of the property's "market value," a seller with a high payoff from moving will adjust the asking price down, aiming for a quick sale; conversely, a seller with a low payoff from moving will be more aggressive in their pricing strategy and more patient. When estimating "market value," sellers (often assisted by real estate agents) must form an understanding of the characteristics that interest buyers and form expectations about their willingness to pay.

A first implication of the search and matching model is that small changes in the very highest anticipated buyer valuations can have large effects on property prices. There are a number of reasons for this. First, sellers price properties taking the top of the market into account, that is, they strategically choose listing prices for their houses, understanding that waiting, though costly, could pay off in the form of a buyer that greatly values their house (Andersen et al. 2022; Badarinza et al. 2025). Put differently, while it is a natural assumption that buyers search for properties, in the housing market, sellers also search for buyers. This means that the distribution of buyer valuations matters—it can pay off to imagine the panoply of potential buyers. Even a rumor that buyers with high willingness-to-pay are lurking in the market can be sufficient to make sellers strategically "fish" for buyers at the cost of waiting a little longer (and potentially revising the price downwards, if that becomes necessary down the road).

With sellers pricing properties strategically in this way, an increase in buyer valuations that affects even a small fraction of the buyer pool can have large consequences. When a local market is more open—that is, when potential buyers come from a more diverse set of locations and countries—this will be reflected in the ultimately realized set of transactions. If sellers are sufficiently motivated to wait to sell, and many buyers cannot afford the high prices that sellers are willing to accept,

transaction volumes can drop and high prices can coexist with high vacancy rates, in turn leading to longer times-on-the-market for both buyers and sellers. In this case the willingness-to-pay of local buyers will also eventually increase because their outside option (that is, the next-best alternative property available on the market) changes, especially if they cannot afford to wait around before buying a property.

Second, while the preference for certain locations and amenities varies across population groups, the supply of high-end properties is limited in the short run (Courant 1978; Van Nieuwerburgh and Weill 2010). Higher foreign demand at the very top end of the housing market therefore means that more potential buyers compete within this segment, generating upward pressure on prices. If other slightly less desirable properties are available, these can experience "spillover" price appreciation if the population of high-end prospective buyers is eventually pushed towards these segments.

Third, participants in the housing search and matching process—buyers, sellers, and the agents that often counsel and represent them in negotiations—often use a comparable set of recent transactions to calculate the "market value" that serves to guide achievable outcomes from any given transaction. Especially in a market that is either fundamentally thin (for example, because properties are heterogeneous and owner preferences are highly segmented as in Piazzesi et al. 2020), or made thin by sellers' strategic behavior, prices can exhibit significant momentum (Guren 2018; Badarinza et al. 2025). Failed transactions are kept private, and successful ones are advertised or published. The realized set of transactions can therefore disproportionally represent the valuations of high-value buyers, in turn biasing perceptions and estimations of value going forward.

While this logic seems compelling, is there empirical evidence that these models could offer a useful framework through which to view the data? We do know that most prospective house buyers search in relatively narrow segments of the housing market. Thus, the model suggests the potential importance of measuring and analyzing house price-to-income ratios at a more granular, subnational level. For example, Figure 3 plots the variation of house price-to-income ratios within countries, for both the United States and the United Kingdom. There is clearly considerable within-country cross-location variation evident in the data. In both countries, the house price-to-income ratios are higher in many more economically dense and prosperous/productive areas, such as major cities in California or leading international cities like London and New York. While this certainly is not conclusive evidence, this pattern is consistent with a search and matching model. In such areas, prospective (foreign and domestic) buyers have deep pockets and a willingness to pay for properties offering high-quality amenities. We would also expect the gains for prospective sellers to be modest, because the payoff from moving away from these highly productive areas is likely to be relatively low.

The Global Wealth Distribution Matters

But why would foreign capital be associated with higher valuations in the first place? In each country, the portion of the income and wealth distribution that is

Figure 3

Within-Country Variation in House Price-to-Income Ratios: US Counties and UK Local Authority Districts
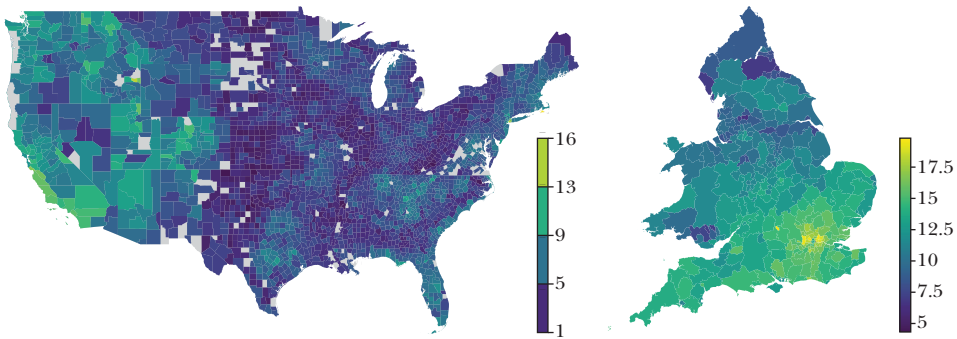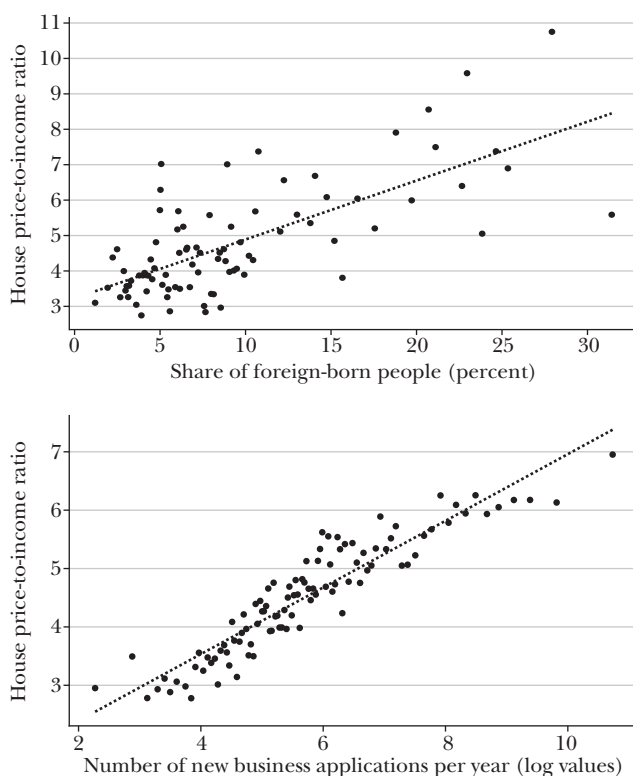

Source: Data for the United States are from Zillow (2000–2025) and the US Bureau of Economic Analysis (2020–2022), and those for the United Kingdom are from HM Land Registry (2021) and the UK Office for National Statistics (1997–2021).
Note: The figure shows the distribution of price-to-income ratios across US counties and local authority districts of England and Wales. We use the Zillow median home value reported at county level, and the gross income level per capita reported by the US Bureau of Economic Analysis for the year 2022. We use the universe of residential property transactions from the Land Registry to compute median transaction prices, and the estimated gross income per capita at the local authority level from the UK Office for National Statistics for the year 2021.

most globally mobile tends to be near the top. For the richest individuals in the world, foreign residential real estate has increasingly become a natural asset to acquire. Therefore, the global distribution of wealth and variations in this distribution will matter for the level of house prices in certain localities, in addition to local inequality. Badarinza, Balasubramaniam, and Ramadorai (2019) and Campbell and Ramadorai (2025) plot the global wealth distribution, and show that the very richest households from China, India, and South Africa are nearly as rich as the very richest US households. If a particular US location is an attractive destination for wealthy Chinese, Indians, or South Africans, prices in this location will be affected by the interest of such high-valuation overseas investors. In further support of this conjecture, Figure 4 provides suggestive evidence from the United States that a higher share of foreign-born people is generally associated with a higher price-to-income ratio.

Studies from other places reveal qualitatively similar findings. In documenting the ownership characteristics of Norwegian households in offshore real estate in Dubai with large-scale administrative data, Alstadsæter et al. (2022) show that while geographical proximity and historic ties appear to be key determinants of cross-border investment, the probability of investing in foreign real estate rises with wealth and peaks at the very top of the wealth distribution. Partially, this pattern is encouraged by local immigration regulations in Dubai and in other countries as well, such as the UK's Tier 1 "golden visa" program, which offered a path to residency in exchange for investments of at least two million pounds—with residential real

Figure 4
US House Price-to-Income Ratios by Fraction of Foreign-Born Population and New Business Applications

estate as a permissible investment vehicle with relatively low transaction costs (and potentially higher liquidity than private businesses).

When purchasing properties in unfamiliar markets, foreign investors can also be prone to overbidding, as they are likely to face more pronounced information asymmetries. In the context of the Paris housing market, Cvijanović and Spaenjers (2021) show that "out-of-country" buyers tend to be wealthier and to negotiate less aggressively on price, with the ultimate consequence that they buy at higher prices and sell at lower prices than domestic buyers. For the case of Singapore, Agarwal, Chia, and Sing (2020) estimate the price premium for foreign investors to be around 3.7 percent and mostly associated with transactions in cash.

An important role here is played by intermediaries—in this case, real estate agents. The quality of information available to a prospective investor, as well as their purchasing power, is likely manifest in their agent selection. Agents often specialize in particular property types, and in most countries, garner commission fees that are a share of the contract price. The proliferation of television reality shows on glamorous high-end international real estate transactions illustrates how this organization of the market could increase the competition among foreign buyers in the top segments of the market.

In addition, in a search and matching setting, the dimension of time is critical. Foreign investors may desire lower search times, because they are unwilling to spend extended amounts of time in temporary accommodations, and therefore they may be prepared to pay higher prices given their relative impatience. For sellers, this can be a tempting proposition, as they have an opportunity to sell fast and for a high price. Social networks can also play a role in the spatial transmission of housing-market information (Bailey et al. 2018), potentially also across borders. If real estate holdings signal social status (Badarinza 2019), the preferences of both domestic and foreign investors can be shaped by (wealth) shocks that happen outside of their country of origin, affecting their propensity to own versus rent, but also changing the perceived desirability of specific locations, amenities, or types of properties.

The Geography of Risk Matters

If patterns of foreign real estate investment were evenly distributed across the world, they might not generate disproportionate distortions of prices in domestic housing markets. In an open global economy with proportionately distributed cross-border real-estate investment, the richest US households and the richest South African households could hypothetically invest in each others' property markets without greatly discernible impacts. In reality, however, international capital flows are far more unevenly spread. Some countries are more likely to originate real-estate capital flows, while others are more frequent net recipients of such flows. There are good reasons for this; some fundamental determinants of the origins and destinations of capital flows include the quality of countries' political and economic environments, the levels of political risk prevailing in different locations, the risks of disenfranchisement or expropriation in specific jurisdictions, and international differences in the reliability of the rule of law (Badarinza, Ramadorai, and Shimizu 2022).

At the other extreme, we might imagine a unipolar situation in which all of the world's richest households allocate their capital to a single safe-haven market that is "the playground of the rich," which would then become unaffordable for local residents. While the past few decades have not been quite so extreme—with the visible exceptions of Monaco and parts of the Middle East—patterns consistent with this scenario have been identified in the data. Focusing on residential properties in London, Badarinza and Ramadorai (2018) find that neighborhoods of London that are more densely populated by specific foreign-born nationalities—potentially attractive capital flow destinations for those of the same nationality—experience visible price appreciation during periods of elevated risk in the related origin

country. This effect also operates on a more aggregate level, as they show by distinguishing between a period of relative calm in the global economy (2001 to 2006), and a later period of greater global turmoil (2007 to 2012); price appreciation is more pronounced in the latter period when foreign capital is "pushed" towards these locations from riskier foreign destinations. Comparing residential properties that are owned by a foreign entity with the local fraction of people born outside the United Kingdom, Sá (2025) also shows evidence consistent with London being viewed as a safe haven over this period.

This tendency of global capital to gravitate towards real estate as a safe haven has only strengthened in recent years. For example, Bomare and Le Guern Herry (2025) find that real estate accounts for a growing share of offshore portfolios, partly due to its exclusion from global transparency reporting requirements. This is reflected in the lower rent-to-price ratios seen in the segments of the market that appeal to such investors (Damen, Korevaar, and Van Nieuwerburgh 2025). It is not surprising that international capital seeking a safe haven demands a relatively low required rate of return, or in other words, such investors are prepared to pay higher prices for safe "boltholes" than economic conditions and the local utility of housing services appear to justify.

More generally, this mechanism is symptomatic of a direct link between broader risk-return tradeoffs in global financial markets and local real estate prices. Bednarek et al. (2021) and Boddin et al. (2024) document evidence from European countries that developments in global bond markets lead to portfolio rebalancing at the individual level, changing investors' required yields (rent-to-price ratios) and potentially leading to increases in property values. Considering the fact that between one-third and one-half of the balance sheet of representative households in developed economies is invested in financial assets (Badarinza, Campbell, and Ramadorai 2016), such effects can be material, with financial asset movements directly affecting households' propensity to enter real estate investment as well as their required rate of return on these investments.

Finally, the global distribution of risk matters for pricing mortgage debt, also affecting the real estate market through this channel. As real estate assets are generally highly collateralized, mortgage interest rates are a key driver of valuations both for owner-occupiers and for "buy-to-let" investors. This is certainly true within countries (Mian, Strain, and Sufi 2020), but also strikingly across countries (Mian and Sufi 2009; Sá and Wieladek 2015), because the savings behavior and financial product choice of any given individual is directly linked to another (potentially very distant) person's cost of borrowing through the structure of modern global capital markets. A case in point is the "savings glut" of Chinese households, which has famously been linked to the US housing boom of the early 2000s, a central element of the resulting financial crisis of 2008.

The Global Distribution of Productivity Matters

On the supply side of the housing market, the research literature has often focused on the dynamics of new construction, which ideally should respond to and at

least partially absorb the pricing impact of movements in demand. However, in most locations around the world, at least in the short-to-medium run, the overwhelming majority of available residential housing supply available for purchase comes from the secondary market—that is, the set of existing homes. For the owners (and potential sellers) of these homes, two critical determinants of the benefit of moving are the productivity of the economy and the degree of labor mobility (Head, Lloyd-Ellis, and Sun 2014). When labor productivity is high or increasing, mobile households can unlock value by exploiting profitable relocation opportunities. The direct benefit of such moves accrue primarily to the household making them, but their labor mobility also increases the available supply of properties on the market where the home is being sold, with a corresponding shift in bargaining power to prospective home-buyers.

In a setting with search and matching frictions, a higher level of labor mobility also increases the diversity of available (idiosyncratic) housing units and housing amenities. This can allow for improvements in match quality—the match between buyer desires and housing available for purchase—creating a more fluid and active market with lower waiting times for both sales and purchases. Instead of competing for a restricted number of units in a small segment of the market, buyers can cast a wider net, which increases the overall efficiency of the search and matching process.

But at the core of this mechanism (and a potential reason for why it breaks down) is the question of the distribution of labor productivity across regions and across countries. If productive opportunities are efficiently distributed across space (Fajgelbaum and Gaubert 2020), the benefits of efficient matching in the labor market are reflected in the housing market through increased mobility. When the geographical distribution of productivity becomes polarized, problems can arise: in low-productivity environments, the payoff from moving to higher environments rises, and housing-market congestion results because many sellers are willing to sell; in high-productivity environments, owners will be less willing to move, and even when they do, they may "fish" for higher house sale prices, thus creating housing-market congestion on the buy side of the market. Suggestive of these effects, Figure 4 shows the strong relationship between US price-to-income ratios and the dynamics of the local business environment, as reflected in the number of firm registration applications in US counties.

Another important aspect of this issue is that the magnitude of the payoff from moving has a strong life-cycle component—it is generally high early in life, and decreases as people approach retirement (Halket and Vasudev 2014). Thus, younger members of the workforce who seek to buy in locations where retirees do not wish to sell can face significant hurdles in accessing housing services. This issue can have significant demographic consequences during a phase of life when high-quality housing services serve as a critical input to family formation (Dettling and Kearney 2014; Charles, Hurst, and Notowidigdo 2018; van Doornik et al. 2024).

Transaction Chains in the Housing Market Can Amplify These Effects

The effects of foreign real estate investment in higher-end properties in desirable cities are by now well-established. In an arithmetic sense, driving up

house price-to-income ratios in a specific segment of certain urban markets will also increase the national house price-to-income ratio. But in addition to the global distributions of risk, productivity, and wealth mattering, as we have argued, the search and matching framework posits another mechanism through which the effects of foreign demand can be amplified. The housing market is linked by what we call "transaction chains," through which direct effects on one segment of the housing market can spread to other parts of the market.

To understand this mechanism, a specific example may be instructive. Proposals are often entertained to reduce house price-to-income ratios in urban areas by having the local government alter its zoning or other rules to allow substantially more building. If builders are not otherwise constrained, they will likely respond by building primarily in high-margin/high-end segments of the market—which at face value seems as if it would blunt the effect of the policy on middle- and low-income households that cannot afford the highest-price housing. However, it is a standard finding in studies that use structural models of housing markets that additional building and additional sellers at the high-end of the market reduce the extent of competition among potential high-end buyers (for example, as discussed in this symposium by Abramson and Landvoigt). As most buyers of one property are sellers of another, usually in a lower market segment (Grindaker et al. 2021; Moen, Nenov, and Sniekers 2021), if greater construction of properties at the highest quality level leads to diminished bidding for these properties, it will also mean diminished bidding for properties at the next-highest property levels, and so on, until in the medium-run greater supply "at the top" influences the market as a whole.

This same logic applies, albeit in the reverse direction, to a surge of foreign demand for upper echelon properties in desirable urban markets. That is, greater bidding pressure in these market segments will also mean greater bidding for properties at the next-highest property levels, and so on, until in the medium-run house price-to-income ratios can rise for the market as a whole.

Such transaction chains of housing can connect seemingly disparate housing markets, especially in combination with a range of other financial and behavioral frictions affecting housing transactions. Consider, for example, someone living in a market where nominal house prices have stayed flat, or even fallen, with a resulting compression in their available home equity. If an opportunity to move to a higher-productivity but more expensive area comes along, they may be forced to pass it up if they are financially constrained, especially if the move requires compromise, say, on house size (Stein 1995). If the potential seller has a fixed rate mortgage (common in countries such as the United States and Germany), but interest rates have risen, then taking out a new mortgage can appear even less attractive, creating mortgage lock-in (Fonseca and Liu 2024). And even in countries where mortgage contract terms are more flexible (such as the United Kingdom, where adjustable-rate mortgages are common) banks can and often do impose high prepayment penalties that can make a mortgage contract difficult to exit (Badarinza, Campbell, and Rama-dorai 2018). To make matters worse, people are very reluctant to sell properties at

nominal losses relative to their original purchase prices, creating "behavioral lock-in" (Genesove and Mayer 2001; Bracke and Tenreyro 2021; Andersen et al. 2022; Badarinza et al. 2025).

These factors mean that when a surge of foreign homebuyers contributes to elevated housing valuations in particular cities, the effects can spread to other areas of the country. Higher price-to-income ratios can coexist with a lower overall supply of vacant properties and a decreased ability of potential buyers to match to desirable properties. Shocks affecting one portion of the housing market can, through chains of (failed or consummated) housing transactions, percolate through the entire market.

What Lies Ahead

Although much of our discussion has emphasized the importance of taking an international perspective when viewing developments in real estate markets, it is clear that national and local factors continue to play a major role in housing markets. That said, we believe that the international dimensions of housing markets have been relatively under-researched. There is substantial work to be done on measuring and identifying the underlying factors driving cross-country or even cross-regional differences in house price-to-income ratios. And more needs to be done to identify the specific channels through which the rise of international buyers directly and indirectly affects various segments of global housing markets. We believe that search and matching models will continue to provide useful guidance on factors to investigate, including the distributions of global productivity, inequality, and risk, as well as the determinants of cross-border mobility and openness to foreign capital, and we look forward to more work on this important topic.

References

Agarwal, Sumit, Liu Ee Chia, and Tien Foo Sing. 2020. "Straw Purchase or Safe Haven? The Hidden Perils of Illicit Wealth in Property Markets." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.3688752.

Alstadsæter, Annette, Bluebery Planterose, Gabriel Zucman, and Andreas Økland. 2022. "Who Owns Offshore Real Estate? Evidence from Dubai." EU Tax Observatory Working Paper 1.

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx, and Tarun Ramadorai. 2022. "Reference Dependence in the Housing Market." *American Economic Review* 112 (10): 3398–440.

Badarinza, Cristian. 2019. "Mortgage Debt and Social Externalities." *Review of Economic Dynamics* 34: 43–60.

Badarinza, Cristian, Vimal Balasubramaniam, and Tarun Ramadorai. 2019. "The Household Finance Landscape in Emerging Economies." *Annual Review of Financial Economics* 11: 109–29.

Badarinza, Cristian, Vimal Balasubramaniam, and Tarun Ramadorai. 2024. "In Search of the Matching

Function in the Housing Market." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4594519.

Badarinza, Cristian, John Y. Campbell, and Tarun Ramadorai. 2016. "International Comparative Household Finance." *Annual Review of Economics* 8: 111–44.

Badarinza, Cristian, John Y. Campbell, and Tarun Ramadorai. 2018. "What Calls to ARMs? International Evidence on Interest Rates and the Choice of Adjustable-Rate Mortgages." *Management Science* 64 (5): 2275–88.

Badarinza, Cristian, and Tarun Ramadorai. 2018. "Home Away from Home? Foreign Demand and London House Prices." *Journal of Financial Economics* 130 (3): 532–55.

Badarinza, Cristian, and Tarun Ramadorai. 2025. *Data and Code for: "International Dimensions of Housing Markets."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E233721V1.

Badarinza, Cristian, Tarun Ramadorai, and Chihiro Shimizu. 2022. "Gravity, Counterparties, and Foreign Investment." *Journal of Financial Economics* 145 (2A): 132–52.

Badarinza, Cristian, Tarun Ramadorai, Juhana Siljander, and Jagdish Tripathy. 2025. "Behavioral Lock-In: Housing Market Taxation with Reference Dependent Agents." Preprint, SSRN. https://dx.doi.org/10.2139/ssrn.4693047.

Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018. "The Economic Effects of Social Networks: Evidence from the Housing Market." *Journal of Political Economy* 126 (6): 2224–76.

Bartik, Alexander, Arpit Gupta, and Daniel Milo. 2025. "The Costs of Housing Regulation: Evidence from Generative Regulatory Measurement." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4627587.

Baum-Snow, Nathaniel. 2023. "Constraints on City and Neighborhood Growth: The Central Role of Housing Supply." *Journal of Economic Perspectives* 37 (2): 53–74.

Baum-Snow, Nathaniel, and Lu Han. 2024. "The Microgeography of Housing Supply." *Journal of Political Economy* 132 (6): 1897–946.

Bednarek, Peter, Daniel Marcel te Kaat, Chang Ma, and Alessandro Rebucci. 2021. "Capital Flows, Real Estate, and Local Cycles: Evidence from German Cities, Banks, and Firms." *Review of Financial Studies* 34 (10): 5077–134.

Boddin, Dominik, Daniel Marcel te Kaat, Chang Ma, and Alessandro Rebucci. 2024. "A Housing Portfolio Channel of QE Transmission." NBER Working Paper 32211.

Bomare, Jeanne, and Ségal Le Guern Herry. 2025. "Avoiding Transparency through Offshore Real Estate: Evidence from the United Kingdom." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.5260099.

Bracke, Philippe, and Silvana Tenreyro. 2021. "History Dependence in the Housing Market." *American Economic Journal: Macroeconomics* 13 (2): 420–43.

Bricongne, Jean-Charles, Alessandro Turrini, and Peter Pontuch. 2019. "Assessing House Prices: Insights from 'Houselev,' a Dataset of Price Level Estimates." European Commission European Economy Discussion Paper 101.

Campbell, John Y., and Tarun Ramadorai. 2025. *Fixed: Why Personal Finance Is Broken and How to Make It Work for Everyone.* Princeton University Press.

Charles, Kerwin Kofi, Erik Hurst, and Matthew J. Notowidigdo. 2018. "Housing Booms and Busts, Labor Market Opportunities, and College Attendance." *American Economic Review* 108 (10): 2947–94.

Collin, Matthew, Florian M. Hollenbach, and David Szakonyi. 2025. "The End of Londongrad? Ownership Transparency and Offshore Investment in Real Estate." EU Tax Observatory Working Paper 28.

Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux. 2012. "The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection." *Econometrica* 80 (6): 2543–94.

Courant, Paul N. 1978. "Racial Prejudice in a Search Model of the Urban Housing Market." *Journal of Urban Economics* 5 (3): 329–45.

Cvijanović, Dragana, and Christophe Spaenjers. 2021. "'We'll Always Have Paris': Out-of-Country Buyers in the Housing Market." *Management Science* 67 (7): 4120–38.

Damen, Sven, Matthijs Korevaar, and Stijn Van Nieuwerburgh. 2025. "An Alpha in Affordable Housing?" Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.5121139.

Davis, Morris A., and Stijn Van Nieuwerburgh. 2015. "Housing, Finance, and the Macroeconomy." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 753–811. North-Holland.

Dettling, Lisa J., and Melissa S. Kearney. 2014. "House Prices and Birth Rates: The Impact of the Real Estate Market on the Decision to Have a Baby." *Journal of Public Economics* 110: 82–100.

Dustmann, Christian, Tommaso Frattini, and Ian P. Preston. 2013. "The Effect of Immigration along the Distribution of Wages." *Review of Economic Studies* 80 (1): 145–73.

Fajgelbaum, Pablo D., and Cecile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *Quarterly Journal of Economics* 135 (2): 959–1036.

Favilukis, Jack, and Stijn Van Nieuwerburgh. 2021. "Out-of-Town Home Buyers and City Welfare." *Journal of Finance* 76 (5): 2577–638.

Fonseca, Julia, and Lu Liu. 2024. "Mortgage Lock-In, Mobility, and Labor Reallocation." *Journal of Finance* 79 (6): 3729–72.

Genesove, David, and Christopher Mayer. 2001. "Loss Aversion and Seller Behavior: Evidence from the Housing Market." *Quarterly Journal of Economics* 116 (4): 1233–60.

Ghent, Andra C., and Eva Steiner. 2024. "The Best Cities for Firms." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4563942.

Glaeser, Edward, and Joseph Gyourko. 2018. "The Economic Implications of Housing Supply." *Journal of Economic Perspectives* 32 (1): 3–30.

Glaeser, Edward L., and Joseph Gyourko. 2013. "How Do We Know When Housing Is 'Affordable.'" In *The Affordable Housing Reader*, edited by J. Rosie Tighe and Elizabeth J. Mueller, 111–15. Routledge.

Gorback, Caitlin S., and Benjamin J. Keys. 2020. "Global Capital and Local Assets: House Prices, Quantities, and Elasticities." NBER Working Paper 27370.

Green, Richard K., and Susan M. Wachter. 2005. "The American Mortgage in Historical and International Context." *Journal of Economic Perspectives* 19 (4): 93–114.

Grindaker, Morten, Artashes Karapetyan, Espen Moen, and Plamen Nenov. 2021. "Transaction Sequence Decisions and Housing Market Volatility: Evidence from Norway." Unpublished.

Gupta, Arpit, Vrinda Mittal, Jonas Peeters, and Stijn Van Nieuwerburgh. 2022. "Flattening the Curve: Pandemic-Induced Revaluation of Urban Real Estate." *Journal of Financial Economics* 146 (2): 594–636.

Guren, Adam M. 2018. "House Price Momentum and Strategic Complementarity." *Journal of Political Economy* 126 (3): 1172–218.

Guren, Adam M., and Timothy J. McQuade. 2020. "How Do Foreclosures Exacerbate Housing Downturns?" *Review of Economic Studies* 87 (3): 1331–64.

Gyourko, Joseph, Christopher Mayer, and Todd Sinai. 2013. "Superstar Cities." *American Economic Journal: Economic Policy* 5 (4): 167–99.

Halket, Jonathan, and Santhanagopalan Vasudev. 2014. "Saving Up or Settling Down: Home Ownership over the Life Cycle." *Review of Economic Dynamics* 17 (2): 345–66.

Han, Lu, and William C. Strange. 2015. "The Microstructure of Housing Markets: Search, Bargaining, and Brokerage." In *Handbook of Regional and Urban Economics*, Vol. 5, edited by Gilles Duranton, J. Vernon Henderson, and William C. Strange, 813–86. North-Holland.

Head, Allen, Huw Lloyd-Ellis, and Hongfei Sun. 2014. "Search, Liquidity, and the Dynamics of House Prices and Construction." *American Economic Review* 104 (4): 1172–210.

HFCS. 2010-2021. *Household Finance and Consumption Survey. European Central Bank. Waves 1-4.*

HM Land Registry. *Price Paid Data.* 2021. https://www.gov.uk/government/statistical-data-sets/price-paid-data-downloads (accessed June 18, 2025).

Howard, Troup, Mengqi Wang, and Dayin Zhang. 2024. "Cracking Down, Pricing Up: Housing Supply in the Wake of Mass Deportation." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4729511.

Johannesen, Niels, Jakob Miethe, and Daniel Weishaar. 2022. "Homes Incorporated: Offshore Ownership of Real Estate in the UK." CESifo Working Paper 10159.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider. 2015. "The Housing Market(s) of San Diego." *American Economic Review* 105 (4): 1371–407.

Li, Zhimin, Leslie Sheng Shen, and Calvin Zhang. 2024. "Local Effects of Global Capital Flows: A China Shock in the US Housing Market." *Review of Financial Studies* 37 (3): 761–801.

Liao, Wen-Chi, Daxuan Zhao, Li Ping Lim, and Grace Khei Mie Wong. 2015. "Foreign Liquidity to Real Estate Market: Ripple Effect and Housing Price Dynamics." *Urban Studies* 52 (1): 138–58.

Metcalf, Gabriel. 2018. "Sand Castles before the Tide? Affordable Housing in Expensive Cities." *Journal of Economic Perspectives* 32 (1): 59–80.

Mian, Atif R., Ludwig Straub, and Amir Sufi. 2020. "The Saving Glut of the Rich." NBER Working Paper 26941.

Mian, Atif, and Amir Sufi. 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the US Mortgage Default Crisis." *Quarterly Journal of Economics* 124 (4): 1449–96.

Moen, Espen R., Plamen T. Nenov, and Florian Sniekers. 2021. "Buying First or Selling First in Housing Markets." *Journal of the European Economic Association* 19 (1): 38–81.

Molloy, Raven, Charles G. Nathanson, and Andrew Paciorek. 2022. "Housing Supply and Affordability: Evidence from Rents, Housing Consumption and Household Location." *Journal of Urban Economics* 129: 103427.

National Association of Realtors. 2024. *International Transactions in US Residential Real Estate.* National Association of Realtors.

OECD. 1960-2024. *Analytical house price indicators.* Retrieved through OECD Data Explorer, Search keyword "Analytical house price indicators". (accessed September 11, 2024).

Pavlov, Andrey, Tsur Somerville, and Jake Wetzel. 2024. "Foreign Buyer Taxes and Housing Affordability." *Real Estate Economics* 52 (3): 928–50.

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel. 2020. "Segmented Housing Search." *American Economic Review* 110 (3): 720–59.

Powell-Smith, Anna. 2021. *New Data on Property in England and Wales Owned by Overseas Individuals.* Centre for Public Data.

Paris Property Group. 2023. "Paris Real Estate Attracts More Diverse Buyers in Recent Years." February 10. https://parispropertygroup.com/blog/paris-real-estate-attracts-more-diverse-buyers-in-recent-years/.

Sá, Filipa. 2025. "The Effect of Foreign Investors on Local Housing Markets: Evidence from the UK." *Journal of Economic Geography* 25 (3): 329–49.

Sá, Filipa, and Tomasz Wieladek. 2015. "Capital Inflows and the US Housing Boom." *Journal of Money, Credit and Banking* 47 (S1): 221–56.

SCF. 2001-2022. *Survey of Consumer Finances.* Board of Governors of the Federal Reserve System. https://www.federalreserve.gov/econres/scf-previous-surveys.htm (accessed May 26, 2025).

Sodini, Paolo, Stijn Van Nieuwerburgh, Roine Vestman, and Ulf von Lilienfeld-Toal. 2023. "Identifying the Benefits from Homeownership: A Swedish Experiment." *American Economic Review* 113 (12): 3173–212.

Stein, Jeremy C. 1995. "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects." *Quarterly Journal of Economics* 110 (2): 379–406.

US Bureau of Economic Analysis. 2020-2022. *Personal Income by County and Metropolitan Area, Table 1.* US Department of Commerce. https://www.bea.gov/sites/default/files/2023-11/lapi1123.xlsx (accessed August 30, 2024).

US Census Bureau. 2005–2024. *Business Formation Statistics.* https://www.census.gov/econ/bfs/xlsx/bfs_county_apps_annual.xlsx (accessed October 4, 2024).

US Census Bureau. 2022. *Foreign-Born CPS Data Tables.* https://www.census.gov/topics/population/foreign-born/data.html (accessed September 26, 2024).

US Office for National Statistics. 1997–2021. *Regional Gross Disposable Household Income (GDHI) at Current Basic Prices.* Retrieved through Nomis. https://www.nomisweb.co.uk/home/detailedstats.asp (acceesed September 14, 2024).

van Doornik, Bernardus, Dimas Fazio, Tarun Ramadorai, and Janis Skrastins. 2024. "Housing and Fertility." Banco Central do Brasil Working Paper 612.

Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill. 2010. "Why Has House Price Dispersion Gone Up?" *Review of Economic Studies* 77 (4): 1567–606.

Zillow. 2000–2025. *Zillow Housing Data. ZHVI All Homes, Time Series, Smoothed, Seasonally Adjusted, County.* https://www.zillow.com/research/data (accessed June 18, 2025).

Zillow. 2018–2025. *Zillow Housing Data. ZHVI All Homes, Time Series, Smoothed, Seasonally Adjusted, Metro & U.S.* https://www.zillow.com/research/data (accessed September 25, 2024).

Interview with Anne O. Krueger

Dylan Matthews

Anne O. Krueger is a Senior Fellow at the School of Advanced International Studies, Johns Hopkins University, Baltimore, Maryland, and the Herald L. and Caroline Ritch Emeritus Professor of Sciences and Humanities, Stanford University, Stanford, California. From 2001 to 2006, she was First Deputy Managing Director of the International Monetary Fund. From 1982 to 1986, she was Vice President, Economics and Research, at the World Bank, Washington, DC. She served as President of the American Economic Association in 1996 and was named a Distinguished Fellow of the AEA in 1997. Dylan Matthews interviewed her at Hunter's Bar and Grill in Potomac, Maryland, on September 2, 2024. The interview has been edited for clarity and length.

Getting into Economics

MATTHEWS: You started learning economics at Oberlin. How did you pick that major? What made you want to get a doctorate?

KRUEGER: Oberlin at that point had one of the better economics departments in the country, graduate or undergraduate. They had Ken Roose there, who had been at UCLA and was a Yale PhD before getting caught up in the loyalty oath

■ *Dylan Matthews is a Senior Correspondent at Vox.com, Washington, DC. His email address is dylan.matthews@gmail.com.*

stuff.[1] They had Ben Lewis, who was a superb lawyer as well as an economist. They were quite a contrast. Ken was very analytical in his approach to things, and Ben was much more on the applied side. He taught a public utilities course, and being a lawyer, he went through actual cases.

It was a small department, so you got to know your professors really well. It was the most interesting thing to major in because it had interesting people. I wanted to go to law school, but when I decided to go to graduate school instead, economics was the obvious field to apply in and the only one I think I was qualified in.

MATTHEWS: Today Oberlin has a reputation as artsy and nonconformist. Was that the feel when you attended?

KRUEGER: It always had a little bit of that. It was an Underground Railroad station during the Civil War. It had a commitment to academic excellence first. At the time I went there, and I think for quite a while afterward, Oberlin produced the highest number of doctors of physics and chemistry—not the highest per capita, the highest *number*. For a little college, it was really quite amazing. The academic excellence was there, along with the commitment to making the world a better place and all that kind of stuff.

MATTHEWS: Were there any academics in your family?

KRUEGER: My father was a professor of medicine at a medical school, and my mother was a professor of education.

MATTHEWS: At Oberlin and Wisconsin, had you specialized yet? Were there specific questions that you were focusing on?

KRUEGER: I was a generalist, but obviously with an international focus. So a generalist rather than a US specialist, might be another way of saying it.

MATTHEWS: Did that manifest as thinking about trade, as thinking about capital flows . . .

KRUEGER: Trade. There was no real big distinction at that time between trade and capital flows. There were only four countries that had Article VIII convertibility at that time.[2] It just wasn't a big issue at all. I taught for years the graduate course at the University of Minnesota in international economics. The first two-thirds were trade and the last third was finance.

Where I was asked to do things—and where things were interesting, and so I learned more and more about them—was focused on trade and exchange rate arrangements. The basic focus was always on the trade part. Countries found they couldn't export anything, but that's because they had the wrong exchange rate, because they were subsidizing everything they produced domestically at high cost, etc. So it was sort of obvious to get into the import substitution part of the story, the development part.

---

[1] Roose was forced out of the University of California at Los Angeles (UCLA) in 1950 for refusing to sign an anti-Communist loyalty pledge, after which he moved to Oberlin College (as told in Blauner 2009, Chapter 10).

[2] Under Article VIII of the IMF Articles of Agreement (available at https://www.imf.org/external/pubs/ft/aa/), countries pledge not to restrict current exchange transactions.

MATTHEWS: What was the debate about import substitution like at this point, in the late 1950s/early 1960s?

KRUEGER: The whole profession believed in import substitution. Almost without question. Even Gottfried Haberler, in his lectures in 1959, said that, of course, infant industry substitution by the developing countries was acceptable. Go back and look at the Cairo lectures.[3] It's in there.

MATTHEWS: Would you say that was how you were thinking about import substitution at the time?

KRUEGER: It didn't quite ring true. More than that, just seeing how import substitution was working made me skeptical. Lawyers who do trade law are more pro free trade than economists, because they know how badly protection works. A distorted economy is terrible. Not just a little bad—import substitution probably cut growth rates in half of what they could have been.

Learning from Turkey and India

MATTHEWS: When did you start traveling and getting experience in countries with these kinds of dysfunctional export and currency rules?

KRUEGER: The first trip was Turkey in 1965. That was also the first year I worked for USAID; I did some papers for them on trade regimes.[4] After I'd done the papers, they said, there are a couple of countries that we want you to go to apply these ideas. One of the countries they gave me was Turkey.

MATTHEWS: Tell me about that trip to Turkey. What were you looking for and what did you find?

KRUEGER: They sent me to the USAID building in Ankara; they sent a cable ahead and told me to show up there. I look around and there's almost nobody there. But there is a general—I think he was a general, it could have been a colonel. I went up to him and he said, "Oh yeah, we knew you were coming. We tried to tell [USAID headquarters] but they didn't listen. There's nobody here but me and I'm going on vacation tomorrow."

He said, "We tried to get a couple of appointments for you with some government officials in Ankara, but they're not interested in seeing you." Which, why should they be? I wouldn't have been either. I said, "Okay, well, I've just got to go and find some people myself. Will you please call the consulate in Izmir?" Don't ask me why I said Izmir, but I did. "Just ask the consulate and he'll make some appointments." The consulate was willing to do so. They found four or five Izmir

---

[3] In the lectures, Haberler (1959 [1988]), then perhaps the world's leading trade economist, stated, "It cannot be denied, I believe, that sometimes well chosen methods of moderate protection of particular industries can help to speed up economic development." In a new introduction to the lectures when republished in 1988, Haberler (p. 11) regretted this position, writing, "now believe that I went too far, especially in my Cairo lectures, in trying to find justification for a certain amount of protectionism in LDCs [less developed countries]."

[4] See, for example, Krueger (1966).

businessmen who were very happy to talk about costs and benefits, what it took to import, what the paperwork was, what the profitability was, and so on. I began getting more ideas as I went.

In those interviews, the businessmen were asking me where I wanted to go next. I said I wanted to go to Istanbul, and the businessmen had contacts there and made introductions. So I spent the next two weeks being wined and dined in Istanbul at the fanciest restaurants, of which there were not very many then. Businessmen who had not been included were coming up to the table and saying they wanted to see me too. I spent about two and a half weeks, by which time I had more than enough data to get a really good read on how inefficient the system was.

As a footnote, two years later, one of my friends and I were a party, and he said, "Oh, I want you to meet my father. He owns such and such company. And that's a company you bankrupted."[5] I was however forgiven because he was doing well in exporting some products that he could not have done in the old regime.

After that, Turkey just became part of my life.

MATTHEWS: So these Turkish businessmen were actively seeking you out and wanted to talk to you. Why do you think that was?

KRUEGER: In Izmir, there weren't so many Americans around. They were happy to see a foreigner because it was interesting and different. But they also did think that I might inform the government in a less biased way, and they had a lot of legitimate complaints. After that, when I went back to Turkey, the people were banging on my door to see me after I'd been there a week. There were people in the government who wanted to know what was going on, because I knew much more than they did.

MATTHEWS: What do you think was behind that? Is there just a lack of a domestic business press? A lack of strong local economics departments?

KRUEGER: There weren't that many people who were educated in the kinds of things that are taught in most Western universities. A business degree there is not the same as an economics degree. It wasn't even a business economics degree. A lot of businesses were in the family and owned and run that way.

Also remember, the intellectual tradition after (and because of) the 1930s was very anti-market. You had a lot of people who thought that if you let the government spend, all will be well, which, as we discovered the hard way, doesn't quite work. Pre-World War I, India had one of the most efficient steel companies in the world, more than Japan.[6] But the Nehru government decided to better tax it and take the revenue to support the public sector. Nobody saw anything wrong with this.

MATTHEWS: When was your first trip to India? You'd written a bit about it before visiting, but how did you get there? Did that trip change your view on how the economy was working there?

---

[5] More specifically, the policies Krueger recommended to reduce restrictions on imports had imposed costs on domestic firms benefitting from limited international competition in the earlier regime, but also opened possibilities for firms to import from abroad.

[6] The dominant firm to which Krueger is referring is Tata Iron and Steel. For a general history of the company and the conglomerate of which it is part, see Raianu (2021).

KRUEGER: Not particularly. I had been working in Turkey and I just went east rather than west when I left in 1969 or 1970. It was USAID again sponsoring. I was trying to do the same thing I did in Turkey, tallying up the costs and benefits of what they're doing. I figured the way to do it was to get some kind of company that's buying lots of parts and things, and then go to those companies and find out who their suppliers were, go to the suppliers, get cross-references, and so on. The interesting industry to do was automobiles, because India wanted its own auto industry, and there were all kinds of parts they hadn't been able to produce. When there were parts they *tried* to produce, some of them worked and some of them didn't.

It was relatively easy to get the prices that were charged by foreign companies in India, the foreign investors in India, and by Indian companies. Tata Steelworks was kind enough to sponsor all this. I went around from guesthouse to guesthouse, and I met lots of interesting people and got quite a bit of data. In one case, two Englishmen who were working there running a company gave me their books quite happily, because they were going to prove to me that the cost of their raw materials was higher than the cost of the finished product in England. But the price they sold it for in India *was* more than twice the cost of the raw materials. They never even noticed that.

The incentives to provide the right information were negative. Most companies had three sets of books. And most companies would show me all three.

MATTHEWS: One for the government, one for themselves. What's the third one?

KRUEGER: For the public. They did keep a set of public books. The true set was for themselves.

Rent-seeking

MATTHEWS: Now we arrive at your 1974 rent-seeking paper.[7] How did you turn these experiences into theory?

KRUEGER: In both countries, I had a number of encounters with people who talked a great deal about black markets and that kind of thing, and I got interested that way, in part. I had a friend in India from graduate school who was working in Delhi at one time when I was there, and he knew I was interested in this, so he invited several of his friends to dinner. He was well connected and invited ministers and vice ministers. At the end of the dinner he sat back and said, "Because Anne is interested in this, and you can trust her, I want each of you to tell exactly how much more you've taken than you should legally have taken in your job." And they did it!

---

[7] Krueger's 1974 paper was titled, "The Political Economy of the Rent-Seeking Society." In a different paper published to commemorate the first 100 years of publication of the *American Economic Review*, the authorial group of Kenneth J. Arrow, B. Douglas Bernheim, Martin S. Feldstein, Daniel L. McFadden, James M. Poterba, and Robert M. Solow (2011) listed Krueger (1974) as one of the top 20 articles the AER had published in its first century.

One said, "I only took enough to get my three daughters each an adequate dowry. And that was ₹12 million" [about $1.6 million in US dollars at that time]. One of my very able Indian friends, then a graduate student in India, was forced to leave India because he was so honest. He couldn't take the money. One of his sons got sick, so he had to work in the States to take care of his son. The corruption had all kinds of human effects that you don't think about.

MATTHEWS: Rereading the paper, what really struck me is how much emphasis you put on the magnitudes: that rent-seeking has quantitatively large effects on the economy.

KRUEGER: What I was taught in graduate school, and received wisdom at the time, was that it's only a transfer, so it doesn't cost. It was harmless. My argument was that it takes real resources. It takes good, smart people and wastes their time and their efforts.

MATTHEWS: I'm curious about the term "rent." Was that in common use as a term for monopoly profits at that time? You get a lot of deserved credit for popularizing "rent-seeking" as a term.

KRUEGER: I don't know why I chose it. It was just simply so obvious to me that that's what was going on. It was one of those things you don't think about.

Rent in economics, as you know, is a part inside the margin. If I'm a facilitator in India, and that was one of the names they were called, what I do is I go and I take some money from you, and then I get the railroad pass or whatever it is you want, and I pay the fee, and I pay a bribe. If I couldn't have gotten a job doing that same kind of thing in the open sector—that's rent. It's a rent that's come about because an artificial kind of scarcity has been created. But I was also trying to make clear that it wasn't only legal or illegal. That was not the distinction.

MATTHEWS: What was an example of legal rent-seeking that you wanted to highlight?

KRUEGER: 5,000 people applying for an NSF grant when there will only be 200 grantees. I like that example.

MATTHEWS: It puts economists in an area they understand. That raises an interesting question: is there a way for the National Science Foundation to spend its budget as allocated by Congress, which is scarce due to factors outside their control, that does not invite rent-seeking? Or is this just inherent in any kind of system?

KRUEGER: Some is probably inherent in any kind of system. That said, at the present time under the IRA [Inflation Reduction Act of 2022], somebody in a good institution can only get a grant if they work with somebody from a lower-ranked institution. I can't even begin to imagine the fun and games that are now going on. In Singapore and in Malaysia, for example, there's a law that says that you can't be at a law firm unless you've got a Malay partner. There aren't a lot of Malay lawyers but there are lots of Malay *names* on those firms. Those guys get a lot of money for their names.

There are things you could do to limit it to some extent, perhaps through some kind of a proof of ability test. Stanford got, I think, about 1,000 applications a year for graduate work in economics when I was there. There were only 40 places. The

secretaries were assigned to eliminate the first 700. We tried to make it clearer what our standards were. And yet people kept on applying. Now colleges are inundated with 50, 60, 70 times the number of applications.

MATTHEWS: How did you see it as relating to work by people like, say, Mancur Olson (1965), who were trying to think about regulatory capture? The paper reads more like it came from personal experience than that sort of theory.

KRUEGER: I knew of his work and I think he knew of mine. But he was basically talking about the institutions, and I was talking about the policies. In that sense, I did not basically look at it very much. I knew about it. Stigler's (1971) work was a little bit closer, and very much about regulators being captured by the regulated.

Joining the World Bank

MATTHEWS: So a couple of decades pass at the University of Minnesota department of economics, and then you're at the World Bank. What was your conception of the bank's role then? Has it changed as the world has evolved since then?

KRUEGER: Well, I don't know. The actual name is the International Bank for Reconstruction and Development, and that was its role. By the 1980s there'd already been a lot of reconstruction in the countries that had been destroyed in World War II. The reconstruction part was diminishing in importance, and the development part of the bank had been, in my view, when I got there, sold too much on the import substitution stuff.

It was fairly easy for me to see that and say, "This is not where we're going, at least as long as I'm here." My conception was and still is, that for low-income countries, there's got to be a lot more attention to the quality of what is being done, and much less than how many dollars are spent.

MATTHEWS: Has the World Bank's lending role changed at all, with the emergence of much larger capital markets for a lot of these countries?

KRUEGER: Until the 1970s, almost all capital flows (ignoring those across the Atlantic) were official. It was not until the 1990s that half of the capital flows in developing countries were from the private sector. Even then, much of the capital for the private sector was in minerals. Even when there was private investment in manufacturing, often it was often in the import substitution sectors. It was there because the companies in developed countries couldn't keep their market share if they didn't do that, and they received protection to cover their higher costs.

The funny thing is that governments didn't like foreign direct investment at that time. They got it all backward; they thought it was a bad thing. Even when I was at the Bank in the early 1980s and there were some private capital flows to emerging markets, they were not that important. Obviously Korea had some; Taiwan did it without too much. Singapore had some, Hong Kong not much. Chile turned it around without much. It did not look as if private capital flows

were all that central at that time. It wasn't until the 1990s when people suddenly realized that private capital flows were important and finally outweighed official flows.

MATTHEWS: Since we're talking for the *Journal of Economic Perspectives*, I have to ask if you saw the Michael Clemens and Michael Kremer (2016) piece in *JEP* arguing that the World Bank's role should be less about lending and more about providing public goods like research and policy coordination that contribute to poverty reduction.

KRUEGER: I not only saw that, that argument actually came up in the '80s, too. I have large misgivings about it. One of the things that both the Bank and the [International Monetary] Fund, and even the parts of the UN where there's any sense, did that was important, was to get insights from comparative experience. Joe Blow goes out and he goes to Nigeria, and John Smith goes out and he goes to Malaysia, and the knowledge that comes from that comparative experience is invaluable. Lo and behold, they discover that these aren't quite so different as they thought.

You've probably seen the stuff that Jagdish Bhagwati and I did (Bhagwati and Krueger 1973). We got ten countries to look at. And by golly, they couldn't be more different in other ways. But the problems that they ran into with exchange controls were so similar that this was something more fundamental than that. That's when I became a great believer in comparisons of cross-country experiences. If you just have a research section where everybody talks to everybody else there and they go out to lunch, and then they talk about, you know, the latest computer program, I don't think you're going to get this comparative perspective.

MATTHEWS: It seems like your USAID trips changed your research interests, or inspired a lot of research on your part. Did the World Bank experience do that for you as well?

KRUEGER: Not to the same degree, certainly. I think quite the contrary. It changed them more than it changed me.

MATTHEWS: Which is what you were there for, I imagine. How do you think you changed them?

KRUEGER: The World Bank has something called the *World Development Report*, and the person who's in charge of research, which was me, oversaw it. The development report comes out around June every year. I came in in September. They'd already decided on the topic for the following year, and they'd already decided on the outline of the chapters, and who would write the chapters. There was one chapter on how to administer price controls. And the title of the whole thing was going to be "Managing Development."

Well, I was there as an economist, and I wasn't going to lose all my credentials in one fell swoop. So we got rid of the price controls and all that. We called it "Management *in* Development." You'd be surprised how many people in the Bank changed their tune then. But then, to my astonishment when I left, some turned back to the former approach.

Managing Crises at the IMF

MATTHEWS: In 2001, you go to the International Monetary Fund. That's right before the crisis in Argentina breaks out.

KRUEGER: I had actually done some work at the request of an Argentine businessmen's group and went there, maybe around 1980. I got a fair idea of what was going on in Argentina because of that, and stayed working with that businessmen's group for a while. So it turned out that the two countries that were in most trouble in 2001 were two I knew more than the average about—Argentina and especially Turkey.

Not only that, but a very good and very able economist, who had been at the World Bank, had been called back to be Turkey's economy minister as the crisis hit the big time. He was one of my closest friends, and we were able to work very well together. Turkey was actually relatively easy to fix. To my regret, Recep Erdogan, who was just beginning as Prime Minister in 2003 [and was later elected President in 2014], took the popularity he got as a result and abused it in all kinds of ways, going way off track in terms of civil liberties. But per capita income in Turkey tripled from 2001 to 2009.

Argentina was a completely different story. Brazil was also in crisis. We pulled them out nicely. It wasn't terribly hard.

MATTHEWS: That's interesting—so in 2003, Lula da Silva takes over in Brazil.

KRUEGER: This was before that. But he certainly was in power when I first visited there. I had worked in Brazil in the late 1960s also. So the countries that turned out to be in trouble after I got to the IMF were countries I knew pretty well.

MATTHEWS: I think of Lula as a very populist figure, but Brazil's outcomes under his leadership don't seem as disastrous as the Kirchner's in Argentina.[8]

KRUEGER: Argentina's much more disastrous than anyone. Per capita income now has gone down about 10 percent over the past ten years. Inflation is still in the 50 percent range.[9] They have a good guy as economy minister and he's trying to do the right things. He is getting quite a bit accomplished, but there remain serious issues that have not yet been sufficiently addressed.

MATTHEWS: Tell me a bit more about the Argentine crisis in 2001–2002. You had written a lot about chronic problems of developing countries: import rules, licensing regimes, things that over the course of decades slow growth. It seems like you were thrown into a situation where you had to tackle very acute problems. Did that feel notably different?

---

[8] Lula da Silva was President of Brazil from 2003 to 2011, and then was reelected to the presidency in 2023. Néstor Kirchner was President of Argentina from 2003 to 2007. He was succeeded by his wife Cristina Fernández de Kirchner, who served as President from 2007 to 2015.

[9] For Argentinean per capita GDP in recent years, see the World Development Indicators website at https://data.worldbank.org/indicator/NY.GDP.PCAP.PP.KD?end=2023&locations=AR&start=1990&view=chart. As of August 2024, monthly inflation was 3.9 percent (as reported by Nessi 2024) implying over 58 percent annual inflation.

KRUEGER: When you have an acute problem, it was normally because they misbehaved. Think of an alcoholic: he's got an acute problem, because he's very sick. But on the other hand, it didn't come up overnight. The big problem is judging, judging how much they need to do to get back to a point where the alcoholic has stopped drinking. If he promises, "I'll go from ten quarts a day to nine," forget it. The other problem is, if he promises to cut from ten to two, is that enough? Will he stick with it?

One of my good economist friends who knows Argentina better than I do says in any country in Latin America, other than Argentina, if your neighbor's farm started burning down, everybody there would be there to help you and pour water on the flames. In Argentina, they would be there to steal the cows. There may be some truth to it. Some of the brightest people I've ever met have been Argentines, but they do not seem to be able to work together politically.

It was actually Peron who came in [he was elected President in 1946] and drove decline. This is almost what scares me most about Trump. Peron took a country that was reasonably sensible—Argentina had very little inflation when Brazil was a highly inflationary country, it was stable, rich, it was the richest country in Latin America. Peron came in, and it was all to help the workers, but he dismantled it. There is a wonderful book on this which, if you're all interested, you should read, by Carlos Díaz Alejandro (1970).

MATTHEWS: You left the Fund before the global financial crisis in 2008–2009. Did you have a sense that that was brewing toward the end of your time there?

KRUEGER: Well, at the Jackson Hole meetings in 2005, Raghuram Rajan (2005) gave a talk where he described what the vulnerabilities were, why helping the poor with mortgages is not going to work. Why are we sitting on a bombshell? I think it was Larry Summers who said we were Luddites.

MATTHEWS: Were you in the audience for that?

KRUEGER: Yes, I was there, and I was Raghuram's boss. We had that one pretty much nailed. Even after the crisis, I said to the person who had been Alan Greenspan's chief assistant during that period, I asked, "Are you sure you don't think that lending to low income groups was a big part of it?" He said, "Oh, no, no, that wasn't part of it." But I don't see how. If one-third of your lending is dubious and doesn't work out, you're going to have trouble. I was not surprised. I was surprised that it took so long, and I was surprised at how people were reacting in a strange way.

MATTHEWS: One legacy of the crisis is that it led to a lot of backlash against economics as a discipline . . .

KRUEGER: Which, in my view, is totally unfair.

MATTHEWS: What were the right lessons to learn from that moment?

KRUEGER: Well, one right lesson is that you don't legislate the maximum interest rate nor allocate credit by fiat. Some lessons that were learned were good ones, like the tests to make sure the financial sector is fine. I haven't looked at them enough to know whether the Fed set the tests at the right levels. My guess is they're not. Raghuram's been writing about it, and he's normally very good on this stuff. So if I really wanted to bone up on it, I'd take a week and read his latest on it.

China and Industrial Policy

MATTHEWS: We haven't talked about China, and a big background question relating to everything we've talked about is, "How do these institutions adjust to the rise of China?" That seems especially important in the context of the IMF, with China taking a great deal of interest in being a lender to various developing economies. How do you think that changes the role of the IMF?

KRUEGER: Let me answer half a question you didn't ask: the Chinese have on occasion tried very hard, I think, to oblige [the United States]. And I think we have missed the signal completely. I think their system is not a good system, and I think they themselves are not quite capable of understanding why it is not a good system, but this idea that they're always stealing from us, that we never steal from them is silly. I mean, the idea that there's no industrial espionage in this country! When I was in Silicon Valley, nobody would send a piece of equipment to any of the electronics fairs, even the big one in Las Vegas, without making sure there was some employee 24 hours a day watching the machinery to make sure nobody could reverse-engineer it. Now we're so mad at the Chinese that they tried to do the same thing. Meanwhile, we've had more than one CIA agent arrested [in China] and pretty much caught dead to rights.

I think that the Chinese have been trying to make things a little better in the past—at least a year, maybe two years. I think they realized how isolated they were and how dependent they were.

Now, to your question. The part of the international system that is not filled in legally is the part about sovereign debt. The trouble with that is that, if a debtor is in trouble, there's going to be a race to see who gets the money out fastest, and that will take any chance they have of recovery and make it much, much smaller, if not nonexistent. We didn't have any good international rules on it. We had the Paris Club, which was an informal organization. It didn't even have a Secretariat, it was just the Bank of France doing the work.[10]

When a country did get in trouble, the Paris Club could meet, the IMF and World Bank could come and say, "Here's what we think their prospects are. Here's what you can do." And then [the lenders] could get an appropriate haircut, and the IMF could lend some money and get it going again. Very often they do need to lend some money to get things going again, because, say, there's no oil in the country and no car can operate or whatever.

Now along come the Chinese. They were invited to join the Paris Club and they said no. So now you've got countries that are in debt [to China] that's become unsustainable, and China has not been willing to participate in determining terms for them. The Chinese are being very bad on the set of issues surrounding debt sustainability, and the Chinese are now big lenders. I think somehow we [the United

---

[10] The website of the Paris Club is https://clubdeparis.org/en. It includes a brief history going back the Club's origins in 1956.

States] have been bad on trade issues, which we should not have been. They've been bad on capital issues, and we need to come together. And that's very hard.

The Japanese, by the way, believe in low interest rates at almost all costs too. So you've got a part of the world that doesn't even believe in higher interest rates for anything. Now they're important enough that somehow or other it has to be sorted out. So I do think there is a problem here.

The Chinese still think their system works as an economy and I don't think it's working. I think they're in big trouble. I think we're going to see some change in [their economic system] in the next couple of years. I think we're better off to wait for that than compromise too much on other things. But we should take off the tariffs because that's hurting us more than them.

MATTHEWS: On the trade side, do you put much credence in the idea that there is an "oversupply" problem with Chinese steel or solar panels?

KRUEGER: Well, on steel, the last thing I read is that the Chinese authorities are very angry because they issued orders to cut back, and they wouldn't allow any new steel plants, and everybody's ignoring them in China. So it's more a principal-agent problem. A lot of it happened for reasons other than deliberate choice. The steel thing is always a problem. The steel industry has been the most protectionist industry in the US since there was first steel. But any resolution needs to be multilateral, involving the EU, the UK, Japan, the US, Canada, and others. We should be using the World Trade Organization to address the issue.

I just finished a piece on solar panels (Krueger 2024). What intrigues me on that one is that the Biden administration delayed the imposition of tariffs for two years, because that's how long it would take to put up a solar panel plant. Well, the Chinese have 80 percent of the raw materials for the whole world. Second, they produce at about $0.08 a panel, and our costs right now are $0.31, almost four times as much. A solar panel installation is good for the environment. There are better jobs in installation.

So why on earth don't we just buy up two years worth of stock at $0.08 and keep two years worth of it? Then we're two years ahead. If at any time the Chinese start dumping, we'll just use up that two years' stock. There would be more installations because of lower costs and more and better jobs because lower costs would increase the number of people installing panels, and installation jobs per roof are more than panel jobs per roof, and the pay of installers is better. Why aren't we doing that now?

When you're hurting employment, you're hurting the environment, and you're hurting inflation, and those are the three objectives—why are you doing that? I have a lawyer friend who just read the solar panel article. He said, well, what happens if the Chinese corner the market? Well, how are they not being able to corner the market as long as they have 80 percent of the raw materials?

MATTHEWS: A broader question that I have here is how we should evaluate noneconomic or nonwelfarist concerns, like national security or concern over forced labor, in thinking through problems like this. Is there a way to do that rigorously and use the tools of economics to do that?

KRUEGER: There are two parts to that answer, one of which is that it's quite possible to form some kind of an expert group that would decide what's truly national security–sensitive stuff. We have the UN, and to my astonishment, even the Americans believe in the International Atomic Energy Agency. But you'd have to have an international panel. One of our mistakes, no matter what we do, is going it alone. Because we're not big enough and we're not important enough. The minute we're with the Europeans and the Japanese, our power over China increases. By the way, the Europeans just put in 100 percent tariff on EVs [electric vehicles]. Why didn't we just do it with them? If we're going to do it at all, do it together.

The second point is that we get together, get some kind of an international commission on what I call military sensitivity. It would be hard to do, but if you can do it on nuclear, you can do it on that. Then sensitive military technology could be deemed off limits. That's fair enough. I'm all for it.

I don't want the Chinese communists taking over either! But I don't want us to pretend that everything in the whole world is leaning that way, because I don't think it is. We keep forgetting that we're a very rich, powerful country. We keep forgetting that they are much weaker than we are.

MATTHEWS: I went back and read a piece you wrote on regional trade alliances in 1999 for the *JEP* (Krueger 1999). It's funny to read now because the vision you advocate of making steady progress through the World Trade Organization just feels worlds away.

KRUEGER: Well, that's where we should be. Regional is still third-best.

An argument that's come up is that the first arrangements were just trade and now they're about so, so much more than trade. There's some truth to that. But let's get more done on those fronts. The European Union is the only pact with mutual recognition of regulations. We could do that with a lot of countries and that could be done bilaterally. It doesn't have to be through any organization at all. We could say, "We trust the Dutch not to cheat us on this, and we trust the French not to cheat us on wines," stuff like that. The only thing that scares me about that is the regulatory agencies might be too much of a headache, but I still think that's better than doing it by regional trading arrangements.

MATTHEWS: Why do you think the GATT/WTO era of liberalization ended? Is there a way to reverse the processes that led to that?

KRUEGER: I think a reversal will happen. Maybe it happens because some countries are more sensible. Remember, already we have dismantled the dispute settlement mechanism, which was arguably the most successful thing that the WTO had done. Twenty-seven countries, of which one is the European Union, are taking the WTO mechanism and they're using among themselves legally and settling disputes with it.[11]

Meanwhile, we [the United States] stayed out of the Trans-Pacific Partnership. And what happened? We got complaints that it was unfair because the Australian

---

[11] For the website of the Multi-Party Interim Appeal Arbitration Arrangement, see https://wtoplurilaterals.info/plural_initiative/the-mpia/.

farmers got their wheat into Japan more cheaply than we did—that's what the farmers are saying. And yet they don't say that we should have been in the TPP.[12]

China is there to stay. The better thing to do is to get them to act within the WTO. Take the environment. The Chinese emit too much. That's it. Why not use that as a model of how it [the WTO] can get them to do the right thing?

MATTHEWS: What do you think works in terms of incentivizing countries to change behavior on things like that?

KRUEGER: It depends on what the issue is and what we're doing. I don't know the answer to that in any general sense, but for example, the Russians and the Chinese worked very hard to get "most favored nation" status. That did work [to change behavior]. There's no doubt about it.

MATTHEWS: It must be surreal, having seen President Nixon go to China in 1972 and the US successfully cleave China and Russia apart from each other, to see it all reversed 50 years later.

KRUEGER: That was unbelievable. I was with that first National Academy of Science economics group that went to China after the opening in the early 1980s. And on the streets, little kids would be coming and trying to practice their English. Everyone would want to talk to us. It was just unbelievable to watch the enthusiasm among the people for that kind of an opening. And we just threw it all away.

MATTHEWS: Another *JEP* paper of yours that I went back to was on government failure as a concept in development (Krueger 1990). One way I understand what the Biden administration is doing is that they're almost copying an idea from China: that market failure is more common than we thought it was, that government failure is less of an issue than we thought it was.

KRUEGER: I don't think they thought that second part through. Some of these arguments about the market assume that if there are market failures, then whatever the government will do will be better. Maybe the market failures are huge, but that does not persuade me that government failures will not automatically be *as* huge. That's the part that's wrong. I still think that when you're talking about lots of economic activities, you want to just look at incentives. If there's something wrong with the market, get the incentives right. Giving bureaucrats the incentive to regulate is not the incentive that will work best in most cases.

MATTHEWS: One counter I sometimes hear from people is, if that's true, then why have these Asian mixed economies that involve a lot of state-owned enterprises and government intervention done so well? Why is China growing so fast? Why did the East Asian "tigers" [Hong Kong, Singapore, South Korea, and Taiwan] grow so fast? Aren't they evidence that industrial policy can work?

KRUEGER: Well, the answer to that is obviously "no." One of the really sad cases is Korea. What they did that was so important was to make sure that nobody was going to get any kind of incentives if they didn't make exports. In the export

---

[12] The United States withdrew from negotiations to enter this trade agreement in 2017. For an overview of the Comprehensive and Progressive Agreement for TransPacific Partnership (CPTPP), as the eventual agreement was named, Congressional Research Service (2023) offers a starting point.

you had a standard of proof, a way to tell, are you performing well or not? There was competition, and competition was not necessarily domestic. It was international. If you exported a dollar's worth of goods, you got $0.10 worth of tax credits. You got interest rates at 1 percent, less than market, and stuff like this. Yes, there was government intervention, but at a level playing field for all goods within that framework. Infrastructure was made to serve all producers (especially exporters) and the government for the most part did not pick winners, and expected export performance.

Within that framework, government bureaucrats were scared stiff. They had to behave even more than the private bureaucrats and the private bureaucrats. [South Korean President from 1962 to 1979] Park Chung Hee had his car driven around town and he would see some kind of fairly affluent house being built. He would stop the car, get out, and find out about the house. He'd go and look to see if the property was legitimate, and discourage opulent spending.

This was not a government driven economy except in the sense the government said, "we need more exports." Because it was terribly unbalanced, where you had 3 percent exports as a share of GDP in 1960 and 13 percent imports.

When you're that poor and you have 80 percent in farming, which is about what they had, then it's easier to figure out what to do than it is when you've got a more complex economy. As the economy grew more complex, it grew more hands off. I really find this argument that the government did it all is just wrong. It was lots and lots of people. It was the little kiosks where parents went and worked for the day, and the eight-year-old son sold the morning paper. It was everybody [that made Korea rich], and there were incentives for everybody.

The same was true of Taiwan, Hong Kong, Singapore. In Hong Kong in particular, you simply cannot make the argument that it was all [government]. They didn't even keep balance of payments records. The first time I went to Singapore, I'm meeting with the Economic Development Board—I didn't know I was going until the last minute. I get there and these five guys were sitting across from me. They said, "Well, what do you want to know?" And I hadn't even thought about it. I said, "Well, by what criteria do you decide what to do?" And they said, "We look at it and see what makes sense." It was just a completely different mentality [from India].

India tried to do industrial policy with many industries including machine tools, tractors, chemicals, autos, and trucks in the 1960s and 1970s. I was there at the time, and a couple of my students were very much in the thick of the argument, and it did not work. It just didn't. Even in South Korea, they decided they had to have foreign investment.

Later, I was on a Korean commission; it had five people and was called the "wise men's commission." I was one of the five wise men. We were to look at the Heavy and Chemical Industry Drive, which Park had said they had to do for defense reasons when the US cut off their petrol supply.[13] We went to a ball bearing plant, which was

[13] For an overview of Korea's Heavy and Chemical Industry drive, a starting point is Horikane (2005). The Energy Policy and Conservation Act (EPCA) of 1975 was passed in the wake of the OPEC embargo

in a beautiful shiny building. It had two workers. It was the size of maybe ten basket-ball courts or more, and had all this shiny machinery that obviously never needed oiling because it was never used. They put in a commission of foreigners to come in and see what had gone wrong and what they should do about it.

Dani Rodrik says that that was a success because now they do some of these things.[14] But if someone proposed [an investment] where for 50 years you pay in, and you get no return for 40 years, I'll bet you wouldn't want to do it. Korea's Heavy and Chemical Industry Drive was a flop. It sucked out all the engineering talent from industries that were exporting. Engineering salaries tripled. It's the only time exports—traditional goods or exports that were booming—fell. There were [perverse] incentives in it the whole way.

The Future of Economics

MATTHEWS: My last two topics are more about the economics profession—I'm sorry for the abrupt transition. You were the chair of a panel on graduate education in economics that concluded in 1991 . . .

KRUEGER: It was a commission. It was not a panel. It was appointed by the AEA.

MATTHEWS: The quotation I have from that report is, "graduate programs may be turning out a generation with too many idiots savants, skilled in technique but innocent of real economic issues. At the margin, increased attention to appli-cations of technique to economic problems would improve their productivity" (Krueger 1991). It seems like if anything, in the 30 years since . . .

KRUEGER: It's gotten worse.

MATTHEWS: Right. Are there ways to course correct?

KRUEGER: Well, I don't know. Who knows? Remember, I've been out of academia now, really, since 2006. In any event, it used to be that we viewed depart-ments in economics like baseball teams: first basemen, second basemen, pitchers, catchers, etc. You wanted the theorist, you wanted the econometricians, you wanted the labor economist, trade, all the fields. A good department had all these things.

Gradually, I think the technique guys took over, theory and econometrics. It weakened the field in several ways, one of which, it set the bar so high that some people just couldn't do it or didn't want to do it. Some, who were good, got out, even though they could do it, because they didn't want [that focus]. In a sense, you've gotten the idiosyncrasies of that group. If the price of milk doubled, they wouldn't even begin to wonder why. That's a big part of it. The fact that we have

---

on oil exports in 1973–1974. The law had a number of provisions, including the enactment of fuel economy standards for cars, but the provision of interest here was a prohibition on nearly all US exports of crude oil—which affected South Korea. For an overview of the ban on US crude oil exports, which lasted from 1975 to 2015, a useful starting point is US Government Accountability Office (2020).
[14] See, for example, Juhász, Lane, and Rodrik (2024).

supercomputers sets everybody in awe, and all I can see is higher caliber regression, *maybe*. And I still think you need some theory before you do it.

I think at some point it has to reverse a bit. Part of the way it's reversing is, in my view, a bad way. Namely, the public policy schools are doing economics without enough analytical underpinning. So you get things like the advisers supporting industrial policy in the Biden administration and stuff like that, precisely because they haven't learned what happened.

One thing that the Chicago school, for instance, was so good at was making everybody think in terms of, "What's the alternative? What's the opportunity cost? What are the incentives that are created?" One of my eureka days was when Milton Friedman was visiting Minnesota sometime in the 1960s, around when the first law came in restricting auto emissions for new cars, making them more expensive. Milton's argument was, well, that would mean more pollution. Why? Because, obviously, then people keep their cars longer. To him, it was so obvious. But to me, until he said it once, it wasn't obvious.

There was so much of that in Chicago. Students learned to think that way. The rest of us had to figure it out the hard way. But Chicago often got it right.

MATTHEWS: I think a lot of the people reading this interview might be in grad school or early in their careers, and want to have an influence in policy. You managed to have a long career in both academia and in policy, and I think had a very positive impact in both places. What advice do you have for someone who wants to follow a similar path? What are the skills to develop, or the opportunities they should take?

KRUEGER: I really do think that depends on the individual and what they want, but I certainly am a very strong believer in having a good analytical framework. That to me is essential, and that means learning what questions to ask, which is often, "Why are people behaving as they are?" and then looking for where the incentives are. Finding out "why" is the important thing first. In the 1960s, the first answer in development economics was "peasants are irrational." But of course, Ted Schultz (1964) put paid to that very fast.[15]

So many of these things, when they say the market isn't working, it turned out there were some bad incentives. It wasn't true that there were just stupid people who didn't make rational choices. So this kind of thinking about why people are responding the way they are is important.

My advice to a young person would be to keep on taking on opportunities as they come and the right things will ultimately find you! I knew (as I said) that my strongest interests were international trade. After I began teaching, I accepted the assignments that came my way. I worked on a study of the impact of the St. Lawrence Seaway on the upper Midwest (Krueger 1963), did a number of commissioned papers on various issues, wrote some academic papers—including one on India

---

[15] Schultz (1964) famously argued that peasants in developing countries were not irrational, but instead were "efficient but poor." For an overview, see Schultz's (1979) Nobel lecture, which cites and discusses the earlier work.

(Krueger 1962) and several on trade theory (for example, Krueger 1969, 1972) and did some consulting assignments. It wasn't until the USAID Turkey assignment came along that I was able to zero in better.

And it was another five or six years before I was really focused. I'd done a reasonably important paper on development (Krueger 1968) and would later do a commissioned book on the balance of payments for Cambridge U press (Krueger 1983), before I was finally almost totally engrossed in the trade-development issues.

I always told my PhD students that when they received that degree, they were getting their union card. It gives a license to take on assignments until the right one comes along. Then they can focus on the area that most interests them.

References

Arrow, Kenneth J., B. Douglas Bernheim, Martin S. Feldstein, Daniel L. McFadden, James M. Poterba, and Robert M. Solow. 2011. "100 Years of the *American Economic Review*: The Top 20 Articles." *American Economic Review* 101 (1): 1–8.

Bhagwati, Jagdish N., and Anne O. Krueger. 1973. "Exchange Control, Liberalization, and Economic Development." *American Economic Review* 63 (2): 419–27.

Blauner, Bob. 2009. *Resisting McCarthyism: To Sign or Not to Sign California's Loyalty Oath.* Stanford University Press.

Clemens, Michael A., and Michael Kremer. 2016. "The New Role for the World Bank." *Journal of Economic Perspectives* 30 (1): 53–76.

Congressional Research Service. 2023. "CPTPP: Overview and Issues for Congress." June 16. https://www.congress.gov/crs-product/IF12078.

Díaz Alejandro, Carlos F. 1970. *Essays on the Economic History of the Argentine Republic.* Yale University Press.

Haberler, Gottfried. 1959 [1988]. *International Trade and Economic Development.* National Bank of Egypt.

Horikane, Yumi. 2005. "The Political Economy of Heavy Industrialisation: The Heavy and Chemical Industry (HCI) Push in South Korea in the 1970s," *Modern Asian Studies* 39 (2): 369–97.

Juhász, Réka, Nathan Lane, and Dani Rodrik. 2024. "The New Economics of Industrial Policy," *Annual Review of Economics* 16: 213–42.

Krueger, Anne. O. 1962. "Interrelationships between Industry and Agriculture in a Dual Economy." *Indian Economic Journal* 10 (1): 1–13.

Krueger, Anne O. 1963. *The Impact of the St. Lawrence Seaway on the Upper Midwest.* Upper Midwest Economic Study.

Krueger, Anne O. 1966. "Some Economic Costs of Exchange Control: The Turkish Case." *Journal of Political Economy* 74 (5): 466–80.

Krueger, Anne O. 1968. "Factor Endowments and Per Capita Income Differences among Countries." *Economic Journal* 78 (311): 641–59.

Krueger, Anne. O. 1969. "Balance-of-Payments Theory." *Journal of Economic Literature* 7 (1): 1–26.

Krueger, Anne. O. 1972. "Evaluating Restrictionist Trade Regimes: Theory and Measurement." *Journal of Political Economy* 80 (1): 48–62.

Krueger, Anne O. 1974. "The Political Economy of the Rent-Seeking Society." *American Economic Review* 64 (3): 291–303.

Krueger, Anne O. 1983. *Exchange-Rate Determination.* Cambridge University Press.

Krueger, Anne O. 1990. "Government Failures in Development." *Journal of Economic Perspectives* 4 (3): 9–23.

Krueger, Anne O. 1991. "Report of the Commission on Graduate Education in Economics." *Journal of Economic Literature* 29 (3): 1035–53.

Krueger, Anne, O. 1999. "Are Preferential Trading Arrangements Trade-Liberalizing or Protectionist?" *Journal of Economic Perspectives* 13 (4): 105–24.

Krueger, Anne O. 2024. "Biden's Solar Tariffs Undermine His Own Climate Agenda." *Project Syndicate*, August 28. https://www.project-syndicate.org/commentary/solar-energy-trade-war-jeopardizes-biden-agenda-by-anne-o-krueger-2024-08.

Nessi, Hernan. 2024. "Argentina Inflation Seen at 31-Month Low of 3.9% in August." *Reuters*, September 9. https://www.reuters.com/markets/argentina-inflation-seen-31-month-low-39-august-2024-09-09.

Olson, Mancur. 1965. *The Logic of Collective Action.* Harvard University Press.

Raianu, Mircea. 2021. *Tata: The Global Corporation That Built Indian Capitalism.* Harvard Univeristy Press.

Rajan, Raghuram G. 2005. "Has Financial Development Made the World Riskier?" In *The Greenspan Era: Lessons for the Future, Jackson Hole Economic Policy Symposium*, 313–69. Federal Reserve Bank of Kansas City.

Schultz, Theodore W. 1964. *Transforming Traditional Agriculture.* Yale University Press.

Schultz, Theodore W. 1979. "The Economics of Being Poor." Speech, lecture to the memory of Alfred Nobel, December 8. https://www.nobelprize.org/prizes/economic-sciences/1979/schultz/lecture/.

Stigler, George J. 1971. "The Theory of Economic Regulation." *Bell Journal of Economics and Management Science* 2 (1): 3–21.

US Government Accountability Office. 2020. *Crude Oil Markets: Effects of the Repeal of the Crude Oil Export Ban.* GAO-21–118.

Protecting Antiquities: A Role for Long-Term Leases?

Michael Kremer and Tom Wilkening

From 1801 to 1812, Lord Elgin, the British ambassador to the Ottoman Empire, had workers remove and ship to England 75 meters of the original 160-meter Parthenon Frieze, a marble sculpture created to adorn the upper part of the temple constructed in Athens about 2,500 years ago. This step structurally weakened the remainder of the Parthenon. The Elgin Marbles, as they are often called, have long been on display at the British Museum in London. Elgin claimed he received permission from Ottoman officials, but the Greek government has long sought their return, challenging the legitimacy of the permissions allegedly granted by Ottoman officials long ago. The governments of Greece and the United Kingdom are to this day involved in talks about whether, or how, or under what conditions the Elgin Marbles might be returned.

Like Greece, many countries view antiquities as a part of their cultural heritage and seek to retain them for the benefit of current and future generations. A standard approach to this objective is to ban the private export of antiquities and to nationalize antiquities from unexcavated sites. However, export bans can be difficult to implement, and the ability of governments to protect archeological sites is often limited. As a result, illegal and clandestine trade in antiquities continues. These practices can be particularly damaging to cultural heritage, because those engaging in this trade work quickly and surreptitiously. As depicted in Figure 1, looters use fast methods of excavation such as pneumatic drills and dynamite (Coggins 1972;

■ *Michael Kremer is Kenneth C. Griffin Department of Economics, University of Chicago, Chicago, Illinois. Tom Wilkening is Professor of Economics, University of Melbourne, Melbourne, Australia. Their email addresses are kremermr@uchicago.edu and tom. wilkening@unimelb.edu.au.*

Figure 1
Damage from Looting: An Example from Cambodia

Note: Before and after picture of a pediment damaged by looting in the Banteay Chhmar complex in Cambodia. Large parts of the Banteay Chhmar complex were systematically looted in 1998 and 1999. Stolen antiquities included two twelve-meter wall sections that were cut into blocks and intended for international sale.

Bator 1982; Prott and O'Keefe 1984). They work to keep site locations secret, and often obscure the origin of antiquities by intentionally damaging sites and breaking antiquities into fragments to pass international borders. Seeing antiquities in place and in context provides archeologists and historians with more insights than seeing pieces in isolation; thus, clandestine trade in antiquities jeopardizes scientific inquiry and limits understanding of past civilizations.

Many who study the antiquities market argue for either stricter enforcement of export bans or for these laws to be repealed in favor of free markets. However, neither of these polar policies are likely to be fully implementable in practice. Artifact-rich countries often have limited resources to monitor potential looting sites or to detect transfers of objects out of the country, as well as to maintain, preserve, and protect objects already in government control. For instance, Nigeria's total budget for its museums and monuments was just over $26 million in 2013, with security being only a subset of those costs; large museums in Nigeria have been the victims of major robberies seven times in the last three decades, with estimated losses of up to $200–250 million from a single case (Akinade 1999; Shyllon 2000).

In cases where the difficulties of enforcement can be overcome, even perfectly enforced export bans may weaken incentives to maintain antiquities that are not under direct government control. Despite their durable-goods nature, antiquities are fragile and may be damaged or permanently destroyed with improper care. If those in possession of artifacts do not have the means or incentive to maintain the

antiquities and are unable to sell them to those who do, the antiquities may be at risk of damage or destruction.

Noting the difficulties inherent in restricting trade in antiquities, some economists have argues that export bans should be abandoned and replaced with free trade (for example, see Posner 2007). However, the idea of a global market in selling and buying antiquities is strongly opposed by most source countries, by international treaties, and by the overwhelming majority of archaeologists (Renfrew and Elia 1993; Brodie, Doole, and Watson 2000). In this journal, Roth (2007) argues that repugnance is a constraint on the operation of some markets. Such repugnance constraints may make the widespread cross-border sale of certain antiquities infeasible.

In this essay, we will argue that leases in which antiquities can be transported outside the source country for a limited time offer an intermediate course between an impractical export ban and a repugnant policy of sales. Like sales contracts, leases generate revenue that the source country can use for a variety of purposes: compensating those who make discoveries, supporting excavations, providing security for at-risk sites, and funding the domestic display of antiquities to the public. Further, leases automatically link the value of rents to the value of the antiquity, which is important for encouraging individuals to reveal and maintain antiquities. However, unlike sales, leases provide an explicit path for antiquities to return to the country of origin and are less sensitive to corruption. Leases are also likely to be more politically feasible than sales, especially if the proceeds are used for protection of antiquities or are in-kind and involve protection and restoration.

The use of leases to facilitate the movement of antiquities is not new, and a number of precedents suggest that leases are feasible. These precedents also suggest that the legal and auxiliary institutions currently used for loans between museums can be adapted to ensure that leaseholders exercise proper care to antiquities.[1] A well-known example is the King Tut exhibit, more formally known as the "Treasures of Tutankhamun," which circulated in the United States and London from 2005 to 2008. The exhibit, which displayed artifacts from the tomb of a boy-king of Egypt more than 3,000 years ago, was leased to a private company which charged $5 million per city and generated proceeds that went to the renovation of the Egyptian Museum in Cairo.[2] The lease agreements for the King Tut exhibit specified transportation, display, and storage conditions, and required insurance for $650 million costing roughly $1 million per city (Boehn 2005). A second example is the long-term lease of two thirteenth-century Byzantine frescoes that were previously stolen from the Church of Cyprus and recovered by the Menil Foundation

---

[1] From a legal perspective, Beltrametti (2013) studies whether current Italian and Greek laws would allow leases of antiquities that involve monetary compensation. Italian law does not explicitly prohibit monetary compensation and there are some indications that practitioners in the legal system would be in favor of such leases. Greek law is less clear on this issue but does not explicitly rule out monetary transfers in reciprocal exchange agreements.

[2] The exhibit circulated again between 2008 and 2011 with additional stops in the United States and Australia.

from the illicit antiquity market. The Menil Foundation restored the frescoes as part of the lease requirements and displayed them from 1992 to 2012, when they were returned to Cyprus. Leases have also been successful in the exchange of artwork. In 2007, the Louvre agreed to lease 200–300 artworks to its counterpart in Abu Dhabi over a ten-year period for $247 million (Riding 2007).

We begin this essay with an overview of the market for illicit antiquities. Next, we highlight some of the challenges that exist in managing antiquities with export bans and some of the drawbacks that exist in using a free market. We discuss how the policy objectives that underlie trade restrictions can be achieved with a lease and discuss some of the remaining challenges.

We also explore some of the practical issues that are likely to arise if leases become more widespread. Here, we discuss the need for standardized procedures for lease auctions and contracts and the potential for international organizations to oversee these transactions, both to reduce the fixed costs of renegotiating these terms from scratch for every transaction and to limit discretion over parts of the contract that could be abused by corrupt officials. We also discuss issues related to antiquities where existing property rights are unclear, and how a combination of leases and a registration system might address these issues. In the conclusion, we suggest that the next step toward a market in leases for antiquities would need to involve active discussions among potential participants and stakeholders.

The Illicit Antiquities Market

Estimates of the size of the black market in antiquities are imprecise and vary widely from $300 million up to $7.4 billion per year (Atwood 2006; Calvani 2009). While these monetary values are relatively small, the damage that the trade has done to the archeological record is extensive. Table 1 provides evidence collected primarily by Brodie, Doole, and Watson (2000) and Brodie and Renfrew (2005) on the scale and scope of damage in selected settings where looted sites were known to researchers and where the damage could be documented (for additional discussion, see Atwood 2006; Gill and Chippindale 1993; O'Keefe 1997; Toner 2002). Given that many sites and objects are unknown, these examples probably represent only a small fraction of the damage done by the illicit trade of antiquities.

The illicit trade in antiquities persists despite export bans, international conventions, and bilateral treaties. Most source countries now see at least certain antiquities as constituting national patrimony—and ban their export. Beginning with restrictions imposed by the Vatican on the right of churches to sell off relics and art, export bans have spread to 140 countries around the world. A total of 119 countries, including the United States and most of Europe, have ratified either the 1970 UNESCO (United Nations Educational, Scientific and Cultural Organization) Convention on the Means of Prohibiting and Preventing the Illicit Import, Export and Transfer of Ownership of Cultural Property or the 1995 UNIDROIT (International Institute for the Unification of Private Law) convention designed to bolster enforcement of these

Table 1

Scope of Damage to Archeological Record by Illegal Trade in Antiquities (Selected Examples)

Italy	Italian Carabinieri recovered more than 326,000 antiquities from clandestine excavations between 1969 and 1999, of which 100,000 were recovered between 1994 and 1999 (Pastore 2001, p. 159).
Greece	Greek police recovered over 23,000 artifacts between 1987 and 2001 (Doole 2001).
Spain	A 2000 study of Andalusia, Spain, reported damage due to illicit excavation in 14 percent of known archeological sites (Fernández Cacho and Sanjuán 2000).
Belize	A 1983 study of Mayan sites in Belize found that 58.6 percent of sites surveyed were damaged by looters (Gutchen 1983).
Ecuador	A single Italian dealer was found to have illegally removed nearly 12,000 antiquities from Ecuador, where hundreds of sites have been damaged (Brodie, Doole, and Watson 2000).
Cyprus	In 1997, German police in Munich revored 50–60 crates containing 138 icons, 61 frescoes, and 4 mosaics that had been torn from the walls of a north cypriot church (Watson 1998, p. 11).
Turkey	Between 1993 and 1995, there were over 17,500 official investigations into stolen antiquities. A recent Turkish government lists antiquities smuggling as the fourth largest source of illicit income, after arms, drug smuggling, and fraud (Kaye 1995).
Iraq, Afghanistan, and Syria	Iraq, Afghanistan, and Syria have experienced widespread looting both from their museums and from the archaeological sites that once were under government control (Lawler 2003; Feroozi and Tarzi 2004; Parkinson, Albayrak, and Mavin 2015).
Pakistan	A survey in a district of northern Pakistan showed that 45 percent of the Buddhist shrines, stupas, and monasteries had been badly damaged or destroyed by illegal excavations (Ali and Coningham 1998).
Cambodia	In Cambodia, decorative friezes and sculptures belonging to Khmer period temples have been systematically looted. A single lorry stopped on the Cambodian-Thai border was found to contain 117 sandstone carvings from the twelfth-century AD temple of Banteay Chhmar (Brodie, Doole, and Watson 2000). Looting in Cambodia has been so extensive that the government has replaced many of the statues in the region of Angkor with replicas so that the originals can be more securely stored (Jessup 2004).
Mali	Between 1989 and 1991, a regional survey in Mali discovered 830 archaeological sites, but 45 percent had already been damaged, 17 percent badly. In 1996, a sample of 80 were revisited and the incidence of looting had increased by 20 percent (Bedaux and Rowlands 2001).

export bans by committing to return all antiquities declared national property and illegally exported after the ratification dates of the convention.

In the United States, the Convention on Cultural Property Implementation Act (1983) implements portions of the 1970 UNESCO convention and prohibits the import of stolen antiquities that have been documented in the inventory of a public or secular institution in countries that are signatories to the convention. The legislation also allows for bilateral agreements with individual countries to restrict importation of specific classes of archaeological artifacts. The United

States currently has agreements with Cyprus, Italy, Bolivia, El Salvador, Guatemala, Honduras, Nicaragua, Peru, Mali, and Cambodia (Brodie and Renfrew 2005). The US legislation provides for civil forfeiture of antiquities but does not include criminal penalties. The United Kingdom's 2003 Dealing in Cultural Objects (Offences) Act goes further and makes it a criminal offense to deal in antiquities that have been illegally excavated or removed from an archaeological site anywhere in the world after the Act came into force (Brodie and Renfrew 2005). Many source countries have adopted a de facto policy that antiquities in government hands cannot be sold.

In addition, many countries have declared unexcavated antiquities to be national property and have made it illegal for private citizens to search for and extract antiquities still in the ground. Such laws increase the ability of a country to prosecute smugglers in importing countries (for a broader legal discussion, see Bator 1982; Borodkin 1995; Phelane 1993). Because international treaties are not retroactive, antiquities exported prior to a treaty's ratification therefore enjoy a different legal status than those exported after a treaty's passage. Thus, those involved in the antiquities trade may obfuscate or destroy information about the origin and context surrounding antiquities so that illicit antiquities can more easily be passed off as legal ones. The pooling of licit and illicit material makes policing the antiquities market difficult and has made the systematic identification of stolen antiquities difficult, if not impossible.

The mixing of licit and illicit objects has also created scope for forgeries in the antiquities market. These forgeries have become increasingly sophisticated over time, and many are hard to identify even by well-trained archaeologists. As one example, marble cannot be carbon-dated and there is no easy way to test whether an object made of marble is genuine (Brodie, Doole, and Watson 2000). The potential that objects are forged makes it difficult for archaeologists to draw scientific conclusions from objects that have not been excavated from pristine sites (Chippingdale and Gill 2000; Gerstenblith 2007).

The trade of illicit antiquities typically flows from lower-income to higher-income countries, with a long chain of intermediaries. As one might suspect, trade flows are also influenced by the extent to which corruption is possible within a source country. Using the gap between the reported value of exports of cultural property and the reported value of imports of receiving countries as a proxy for the size of illicit trade in this market, Fisman and Wei (2009) find that more corrupt countries have a larger import-export gap and that this relationship is strongest in artifact-rich source countries. The intermediaries manage the smuggling of antiquities across borders and can try to obfuscate the origins of antiquities by shifting them through multiple jurisdictions.

Studies that trace the supply chain of an illegally traded antiquity estimate that the original holders of antiquities typically receive less than 2 percent of the antiquities' final sales price (Beech 2003; Brodie 1998). For example, the Euphronius Krater, an ancient Greek vase at the heart of a long-standing dispute between New York's Metropolitan Museum of Art and Italy, was bought by the museum for $1 million in 1972—but is reported to have been purchased by the looter who first

illegally excavated the vase from a previously unknown site for roughly $21,500 (8,800 pounds) one year earlier (Slayman 1998).

This markup of the prices of antiquities along the chain of intermediaries is due, in part, to the cost and risk associated with illegally exporting antiquities from the source country, importing them into the receiving countries, and finding a buyer. Moreover, there is limited competition among intermediaries and they are better informed than sellers, and hence likely to extract a large portion of the rents associated with looted antiquities.

Most museums will not knowingly accept antiquities that have been illegally excavated after 1971, based on the legal framework developed in the 1970 UNESCO convention. There is controversy over whether some museums are doing enough to ensure their collections do not have illicit material, or even may be attempting to relax their criterion for accepting material without provenance (Brodie and Renfrew 2005; Gerstenblith 2003). With such concerns duly noted, the rules are intended to reduce the demand for excavation and to prevent private collectors from obtaining tax breaks or public recognition for donating illicit material (Brodie and Renfrew 2005). However, an unwanted byproduct of these laws is that many antiquities now in private collectors hands do not have a path by which they can reenter the public domain.
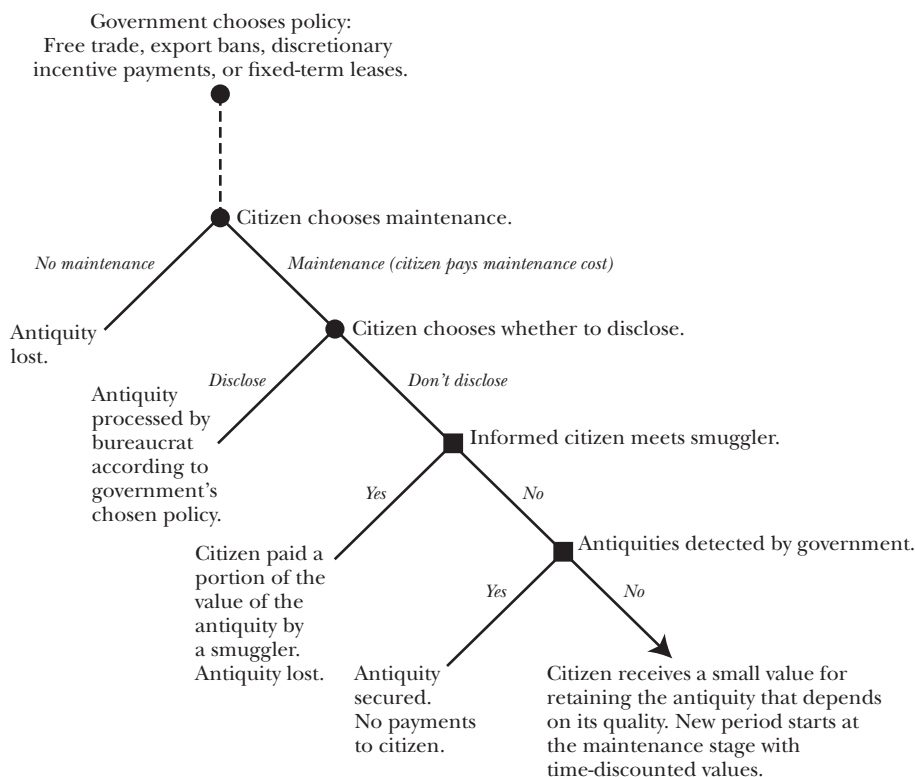
Antiquities Unknown to the Government

Governments that are seeking to preserve antiquities face a policy challenge. Many antiquities and sites are known to private citizens, but are either unknown to the government or too remote to protect directly. In addition, antiquities require maintenance spending in order to preserve them. Such settings naturally lead to an "agency" problem where the policies put into place by the government can influence the decisions of individuals who have *de jure* ownership over antiquities, but the government cannot fully control their decisions.

In economics, problems of agency are typically studied using tools from information economics and contract theory. In these areas, the private information of individuals plays an important role in understanding how policies shape outcomes. In particular, an individual who has private knowledge about the location of an object has the choice of whether to reveal it. As this knowledge has an implicit value, private citizens must be compensated for the costs of maintaining antiquities and revealing their knowledge of the antiquities to the government.

At least some countries, for some antiquities, may be amenable to short-term leases. We argue that these leases are likely to be a useful tool in providing information rents and overcoming the agency problem. In particular, leases can provide financial resources to individuals that reveal information to the government and directly link the value of rents to the value of the antiquity. At the same time, the short-term nature of leases reassures the government (and the population) that antiquities are owned by the country and will eventually be returned.

Figure 2
Timing Chart

Government chooses policy:
Free trade, export bans, discretionary
incentive payments, or fixed-term leases.

Citizen chooses maintenance.

No maintenance

Maintenance (citizen pays maintenance cost)

Antiquity
lost.

Citizen chooses whether to disclose.

Disclose

Don't disclose

Antiquity
processed by
bureaucrat
according to
government's
chosen policy.

Informed citizen meets smuggler.

Yes

No

Citizen paid a
portion of the
value of the
antiquity by
a smuggler.
Antiquity lost.

Antiquities detected by government.

Yes

No

Antiquity
secured.
No payments
to citizen.

Citizen receives a small value for
retaining the antiquity that depends
on its quality. New period starts at
the maintenance stage with
time-discounted values.

Source: Authors' creation.

Our arguments come from the study of a dynamic game that models the flow
of antiquities currently in the possession of informed private citizens into different
potential final outcomes. We use the game to explore how government policies
affect outcomes and to understand the advantages and disadvantages of different
policies. The timing chart for this model is provided in Figure 2.[3]

We begin by assuming that "informed citizens" have antiquities in their posses-
sion or know the location of an unexcavated site and that the existence of the
antiquity or site is unknown to the government. Because a central concern in the
antiquity market is the preservation of antiquities, we consider situations in which
the preservation of an antiquity requires a maintenance payment at the beginning of
each period by the owner. For unexcavated antiquities, this could be the opportunity

---

[3]The online Appendix provides an algebraic formulation of the model described here. In particular,
it shows how to derive conclusions about the superiority of leases over other policy choices in this
framework—and what assumptions are needed to support that conclusion.

cost of not using the land and the cost of protecting the site from damage by vandals or theft by looters. For antiquities that are already excavated, maintenance could involve proper storage that controls for heat, humidity, and sunlight and that mitigates the risk of damage from fire, flood, vandalism, or theft. For convenience, and rather than seeking a relationship between maintenance effort and costs of damage, we proceed with the assumption that maintenance is binary; that is, if it is not paid, the antiquity or site is lost or destroyed.

The next decision for a citizen who has made the maintenance payment is whether to disclose the antiquity to the government. If the citizen discloses the antiquity, then the government policy with regard to antiquities applies. We will consider four policies: (1) a *free trade* policy, in which the antiquity is sold to a foreign collector or government, with the proceeds paid to the informed citizen who discloses the antiquity; (2) an *export ban* policy, in which the government confiscates all antiquities it can find and makes no payments to informed citizens; (3) a *discretionary income payments* policy, in which the value of an antiquity is assessed by a potentially corrupt government bureaucrat and a payment is made to the informed citizen based on the assessed value; and (4) a *fixed-period leases* policy, in which a foreign collector pays for the item to be displayed outside the country for a time, with a share of the proceeds paid to the informed citizen who discloses the antiquity.

However, faced with the policy chosen by the government, the informed citizen may alternatively decide not to disclose the item. In this case, the informed citizen might meet a smuggler, who is part of the chain of intermediaries that can ship the item to a foreign buyer, at which point the antiquity is lost. Smugglers pay the informed citizen less that the citizen would receive under a free market.

Objects not sold to the smuggler may eventually be detected by the government. We will assume that the government can detect a small subset of antiquities or sites each period and can excavate (if still in situ) or confiscate them without a payment to the individual.

Finally, we assume that objects are not seized by the government are retained by the informed citizen who may receive some use value for retaining the antiquity. The game then repeats itself starting with the maintenance decision and with all values time discounted. In this way, we can study the transition of antiquities over time and identify how different policies influence the share of objects that are (1) destroyed due to nonmaintenance, (2) lost to the illicit market, or (3) secured by the government. We compare government policies based on the discounted utility that it receives from secured antiquities minus any transfers it must make to secure them.

We also allow for there to be a mix of high- and low-quality objects and allow for low-quality objects to be forgeries. Whether antiquities are of high or low quality is not perfectly observable and can only be assessed by experts. These experts could be either foreign collectors interested in buying the antiquity or bureaucrats in the home country.

Finally, we do not explicitly model the illicit market, but we do make some assumptions about how the enforcement capabilities of the home government

influence the value and likelihood of illicit trade. First, we assume that as enforcement capabilities increase, the cost of smuggling increases and the number of potential intermediaries decreases. This implies that the probability of matching with a smuggler and the proportion of final value that the citizen receives from the smuggler are decreasing in enforcement capabilities. Second, we assume that the probability that an antiquity is detected by the government each period is increasing in its enforcement capability.

The Inefficiencies of Free Trade and Export Bans for Antiquities

As the traditional debate about antiquities centers around the polar policies of free trade and export bans, it is useful to first analyze these policies through the lens of the model. Under free trade, objects that are revealed are free to be sold abroad or to the government. Because the informed citizens receive the maximum possible value of an antiquity revealed, free trade gives strong incentives to both maintain and reveal antiquities. However, if the government wishes to secure an object, it must compete on the international market for every antiquity and therefore faces potentially high costs for securing its cultural heritage. As citizens who have private information about an antiquity or site are not necessarily owners in a legal sense, large transfers to these individuals are likely to be seen as waste from the perspective of the government. Thus, free trade is likely to be a very costly policy for the government because the government must compete for resources it already legally owns.

Under an export ban, the government can exercise its legal rights to seize objects it finds but faces two potential issues. First, because informed citizens do not receive anything for objects that are revealed, there is no incentive for revealing antiquities to the government and citizens will always sell antiquities to the illicit market when given the chance. As such, securing objects is a slow process and a subset of antiquities are likely to be smuggled abroad.

Second, because informed citizens must invest to maintain their antiquity, there is a limit to what the government can recover through enforcement. If enforcement is toughened and the expected return of holding the antiquity falls below the maintenance costs, informed citizens may lose their incentive to protect and maintain antiquities, leading to their destruction.

Thus, export bans are likely to lead to a large proportion of antiquities lost to smuggling in environments where there is limited enforcement—and may also lead to destruction in environments where the maintenance cost is high and the private benefits for antiquities are small. In this way, an export policy based purely on enforcement (as opposed to a policy that also rewards for revelation) may generate fewer incentives for protection and ultimately lead to inefficient social outcomes.

In settings where private individuals are holding illegal antiquities and direct property rights are clear, one option is to design incentives for revelation based primarily or exclusively on a threat of punishment. However, such policies are unlikely to be effective in most settings. For antiquities that are unexcavated, it would be inherently difficult to prove an individual is aware of a site and hold that person directly accountable for the site's outcome. For other antiquities, it is likely that many

innocent individuals would get swept up in these policies, because it may be difficult to know if antiquities they come into contact with are licit or illicit, or forgeries.

Incentives from Discretionary Payments and Leases

The preceding discussion suggests that augmenting an export ban policy with explicit incentives for revealing the location of antiquities may improve social welfare. Payments for revelation may not only resolve the information asymmetry, but also provide incentives for citizens to maintain their antiquities in the first place. The challenge, then, is how to provide these incentives in a way that (1) secures and retains objects in the long run and (2) is not overly vulnerable to corruption.

A first policy that might come to mind is to provide the informed citizen with a discretionary incentive payment high enough to ensure that the object is maintained, and just high enough that the citizen prefers to reveal their object to the government rather than wait and sell it to a smuggler. However, the policy may run into difficulties, as antiquities vary in quality and assessing quality requires expertise. In these cases, there is scope for bureaucrats who might be assigned to the task of making discretionary payments to misidentify the quality of objects in exchange for bribes. Misidentifying the quality of objects can make incentive payment programs very expensive, as low-quality antiquities may be passed off as high-quality ones. Forgeries can complicate the situation further: if there are rewards for individuals who provide antiquities, there will potentially be rewards for forgeries as well.

By requiring valuation of antiquities by bureaucrats who can gain privately from misreporting, discretionary incentive payments will lead to distortions both in the amount paid to secure antiquities and in the quality of the antiquities secured. Corruption may also lead to the purchase of some forgeries in programs that use discretionary incentive payments, and the payments associated with these forgeries can vastly diminish the efficacy of discretionary payment programs.

In comparison, the advantage of a lease program is that the information rents generated by an antiquity can be linked directly to its value without relying on government assessments. Private citizens who reveal their objects can receive the proceeds that are generated from leasing the objects externally. This option can induce informed citizens to maintain their antiquity and to reveal them to the government, because the alternative of waiting to sell on the black market requires informed citizens to share rents with smugglers and exposes them to the risk that the government will confiscate an unreported antiquity before the citizen can match with a smuggler. Further, the social welfare loss from temporarily sending antiquities abroad as part of a lease program is likely small relative to the costs of permanently losing some antiquities under an export ban without revelation or paying market rates under a free market.[4]

---

[4] The formal proof of the superiority of leases over the alternatives does require that the discount rates for the government of the source country are sufficiently small. If discount rates are very high, a very impatient source country may prefer to maintain export bans over leases because such bans allow them to capture and use at least some antiquities in the current period.

The necessary length of the lease will need to be adjusted depending on the value that a smuggler can provide to an informed citizen relative to the market price for legally transferred leases. The necessary lease length will be shorter given the greater the ability of the country to police illegal markets and the smaller the value of illicit antiquities is to licit antiquities to final purchasers. We thus see leases as a complement to enforcement policies and policies that actively reduce foreign demand for illicit material. Leases are especially likely to outperform discretionary incentive payments for high-quality objects, where the potential shortfalls of government payments are greatest.[5]

In short, even if a source country places a very high value on keeping antiquities at home, because the social value of the antiquity at home is greater than the value to foreign collectors, a government constrained by a lack of information, or resources, or concerns about incentives for information revelation and private maintenance, as well as risks of corruption, might sensibly choose to allow fixed-term and thus temporary leases for display outside the country.

Antiquities in the Hands of the Government

In this section, we consider the case of antiquities that have been previously secured by the government: thus, the issues of incentives for "informed citizens" to reveal what they know about the location of antiquities do not apply. Instead, we explore why trade restrictions might be placed on antiquities already owned by the government and explore how fixed-duration leases might improve outcomes.

As a starting point, consider a situation in which a country is poor today but may become richer in the future. In this setting, the net value of an antiquity to a foreign collector might be greater than the net value domestically in the short term, but there may be value in securing an object in the long run. Thus, one might ask whether it makes sense for a government to sell objects in the short term, use the money to protect other parts of its collection, and then buy them back in the long run.

A first issue with such a policy is that it may be very difficult to repurchase objects sold abroad in the future. Antiquities that end up being donated to foreign museums often have very strict covenants placed on them that make it difficult for them to be resold, even back to the country of origin. Further, even if an antiquity is for sale, the cultural significance of an antiquity may be hard to quantify, leading to asymmetric information. This can result in a hold-up problem that makes sale-and-repurchase schemes inefficient.

---

[5]A government sales tax on antiquities could also allow the value of antiquities to be split and could automatically link incentives for informed citizens to reveal antiquities to the antiquities' values without government discretion. However, sales of antiquities are problematic, for reasons discussed in the text. Also, tax programs may be easier to game than leases. For example, antiquities may be broken into parts, sold in separate lots at low prices, and reassembled by a colluding foreign collector to bypass taxation.

A potentially bigger problem is that allowing for sales creates new avenues for corruption. If sales by the government are legal, then a corrupt official who is appointed to oversee antiquities may sell the cultural patrimony of the country and expropriate some of its value. As such, even if the chance of corrupt officials coming to power is small, the country may adapt laws, such as blanket export bans, that are difficult to circumvent.

In the context of antiquities owned by the country, an export ban can be viewed as a blunt tool to constrain corrupt future officials from acting in a malevolent way. A byproduct, however, is that this tool also prevents honest officials from managing their cultural heritage effectively and generating revenue that might be used to secure at risk sites.

Again, we see leases as a way of balancing concerns about corruption with efficiency considerations. Leases can provide resources for maintenance; indeed, even the possibility of future leases can improve current maintenance incentives.[6] Further, even if a corrupt official is paid for negotiating a low-price lease, at least that official can only appropriate the value of the antiquity during the lease period: after the lease, the antiquity will either be returned to its own home country or the terms will be renegotiated with a different official, who has some chance of not being corrupt.

Leases also offer discretion to honest officials in settings where the source country is experiencing civil turmoil. If for instance, there is some chance that antiquities are threatened or the state is unable to maintain or protect them, officials could move antiquities abroad to protect them. In cases of armed conflict, restricting all transactions to leases also reduces incentives for combatants to search for antiquities for the purpose of selling them to fund war efforts. Looting from museums and illegal excavations have been common in the recent conflicts in Iraq and Syria (Brodie, Doole, and Watson 2000; Parkinson, Albayrak, and Mavin 2015).

Leases offer a way of balancing concerns about corruption with efficiency considerations. In particular, short-term leases can restrict the long-term damage by corrupt officials, while still giving benevolent ones the ability to make Pareto-improving short-term trades. In the next section, we discuss how standardized contracts may help to sustain lease arrangements for antiquities.

Creating Standardized Contracts for Leasing Antiquities

There are already some prominent examples of leases of antiquities, as we noted earlier, including the King Tut exhibit from 2005 to 2008, the Menil collection from

---

[6]The loan program used by the Menil Collection in Houston, discussed earlier, suggests that the party receiving the antiquity for a time might provide in-kind payments—like building exhibition and curation facilities in the source country—rather than cash payments. This preference might signal something about the characteristics of the relevant government officials (for discussion, see Roth 2007; Bénabou and Tirole 2011). For example, cultural officials in the source country may have de facto veto rights and may prefer in-kind transactions rather than cash transactions, because the proceeds of a lease transaction may otherwise go outside of their ministries.

Cyprus, and the artworks leased from the Louvre to Abu Dhabi. The broader feasibility of lease markets is being discussed actively in the art disciplines; for example, lease contracts have been briefly mentioned in press by Butcher and Gill (1990), Asgari (1993), and Gerstenblith (2001).[7]

However, if one rationale for leases is to provide incentives and methods of valuation in a setting where there is a risk that source-country officials may be corrupt, then both the structure of the leases and the leasing process itself need to be designed to limit the power of these same corrupt officials. A primary protection against sweetheart deals is to standardize provisions in a contract for leasing antiquities, such as the length of the lease, the storage conditions, and insurance provisions. In dimensions where some flexibility may be required, such as rules governing the qualification of bidders for the lease, it is important to insist on transparency in the process. Here, we discuss a series of standardizations and safeguards that could be implemented to reduce the scope for corruption in lease contracts.

First, an important consideration in the design of lease contracts for antiquities is to assure that proper care is taken with the items. Both the Menil Collection and King Tut exhibit shed light on the current practice. In common law countries, museum loan programs are considered "bailments," where one party gives possession of the antiquity to another for safekeeping. Museum contracts augment the common law by specifying the arbitration process for damaged goods and providing requirements for care: for details of the contracts used in museum exchange and lending programs, see Simpson (n.d.). In the King Tut case, the lease agreements specified transportation, display, and storage conditions. Based on the success of traveling exhibits and the extensive level of museum-to-museum lending that currently exists, our sense is that these issues could be adequately addressed contractually, as long as the legal system in the receiving country is sufficiently well functioning.

A second set of concerns is to reduce the risk of collusion between bidders and officials in the leasing process itself. Transparent auctions with fixed rules make it harder for the official running the auction to collude with a specific bidder. This not only helps mitigate concerns about corruption, but also may mitigate concerns by bidders that winning an auction could lead to bad public relations. Standardized procedures to address corruption would include setting initial consultation periods that would allow qualified bidders to verify the condition and quality of goods prior to auction. Qualified bidders would need to meet some minimal level of trustworthiness along with demonstrating the financial resources to provide the proper level of insurance and care. To reduce potential corruption, the process of qualifying bidders and the scoring rules used to select winning bids would also likely need to be standardized.

Third, the length of leases probably need to be adapted to circumstances. If leases are to provide incentive for private individuals to make information public

---

[7]Asgari (1993) (as quoted in Erdem 1993) argues that ten-year leases may be used between major museums to reduce incentives to purchase illicit artifacts. Gerstenblith (2001) proposes leases between museums to decrease demand for new pieces from foreign countries.

rather than to deal with smugglers, then leases should be long enough to provide meaningful compensation for private individuals. In the case of leases used for objects in government control, lease lengths must be long enough to make trans-action costs worthwhile, but short enough to limit rent extraction from future generations by corrupt officials. For unexcavated sites, where both information rents and site protection must be organized prior to excavation, relatively longer leases may be desired. By contrast, for already excavated antiquities, we believe shorter leases are likely optimal. Lease lengths may also need to be longer for objects where transportation costs are high relative to the value of the antiquity.[8]

Finally, it may be desirable for an international organization to facilitate a lease system by serving as a standard intermediary and for contracts to be written in a stable country with strong contract enforcement capabilities. An international organization is more robust to regime change in the source- and renting-countries and is likely to have an easier time committing to follow standard policies. The participation of an established international organization could also provide legitimacy and address concerns that officials in source countries proposing such arrangements were pursuing this approach out of corrupt motives. To reduce costs and increase the number of international bidders, an international organization could also work with several established auction houses to run the auctions.

Shifting the norm of antiquities markets to a lease system, managed through an international organization, could also provide for safeguards in times of political upheaval and war (as in Kremer and Jayachandran 2006; Pogge 2001). It is well-documented that looting of antiquities is widespread in times of war and afterwards; for example, Russell (1997) and Feroozi and Tarzi (2004) discuss the situation in the wars in Iraq and Afghanistan. The 1954 Hague convention on cultural prop-erty and its 1954 First Protocol and 1999 Second Protocol try to provide additional protections for antiquities in times of war, but in practice fall short. In this setting, restricting market participants in antiquities to standardized leases may be useful, because even a suboptimal lease has an end date, so that with a lease, the transfer of antiquities as a result of conflict would be limited to a single generation.

Unexcavated Sites, Repatriation, and Other Applications of Leases

For unexcavated antiquities, property rights are often ill-defined or weakly enforced. In these cases, there is concern that increasing the value of an antiquity—say, by having a well-recognized leasing system—could increase incentives to

---

[8]There are, of course, cases where the government may wish to compensate for the revelation of objects and also generate revenue from them. In these cases, allowing for a longer lease length and taxing the proceeds is likely easier than running an initial auction for information rents and a subsequent auction for revenue. This approach is especially relevant for objects with high transportation costs where bidders may not value very short leases.

appropriate the antiquity. This possibility has implications both for the applicability of leases and for implementation.

When antiquities are unexcavated, safeguards must be taken to prevent appropriation risk. One possibility is to grant payments (from leases or other sources) only to individuals who report the location of sites and to make the reward contingent on confirmation by archeologists that antiquities in the site have been undisturbed. This approach would involve no amnesty on objects that had been illegally excavated. An alternative option would be to allow for reasonably long lease lengths for antiquities recovered from legal excavations, but only short leases for individual antiquities without such a provenance. This two-tiered policy can create incentives for disclosing both unexcavated sites and antiquities that are currently illegally in private hands, without increasing incentives to loot unexcavated sites.

The first approach is likely to be favored by archaeologists, who stress the importance of professional excavation and who may view it as repugnant to pay an informed citizen a reward for revealing objects that might have resulted from past illegal digging. However, the alternative two-tiered approach is likely to recover more objects and may generate information on sites that are currently in the process of being looted.

Unexcavated sites may be well known to a number of individuals and granting exclusive property rights to one of them may create incentives to appropriate the antiquity by others. It may make sense for some share of the proceeds of lease auctions to go to the broader community. There is likely to be reasonable scope to provide sufficient payments to multiple interests: remember, in the current environment, the price at which antiquities are purchased from informed citizens is often less than 2 percent of the final price, which itself is depressed by the illicit nature of illegally traded objects.

Given the limited budgets in many source countries for protecting antiquities, reported sites are likely to require external funding for security and proper excavation. In these cases, firms or foreign entities could provide financial assistance for security and excavation in return for rights to lease or hold a subset of antiquities that are excavated for a fixed number of years. In the early twentieth century, foreign archaeological expeditions often agreed to work for a share in the excavated antiquities. For instance, Iraq had a policy where half of the duplicate antiquities from an expedition would be allowed to leave with the excavating party, but all unique items went to the central museum (Bernhardsson 2005). Using a lease policy to finance security at a site could have similar effects: it ties the value of the security contract to the value of antiquities at the site and thus provides proper incentives across a broad range of sites with varying value to looters.

A lease system could also be linked to a registration system, which could help reduce incentives for extra-legal appropriation. Leases for individual antiquities could be made contingent on registering a description of the institution or person who is seeking to lease the antiquity and a set of photos that could be used to check against a registry of previously registered or stolen antiquities. There could also be

a waiting period under which anyone who wants to dispute ownership of the object could do so.

We believe that there are likely complementarities between lease and registration systems. On the one hand, informed citizens would have limited incentives to use a registration system without the financial incentive provided by the prospect of leases. Thus, a lease system makes the registration system more workable. On the other side, a registration system increases documentation in a way that makes it more difficult to pass illegal antiquities off as legal ones. An increase in legally leasable antiquities should reduce the price of illegally traded antiquities.

While leases are likely suitable for museum-quality antiquities, for many lower-quality antiquities (say, those where the government does not foresee a future value), the transaction costs associated with leases would be prohibitively expensive. For unexcavated antiquities, perhaps low-quality antiquities could be combined into lots and leased to researchers based on their scientific value. For already excavated antiquities, however, it may be infeasible to combine lots and the price per antiquity may end up being very low. In these cases, the government may wish to allow very low-quality antiquities to be sold abroad, but may be concerned that allowing sales contracts could give informed citizens incentives to misrepresent the value of antiquities and collude with foreign buyers to sell high-quality antiquities abroad. In fact, many antiquities are currently smuggled by painting them to look like replicas and then reversing the process once imported into the other country.

In this context, we view option contracts as potentially useful where antiquities are allowed to be sold but where the government has the right to purchase the antiquity in the future at an inflation-indexed multiple of the original sale price. Such contracts would create a path for repatriation of antiquities that were clearly misreported and would likely require less overhead than a long-term lease.[9] An alternative approach, currently being used on a small scale in Mali, is to use low-quality antiquities as collateral for loans. The Culture Bank program has been running in Mali since 1997 and gives small-scale loans to individuals who allow their objects to be conserved and exhibited in a community museum collection. Individuals are granted larger loans depending on how much they know of the history of the antiquities being deposited (Deubel 2006; Kersel 2012).

A final use of leases could be to create a pathway for either the repatriation of objects that are legally owned in other countries but now considered to be repugnant or for illegally exported to reenter the public domain. In the first case, many objects that were looted in past conflicts are legally owned by foreign museums. These objects are increasingly seen as repugnant and curators are often uncomfortable displaying them. Leases offer a useful way for both the foreign museum and the home country to bargain to a better overall outcome. By transferring ownership to the source country but simultaneously signing a long-term lease that

---

[9]We study option contracts in detail in an earlier version of the paper (Kremer and Wilkening 2012). Option contracts may also be useful for governments who are credit-constrained and want to collateralize antiquities to fund other essential services.

explicitly allows objects to be displayed, the current owner can remove the repugnance associated with the object, while the home country can secure the object for future generations.

A similar approach could be used for individuals who are holding antiquities with a murky provenance. Individuals who reveal illegally moved antiquities could be granted the right to lease them to others, and this arrangement could potentially improve outcomes because illicit antiquities are typically not in the possession of the person or institution who values them the most. However, this type of lease is equivalent to a partial amnesty program and has the disadvantage that collectors might purchase an illegal antiquity with the belief that it could be converted to a lease in the future.

Any amnesty program, in which a past illegality is excused, has a mixture of advantages (revealing previously hidden antiquities) and disadvantages (reducing credibility). Here, we do not seek to assess the overall costs and benefits of amnesties in this setting. But if amnesties are to be provided, leases offer a useful way to structure the transactions.

In some cases, the "leases" involved in this partial amnesty approach may not involve cash payments, and these in-kind leases are likely to have smaller behavioral and social effects than programs that use cash (Bénabou and Tirole 2011). As one example, both the Getty and the Metropolitan Museum of Art in New York City recently repatriated portions of their Italian collection back to Italy in exchange for long-term loans of Italian art with similar value. In the case of the Metropolitan Museum of Art, this agreement ended a legal battle which had been fought for almost 30 years (Kennedy and Eakin 2006).

More recently, ownership of a collection of 161 Cycladic artifacts assembled by Leonard Stern was transferred back to Greece, with a provision that the majority of the collection would be put on public display at the Metropolitan Museum of Art for 25 years. This compromise allowed the Metropolitan Museum of Art to show the collection (the preferred outcome of the collector) but ensured that the collection remained intact and was returned to Greece in the long run. The deal also involved the donation of the collection to a Delaware-based Institute of Ancient Greek Culture, which provided the collector with the tax benefits that exist by US law, along with cash contributions to the Metropolitan Museum of Art and the Museum of Cycladic Art in Athens for conservation of the collection.[10] New legislation was required to be passed in Greece to ratify the agreement and grant the Metropolitan Museum of Art interim lease rights. For a discussion of this deal and arguments for and against it, see Ioannidis (2022).

---

[10] Once donated to museums, antiquities in the foreign collector's country may be difficult to return to their country of origin. Museums are often constrained in selling antiquities in their collection due to restrictions placed on the collection by benefactors, by private charters, and by accreditation requirements to museum associations. For instance, the American Association of Museums allows for sales used only for the "direct care" of a collection (Pogrebin 2010). Museums may also fear that selling donated antiquities may deter future donations. Leases circumvent these issues by having the antiquity automatically return to the home country without the need for renegotiation.

Conclusion

When a market is viewed as repugnant, like the market for antiquities sold and transported out of source countries, a natural political response is to ban such activity and to back up the ban with a push to reduce international demand for antiquities by making collecting them or moving them out of the country socially unacceptable (Elia 1997; Gustafsson 2010; Renfrew and Elia 1993). However, when there is demand for antiquities from high-income countries and supply of antiquities in lower-income source countries, both historical and present experience show that the incentives lead to antiquities moving across national borders. Allowing antiquities to be exported under time-limited leases allows these demand and supply incentives some scope to function, while also offering an opportunity to protect antiquities and to keep ultimate ownership in the source country.

In this essay, we have argued that a lease-based policy offers a flexible approach that will often be superior to the other policies including export bans, discretionary incentive payments, and outright sales. The advantages of leases arise because they can address a variety of issues: the incentives of private individuals with knowledge of antiquities to make their discoveries public; concerns about corruption and shortcomings of enforcement in source countries; resource shortages in source countries that can make it difficult for them to protect sites or to maintain and display antiquities; political imperatives in source countries that may allow antiquities to be displayed elsewhere as long as they are scheduled to return; a potential resolution of long-standing conflicts that involve private and museum collections; and other issues.

Although our discussion has been rooted in practical concerns surrounding the market for antiquities, other concerns have gone unmentioned. For example, some collectors might prefer a lease that lasts as long as their lifetime, to mitigate concerns about the maintenance of property bequeathed to heirs who do not share the same tastes as the original collector. As another example, many countries encourage donation of private collections to domestic museums through preferential tax treatments. These preferential tax breaks can be over one-third of the value of an antiquity and are greatest for antiquities that have appreciated in value over time (Fullerton 1991). Thus, it might be useful to combine a lease that allows private ownership for part of the lease, together with donation to a museum for the remainder of the lease.

We believe that leases for antiquities are likely to be legally, administratively, and politically useful, but we are also keenly aware that such arrangements will succeed or fail base on how well they address practical concerns. If lease markets for antiquities are to become widespread, we believe an interdisciplinary panel of archeologists, museum officials, lawyers, economists, and government representatives from source nations is likely necessary to design standardized contracts and develop feasible institutional details.

References

Akinade, Olalekan Ajao. 1999. "Illicit Traffic in Cultural Property in Nigeria: Aftermaths and Antidotes." *African Study Monographs* 20 (2): 99–107.

Ali, Ihsan, and Robin Coningham. 1998. "Recording and Preserving Gandhara's Cultural Heritage." *Culture without Context* 9 (3): 10–16.

Asgari, Nusin. 1993. "Speech." Quoted in Suna Erdem, "New Trojan war highlights pillage of Turkey's past," *Reuters*, October 13, available on Factiva as document lba0000020011121dpad02wzo.

Atwood, Roger. 2006. *Stealing History: Tomb Raiders, Smugglers, and the Looting of the Ancient World.* St. Martin's Press.

Bator, Paul M. 1982. "An Essay on the International Trade of Art." *Stanford Law Review* 34 (2): 275–384.

Bedaux, R. M. A., and M. Rowlands. 2001. "The Future of Mali's Past." *Antiquity* 75 (290): 872–76.

Beech, Hannah. 2003. "Spirited Away." *Time*, October 13. https://time.com/archive/6956423/spirited-away-3/.

Beltrametti, Silvia. 2013. "Museum Strategies: Leasing Antiquities." *Columbia Journal of Law and the Arts* 36 (2): 203–60.

Bénabou, Roland, and Jean Tirole. 2011. "Identity, Dignity and Taboos: Beliefs as Assets." *Quarterly Journal of Economics* 126 (2): 805–55.

Bernhardsson, Magnus T. 2005. *Reclaiming a Plundered Past Archaeology and Nation Building in Modern Iraq.* University of Texas Press.

Boehn, Mike, guest. 2005. "New Blockbusters: Traveling Exhibits." *Talk of the Nation*, NPR, June 1. https://www.npr.org/2005/06/0¼675421/new-blockbusters-traveling-exhibits.

Borodkin, Lisa J. 1995. "The Economics of Antiquities Looting and a Proposed Legal Alternative." *Columbia Law Review* 95 (2): 377–417.

Brodie, Neil. 1998. "Pity the Poor Middlemen." *Culture without Context* 3: 4–6.

Brodie, Neil, Jennifer Doole, and Peter Watson. 2000. *Stealing History: The Illicit Trade in Cultural Material.* McDonald Institute for Archaeological Research.

Brodie, Neil, and Colin Renfrew. 2005. 'Looting and the World's Archaeological Heritage: The Inadequate Response." *Annual Review of Anthropology* 34: 343–61.

Brouwer, Andy. 2008. "Valmiki and Brahma." Personal blog, May 9. http://blog.andybrouwer.co.uk/2008/05/valmiki-and-brahma.html.

Butcher, Kevin, and David Gill. 1990. "Mischievous Pastime or Historical Science?" *Antiquity* 64 (245): 946–50.

Calvani, S. 2009. Frequency and Figures of Organized Crime in Art and Antiquities." In *Organised Crime in Art and Antiquities*, edited by Stefano Manacorda. International Scientific and Professional Advisory Council.

Chippingdale, Christopher, and David W. J. Gill. 2000. "Material and Intellectual Consequences of Contemporary Classical Collecting." *American Journal of Archaeology* 104 (3): 463–511.

Coggins, Clemency. 1972. "Archeology and the Art Market." *Science* 175 (4019): 262–66.

Deubel, Tara F. 2006. "'Banking on Culture': Microcredit as Incentive for Cultural Conservation in Mali."

In *Microfinance: Perils and Prospects*, edited by Jude L. Fernando, 115. Routledge.

Doole, Jenny. 2001. "In the News." *Culture without Context* 9: 16–23.

Elia, Ricardo J. 1997. "Looting, Collecting, and the Destruction of Archaeological Resources." *Nonrenewable Resources* 6 (2): 85–98.

Erdem, S. 1993. "New Trojan war highlights pillage of Turkey's past." *Reuters*, October 13, available on Factiva as document lba0000020011121dpad02wzo.

Fernández Cacho, Silvia, and Leonardo García Sanjuán. 2000. "Site Looting and the Illicit Trade of Archaeological Objects in Andalusia, Spain." Culture without Context 7: 17–24.

Feroozi, Abdul Wasey, and Zemaryalai Tarzi. 2004. "The Impact of War upon Afghanistan's Cultural Heritage." Archaeological Institute of America. https://www.archaeological.org/pdfs/papers/AIA_Afghanistan_address_lowres.pdf.

Fisman, Raymond, and Shang-Jin Wei. 2009. "The Smuggling of Art, and the Art of Smuggling: Uncovering the Illicit Trade in Cultural Property and Antiques." *American Economic Journal: Applied Economics* 1 (3): 82–96.

Freeman, Michael. 1997. "Photo: Harp Players below Worshippers of Brahma." In *Khmer Mythology*, edited by Vittorio Roveda, 153. Thames and Hudson. http://www.michaelfreemanphoto.com/.

Fullerton, Don. 1991. "Tax Policy toward Art Museums." In *The Economics of Art Museums*, edited by Martin Feldstein, 195–236. University of Chicago Press.

Gerstenblith, Patty. 2001. "The Public Interest in the Restitution of Cultural Objects." *Connecticut Journal of International Law* 16: 197–246.

Gerstenblith, Patty. 2003. "Acquisition and Deacquisition of Museum Collections and the Fiduciary Obligations of Museums to the Public." *Cardozo Journal of International and Comparative Law* 11: 409–65.

Gerstenblith, Patty. 2007. "Controlling the Internation Market in Antiquities: Reducing the Harm, Preserving the Past." Chicago Journal of International Law 8 (1): 169–95.

Gill, David W. J., and Christopher Chippindale. 1993. "Material and Intellectual Consequences of Esteem for Cycladic Figures." *American Journal of Archaeology* 97 (4): 601–59.

Gustafsson, Amanda. 2010. "Beware the Invisible." *Papers from the Institute of Archaeology* 20 (1): 97–110.

Gutchen, Mark. 1983. "The Destruction of Archaeological Resources in Belize, Central America." Included in "The Antiquities Market: News and Commentary on the Illicit Traffic in Antiquities" by Karen D. Vitelli. *Journal of Field Archaeology* 10 (2): 213-28.

Ioannidis, Sakis. 2022. "Debate: Cycladic Idol Deal Signals New Chapter in Hertitage Management." Ekathimerini.com, September 13. https://www.ekathimerini.com/culture/1192954/debate-cycladic-idol-deal-signals-new-chapter-in-heritage-management/.

Jessup, Helen Ibbitson. 2004. *Art and Architecture of Cambodia*. Thames and Hudson.

Kaye, L. 1995. "The View from the Bosphorous." Presentation at Art Theft Control discussed in Brodie et. al. "Looting and the World's Archaeological Heritage: The Inadequate Response." *Annual Review of Anthropology* 34: 344-61.

Kennedy, Randy, and Hugh Eakin. 2006. "The Met, Ending 30-Year Stance, Is Set to Yield Prized Vase to Italy." *New York Times*, February 3. https://www.nytimes.com/2006/02/03/arts/the-met-ending-30year-stance-is-set-to-yield-prized-vase-to-italy.html.

Kersel, Morag M. 2012. "The Value of a Looted Object: Stakeholder Perceptions in the Antiquities Trade." In *The Oxford Handbook of Public Archaeology*, edited by Robin Skeates, Carol McDavid, and John Carman, 253–74. Oxford University Press.

Kremer, Michael, Seema Jayachandran. 2006. "Odious Debt." *American Economic Review* 96 (1): 82–92.

Kremer, Michael, and Tom Wilkening. 2012. "Protecting Antiquities: A Role for Long-Term Leases." Unpublished.

Lawler, Andrew. 2003. "Beyond the looting: what's next for Iraq's treasures?" *National Geographic*, October 2003, 58+. Gale Academic OneFile (accessed March 16, 2025). https://link.gale.com/apps/doc/A338040707/AONE?u=unimelb&sid=bookmark-AONE&xid=b8ec178b.

O'Keefe, Patrick J. 1997. *Trade in Antiquities: Reducing Destruction and Theft*. UNESCO.

Parkinson, Joe, Ayla Albayrak, and Duncan Mavin. 2015. "Culture Brigade: Syrian 'Monuments Men' Race to Protect Antiquities as Looting Bankrolls Terror." *Wall Street Journal*, February 10. https://www.wsj.com/articles/syrian-monuments-men-race-to-protect-antiquities-as-looting-bankrolls-terror-1423615241.

Pastore, Giovanni. 2001. "The Looting of Archaeological Sites in Italy." In *Trade in Illicit Antiquities: The Destruction of the World's Archaeological Heritage*, edited by Neil Brodie, Jennifer Doole, and Colin Renfrew, 155–60. McDonald Institute for Archaeological Research.

Phelane, Marilyn. 1993. "A Synopsis of the Laws Protecting Our Cultural Heritage." *New England Law Review* 28: 63–108.

Pogge, Thomas. 2001. "Achieving Democracy." *Ethics and International Affairs* 15 (1): 3–23.

Pogrebin, Robin. 2010. "Museum Sells Pieces of Its Past, Reviving a Debate." *New York Times*, December 5. https://www.nytimes.com/2010/12/06/arts/design/06sales.html.

Posner, Eric A. 2007. "The International Protection of Cultural Property: Some Skeptical Observations." *Chicago Journal of International Law* 8: 213–31.

Prott, Lyndel V., and Patrick J. O'Keefe. 1984. *Law and Cultural Heritage.* Professional Books Limited.

Renfrew, Colin, and Ricardo Elia. 1993. "Collectors Are the Real Looters." *Archaeology* 46 (3): 16–17.

Riding, Alan. 2007. "The Lourve's Art: Priceless. The Louve's Name: Expensive." *New York Times*, March 7. https://www.nytimes.com/2007/03/07/arts/design/07louv.html.

Roth, Alvin E. 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives* 21 (3): 37–58.

Russell, John Malcolm. 1997. "The Modern Sack of Nineveh and Nimrud." *Culture without Context* 1: 8–20.

Shyllon, Folarin. 2000. "The Recovery of Cultural Objects by African States through the UNESCO and UNIDROIT Conventions and the Role of Arbitration." *Uniform Law Review* 5 (2): 219–40.

Simpson, Shane. n.d. "Collections Law: Legal Issues for Australian Archives, Galleries, Libraries and Museums." Powderhouse Museum. https://registrars.org.au/resources/standards-forms/collections-law-legal-issues-for-australian-archives-galleries-libraries-museums-2/?v=0b3b97fa6688.

Slayman, Andrew L. 1998. "Case of the Golden Phiale." *Archaeology* 51 (3); 36–41.

Toner, Mike. 2002. *The Past in Peril.* Southeast Archeological Center.

Watson, Peter. 1998. "The Sequestered Warehouses." *Culture without Context* 2: 11–14.

Wright, H., T. J. Wilkinson, E. Stone, and M. Gibson. 2003. *The National Geographic Society's Cultural Assessment.*

Basel Endgame: Bank Capital Requirements and the Future of International Standard Setting

Stephen Cecchetti, Jeremy Kress, and Kermit Schoenholtz

Over the past half-century, central bank officials and banking supervisors have developed regulatory standards for internationally active banks through the Basel Committee on Banking Supervision. The most recent round of reforms, known as Basel III, might seem esoteric, but it attracted widespread attention, including litigation threats against federal agencies (Kress 2024), numerous congressional oversight hearings, and even advertisements during nationally televised football games.

In 2023, three US bank regulatory agencies—the Federal Reserve, the Federal Deposit Insurance Corporation, and the Office of the Comptroller of the Currency—proposed the "Basel Endgame" rule to implement the most recent Basel Committee standards in the United States. Although the Basel Committee's framework did not explicitly require raising bank capital levels, the US proposal would have increased capital requirements for the largest banks by close to 20 percent. The banking industry and its political allies fiercely opposed the Basel Endgame

■ *Stephen Cecchetti is Rosen Family Chair in International Finance, Brandeis University, Waltham, Massachusetts. He is also Vice-Chair of the Advisory Scientific Committee, European Systemic Risk Board, Frankfurt, Germany; Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts; and Research Fellow of the Centre for Economic Policy Research, London, United Kingdom. Jeremy Kress is Associate Professor of Business Law, Stephen M. Ross School of Business, University of Michigan, Ann Arbor, Michigan. Kermit Schoenholtz is Clinical Professor Emeritus, Leonard N. Stern School of Business, New York University, New York City, New York. Their email addresses are cecchetti@brandeis.edu, kressj@umich.edu, and kschoenh@stern.nyu.edu.*

proposal, effectively killing it in its original form. At this writing (in mid-2025), the fundamental question of how or whether the United States will implement the final Basel III standards remains unresolved.

The stakes of this debate are significant. The US failure to implement the final Basel III agreement could undermine international cooperation and multilateral standard setting in bank regulation. The United States' noncompliance also could jeopardize the ability of US banks to compete abroad. There is thus a compelling case that implementing the Basel III agreement serves everyone's interest, despite industry's opposition to the 2023 Basel Endgame proposal.

This essay begins by explaining what bank capital is and why policymakers set minimum bank capital requirements. We review the history of the Basel Committee and international banking standards, including the US implementation of the initial Basel III reforms in 2013 and the core provisions of the final Basel III framework announced in 2017. We then analyze the 2023 Basel Endgame proposal and discuss how the debate conflated two distinct issues: (1) whether to comply with international regulatory standards, and (2) whether to raise capital requirements for large banks. US authorities gambled that they could achieve these goals simultaneously—and so far have achieved neither.

While implementing regulations that address both issues would likely make the US banking system safer and more competitive globally, these questions are logically separable. The United States could implement the final Basel III standards while maintaining current capital levels by making offsetting adjustments to existing capital requirements. Implementing international standards in a capital-neutral manner would preserve global cooperation in banking regulation, leaving the separate question of raising capital requirements for future consideration.
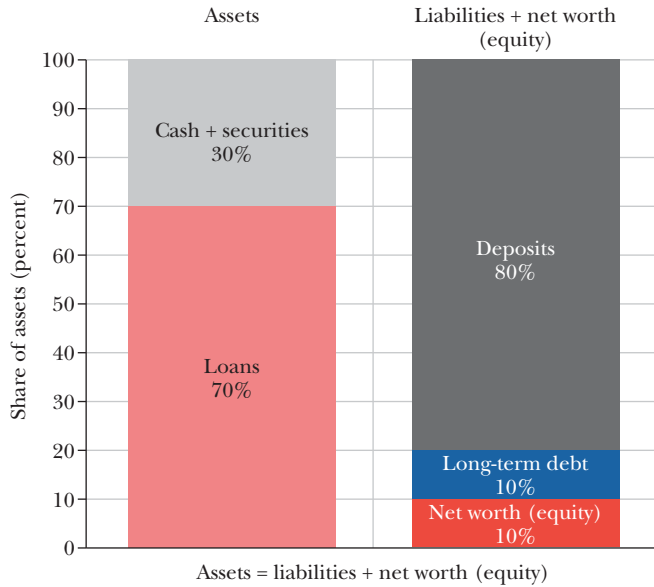
What Is Bank Capital?

Before delving into the details of international bank capital standards, it is useful to explain what bank capital is and why the debate over bank capital requirements is so contentious.

Figure 1 depicts a simplified bank balance sheet. Most bank assets are loans, which can fluctuate in value depending on financial and economic conditions. Funding comes primarily from deposits, with the remainder composed of long-term debt and net worth (equity). Of course, bank assets must always equal liabilities plus net worth, so when the value of assets declines, equity owners take the first loss.

Figure 1 illustrates two consistent definitions of a bank's capital (or, equivalently, its net worth). First, capital is the accounting residual that remains after subtracting a bank's liabilities (deposits and long-term debt) from its assets (cash, securities, and loans). If a bank's liabilities are fixed, when its assets decline in value, its capital shrinks. For example, if a borrower fails to repay a loan or a security loses value, the bank's capital will decline commensurately. Second, capital is the buffer that separates the bank from insolvency, which occurs when a bank's liabilities exceed its assets.

Figure 1
A Stylized Bank Balance Sheet

Source: Authors' creation.
Note: Percentages are roughly based on aggregate data for the US banking system.

Contrary to a common misconception, bank capital is neither cash locked away in a vault nor reserves held at the central bank. Rather, bank capital is a source of funds that banks routinely deploy for lending, trading, and other activities—it is never idle. Accordingly, as Figure 1 highlights, capital appears on the liability side, *not* the asset side, of a bank's balance sheet. Furthermore, new equity is an instrument that the bank "issues," rather than an asset the bank "holds." A bank also can add to its equity by retaining earnings, rather than making payouts to shareholders.

Bank managers (accurately) view equity capital as a relatively costly source of funding because the owners of equity demand a higher return than the interest rate that banks pay to depositors.[1] Thus, banks prefer to fund their assets by borrowing—issuing deposits or bonds—rather than by issuing equity (or through retained earnings). Stating it slightly differently, to increase their return on equity, profit-seeking banks will try to increase their leverage (the ratio of assets to equity).

---

[1] According to the Modigliani-Miller theorem (as discussed in this journal by Miller 1988), a firm's value is independent of its liability structure, so owners and managers should be indifferent between debt and equity financing. However, the theorem assumes that there are no subsidies, taxes, or bankruptcy costs. As we will see, it is precisely these deviations from the Modigliani-Miller assumptions that lead banks to favor debt over equity financing.

But banks' preference for debt over equity financing largely reflects public policies that create a gap between the *private* and *social* costs of capital. Three key factors drive this wedge. First, corporate tax law favors debt: interest payments to depositors and bondholders are treated as an expense that lowers a bank's tax burden, but a bank must pay dividends to equity investors from after-tax profits. Second, limited liability creates perverse incentives when bank capital is low. Shareholders enjoy the upside if risky bets pay off, but creditors and taxpayers bear the losses if the bank fails.

Third, banks benefit from explicit and implicit government guarantees that reduce their borrowing costs. Deposit insurance represents an explicit guarantee that protects smaller depositors and prevents runs and panics to which banks are inherently vulnerable (Diamond and Dybvig 1983; Anginer and Demirguc-Kunt 2018). "Too big to fail" policies represent an *implicit* guarantee: while governments do not formally promise to rescue large banks, their investors expect systemically important institutions to receive public support during a crisis, because a large bank failure could undermine the financial system. Both types of guarantees make banks appear safer to creditors, allowing them to borrow at lower cost than they would without government backing.

These policy distortions create a classic moral hazard problem. Shielded from downside risk by government guarantees, bank shareholders and managers have strong incentives to maximize leverage and pursue risky strategies. If high-risk bets succeed, shareholders capture the upside, but if the bets fail and the bank collapses, the Federal Deposit Insurance Corporation and taxpayers foot the bill through deposit insurance payouts and bailouts. This asymmetric payoff structure encourages banks to operate with minimal capital while taking on dangerous levels of risk.

Bank capital requirements directly counter these warped incentives. That is, capital requirements limit the moral hazard created by government guarantees. While government-supplied downside protection motivates banks to take more risk, capital requirements limit leverage and constrain risk-taking. Put differently, capital requirements support optimal contracting in the presence of government guarantees (Cooper and Ross 2002).

Bank capital requirements also make the financial system safer by creating a cushion that absorbs losses when asset values decline (as illustrated in Figure 1). Larger capital buffers reduce the risk of bank runs by uninsured depositors, who are prone to panic when a bank appears financially vulnerable. The 2023 failures of Silicon Valley Bank, Signature Bank, and First Republic Bank illustrate this dynamic: as rising interest rates eroded these banks' asset values, uninsured depositors rushed to withdraw their funds, accelerating the banks' collapse (Cecchetti and Schoenholtz 2023; Acharya et al. 2023; for a discussion in this journal, see Metrick 2024).

In addition to making the financial system safer, bank capital requirements also make it more efficient. When bank shareholders have sufficient "skin in the game," they are more likely to insist that their agents running the bank manage risk prudently. As a result, well-capitalized banks are less likely to lend to unprofitable

and highly indebted firms—sometimes called "zombies"—that may be willing and eager to promise high interest payments but can only survive if someone provides them with even more credit (Acharya et al. 2022).

While banks face significant *private* costs when funding operations with capital, the *social* costs are far lower. Critics contend that capital requirements reduce the willingness of banks to lend and make markets (that is, to provide liquidity for securities trading). However, the social benefits of bank capital likely outweigh these potential costs, because well-capitalized banks can maintain credit flows and market liquidity precisely when they are needed most—during economic downturns.

Moreover, the various legal provisions that make debt financing artificially cheap for banks—namely, the tax deductibility of interest payments, the limited liability of shareholders, and the range of public guarantees for deposits and other bank liabilities—do not affect the social cost. By forcing banks to internalize more risk, capital requirements reduce these distortionary subsidies. Put differently, higher capital requirements help align private incentives with those of society (for more discussion, see Admati and Hellwig 2024; Cecchetti and Schoenholtz 2018, 2020).

International Capital Standards: The Implementation Challenge

The early 1970s were a time of considerable disruption in international financial markets. The collapse of the Bretton Woods exchange rate system in 1973 coincided with rising global inflation, creating widespread currency volatility. Foreign exchange losses brought down West Germany's prominent Bankhaus Herstatt in 1974, sending shock waves through the global banking industry. Partly in response, central bank leaders from the Group of Ten countries formed the Basel Committee on Banking Supervision that same year as a forum to promote more collaborative oversight of internationally active banks.[2]

In the 1980s, bankers saw the opportunity to expand their operations across national boundaries as restrictions on cross-border capital flows eased. While this development generally was welcomed by bank customers, it gave rise to two related concerns: the potential for financial instability due to inadequately capitalized, globally active banks; and the then-perceived competitive threat from the rapid growth of banks with relatively lax capital requirements (and hence low financing cost) in jurisdictions like Japan (Tarullo 2008).[3]

---

[2] The Group of Ten—actually composed of eleven high-income countries—met regularly to discuss economic, monetary, and financial matters at the Bank for International Settlements (BIS). The countries are Belgium, Canada, France, Germany, Italy, Japan, the Netherlands, Sweden, Switzerland, the United Kingdom, and the United States. Over the years, the membership of the Basel Committee has increased so that now it includes 27 countries plus the European Union (for membership, see https://www.bis.org/bcbs/membership.htm).

[3] The perceived advantage of Japanese banks ended abruptly in the early 1990s when Japan's asset-price collapse highlighted the catastrophic risks of inadequate capitalization (Hoshi and Kashyap 1999).

These concerns led to a movement to create an international regime of minimum regulatory standards that would promote financial stability within and across countries, while maintaining a competitive balance among banks operating globally. Officials turned to the Basel Committee to develop such standards, and the result was the 1988 Basel Accord. Although the Accord was merely a nonbinding recommendation that each jurisdiction could choose to enact, it represented an unprecedented effort at international regulatory cooperation.

The original Basel Accord revealed the challenges of designing bank capital standards. The 1988 framework required internationally active banks to maintain capital equal to at least 8 percent of their "risk-weighted assets"—a measure that adjusts asset values based on their perceived riskiness. Under this system, banks' assets were sorted into four risk categories: (1) sovereign debt issued by industrialized countries (assigned a 0 percent risk weight); (2) claims on banks of industrialized countries (20 percent weight); (3) residential mortgages (50 percent weight); and (4) consumer and corporate loans (100 percent weight).

Banks calculated their required capital by multiplying each asset's value by its risk weight. For example, because residential mortgages had a risk weight of 50 percent, a bank holding a $100,000 mortgage would require (8 percent × 50 percent × $100,000) = $4,000 of equity funding, and the remainder of its financing for the loan ($96,000) could come from a combination of deposits and long-term debt. A bank holding 100 percent of its assets in sovereign debt would not need to have any capital at all.

The 1988 agreement had serious practical limitations. We highlight two of the most significant. First, it treated all sovereign bonds issued by governments of OECD countries equally, regardless of issuer or maturity. That is, it assigned a ten-year US Treasury bond, a three-month US Treasury bill, and a five-year Turkish government bond the same risk weight: zero. Second, a corporate bond received a 100 percent risk weight, regardless of whether it was a highly-rated (low-yield) AAA bond or a lowly-rated (high-yield) "junk" bond (BB+ or below). The same problem arose in using a single risk weight for all commercial loans. These shortcomings facilitated regulatory arbitrage: banks could shift their investments toward assets with inappropriately low risk weights, allowing them to own riskier assets (with higher yields) without increasing their required level of capital funding.

Starting in 1998 the Basel Committee began negotiating a revised framework. In 2004, the Committee published the Basel II standards. The new standards did not address the failings associated with sovereign debt in the initial Accord. However, they made substantial changes in the method used to compute banks' risk-weighted assets. Banks could use an updated "standardized approach" where the risk weights would be sensitive to borrowers' external credit ratings. In addition, under Basel II, large banks were given the option of using their own internal model estimates of default to calculate risk weights for their assets.[4]

---

[4] Haubrich (2020) provides a brief history of capital requirements in the United States.

In addition to specifying capital requirements for *credit risk*—the risk that a borrower might not repay a loan—the Basel II framework also captured another type of risk that was not in the original Basel I standards. Specifically, Basel II established capital requirements for *operational risk*—the risk of loss arising from (1) inadequate or failed internal processes, people, and systems, or (2) external events. However, Basel II also allowed banks to include as capital various liabilities that would be less effective than common equity for absorbing losses in a crisis (Cecchetti and Schoenholtz 2014).

In November 2007, the US banking agencies adopted a rule implementing Basel II that was to be introduced gradually starting in April 2008. That is, implementation of Basel II in the United States started between the failure of Bear Stearns in March 2008 and the failure of Lehman Brothers six months later. Thus, when the global financial crisis hit, Basel II was not yet in place in the United States. This US timing was not unusual: when the crisis hit, Basel II implementation was still underway in much of the world (Yetis 2008).

The global financial crisis in 2008–2009 exposed fundamental weaknesses in bank regulation that would have persisted even if Basel II had been fully implemented. Most importantly, many large banks lacked sufficient equity capital to absorb the losses they incurred on their risky assets. As capital cushions eroded, depositors and creditors lost confidence, triggering runs on the most vulnerable institutions, widespread insolvencies, and government bailouts. (This pattern did not stop with banks, as governments around the world also bailed out numerous large nonbank financial institutions that failed or were severely weakened.)

To address many of the shortcomings exposed by the global financial crisis, in 2010 the Basel Committee issued the initial Basel III agreement. The new agreement emphasized a specific type of capital called "common equity tier 1," composed primarily of retained earnings and funds raised by issuing common equity. This stricter definition of capital excluded liabilities that would be less able than common equity to absorb losses in a crisis. In addition, Basel III raised the risk weights applied to certain asset classes.

This initial Basel III agreement also established new minimum capital requirements that require recapitalization when a bank's capital levels fall below certain thresholds. In effect, the common equity tier 1 requirement for internationally active banks under Basel III is 7 percent of risk-weighted assets. When a bank's capital level falls below this level, the standards state that the bank must limit dividend payments, stock buybacks, and management bonuses until its capital is rebuilt. Additional risk-based capital buffers in excess of the 7 percent minimum are applied to banks deemed systemically important.

As a backstop and complement to the risk-based capital ratios that are prone to gaming, Basel III introduced a "leverage" capital requirement: a minimum ratio of capital divided by total assets. (As a technical matter, total assets in this calculation include off-balance-sheet exposures such as measures associated with derivatives exposures.) The leverage ratio looks at total assets, without risk-weighting, and thus does not depend on the limited ability of banks and regulators to anticipate asset

risk (Klein 2016). The leverage ratio also is less likely to be distorted by regulatory arbitrage. As a result, it is useful to assess the evolving riskiness of the banking system as a whole.

Policymakers never intended the initial 2010 version of Basel III to be the final word. Instead, to facilitate a quick initial agreement and a timely response to the global financial crisis, authorities postponed consideration of numerous issues. This approach led to an iterative process consistent with the Basel Committee's longstanding approach. First, standards are agreed and implemented. Then, as financial systems evolve and the effects of the standards are assessed, revisions are proposed, and the process starts again.

In December 2017, the Basel Committee issued nearly comprehensive Basel III capital standards, followed by a January 2019 enhancement focusing on market risk (the risk associated with banks' trading activities). These reforms were intended to "restore credibility in the calculation of risk-weighted assets . . . and improve the international comparability of banks' capital ratios."[5] The most important refinements limit the use of internal models in the calculation of large banks' capital ratios and improve the methodologies for assessing credit risk, operational risk, and market risk.

Partial Implementation of Basel III in the United States

US bank regulators implemented the initial Basel III reforms in 2013. In some cases, the US banking agencies "gold-plated" the international standards, requiring capital buffers that exceeded the Basel Committee minima. With memories of the global financial crisis still fresh, these policy changes were relatively noncontroversial.

As the United States phased in the initial Basel III reforms over the ensuing years, capital levels rose steadily. Figure 2 illustrates capital ratios for two samples of large banks. The black line shows the common equity tier 1 risk-weighted capital ratio for the twelve large banks in the Americas that are part of the Basel Committee's annual implementation monitoring. The red line shows the leverage ratio for the eight largest US banks—designated as "global systemically important banks." The shaded areas denote the recessions associated with the global financial crisis and the Covid-19 pandemic.

From 2011 to 2017—which includes the initial Basel III period of implementation—the risk-weighted common equity tier 1 capital ratio rose from 7 to 12 percent. The leverage ratio for US banks rose from roughly 6 percent in 2008 to 9 percent in 2016. But this upward trend did not persist. By 2021, the leverage ratio at the global systemically important banks had sunk back to post-crisis lows, resulting in a wide gap between the risk-weighted and leverage measures. This growing gap may partly reflect the fact that as banks' total assets rose, they increased holdings of reserves

---

[5] A fuller description is available at the Bank of International Settlements website at https://www.bis.org/fsi/fsisummaries/rcrf.htm.

Figure 2
**Capital Ratios for Large Banks (Quarterly or Semi-annually),
June 2002–December 2024**
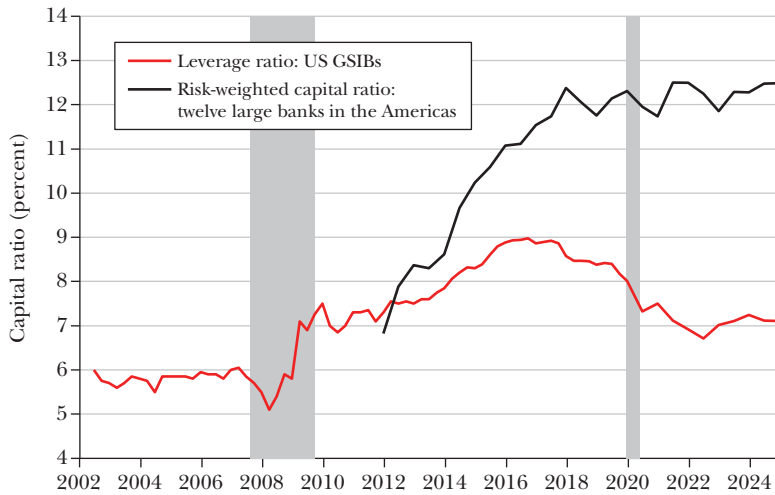
Source: The source for the Leverage Ratio is Pellerin (2025, Chart 1). The source for the Risk-weighted capital ratio is Graph 17 in the Basel III Monitoring Report (Basel Committee on Banking Supervision 2024a). Banks' risk-weighted capital ratios are not available prior to 2012 because banks were not required to calculate their capital under Basel III standards before then. This sample of twelve banks likely includes a few banks from Canada and Brazil. However, it is the only consistent measure of risk-weighted capital ratios available over this full period.
Note: "Leverage Ratio: US GSIBs" is the Tier 1 leverage ratio for the eight US banks designated as global systemically important banks. The "Risk-weighted capital ratio: twelve large banks in the Americas" is the common equity tier 1 ratio for the consistent sample of twelve Category 1 banks in the Americas. Category 1 banks include, but are not limited to, the global systemically important banks. Shaded areas denote NBER-dated recessions. The fact that the two lines cross is a consequence of differences in the samples of banks.

and other zero-risk-weight assets. Nevertheless, a portion of this reversal and of the divergence between the two measures should be attributed to banks' efforts to offset the impact of the initial Basel III reforms through regulatory arbitrage and the willingness of US bank supervisors after 2016 to accept this behavior.

The Basel Endgame Proposal: Deviations from the Basel III Framework

In early 2023, the extraordinary failures of Silicon Valley Bank, Signature Bank, and First Republic Bank in the United States, as well as Credit Suisse in Switzerland, necessitated emergency government interventions and refocused policymakers' attention on financial reform (Metrick 2024). Against this backdrop, the US banking agencies—the Federal Reserve, the Federal Deposit Insurance Corporation, and

the Office of the Comptroller of the Currency— proposed the "Basel Endgame" in July 2023.

The Basel Endgame proposal mostly tracked the 2017 updated Basel III standards, including restrictions on banks' use of internal models to calculate their capital ratios and improvements to standardized risk-assessment methods. However, the US proposal diverged from international standards in several important ways. We highlight three significant differences.

First, the Basel Endgame would have applied enhanced capital requirements to more US banks than prior iterations of Basel standards. By their terms, Basel standards target "internationally active banks."[6] Before 2023, the United States applied many of the key provisions of the initial Basel III rules only to the eight US global systemically important banks: namely, Bank of America, Bank of New York Mellon, Citigroup, Goldman Sachs, JP Morgan Chase, Morgan Stanley, State Street, and Wells Fargo, as well as to one bank that serves as a secure holder (custodian) of securities (Northern Trust). The Basel Endgame proposal would have expanded the application of certain Basel III provisions to all US banks with more than $100 billion in assets—a total of 37 firms. This proposed expansion would have subjected regional banks similar in scale to Silicon Valley Bank and First Republic Bank to some of the same capital standards as the very largest US banks like JP Morgan Chase and Bank of America.

In a second major difference, the Basel Endgame proposal sought material increases in capital requirements for US banks. When the Basel Committee published its final 2017 standards, it asserted that the revisions would *not* "significantly increase overall capital requirements" (Basel Committee on Banking Supervision 2017, p. 1). However, according to estimates by the US regulatory agencies, Basel Endgame would have raised the common equity tier 1 capital requirements for covered banking organizations by an average of 16 percent (Office of the Comptroller of the Currency et al. 2023, page 64169).[7] By contrast, the European Banking Authority (2023) originally projected that its finalization of Basel III would raise capital requirements on its banks by 9 percent, while the Bank of England (2023) anticipated just a 3 percent increase.

Much of the proposed capital increase in the Basel Endgame was attributable to a third major divergence from the Basel III standard: the calculation of capital required to provide a buffer against losses associated with operational risk (as defined earlier). Before the final Basel III reforms, both the Basel Committee and US policymakers allowed banks to determine what is in effect an addition to risk-weighted assets using their own estimates of future operational risk losses. In

---

[6] For the list of all 29 global systemically important banks, see Financial Stability Board (2024).
[7] With the median capital requirement for large banks running about 8 percent of risk-weighted assets, this would have meant an increase of roughly 1.25 percentage points. This is the median of the common equity tier 1 capital ratio requirement for the 32 largest US banks in 2024 (Board of Governors of the Federal Reserve 2024, p. 4).

practice, this methodology proved complex and generated highly variable estimates across similar banks.

The final Basel III reforms offered a standardized formula in which a bank's operational risk is a function of (1) the bank's revenues from different business lines, and (2) an "internal loss multiplier" based on the bank's operational risk losses over the previous ten years. Under this framework, a bank's internal loss multiplier could range from less than one (for banks with low historical operational losses) to greater than one (for banks with high historical operational losses). By contrast, the Basel Endgame would have set a floor of one on the internal loss multiplier and thereby effectively applied the standard for high historical losses to all banks. According to a report from the Bank Policy Institute (a trade association representing the largest US banks), "the new operational risk charge accounts for nearly 90 percent of the increase in banks' capital requirements" under the Basel Endgame proposal (Covas 2023).

The Basel Endgame Proposal: Key Issues

The Basel Endgame proposal provoked intense public debate. Supporters and opponents clashed over a host of issues: the appropriate risk weights for various asset classes, the calibration of market-risk formulas, the proposed approach to operational risk capital, and others. At its core, however, the controversy centered on the optimal level of bank capital requirements, the effects of bank capital requirements on lending and market-making (that is, the perceived social costs of bank capital), and the structure of the global financial system. In this section, we summarize the biggest points of contention.

First, numerous commenters objected to the Basel Endgame's proposed risk weights. The agencies' treatment of residential mortgage loans generated especially vociferous pushback. Under the US implementation of the initial Basel III rules, all residential mortgage loans that meet prudent underwriting criteria receive a 50 percent risk weight. The final 2017 Basel III reforms introduced differentiated risk weights ranging from 20 percent to 70 percent, depending on a mortgage's loan-to-value ratio. The Basel Endgame proposal would have incorporated differentiated risk-weights for residential mortgage loans that were 20 percentage points *higher* than the Basel framework for each loan-to-value category.

US financial regulators stated that the rationale for the proposed increase over the Basel standard was that, without it, larger US banks would benefit from lower risk weights than smaller community banks not subject to the same capital rules. Large banks and their allies vigorously opposed the rule, arguing that it would make mortgage lending uneconomical, cede market share to nonbank mortgage lenders, and reduce access to affordable housing.

Second, under the US implementation of the initial Basel III standards, large banks calculated their market-risk requirements using a "value-at-risk" methodology, which measures the worst possible loss over a specific time horizon for a given

probability. (Recall that the market-risk framework refers to the risk of losses associated with banks' trading activity.) Consistent with the final Basel III agreement, the Basel Endgame proposal would have shifted from value-at-risk methodologies to an "expected shortfall" approach that reflects a bank's expected or average loss during an extreme downturn. Regulators reasoned that the expected shortfall approach would better reflect banks' risks in the worst possible circumstances. The regulatory agencies estimated that shifting from value-at-risk to the expected shortfall methodology, along with other market-risk changes, would increase required capital ratios by 67 basis points for affected banks (Office of the Comptroller of the Currency et al. 2023, p. 64170).

Banks raised two main objections to these changes in the treatment of market risk. First, they argued that the new requirements would duplicate standards already built into in the Federal Reserve's annual "stress tests." Second, the industry asserted that simply adding up risks across individual trading desks would ignore the risk-reducing benefit of diversification. However, while diversification provides some protection during normal market conditions, this benefit can disappear during financial crises when prices of risky assets tend to fall in unison.

These important details clearly matter, but the Basel Endgame debate also focused on more fundamental questions about the optimal levels of bank capital and its effect on the macroeconomy. Critics maintained that US capital requirements are already high, so raising them even modestly would constrain the supply of credit and restrain economic growth (for example, Bank Policy Institute Staff 2023). Critics further argued that the costs of increased capital requirements would be borne primarily by households and businesses that would have less access to—and would have to pay more for—loans and other banking services. Finally, banks contended that raising capital requirements is unnecessary because capital levels are already within the range that correctly balances trade-offs between promoting economic growth, on the one hand, and making banks safer, on the other (Covas and Nelson 2023).

By contrast, Basel Endgame supporters asserted that the increase in bank capital requirements was relatively modest and would enhance bank safety in a manner that was unlikely to constrain bank lending or market-making significantly. Indeed, banks had already adapted to much larger capital increases over the past decade without major disruptions to credit availability. However, the limited nature of the changes cuts both ways. A moderate increase in capital requirements could produce only moderate improvements in banking system safety—not the transformational change that some advocates might wish.

Admittedly, there is no theoretical consensus on *optimal* bank capital levels, and studies on the economic effects of higher capital requirements are sensitive to assumptions and modeling techniques. However, the weight of the evidence suggests that setting large-bank capital requirements in the range contemplated by the Basel Endgame proposals would have little if any impact on the supply of credit over the long-term (for discussion of credit, see Gambacorta and Shin 2018; for assessment of the optimal capital range, see Birn et al. 2020). Better-capitalized banks tend to

Figure 3

United States and Europe: Large Banks' Annualized Return on Equity, June 2015–December 2024 (Quarterly Data)
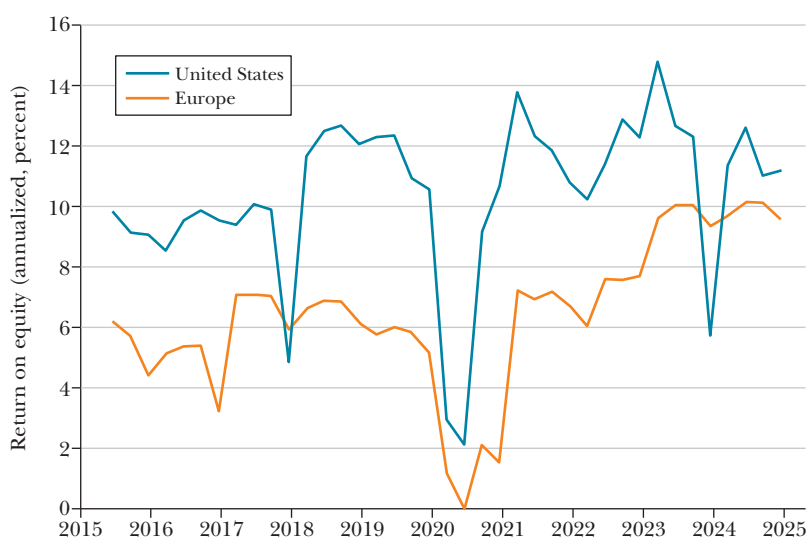
Source: FDIC Quarterly Banking Profile, European Central Bank, and authors' calculations.
Note: US banks include those 158 institutions with assets exceeding $10 billion. The European sample includes the 114 euro-area institutions supervised by the Single Supervisory Mechanism of the European Central Bank. In both cases, the banks in the sample hold more than 80 percent of banking assets in their jurisdiction. The only two quarters when US banks underperformed reflected temporary US government-driven distortions: a tax change in 2017 and an FDIC fee in 2023 to restore its crisis-depleted insurance fund.

(1) act countercyclically, lending more during economic and financial stress when other sources of financing dry up (Basel Committee on Banking Supervision 2019, p. 10), and (2) lend to healthier borrowers, favoring more efficient economywide use of resources (Caballero, Hoshi, and Kashyap 2008).

Even more broadly, the Basel Endgame debate concerned the role of US banks in the global financial system. Opponents argued that if the United States were to implement stronger capital requirements than other jurisdictions, financial activity would migrate abroad (Kroszner 2024). In response, supporters countered that higher capital requirements can be a competitive advantage, because robust capitalization makes the domestic financial system strong and resilient.

The latter view is supported by the fact that *after* the United States "gold-plated" the initial 2010 Basel III standards, the relatively well-capitalized US banks outperformed their European competitors (Simoens and Vander Vennet 2021). As Figure 3 highlights, large US banks' return on equity has significantly and persistently exceeded that of large European banks since 2015 (the period for which pan-European data are available).

The Basel Endgame debate also exposed diverging views about the provision of financial services by banks vis-à-vis nonbank financial institutions. Critics alleged that increasing large-bank capital requirements would shift financial activity (including market making) to nonbanks such as broker-dealers, hedge funds, private credit funds, and nonbank mortgage lenders. In their view, migration of financial activity away from the banking sector weakens the financial system because nonbanks are subject to less stringent (if any) prudential regulation (Kroszner 2024).

On the other hand, Basel Endgame supporters highlighted evidence that counters the predicted shift to nonbanks: as the United States implemented the large initial Basel III increase in capital requirements, from 2011 to 2016 banks *increased* their share of credit provision to the nonbank sector (Cecchetti and Schoenholtz 2020). They also argued that well-capitalized intermediaries buttress financial market resilience, and vice versa (Cecchetti and Schoenholtz 2015). Furthermore, supporters insisted that even if higher capital requirements did push some activity toward nonbanks, this shift could be desirable if risks shift from global systemically important banks to nonbanks that are not as systemically important. Finally, if nonbank risks did increase as a result, the appropriate policy response would be better oversight of nonbanks, not weakening bank capital standards.

The Path Forward

The 2023 Basel Endgame proposal faced an uphill battle from the outset as it elicited rare dissenting votes from members of the Federal Reserve and boards of directors of the Federal Deposit Insurance Corporation boards of directors (Bowman 2023; Waller 2023; Hill 2023; McKernan 2023). Even regulators who formally supported issuing the proposal for public comment, including Federal Reserve Chairman Jerome Powell (2023), expressed misgivings that foreshadowed a rocky path ahead. As the public comment period progressed, the tenor of the criticism made it clear that the proposal would not survive as originally conceived.

By mid-2024, Federal Reserve Chair Powell (2024) and Vice Chair Barr (2024) promised Congress that the agencies would make "broad and material changes," essentially starting over and re-proposing the Basel Endgame rules before finalization. The likely revisions would include three major concessions: (1) exempting banks with less than $250 billion in assets from most Basel Endgame reforms; (2) reducing proposed risk weights for residential mortgage loans and other asset classes; and (3) lowering operational risk capital requirements for banks with low prior operational losses. According to Barr, these changes would cut the proposed capital increases by more than half compared to the original 2023 proposal.[8]

---

[8] The agencies' originally estimated that the July 2023 Basel Endgame proposal would increase capital requirements for Category 1 and 2 banks, which include global systemically important banks and other banks with more than $700 billion in assets or international activity of more than $75 billion, by 19 percent (Office of the Comptroller of the Currency et al. 2023, p. 64169).

But at this writing in mid-2025, the US bank regulatory agencies have not yet issued a revised Basel Endgame proposal, leaving US banking rules out of compliance with international Basel III standards. Policymakers now face two decisions: Should the United States adopt capital rules that comply with Basel III? And should regulators raise capital requirements on large banks? While the case for an affirmative answer to both questions is strong, the issues are logically separable. If implementing Basel III standards would increase US capital requirements, regulators could make offsetting adjustments elsewhere in the capital framework to maintain current levels—achieving international compliance without raising overall capital requirements.

The Case for US Compliance with the Final Basel III Standards
A capital-neutral implementation of the Basel Endgame proposal would allow the United States to comply with international standards while minimizing costs to banks and preserving the benefits of global regulatory cooperation. The advantages of implementing the final Basel III standards in this way almost certainly outweigh the costs.

Although the agreed-upon Basel standards lack a formal enforcement mechanism, policymakers can employ a "name and shame" process to pressure noncompliant jurisdictions. In addition, individual jurisdictions can take punitive actions against internationally active banks from countries that do not comply with Basel standards. Indeed, the Basel framework contemplates that jurisdictions could opt to impose higher risk weights on exposures to banks whose home regulators are not Basel III–compliant.[9] Some European lawmakers have already suggested that the European Union should take steps to punish American banks if the United States does not implement the 2017 Basel reforms (Wilkes 2025). Such punitive treatment could significantly impair the ability of US banks to compete abroad. Finally, a failure by the United States to implement the 2017 Basel reforms could jeopardize the United States' ability to influence global standards in the future.

If US authorities ultimately choose not to comply with the Basel framework, then foreign jurisdictions will have far less incentive to achieve or maintain compliance. To be sure, several jurisdictions—including Canada, Hong Kong, Japan, Singapore, and Switzerland—have already achieved full compliance with the Basel III capital standards. And China meets Basel capital requirements with the exception of certain margin rules.[10] However, others are now wavering. The United Kingdom

---

[9] Under the final Basel III standards, exposures to healthy banks generally receive risk weights between 20 and 50 percent. However, to qualify for this favorable treatment, internationally active banks must be subject to "appropriate prudential standards (e.g. capital and liquidity requirements) and level of supervision . . . *in accordance with the Basel framework*" (Basel Committee on Banking Supervision 2017, p. 7; authors' emphasis).

[10] Geopolitical concerns may also be relevant for US policymakers. Despite its extensive capital controls, it is conceivable that—if the United States were to renege on its commitment to implement the final Basel III standards—China could seek to supplant the United States as the leader in international financial standard setting.

has yet to implement its final rules on credit, market, and operational risk, while the European Central Bank and the Bank of England have delayed their Basel III implementation, citing US inaction (Bank of England 2025; Jones 2024; Canepa, Strupczewski, and Fonte 2025). The potential unraveling of Basel standards could generate a regulatory race-to-the-bottom, increasing the risk of future financial crises. As Zaring (2019) argues, the resulting failure of multilateral standard-setting could put global banking stability at risk.

The Case for Higher US Capital Requirements for Large Banks

Implementing the final Basel III standards in a capital-neutral way would ensure the continued US compliance with international norms, but it would side-step another question: Are current US bank capital levels sufficient? The academic literature on the optimal level of bank capital requirements offers far too wide a range of answers to settle this issue: various studies have estimated the optimal capital ratio in the range of 6 percent to 25 percent of risk-weighted assets, and 5 percent to 19 percent of total assets.[11] Nevertheless, there are compelling reasons to believe that if US policymakers choose to fulfill the Basel III requirements in a capital-neutral way, capital levels in the largest banks would still be too low.[12] We list five.

First, the average leverage ratio of US global systemically important banks declined from a 2016 peak of 9 percent to about 7 percent in 2023 (see Figure 2). This is roughly in line with the level prevailing prior to the implementation of the more stringent regulations embodied in the Wall Street Reform and Consumer Protection Act of 2010, commonly known as the Dodd-Frank Act (Dugan and Nonaka 2017). To be sure, banks increased holdings of assets devoid of default risk when the Federal Reserve rapidly expanded reserves during the Covid-19 pandemic. However, much of the leverage ratio decline predated that episode. Against this background, increasing capital requirements would help reverse the backsliding in large banks' capital standards.

Second, as we noted earlier, there is little evidence that the rise of US global systemically important banks' capital ratios from 2009 to 2016 had a negative impact on bank credit supply, either overall or relative to nonbanks. Even after 2016, the fraction of US total credit to the private nonfinancial sector provided by banks was roughly stable at about one-third until the start of the pandemic (Cecchetti and

---

[11] Using a general equilibrium model, Elenev, Landvoigt, and Van Nieuwerburgh (2021) estimate that a risk-weighted capital requirement of 6 percent would maximize welfare, while Firestone, Lorenc, and Ranish (2017) estimate an optimal Tier 1 risk-weighted capital ratio as high as 26 percent. Using Swedish data, Almenberg et al. (2017) estimate the optimal leverage ratio as low as 5 percent, while Barth and Miller (2018) put the optimal leverage ratio for US banks at 19 percent.

[12] These arguments would be even more compelling if, as recently proposed, US regulators act to lower the leverage requirement on the largest banks (Office of the Comptroller of the Currency et al. 2025).

Figure 4

Global and Systemically Important Banks in Europe, Japan, and the United States: Supplementary Leverage Ratio versus Price-to-Book Ratio, June 30, 2024
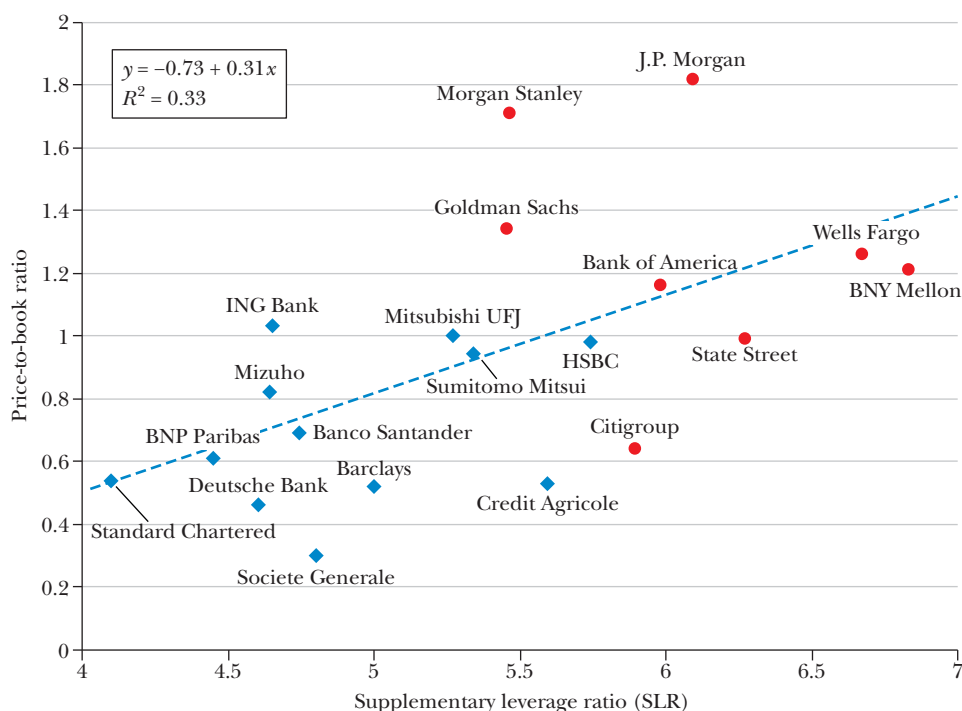
Source: Pellerin (2025, Table 1).
Note: US banks are red circles; foreign banks are blue diamonds.

Schoenholtz 2020).[13] Similarly, there is no indication that higher capital require-ments have a negative impact on the profitability of large US banks (as illustrated earlier in Figure 3).

   Third, looking at a cross-section of global systemically important banks, the leverage ratio is *positively* correlated with the price-to-book-value ratio, as shown in Figure 4. Given that US banks generally have both higher leverage and higher price-to-book ratios, this correlation may partly reflect differences between US and non-US banks (as well as between US and non-US stock markets). However, it is reasonable to think of the leverage ratio as a choice variable for bank manage-ment. Analyzed in that way, the relationship in Figure 4 suggests that boosting

[13] To be sure, the ratio of private nonfinancial sector credit to GDP fell from its high above 170 percent in 2008 to 150 percent in 2015. However, the pre-crisis level circa 2008 was clearly unsustainable. The latest observation for the third quarter of 2024 is 145.2 percent. For this data, see the BIS Data Portal at https://data.bis.org/topics/TOTAL_CREDIT/BIS,WS_TC,2.0/Q.US.P.A.M.770.A).

equity capitalization—at least by retaining earnings—raises market capitalization more than one for one. If that is the case, it raises the question why banks do not voluntarily boost their capitalization in the absence of higher requirements. One reason may be that equity investors are prone to distrust bank managers who opt to retain earnings, even if the investors reward banks when higher capital requirements substitute for managerial choice.

Fourth, more broadly, when banks are short of capital, their clients tend to suffer. Banks with insufficient capital buffers may do some combination of three undesirable things: (1) they could withdraw from lending and market-making activities, increasing their capital ratio by shrinking their balance sheet rather than issuing more equity to investors; (2) they could extend the maturity of existing loans to low-quality borrowers ("zombie firms") to avoid recognizing losses on their capital-scarce balance sheets (for example, Acharya et al. 2022); and (3) they may "gamble for resurrection" by taking greater risks across a range of activities, hoping to boost returns and capital buffers in this way. Conversely, healthy banks have a greater capacity to lend and to make markets in recessions, as well as to acknowledge nonperforming loans in a timely way. In troubled times, the robust reputation of well-capitalized banks makes them especially attractive to clients. Their skin in the game also generates an incentive to lend to high-quality borrowers. Overall, banks with high levels of capital funding support economic stability, acting as a source of funds during downturns and episodes of financial turmoil.

Fifth, and closely related to the previous point, there is a link between capital shortfalls in the banking system and extended periods of economic stagnation. The classic example is Japan in the 1990s. It was not until the government of Japan moved to recapitalize banks in the early 2000s that economic growth began to recover (for discussion in this journal, see Hoshi and Kashyap 2004). Conversely, we know of no links between episodes of stable, robust bank capitalization and periods of economic stagnation.

Taken together, these considerations argue strongly for restoring US large bank capital ratios to at least the 2016 level. Policymakers wishing to generate momentum in that direction would be well-advised to conduct and publish a comprehensive quantitative cost-benefit analysis of raising capital requirements. Granted, it can be difficult to quantify the costs and benefits of financial regulations (for discussion, see Coates 2015; Gordon 2014), and there is no legal requirement to do so. However, a careful and independent quantitative impact study can illuminate the economic and financial trade-offs associated with changes in bank capital requirements. We expect that such a study would foster public discussion and serve as a guide (and a goal) for policy action.

References

Acharya, Viral V., Matteo Crosignani, Tim Eisert, and Sascha Steffen. 2022. "Zombie Lending: Theoretical, International, and Historical Perspectives." *Annual Review of Financial Economics* 14: 21–38.

Acharya, Viral V., Matthew P. Richardson, Kermit L. Schoenholtz, and Bruce Tuckman, eds. 2023. *SVB and Beyond: The Banking Stress of 2023*. Volatility and Risk Institute, New York University Stern School of Business.

Admati, Anat, and Martin Hellwig. 2024. *The Bankers' New Clothes: What's Wrong with Banking and What to Do about It—New and Expanded Edition*. Princeton University Press.

Almenberg, Johan, Markus Andersson, Daniel Buncic, Cristina Cella, Paolo Giordani, Anna Grodecka, Kasper Roszbach, and Gabriel Söderberg. 2017. "Appropriate Capital Ratios in Major Swedish Banks—New Perspectives." Riksbank Staff Memo, May. https://www.riksbank.se/globalassets/media/rapporter/staff-memo/engelska/2017/staff_memo_170519_eng.pdf.

Anginer, Deniz and Asli Demirguc-Kunt. 2018. "Bank Runs and Moral Hazard: A Review of Deposit Insurance." World Bank Policy Research Working Paper 8589.

Bank of England. 2023. "Implementation of the Basel 3.1 Standards Near-Final Part 1." December 12. https://www.bankofengland.co.uk/prudential-regulation/publication/2023/december/implementation-of-the-basel-3-1-standards-near-final-policy-statement-part-1.

Bank of England. 2025. "The PRA Announces a Delay to the Implementation of Basel 3.1." News Release, January 17. https://www.bankofengland.co.uk/news/2025/january/the-pra-announces-a-delay-to-the-implementation-of-basel-3-1.

Bank Policy Institute Staff. 2023. "Basel Endgame: Background and Key Issues." Bank Policy Institute, September 5. https://bpi.com/basel-endgame-background-and-key-issues/.

Barr, Michael S. 2022. "Why Capital Matters." Speech, American Enterprise Institute, Washington, DC, December 1, 2022. https://www.bis.org/review/r221202b.pdf.

Barr, Michael S. 2024. "The Next Steps on Capital." Speech, Brookings Institution, Washington, DC, September 10. https://www.bis.org/review/r240910b.htm.

Barth, James R., and Stephen Matteo Miller. 2018. "Benefits and Costs of a Higher Bank 'Leverage Ratio.'" *Journal of Financial Stability* 38: 37–52.

Basel Committee on Banking Supervision. 2010. *Results of the Comprehensive Quantitative Impact Study*. Bank for International Settlements.

Basel Committee on Banking Supervision. 2017. *Basel III: Finalising Post-crisis Reforms*. Bank for International Settlements.

Basel Committee on Banking Supervision. 2019. "The Costs and Benefits of Bank Capital—A Review of the Literature." Bank for International Settlements Working Paper 37.

Basel Committee on Banking Supervision. 2024b. *The Basel Framework*. Bank for International Settlements.

Basel Committee on Banking Supervision. 2025. "RCAP on Timeliness: Basel III Implementation Dashboard." Bank for International Settlements, May 16. https://www.bis.org/bcbs/implementation/rcap_reports.htm.

Birn, Martin, Olivier de Bandt, Simon Firestone, Matías Gutiérrez Girault, Diana Hancock, Tord Krogh, Hitoshi Mio, Donald P. Morgan, Ajay Palvia, Valeria Scalone, Michael Straughan, Arzu Uluc, Alexander H. von Hafften, and Missaka Warusawitharana. 2020. "The Costs and Benefits of Bank Capital—A Review of the Literature." *Journal of Risk and Financial Management* 13 (4): 74.

Board of Governors of the Federal Reserve System. 2024. "Large Bank Capital Requirements." https://www.federalreserve.gov/publications/files/large-bank-capital-requirements-20240828.pdf.

Bowman, Michelle W. 2023. "Statement by Governor Michelle W. Bowman." Press release, Board of Governors of the Federal Reserve System, July 27. https://www.federalreserve.gov/newsevents/pressreleases/bowman-statement-20230727.htm.

Caballero, Ricardo J., Takeo Hoshi and Anil K. Kashyap. 2008. "Zombie Lending and Depressed Restructuring in Japan." *American Economic Review* 98 (5): 1943–77.

Canepa, Francesco, Jan Strupczewski, and Giuseppe Fonte. 2025. "Exclusive: EU to Delay Bank Rules as It Waits for Trump's Deregulation Moves, Sources Say." *Reuters*, May 22.

Cecchetti, Stephen G., and Kermit L. Schoenholtz. 2014. "Making Finance Safe." Money and Banking, October 6. https://www.moneyandbanking.com/commentary/2014/10/6/making-finance-safe.

Cecchetti, Stephen G., and Kermit L. Schoenholtz. 2015. "Bond Market Liquidity: Should We Be Worried?"

Money and Banking, August 17. https://www.moneyandbanking.com/commentary/2015/8/17/bond-market-liquidity-should-we-be-worried.

Cecchetti, Stephen G., and Kermit L. Schoenholtz. 2018. "Understanding Bank Capital: A Primer." Money and Banking, February 12. https://www.moneyandbanking.com/primers/2018/2/11/understanding-bank-capital-a-primer.

Cecchetti, Stephen G., and Kermit L. Schoenholtz. 2020. "Settling Bank Capital Requirements." Money and Banking, October 12. https://www.moneyandbanking.com/commentary/2020/10/11/setting-bank-capital-requirements.

Cecchetti, Stephen G., and Kermit L. Schoenholtz. 2023. "The Extraordinary Failures Exposed by Silicon Valley Bank's Collapse." Money and Banking, March 20. https://www.moneyandbanking.com/commentary/2023/3/20/the-extraordinary-failures-exposed-by-silicon-valley-banks-collapse.

Cooper, Russell, and Thomas W. Ross. 2002. "Bank Runs: Deposit Insurance and Capital Requirements." *International Economic Review* 43 (1): 55–72.

Coates, John C., IV. 2015. "Cost-Benefit Analysis of Financial Regulation: Case Studies and Implications." *Yale Law Journal* 124 (4): 882–1051.

Covas, Francisco. 2023. "About Excessive Calibration of Capital Requirements for Operational Risk." Bank Policy Institute, October 30. https://bpi.com/about-excessive-calibration-of-capital-requirements-for-operational-risk/.

Covas, Francisco, and Bill Nelson. 2023. "US Bank Capital Levels: Aligning with or Exceeding Midpoint Estimates of Optimal." Bank Policy Institute, September 18. https://bpi.com/u-s-bank-capital-levels-aligning-with-or-exceeding-midpoint-estimates-of-optimal/.

Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (3): 401–19.

Dugan, John C., and Michael Nonaka. 2017. "President Trump Begins Efforts to Roll Back Financial Regulations." Harvard Law School Forum on Corporate Governance, February 6.

Elenev, Vadim, Tim Landvoigt, and Stijn Van Nieuwerburgh. 2021. "A Macroeconomic Model with Financially Constrained Producers and Intermediaries." *Econometrica* 89 (3): 1361–418.

Ernst and Young. 2024. "House Financial Services Committee Questions Federal Reserve Chair Powell on Basel III, Stress Tests, Liquidity Rules." July 11. https://taxnews.ey.com/news/2024-1354-house-financial-services-committee-questions-federal-reserve-chair-powell-on-basel-iii-stress-tests-liquidity-rules.

European Banking Authority. 2023. *Basel III Monitoring Exercise Results based on Data as of 31 December 2022.*

Ferguson, Keegan. 2024. "Chopra's Risky Basel Endgame Gamble." Capstone, October 14. https://capstonedc.com/insights/chopras-risky-basel-endgame-gamble/.

Firestone, Simon, Amy Lorenc, and Ben Ranish. 2019. "An Empirical Economic Assessment of the Costs and Benefits of Bank Capital in the United States." *Federal Reserve Bank of St. Louis Review* 101 (3): 203–30.

Financial Stability Board. 2024. "2024 List of Global Systemically Important Banks (G-SIBs)." November 26. https://www.fsb.org/2024/11⁄2024-list-of-global-systemically-important-banks-g-sibs/.

Gambacorta, Leonardo, and Hyun Song Shin. 2018. "Why Bank Capital Matters for Monetary Policy." *Journal of FInancial Intermediation* 35 (B): 17–29.

Gordon, Jeffrey N. 2014. "The Empty Call for Benefit-Cost Analysis in Financial Regulation." *Journal of Legal Studies* 43 (S2): S351–78.

Haubrich, Joseph G. 2020. "A Brief History of Bank Capital Requirements in the United States." Federal Reserve Bank of Cleveland Economic Commentary 2020-05.

Heeb, Gina. 2024. "Trump Advisers Seek to Shrink or Eliminate Bank Regulators." *Wall Street Journal*, December 12. https://www.wsj.com/finance/regulation/trump-advisers-bank-regulations-fdic-efa761dc.

Hill, Travis. 2023. "Proposal to Revise the Regulatory Capital Requirements for Large Banks." Statement at the FDIC Board Meeting, July 27. https://www.fdic.gov/news/speeches/2023/spjul2723b.html.

Hoshi, Takeo, and Anil Kashyap. 1999. "The Japanese Banking Crisis: Where Did It Come from and How Will It End?" In *NBER Macroeconomics Annual 1999*, Vol. 14, edited by Ben S. Bernanke and Julio J. Rotemberg, 129–201. MIT Press.

Hoshi, Takeo, and Anil K. Kashyap. 2004. "Japan's Financial Crisis and Economic Stagnation." *Journal of Economic Perspectives* 18 (1): 3–26.

Jones, Huw. 2024. "EU to Delay Core Element of Basel Bank Capital Reforms by One Year." *Reuters*, June 18. https://www.reuters.com/business/finance/eu-says-delay-core-element-of-basel-bank-capital-rules-2024-06-18/.

Klein, Aaron. 2016. "Risk Weights or Leverage Ratio? We Need Both." Brookings, December 22. https://www.brookings.edu/articles/risk-weights-or-leverage-ratio-we-need-both/.

Kress, Jeremy. 2024. "Basel Endgame Supporters Need a Litigation Strategy of Their Own." *American Banker*, July 11. https://www.americanbanker.com/opinion/basel-endgame-supporters-need-a-litigation-strategy-of-their-own.

Kroszner, Randall S. 2024. "White Paper on Basel III Endgame Proposal," February 5. https://www.federalreserve.gov/SECRS/2024/February/20240209/R-1813/R-1813_020724_158313_294423925299_1.pdf.

McKernan, Jonathan. 2023. "Proposed Amendments to the Capital Framework." Statement at the FDIC Board Meeting, July 27. https://www.fdic.gov/news/speeches/2023/spjul2723c.html.

Metrick, Andrew. 2024. "The Failure of Silicon Valley Bank and the Panic of 2023." *Journal of Economic Perspectives* 38 (1): 133–52.

Miller, Merton H. 1988. "The Modigliani-Miller Propositions after Thirty Years." *Journal of Economic Perspectives* 2 (4): 99–120.

Office of the Comptroller of the Currency, Federal Reserve System, and Federal Deposit Insurance Corporation. 2023. "Regulatory Capital Rule: Large Banking Organizations and Banking Organizations with Significant Trading Activity." *Federal Register* 88 (179): 23671.

Office of the Comptroller of the Currency, Federal Reserve System, and Federal Deposit Insurance Corporation. 2025. *Regulatory Capital Rule: Modifications to the Enhanced Supplementary Leverage Ratio Standards for U.S. Global Systemically Important Bank Holding Companies and Their Subsidiary Depository Institutions; Total Loss-Absorbing Capacity and Long-Term Debt Requirements for U.S. Global Systemically Important Bank Holding Companies.* Notice of Proposed Rulemaking, 12 CFR Part 324, June 25. https://www.federalreserve.gov/aboutthefed/boardmeetings/files/frn-leverage-ratio-20250625.pdf.

Pellerin, Sabrina. 2025. "Bank Capital Analysis Semiannual Update." Federal Reserve Bank of Kansas City, May 21. https://www.kansascityfed.org/documents/10877/Bank_Capital_Analysis_Report_-_4Q_2024_-_final.pdf.

Powell, Jerome. 2023. "Statement by Chair Jerome H. Powell." Press release, Board of Governors of the Federal Reserve, July 27. https://www.federalreserve.gov/newsevents/pressreleases/powell-statement-20230727.htm.

Powell, Jerome. 2024. "Transcript of Chair Powell's Press Conference." September 18. https://www.federalreserve.gov/mediacenter/files/FOMCpresconf20240918.pdf.

Quarles, Randal K. 2018. "A New Chapter in Stress Testing." Speech, Brookings Institution, Washington, DC, November 9, 2018. https://www.federalreserve.gov/newsevents/speech/quarles20181109a.htm.

Quarles, Randal K. 2021. "Between Hither and the Farther Shore—Thoughts on Unfinished Business." Speech, American Enterprise Institute, Washington, DC, December 2, 2021. https://www.bis.org/review/r211203a.htm.

Simoens, Matheiu, and Rudi Vander Vennet. 2021. "Bank Performance in Europe and the US: A Divergence in Market-to-Book Ratios." *Finance Research* Letters 40: 101672.

Tarullo, Daniel K. 2008. *Banking on Basel: The Future of International Bank Regulation.* Peterson Institute for International Economics.

Waller, Christopher J. 2023. "Statement by Governor Christopher J. Waller." Press Release, Board of Governors of the Federal Reserve System, July 27.

Wilkes, Samuel. 2025. "US Basel Equivalence Questioned as EU Patience Wears Thin." Risk.net, May 1. https://www.risk.net/regulation/7961433/us-basel-equivalence-questioned-as-eu-patience-wears-thin.

Yetis, Ahmet. 2008. "Regulators in Accord." *Risk Magazine*, January 1. https://www.risk.net/regulation/basel-committee/1526824/regulators-accord.

Zaring, David. 2019. *The Globalized Governance of Finance.* Cambridge University Press.

Carbon Rollercoaster: A Historical Analysis of Decarbonization in the United States

Karen Clay, Akshaya Jha, Joshua Lewis, and Edson Severnini

The United States was the world's largest source of carbon emissions for more than a century. In 1888, the United States surpassed Great Britain in terms of annual carbon emissions, and it remained the leading source of emissions until 2005, when as shown in Figure 1, it was overtaken by China.

The dominance of the US economy in carbon emissions was not simply because of its large population. Black lines in Figure 2, panels A and B, show the trajectories of total and per capita carbon emissions from fossil fuels and industry in the United States. In 1888, the US economy emitted 370 million tons of carbon dioxide (CO_2). In 2005, the US economy emitted 6.1 billion tons, an increase of more than a factor of fifteen. Land use changes can also affect carbon emissions, primarily through deforestation for land clearing, fuel wood, and wood products.[1]

■ *Karen Clay is Teresa and H. John Heinz III Professor of Economics and Public Policy and Akshaya Jha is Associate Professor of Economics and Public Policy, both at Heinz College of Information Systems and Public Policy, Carnegie Mellon University, Pittsburgh, Pennsylvania. Joshua Lewis is Associate Professor in Economics, University of Montreal, Montreal, Canada. Edson Severnini is Associate Professor of Economics and Core Faculty of the Schiller Institute for Integrated Science and Society, Boston College, Boston, Massachusetts, and Visiting Professor of Economics at Nova School of Business and Economics, Lisbon, Portugal. Clay and Severnini are Research Associates and Jha is a Faculty Research Fellow at the National Bureau of Economic Research, Cambridge, Massachusetts. Their email addresses are kclay@andrew.cmu.edu, akshayaj@andrew.cmu.edu, joshua.lewis@umontreal.ca, and edson.severnini@bc.edu.*

[1] The net carbon emissions from land use and forestry changes include carbon released from activities such as deforestation, planting trees, logging, and the degradation of forests (including timber

Figure 1
Carbon Dioxide Emissions for Selected Countries
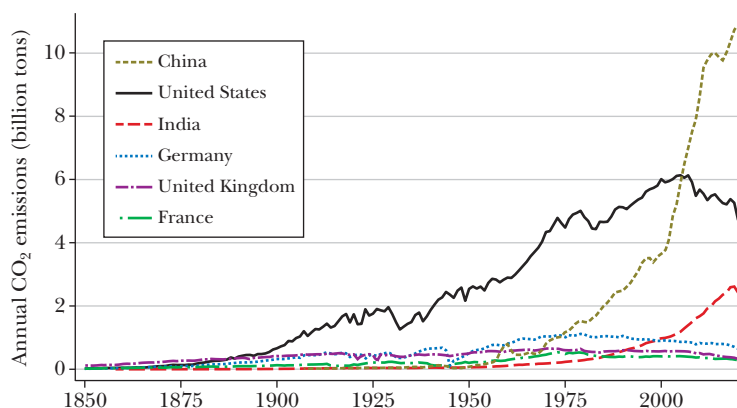
Source: The data underlying this figure come from Ritchie, Rosado, and Roser (2023).
Note: This figure plots annual total carbon dioxide emissions from fossil fuels and industry in selected countries (in billions of tons).

The dashed line in Figure 2, panel A, shows that land use changes were relatively important initially but contributed an increasingly small share of carbon emissions after 1940. Thus, this paper focuses on emissions from fossil fuels and industry. As shown in Figure 2, panel B, on a per capita basis, US carbon emissions were 6.3 tons in 1888, peaked at 23.1 tons in 1973, and declined to 20.7 tons by 2005. On a per capita basis, the United States surpassed Great Britain in 1903 and was surpassed by Australia in 2009.[2]

The growth rate of US carbon emissions and carbon emissions per capita shows striking variation over time. Although carbon emissions displayed a rising general trend over the twentieth century, Figure 2 indicates that the United States also experienced sustained periods of acceleration and deceleration in both total carbon emissions and per capita emissions. Our understanding of this historical variation in carbon emissions is surprisingly limited. Most studies do not go back very far in time. Those that do cover long historical periods tend to focus on decompositions and high-level trends in fuel use (Tol, Pacala, and Socolow 2009; Henriques and Borowiecki 2017).

---

harvesting). It also covers practices like shifting cultivation (where forests are cleared for farming and then abandoned), forest regrowth after logging or agricultural abandonment, as well as emissions from burning or draining peatlands (Friedlingstein et al. 2025). It is important to point out, however, that "[c]omparison of estimates from multiple approaches and observations shows . . . a persistent large uncertainty in the estimate of land-use change emissions" (Friedlingstein et al. 2025, p. 968).

[2] A number of small fossil-fuel-rich countries in the Middle East have even higher per capita carbon emissions, but are not comparable to the United States along other dimensions.
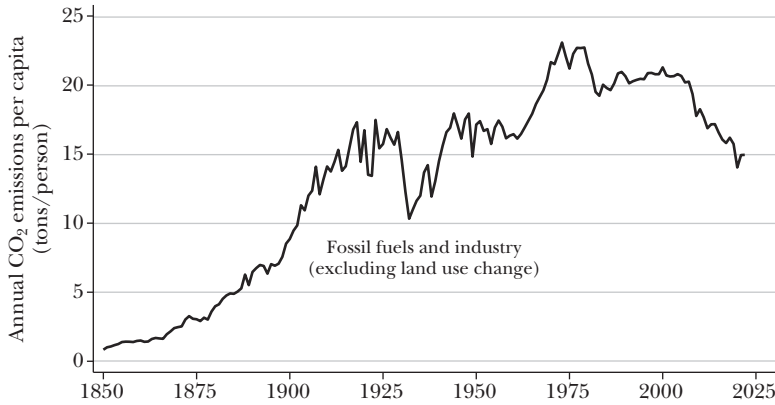
Figure 2
CO$_2$ Emissions and Per Capita CO$_2$ Emissions in the United States

Panel A. Annual CO$_2$ emissions
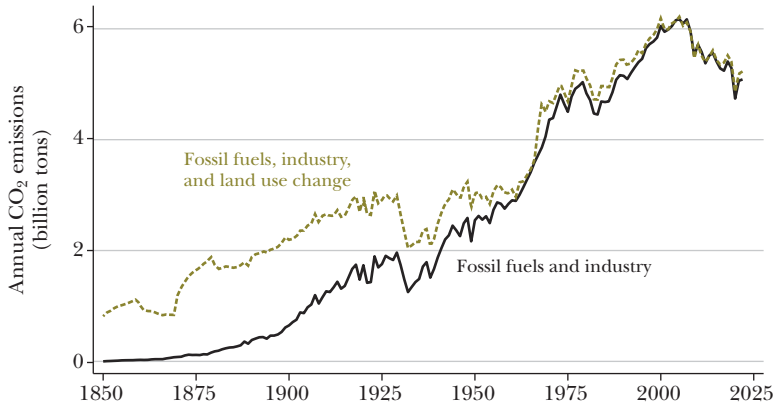
Panel B. Annual CO$_2$ emissions per capita

Source: The data underlying this figure come from Ritchie, Rosado, and Roser (2023).
Note: The top panel of this figure plots annual total carbon dioxide emissions in the United States from fossil fuels and industry as well as from fossil fuels, industry, and land use change (in billions of tons). The bottom panel of this figure plots annual carbon dioxide emissions per capita (in tons per person).

This paper documents the changing trends in US carbon emissions and discusses the main factors that contributed to the historical carbon emissions rollercoaster. We divide the discussion into four periods: up to 1920, 1920–1960, 1960–2005, and after 2005. For each period, we discuss the main drivers of national carbon emissions. We then discuss trends in carbon emissions in the electricity sector. Electricity sector emissions were initially very small, but would become the largest source of US carbon emissions over the period 1980–2015, and the largest contributor to decarbonization since 2007. In the last section, we offer some lessons for what developing economies might learn from the US experience.

Increases in Carbon Emissions up to 1920

The Economy up to 1920

Both total and per capita US carbon emissions rose rapidly up to about 1920, as shown in Figure 2. Two interrelated factors played key roles. The first was early US development of fossil fuels. Other economies had larger endowments, but the United States invested in extracting them earlier. The second was the use of these fossil fuels to support economic activity. In the absence of development of fossil fuels and the subsequent expansion of economic activity, carbon emissions would have remained low.

By the early twentieth century, the United States was the world's leading producer of fossil fuels (Wright 1990), producing 39 percent of world coal, 65 percent of world petroleum and 95 percent of world natural gas. This is despite having about 23 percent of the total world endowment in 1989 (reserves plus historical cumulative production) of coal and 20 percent of world endowment in petroleum. The US economy was also a top producer across a broad range of industrial minerals including iron ore, copper, lead, and zinc, among others. Other countries also had large natural endowments of minerals, including fossil fuels. What distinguished the United States was early development of its resources through state mineral surveys, state investment in mining education, and the incentives for discovery provided in federal mining laws (David and Wright 1997). State and federal policy was a catalyst that increased both economic activity and carbon emissions.

Coal became the largest source of energy for the US economy starting around 1880, as shown in Figure 3, when it surpassed renewables. In this era, renewables were wood and hydropower. In 1920, coal remained the dominant fossil fuel, although petroleum and to a lesser extent natural gas were beginning to play small roles in consumption.

Coal and other fossil fuels were used to support economic activity, such as transportation. For example, coal represented 90 percent of railroad fuel in 1880 (Atack 2024). In 1915, railroads consumed slightly more than one-quarter of all bituminous coal consumed by the US economy and a smaller share of anthracite coal. Petroleum was initially refined into kerosene for lighting, but became increasingly important as petroleum fueled cars, trucks and airplanes.

Electricity and steam generated by fossil fuels powered US manufacturing during this time (Lafortune et al. 2021). Up to 1880, industrial production grew steadily (Davis 2004, 2006). After 1880, it began to grow more rapidly. And after 1910, it began to exhibit increasing returns to scale (Lafortune et al. 2021). Consistent with these trends, carbon emissions per dollar of GDP rose up to World War I (Muller 2022).

Fossil fuels helped the United States surpass the United Kingdom to become the leader in world industrial production (Wright 1990). One of the distinctive characteristics of US manufacturing exports between 1880 and 1920 was their high and rising intensity in nonreproducible natural resources. Coal was a key input
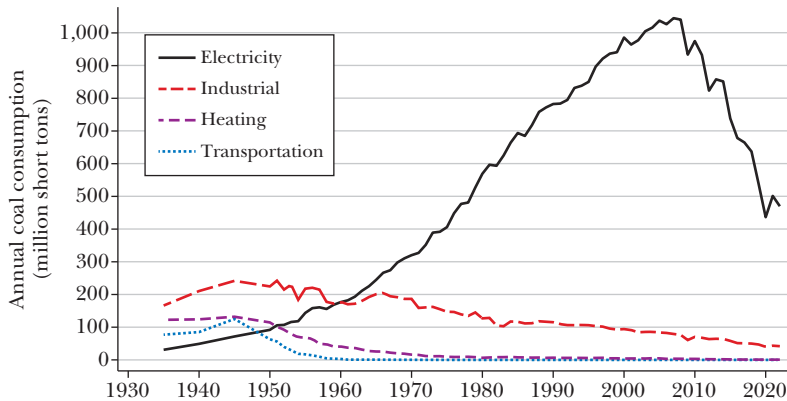
Figure 3
Energy and Coal Consumption in the United States

Panel A. Energy consumption by fuel type
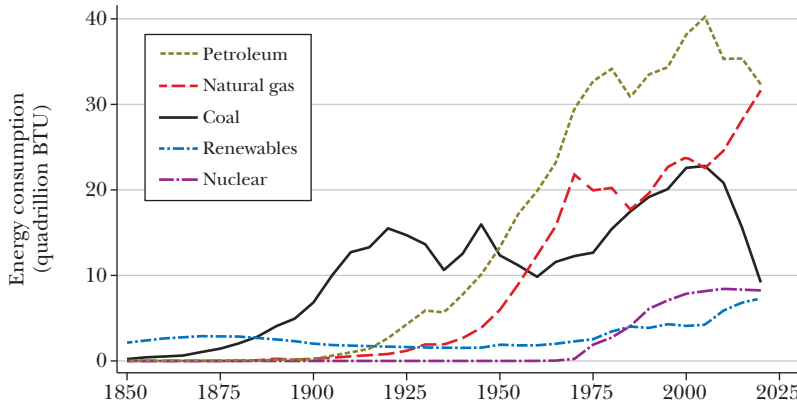
Panel B. Coal consumption by sector

Source: The data underlying this figure are from US Energy Information Administration (2024b). Panel B is based on data from the US Geological Survey (various years) and Cintron (1995).
Note: Panel A of this figure plots annual total primary energy consumption (in quadrillion British thermal units) by source. Renewables include wood, hydropower, and more recently wind and solar power. Panel B of this figure plots annual total coal consumption (in million short tons) by sector.

into iron and steel products, and iron and steel products along with petroleum products were increasing shares of US exports. By 1913, these two categories represented 38 percent of US exports, and the shares would climb above 50 percent in the 1920s. Petroleum exports were initially kerosene and later included gasoline, diesel, aviation gasoline, and other products. During World War I, petroleum products took on enormous strategic importance as motorcycles, cars, trucks, and airplanes played critical roles.

The high fossil fuel intensity of early exports raises an important point. Our discussion will focus on territorial emissions from the burning of fossil fuels. These do not include the carbon emissions associated with imported and exported goods or services. Data series on consumption-based emissions are only available from 1990 onward. Over that period, consumption-based emissions for the United States were generally not more than 10 percent higher than territorial emissions (Ritchie, Rosado, and Roser 2023).

The Electricity Sector up to 1920

Total coal consumption by US electric utilities increased dramatically from 1882 to 1920. Coal consumption in the electricity sector was small in the first few decades following the opening of the first commercial electricity plant, Thomas Edison's Pearl Street power plant in New York City, in 1882. By the turn of the twentieth century, electric utilities consumed roughly nine million tons of coal per year, accounting for just 3 percent of total US consumption. Annual coal consumption by electric utilities rose to 37 million tons by 1920. Despite this increase, however, the electricity sector accounted for just 7 percent of total US consumption by 1920 (US Census Bureau 1976).

The growth in coal consumption by electric utilities reflected a national rise in electricity generation. Between 1900 and 1920, total generation by electric utilities increased from less than two billion kilowatt-hours to nearly 40 billion kilowatt-hours (US Census Bureau 1976). This increase was driven by rising demand for electricity. At the turn of the century, less than 5 percent of homes were equipped with electric lights and less than 5 percent of manufacturing was powered by electricity (Lebergott 1976; Devine 1983). Over the next four decades, there was a sharp increase in coal-fired generation as electricity replaced steam power in manufacturing and the nation's homes began to electrify.[3]

Despite the rise in generation, the electricity sector remained a minor contributor to total US carbon emissions through the early twentieth century. Indeed, Figure 3, panel B, highlights that electricity was the smallest of the coal-consuming sectors in 1933, when the series on coal consumption by sector begins.

In summary, over the period up to 1920, carbon emissions grew rapidly. Growth was driven by early development of fossil fuels and their use to support economic activity, including transportation, manufacturing, and exports. While carbon emissions from the electricity sector would later be large, its emissions were small during the period up to 1920. Had the development of fossil fuels and the expansion of economic activity not occurred, carbon emissions would have remained low.

---

[3] By 1920, just over one-third of US homes were wired for electricity. Even in electrified homes, electricity consumption was limited, as few families could afford to purchase major electric appliances. As a result, residential consumption accounted for just 7 percent total electricity use in 1920 (Lebergott 1976; US Census Bureau 1976).

Slower Growth in Carbon Emissions from 1920 to 1960

The Economy from 1920 to 1960

If the trends from 1900 to 1920 had continued, there would have been more rapid increases in carbon emissions overall and a continuation in growth of per capita emissions up to 1960. Instead, we see slower growth overall, a flattening of emissions per capita, and sharp declines during the Great Depression decade (1929–1939). What drove the changes in the non–Great Depression years? Three factors appear to have contributed. The first is the changing energy mix, which had implications for carbon emissions, because petroleum and natural gas have lower carbon emission factors per unit of energy than coal. The second is thermal efficiency improvements, which allowed more output to be created from the same amount of energy. The third was the continuation of a preexisting trend in which services made up an increasing share of employment and GDP. All three slowed the growth of carbon emissions. In their absence, carbon emissions would have been higher.

Figure 3, panel A, shows changes in the energy mix that began around 1920. Coal consumption, which had been rising rapidly, began to flatten and even fall. This pattern continued until about 1960, when coal consumption began to rise again. Figure 3, panel B, highlights shifting sectoral use of coal. Electricity sector consumption was rising and by 1960 was equal to industrial consumption. Industrial consumption was fairly flat, and consumption for heating and transportation were falling. The declines in heating and transportation reflected a shift to petroleum and natural gas for fuel. Oil and natural gas began to replace coal for residential and commercial heating in many parts of the country. Locomotives and ships shifted from coal-fired to diesel engines.

The period from 1920 to 1960 was one of rapid improvements in thermal efficiency and growth in total factor productivity in the US economy (Jorgenson 1984; Gordon 2017; Field 2018; Bakker, Crafts, and Woltjer 2019). The increases in thermal efficiency were due to a number of factors, including increases in the efficiency with which coal was converted into electricity and possibilities that electricity opened for more efficient use of capital and labor in many industries (Jorgenson 1984; Rosenberg 1998).[4] Although output increased, greater efficiency in converting input heat to output slowed the growth in carbon emissions.

The service sector as a share of employment and services as a share of GDP were both rising over the period 1920–1970. These were continuations of a preexisting trend. In 1880, 26 percent of civilians were employed in the service sector (Fuchs 1980). By 1920, the share was 35 percent; by 1960, it was 52 percent. In line with this trend, the share of GDP accounted for by service-producing industries increased steadily since 1947, when data first became available (Yuskavage and

---

[4] There is some evidence that thermal efficiency was improving even before 1920, especially in the electricity sector. The electricity sector was, however, small before 1920.

Fahim-Nader 2005). These changes are important because the service sector has lower average carbon emissions than goods-producing sectors (Suh 2006).

The Electricity Sector from 1920 to 1960

Between 1920 and 1960, growth in carbon emissions in the electricity sector slowed as well. Unlike the overall economy, this change was largely not driven by the changing energy mix in electricity. Figure 3, panel B, shows the expanding role of electricity generation in total US coal consumption. By 1960, the electricity sector consumed the most coal of any sector. Further, despite increases in federal investments in hydroelectric capacity, the shares of electricity generation from coal and natural gas increased over this time period. In 1960, coal was used for 55 percent of all electricity generation and 69 percent of fossil fuel electricity generation. Instead, the slowdown in carbon emissions in the electricity sector, like the overall economy, was driven by increases in productivity.

Increases in productivity in the electricity sector were due to interrelated changes in technology, transmission, and state and federal policy. Coal-fired electricity plants generate electricity by heating water into high-pressure steam in order to spin turbines and ultimately to generate electrical energy. The period from 1920 to 1960 was characterized by technological advancements that allowed new plants to utilize steam with higher temperature and pressure—with the associated improvements in thermal efficiency. In particular, in the post–World War II era, the electricity industry benefited from wartime-related improvements in design and metallurgy that supported larger scale (Lovell 1941; Rosenberg 1998).[5]

The continued development and interconnection of high voltage transmission infrastructure also played a crucial role by reducing the reliance of communities on small, local electricity producers. This interconnected grid enabled utilities to build larger and more efficient electricity plants farther away from demand centers (Clay et al. 2025).

State and federal policies seeking to broaden access to electricity also contributed to the improvements in productivity in the pre–World War II era, but mostly as an unintended consequence. Beginning in 1907, states began regulating the electricity sector, establishing rate-of-return regulation and enacting territorial monopolies. These new regulations reduced the borrowing costs for private electric utilities and allowed them to exploit economies of scale (Hausman and Neufeld 2002; Knittel 2006). These policies reinforced the shift away from smaller and less efficient municipal plants to larger-scale privately owned utilities. At the same time, they promoted consolidation of privately owned utilities, which were better able to exploit economies of scale under the new regulation. Thus, state and federal policies helped reduce carbon emissions per unit of electricity.

---

[5] The effects on scale are evident in Supplemental Appendix Figure A.1, which documents the size of the largest generating unit and the average unit size in each year (Hales 1976). The size of the largest unit increased dramatically around 1950, and the average size installed in a given year began to trend up around the same time.

Figure 4
Coal Efficiency and Carbon Emissions per Gigawatt-Hour
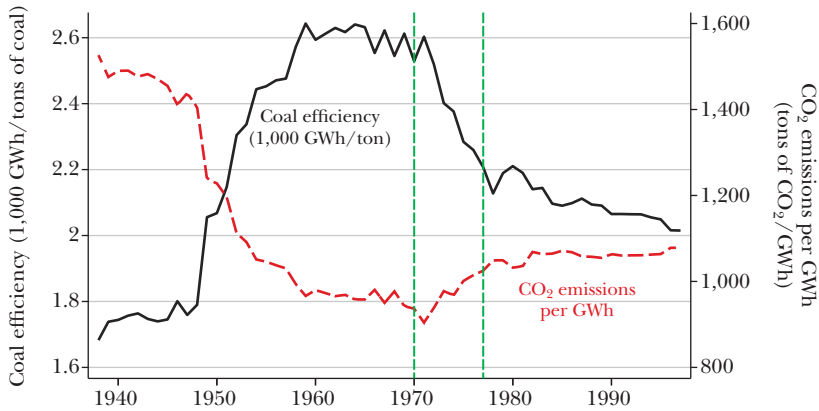
Source: This figure is based on data on power plant operations from 1938 to 1994 from Clay et al. (2025).
Note: In this figure, we plot annual nationwide coal efficiency: annual total electricity production from coal-fired power plants (in 1,000 GWh) divided by annual total tons of input coal burned (in tons). In addition, we plot annual total carbon dioxide emissions divided by coal-fired electricity production (in tons of emissions per GWh of output). We translate input heat energy from coal to carbon emissions using a conversion factor of 211.63 pounds of CO_2 per million BTU of coal (EIA 2024a). The green vertical dashed lines represent the 1970 Clean Air Act and the 1977 amendments.

As a result of these changes in technology, transmission, and state and federal policy, the electricity sector witnessed substantial improvements in coal efficiency (coal-fired electricity generation divided by tons of coal burned for generation). Drawing on aggregate data on electricity generated and coal burned by electric utilities from 1920 to 1960 (US Census Bureau 1976), we calculate that coal efficiency increased by more than 400 percent. The patterns are consistent with the findings of Bakker, Crafts, and Woltjer (2019).

Beginning in 1938, we have annual power-plant-level data, which allow us to examine changes in coal efficiency and carbon emissions in more detail (Clay et al. 2025). The coal efficiency of coal-fired electricity plants improved substantially between 1938 and 1960, as illustrated in Figure 4. In 1940, it took 573 tons of coal to generate a gigawatt-hour of electricity. By 1960, this had fallen to 386 tons per gigawatt-hour.

These improvements in coal efficiency resulted in a 34 percent decline in carbon emissions per gigawatt hour of electricity between 1940 and 1960.[6] To put this decline in perspective, annual total carbon emissions from coal-fired generation in 1960 was 387 million tons, but would have been 589 million tons if total coal-fired

---

[6] Throughout the paper, we translate input heat energy from coal to carbon emissions using a conversion factor of 211.63 pounds of CO_2 per million BTU of coal (US Energy Information Administration 2024a).

generation in 1960 was produced at the aggregate carbon emissions per gigawatt hour of electricity in 1940.

Beginning in the 1930s, the federal government was a major driver of the expansion in hydroelectric capacity. This period, known as the "Big Dam Era," witnessed the construction of major federal hydroelectric generation stations, including the hydroelectric plants run by the Tennessee Valley Authority and the Bonneville (Oregon) Power Administration (Billington, Jackson, and Melosi 2005). The ability to develop these new sources of hydropower also depended on the technological innovations in transmission technology, which allowed for cost-effective use of remote sources of hydropotential (Lovell 1941). Thus, federal policy with respect to hydroelectric generation helped slow the growth in carbon emissions, because hydroelectric generation does not emit carbon.

This expansion in hydroelectric capacity reduced reliance on coal-fired generation and contributed to the decline in the share of electricity production from coal over this period. Hydroelectricity resources produced roughly 16,000 gigawatt-hours of electricity in 1920 and roughly 47,000 gigawatt-hours of electricity in 1940. To put these magnitudes in perspective, if the 31,500 additional gigawatt-hours produced by hydro in 1940 relative to 1920 was instead produced by coal, carbon emissions from coal-fired electricity production would have been 41 percent higher in 1940 (an increase of roughly 38.5 million tons in carbon emissions).[7]

Rising coal consumption in the electricity sector brought with it concerns about air pollution, because emissions of conventional air pollutants were large, highly visible, and increasingly concentrated in specific locations. Public concern about the impacts of local air pollution were spurred by high-profile episodes such as the Donora Smog in 1948, which killed 20 people and caused respiratory problems for thousands more in Pennsylvania, and the Great Smog of London in December 1952, which led to about 12,000 deaths (Bell and Davis 2001). The perceived value of having electricity produced by a nearby coal-fired plant evolved over this time. Looking at data for local areas surrounding all major US coal-fired plants from 1938 to 1962, Clay, Lewis, and Severnini (2024) suggest that earlier in this period, expansions in coal capacity for electricity production led to decreases in local infant mortality, but later additional increases led to increases in infant mortality.

Despite these concerns, there was little federal or state action to address air pollution during this time frame. The Air Pollution Control Act of 1955 allocated

---

[7] Aggregating over the coal-fired fleet in 1940, burning 639 short tons of coal resulted in one gigawatt-hour of electricity. Therefore, a 31,500 gigawatt-hours increase in coal-fired generation corresponds to a 20.1 million tons increase in coal burned. Applying a conversion rate of 3,826.88 pounds of CO_2 per short ton of coal burned (US Energy Information Administration 2024a), we obtain that roughly 38.5 million more tons of carbon would have been emitted if the increase in hydro production between 1920 and 1940 was instead produced by the 1940 coal-fired fleet. Because carbon emissions from the coal-fired fleet in 1940 were roughly 94 million tons of CO_2, the 38.5-million-tons increase from counterfactually replacing hydro output with coal output amounts to a 41 percent increase in CO_2 emissions. All of the numbers underlying this calculation come from the US Energy Information Administration. They are adjusted prior to 1949 based on US Bureau of the Census (1960, p. 507).

funding for studying the issue, but did not implement regulatory measures. Electric utilities faced no immediate threat of government regulation and made minimal efforts to reduce local air pollution emissions. Environmental considerations were scarcely mentioned in annual reports of the Federal Power Commission. Its 1966 report, "Steam-Electric Plant Construction Cost and Annual Production Expenses," included, for the first time, a section on "environmental influences on plant design, construction, and operation." The report highlighted that air pollution, among other issues, was becoming a significant socioeconomic concern for the electric power industry (Federal Power Commission 1967, p. ix). Although experimentation with "baghouses" (which use fabric bags to trap pollutants as they are emitted) and other pollution-abatement technologies began in the 1940s, few coal-powered electricity plants had installed such systems before 1960.

Over the period 1920–1960, the economy experienced changes in energy mix, thermal efficiency improvements, and increases in the service sector as a share of overall employment and GDP. The electricity sector experienced rapid increases in productivity as measured by coal efficiency, and the federal government made large investments in hydroelectric dams. Taken together, these changes slowed the growth in carbon emissions overall and led to the flattening of carbon emissions per capita. In the absence of these changes, carbon emissions would have been much higher.

Faster Growth in Carbon Emissions from 1960 to 2005

The Economy from 1960 to 2005

US carbon emissions grew rapidly from the 1960s up to 2005, except for a period in the 1970s and early 1980s, as shown earlier in Figure 2. Why do we see faster growth over the period 1960–2005 as compared to 1920–1960? The primary story, both economy-wide and for the electricity sector specifically, is a changing energy mix driven in part by crises and regulation. Figure 3, panel A, shows the increase in consumption of all fuel types between 1960 and 2005. Petroleum and natural gas rose, fell during the energy crises, and then rose again. Coal, after falling in the years before 1960, began to increase steadily. Nuclear began to contribute significantly around 1970. Renewables increased more slowly. In the absence of crises and regulation and holding total energy consumption constant, coal consumption would likely have grown more slowly or even declined. Given that alternative fossil fuels, nuclear, and renewables have lower carbon emissions, total carbon emissions would have increased more slowly if coal had grown more slowly or declined.

Two other factors, which had contributed to slower growth in previous periods, made smaller contributions in this period. Thermal efficiency improved more slowly relative to the previous period and output continued to expand, increasing overall emissions (Jorgenson 1984; Metcalf 2008). Further, gains from ongoing shifts to the service sector were smaller, because of smaller differences in emissions in the service and goods sector.

By the 1960s and early 1970s, US oil and gas markets faced challenges. On the oil side, consumption and imports were rising, because domestic production was too low to meet demand (Akins 1973; Painter 2012; Hamilton 2013). President Richard Nixon created a Cabinet Task Force on Oil Import Controls to investigate oil imports in 1969 (Bohi and Russell 2013). However, the task force's final report, *The Oil Import Question*, in early 1970 suggested that price increases would be limited and the market share of oil imports would remain small. In early 1973, price controls associated with rising inflation, together with the end of the Bretton Woods fixed-exchange-rate agreement and the resulting depreciation of the dollar, were causing gasoline shortages.

Natural gas was facing significant supply and demand imbalances as early as the 1950s, which caused prices to rise (Akins 1973; Davis and Kilian 2011). The rapid increase in prices led to the 1954 Supreme Court case *Phillips Petroleum Company v. Attorney General of Wisconsin* (347 US 672 [1954]), which held that natural gas prices were subject to federal regulation. The court ruling led to declining investment in gas production, and by 1970, this had begun the curtailment of natural gas deliveries to industrial customers.

These preexisting issues were greatly exacerbated by how global energy markets were affected by aftereffects of the Arab-Israeli Yom Kippur War in 1973 and the Iranian Revolution in 1978–1979 (Painter 2012; Hamilton 2013; Brew 2019). The Organization of Petroleum Exporting Countries (OPEC) had been founded in 1960 for its members to obtain greater leverage with the major multinational oil companies. In 1971, the major oil companies and six OPEC members in the Persian Gulf signed the Tripoli Agreement. In October 1973, Arab and some non-Arab members of OPEC declared an embargo against the United States and other countries that supported Israel in the 1973 war and subsequently cut production. The real price of imported oil rose dramatically, from $10.67 per barrel in 1972 (in 2007 US dollars) to $36.05 in 1974 (Seiferlein 2007, p. 171). Turbulence in the Middle East kept prices high. Unrest in Iran and the Iran-Iraq War caused further disruption, driving oil prices to $62.71 per barrel in 1980.

Figure 5 shows carbon emissions by sector starting in 1973, when the data series begins.

In 1973, emissions from the industry, transportation, and electricity sectors were quite similar and much higher than emissions from the commercial and residential sectors.[8] By 1980, the electricity sector had the largest emissions. It remained the largest sector until 2005. Emissions from the transportation sector rose over the period as well, although more slowly than the electricity sector. In contrast to the electricity and transportation sectors, industrial emissions fell somewhat up to 1985 and were fairly flat thereafter. By 2005, industrial emissions were half or less than half of the transportation sector and the electricity sector.

---

[8] The commercial sector consists of service-providing businesses and organizations, while the industrial sector is made up of goods-providing businesses and organizations.

Figure 5
Carbon Emissions by Sector
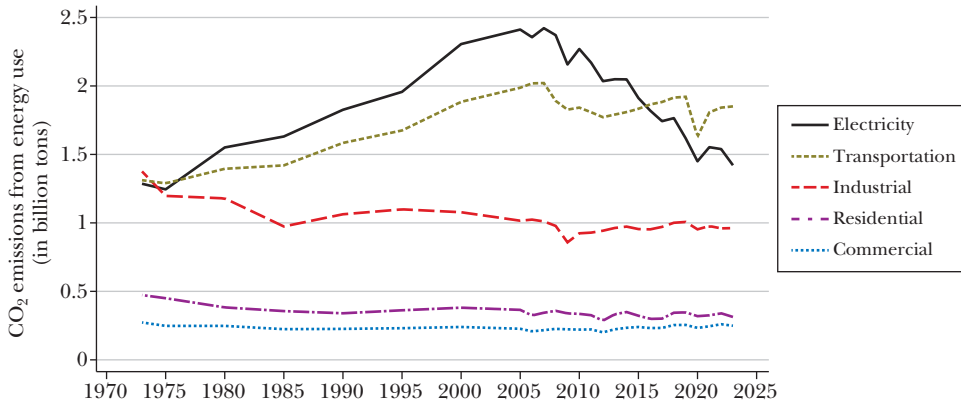
Source: The data underlying this figure come from US Energy Information Administration (2025b).
Note: This figure plots annual total carbon dioxide emissions from energy use by sector from 1973 to 2024 (in billion tons).

The Electricity Sector from 1960 to 2005

Like national carbon emissions, emissions from the electricity sector rose rapidly. This faster growth was driven by three factors. First, coal consumption in the sector continued to rise, as documented in Figure 3, panel B. By 2005, coal consumption was five times what it had been in 1960. This in part reflects the fact that legislation during the energy crisis of the 1970s and the accident at the Three Mile Island nuclear power plant in Pennsylvania made it difficult to build new plants that utilized other fuel sources. Second, there were limited additional technological improvements after 1960. Third, the electricity industry was particularly affected by the 1970 Clean Air Act and 1977 amendments, such that the resulting actions taken to reduce local air pollution contributed to increasing carbon emissions.

Various regulations passed during and after the crisis reinforced the continued use of coal in electricity and other sectors. The first major piece of legislation was the Energy Supply and Environmental Coordination Act of 1974, which required that, if feasible, electric power plants burning oil and natural gas would have to convert to coal (Meltz 1975). This law was then largely superseded by the Fuel Use Act of 1978. Edward Lublin, Acting Deputy Assistant General Counsel for Coal Regulations in the Department of Energy, wrote: "The Fuel Use Act prohibits new facilities and allows DOE to prohibit existing facilities, from using petroleum or natural gas as a primary energy source unless DOE determines to grant to such facility an exemption from the Fuel Use Act's prohibitions (Lublin 1981, p. 355)." This pro-coal legislation was often justified in terms of energy independence, given the abundant US reserves of coal. The legislation covered both electric utilities and major industrial fuel-burning installations, although much of the Department of

Energy enforcement effort focused on electric utilities. While the legislation did not induce widespread conversion to coal at existing oil and gas plants—most utilities delayed any such conversion due to high cost, even when faced with conversion orders—new plants built after this legislation almost exclusively burned coal. Thus, these regulations increased carbon emissions.

Roughly concurrent with the energy crises of the 1970s, many nuclear power plants opened (Joskow 2006; Davis 2012). However, there was increasing resistance to nuclear power even beginning in the 1960s. Groups opposed to nuclear power filed legal cases. Federal agencies responded with safety-related changes and permitting and regulatory processes, and construction of nuclear power plants was delayed. These federal actions drove up costs, in some cases even among plants already under construction. In light of the pending cost increases, many proposed nuclear plants were canceled in the mid- to late 1970s.

Despite tightening standards, enough nuclear plants opened in the 1970s that nuclear power's share of total net electricity generation increased to around 10 percent. In 1979, a serious accident at Three Mile Island in Pennsylvania further highlighted the risks of nuclear power. After Three Mile Island, no new nuclear power plant construction was authorized until 2012. Because nuclear plants displaced coal-fired electricity generation—one gigawatt-hour of nuclear generation resulted in a roughly 0.8 gigawatt-hour decrease in coal-fired generation historically (Adler, Jha, and Severnini 2020)—the nuclear upheaval kept coal consumption higher than it would otherwise have been.

Electricity generation from coal as a share of all generation declined starting in the mid-1960s and then flattened in the 1970s. The share increased in the 1980s and remained above 50 percent up to 2005. Both oil and natural gas fell as shares of energy production over 1960–2005. The effects of Three Mile Island are less evident, because many nuclear plants were under construction and so would come online in the years following the accident. In the absence of the accident, however, it is likely that additional plants would have been built.[9]

After tremendous improvements in technology between 1940 and 1960, there were limited additional technological improvements to increase thermal efficiency. Until then, coal-fired power plants were operated at relatively low pressures. So-called "supercritical" boilers, developed after 1959, represented a major technological shift, using higher pressures to generate more efficient steam. They offered efficiency gains through reducing the need for some equipment and benefiting from economies of scale. Initially, the industry moved toward adopting supercritical technology, and by the mid-1970s, it seemed poised to replace older designs. However, by the 1980s, this shift reversed, and supercritical technology was largely abandoned (Joskow and Rose 1985). As noted in Joskow and Rose (1985), this reversal may have stemmed from higher outage rates and maintenance costs among supercritical units, as well as flagging electricity demand growth in the 1980s. These

---

[9] The effects of the energy crisis on the mix of fuels in the electricity sector is illustrated in Supplemental Appendix Figure A.2.

factors made it less attractive to build very large supercritical units that benefit from economies of scale.

Electrical utilities did not face significant pressure to curb emissions prior to 1960, although air pollution had been a source of local concern before that time. The period from 1960 to 2000 witnessed the emergence of the federal government's role in combating local air pollution. The 1963 Clean Air Act marked the first legislation granting the federal government authority to regulate air pollution. However, this law mainly created a small number of "abatement conferences," which brought together state and local officials, industry representatives, and federal authorities to discuss pollution sources, assess their impact, and seek voluntary agreements to reduce emissions before resorting to enforcement actions or legal steps. They were an early attempt to resolve air quality issues collaboratively. In the end, the 1963 legislation resulted in minimal enforcement to reduce air pollution.

The regulatory landscape changed dramatically with the passage of the 1970 Clean Air Act, which represented the first comprehensive federal effort to address local air pollution on a national scale. The 1970 Clean Air Act focused on reducing major local air pollutants—such as ambient ozone, particulate matter, and carbon monoxide. The United States saw sharp decreases in local air pollution emissions after 1972 (Cropper et al. 2023). The 1970 Clean Air Act had indirect effects on carbon emissions, even though CO_2 was not regulated at the time. As we describe below, efforts to cut local air pollution often increased carbon emissions. The 1970 Clean Air Act and subsequent amendments in 1977 coincided with less efficiency in converting coal to electricity sold and higher carbon emissions, as illustrated earlier in Figure 4. The aggregate implications of this shift from 1970 to 1990 are meaningful: annual total carbon emissions in 1990 from coal-fired generation was 1,607 million tons, but would have been 1,415 million tons if the same amount of coal-fired electricity had been generated at 1970 levels of carbon emissions per gigawatt-hour. Similarly, the aggregate kilowatt hours of electricity sold per ton of coal burned decreased from 2,529 in 1970 to 2,065 in 1990. Thus, regulation increased coal consumption and carbon emissions.

Compliance with environmental regulations reduced the efficiency with which coal was converted to electricity and increased carbon emissions for three reasons. First, the adoption of pollution-abatement technologies contributed to the declines in coal efficiency. These technologies (such as scrubbers, selective catalytic or noncatalytic reduction, and electrostatic precipitators) were implemented to control emissions of sulfur dioxide, nitrogen oxides, and particulate matter, respectively. However, these devices also require nontrivial amounts of energy to operate. Because the power plant must devote some of the energy it produces to operating pollution-abatement technology, it can sell less energy to the market. In addition, using some types of pollution-abatement technology lowers the thermal efficiency of power generation. For both reasons, the installation and operation of these technologies contribute to decreases in gigawatt-hours of electricity sold per tons of coal burned.

Second, in order to comply with the new regulations, many plants shifted towards burning coal with lower sulfur and heat content from the western United

States (Clay et al. 2025). Burning lower-sulfur coal reduced emissions of sulfur dioxide. But the fuel switching may also have lowered thermal efficiency, because plants are optimized to burn a particular type of coal (Joskow 1985). In addition, low-sulfur coal has lower heat content: more tons of coal are required to produce the same amount of heat input. This only had a small effect on carbon emissions, however, because carbon emissions per unit of heat are only slightly higher for western sub-bituminous coal than eastern bituminous coal (US Energy Information Administration 2025a).

Third, environmental legislation involved looking at emissions of criteria pollutants on a county-by-county basis. When a county was designated as "nonattainment," and not meeting emissions goals, power plants often responded by reducing output from that plant (Clay et al. 2025). These reductions may have led to plants producing at suboptimal efficiency levels.[10]

These three compliance methods highlight a significant trade-off between local pollution and carbon emissions. While the Clean Air Act effectively reduced local pollution, it may have inadvertently led to an increase in carbon emissions. The increase in carbon emissions per gigawatt-hour of electricity sold after 1970 underscores the challenge of balancing local air quality improvements with broader environmental impacts.

In summary, faster growth in carbon emissions at its most basic level was the result of higher energy consumption, including increases in coal consumption, a high carbon-emission fuel. In contrast to 1920–1960, there were limited improvements in thermal efficiency and limited gains from the ongoing shift to services. Both had helped offset the rise in carbon emissions prior to 1960. Unlike the previous period, the electricity sector was also a major source of emissions and the largest source of emissions from 1980 to 2005. Part of this was driven by the overall growth of the electricity sector. The energy crisis and the Three Mile Island accident reinforced the use of coal to produce electricity. This led to higher emissions, because coal has higher carbon emissions per unit of energy than oil, natural gas, and especially nuclear and renewables.

Declining Carbon Emissions from 2005 onward

After 150 years of nearly continuous increases in US carbon emissions, why do we suddenly see declines in carbon emissions after 2005 (as shown in Figure 2)? Given the importance of the electricity sector for emissions by 2005, we will discuss the economy and the electricity sector together in this section.

---

[10] Consistent with this insight, Supplemental Appendix Figure A.3 documents that aggregate coal efficiency decreased by more from 1970 to 1990 among coal-fired plants located in counties that ever faced nonattainment between 1972 and 1994 relative to coal-fired plants located in counties that never faced nonattainment between 1972 and 1994.

The patterns in Figure 3, panel A, presented earlier, suggest that these decreases in carbon emissions may have been driven by the substantial declines in petroleum and coal consumption and modest increases in renewables. Natural gas consumption increased, but this increase did not offset declines in the other two fossil fuels. In the absence of these declines in petroleum and especially coal, carbon emissions would have been substantially higher.

The key to understanding the decline in coal and the increase in natural gas is the natural gas fracking revolution. Coal had been the dominant energy source for electricity generation for decades, reaching its peak in the mid-2000s with over one billion short tons consumed annually. However, coal consumption saw a dramatic drop in the 2010s, falling by almost half by the early 2020s. Figure 3, panel B, highlights a sharp decline in US coal consumption for electricity generation starting around 2008. Not coincidentally, Figure 2 shows a drop in both total and per capita carbon emissions for the US economy starting about the same time.

Hydraulic fracturing, known as "fracking," involves pumping water, sand, and chemicals into underground rock formations. This creates fractures in the rock that allow drilling for natural gas. The "fracking revolution" unlocked vast reserves of natural gas in the United States, which in turn led to a drop in US natural gas prices of over 60 percent between 2008 and 2015. As a result, natural gas quickly displaced coal as the primary fuel for electricity generation. This switch reduced carbon emissions, because natural gas emits roughly 50–60 percent less carbon than coal when burned for energy.

The rise of renewable energy sources like wind and solar also contributed to the reduction in the use of coal for electricity generation in the United States. Technological innovation and declining costs helped make renewables more attractive. Government policies, including tax incentives and state-level renewable portfolio standards, further encouraged the shift to renewables. In 2010, only 2 percent of US total electricity generation came from wind and solar resources (US Energy Information Administration 2024b). However, wind power alone accounted for about 10 percent of US electricity generation in 2023, over quadruple its contribution from a decade earlier. Solar power accounted for about 4 percent of electricity generation in 2023, as compared to only 0.03 percent in 2010.

At the same time renewables were becoming more competitive from a cost perspective, regulation was making coal-fired generation more expensive. We have already seen that regulation—the 1970 Clear Air Act and the 1977 Amendments—reduced the efficiency of coal-fired power plants. Further regulation including the 1990 Clean Air Act Amendments made coal even more unattractive.

Changes were also occurring in the transportation sector. After the early 2000s, the US government took significant legislative steps to improve fuel efficiency in response to energy security concerns, rising oil prices, and growing awareness of climate change. Two landmark laws—the Energy Policy Act of 2005 and the Energy Independence and Security Act of 2007—played a central role in tightening fuel economy standards. The Energy Policy Act of 2005 laid the groundwork by promoting alternative fuels and energy efficiency, including incentives for hybrid

and alternative fuel vehicles. However, the more transformative legislation was the Energy Independence and Security Act of 2007, which mandated the first substantial increase in the Corporate Average Fuel Economy (CAFE) standards since the 1970s. It required automakers to raise the average fuel economy of new passenger cars and light trucks to 35 miles per gallon by 2020, up from roughly 25 miles per gallon at the time. The Act also expanded the Renewable Fuel Standard, pushing more ethanol and biofuels into the gasoline supply. These laws helped limit the growth of carbon emissions in the transportation sector.

Carbon emissions in the electricity sector began to fall around 2005, as shown in Figure 5. By 2016, transportation became the largest carbon-emitting sector. By 2024, carbon emissions in the electricity sector had fallen almost one billion tons from their peak in 2005. Thus, the electricity sector was responsible for virtually all of the decline in carbon emissions in the economy in the 2000s.

At present, the electricity sector faces a new challenge. Electricity consumption from data centers is rising rapidly with the large-scale development and use of artificial intelligence. Data centers accounted for just 2 percent of total US electricity consumption in 2018 but over 4 percent in 2023—and could consume 7 to 12 percent of total US electricity by 2028 (Shehabi et al. 2024). Meeting this growing demand will require substantial investment in new generation resources—with significant carbon-emissions implications depending on whether the added load is met by new gas-fired plants, expanded wind and solar capacity, or delayed coal plant retirements.

In summary, US carbon emissions began declining around 2005 after more than a century of growth, largely due to major shifts in the electricity sector. The widespread adoption of hydraulic fracturing made natural gas cheaper, allowing it to replace coal—which emits significantly more carbon—as the dominant source of electricity. Simultaneously, wind and solar power expanded due to falling costs, technological innovation, and supportive policies. Environmental regulations also made coal generation more expensive and less efficient. As a result, electricity-sector emissions fell by nearly one billion tons between 2005 and 2024, accounting for almost the entire national decline. However, rising electricity demand from data centers, driven by artificial intelligence, presents a new challenge for future decarbonization.

Implications for Developing Countries

Our analysis underscores several lessons from the US experience that may inform ongoing decarbonization efforts, particularly in developing economies where energy demand is rising rapidly. While each country will face unique institutional, political, and technological constraints, historical patterns offer useful guidance on avoiding carbon-intensive development paths and navigating the complex trade-offs between local and global environmental objectives.

First, coal remains a major fuel source for electricity generation worldwide, particularly in low- and middle-income countries (IEA 2025a). As electricity demand

from industry and households rises, the electricity sector is likely to play an increasingly central role in global carbon emissions (IEA 2024, 2025b). The US experience offers a precedent: by the mid-twentieth century, electricity generation had become the dominant consumer of coal. A similar pattern may emerge elsewhere unless countries adopt cleaner electricity-generation technologies.

Second, our findings underscore the central role of technological progress and coal-efficiency gains in curbing carbon emissions. In the United States, rapid innovation significantly contributed to efficiency gains between 1920 and 1960. In the electricity sector specifically, major advances—such as more efficient generators, high-voltage transmission lines, and economies of scale from centralized production—substantially lowered the amount of coal needed to generate electricity. By adopting proven technologies and investing in emerging innovations, research, and local capacity-building, developing economies may be able to expand electricity access while limiting carbon emissions.

Third, the US experience demonstrates how geopolitical shocks and policy responses can have unintended and enduring effects on the energy mix. The energy crises of the 1970s, coupled with efforts to enhance energy security, inadvertently deepened the nation's reliance on coal. Similarly, public backlash to nuclear accidents—Three Mile Island in 1979, Chernobyl in 1986, and Fukushima in 2011—led the United States, Germany, and others to scale back nuclear power in favor of fossil-fuel generation (Severnini 2017; Adler, Jha, and Severnini 2020; Jarvis, Deschenes, and Jha 2022). Given the long lifespans of energy infrastructure, these short-term policy responses had enduring consequences for carbon emissions. These historical episodes underscore the need to anticipate the long-run environmental and climate implications of energy decisions made in response to crises.

Finally, our analysis highlights the persistent challenge of designing environmental regulations that address both local and global pollutants. In the United States, the Clean Air Act of 1970 and its 1977 and 1990 amendments led to major reductions in harmful local pollutants such as sulfur dioxide and particulate matter. However, some of the compliance strategies, including the installation of scrubbers and a shift to low-sulfur coal with lower heat content, inadvertently increased overall coal consumption and thus carbon emissions. While many advanced economies are now seeking to transition toward renewables that do not emit local or global pollution in the generation process, developing countries that continue to rely on fossil fuels may face similar trade-offs. As emerging technologies such as carbon capture, utilization, and storage (CCUS) and hydrogen production become more widespread, it is essential to evaluate their broader environmental impacts. For instance, while CCUS can lower carbon emissions, it may increase releases of other pollutants, including ammonia (Waxman, Huber-Rodriguez, and Olmstead 2024).

Moving forward, effective climate and energy policy must draw on historical experience to design frameworks that minimize unintended consequences. Developing countries can avoid some of the pitfalls encountered by the United States by investing in efficient technologies, avoiding overreliance on coal, responding

strategically to geopolitical shocks, and crafting environmental regulations that reflect both local and global pollution objectives. By fostering innovation, maintaining regulatory flexibility, and holistically accounting for environmental externalities, policymakers can more effectively meet their climate goals while ensuring energy access and fostering economic growth.

References

Adler, David B., Akshaya Jha, and Edson Severnini. 2020. "Considering the Nuclear Option: Hidden Benefits and Social Costs of Nuclear Power in the U.S. since 1970." *Resource and Energy Economics* 59: 101127.

Akins, James E. 1973. "The Oil Crisis: This Time the Wolf Is Here." *Foreign Affairs* 51 (3): 462–90.

Atack, Jeremy. 2024. "Railroads." In *Handbook of Cliometrics*, edited by Claude Diebolt and Michael Haupert, 2329–57. Springer Nature.

Bakker, Gerben, Nicholas Crafts, and Pieter Woltjer. 2019. "The Sources of Growth in a Technologically Progressive Economy: The United States, 1899–1941." *Economic Journal* 129 (622): 2267–94.

Bell, Michelle L., and Devra Lee Davis. 2001. "Reassessment of the Lethal London Fog of 1952: Novel Indicators of Acute and Chronic Consequences of Acute Exposure to Air Pollution." *Environmental Health Perspectives* 109 (S3): 389–94.

Billington, David P., Donald C. Jackson, and Martin V. Melosi. 2005. *The History of Large Federal Dams: Planning, Design, and Construction in the Era of Big Dams.* US Department of Interio, Bureau of Reclamations.

Bohi, Douglas R., and Milton Russell. 2013. *Limiting Oil Imports: An Economic History and Analysis.* RFF Press.

Brew, Gregory. 2019. "OPEC, International Oil, and the United States." In *Oxford Research Encyclopedia of American History.* Oxford University Press. https://doi.org/10.1093/acrefore/9780199329175.013.719.

Cintron, Lizette, ed. 1995. *Historical Statistics of the Electric Utility Industry through 1992.* Edison Electric Institute.

Clay, Karen, Akshaya Jha, Joshua A. Lewis, and Edson R. Severnini. 2025. "Impacts of the Clean Air Act on the Power Sector from 1938–1994: Anticipation and Adaptation." NBER Working Paper 28962.

Clay, Karen, Joshua Lewis, and Edson Severnini. 2024. "Canary in a Coal Mine: Infant Mortality and Tradeoffs Associated with Mid-20th Century Air Pollution." *Review of Economics and Statistics* 106 (3): 698–711.

Cropper, Maureen, Nicholas Muller, Yongjoon Park, and Victoria Perez-Zetune. 2023. "The Impact of the Clean Air Act on Particulate Matter in the 1970s." *Journal of Environmental Economics and Management* 121: 102867.

David, A. Paul, and Gavin Wright. 1997. "Increasing Returns and the Genesis of American Resource Abundance." *Industrial and Corporate Change* 6 (2): 203–45.

Davis, Joseph H. 2004. "An Annual Index of US Industrial Production, 1790–1915." *Quarterly Journal of Economics* 119 (4): 1177–215.

Davis, Joseph H. 2006. "An Improved Annual Chronology of US Business Cycles since the 1790s." *Journal of Economic History* 66 (1): 103–21.

Davis, Lucas W. 2012. "Prospects for Nuclear Power." *Journal of Economic Perspectives* 26 (1): 49–66.

Davis, Lucas W., and Lutz Kilian. 2011. "The Allocative Cost of Price Ceilings in the US Residential Market for Natural Gas." *Journal of Political Economy* 119 (2): 212–41.

Devine, Warren D., Jr. 1983. "From Shafts to Wires: Historical Perspectives on Electrification." *Journal of Economic History* 43 (2): 347–72.

Field, Alexander J. 2018. "Manufacturing Productivity and US Economic Growth." In *Oxford Handbook of American Economic History*, edited by Louis Cain and Paul Rhode, 213–34. Oxford University Press.

Federal Power Commission. 1967. *Steam-Electric Plant Construction Cost and Annual Production Expenses: 1966.* US Government Printing Office.

Friedlingstein, Pierre, Michael O'Sullivan, Matthew W. Jones, Robbie M. Andrew, Judith Hauck, Peter Landschützer, Corinne Le Quéré, et al. 2025. "Global Carbon Budget 2024." *Earth System Science Data* 17 (3): 965–1039.

Fuchs, Victor R. 1980. "Economic Growth and the Rise of Service Employment." NBER Working Paper 0486.

Gordon, Robert J. 2017. *The Rise and Fall of American Frowth: The US Standard of Living since the Civil War.* Princeton University Press.

Hales, Jeremy M. 1976. *Tall Stacks and the Atmospheric Environment.* US Environmental Protection Agency.

Hamilton, James D. 2013. "Historical Oil Shocks." In *Routledge Handbook of Major Events in Economic History*, edited by Randall Parker and Robert Whaples, 239–65. Routledge.

Hausman, William J., and John L. Neufeld. 2002. "The Market for Capital and the Origins of State Regulation of Electric Utilities in the United States." *Journal of Economic History* 62 (4): 1050–73.

Henriques, Sofia Teives, and Karol J. Borowiecki. 2017. "The Drivers of Long-Run CO2 Emissions in Europe, North America and Japan since 1800." *Energy Policy* 101: 537–49.

IEA. 2024. *Global EV Outlook 2024.* International Energy Association.

IEA. 2025a. "Coal." International Energy Association. https://www.iea.org/energy-system/fossil-fuels/coal.

IEA. 2025b. *Energy and AI.* International Energy Association.

Jarvis, Stephen, Olivier Deschenes, and Akshaya Jha. 2022. "The Private and External Costs of Germany's Nuclear Phase-Out." *Journal of the European Economic Association* 20 (3): 1311–46.

Jorgenson, Dale W. 1984. "The Role of Energy in Productivity Growth." *Energy Journal* 5 (3): 11–26.

Joskow, Paul L. 1985. "Vertical Integration and Long-Term Contracts: The Case of Coalburning Electric Generating Plants." *Journal of Law, Economics, and Organization* 1 (1): 33–80.

Joskow, Paul L. 2006. "The Future of Nuclear Power in the United States: Economic and Regulatory Challenges." MIT Center for Energy and Environmental Policy Research Paper 06-019.

Joskow, Paul L., and Nancy L. Rose. 1985. "The Effects of Technological Change, Experience, and Environmental Regulation on the Construction Cost of Coal-Burning Generating Units." *RAND Journal of Economics* 16 (1): 1–27.

Knittel, Christopher R. 2006. "The Adoption of State Electricity Regulation: The Role of Interest Groups." *Journal of Economic History* 54 (2): 201–22.

Lafortune, Jeanne, Ethan G. Lewis, José Pablo Martínez, and José Tessada. 2021. "Changing Returns to Scale in Manufacturing 1880–1930: The Rise of (Skilled) Labor?" NBER Working Paper 28633.

Lebergott, Stanley. 1976. *The American Economy: Income, Wealth and Want.* Princeton University Press.

Lovell, Alfred H. 1941. *Generating Stations: Economic Elements of Electrical Design.* 3rd ed. McGraw-Hill.

Lublin, Edward L. 1981. "The Future of the Department of Energy's Coal Conversion Program." *Energy Law Journal* 2: 355–68.

Meltz, Robert. 1975. "The ESECA Coal Conversion Program: Saving Oil the Hard Way." *Environmental Law Reporter* 5: 50146.

Metcalf, Gilbert E. 2008. "An Empirical Analysis of Energy Intensity and Its Determinants at the State Level." *Energy Journal* 29 (3): 1–26.

Muller, Nicholas Z. 2022. "War, Influenza, and U.S. Carbon Intensity." NBER Working Paper 30522.

Painter, David S. 2012. "Oil and the American Century." *Journal of American History* 99 (1): 24–39.

Ritchie, Hannah, Pablo Rosado, and Max Roser. 2023. "CO2 and Greenhouse Gas Emissions." Our World in Data. https://ourworldindata.org/co2-and-greenhouse-gas-emissions.

Rosenberg, Nathan. 1998. "The Role of Electricity in Industrial Development." *Energy Journal* 19 (2): 7–24.

Schurr, Sam H., and Bruce C. Netschert. 1960. *Energy in the American Economy, 1850–1975: An Economic History of Its History and Prospects.* Johns Hopkins University Press.

Seiferlein, Katherine E. 2007. *Annual Energy Review 2006.* Energy Information Administration.

Severnini, Edson. 2017. "Impacts of Nuclear Plant Shutdown on Coal-Fired Power Generation and Infant Health in the Tennessee Valley in the 1980s." *Nature Energy* 2 (4): 17051.

Shehabi, Arman, Sarah J. Smith, Alex Hubbard, Alex Newkirk, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet, and Dale Sartor. 2024. *2024 United States Data Center Energy Usage Report.* Lawrence Berkeley National Laboratory.

Suh, Sangwon. 2006. "Are Services Better for Climate Change?" *Environmental Science and Technology* 40 (21): 6555–60.

Tol, Richard S. J., Stephen W. Pacala, and Robert H. Socolow. 2009. "Understanding Long-Term Energy Use and Carbon Dioxide Emissions in the USA." *Journal of Policy Modeling* 31 (3): 425–45.

US Census Bureau. 1976. *Historical Statistics of the United States, Colonial Times to Present.* US Census Bureau.

US Energy Information Administration. 2024a. "Carbon Dioxide Emissions Coefficients." Energy Information Administration, September 18. https://www.eia.gov/environment/emissions/co2_vol_mass.php.

US Energy Information Administration. 2024b. *Monthly Energy Review August 2024.* US Department of Energy.

US Energy Information Administration. 2025a. *Assumptions to the Annual Energy Outlook 2025: Coal Market Module.* US Department of Energy.

US Energy Information Administration. 2025b. *Carbon Dioxide Emissions from Energy Consumption in the United States from 1975 to 2024, by Sector.* Distributed by Statista. https://www.statista.com/statistics/204879/us-carbon-dioxide-emissions-since-1975-by-sector/ (accessed April 25, 2025).

US Geological Survey. Various Years. Minerals Yearbook. US Government Printing Office.

US Bureau of the Census. 1960. *Historical Statistics of the United States, Colonial Times to 1957.* US Government Printing Office.

Waxman, Andrew R., H. R. Huber-Rodriguez, and Sheila M. Olmstead. 2024. "What Are the Likely Air Pollution Impacts of Carbon Capture and Storage?" *Journal of the Association of Environmental and Resource Economists* 11 (S1): S111–55.

Wright, Gavin. 1990. "The Origins of American Industrial Success, 1879–1940." *American Economic Review* 80 (4): 651–68.

Yuskavage, Robert E., and Mahnaz Fahim-Nader. 2005. "Gross Domestic Product by Industry for 1947–86." *Survey of Current Business*: 70–84.

Text as Data in Economic Analysis

Tarek A. Hassan, Stephan Hollander, Aakash Kalyani, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun

Advancements in information technology, coupled with the rapid growth of digitally accessible text, have greatly enhanced economists' ability to use text as data in their research (Baker, Bloom, and Davis 2016; Gentzkow, Kelly, and Taddy 2019). Today, the range of such text sources is virtually limitless, including books, contracts, transcripts, news articles, social media feeds, and many more. Of course, economists have long engaged in analyzing text, as evidenced, for example, by Romer and Romer (1989, 2023), who analyzed historical transcripts of Federal Reserve policymaking meetings. This article focuses on the technological advancements that have transformed this practice. We explore how these innovations enable faster, larger-scale, and more consistent processing of text, while also facilitating the measurement of multiple textual dimensions.

■ *Tarek A. Hassan is Professor of Economics, Boston University, and Cofounder and Chief Executive Officer, NL Analytics Inc., both in Boston, Massachusetts. Stephan Hollander is Professor of Financial Accounting, Tilburg University, Tilburg, the Netherlands. Aakash Kalyani is Research Economist, Federal Reserve Bank of St. Louis, St. Louis, Missouri. Laurence van Lent is Professor of Accounting and Economics, Frankfurt School of Finance and Management, Frankfurt, Germany. Markus Schwedeler is cofounder and Chief Executive Officer, NL Analytics Inc., Boston, Massachusetts. Ahmed Tahoun is Professor of Accounting at London Business School, London, United Kingdom. Hassan, Kalyani, Schwedeler, and Tahoun are cofounders of NL Analytics Inc., a startup that provides tools for earnings-call analysis to researchers. Their respective e-mail addresses are thassan@bu.edu, s.hollander@tilburguniversity.edu, aakash.kalyani@stls.frb.org, l.vanlent@fs.de, markus@nlanalytics.tech, and atahoun@london.edu.*

Specifically, we focus on how analyzing largely unstructured corporate text—produced and utilized daily for decision-making by entrepreneurs and their firms' shareholders, financial analysts, and other stakeholders—can enhance economic research for academic scholars, policymakers, and corporate management, as well as business and economics students. Traditional quantitative economic research has typically relied on survey data or accounting metrics, such as firms' profit, investment, or employment figures. The use of text opens up a wealth of new data, offering insights into expectations, perceived risks, costs, and opportunities that are difficult or impossible to extract from conventional data sources.

To illustrate the power of text analysis, we will focus much of this paper on the example of analyzing text from corporate earnings calls, a particularly active area of current research. These calls provide a real-time window into how companies respond to unfolding events, offering insights into their strategies, concerns, and expectations. To illustrate, consider Russia's large-scale invasion of Ukraine on February 24, 2022. Within hours of this escalation, Geox S.p.A., an Italian shoe manufacturer, held its quarterly earnings call. The call focused on the unfolding crisis, with analysts and investors inquiring about contingency plans, revenue dependencies, and strategic options in light of potential global embargoes, disrupted supply chains, and interrupted interbank communications. Geox was not unique in this regard—management of numerous other corporations engaged in similar discussions during earnings calls in the following days, anticipating many of the impacts economic policymakers had to confront over the coming years, ranging from high natural gas prices to increased defense budgets. Even beyond geopolitical crises, earnings calls routinely bring policy-relevant issues to the surface. Recent examples include ongoing concerns about the dangers of artificial intelligence (AI), the sources of persistent inflation, and trade policy uncertainty.

In this paper, we introduce a set of methods that allow researchers to systematically analyze the content of these earnings calls (and other texts) at scale. These methods enable researchers to measure the risks, costs, and opportunities that firms around the globe attribute to specific shocks, policies, and regulations. Building on these techniques, we show how to trace shock-induced risk transmission and potential contagion across firms, sectors, and regions; and how to measure present and past expectations regarding specific events, such as the 2003 SARS outbreak or the 2017 Tax Cuts and Jobs Act, in this way overcoming the limitations of retrospective surveys that attempt to reconstruct pre-event expectations (Gallemore et al. 2025).[1] Moreover, we demonstrate how to make cardinal comparisons of different risks, showing, for example, that in 2023, US-based firms allocated twice as much attention to inflation-related risks than to risks from geopolitics or trade wars—providing insights that are difficult or impossible to glean from other data.

---

[1] One possible reason why measuring past expectations related to historical events is important is provided by Hennessy and Strebulaev (2020). They demonstrate analytically that if firms' behavior is influenced by expectations about a policy, neglecting to account for these expectations in empirical analyses can distort inferences about the causal effects of that policy.

Importantly, the transcripts of these earnings calls are widely available almost immediately, allowing quantitative analysis without the months-long lags typical of accounting or survey data.

Beyond the specific example of earnings calls, we argue that analyzing the full range of corporate text data—including job postings and patents—represents a sea change in our ability to track and understand economic shocks. At a narrow level, discerning why and how firms respond to proposed reforms or technological changes can help shape more effective policy design and facilitate smoother adaptation to economic shocks. On a broader scale, tracing the propagation of supply, demand, and other economic shocks across firms, sectors, and countries is key for testing and refining theories of how the economy functions.

In the subsequent sections, we begin by introducing earnings conference call transcripts as a data source, then outline the building blocks of our approach to creating measures of firms' exposure to specific shocks using these data. We discuss the options available to researchers in text classification tasks, including keyword-based approaches, the use of "training libraries," and large language models (LLMs). We then introduce and discuss straightforward techniques that enable the measurement of the risks, costs, and opportunities that decision-makers associate with specific shocks and policies, emphasizing the advantages they provide in terms of cardinal comparisons of risk, differentiating between the effects of bad news versus increased risks, and analyzing the transmission of shocks across countries and firms. Finally, we show how to expand these techniques to other corporate text sources, including patents and job postings. We demonstrate how to use these sources to trace the origin and spread of technological innovations across geography, industries, firms, and occupations.

In sum, this article offers a glimpse of how scholars, central banks, think tanks, and international organizations are already integrating text-analysis techniques for a range of purposes in their research and policymaking. Throughout, we contend that the main frontier of text as data in economic analysis is not in developing more advanced text-analysis techniques, but in making vast, untapped text corpora accessible to scholars, economists, and policymakers (Bae, Hung, and van Lent 2023).

Earnings Conference Calls: A Marketplace of Information

Corporate communication has changed substantially over the past decades. Before 2000, publicly listed firms primarily communicated through accounting numbers. However, these accounting numbers are inherently backward-looking and slow to reflect new information: they document past events, take months to prepare, and incorporate changes into earnings figures only gradually due to accounting principles that prioritize caution. For this reason, firms' management would typically provide more information and context for their accounting disclosures in private conversations with selected journalists, analysts, and investors.

In 2000, the Securities and Exchange Commission (SEC) made this practice illegal and instead required publicly traded companies to share any such context with the general public. While earnings conference calls—held each quarter after the release of a firm's financial results—were practiced prior to this regulation, the new requirements for simultaneous public disclosure accelerated their widespread adoption for discussing a firm's key financial information. This shift made earnings calls the primary channel through which market participants—institutional investors, financial analysts, and individual shareholders alike—could access management's interpretation of their firms' performance.

During a typical corporate earnings call, an executive of the firm (often the chief executive officer or the chief financial officer) reviews the firm's recent financial performance, discusses operational challenges, and outlines any issues that impact the firm's future prospects. Financial analysts then pose questions during a question-and-answer session, prompting further discussion of the company's strategies. These conversations last up to an hour. They are livestreamed online, and then later transcribed and made available. Primarily due to this interactive nature, earnings calls offer substantial complementary textual narratives for nearly every set of accounting numbers available through databases like S&P's Compustat, and provide valuable context and additional information about firms' current and future situations.[2]

Economists have historically been skeptical of using information on what economic decision makers *say*, favoring instead evidence from what they *do* (choices and behaviors), because of concerns about "cheap talk"—the perception that decision makers can spin information or outright lie to gain an economic advantage. Notwithstanding this, we argue that earnings calls are a valuable venue where what participants say is highly informative about economic behavior. Indeed, unlike casual statements or survey responses, communications during earnings calls carry real consequences. For instance, securities law enforces significant penalties for making materially misleading statements, while the scrutiny of sophisticated analysts, who can challenge management's claims, generates reputational costs for evasive or inconsistent responses (Hollander, Pronk, and Roelofsen 2010; Rogers, Van Buskirk, and Zechman 2011). Perhaps most importantly, market participants have the option (and are known) to follow up and verify statements made during earnings calls by comparing them to the firm's subsequent behavior (for example, Suijs 2007).

These institutional features distinguish earnings calls from other contexts where talk may be cheap. Understanding this distinction helps in determining when text data will likely offer reliable economic insights. In general, corporate communication is most informative when statements involve explicit or implicit costs and can

---

[2] Even when accounting numbers are available, they provide limited insight into how or why shocks affect firms. For example, when profits drop due to lower sales, identifying the underlying causes—such as supply chain disruptions, demand contractions, or policy interventions—necessitates additional context. Earnings conference calls provide an opportunity to obtain this additional information by questioning management.

be cross-validated against observable outcomes, and when institutional structures create incentives for truthful communication.

Many international firms, listed on a stock exchange, either conduct their earnings calls in English or provide translations, enabling rapid assembly of directly comparable corpora across more than 12,500 firms from nearly 90 countries. Primary sources for earnings-call transcripts include corporate investor relations websites and public financial platforms like Yahoo! Finance and Seeking Alpha. Commercial data providers, including the London Stock Exchange Group (LSEG) and S&P, standardize and aggregate these transcripts, making them available through institutional subscriptions commonly held by universities. The earnings call data used in this paper come from LSEG through NL Analytics, a startup founded by some of the authors of this paper, which offers online tools for earnings call analysis to researchers, central banks, and international organizations.

Decoding Shocks: Exposure, Risk, and Sentiment

Exposure

The time people spend discussing a topic reveals its importance to them—a simple yet powerful insight that opens new possibilities for measuring the importance of different topics in corporate communication. This observation becomes particularly valuable in the context of earnings conference calls, where executives have less than an hour to address their most pressing concerns (Matsumoto, Pronk, and Roelofsen 2011). The time constraint creates a natural prioritization: topics that receive more attention in these discussions are, by necessity, those that executives and noncorporate participants consider most significant.

This relation between discussion time and importance allows us to calculate a precise measure of firms' exposure to specific events or issues, which we call "topic exposure." Specifically, by calculating the proportion of sentences devoted to a particular topic during an earnings call, we can quantify its relative significance to the firm. For example, counting sentences about Russia's invasion of Ukraine as a percentage of total sentences reveals which firms face meaningful exposure to the conflict and how this exposure compares to other concerns. Extensive audits with human coders strongly support the validity of this classification approach (Hassan et al. 2019, 2023, 2024a, 2024b).

Risk

One application of this idea is measuring the overall risk a given firm faces at a particular moment by calculating the proportion of the earnings-call discussion devoted to risk or uncertainty. While economists define "risk" as the variance of outcomes, participants in earnings calls may use the term more loosely, frequently focusing on adverse scenarios ("downside risk"). To bridge this gap, we construct a comprehensive dictionary of risk-related terms by identifying single-word synonyms of "risk," "risky," "uncertain," and "uncertainty" from the Oxford English

Dictionary, excluding only the terms "question" and "questions," which of course carry a different meaning in earnings calls. This systematic approach yields terms across a spectrum, from those suggesting volatility ("variable," "variability"), to those indicating upside and downside potential ("chance," "prospect," "pending"). This breadth of terms ensures we capture the full range of risks firms face, not just their concerns about adverse scenarios.[3]

Applying this dictionary to analyze earnings-call discussions yields a firm-level measure of "overall risk": the percentage of sentences in a given firm's earnings call that mention a synonym for risk or uncertainty. Figure 1, panel A, shows the average of overall risk across all firms with a US stock market listing. This measure of economy-wide risk reveals distinct trends. During the "Great Moderation," from 2002 to 2007, risk mentions remained relatively subdued, with fewer than 1.6 percent of sentences containing risk-related terms (Benati and Surico 2009; Galí and Gambetti 2009). The global financial crisis, however, marked a sharp departure from this pattern, with risk discussions peaking at 2.2 percent of earnings call content in late 2008. Subsequent spikes emerged during the European sovereign debt crisis of 2011–2012, and risk-related discussions peaked at 2.8 percent during the Covid-19 pandemic.

At this aggregate level, the pattern of risk-related discussions in earnings calls correlates strongly with market-based measures of risk such as the Chicago Board Options Exchange volatility index, known as the VIX. (The VIX reflects the likelihood of large price swings in the S&P 500 index over the next 30 days, based on prices of stock options.) However, compared to the VIX and other aggregate market-based indicators, a text-based measure offers a distinct advantage: straightforward disaggregation. Indeed, starting with the firm-level risk-related discussions, we can construct detailed, disaggregated metrics by sector and region (who is perceiving risk), and even analyze risk patterns for specific topics (what is the risk about).

To measure "topic risk," we calculate the proportion of earnings-call sentences that mention both risk-related terms and a specific subject, relative to the total number of sentences in the call. Figure 1, panel B, illustrates this topic-specific approach by tracing Brexit-related risks discussed by firms around the world (Hassan et al. 2024a). The time-series pattern shows notable spikes in "Brexit risk" that correspond with key political developments during that period. To illustrate, discussions of Brexit-related risk surged immediately after the 2016 UK referendum vote to leave the European Union. While this initial uncertainty subsided somewhat in 2017, it intensified sharply in the second half of 2018 as the details of Prime Minister Theresa May's agreement with the European Union became clear, along with the challenges of securing parliamentary approval. The prospect of a no-deal Brexit under the subsequent Prime Minister, Boris Johnson, kept uncertainty elevated through 2019, until Brexit finally took effect on January 31, 2020.

---

[3] In comparison, Baker, Bloom, and Davis (2016) construct their measure of economic policy uncertainty by including only the terms "uncertain" and "uncertainty," which focuses their analysis on uncertainty rather than broader notions of risk.

Figure 1
Text-Based Risk
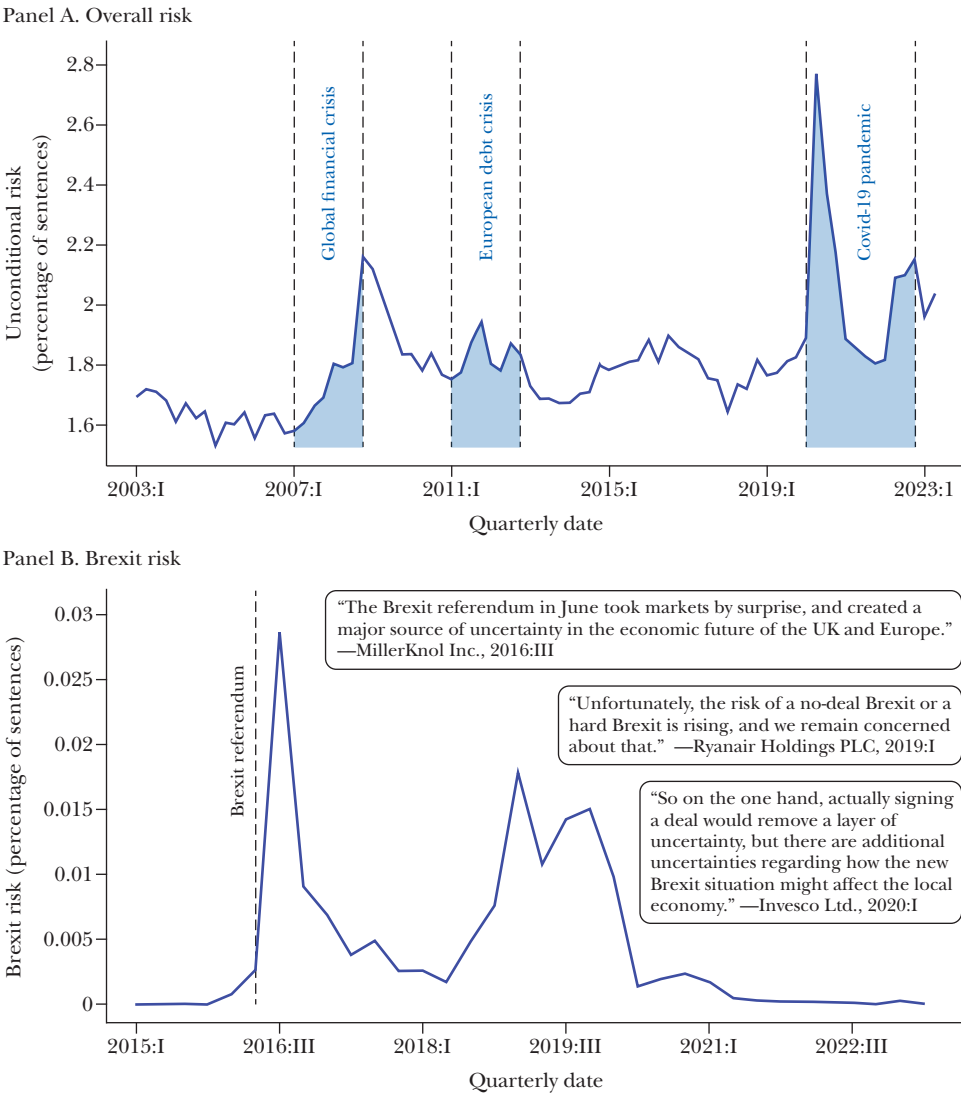
Panel A. Overall risk

Panel B. Brexit risk

Note: This figure displays overall risk (panel A) and Brexit risk (panel B) as a percentage of sentences in earnings calls. The sample comprises 379,227 earnings conference calls, representing 12,805 firms headquartered in 89 countries.

By intersecting discussions of risk with discussions of specific topics, we can thus assess how much risk a given firm associates with a given specific topic (such as Brexit).

The figure highlights another key advantage of this text-based measurement approach: its capacity for self-labeling. In other words, the reasons behind observed fluctuations in text-based risk measures become apparent by examining the underlying discussions from the transcripts that contribute to each spike. As the figure illustrates, excerpts from high-scoring earnings-call transcripts enable us to directly link periods of elevated Brexit risk to specific political events and their anticipated economic impacts.

Sentiment

Measuring sentiment from text involves identifying terms that indicate how positively or negatively firms discuss specific topics. The most straightforward method to flag earnings-call sentences as positive is if they contain words like "good," "strong," or "great" without negative terms, and negative if they include words like "loss," "decline," or "difficult" without positive ones. Loughran and McDonald (2011) provide a comprehensive sentiment lexicon specifically calibrated for financial text.

Recent advances in natural language processing allow for understanding words in their full context without relying on specialized dictionaries. However, systematically identifying positive- and negative-toned discussions naturally complements the earlier discussion on measuring topics and risk through word frequency. Indeed, similar to how a simple count of topic-specific terms reveals firm exposure and risk-related terms capture uncertainty, tracking sentiment-related words indicates whether firms perceive economic shocks as having a positive or negative impact. Following this logic, "topic sentiment" counts sentences that mention a given topic with positive sentiment, subtracts the count of those with negative sentiment, and normalizes by the total number of sentences in the call.

Figure 2 shows how topic-sentiment measures capture firms' responses to major economic shocks. Specifically, panel A compares average sentiment across several significant events: Russia's 2022 invasion of Ukraine, the Covid-19 outbreak in 2020, Brexit-related developments in 2016, and breakthroughs in artificial intelligence, particularly the release of large language models, in 2023. The measures indicate systematically negative sentiment in discussions related to Brexit, Covid-19, and Russia's invasion, while discussions related to AI trend positive, with 42.1 percent of sentences related to AI containing positive terms, compared to 10.2 percent negative, yielding a net positive sentiment of +32 percent.

These aggregate patterns mask substantial heterogeneity at the firm, sector, and country levels. Panel B shows firm-level variation in discussions related to artificial intelligence. Technology firms heavily invested in development of AI, such as Alphabet, NVIDIA, and Microsoft, express strongly positive sentiment. In contrast, content producers like News Corp, Universal Music Group, and Warner Music Group discuss AI with notably negative sentiment, reflecting concerns about intellectual property protection.[4]

---

[4] For additional examples of AI-related excerpts from earnings-call conversations, see Supplemental Appendix Table 1.
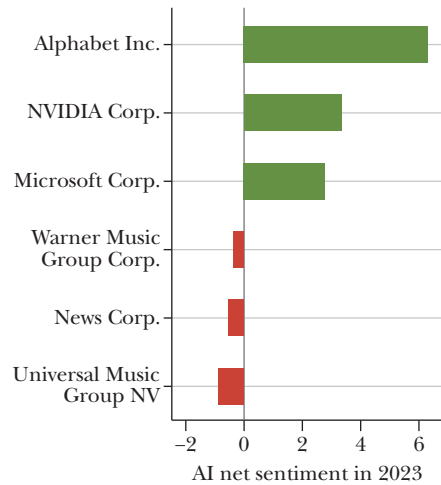
Figure 2
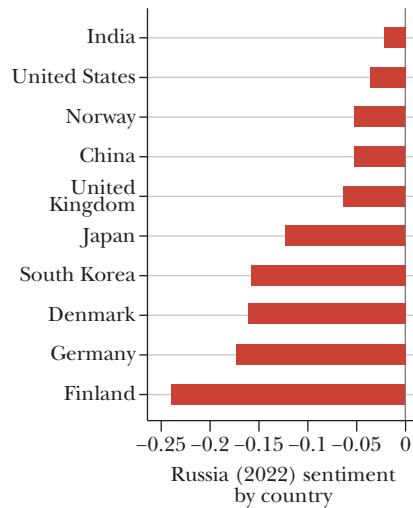**Text-Based Sentiment**

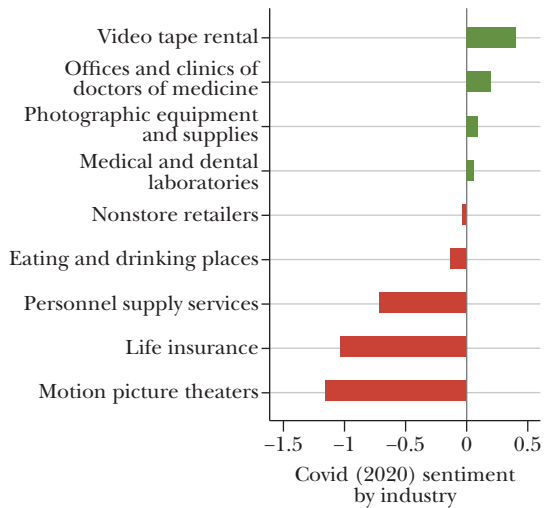Panel A. Sentiment by event
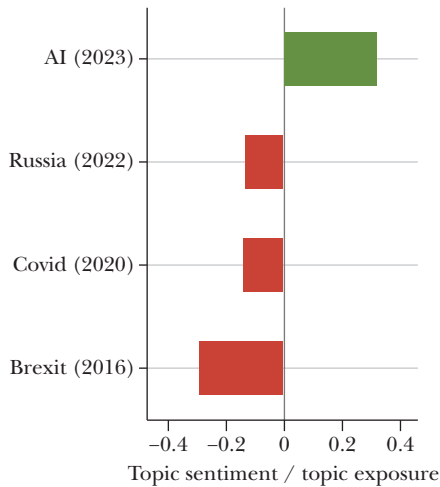
Panel B. AI (2023) sentiment by firm

Panel C. Russia (2022) sentiment by country

Panel D. Covid (2020) sentiment by industry

Source: Data sourced from www.nlanalytics.tech.
Note: This figure illustrates topic sentiment as follows: panel A shows topic-specific (namely, Russia (2022), Covid-19 (2020), Brexit (2016), and AI (2023)) sentiment divided by topic-specific exposure; panel B shows AI (2023) sentiment, as a percentage of sentences, by firm; panel C shows Russia (2022) sentiment, as a percentage of sentences, by country; and panel D shows Covid-19 (2020) sentiment, as a percentage of sentences, by three-digit SIC sector. To calculate topic-specific sentiment on the aggregate or by firm, country, or industry, we take an average across earnings calls of topic sentiment where comprises of topic-specific keywords. To illustrate, the topic Russia includes keywords that relate to the country, such as "Russia," "Russian," "Ruble," and "Mr Putin." Similarly, for the topic of AI, in addition to "AI" and "artificial intelligence" we included terms like "ChatGPT," "LLM," "deep learning," and "machine learning."

Panels C and D display variation in sentiment across geographic regions and sectors, respectively. Panel C shows that Finnish, German, and Danish firms expressed particularly negative sentiment about Russia-related developments in 2022, with discussions centered on regional security, energy markets, and economic relationships. The transcripts show Japanese and South Korean firms focusing their concerns specifically on trade access through Russian territory.

Even during the same economic shock, sentiment varies substantially across sectors. Panel D shows that early in the Covid-19 pandemic, the Video Tape Rental and Medical Services sectors anticipated benefits, reflecting expected increases in streaming and healthcare demand. In contrast, Life Insurance and Motion Picture Theater sectors projected strongly negative impacts, anticipating increased mortality-related payouts and restricted public gatherings. Collectively, these findings suggest that topic-sentiment analysis is a powerful tool, enabling us to characterize the heterogeneous impact of economic shocks at the firm, sector, and country levels.

One important caveat is that earnings-call language has become more positively toned over time, as executives learned to steer clear of words that might trigger algorithmic trades (Cao et al. 2023). This upward drift inflates raw sentiment, but it is straightforward to purge with time fixed effects.

Lexicon Construction

Each of our measures of topic-specific exposure, risk, and sentiment involves a systematic procedure for identifying sentences that mention a particular topic. For this task, researchers have three options: keyword-based searches, training libraries, and machine learning. These options differ in the degree to which they depend on human judgement, but they all share a common goal: systematically capturing all relevant content (that is, minimizing false negatives), while avoiding spurious mentions (that is, minimizing false positives).

Keyword-Based Approach

The simplest approach to identifying discussions on a given topic is by selecting relevant keywords (Baker, Bloom, and Davis 2016). In our experience, this approach is particularly effective when the topic of interest pertains to a specific, narrowly defined event or policy. For instance, in Hassan et al. (2024a), the single keyword "Brexit" sufficed to accurately capture conversations related to the United Kingdom's exit from the European Union. Similarly effective examples include Covid-19, supply chain disruptions, and cyber risk, where a handful of keywords suffice.

The process of choosing keywords usually begins with a set of annotated sentences obtained from human readers, flagging those that discuss the topic of interest.[5] Next, we select keywords or phrases that mimic this human classification

---

[5] This systematic approach requires the topic to have existed long enough for researchers to annotate a sufficiently large number of sentences, which in some cases can present challenges. In the case of

within the sample of previously annotated text. These keywords or phrases are frequent in sentences discussing the topic of interest and infrequent in those that do not. To check their effectiveness out-of-sample, we sample new sentences from our corpus and assess the rate of false positives. This process may prompt changes to our set of keywords, followed by further sampling and iteration until we identify a set that performs well on our corpus. For example, to identify supply chain-related discussions in Hassan et al. (2023), we started with "supplier" and "supply" and, thus, iteratively built a comprehensive set of keywords.

Training Libraries

For a broader or less precisely defined event or policy, a keyword-based approach can become unwieldy. In such cases, relying on "training libraries" rather than human judgement can be a more effective solution. A straightforward approach in this context is to use the numerical statistic known as term frequency–inverse document frequency (Manning, Raghavan, and Schütze 2008; Song and Wu 2008). We apply this approach in Hassan et al. (2019, 2024b) to differentiate between political and nonpolitical risk discussions and to assess the risks firms associate with specific countries.

This approach requires a sufficiently large corpus of text representative of discussions on the given topic. For example, to differentiate between political and nonpolitical discussions, we use two textbooks: one focused on politics and the other on a nonpolitical subject (in our context, we opted for an accounting textbook). From these sources, we extract and tally the occurrences of bigrams (that is, two-word combinations), identifying those that are frequent in one training library but not in the other. For example, the bigrams "Joe Biden" and "Congress passed" might be frequent in discussions related to politics, but never or rarely occur in discussions on nonpolitical subjects. This typically results in a list of several thousands of topic-related bigrams, which are assigned weights proportional to their frequency in the training library.

Compared to a keyword-based approach, training libraries generally capture a much broader range of language used to discuss a topic, allowing for the detection of more subtle and indirect references to the topic in question.

Machine Learning

The third approach leverages machine learning to aid in keyword selection. In this method, the common practice is for a machine learning model to propose keywords to the researcher, who then decides whether to accept or reject them based on how well they align with their definition of the topic in question. For example, Sautner et al. (2023) utilize training libraries to train a machine learning model to distinguish between climate- and non-climate-related discussions in earnings calls.

Covid-19, as studied in Hassan et al. (2023), no such training data were available at the onset of the outbreak. As a result, we relied on a list of Covid-19 synonyms provided by the World Health Organization, supplemented by cross-referencing contemporaneous newspaper articles and online resources.

The model then proposes phrases that most effectively make this distinction. For each suggestion, they randomly sample from the text of earnings calls that mention the proposed phrase and accept it if the false positive rate falls below a certain threshold.

In contrast, Kalyani et. al. (2025a) uses a large language model to make suggestions. To identify job postings for research and development positions, the authors start with a set of seed phrases (for example, "research," "and develop") and let the LLM suggest additional phrases commonly used in similar contexts. They then randomly sample excerpts from job postings for each suggestion to check for false positives and include those suggested phrases that meet the threshold for true positives.

A more evolved version of this approach completely delegates the keyword selection-rejection decision to a large language model. In this approach, the researcher provides a verbal description of the topic and allows the LLM to determine whether a given sentence or paragraph from an earnings-call transcript aligns with this definition (Jha et al. 2024). The drawback of this fully automated approach, in our view, is that the researcher effectively relinquishes control over the definition of the topic and hands it over to the LLM. This can lead to issues if not handled carefully. One issue is replicability—different LLMs may make different selection-rejection decisions for reasons that are difficult, if not impossible, to interpret afterward. Moreover, particularly when dealing with historical data, LLMs are prone to hindsight bias. For example, an LLM might link "Fukushima" to a nuclear disaster, even when analyzing text from the time before the disaster occurred (Sarkar and Vafa 2024).[6]

In conclusion, depending on the specific application, a researcher should thoughtfully choose one of the approaches outlined above to build a lexicon. Our experience suggests the virtues of erring on the side of simplicity and opting for a more straightforward approach whenever feasible. Sample code illustrating each approach is available in Hassan et al. (2025b), which can be found at the landing page for this article at the American Economic Association's website.

Text Analysis Applications in Economic Research

We next present three applications that show how our text-based measures of exposure, risk, and sentiment allow us to track the impact of economic shocks in novel ways that would be hard or impossible to achieve with conventional data. These applications demonstrate how to decompose the relative importance of different risks, separate direct effects from uncertainty impacts, and track the transmission of shocks across global markets.

---

[6] While current large language models offer promising capabilities for keyword selection, their relative advantage over lower-tech methods may evolve as both technologies advance. Future improvements in model efficiency, interpretability, and cost-effectiveness could shift the balance, making it important to regularly reassess their value in comparison to traditional approaches.

Figure 3
Risk Decomposition
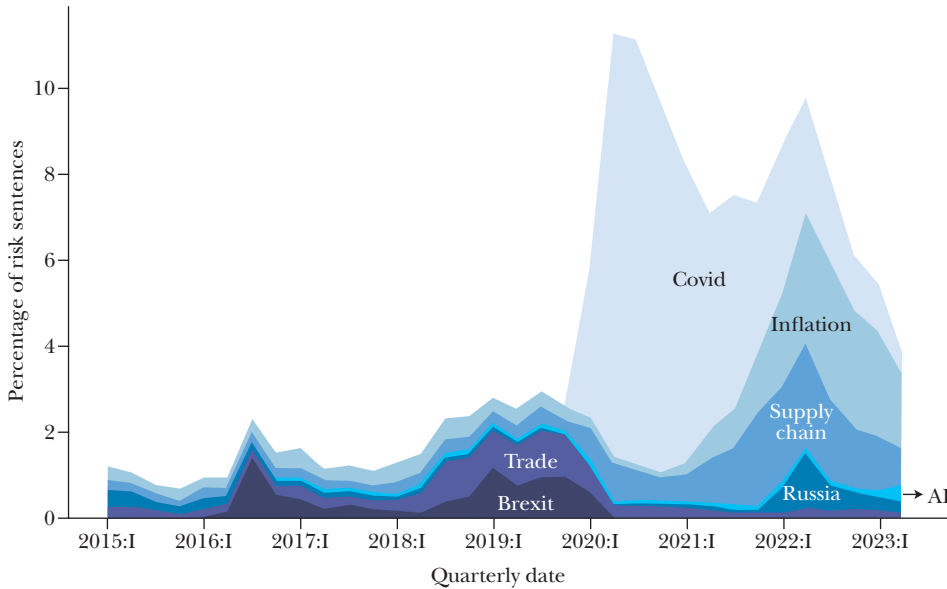
Source: Data sourced from www.nlanalytics.tech.
Note: This figure shows, as a percentage of sentences, a breakdown of risk related to AI, Brexit, Covid-19, inflation, Russia, the supply chain, and trade for our sample of US-listed firms.

Cardinal Risk Decomposition

Our first application illustrates how text analysis facilitates a cardinal comparison of different sources of risk. This approach provides a more precise understanding of how each specific risk contributes to a firm's overall risk profile and can help policymakers and analysts identify which risks are most pressing.

To fix ideas, note from Figure 1 that in 2016:III, firms devoted 1.90 percent of all sentences to discussing risks and 0.029 percent specifically to discussing Brexit risk. We can thus conclude that Brexit accounted for 1.53 percent of overall risk-related discussions in that quarter. This decomposition of overall risk into specific topics is an important addition to the literature. The closest analog is the attribution of excess returns to risk factors in empirical asset pricing. However, these methods generally do not allow an assessment of which specific shock contributes to the overall uncertainty faced by a firm at a given time, even when paired with an event study.

Figure 3 applies this decomposition more broadly to track the evolution of different sources of risk over time. The data reveal that trade policy uncertainty reached its peak during the trade war in the first term of the Trump administration, representing 0.98 percent of risk mentions in 2018:III—comparable to Brexit-related concerns. The Covid-19 crisis subsequently emerged as the dominant source of risk, accounting for 9.88 percent of the risk-related discussions in 2020:II.

As the pandemic's direct effects subsided, firms' focus shifted toward inflation and supply chain disruptions. Although Russia's invasion of Ukraine garnered substantial media coverage, its effect on firm-level risk discussions was generally limited, with notable exceptions in certain regions. For example, German firms devoted 4.02 percent of their risk discussions to Russia in 2022:I.

This cardinal decomposition provides two key advantages. First, it enables a systematic comparison of the relative importance of different (sources of) risks at any given time, helping policymakers understand which forces are the primary drivers of risk for specific (groups of) firms. Second, the ability to track the evolution of these proportions over time allows for the analysis of how the importance of different risk sources rises and falls. The consistent measurement framework allows us to make meaningful statements about how firms' concerns with one source of risk compare to another, both at a given point in time (that is, cross-sectionally) and over time, even as the overall volume of risk discussion fluctuates.

Distinguish Bad News from Uncertainty Effects

Our second application demonstrates how to separate the effects of "bad news" from those of uncertainty in firm-level data. To illustrate the point, consider the management of a music company learning about advancements in artificial intelligence capable of imitating artists' voices. These advancements represent both bad news (that is, a potential decline in the company's expected future cash flows due to competition from content generated by AI) and increased uncertainty (for example, from the possibility of legal action or changes in legislation to protect copyright). If we observe a subsequent decline in this music company's hiring or planned investments, it raises a question: Is this decrease driven by the anticipated negative impact of the AI advancements on its business, or is it the result of increased uncertainty regarding future legislation or court rulings? Separating both effects in the data is typically challenging, as negative news about a firm's prospects often coincides with increased uncertainty about its future (Berger, Dew-Becker, and Giglio 2020).

With our text-based measures of AI sentiment and AI risk, we can readily separate between the two. For example, Supplemental Appendix Table 1 shows an excerpt from one of Chegg's earnings calls, in which the education technology company stresses the negative impact of ChatGPT's launch: students were abandoning the firm's subscription products for it. For Chegg, the boom in AI is unambiguously bad news, yielding a strongly negative measure of AI sentiment. In contrast, Spotify and News Corp emphasized the uncertainty effects in their earnings calls, particularly focusing on the future intellectual property landscape and potential regulatory changes—yielding higher levels of AI risk.

Suppose a researcher wants to examine how firms adjust their investment decisions in response to the artificial intelligence boom. The dependent variable is the investment rate for a given firm in a given quarter. Textual classification of firms' earnings-call transcripts provides three key independent variables: a measure of overall risk (the percentage of sentences discussing risk, irrespective of AI), topic-specific risk measures for AI, and AI sentiment expressed in discussions of AI.

Table 1

Risk, Sentiment, and Investment Rate

	Investment Rate$_{i,t}$ $\times$ 100		
	(1)	(2)	(3)
Unconditional Risk$_{i,t}$	−0.039	−0.039	−0.039
	(0.006)	(0.006)	(0.006)
AI Risk$_{i,t}$		−0.333	
		(0.125)	
AI Positive Sentiment$_{i,t}$		0.073	
		(0.041)	
AI Negative Sentiment$_{i,t}$		−0.002	
		(0.077)	
Brexit Risk$_{i,t}$ pre-2019			−0.298
			(0.136)
Brexit Risk$_{i,t}$ post-2019			0.119
			(0.326)
Brexit Negative Sentiment$_{i,t}$ pre-2019			0.095
			(0.118)
Brexit Negative Sentiment$_{i,t}$ post-2019			−0.256
			(0.154)
Sector fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Observations	84,500	84,500	84,500

Note: This table reports estimation results from a regression of firm *i*'s investment rate at time (quarterly) *t* on text-derived risk and sentiment variables. Investment rate is calculated as the ratio of capital expenditures reported by firm *i* at time *t* to the stock of property, plant, and equipment reported at time *t*−1. All risk and sentiment variables are in units of percentage of the total number of sentences in earnings calls. *Unconditional Risk$_{i,t}$* is the percentage of sentences with risk synonyms in an earnings call hosted by firm *i* at time *t*. *AI (Brexit) Risk* and *Sentiment* is the percentage of sentences with AI (Brexit) keywords mentioned in conjunction with a risk synonym as defined on page 5. *Brexit Risk* pre-2019 (similarly for post-2019) is the interaction of *Brexit Risk* with a dummy equal to one for the period before (post) 2019; zero otherwise. Standard errors are clustered by firm. All specifications include controls for sector (2 digit SIC) fixed effects and time (quarter) fixed effects. Only firms with at least ten earnings calls during 2015–2023 are included.

Regression results for such an analysis can be found in Table 1. As a general pattern, firms tend to reduce investment significantly when faced with higher overall risk, supporting the notion that they prefer to wait for more clarity before committing resources. The response to AI, however, reveals more nuanced patterns. While firms that discuss AI-related risks typically invest less, those expressing positive sentiments about AI's impact on their business significantly increase their investment rates.

The following column of Table 1 repeats the same analysis for Brexit risk and sentiment. Up until 2019, higher Brexit-related risk was associated with significantly lower investment across firms. However, after Brexit's implementation on January 31, 2020, it no longer exhibits a statistically significant relation with investment patterns. In contrast, negative sentiment about Brexit's impact on business operations now emerges as the key explanatory variable: firms that view Brexit as detrimental to their

business reduce investment, while other firms do not. This pattern suggests that, once Brexit's form became clearer, firms' assessment of its future cash flow effects, rather than uncertainty surrounding the policy, guided their investment decisions.

Transmission of Shock-Induced Risks across Countries

Textual analysis can also be used to assess the transmission of economic shocks across the globe. By analyzing how firms from different countries discuss risks related to specific foreign countries in their earnings calls, we can trace the propagation of shocks from one country to another. Figure 4 depicts this by presenting two contrasting examples: discussions of risks relating to the United Kingdom ("UK Risk") during the Brexit crisis in 2016 and discussions of risks relating to the US economy ("US Risk") during the global financial crisis of 2008.

Panel A of Figure 4 shows how discussions on UK-related risk evolved across countries after the Brexit referendum.[7] The *x*-axis represents the average level of UK-related risk discussions among firms in each country in noncrisis times, while the y-axis indicates the level of UK-related risk in the two quarters after the 2016 Brexit referendum. To analyze these data systematically, we focus on three key metrics. First, the "global impact" measure of 1.31 quantifies how much the median country's UK-related risk discussion surpassed expected precrisis average, measured by the median vertical distance from the 45-degree line (that is, the median country increased its UK-related risk discussions by 31 percent.) Second, the steep slope coefficient of 1.51 signifies the strength of "bilateral transmission," indicating that countries where firms previously engaged more in UK-related risk discussions saw significantly greater increases in such discussions. Finally, the high R^2 of 0.86 reflects strong "regularity" in transmission, demonstrating how predictably the increase in risk discussions followed patterns based on preexisting economic relationships.

Together, these three metrics provide insights into the nature of Brexit as an economic shock. The relatively moderate global impact combined with strong bilateral transmission suggests that the shock propagated primarily through existing economic relationships rather than leading to widespread global disruption. The high regularity indicates that firms' precrisis discussions of UK-related risks largely predicted their post-referendum UK–related concerns. This pattern is evident in the data: Irish firms, with significant precrisis exposure, show the highest risk concerns due to their deep integration with the United Kingdom, while firms in France, Germany, and Sweden, with more limited precrisis exposure, exhibit increases in risk in sync with their pre-referendum exposure. That is, the Brexit shock largely impacted firms through established economic channels rather than creating systemic global risk.

Panel B of Figure 4 shows how US-related risk discussions spread during the global financial crisis of 2008. The *x*-axis represents precrisis levels of US-related

---

[7] To identify discussions of specific foreign countries, we use training libraries for each of 45 countries sourced from the Economist Intelligence Unit Country Commerce Reports, in combination with term frequency–inverse document frequency as described above.

Figure 4
Country Risk Transmission

Panel A. Transmission of UK country risk after the Brexit referendum
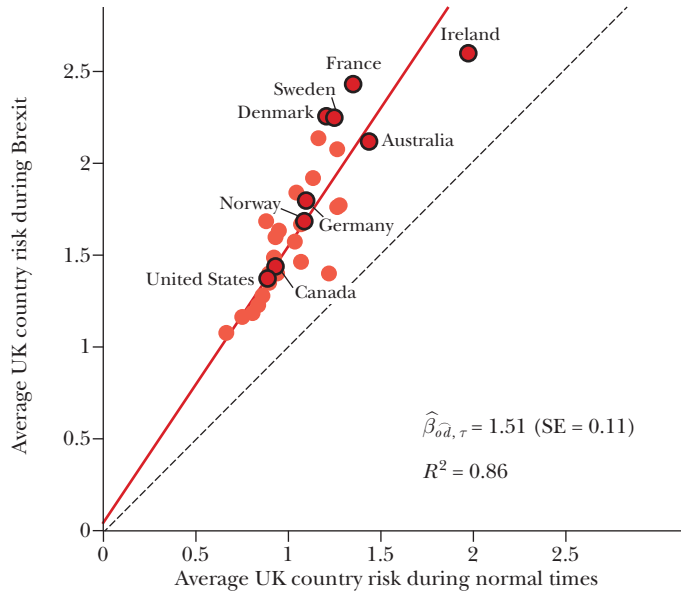

risk discussion among firms in each country, while the *y*-axis indicates the levels during Q1–Q3 of 2008. (We focus on US-related risk because the 2008 financial crisis notably accelerated with the US housing bubble collapse and the subsequent subprime mortgage crisis, eventually leading to a recession in the United States beginning in December 2007.) The pattern differs markedly from the Brexit case shown in panel A. The reduced slope coefficient (0.92) indicates weaker bilateral transmission, meaning that a country's precrisis exposure to US-related risks has little predictive power for its increase in risk discussion during the crisis. The lower R^2 of 0.55 suggests much less regularity in transmission, implying that the crisis spread through channels that precrisis economic relationships were unable to predict.

Despite weaker bilateral transmission, the 2008 financial crisis had a larger global effect than Brexit (2.27), with substantial increases in US-related risk discussion across *all* countries, regardless of their precrisis exposure. This is evident in the data: countries like China, Mexico, and Norway are grouped at higher levels of risk discussion during the crisis, despite differing widely in their precrisis exposure. The more scattered pattern of observations and their overall upward shift, compared to the concentrated pattern seen with Brexit (panel A), suggests that the financial crisis generated systemic global risk concerns that transcended traditional economic relationships.

Figure 4 (continued)
Country Risk Transmission

Panel B. Transmission of US country risk during the global financial crisis

$\widehat{\beta}_{\widehat{od},\tau} = 0.92$ (SE = 0.15)

$R^2 = 0.55$
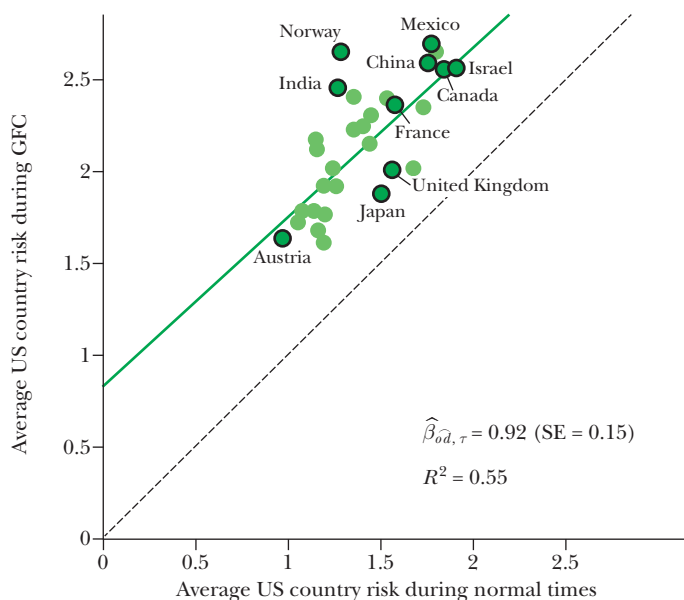
Source: Hassan et. al. (2024b).
Note: Panel A plots each country's average UK risk from earnings calls in the two quarters after the Brexit referendum (June 23) against its average UK risk in all other periods. In panel B, we repeat the same plot for US risk with and without the first three quarters of the global financial crisis: Q1–Q3 of 2008. To identify US-related risk (and similarly for UK risk) we use terms from Hassan et al. (2024b) that identify US-related conversations—such as "United States," "US patent," "USA," and "the DoJ"—using training libraries of archetypical text for multiple different countries. For each panel, we report the coefficient along with standard errors and the R^2.

This contrast between the two crises, shown in panels A and B of Figure 4, illustrates how these simple text-based metrics can capture fundamental differences in shock transmission. While the Brexit shock event primarily affected firms through established dependencies, the financial crisis spread through intricate financial linkages, creating widespread uncertainty that existing economic relationships could not predict.

Text as Data: Lessons and Pitfalls

In this section, we identify three common challenges researchers face when working with text as data and discuss effective strategies for overcoming them.

Preprocessing

Textual analysis starts with preprocessing—cleaning and transforming raw text into a structured format suitable for analysis. The first challenge revolves around preprocessing choices that can affect the quality of measurements. Unlike many other text-analysis applications in computational linguistics, our experience with earnings-call transcripts suggests that invasive preprocessing techniques such as stemming (treating words with the same stem as identical), lemmatizing (standardizing variant forms of a word), and removing punctuation are generally not beneficial. Furthermore, although we usually remove capitalization, we find that preserving it can sometimes be useful—for example, when identifying discussions of different countries, we might want to distinguish the animal "turkey" from the country "Turkey."

Because of these considerations, we prefer to preserve the original text as much as possible. In fact, we structure our code to allow us to access and read the text underlying each data point at any time. This practice has been particularly useful when interpreting over-time variation in our constructed measures, as it enables us to review any fluctuations in text-based measures through a targeted reading of the underlying text.

Cheap Talk

The second challenge relates to the potential for strategic behavior in management communications. Although the structure of earnings calls imposes institutional constraints on cheap talk, as previously discussed, researchers must still account for any potential residual strategic behavior in their empirical estimations. For example, when firms face declining operational performance, executives may expand their discussion to include external factors like climate change or political uncertainty in an effort to shift the focus of financial analysts away from company-specific issues. Researchers can address this through various methods. One approach involves examining whether estimation effects of text-based measures remain robust when controlling for earnings surprises and pre-call stock returns (for applications, see Hassan et al. 2019; Sautner et al. 2023). Other approaches aim to detect potential strategic communication through linguistic patterns (Dikolli et al. 2020) or vocal traits during the call (Hobson, Mayew, and Venkatachalam 2012), each offering distinct signals of possible misrepresentation.

Measurement Error

The third challenge concerns how accurately text-derived measures reflect the economic reality of a given firm. Like any measurement method (including those using conventional data sources), our text-derived measures of exposure, risk, and sentiment are subject to measurement error—random deviations from the true values. Generally, we can expect measurement error to be most pronounced at the granular (firm-quarter) level and to reduce considerably when aggregating across multiple observations, such as by averaging over multiple time periods for a given firm, all firms in a sector, or all firms in a country.

Several methods have been developed to quantify the remaining measurement error directly. For instance, Hassan et al. (2019) construct two distinct measures of political risk: one based on earnings-call transcripts and another derived from text in regulatory filings—specifically, the Form 10-K that publicly traded firms must file annually with the Securities and Exchange Commission.[8] One can then serve as an instrument for the other in an instrumental variables regression to gauge the degree of random noise in each. Under the assumption that measurement errors are uncorrelated between the earnings-call-based and the Form 10K–based measures, we find that approximately half of the firm-quarter variation in political risk measures results from measurement error, indicating a signal-to-noise ratio of about 1:1. This level of precision is comparable to the one reported by Bloom et al. (2018) when using accounting data to measure total factor productivity at the firm-quarter level.[9] While the extent of measurement error will always depend on the specific application, we can conclude that text-based measurements are not inherently noisier than those based on conventional data sources.

Other Sources of Text: Patents and Job Postings

In addition to corporate earnings call transcripts and 10-K filings, economists increasingly draw on other sources of text, including the use of patent documents and online job postings to study innovation and technological change.[10] Patents have long been used to measure innovation—for example, by counting filings or citations (Lerner and Seru 2022; Hall, Jaffe, and Trajtenberg 2005). Recently, researchers have shifted toward analyzing the content of the patent text itself, thereby measuring how patents differ or how they incorporate new technical vocabulary. Job postings, by contrast, can help to capture the labor market dimension of innovation and technology adoption. As job advertisements have steadily moved online, data vendors have accumulated large databases of postings containing rich information on skills demanded and job locations (Hershbein and Kahn 2018). By analyzing the wording of these postings, economists can study

---

[8] Earnings calls differ from regulatory filings, like the Form 10-K filing (Handley and Li 2020), in a number of ways. While the latter provide standardized quantitative information under strict accounting rules, earnings calls create an interactive forum where market participants can probe management's understanding and interpretation of firm performance. The question-and-answer format enables analysts to press management for clarification on vague statements, challenge overly optimistic projections, and request specific details about business conditions—imposing reputational costs for evasive or inconsistent responses. Under the protection of "safe harbor" provisions, management can discuss forward-looking information more extensively than in mandated filings, addressing emerging market conditions, anticipated economic shocks, and strategic responses not yet reflected in accounting metrics.

[9] Sautner et al. (2023) find a lower level of measurement error in their text-based measures of climate change risk, possibly due to the more standardized vocabulary used in climate-related discussions.

[10] We source full text of online job postings from Burning Glass, which aggregates them from online job boards, employer websites, and other sources. Patent text is obtained from the United States Patent and Trademark Office (USPTO).

which jobs are involved with using and developing which technologies, and where the jobs are located.

Broadly speaking, there are two complementary ways to use text analysis in these domains. One approach classifies the content of patents themselves, grouping them by linguistic similarity or by alignment with specific occupations or tasks. Another approach identifies specific technological concepts from within patent text, and then tracks their subsequent diffusion across other textual sources.

Classifying Patent Text

In the first category, for instance, Kelly et al. (2021) use a term frequency–inverse document frequency based measure of textual similarity between patents to identify "breakthrough" patents by comparing each patent's language to both prior and subsequent patents. Patents featuring low backward similarity (previously uncommon language) but high forward similarity (strong influence on subsequent patents) are labeled breakthroughs. Kalyani (2023) instead focuses on the number of new technical phrases within the patent's text to distinguish more creative patents (substantial new language) from derivative ones. He finds that creative patents—those introducing many new technical concepts—are strongly linked to firm- and sector-level growth in total factor productivity, as well as to abnormally positive stock returns. In contrast, derivative patents, which lack new technical language, show no such associations.[11]

Autor et al. (2024) and Kogan et al. (2021) use embeddings-based textual similarity between patents and descriptions of tasks from the Dictionary of Occupational Titles to classify inventions as "labor-saving" or "labor-augmenting." Inventions with higher overlap with routine tasks are classified as labor-saving, and those with higher overlap with nonroutine tasks as labor-augmenting. They then construct an occupation-by-year "exposure" measure to examine the wage and employment effects of these inventions.

Identifying Novel Technologies and Tracking Them across Multiple Corpora

Rather than classifying patents, the second approach uses patent text to identify a list of specific new technologies. In Kalyani et al. (2025a), we first analyze the full text of US patents from 1976 to 2014 to identify phrases that appear in multiple patents but did not exist before 1970, thus pinpointing new terminology specific to influential innovations over the past decades. Next, to map these phrases to specific new technologies, we search for these terms on Wikipedia, leveraging the standardized structure of Wikipedia technology pages to identify how a new phrase maps to a specific technology. This procedure leads us to identify a set of phrases associated with each of 1,286 unique new technologies developed since 1976. Examples include Computer Program, Monoclonal Antibody, Molecular Cloning, and Mobile Device, each linked to a set of new phrases associated with these new technologies.

---

[11] Mann and Püttmann (2023) classify patents as either automation-related or nonautomation, helping measure the extent to which industries face labor-replacing technologies.

Figure 5
Machine Learning/AI Jobs and Patents

Panel A. Core-based statistical area level
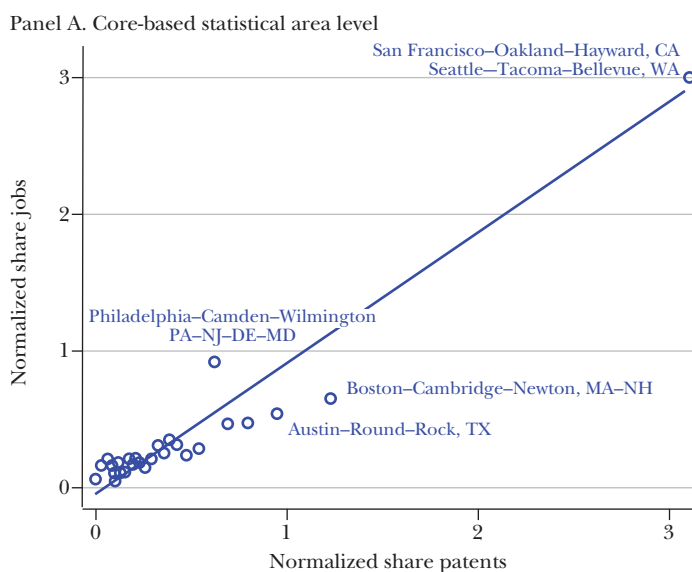
(continued)

With this list of technologies in hand, we then search for them in patents and job postings. By examining inventor addresses of patents with early mentions of the technology, we determine the locations where each technology was developed and, by tracking patent application years, identify when these technologies saw significant increases in patent citations (their "emergence year"). Finally, we cross-reference the technology phrases with the full text of online job postings, identifying millions of job listings from 2010 to 2019 that mention the new technologies.

These detailed data enable us to track the number, location, and skill requirements of job postings linked to each new technology. From these data, we observe that "pioneer locations"—areas where a technology was invented—enjoy large and persistent job growth in roles using or producing that technology. This advantage for pioneering locations is especially strong in high-skill jobs and persists for decades.

Figure 5 illustrates these relations for one of the 1,286 new technologies: artificial intelligence. Panel A plots a given region's share of patents mentioning AI prior to 2014 (AI emerged in 2014 as a new technology in patents) against its share of jobs mentioning AI post-2014. In each case, we normalize by the size of the local labor market so that the "normalized share" of jobs and patents indicates the local over- or underrepresentation of AI patents and jobs in a given area. We observe a strong association between AI's early development and subsequent job creation related to the technology.

Figure 5 (continued)
Machine Learning/AI Jobs and Patents

Panel B. Firm level
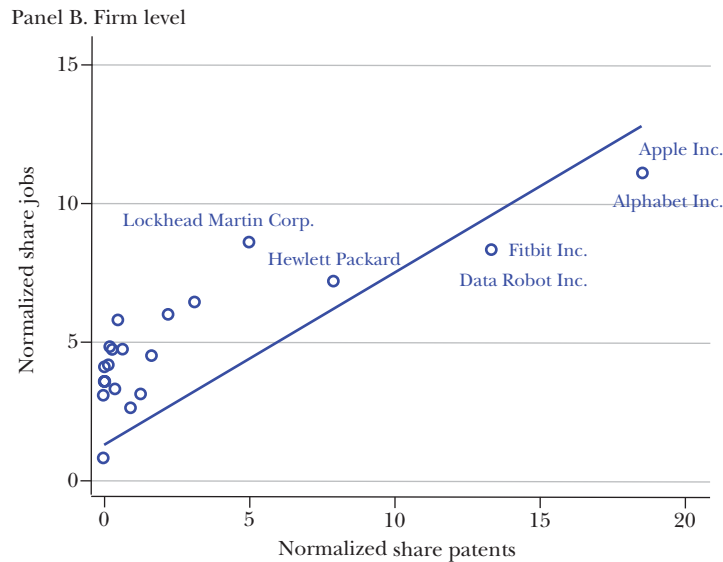
Note: Panel A shows a binned scatter plot depicting the relation between the normalized share of machine learning/AI-related job postings and patents across CBSAs (core-based statistical areas, defined by the Office of Management and Budget). The *x*-axis represents the normalized share of patents filed in a CBSA between 1976 and 2014 that mention machine learning/AI keywords, while the *y*-axis represents the normalized share of job postings mentioning these keywords in the same CBSA from 2015 to 2019. Specifically, the normalized share of jobs for each CBSA is calculated as:

$$Normalized\ share\ jobs_{c,\tau} = \frac{Share\ of\ machine\ learning\ or\ AI\text{-}related\ job\ postings_{c,\tau}}{Share\ of\ all\ job\ postings_{c,\tau.}}$$

This represents the proportion of machine learning/AI-related job postings within a CBSA relative to the national total, adjusted for the CBSA's overall share of job postings in the United States during 2015–2019. Similarly, the normalized share of patents for each CBSA is calculated as the proportion of patents filed between 1976 and 2014 that mention machine learning/AI keywords, relative to the total number of patents filed in that CBSA during the same period. Panel B conducts the same analysis, but at the firm level instead of the level of the core-based statistical area.

Places like San Francisco, Boston, Philadelphia, and Austin pioneered artificial intelligence patents, and later hosted the lion's share of AI-related jobs. A similar pattern holds true for the remaining 1,285 other technologies, where the vast majority of cases show a strong link between early patenting in the technology and subsequent job growth in using and producing that technology.

Panel B illustrates the same relationship at the firm level. Companies like Apple and Alphabet, which pioneered early development of the technology, show high levels of artificial intelligence patenting before 2014 and subsequent job creation. However, the slope of the regression line is now significantly flatter, indicating

positive spillovers from these companies to other local companies using and producing the new technology. This finding suggests that the relationship is steeper at the region level rather than at that of the firm.

This combined textual source approach thus offers valuable insight into how technology is adopted and spreads: Hassan, Kalyani, and Restrepo (2025) use these data to show a distinct boom-bust pattern in the number of technologies emerging over time and large regional discrepancies in the age of the median technology used. For example, the most frequently mentioned technology in job postings for positions in New York is about 34 years old, whereas the most frequently mentioned technologies in rural areas tend to be around 52 years old. These patterns have important implications for job opportunities and wage inequality across regions and over time.

Concluding Remarks

A text-based approach to economic analysis offers flexibility in tackling a variety of new and previously unexplored research questions. Our discussion has focused on corporate texts—such as earnings-call transcripts, but also patents and job postings—that are increasingly used in economic research. These sources offer valuable insights to researchers and policymakers on both historical and current events. Perhaps the most appealing feature of earnings calls is that they systematically capture the opinions, preferences, and expectations of key corporate decision-makers. They are available not only for a large sample of US publicly traded companies but also for numerous international firms, which makes earnings-call analysis a useful resource for gaining a global perspective on corporate strategies, risks, and expectations. Likewise, patents and job postings provide detailed and systematic insights into emerging technologies and their spread.

A wide array of other corporate text sources are available, which have so far received less attention, including but not limited to regulatory filings (such as annual or quarterly financial reports), investor presentations, press releases, analyst reports, corporate blogs, and contracts. Systematic research efforts will be needed to bring these additional corpora online for economic research and to determine how best to parse, structure, and interpret the information they contain.

The text-as-data techniques we discussed in this paper are increasingly informing policy work at central banks and international organizations. Researchers now use earnings calls to quantify trade-policy uncertainty (Correa et al. 2023), firms' cross-country inflation expectations (Albrizio, Dizioli, and Simon 2023), and supply chain "friendshoring" or "reshoring" (Alfaro and Chor 2023). Central bank economists build realtime dashboards of supply/demand imbalances, corporate sentiment, and risk from the same data (Gosselin and Taskin 2023; Andersson, Neves, and Nunes 2023), and exploit patent text to gauge how monetary shocks shape innovation (Ma and Zimmermann 2023). Several recent issues of the IMF's *World Economic Outlook* and similar publications include chapters that incorporate earnings-call

data to analyze supply chains, inflation expectations, and other macroeconomic risks (Faella et al. 2022; International Monetary Fund 2022; Ahn et al. 2023; Albrizio et al. 2023).

These applications are generally straightforward and easily interpretable, demonstrating that meaningful results in this field do not always require intricate linguistic and econometric models. Given the wealth of information available, through various widely accessible text sources, we anticipate that text analysis will become a standard part of the future tool kit for economic research.

While our analysis highlights the power of textual analysis in extracting meaningful signals from unstructured data, a crucial question is how to integrate these signals into economic theories, whether through structural modeling, causal inference, or predictive frameworks. Key open questions remain, particularly in refining textual measures to better align with economic concepts (such as the measurement of "risk"), and ensuring robustness in cross-context applications. Addressing these challenges will be essential for advancing the field.

References

Ahn, JaeBin, Benjamin Carton, Ashique Habib, Davide Malacrino, Dirk Muir, and Andrea Presbitero. 2023. "Geoeconomic Fragmentation and Foreign Direct Investment." In *World Economic Outlook*, 91–114. International Monetary Fund.

Albrizio, Silvia, John Bluedorn, Allan Dizioli, Christopher Koch, and Philippe Wingender. 2023. "Managing Expectations: Inflation and Monetary Policy." In *World Economic Outlook*, 49–69. International Monetary Fund.

Albrizio, Silvia, Allan Gloe Dizioli, and Pedro Vitale Simon. 2023. "Mining the Gap: Extracting Firms' Inflation Expectations from Earnings Calls." International Monetary Fund Working Paper 202.

Alfaro, Laura, and Davin Chor. 2023. "Global Supply Chains: The Looming 'Great Reallocation.'" NBER Working Paper 31661.

Andersson, Malin, Pedro Neves, and Carolina Nunes. 2023. "Earnings Calls: New Evidence on Corporate Profits, Investment, and Financing Conditions." *European Central Bank Economic Bulletin* 4: 56–62.

Autor, David, Caroline Chin, Anna Salomons, and Brayn Seegmiller. 2024. "New Frontiers: The Origins and Content of New Work, 1940–2018." *Quarterly Journal of Economics* 139 (3): 1399–1465.

Bae, Jihun, Chung Yu Hung, and Laurence van Lent. 2023. "Mobilizing Text as Data." *European Accounting Review* 32 (5): 1085–1106.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–1636.

Benati, Luca, and Paolo Surico. 2009. "VAR Analysis and the Great Moderation." *American Economic Review* 99 (4): 1636–52.

Berger, David, Ian Dew-Becker, and Stefano Giglio. 2020. "Uncertainty Shocks as Second-Moment News Shocks." *Review of Economic Studies* 87 (1): 40–76.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 (3): 1031–65.

Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun. 2021a. "The Diffusion of Disruptive Technologies." NBER Working Paper 28999.

Bloom, Nicholas, Tarek Alexander Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun. 2021b. *Data and Code for: "The Diffusion of Disruptive Technologies."* https://www.techdiffusion.net.

Bochkay, Khrystyna, Jeffrey Hales, and Sudheer Chava. 2020. "Hyperbole or Reality? Investor Response to Extreme Language in Earnings Conference Calls." *Accounting Review* 95 (2): 31–60.

Cao, Sean, Wei Jiang, Baozhong Yang, and Alan L. Zhang. 2023. "How to Talk When Machines Are Listening: Corporate Disclosure in the Age of AI." *Review of Financial Studies* 36 (9): 3603–42.

Correa, Ricardo, Julian di Giovanni, Linda S. Goldberg, and Camelia Minoiu. 2023. "Trade Uncertainty and U.S. Bank Lending." Federal Reserve Bank of New York Staff Reports 1076.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica* 50 (6): 1431–51.

Dikolli, Shane S., Thomas Keusch, William J. Mayew, and Thomas D. Steffen. 2020. "CEO Behavioral Integrity, Auditor Responses, and Firm Outcomes." *Accounting Review* 95 (2): 61–88.

Dixit, Avinash K., and Robert S. Pindyck. 1994. *Investment under Uncertainty.* Princeton University Press.

Faella, F., B. Javorcik, L. Kitzmüller, V. Langella, and H. Schweiger. 2022. "Global Supply Chains in Turbulence." In *European Bank for Reconstruction and Development Transition Report*, 59–61. European Bank for Reconstruction and Development.

Flynn, Joel P., and Karthik Sastry. 2024. "Attention Cycles." NBER Working Paper 32553.

Frankel, Richard, Marilyn Johnson, and Douglas J. Skinner. 1999. "An Empirical Examination of Conference Calls as a Voluntary Disclosure Medium." *Journal of Accounting Research* 37 (1): 133–50.

Galí, Jordi, and Luca Gambetti. 2009. "On the Sources of the Great Moderation." *American Economic Journal: Macroeconomics* 1 (1): 26–57.

Gallemore, John, Stephan Hollander, Martin Jacob, and Xiang Zheng. 2025. "Tax Policy Expectations and Investment." *Journal of Accounting Research* 63 (1): 363–412.

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy. 2019. "Text as Data." *Journal of Economic Literature* 57 (3): 535–74.

Gosselin, Marc-André, and Temel Taskin. 2023. "What Can Earnings Calls Tell Us about the Output Gap and Inflation in Canada? Bank of Canada Staff Discussion Paper 2023-13.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2005. "Market Value and Patent Citations." *RAND Journal of Economics* 36 (1): 16–38.

Handley, Kyle, and J. Frank Li. 2020. "Measuring the Effects of Firm Uncertainty on Economic Activity: New Evidence from One Million Documents." NBER Working Paper 27896.

Hassan, Tarek A., Stephan Hollander, Aakash Kalyani, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun. 2025a. *Data and Code for: "Text as Data in Economic Analysis."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E228102V1.

Hassan, Tarek A., Stephan Hollander, Aakash Kalyani, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun. 2025b. *Tutorial Package for: "Text as Data in Economic Analysis."* https://doi.org/10.34894/KNDZ9T, Dataverse NL, V1.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun. 2023. "Firm-Level Exposure to Epidemic Diseases: Covid-19, SARS, and H1N1." *Review of Financial Studies* 36 (12): 4919–64.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2019. "Firm-Level Political Risk: Measurement and Effects." *Quarterly Journal of Economics* 134 (4): 2135–202.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun. 2024a. "The Global Impact of Brexit Uncertainty." *Journal of Finance* 79 (1): 413–58.

Hassan, Tarek A., Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun. 2024b. "Sources and Transmission of Country Risk." *Review of Economic Studies* 91 (4): 2307–46.

Hassan, Tarek A., Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun. 2024c. *Replication Package for: "Sources and Transmission of Country Risk."* https://doi.org/10.5281/zenodo.7783429.

Hassan, Tarek A., Aakash Kalyani, and Pascual Restrepo. 2025. "New Technologies and the College Premium." Working Paper.

Hennessy, Christopher A., and Ilya A. Strebulaev. 2020. "Beyond Random Assignment: Credible Inference and Extrapolation in Dynamic Economies." *Journal of Finance* 75 (2): 825–66.

Hershbein, Brad, and Lisa B. Kahn. 2018. "Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings." *American Economic Review* 108 (7): 1737–72.

Hobson, Jessen L., William J. Mayew, and Mohan Venkatachalam. 2012. "Analyzing Speech to Detect Financial Misreporting." *Journal of Accounting Research* 50 (2): 349–92.

Hollander, Stephan, Maarten Pronk, and Erik Roelofsen. 2010. "Does Silence Speak? An Empirical Analysis of Disclosure Choices during Conference Calls." *Journal of Accounting Research* 48 (3): 531– 63.

Huang, Allen H., Reuven Lehavy, Amy Y. Zang, and Rong Zheng. 2018. "Analyst Information Discovery and Interpretation Roles: A Topic Modeling Approach." *Management Science* 64 (6): 2833–55.

International Monetary Fund. 2022. "Global Prospects and Policies." In *World Economic Outlook: Countering the Cost-of-Living Crisis*, 1–49. International Monetary Fund.

Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang. 2024. "ChatGPT and Corporate Policies." NBER Working Paper 32161.

Kalyani, Aakash. 2023. "The Creativity Decline: Evidence from US Patents." Preprint, SSRN. http://dx.doi.org/10.2139/ssrn.4318158.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Alexander Hassan, Josh Lerner, and Ahmed Tahoun. 2025a. "The Diffusion of New Technologies." *Quarterly Journal of Economics* 140 (2): 1299–365.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Alexander Hassan, Josh Lerner, and Ahmed Tahoun. 2025b. *Replication Package for: "The Diffusion of New Technologies."* https://doi.org/10.7910/DVN/AYRHYN.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights* 3 (3): 303–20.

Kim, Alex G., Maximilian Muhn, and Valeri V. Nikolaev. 2023. "Bloated Disclosures: Can ChatGPT Help Investors Process Information?" Chicago Booth Research Paper 23-07.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. 2021. "Technology-Skill Complementarity and Labor Displacement: Evidence from Linking Two Centuries of Patents with Occupations." NBER Working Paper 29552.

Lee, Joshua. 2016. "Can Investors Detect Managers' Lack of Spontaneity? Adherence to Predetermined Scripts during Earnings Conference Calls." *Accounting Review* 91 (1): 229–50.

Lerner, Josh, and Amit Seru. 2022. "The Use and Misuse of Patent Data: Issues for Finance and Beyond." *Review of Financial Studies* 35 (6): 2667–704.

Loughran, Tim, and Bill McDonald. 2011. "When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks." *Journal of Finance* 66 (1): 35–65.

Ma, Yueran, and Kaspar Zimmermann. 2023. "Monetary Policy and Innovation." NBER Working Paper 31698.

Mann, Katja, and Lukas Püttmann. 2023. "Benign Effects of Automation: New Evidence from Patent Texts." *Review of Economics and Statistics* 105 (3): 562–79.

Manning, Christopher D., Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval.* Cambridge University Press.

Matsumoto, Dawn, Maarten Pronk, and Erik Roelofsen. 2011. "What Makes Conference Calls Useful? The Information Content of Managers' Presentations and Analysts' Discussion Sessions." *Accounting Review* 86 (4): 1383–414.

Mayew, William J., Mani Sethuraman, and Mohan Venkatachalam. 2020. "Individual Analysts' Stock Recommendations, Earnings Forecasts, and the Informativeness of Conference Call Question and Answer Sessions." *Accounting Review* 95 (6): 311–37.

Min, Sarah. 2023. "Chegg Shares Drop More than 40 Percent after Company Says ChatGPT Is Killing Its

Business." *CNBC*, May 2. https://www.cnbc.com/2023/05/02/chegg-drops-more-than-40percent-after-saying-chatgpt-is-killing-its-business.html.

Pindyck, Robert S. 1991. "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature* 29 (3): 1110–48.

Rogers, Jonathan L., Andrew Van Buskirk, and Sarah L. C. Zechman. 2011. "Disclosure Tone and Shareholder Litigation." *Accounting Review* 86 (6): 2155–83.

Romer, Christian D., and David H. Romer. 1989. "Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz." In *NBER Macroeconomics Annual*, Vol. 4, 121–70. MIT Press.

Romer, Christina D., and David H. Romer. 2023. "Presidential Address: Does Monetary Policy Matter? The Narrative Approach after 35 Years." *American Economic Review* 113 (6): 1395–423.

Sautner, Zacharias, Laurence Van Lent, Grigory Vilkov, and Ruishen Zhang. 2023. "Firm-Level Climate Change Exposure." *Journal of Finance* 78 (3): 1449–98.

Sarkar, Suproteem K., and Keyon Vafa. 2024. "Lookahead Bias in Pretrained Language Models." http://dx.doi.org/10.2139/ssrn.4754678.

Song, Min, and Yi-fing Brook Wu. 2008. *Handbook of Research on Text and Web Mining Technologies.* Information Science Reference, IGI Global.

Suijs, Jeroen. 2007. "Voluntary Disclosure of Information When Firms Are Uncertain of Investor Response." *Journal of Accounting and Economics* 43 (2–3): 391–410.

Taskin, Temel, and Franz Ulrich Ruch. 2023. "Global Demand and Supply Sentiment: Evidence from Earnings Calls." Bank of Canada Staff Working Paper 2023-37.

Vafa, Keyon, Justin Y. Chen, Ashesh Rambachan, Jon Kleinberg, and Sendhil Mullainathan. 2024. "Evaluating the World Model Implicit in a Generative Model." Preprint, arXiv. https://doi.org/10.48550/arXiv.2406.03689.

How Congress Designed the Federal Reserve to Be Independent of Presidential Control

Gary Richardson and David W. Wilcox

Scholars typically trace the independence of the Federal Reserve to the famous Treasury-Fed Accord of 1951 (Friedman and Schwartz 1963; Goodfriend 1994; Meltzer 2003; Romero 2013). At that time, President Harry Truman, his staff, and the Federal Open Market Committee (FOMC) met in the Oval Office. The Treasury wanted to minimize the cost of government borrowing and asked the Fed to keep interest rates low. The president's subordinates released a statement to the press indicating that the Fed agreed to do as the Treasury asked. A courageous member (and former chair) of the Federal Reserve Board, Marriner Eccles, defied Truman and undercut the administration's position by leaking to the press a summary of the meeting indicating the Fed disagreed with the Treasury. The two agencies then released a joint statement in which they concurred the Fed would manage the money supply as it saw fit and with an eye to keeping inflation in check.

In this version of history, the Treasury-Fed Accord is the foundation for the Federal Reserve's independence from the Executive Branch, and Marriner Eccles is its champion. Those inferences are inaccurate, however, except for the bit about the leak. Eccles did release the summary. This paper sets the record straight.

We demonstrate, based on primary source documents from the day, that the independence of the Federal Reserve traces back to the Banking Act of 1935. Eccles played a leading role in drafting the initial version of the 1935 legislation, but

■ *Gary Richardson is Professor of Economics, University of California, Irvine, Irvine, California and a research associate at the National Bureau of Economic Research. David Wilcox is Senior Fellow, Peterson Institute for International Economics, Washington, DC, and director of US economic research, Bloomberg Economics. Their email addresses are garyr@uci.edu and DWilcox@piie.com.*

contrary to his image as a champion of Fed independence, Eccles argued that the Fed and monetary policymaking should be directly controlled by the president. As the bill worked through the House, the Senate, a conference committee, and then to a final vote, however, Congress explicitly rejected the Eccles position and devised the institutional features that today serve as the bulwarks of Fed independence.

This historical record remains relevant for current policy and political debates. If Fed independence is built on a 1951 agreement between the Fed and the Treasury, then any US president could direct that the earlier agreement be overturned, and a new agreement put in its place. However, if legislation undergirds the Fed's independence, then the strong presumption would be that only an act of Congress changing the law or a decision of the Supreme Court ruling the law unconstitutional could fundamentally reshape the president's ability to influence monetary policy.

The Meaning of Fed Independence

The congressional architects of the Federal Reserve wanted the Fed's leaders to act independently and structured the System to foster this independence. The Fed's architects used the term "independent" with a specific definition in mind. Independent policymakers acted in the public's interest (as defined by Congress) and were meant to make decisions free from (or independent of) personal, private, and political interests. In 1913, as the legislation creating the Federal Reserve was taking shape, discussions on this topic emphasized ensuring the Fed's leaders would not make decisions from which they would profit personally; decisions that would benefit private parties, firms, financial institutions, "bankers," "a clique of banks," "Wall Street," "the Money Trust," "rich men," or "a few men in New York" at the expense of the public; or decisions that were influenced by partisan politics; or decisions that would benefit one section of the country relative to another (Senate 1913, pp. 2071–2, 2604–6, 2655–65, 2702, 2755, 2837, 2840, 2957, 3087, 3114–5, 3175).

In particular, the relationship between the Federal Reserve System and the executive branch of the federal government was on the minds of the architects of the Fed. The original design protected the independence of regions of the United States from each other and the federal government. The United States was divided into Federal Reserve districts, and in the original design of the Fed, each district could enact policies it preferred. A Federal Reserve Board supervised the Fed districts, but the Board had limited authority, could not compel districts to act, and had no ability to act on its own. The Board's members were appointed by the president and included the Secretary of Treasury, who served at the will of the president, but Congress limited the Board's powers because a Fed Board with too much power could become a "tyrant" and a president with power over the Fed raised risks of "too much politics in the bad sense of the word" (Senate 1913, pp. 2655, 3168).

The relationship between the Fed and Congress was also on the minds of the Fed's architects. We have found no discussion in the Congressional Record from 1913 until today indicating that members of Congress thought the Fed should

operate independent of congressional oversight or intrusion. On the contrary, the original architects of the Fed, like Congress today, regularly compelled leaders of the Fed to testify before Congress, to explain their decisions, and to discuss ways to improve the performance of the system. If the Fed did not act as Congress intended, then Congress admonished its leaders and amended the law to correct the problem. Congress, in other words, has always insisted that the Fed was accountable to the legislature.

The Fed's independence from the president became a pressing issue in 1935, when Congress redesigned the Federal Reserve, eliminated regional independence, consolidated the regional Feds into a unified system, and centralized monetary-policy decision making. At that time, discussions of independence focused on ensuring monetary policy remained independent from the president, whose political and personal incentives did not always coincide with the interests of the public. Since 1935, the term "Fed independence" typically indicates the independence of monetary policy from presidential control.

The Eccles Bill and the Banking Act of 1935

Title II of the Banking Act of 1935 redesigned the Federal Reserve System's leadership structure. Marriner Eccles, who served as Governor of the Federal Reserve Board in 1934 and 1935 and chair of the Board of Governors of the Federal Reserve System from 1936 to 1948, led a team of four Fed staffers who drafted Title II.[1] While this group received feedback from a few members of the Roosevelt administration, it neither solicited feedback from the Federal Reserve Board, the Federal Open Market Committee, or Federal Advisory Council, nor from senators or congressmen or their staffs, nor from regulatory or policymaking professionals in the federal or state governments, nor from academics, businessmen, or bankers (House of Representatives 1935a, pp. 351–3; Senate 1935, pp. 550, 564). Except for Eccles himself, members of the Fed Board did not see the bill until it was "presented [to Congress] and printed" (Senate 1935, pp. 554). Eccles presented the draft to Congress, argued for its provisions, and claimed authorship. Contemporary Congressmen and commentators referred to Title II as "the Eccles Bill."

Eccles proposed to reconfigure monetary policymaking by the Fed in two fundamental ways. First, monetary policy would be nationalized, and control of it would be centralized in Washington, DC. The Fed's leaders would be given tools enabling them to adjust the nationwide supply of money and credit and thus to influence nationwide interest and inflation rates. Second, monetary policy would be directed by the president. Eccles (House of Representatives 1935a, p. 191) asserted:

---

[1] From 1913 to 1935, the chair of the Federal Reserve Board (who was also the Secretary of Treasury) was a distinct position from the Governor of the Board (who served as the Board's chief executive officer).

> [A]n administration is charged, when it goes into power, with the economic and social problems of the Nation. Politics are nothing more or less than dealing with economic and social problems. It seems to me that it would be extremely difficult for any administration to be able to succeed and intelligently deal with them entirely apart from the money system. There must be a liaison between the administration and the money system—a responsive relationship.

President Franklin D. Roosevelt had controlled monetary policy since the depths of the Great Depression in 1933 when passage of a series of laws temporarily gave him this authority on an emergency basis. The Banking Act of 1935 was an opportunity to make those temporary changes permanent (House of Representatives 1935a, pp. 72, 181–3).

National monetary policy and presidential control were sea changes from how the Fed had operated previously. From its founding in 1913 until 1933, decisions about monetary and credit policy had rested with heads of the regional Fed banks (then called governors, the same title as the head of the Federal Reserve Board), who had set policies for their own districts, and whose actions were constrained by rules like the gold standard that limited their discretion over aggregate money and credit. The regional Feds struggled to coordinate policies. Those struggles inspired the creation of the Federal Open Market Committee by the Banking Act of 1933. The original FOMC consisted of the twelve heads of the regional Fed banks. It set open-market policies for the entire Fed System, although the Federal Reserve Board could veto its decisions, and individual Federal Reserve Banks could decline to participate in FOMC-dictated policies—although dissenting banks could not initiate countervailing policies on their own accord.[2]

To nationalize policy under the president, Eccles proposed to restructure the Federal Open Market Committee, increase its authority, and change the relationship between its members and the president. In the Eccles proposal, members of the Federal Reserve Board would serve at the pleasure of the president and could be fired at any time and for any reason. The heads of the twelve Federal Reserve banks would serve one-year terms. They would be appointed by their bank's board of directors, but their initial appointment and all reappointments would have to be approved by the Fed Board. The FOMC would consist of the head of the Fed Board (who would serve as chair of the committee), two other members of the Fed Board (possibly including the Secretary of Treasury and Comptroller of Currency, both of whom sat *ex officio* on the Fed Board), and the heads of two Federal Reserve banks (Senate 1935, pp. 196, 313, 395, 535). This committee would have the authority to devise and execute open-market operations for the entire Fed system. The Fed Board would have the power to determine reserve requirements, which had

---

[2] The Banking Act of 1935 gave the FOMC something close to its current form. The FOMC's structure was further amended in 1942, which introduced the current voting structure. For discussion, see https://www.minneapolisfed.org/article/1999/the-federal-reserves-beige-book-a-better-mirror-than-crystal-ball.

previously been determined by Congress, as well as the discount rate, which had previously been set by regional Fed banks with the consent of the Federal Reserve Board. In sum, under Eccles' proposal, the Fed Board and the FOMC would control the three principal tools of monetary policy, and the president could swiftly replace the Fed's leaders (in the case of the Board) or influence their selection (in the case of the heads of the Reserve Banks) if the president disliked their policies.

The House Debate

The House of Representatives held hearings on the initial Eccles bill over 25 days between February 21 and April 8, 1935. For ten days, Eccles was the sole witness called. When Eccles testified, he introduced a series of amendments that he had written himself, which further concentrated monetary policymaking. The Fed Board would have authority over all three levers of monetary policy: open-market operations, discount lending, and reserve requirements. The chair, governor, and all members of the Fed Board would serve at the president's discretion. The Federal Open Market Committee would be replaced by an advisory committee consisting of representatives of five of the twelve Federal Reserve banks, which could recommend policies to the Board, but which would not vote upon or have authority over monetary policymaking (Senate 1935, p. 699; House of Representatives 1935a, p. 181–3).

At this stage of the legislative process, most members of the House seemed supportive of the Eccles bill. During the hearings, representatives voiced occasional concerns, but they called witnesses that criticized the Eccles proposal on only one issue: the lack of direction about how to use the powers of monetary policy. Both Irving Fisher, the famous economist, and Robert Latham Owen, a politician who was one of the sponsors of the Federal Reserve Act in 1913, spoke in favor of clear guidance and a price-level mandate. Fisher argued:

> [The] bill does not specifically state what is right and what is wrong. In all probability it cannot be rightly administered. There is too much discretion in it and too little guidance, too little in the nature of a criterion (Senate 1935, p. 519).

Fisher thought the Congress should decide on a price index and target level of prices and order the Federal Reserve to use open-market operations and the other tools of monetary policy to stabilize prices around that target. The bill, as then drafted, gave "tremendous and dangerous power to [the Federal Reserve], and some day it will be abused. . . . it would be manipulated for some ulterior purpose" (House of Representatives 1935a, p. 520).

While hearings in the House provided Eccles with a friendly forum, it soon became clear that a number of representatives opposed his plan. Seven of the twenty-five members on the House Committee on Banking and Currency issued a minority report whose lead author was Representative John Hollister. The minority

argued that making members of the Fed Board removable at the will of the president took away their independence and subordinated monetary policy to the president. Central banks in prosperous democracies, like England, operated independently. Central banks in dictatorships, like the Soviet Union, were controlled by their governments. The minority argued that the Eccles bill should not pass as written and that monetary policymaking's structure should be changed only after careful study by experts inside and outside the federal government.

Despite these criticisms, a majority of the House of Representatives supported the Eccles bill, mostly in the form he recommended. The bill then passed to the Senate.

The Senate Debate

The Senate held hearings from April 19 to June 3, 1935, on both the Eccles original bill and the amended version that passed the House. Carter Glass played a leading role. As a member of the House of Representatives, Glass had been one of the principal sponsors of the original Federal Reserve Act in 1913. He entered the Senate in 1920 and had been a contributor to most monetary, banking, and financial legislation that passed Congress between 1913 and 1935. Glass set the tone when, rather than referring to the legislation that had been passed over from the House as "the administration bill," he noted "it is Governor Eccles' bill" (Senate 1935, p. 357).

A Senate subcommittee chaired by Glass called a lineup of luminaries to scrutinize the Eccles plan. While most came down in favor of national policy and centralized authority, almost all criticized lodging that authority in a body beholden to the president. Critics included many members of the Federal Reserve Board and Federal Advisory Council, directors of many Federal Reserve banks, the head of the American Bankers Association, and even the Secretary of Treasury, Henry Morgenthau, who was at that time also chair of the Federal Reserve Board and a personal confidant of President Franklin Roosevelt.

Morgenthau argued that he would like to have monetary authority "concentrated in an independent Government agency" (Senate 1935, p. 505). The agency should operate independent of "all outside influence—just as independent as you can make it . . . . [like] the Supreme Court . . . . independent of the President. . . . No member of the board could be removed except by impeachment" (Senate 1935, p. 506).

Adolph Miller, a member of the Federal Reserve Board since its founding in 1914, also supported the centralization of monetary authority (Senate 1935, pp. 750–1). The Fed's original structure divided authority and responsibility among the Federal Reserve banks, making it difficult to discern who was the "responsible agent," and leaving policymaking susceptible to being influenced by special interests (Senate 1935, p. 687). Miller criticized subordination of monetary policy to the president. Miller feared "political control" as well as "banker control" (Senate

1935, p. 687). He believed the Federal Reserve Board needed to be "independent" with members who regard service "as a great public responsibility which runs to the public rather than to an official of the administration of the day" (Senate 1935, p. 729–30). Miller suggested many of the institutional features adopted in 1935 that underlie the Fed's independence today. These include removing the Secretary of Treasury and Comptroller of the Currency from the Federal Reserve Board, changing that organization's name to the Board of Governors of the Federal Reserve System (Senate 1935, p. 754), setting membership of the Board at seven (Senate 1935, p. 758), and writing into law the provision that members of the Board of Governors could be dismissed only "for cause."

During the hearing, the Senate called 60 witnesses. Almost all criticized the Eccles proposal to place the president in control of monetary policy and advocated that the central bank operate independent of the president. Senators picked up on ideas they liked and asked later witnesses what they thought of them. An example was Miller's idea for changing the Federal Reserve Board's name to the Board of Governors of the Federal Reserve System. "Governor" was the traditional term for the chief executive officer of a national or central bank. The Bank of England had a governor, as did the other central banks in Europe. To highlight the shift in authority to the Board from the Reserve Banks, Miller thought that all members of the Board should have the title "governor" and the name of the board itself should include the word "governors." William McAdoo, a senator from California who had been Secretary of Treasury when the Fed was founded 22 years before, then suggested relabeling the heads of the regional banks—who had been called governors—as something else. Later discussions led to the idea of giving them the title of "president," which was the traditional title for the head of a commercial bank.

When Eccles testified before the Senate Subcommittee, he defended his proposal. He insisted that the Board was and would always be political. Even if the president lacked legal authority to remove Board members (Senate 1935, p. 282), Eccles argued,

> [N]o man would stay on the Board if the President of the United States wished to appoint someone else in his place. . . . It seems to me to be immaterial whether a Governor has or has not a technical right to stay on the Board, if the President prefers to have someone else as Governor, because no person who is qualified for that position would choose to remain in these circumstances.

Eccles would demonstrate his conviction about this proposition in later decades. However, the Senate strove to ensure that Eccles was wrong.

At the end of the debate, the Senate removed almost all the specifics of the Eccles bill from Title II and replaced them with concepts drawn from previous proposals, witnesses' suggestions, and the senators' own insights. The Senate vested monetary policymaking in a Federal Open Market Committee and Board of Governors designed to operate independent of the Executive Branch. The Senate added many provisions to shield the Federal Reserve System from presidential pressure.

Members of the Board of Governors serve 14-year terms, staggered so that one term expires on January 31 of every even-numbered year. A member may continue to serve after the expiration of their term until their successor has been confirmed by the Senate. The president nominates a chair of the Board from among the members of the Board, and that person receives a four-year term as chair (separate from his or her term as governor) that starts on the date of their confirmation.[3] Members of the Board can be removed only "for cause." The heads of the Federal Reserve Banks, now labelled presidents, serve five-year terms. The boards of directors of each Fed bank appoint the presidents, subject to the approval of the Board of Governors. The FOMC makes the principal decisions concerning open-market policy. The FOMC elects its own chair and vice chair. By tradition, the FOMC elects the chair of the Federal Reserve Board as the chair of the committee, and the president of the Federal Reserve Bank of New York as its vice chair, but this arrangement is not written in law. The FOMC's twelve voting members consist of the seven members of the Board of Governors, the president of the New York Fed, and presidents of four other Fed banks who serve on a rotating basis. The other Fed bank presidents serve as nonvoting participants but technically are not "members" in years when they do not vote. Under this structure, a single president of the United States serving two terms can appoint only four members of the FOMC if the members of the Board of Governors all serve their full terms.

The President's Power to Dismiss Fed Board Members

A focus of the Senate's discussion was how to prevent the president from dismissing members of the Federal Reserve Board for mere policy disputes. Winthrop Aldrich, chair of Chase National Bank (antecedent of today's JPMorganChase) and son of Nelson Aldrich (who had spearheaded the initial drive a quarter-century earlier to create what became the Fed), suggested limiting dismissal to cases where a Board member had "become permanently incapacitated or has been inefficient, or guilty of neglect of duty, or of malfeasance in office, or of any felony or conduct involving moral turpitude, and for no other cause and no other manner except by impeachment" (Senate 1935, pp. 396–7). Frank Vanderlip, former Assistant Secretary of Treasury, former president of National City Bank of New York (today's Citibank), and a contributor to the original Federal Reserve Act in 1913, argued that members of the Fed board "should not be removable by the President" (Senate 1935, pp. 916–7). As noted earlier, Morgenthau argued that no member of the Fed board should be removable except by Congress via impeachment, just like members of the Supreme Court (Senate 1935, p. 506). McAdoo concurred on this point, although his analogy was to congressional removal of federal judges (Senate 1935,

---

[3] Until the Federal Reserve Reform Act of 1977, the chair and vice chair did not need to be separately confirmed to their respective roles. They were simply designated as such by the president. https://www.federalreservehistory.org/essays/fed-reform-act-of-1977.

p. 755). Miller favored the phrase "no member of the Board shall be removable from office during the term for which he was appointed . . . except for malfeasance" (Senate 1935, p. 754).

The extended discussion during the Senate hearings occurred, in part, due to uncertainty about the law. In 1926, the Supreme Court held in *Myers v. United States* (272 US 52) that the president had the sole power to dismiss executive branch officials and that restrictions on the president's power to dismiss were unconstitutional. On May 1, 1935, while the Senate debated the Eccles bill, the Supreme Court revisited *Myers* when it heard arguments in *Humphrey's Executor v. United States* (95 US 602 [1935]). The circumstance that differed between the two cases was that *Humphrey's Executor* involved a member of the Federal Trade Commission (FTC)—an independent agency, not unlike the Fed—whereas *Myers* concerned the president's powers over employees within the executive branch itself.

Senators and witnesses discussed these cases, whether they thought *Myers* or *Humphrey's Executor* applied, and whether the Supreme Court would find for Humphrey's executor.[4] Most thought (1) the Supreme Court would find for Humphrey's executor, meaning that Congress could limit the president's ability to dismiss personnel from agencies like the Federal Trade Commission or the Fed; and (2) the Senate should wait for the court to hand down its decision in *Humphrey's Executor* before finalizing the language in the legislation. The Supreme Court announced its decision in favor of Humphrey's executor on May 27, near the end of the Senate hearings, and witnesses soon noted its relevance for the independence of the Federal Reserve (Senate 1935, p. 998).

A surprising issue arose during these discussions. From the founding of the Fed in 1913 until 1933, members of the Fed Board who were appointed and confirmed by the Senate served for fixed terms "unless sooner removed for cause by the President." The Banking Act of 1933, however, removed the phrase "for cause" from the law (Senate 1935, p. 396). Aldrich raised this issue during his testimony. Aldrich stated that Eccles and the other members of the Fed Board were unaware of the change. Glass said "I must have been asleep when that was eliminated from the act. I have no recollection of it." Aldrich replied that the unnoticed change illustrated the dangers of "hasty legislation." Glass retorted, "I do not know that it was due to hasty action. It might have been due to covert action" (Senate 1935, p. 398).

House and Senate Agree on Federal Reserve Independence

The Senate bill differed substantially from the Eccles bill that had emerged from the House. When the House and Senate pass differing versions of legislation,

---

[4] The case was called "*Humphrey's Executor*" because Humphrey had died by the time of the Supreme Court argument. The issue at stake in the case was whether his executor should be allowed to recover the salary that would have been paid to him had he remained on the Federal Trade Commission between the time of his dismissal and his death.

they typically convene a temporary, ad hoc committee to reconcile their differences and settle on a single version of the legislation that the two chambers will adopt and send to the president. The leadership of each chamber chooses its own members for this conference committee, usually from among the members of the bill's committee of jurisdiction. The number chosen is not set but determined by the leadership, and it can differ across chambers. The conferees meet, discuss, and negotiate the text of the final legislation, but each chamber's delegation votes separately. The Senate conferees vote whether to accept the text of the conference report, and the House conferees do likewise. Agreement is reached when both the Senate and House delegations to the conference committee vote to accept the final text. By selecting members of their delegation whose views they know, the leadership of each chamber can influence the report crafted by the conference committee.

The intent of House leaders becomes clear when one sees the representatives they chose for their unusually small delegation of only three members (House of Representatives 1935b). Henry Steagall, chair of the Committee on Banking and Currency, was a natural choice given his stewardship of the legislation and knowledge of the subject. T. Alan Goldsborough, a Democrat, and John Hollister, a Republican, were two of the most prominent and strident critics of the Eccles proposals. They did not support the version of Title II that passed the House. Appointing both of them and only one other individual to the conference delegation ensured that the Senate's version of Title II, where monetary policy was independent of the president, would become law.

A Slow Transition to Independence

Passage of the Banking Act of 1935 did not transform the Fed into an independent central bank overnight. The provisions phased in over time. On January 1, 1936, the Act removed the Secretary of Treasury and Comptroller of Currency from the Federal Reserve Board and ended the terms of the Board's six other members. The Roosevelt administration reappointed Eccles and one other member to the new Board of Governors and nominated five new governors. Their terms were staggered to end after 2, 4, 6, 8, 10, 12, and 14 years.[5] All subsequent terms would last 14 years. These arrangements put in place the restraint that no president after Roosevelt would be able to appoint a majority of the Federal Open Market Committee, as long as Board members served their full terms.

This transition gave the Roosevelt administration what it wanted, a Federal Reserve Board that it had chosen. For supporters of Fed independence, it was also a sensible compromise, because under the law existing in 1935, the Roosevelt administration could have dismissed all members of the Board before the Banking Act of 1935 came into effect. The process of confirming seven new members could have taken

---

[5] A governor who is appointed to and serves out the remaining portion of an unexpired 14-year term may be appointed to a new 14-year term, and so may end up serving even longer than 14 years in total.

months (or longer, if disputes had arisen between the president and the Senate), possibly leaving the Board with too few members to function. The Banking Act of 1935 avoided these complications by ending the terms of all current board members, which resolved potential employment and pension complications, and by allowing current members to continue serving on the Board until the Senate approved their replacements. The compromise also gave the Senate substantial influence over Roosevelt's selections to the Board, because the Senate could keep current Board members in place until the president appointed replacements the Senate deemed acceptable.

With Roosevelt's new Board of Governors in place, the Federal Reserve continued to implement Roosevelt's monetary program, as Eccles had intended. The Fed collaborated closely with the Treasury, particularly during World War II, when the central bank helped to finance the war, fund US allies, embargo enemies, stabilize the domestic and world economy, and plan postwar economic policies (Richardson 2013). During World War II and for years afterwards, the Fed pegged the price of US government bonds, keeping interest rates low and stable. This practice helped the government finance the war, but it also increased currency held by the public and reserves held by banks. The low rates and buildup of these balances triggered inflation after the end of wartime rationing and price controls. Inflation exceeded 8 percent in 1946 and 14 percent in 1947 before subsiding during the recession that began in 1948.

Eccles served as the chair of the Board of Governors throughout this period and remained a faithful executor of the president's monetary program, which he believed was the proper role for a person in his position. Roosevelt reappointed Eccles as chair in 1940 and 1944.

When Roosevelt died in 1945, Harry Truman became president. True to his beliefs, Eccles offered to resign and told Truman that it was his "feeling that the Chairman, who is designated by the President, should serve at the pleasure of the President." Truman (1948) rejected Eccles's offer, told him that "there was no one I [Truman] desired to appoint in your place," and asked Eccles to complete his term as chair.

In 1948, when his third term as chair expired, Eccles wrote to Truman that "I have not altered my conviction that the Chairman of this Board should serve at the pleasure of the President, and I sought to have such a provision included in the Banking Act of 1935" (Eccles 1948). This time, Truman told Eccles that he now desired to appoint to the Board of Governors a new member, Thomas McCabe, who would be designated as chair. Truman (1948) urged Eccles to "remain as a member of the Board and accept the Vice Chairmanship so that the benefit of your long experience and judgment will continue to be available and so that you may carry forward legislative proposals now pending in Congress." Eccles remained on the Board, initially accepting but later declining the position of vice chair.[6]

---

[6] Eccles stayed on the Board for an additional three years after his term as chair expired. During that time, he completed a memoir. Eccles left the Board to launch a campaign for one of Utah's Senate seats. He did

With the transition from Roosevelt to Truman and from Eccles to McCabe, the Federal Reserve's leadership gradually became less tied to the current president, as Congress envisioned when it designed the Fed's leadership structure in 1935. By mid-1950, the FOMC's members consisted of two governors (McCabe and Thomas Vardaman) appointed by the sitting president, four members appointed by the preceding president, a vacancy due to the death of a governor, and five presidents of Reserve Banks, including Alan Sproul from the New York Fed.

The 1951 Treasury-Fed Accord

By 1951, McCabe served as the chair of the Federal Open Market Committee, and Sproul served as the vice chair. In the summer of 1950, McCabe and Sproul had led efforts to confront the incipient inflation. An economic rebound began at the start of the year. At that time, an expanding economy and low interest rates pushed up prices. Inflationary pressures increased in June with the onset of the Korean War, fueled partly by a rush to buy durable goods to get ahead of presumed wartime shortages (Gagnon and Rose 2025). Sproul, McCabe, and their staff advocated slowing purchases of government bonds, which would raise interest rates, curtail lending, reduce reserve balances of banks that form the foundation for the money supply, and limit the financial and economic activity fueling inflation. A majority of the FOMC, including all five Reserve Bank presidents and a few governors, supported this position.

Other members of the Federal Open Market Committee, including Eccles, felt the Fed should accommodate the Treasury Department's request that the Fed continue to purchase government bonds at a fixed price. The guarantee limited the cost of government borrowing, which increased rapidly with the expansion of defense expenditures necessitated by the fighting in Korea and fears that the conflict could spread to other parts of the globe. Proponents of this policy thought the Fed should try to limit inflation by raising reserve requirements and controlling credit. The former would force commercial banks to hold a larger proportion of their resources as non-earning deposits at Federal Reserve banks. The latter would prevent commercial banks from expanding lending to consumers and businesses, compelling banks to purchase bonds (both domestic and foreign) or to hold increasing quantities of cash. This faction conceded that current laws might not give the Fed sufficient power over reserves and credit, but thought the Fed should ask Congress to give it stronger powers to control commercial bank lending.

From the summer of 1950 through the winter of 1951, the Federal Open Market Committee debated which policy would work best. Treasury Secretary John W. Snyder and President Truman himself met with the FOMC's leaders on several occasions in the fall of 1950 and early 1951. After one of these meetings, McCabe

not prevail in the primaries and then resumed his career as a businessmen and banker with his family's industrial and financial conglomerate.

reported that the members of the FOMC should implement the policies outlined by the Treasury and the president or they could resign (FOMC, January 31, 1951, p. 16).[7] Sproul continued to advocate the majority's position. He argued that Congress had given the Federal Reserve and the Treasury different jobs to do, and that the Federal Reserve should do its job as outlined by Congress (FOMC, January 31, 1951, pp. 18–9, 23).

This policy dispute in 1950 and 1951 seems to be a scenario envisioned by Congress when it designed the Fed's leadership structure 15 years earlier. The most powerful monetary policy tool, open market operations, was controlled by the Federal Open Market Committee. The current president had appointed a minority of the governors, and even if both of those advanced the president's preferred policies, they could not force the FOMC to act if the governors appointed by the previous president disagreed. When the governors disagreed, then the Federal Reserve bank presidents—who were not political appointees, usually extremely experienced, and held five seats on the FOMC—would settle the dispute. The president could not compel the FOMC to implement presidential directives. The president could not fire FOMC members merely due to policy disputes. The president could ask FOMC members to resign, but they could refuse, or ignore the suggestion, as they did in this case.

As the Federal Open Market Committee and the Treasury debated the issue during 1950 and 1951, both organizations looked to the law and to Congress for guidance. In July 1950, in a letter to the Secretary of Treasury, the FOMC argued that the task given to the FOMC by Congress was "maintenance of public confidence in the value of money" and limiting inflation (FOMC, August 18, 1950, p. 4). The FOMC believed the Treasury's proposal would be highly inflationary. "Logic, as well as the bitter experience of recent years," proved their point (FOMC, August 18, 1950, pp. 2–4). Treasury Secretary John Snyder replied that "he sometimes thought it might be wise to ask the Congress to put the responsibility for handling the debt in the Federal Reserve which had the money market tools, adding that, of course, he realized the Congress wouldn't do that" (FOMC, August 18, 1950, p. 7).[8] The Treasury Secretary, in other words, acknowledged that federal law did not authorize the FOMC to pursue the policies advocated by the president, and Congress would not change the law even if the president asked it to do so.

In September 1950, the Treasury Secretary and the leaders of the Federal Open Market Committee, McCabe and Sproul, again discussed the "relative responsibilities

[7] While McCabe reported the resignation option to the FOMC, and he himself resigned shortly afterwards, no extant evidence indicates that President Truman or the Secretary of Treasury requested or demanded the FOMC's resignation. The resignation option may have been McCabe's opinion, or it may have been suggested to him by others, perhaps during unrecorded conversations with the president, Secretary of Treasury, or their staff.

[8] Snyder discussed this issue in detail on several occasions. His discussions indicate that by "handling the debt," Snyder meant that the Fed should keep the costs of the federal borrowing low by keeping interest rates low.

of the Treasury and the Federal Reserve" (FOMC, September 28, 1950, p. 9). The Treasury Secretary stated:

> [H]e thought it highly desirable at some time for this whole question of responsibility to be reviewed by the proper authorities because conditions had changed so materially since action by the Congress creating the Federal Open Market Committee as now constituted. At the time of that action, he said, we did not have a public debt of anything like the present magnitude, and in view of the problems involved in the handling of the debt, he felt there should be a review of the Congressional authority and the responsibilities of the agencies should be re-defined.

Neither the Treasury Secretary nor the president, however, asked Congress to review the situation or change the law, presumably because they did not think that Congress would do so.[9]

In January 1951, the Federal Open Market Committee again discussed the law and the intent of Congress. Sproul argued that "the main job Congress had given the System to do; in terms of the importance of the matter to the Federal Reserve and the nation, the most important job was to maintain public confidence in the real value of the dollar" and not maintaining "a fixed interest rate" or "fixed prices of Government securities." (FOMC, January 31, 1951, p. 18). "In light of the responsibilities placed on the System by the Congress," pursuing policies advocated by the Treasury that conflicted with those responsibilities "would be both impossible and improper" (FOMC, January 31, 1951, p. 23). The FOMC, Sproul argued, should not resign due to its dispute with the president. Resignation "would be an admission of failure or inability to carry out our statutory responsibilities, without giving the Congress an opportunity to review our performance and to express its will . . . [and a] complete abdication of the responsibility of the members of the Federal Open Market Committee" (FOMC, January 31, 1951, p. 19).

The battle over Fed independence came to a head in early 1951 in a meeting between the Federal Open Market Committee and President Truman and his staff. The president explained the political problems that he faced collecting sufficient taxes to fund the federal government and fight the communists in Korea and elsewhere. He emphasized that maintaining confidence in the value of federal government securities would help combat communism (FOMC, January 31, 1951,

---

[9] While the FOMC minutes refer on numerous occasions to the law, statutes, statutory responsibilities, and Congress in general, the minutes seldom mention specific laws or statutes. An exception appears in the minutes for the FOMC executive committee meeting on February 14, 1951. Chairman McCabe noted that in response to a comment from Senator O'Mahoney, he asked the Board of Governor's general counsel to research the law regarding the Treasury and the Fed's roles in management of the public debt. The general counsel found that nothing had changed "since the Banking Act of 1935" (FOMC 1933–2025; February 14, 1951, p. 5). In the minutes for the preceding FOMC meeting, McCabe had reported that he asked the legal division to research the law regarding the Fed, Treasury, and government debt management from the Banking Act of 1935 to the present.

pp. 24, 26). "He did not indicate exactly the details of what he had in mind, but he reiterated that we should do everything possible to maintain confidence in the Government securities market" (FOMC, January 31, 1951, pp. 26).

McCabe responded that the Federal Open Market Committee had statutory responsibilities to the "general economy" that it had to fulfill "principally to promote stability in the economy by regulating the volume, cost and availability of money, keeping in mind at all times the best interests of the whole economy" (FOMC, January 31, 1951, pp. 24–6). The FOMC could not ignore these legal imperatives or assign them to other parties. Given that it fulfilled its legal responsibilities, McCabe indicated that the FOMC had in the past and would continue in the future help to ensure an orderly market for federal government bonds.

After the meeting, the Truman administration released a statement to the press indicating that the Fed had agreed to peg the price of government bonds as the Treasury directed. In response, Eccles released the memorandum from the Federal Open Market Committee of the meeting, talked to reporters, and told the press that the FOMC had not, in fact, agreed to do what the administration had claimed. According to FOMC members and the FOMC's notes of the meeting, the president spoke vaguely, never directly discussing the policy dispute between the Fed and the Treasury, and never asking the FOMC to commit to specific policies or actions.

In the next meeting of the Federal Open Market Committee, Eccles defended, explained, and accepted responsibility for his actions. The FOMC minutes presented Eccles's 1,650-word-long statement in single-spaced text with narrow margins, which often indicates quotation of a written document that the author presented to the FOMC and read at the meeting. Eccles stated that he acted on his own to discharge his "public responsibility," to protect the Federal Reserve, and to protect "my own record" (FOMC, February 6–8, 1951, p. 16). During the Great Depression of the 1930s and World War II, the Fed had followed the Treasury's lead and pegged interest rates on government bonds at a low level. The situation now was different. The economy was booming; international conflicts were limited. The problem was inflation due to low interest rates, cheap money, and monetary expansion. The Fed had "provided the means for the growth in the money supply, which had been directly related to the increase in the cost of living and the price level" (FOMC, February 6–8, 1951, p. 18). The Fed was "almost solely responsible for this inflation" (FOMC, February 6–8, 1951, p. 18). In the past year, Eccles had advocated asking Congress for increased powers so that the Fed could peg interest rates as the Treasury asked, while raising the reserves of and limiting lending by commercial banks (FOMC, February 6–8, 1951, p. 17).

Eccles ended his statement with a mea culpa (FOMC, February 6–8, 1951, p. 19):

> I think that I, as a member of this Committee, have not been as aggressive as I should have been. I think I have not made the record I should have. As I look back to 1946 and 1947, when the Treasury and a budgetary surplus and the war was over [sic], particularly when we were having our troubles with Secretary of the Treasury Vinson, we should have taken a stronger stand. If we

had had a row, I could have resigned. As I look back on it, I regret I did not take a stronger stand for obtaining substitute authority from the Congress. But until we get that authority, it is up to us to use what we have.

In retrospect, Eccles realized that as Fed chair, he should have acted to curtail inflation after World War II. Then, he could have asked Congress to expand the Fed's powers over commercial-bank lending so that the Fed could attempt the strategy that he advocated now. If Congress said no, he could have resigned. Now, the Fed had to act to control inflation and needed to employ the strategy championed by Sproul and McCabe, as that was the only arrow in the Fed's quiver.

The day after the Eccles soliloquy, the FOMC sent two letters: one to the Treasury Secretary described details of their proposed policies, while the other to the president described the problem that the country faced and the Fed's solution. The latter indicated (FOMC, February 6–8, 1951, p. 26):

In your recent meeting with us you clearly stated as your objective one which underlies Federal Reserve operations—the maintenance of confidence in the integrity of the dollar and therefore in Government securities. In your recent economic report to the Nation you said: "If inflation continues to gain cumulative force it will multiply the cost of the defense program. It will undermine production, destroy confidence, generate friction and economic strife, impair the value of the dollar, dissipate the value of savings and impose an intolerable burden upon fixed income groups. This must not happen."

The Federal Open Market Committee explained that their objective was the same as the president's: to stabilize the value of the dollar and eliminate inflationary pressures: "There is no effective way of stemming this tide that will not reflect itself in interest rates" (FOMC, February 6–8, 1951, p. 27). The FOMC did not favor higher interest rates on government borrowing per se, as the Treasury had insinuated. The FOMC favored the lowest interest rate on government borrowing consistent with a stable dollar and "that will cause true investors to buy and hold these securities" (FOMC, February 6–8, 1951, p. 27). In inflationary times, Federal Reserve purchases of government securities at below-market prices actually undermined confidence. "The inevitable result is more and more money and cheaper and cheaper dollars. This means less and less public confidence" (FOMC, February 6–8, 1951, p. 27). In their meeting, the president had not asked the FOMC to continue on this dangerous road, and the FOMC would not do so, because the policy advocated by the Treasury would weaken the financial structure of the United States.

These letters sparked a month of intense discussion between the staffs of the Treasury and Federal Reserve and between leaders of those organizations. In March 1951, reports presented to the Federal Open Market Committee indicated that "both sides agreed that monetization of debt must be stopped as far as possible" (FOMC, March 1–2, 1951, p. 11), but the two sides continued to disagree on how this might be done. These discussions led to the release later that month of the

Treasur-Fed Accord, a brief document which succinctly summarizes the different jobs given to the Fed and Treasury by the Congress.

Following the Accord, Fed chair McCabe resigned, and in his place, Truman appointed William McChesney Martin, an assistant secretary of Treasury who had played an important role as the Treasury's representative negotiating the text of the Treasury-Fed Accord.[10] Truman may have thought that by changing the Fed's chair, he would change the Fed's policies. The majority of the Federal Open Market Committee, however, led by Sproul, continued to believe that the Fed should act as Congress intended. The new chair, Martin, agreed. So, the Fed fought inflation. The effectiveness of the leadership structure crafted by Congress was clearly demonstrated to all parties.

Implications

The congressional debate over the Banking Act of 1935 leaves a voluminous record of Congress's intent when it crafted the modern Fed. Congress clearly posed and thoroughly debated the issue of Fed independence from the president. The record demonstrates that Congress wanted the president's hands far from the levers of monetary policy. This record may be a vital piece of evidence for current and future disputes about the legal basis for Federal Reserve independence.

References

Eccles, Marriner. 1948. "Letter to President Truman." Marriner S. Eccles Papers, Box 81, Folder 18, Item 3. https://fraser.stlouisfed.org/archival-collection/marriner-s-eccles-papers-1343/letter-president-truman-464934.

Eccles, Marriner S. 1951. *Beckoning Frontiers, Public and Personal Recollections.* Edited by Sidney Hyman. Alfred A. Knopf.

FOMC (Federal Open Market Committee). 1933–2025. "Federal Open Market Committee Meeting Minutes, Transcripts, and Other Documents." https://fraser.stlouisfed.org/title/677.

Freidman, Milton, and Anna Jacobson Schwartz. 1963. *A Monetary History of the United States, 1867–1960.* Princeton University Press.

Gagnon, Joseph E., and Asher Rose. 2025. "Why Did Inflation Rise and Fall So Rapidly? Lessons from the Korean War." Peterson Institute for International Economics Working Paper 25-1.

---

[10] McCabe's resignation may have been delayed, rather than hastened, by the dispute with the Treasury. In his resignation letter, McCabe reminded President Truman that he had not planned on serving his full term as Fed chair when Truman nominated him in 1948, since he wished to return to the private sector. McCabe had prolonged his service to handle the resurgence of inflation, the onset of the Korean conflict, and the controversy over how the Fed should handle these twin crises. After resolving these problems, McCabe stated that he now wished to resign from government service and resume his career in the private sector.

Goodfriend, Marvin. 1994. "Why We Need an 'Accord' for Federal Reserve Credit Policy: A Note." *Journal of Money, Credit and Banking* 26 (3): 572–80.

Hetzel, Robert L., and Ralph Leach. 2001. "The Treasury-Fed Accord: A New Narrative Account." *Federal Reserve Bank of Richmond Economic Quarterly* 87 (1): 33–56.

House of Representatives. 1935a. *Banking Act of 1935: Hearings before the Committee on Banking and Currency, House of Representatives, Seventy-Fourth Congress, First Session, on H.R. 5357, February 21, 22, 26, 27, 28, March 1, 4, 5, 6, 11, 12, 13, 14, 15, 18, 19, 20, 21, 22, 25, 26, 27, 28, April 2, 8, 1935.* Government Printing Office.

House of Representatives. 1935b. *Conference Report No. 1822 (Report to Accompany H.R. 7617) Seventy-Fourth Congress, First Session.* Government Printing Office.

Meltzer, Allen. H. 2003. *A History of the Federal Reserve, Volume 1: 1913–1951.* University of Chicago Press.

Richardson, Gary. 2013. "The Federal Reserve's Role During WWII." Federal Reserve History, November 22. https://www.federalreservehistory.org/essays/feds-role-during-wwii.

Romero, Jesse. 2013. "The Treasury-Fed Accord." *Federal Reserve History*, November 22. https://www.federalreservehistory.org/essays/treasury-fed-accord.

Senate, Committee on Banking and Currency. 1913. *Public Law 63-43, 63d Congress, H.R. 7837.* Government Printing Office.

Senate, Subcommittee of the Committee on Banking and Currency. 1935. *Banking Act of 1935: Hearings before a Subcommittee of the Committee on Banking and Currency, United States Senate, Seventy-Fourth Congress, First Session, on S. 1715 and H.R. 7617, April 19 to June 3, 1935.* Government Printing Office.

Truman, Harry. 1948. "Letter to Marriner Eccles." Marriner S. Eccles Papers, Box 2, Folder 23, Item 3. January 27, 1948. https://fraser.stlouisfed.org/archival-collection/marriner-s-eccles-papers-1343/letter-mr-eccles-461505.

Recommendations for Further Reading

Timothy Taylor

This section will list readings that may be especially useful to teachers of undergraduate economics, as well as other articles that are of broader cultural interest. In general, with occasional exceptions, the articles chosen will be expository or integrative and not focus on original research. If you write or read an appropriate article, please send a copy of the article (and possibly a few sentences describing it) to Timothy Taylor, preferably by e-mail at <taylort@ macalester.edu>, or c/o Journal of Economic Perspectives, Macalester College, 1600 Grand Ave., Saint Paul, MN 55105.

Smorgasbord

Leonardo Gambacorta and Vatsala Shreeti describe "The AI supply chain" (Bank of International Settlements Papers No. 154, March 2025, https:// www.bis.org/publ/bppdf/bispap154.htm). They provide "a comprehensive analysis of the market structure, economic forces and challenges along the five key input layers of the AI supply chain . . . The first layer consists of specialised hardware— most notably specialised microprocessors or chips that perform the complex computations needed for AI model training and inference. The second layer is cloud computing, which provides the infrastructure required to build, store and use AI models. Training data are the next input layer: AI models feed on vast data sets, which include everything from text to images and videos, typically sourced

■ *Timothy Taylor is Managing Editor, Journal of Economic Perspectives, based at Macalester College, Saint Paul, Minnesota.*

from both public and proprietary repositories. Foundation models, the fourth layer, are large, pre-trained models that can be adapted to many different uses. These form the base for the last layer: downstream AI applications. The market structure and the economic forces shaping the layers are different. The first two layers of the AI supply chain (hardware and cloud) are characterised by high fixed costs, economies of scale and scope, high switching costs and consumer inertia, and network effects. The substantial investment required for developing and maintaining AI infrastructure creates barriers to entry and favours larger firms with significant financial resources. In contrast, the market for training data, foundation models and downstream AI applications may currently be more contestable but could still be prone to 'winner takes all' dynamics."

Cass R. Sunstein discusses "The use of algorithms in society" (*Review of Austrian Economics*, December 2024, 37: 399–420, https://link.springer.com/article/10.1007/s11138-023-00625-z). "Across a wide range of settings, use of algorithms is likely to improve accuracy, because algorithms will reduce both bias and noise. Indeed, algorithms can help identify the role of human biases; they might even identify biases that have not been named before. . . . In important cases, algorithms struggle to make accurate predictions, not because they are algorithms but because they do not have enough data to answer the question at hand. Those cases often, though not always, involve complex systems. (1) Algorithms might not be able to foresee the effects of social interactions, which can depend on a large number of random or serendipitous factors, and which can lead in unanticipated and unpredictable directions. (2) Algorithms might not be able to foresee the effects of context, timing, or mood. (3) Algorithms might not be able to identify people's preferences, which might be concealed or falsified, and which might be revealed at an unexpected time. (4) Algorithms might not be able to anticipate sudden or unprecedented leaps or shocks (a technological breakthrough, a successful terrorist attack, a pandemic, a black swan). (5) Algorithms might not have 'local knowledge,' or private information, which human beings might have. Predictions about romantic attraction, about the success of cultural products, and about coming revolutions are cases in point. The limitations of algorithms are analogous to the limitations of planners, emphasized by Hayek in his famous critique of central planning."

David Schleicher and Nicholas Bagley tackle a key problem of federalism in "The State Capacity Crisis" (Niskanen Center, January 1, 2025, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5188510). "Three areas are of particular concern to us. First, . . . [p]olarization notwithstanding, in 39 out of 50 states, both houses of legislatures and the governor come from the same party and only rarely have institutional limits like the filibuster. As a result, majority parties can usually do what they want. Gridlock is not the problem. Yet state legislatures are in an even worse spot than Congress. Voters know almost nothing about what is happening in state politics, and increasingly vote for the same party for state legislature that they do for president and Congress. This pervasive nationalization of state and local elections means that state legislative performance has little

connection to electoral outcomes. Gerrymandering is also a much worse problem at the state level than at the federal level . . . as is the lack of staff capacity and resources. Second, . . . [s]tate administrative law is as strict, and often stricter, than federal administrative law, both with respect to the procedures it imposes and the intensity of judicial review. State and, in particular, local governments have extremely powerful rules requiring lots of public participation in administration. Because small groups with members that care intensely about state and local decisions are much easier to form than groups representing a diffuse public interest, unrepresentative private interests—whether that's the Chamber of Commerce or neighborhood NIMBYs—overwhelm administrative process at the expense of majoritarian preferences. Third, . . . [e]very state (save Vermont) is legally required to balance its budget, no state can print money to inflate away its deficits, and all states face both legal and market limits on their capacity to borrow. When a recession depletes tax revenue, states have few choices except to increase taxes or reduce spending—right when public services are needed most."

Sandra Baquie, Yueling Huang, Florence Jaumotte, Jaden Kim, Rafael Machado Parente, and Samuel Pienknagura from the IMF have published a Staff Discussion Note summarizing some lessons in "Industrial Policies: Handle with Care" (IMF, SDN/2025/002, March 2025). "After falling out of favor in the 1990s, IPs [industrial policies] have been widely used, especially after 2017. Arguments in IPs' favor relate to market failures, economies of scale, and collective action problems, but factors like limited state capacity and capture by private and political actors can hamper their effectiveness. Moreover, IPs' historical track record has been mixed. . . . IPs are associated with moderate and uneven improvements in the performance of targeted sectors, but this relationship varies with sector characteristics and instruments. . . . IPs targeting highly distorted sectors (those with high markups and external financial dependence making firms vulnerable to credit market imperfections) are linked to improvements in value added that are four times as large in the medium term as those targeting less distorted sectors. Similarly, IPs targeting upstream sectors (those providing inputs for other sectors) are associated with broader benefits to the economy through positive spillovers along the supply chain. Targeting products that are closer to the frontier, as gauged by a high initial revealed comparative advantage, also yields faster and larger gains, while IPs targeting products in which the country is not competitive do not show definite gains over the horizon considered. Turning to instruments, export incentives are linked to more sustained competitiveness and productivity improvements (albeit after an adjustment period) compared to domestic subsidies, which are strongly associated with increases in capital accumulation. . . . Structural reforms have, on average, much larger effects than IPs, pointing to their fundamental role. IPs are accompanied by smaller economic benefits than 'horizontal policies' focused on lowering corruption, improving governance and enhancing access to credit. Even when IPs may be desirable, horizontal policies are key. IPs are more effective in countries with better institutions, business environment, and financial market conditions, and a more educated workforce."

The World Bank has published *21st-Century Africa: Governance and Growth*, a collection of eight chapters on different aspects of development in the countries of sub-Saharan Africa, edited by Chorching Goh (May 2025, https://openknowledge. worldbank.org/entities/publication/03bdf5e3-daba-404a-a638-25c5511c2376). From the "Main Message" section at the start of the report: "Over the past 25 years, Africa has achieved notable progress . . . Mortality rates have fallen, with life expectancy rising from 50 years in 1998 to 61 years in 2022. School attendance has improved, with primary school enrollment increasing from 80 percent in 1999 to 99 percent in 2022 and secondary school enrollment increasing from 26 percent to 45 percent over the same period. The early 2000s saw strong economic growth fueled by high commodity prices. China emerged as a trade and investment partner, and the continent experienced a massive inflow of foreign capital from 17.6 percent of gross domestic product (GDP) in 1998 to 38.1 percent in 2018. Consequently, African countries have shown significant growth performances: from 2000 to 2019, 7 of the world's 10 fastest-growing economies were in Africa. Aid dependence has declined, tax revenues have increased, and the median poverty rate fell by about 10 percentage points to about 43 percent." However, the chapters focus on future challenges. For example, Cesar Calderon and Ayan Qu point out in Chapter 3 on "Productivity": "When compared with the average living standards of the rest of the world, GDP per capita in Sub-Saharan Africa has declined over the past three decades. . . . The region's lack of convergence in living standards with the rest of the world largely results from its inability to sustain growth over time. If Sub-Saharan Africa had grown (in per capita terms) at the same pace as the global economy since 1990, its level of income per capita in 2022 would have been more than 40 percent higher than its actual level. If it had grown at the same pace as emerging East Asia, the region's income per capita would have been nearly three times its 2022 level."

Simon Pittaway grabs hold of the arguments in "Yanked away: Accounting for the post-pandemic productivity divergence between Britain and America" (Resolution Foundation, April 2025, https://www.resolutionfoundation.org/publications/ yanked-away). "Britain isn't alone in its recent struggles: productivity in Canada, France, and Italy has also fallen below pre-pandemic levels. But the US has bucked this trend entirely. Although recent US economic data has been soft, its productivity growth since the pandemic has been in a league of its own: between 2019 and 2024, US productivity has grown by 9.1 per cent. It is the only G7 economy where productivity growth has accelerated in recent years, increasing at by 1.9 per cent a year on average, compared to 1.0 per cent in the 2010s . . . Meanwhile, US productivity growth has been boosted by a continued boom in oil and gas extraction. . . . US tech companies have become world-beating, but it is the way that the rest of the US economy uses this tech that drives its productivity overperformance. Tech-using rather than tech-producing services account for more of its recent productivity gains relative to Britain. Professional, scientific and technical services accounts for another sixth (17 per cent) of the post-pandemic gap in productivity growth between the US and the UK—twice as much as the tech (ICT) sector (8 per cent). What is driving this broad outperformance? Business investment is a critical part of the story . . . "

International Trade

Pol Antràs surveys the field in his "FBBVA Lecture 2024: The Uncharted Waters of International Trade" (*Journal of the European Economic Association*, February 2025, pp. 1–51, https://academic.oup.com/jeea/article-abstract/23/1/1/7931801). "[I]nternational trade research has undergone transformative developments over the past 25 years, fueled by breakthroughs in data, methodologies, and theory. Notably, the seminal work of Bernard and Jensen (1999) and Melitz (2003) revolutionized the field by shifting the focus from industries and products to firms as the central units of empirical and theoretical analyses of international trade flows. This firm-level perspective has significantly broadened the scope of organizational decisions examined in global firms. Researchers have highlighted that firms engage not only in exporting but also in importing (Antràs, Fort, and Tintelnot 2017), multinational activity (Helpman, Melitz, and Yeaple 2004), and the organization of global value chains (GVCs) (Antràs and Chor 2013). In the process, firms make key strategic organizational design decisions to manage cross-border connections with offshore production units efficiently (Antràs and Helpman 2004). Simultaneously, the field has witnessed a quantitative revolution, marked by the development of medium-scale models that allow for quick, back-of-the-envelope estimates of the implications of trade cost shocks on real income (Eaton and Kortum 2002; Arkolakis, Costinot, and Rodríguez-Clare 2012). Additionally, a growing body of empirical work in international trade has been 'unshackled' from the constraints of traditional trade theory, leveraging new data sources and innovative methods to provide fresh insights into the nature of global trade, in many cases exploring topics that had not yet been developed theoretically (e.g., Autor, Dorn, and Hanson 2013; Atkin, Khandelwal, and Osman 2017). . . . The primary goal of this paper is to highlight several underexplored areas—what I call 'uncharted waters'—for young researchers trying to leave an imprint in the international trade field."

Maurice Obstfeld interrogates "The U.S. Trade Deficit: Myths and Realities" (*Brookings Papers on Economic Activity*, Spring 2025, with video, comments, and discussion, https://www.brookings.edu/articles/the-us-trade-deficit-myths-and-realities). "[T]hree prominent myths locate the principal sources of U.S. deficits beyond America's borders. The first myth is that U.S. deficits originate mostly in unfair foreign trade practices to which America has exposed itself through ill-advised trade liberalization. On this theory, tariffs provide a ready remedy. A second myth is that the world's desire to hold the dollar as its main reserve currency is a prime determinant of U.S. foreign deficits. One variant of this view, which is entirely false, is that U.S. current account deficits are necessary for foreign official holders to acquire dollars; another variant, of limited quantitative importance but more accurate, is that global dollar preference has asset price effects that make the U.S. deficit bigger. A final myth is that U.S. deficits result entirely from excessive saving by our trade partners, which forces the U.S. to borrow from them and spend the proceeds on extra imports. . . . This paper shows that the realities are more nuanced. While foreign and domestic trade policies can affect both imports

and exports separately, they are not principal drivers of their difference, the trade deficit. The U.S. can supply the world with dollars without trade deficits. Finally, the trade deficit reflects the interplay of foreign and U.S. macroeconomic factors (including China's saving rate and the U.S. government budget deficit) and often U.S. factors are dominant. Higher Federal fiscal deficits, for example, will likely raise U.S. trade deficits despite more import tariffs."

Kyle Pomerleau and Erica York offer a primer for "Understanding the Effects of Tariffs" (*AEI Economic Perspectives*, April 2025, https://www.aei.org/research-products/report/understanding-the-effects-of-tariffs/). "Tariffs do not directly affect the trade deficit, which is driven by net lending and borrowing between the US and the rest of the world. Tariffs would cause the US dollar to appreciate, which would increase untaxed imports, reduce exports, and transfer wealth from Americans to foreign holders of US assets. As an indirect tax, tariffs could prompt the Federal Reserve to increase the price level to prevent unemployment. Tariffs would reduce the after-tax return on work and investment, distort the allocation of resources in the economy, and ultimately reduce economic output in the long run. Broad-based tariffs can raise revenue for the federal government, but that revenue is partially offset by significant behavioral and economic responses. Tariffs would reduce after-tax income for households at all income levels but be slightly regressive."

Interviews

Jon Hartley serves as interlocutor in "Revisiting Empirical Macroeconomics with Robert Barro" (Hoover Institution, *Capitalism and Freedom Podcast*, March 25, 2025, audio and transcript available https://www.hoover.org/research/revisiting-empirical-macroeconomics-robert-barro-harvard-economics-professor). "So if you have a massive increase [in government spending], such as the transfer payments, even more under Biden than under the first Trump administration, a way to avoid paying for that by cutting other spending or by raising taxes is by having a inflation that's surprising from a perspective of a pre-crisis period. And that basically wipes out a lot of real value of the government bonds that are outstanding and it amounts to a very large temporary source of revenue which can be something like 10, 15 percent of the GDP. So it's not a minor deal. And that's empirically about what happened in the US and also in other places. . . . [P]eople don't like the idea that it might not have been completely crazy to pay for the expenditure in substantial part through this surprise inflation. I get a lot of grief on that point from people who normally, who normally are on my side about things because they just want to think of the inflation as being stupid and being dramatically harmful. So I think what was harmful is the excessive fiscal expansion, particularly under Biden. It was unnecessary to have that vast increase in transfer payments, but given that you had it, we effectively paid for most of it through the surprise inflation. And maybe that part was not so crazy because the alternatives would have been also very costly."

"Paul Krugman talks trade, industrial policy, and Trump," with Chad Bown (*Trade Talks Podcast*, March 16, 2025, https://tradetalkspodcast.com/podcast/206-paul-krugman-talks-trade-industrial-policy-and-trump/). On the distinction between industrial policy and tariffs: "Max Corden's 1974 book *Trade Policy and Economic Welfare* remains relevant. And what Corden and others said was, if there's something that you think you need to be producing, then encourage production. The answer is industrial policy. The answer is to subsidize or otherwise promote. In general, a tariff has side effects that may not be what you want. If you were worried that too many of the world's semiconductors are being produced within striking range of China, then you want to subsidize production of high-end semiconductors in the United States. But that's not a good reason to raise the cost of semiconductors to the U.S. downstream industry. So, there's a really pretty strong case for industrial policy here. That's the generic principle. Now actually implementing it is tricky, by the way. The thing about these agglomeration economies is that, once they're well established, they're really hard to break. . . . And so if you want to develop rival agglomerations to the existing agglomerations that you think are in the wrong place, it's going to be expensive and hard, which doesn't mean you shouldn't do it, but you should realize that it's not something you do by throwing a few dollars at the problem."

Discussion Starters

The OECD discusses "The Ocean Economy to 2050" (March 31, 2025, https://www.oecd.org/en/publications/the-ocean-economy-to-2050_a9096fb1-en.html). From the "Executive Summary": "The ocean covers 71 percent of Earth's surface, comprises 90 percent of the biosphere, provides food security for over three billion people, enables the transportation of over 80 percent of global goods, and hosts sea cables carrying 98 percent of international Internet traffic. . . . If considered a country, the ocean economy would be the world's fifth-largest economy in 2019. From 1995 to 2020, it contributed 3 percent to 4 percent of global gross value added (GVA) and employed up to 133 million full-time equivalents (FTEs). The global ocean economy doubled in real terms in 25 years from USD 1.3 trillion of GVA in 1995 to USD 2.6 trillion in 2020, growing at an annual average rate of 2.8 percent."

Caitlin Myers investigates "From *Roe* to *Dobbs*: 50 Years of Cause and Effect of US State Abortion Regulations" (*Annual Review of Public Health*, 2025, pp. 433–446, https://www.annualreviews.org/content/journals/10.1146/annurev-publhealth-071823-122011). "For more than 50 years—beginning with abortion reforms in the 1960s and continuing through the *Dobbs* decision in 2022—state regulations of abortion were neither uniform nor consistent. States reformed and repealed abortion bans leading up to the *Roe* decision in 1973. Following *Roe*, they enacted both demand-side regulations of people seeking abortions and supply-side regulations of people providing abortions. The resulting laboratory of state policies affords natural experiments that have yielded evidence on the effects of abortion

regulations on demographic, health, economic, and other social outcomes. I present a brief history of state policy variation from 1967 through 2016 and review the empirical scholarship studying its effects. This literature demonstrates that the liberalization of abortion access in the 1960s and 1970s allowed women greater control over their fertility, resulting in increased educational attainment and earnings. Subsequent state restrictions in the 1980s through 2010s had the opposite effect, particularly when they increased the financial and logistical costs of obtaining an abortion."

Alex Tabarrok lays out some steps toward "Pandemic preparation without romance: insights from public choice" (*Public Choice*, published online April 16, 2025, https://link.springer.com/article/10.1007/s11127-025-01277-2). "I examine how standard political incentives—myopic voters, bureaucratic gridlock, and fear of blame—predictably produced an inadequate pandemic response. The analysis rejects romantic calls for institutional reform and instead proposes pragmatic solutions that work within existing political constraints: wastewater surveillance, prediction markets, pre-developed vaccine libraries, human challenge trials, a dedicated Pandemic Trust Fund, and temporary public–private partnerships. These mechanisms respect political realities while creating systems that can ameliorate future pandemics, potentially saving millions of lives and trillions in economic damage."

The American Economic Association

The Journal of
Economic Perspectives

Summer 2025, Volume 39, Number 3

Symposia

Housing Markets

Articles

Feature

AMERICAN
ECONOMIC
ASSOCIATION