

From Wages to Welfare: Decomposing Gains and Losses from Rising Inequality*

Jonathan Heathcote

Federal Reserve Bank of Minneapolis and CEPR

Kjetil Storesletten

Federal Reserve Bank of Minneapolis and CEPR

Giovanni L. Violante

New York University, CEPR, and NBER

PRELIMINARY DRAFT: AUGUST 3, 2010

Abstract

This paper offers a critical evaluation of the large literature that studies the welfare consequences of the recent shift in the wage structure in the United States. Welfare calculations based on changes in the the empirical distribution of consumption and hours worked –analyzed through the lenses of a social welfare function– yield welfare losses of the order of 2 percent of lifetime consumption. We argue that these welfare calculations ignore that the same sources behind the shift in the wage structure –the growth in the skill premium and the rise in wage volatility– can lead to a welfare improvement as individuals adjust their education and labor supply decisions. Quantifying the importance of these channels of adjustment requires a structural model. In our model-based calculations, under a plausible calibration, welfare losses turn into gains over 1 percent of lifetime consumption.

*We are grateful to Chris Tonetti for outstanding research assistance and to Mark Aguiar, Erich Battistin, and Fatih Guvenen for sharing their data. The opinions expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

The structure of relative wages in the US economy has undergone a major transformation in the last thirty years. Wage differentials between more and less educated workers have risen sharply over this period (Katz and Autor (1999); Lemieux (2008)). Within narrow groups of workers defined by education, gender, and birth cohort, the distribution of wages has become much more dispersed (Juhn, Murphy, and Pierce (1993)). This increase in within-group dispersion reflects wider fixed individual wage differentials and more pronounced volatility in both persistent and transitory shocks (Gottschalk and Moffitt (1994, 2009)). Overall, the US wage structure has become much more unequal.

This surge in US economic inequality has generated great interest among labor economists and macroeconomists. A vast theoretical and empirical literature set its sight on the sources of this phenomenon. The leading explanation is that the widespread adoption of new information and communication technologies has raised the relative productivity of more skilled labor – complementary to the new technologies in production – and reduced the demand for less skilled workers employed in tasks easily replaceable by the new machines (Krusell, Ohanian, Rios-Rull, and Violante (2000); Acemoglu (2002); Autor, Katz, and Kearney (2006); Acemoglu and Autor (2010)).¹ A concurrent, but less prominent, role is attributed to falling demand for unskilled-intensive goods produced in the US because of greater openness to trade (Autor, Katz, and Kearney (2006)) and off-shoring of unskilled stages of production (Feenstra and Hanson (1996)). The rise in idiosyncratic volatility is viewed as the result of a more turbulent work environment and faster skill obsolescence (Violante (2002)), changes in wage compressing labor market institutions such as unions (DiNardo, Fortin, and Lemieux (1996)), and the contractual shift towards performance-based and piece-rate pay (Lemieux, MacLeod, and Parent (2009)).

Much of the underlying motivation for this large body of research is that the observed movements in the relative wage structure had sharp welfare consequences for US households. In particular, the low skilled workers were hit especially harshly: when deflated by the official CPI, earnings of the bottom 10 percent of US male workers have not grown since the late 1970s. This striking trend poses a challenge to economists and policy-makers interested in how the government should alter the efficiency-equity trade off through redistributive taxation and other similar instruments.

¹Some authors argue that the simple model of capital-skill complementarity, or skill-biased demand shift, is not fully adequate to explain the most recent dynamics of wage inequality (in the last decade), when wage differentials widened exclusively at the top of the distribution. See Lemieux (2008) for a discussion.

Starting with the pioneering analysis of Cutler and Katz (1991, 1992), Slesnick (1994) and Attanasio and Davis (1996) a parallel growing literature has been aiming at quantifying the welfare implications of the rise in inequality in the US economy. The premise of this body of work is that, given the numerous insurance channels available to US households to absorb wage movements, inferring purely from wage data how standards of living have changed is inappropriate. Wider wage dispersion is indeed detrimental to economic welfare only if it transmits, at least partially, into wider consumption dispersion. Therefore, shifts in the distribution of consumption are a better indicator of shifts in the distribution of household welfare. Of course, since high consumption in the presence of low wages can be achieved at the expense of longer hours worked, one should factor in also shifts in the distribution of leisure.

The goal of this paper is to re-examine the question “What are the welfare gains and losses from the rise in wage dispersion?” by providing a critical overview of the literature, emphasizing the lessons we learned so far, and identifying the open issues. In the rest of this Introduction, we argue that this question is of first-order importance, we explain that answering it is a nontrivial task, and we summarize the key findings.

1.1 Why is this an important question?

First of all, volatility in idiosyncratic income growth in income is quantitatively very large, much larger than aggregate volatility —the input for calculations of the welfare cost of business cycles. To put things in perspective, note that the variance of the growth rate of individual earnings in the cross-section is between one and two orders of magnitude higher than the variance of the growth rate of average earnings at business cycle frequencies. Moreover, the recent increase in wage inequality is, historically, the largest and most prolonged in the postwar period. As a result, one should expect the welfare implications of rising wage dispersion to be quantitatively significant.²

Second, governments have a variety of instruments at their disposal for redistribution and social insurance: progressive taxation, unemployment compensation and welfare benefits, just to name a few. A deep understanding of the welfare consequences of the shift in the wage structure among US workers is a useful step before plunging into the debate on whether and how to intervene in order to equalize standards of living across households through the activation

²Incidentally, even though calculations of the welfare cost of business cycles based on representative agent models are tiny, they become somewhat larger once heterogeneity is introduced and once it is recognized that the dispersion of persistent income risk is countercyclical (see Lucas (2003) for a survey).

of large scale public policies.³

Third, the massive relative wage movements witnessed in the last forty years represent a unique opportunity –a macro natural experiment of sorts– to test our models of risk-sharing and consumption insurance and deepen our understanding of the mechanisms that insulate the distribution of consumption from labor income risk (Blundell and Preston (1998); Attanasio and Davis (1996); Blundell, Pistaferri, and Preston (2008); Kaplan and Violante (2010)).

Finally, the recent US experience provides a textbook example that growth and distribution are closely linked. Most of the measured productivity growth in the last thirty years spurred from the adoption in the workplace of new information technologies embodied in capital equipment (Greenwood, Hercowitz, and Krusell (2000); Jorgenson and Stiroh (1999); Cummins and Violante (2002)). Moreover, large gains from increased product variety and lower prices were achieved through increased openness to trade (Broda and Weinstein (2006)). These same forces –technological progress and globalization– are likely to be behind the observed demand shift away from unskilled labor and the rise in earnings volatility. Therefore, examining this question can also be intended as an investigation of the distributional costs of economic growth.

1.2 Why is the answer nontrivial?

As we will explain in more detail below, it will not suffice to compute the welfare effects of the rising inequality by just measuring the change in dispersion in earnings or consumption and evaluate it with some utility function. The welfare calculation is challenging because there are numerous economic forces at work that cause interactions between inequality and growth and, hence, must be quantified. Some of the underlying forces behind inequality induce welfare losses and others induce welfare gains. A central objective of this paper is to carefully decompose sources of gains and losses. To fix ideas, consider a simple example where the economy is subject to an unexpected, one-time mean preserving spread of the relative wage distribution with two components: a rise in the college premium and an increase in wage volatility within education groups.

Let's begin from the wider skill premium. Since wage differentials attributable to education are permanent and ex-ante uninsurable, they translate one for one into consumption differentials. College graduates are therefore better off and high-school graduates are worse off. However, this argument is incomplete because education is a choice. New cohorts can take

³For example, some commentators have argued that the home-ownership expansion policies of the 2000s were a political response to the lack of income growth for low-skilled households (e.g., Rajan (2010)).

advantage of the opportunities presented by the skill-biased demand shifts –and the associated larger return to education– by increasing their investment in human capital. This behavioral response can be a source of welfare gains as demonstrated by Heathcote, Storesletten, and Violante (2010a).

Consider now increased wage instability. The crucial determinant of the welfare effect of more volatile wages is whether this volatility is (self-) insurable or not. The uninsurable component of volatility transmits to consumption and reduces household welfare. But households have access to a multiplicity of smoothing channels to absorb wage fluctuations (e.g., offsetting individual and spousal labor supply responses, private transfers within networks of friends and extended family, borrowing and saving, progressivity of government redistribution), thus a sizeable component of wage fluctuations is insured.⁴ Interestingly, a rise in insurable wage dispersion is not welfare-neutral, but is welfare improving as long as workers can flexibly adjust their labor supply in response to wage changes (Heathcote, Storesletten, and Violante (2008)).

To summarize, this investigation faces the arduous task of quantifying the relative importance of two competing views in the public policy arena. The first view states that the rise in inequality is beneficial because it largely reflects higher returns to education. Individuals can take advantage of the new wage structure by increasing human capital investment. The second view argues instead that the rise in inequality is harmful because of the higher income instability and the fall in real wages at the bottom of the distribution. More and more households face the risk of very low incomes and poverty or, to maintain the same standard of living, they are forced to work longer hours.⁵

Finally, a methodological issue makes the analysis of this question especially challenging. We are interested in the impact of a shift in the wage distribution on welfare, and welfare is not a function of wages, but of consumption and leisure. How to link movements in relative wages to movements in relative consumption and leisure is the crucial step of the analysis. The literature has followed two strategies. The first is to look directly at the shift in the distribution of consumption (and leisure) in the micro data: we call this strategy the “empirical approach”.

⁴For recent surveys on the transmission of income shocks to consumption and the role of various smoothing mechanisms, see Blundell (2010) and Meghir and Pistaferri (2010).

⁵For example, Lazear (2006) wrote that “While there is no doubt that some people have been left behind,[...] the good news is that most of the inequality reflects an increase in returns to investing in skills - workers completing more school, [...] and acquiring new capabilities.” Conversely, Krugman (2005, 2007) wrote: “Over the past three decades the lives of ordinary Americans have become less secure, and their chances of plunging from the middle class into acute poverty ever larger [...] People aren’t nearly as much better off as they would be if the gains from economic growth had been broadly distributed.”

The second is to lay out a structural model to draw a mapping from wages to consumption and leisure: we call this strategy the “structural approach”. Both approaches have pros and cons which we discuss below.

1.3 A preview of the findings

The empirical and structural approaches yield different answers. Welfare calculations based on the empirical approach compare average utility derived from the empirical distribution of consumption and hours worked before and after the shift in the wage structure. Therefore, these calculations have the great virtue of only requiring assumptions on the specification and parameterization of preferences. Similarly to Krueger and Perri (2003), we estimate that comparing the distribution in 2001-2005 to the one in 1980-1984 (the earliest date available in the Consumer Expenditure Survey) results in a welfare loss of roughly 1.9 percent of lifetime consumption in our baseline.

However, the empirical strategy has a serious drawback: since in comparing the two empirical distributions the data are demeaned, this methodology abstracts from what we label “level effects” on welfare, i.e. effects on average consumption and leisure of those same forces which trigger the rise in dispersion: skill-biased demand shifts influence output through their increased human-capital accumulation, while rising wage volatility impacts productivity through modified labor supply decisions. Because these outcomes are the result of individuals’ optimal response to exogenous forces, only a structural micro-founded model can properly incorporate them. We then lay out a stripped-down version of the partial insurance model with endogenous education and labor supply choices developed in Heathcote, Storesletten, and Violante (2009a). The upshot of this equilibrium framework is that one can obtain a transparent closed-form expression for the welfare change from the shift in wage inequality that highlights all the sources of gains and losses separately.

Our key result is that, according to the model, the aforementioned gains dominate the losses arising from increased dispersion and imperfect consumption insurance. Overall, we find a welfare gain from the shift in the wage structure of 1.4 percent. Our counterfactual experiments indicate that investments in human capital as a response to the surging skill premium is the key source of this gain. When we counterfactually shut off this channel of adjustment, welfare losses re-emerge of similar sizes to those computed under the empirical methodology.

The rest of the paper is organized as follows. Section 2 sets the stage for our welfare

calculations by describing the facts on the changing inequality in wages, consumption and hours worked in the US. Section 3 gives an overview of the empirical approach, its advantages and its limitations. Section 4 lays out our structural economic model, its calibration and our model-based welfare calculation –here we arrive at an expression for the welfare change of rising wage inequality in closed form that can be decomposed in to all the critical forces at work. Section 5 contains some concluding remarks on open research questions and a reflection on public policy.

2 Setting the stage: data and facts

In this section we briefly discuss the data on wages, hours worked and consumption which we use throughout the paper. We start by describing the source of these data –the March Current Population Survey (CPS) and the Consumer Expenditure Survey (CEX), sample construction, and variable definition. Next, we present the salient facts on the evolution of cross-sectional dispersion of wages, hours and consumption.

2.1 The data

CPS The CPS is the source of official US government statistics on employment and unemployment, and is designed to be representative of the civilian non-institutional population. The Annual Social and Economic Supplement (ASEC) applies to the sample surveyed in March, and extends the set of demographic and labor force questions asked in all months to include detailed questions on income. For the ASEC supplement, the basic CPS monthly sample of around 60,000 households is extended to include an additional 4,500 Hispanic households (since 1976), and an additional 34,500 households (since 2002) as part of an effort to improve estimates of children’s health insurance coverage: this is the “SCHIP” sample. We use the March supplement weights to produce our estimates. Our CPS sample covers the period 1967-2005.

CEX The CEX consists of two separate surveys, the quarterly Interview Survey and the Diary Survey, both collected for the Bureau of Labor Statistics by the Census Bureau. Even though its main purpose is that of providing weights for the Consumer Price Index (CPI), it is the only US data set that contains detailed information about household consumption expenditures. The Diary Survey focuses only on expenditures on small, frequently purchased items (such as food, beverages, and personal care items), whereas the Interview Survey aims at providing information on up to 95% of the typical household’s consumption expenditures.

We will focus only on the Interview Survey, but we return to this point below. The CEX Interview Survey is a rotating panel of households that are selected to be representative of the US population. It started in 1960, but continuous data are available only from the first quarter of 1980, which is the start of our sample. Each quarter the survey reports, for the cross section of households interviewed, detailed demographic characteristics for all household members, detailed information on consumption expenditures for the three-month period preceding the interview, and information on income, hours worked, and taxes paid over a yearly period. Each household is interviewed for a maximum of four consecutive quarters. Our CEX sample covers the period 1980-2005.

Sample selection In both data sets we construct the sample following the same criteria, those outlined in Heathcote, Perri, and Violante (2010). From the raw data, we drop records 1) if there is no information on age for either the head or spouse; 2) if no household member is of working age, which we define as between the ages of 25 and 60; 3) if either the head or spouse has positive labor income but zero weeks worked; 4) if either the head or spouse has an hourly wage less than half of the corresponding federal minimum wage in that year. In the CEX, we also drop households whose quarterly equivalized food consumption is below \$100 in 2000 dollars and those flagged as “incomplete income reporters”.⁶

In all data sets, we forecast mean values for top-coded observations by extrapolating a Pareto density fitted to the non-top-coded upper end of the observed distribution. We apply this procedure separately to each component of income in each year. Throughout the paper, unless explicitly mentioned, we express all income and expenditure variables in year 2000 dollars. The price deflator used is the Bureau of Labor Statistics (BLS) CPI-U series, all items.

Variable definition Hours worked in the market are defined as total annual hours worked on all jobs. We define individual wage as annual individual earnings divided by annual hours worked, where annual earnings are defined as wage and salary income plus 2/3 of self-employment income. Our baseline measure of consumption includes expenditures on nondurables, services, small durables, and an estimate of the service flow from vehicles and housing. Household consumption expenditures are adjusted to a per-adult-equivalent basis using the OECD equivalence scale. The OECD scale assigns a weight of 1.0 to the first adult, 0.7 to each additional adult, and 0.5 to each child, defined as an individual age 16 or younger. See Heathcote, Perri, and

⁶Table 1 in the Appendix of Heathcote, Perri, and Violante (2010) summarizes the number of records in each data set that are lost at each stage of the selection process. Their Table 2 contains some summary statistics of the sample.

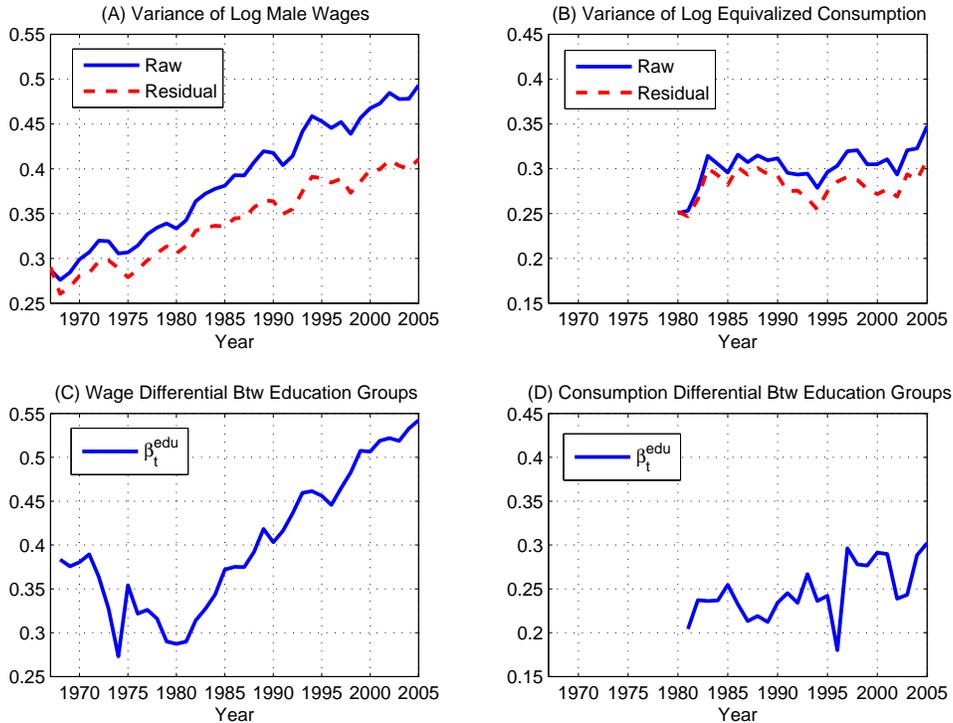


Figure 1: Evolution of inequality in male wages (CPS) and household consumption (Interview Survey of the CEX). Wages are computed as annual earnings (plus 2/3 of self-employment income) divided by annual hours worked. Consumption includes expenditures on nondurables, services, small durables and an imputed flow from vehicles and housing. Consumption is equivalized based on the OECD scale.

Violante (2010) for more details on the sample construction and the variable definition.

2.2 The facts

Panel (A) of Figure 1 plots two lines. The solid line is the variance of log wages (w_{it}) for male workers in the US from 1980 to 2005. Wage inequality rises steadily throughout the period. We focus on male wages to avoid selection issues, but Figure 4 in Heathcote, Perri, and Violante (2010) shows that, perhaps surprisingly, the upward trend in log wage inequality is virtually the same for women.

The dashed line depicts residual (or within-group) wage inequality estimated from the regression

$$\ln w_{it} = D_t + \beta_t^{edu} D_i^{edu} + f(\text{age}_{it}; \beta^{age}) + \varepsilon_{it}, \quad (1)$$

where D_t is a year dummy, D_i^{edu} an education dummy equal to one if the individual has a college degree, and $f(\cdot)$ is a quartic in age. Residual wage inequality is measured as the variance of ε_{it} . Also residual wage dispersion rises steadily over the period. A comparison with the “raw”

variance of wages reveals that the within-group component accounts for about two-thirds of the increase in cross-sectional male wage dispersion since 1980. By design, the remaining one-third is explained by the skill premium: panel (C) plots the value of β_t^{edu} and shows the well known surge in the return to investment in education over this period.

Panels (B) and (D) plot the corresponding variables for equivalized household consumption expenditures (c_{it}). The first noticeable feature of these plots is that, quantitatively, the rise in the variance of log consumption is much less pronounced than the rise in the corresponding inequality measure for wages –less than half (Slesnick (2001); Krueger and Perri (2006)). Second, the increase in the within-group component of consumption dispersion accounts for a smaller part of the increase compared to wages.⁷ Third, education consumption differentials stayed at roughly 2/3 of education wage differentials throughout the period. Put differently, the wider education wage gap has largely translated into wider consumption dispersion, whereas larger within-group wage volatility had a much milder impact on consumption inequality.⁸ Both facts have been emphasized before by Attanasio and Davis (1996) and Krueger and Perri (2003).

Inequality in male and female market hours worked (h_{it}), and its components, are reported in Figure 2. Male wage dispersion is counter-cyclical, but exhibits no obvious long-run trend, whereas female wage dispersion declines significantly. This decline in female hours dispersion toward the level for men reflects the rise in their average hours worked and the fact that more and more women work full time. Education explains virtually nothing of the hours differentials among both men and women, as visualized by the substantial overlap of residual and raw variances. This, together with the fact that the education component of the variance remained flat during this period, while the skill premium doubled, is a strong indicator that household preferences over consumption are close to logarithmic.

2.2.1 Some measurement issues

Consumption It is well known that aggregate consumption expenditures computed from the CEX are lower than Personal Consumption Expenditures (PCE) in the National Income and Product Accounts (NIPA) for a wide number of comparable expenditure categories. More disturbingly, the gap between the two series has grown larger over time. For example, for a

⁷We do find though that this within group component has increased over time, as opposed to Krueger and Perri (2003) who report a decline from 1972 to 2000.

⁸The fact that permanent consumption differentials by education are smaller than permanent income differentials is consistent with an overlapping-generations, incomplete markets model with finite horizon, progressive social security system and wealth accumulation. See, for example, Storesletten, Telmer, and Yaron (2004) and Kaplan and Violante (2010).

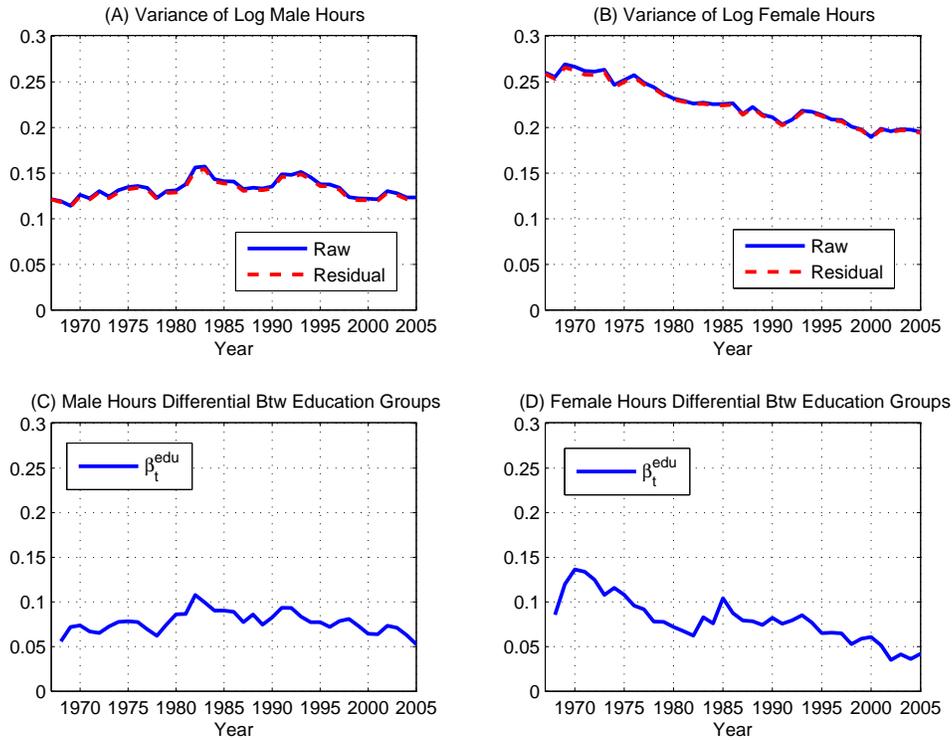


Figure 2: Evolution of inequality in male and female hours worked in the market (CPS).

broad definition of non-durable consumption, the gap grows from 20% in 1980 to 60% in 2005 (Figure 3 in Heathcote, Perri, and Violante (2010)).⁹ This growing discrepancy between survey mean and actual mean casts some doubt on the measurement of inequality trends as well. A number of studies has investigated the reliability of the survey-based consumption inequality statistics by trying to obtain alternative estimates.

Attanasio, Battistin, and Ichimura (2007) note that the Diary Survey (DS) of the CEX is better designed than the Interview Survey (IS) to measure expenditure in goods and services which are frequently purchased (e.g., food, personal care, housekeeping services). The DS, available only from 1986, shows a rise in consumption inequality which is larger than that emerging from the IS. Attanasio, Battistin, and Ichimura (2007) and Attanasio, Battistin, and Padula (2010) combine the two surveys by choosing, for each item, the one reporting

⁹The investigation on the sources of this discrepancy between survey-based and NIPA aggregate consumption is ongoing (Slesnick (2001); Garner, Janini, Passero, Paszkiewicz, and Vendemia (2006)). Conceptual differences between the CEX and the NIPA can account for some of the discrepancy. For example, among medical care expenditures, a rapidly growing item in the NIPA consumption, the BEA includes expenditures by Medicare, Medicaid, and private insurers, whereas the CEX reports only out-of-pocket expenses. However, the growing gap between the CEX and the NIPA applies across a broad range of consumption categories, suggesting that specific definitional differences are only part of the explanation. Another candidate explanation is that the CEX sample under-represents the upper tail of the income and consumption distributions, and that growth in aggregate consumption has been largely driven by these missing wealthy households.

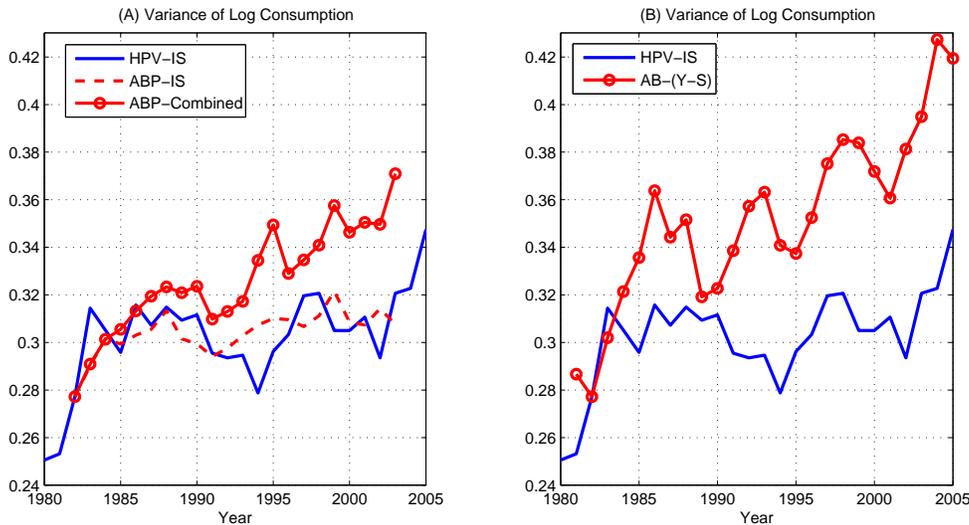


Figure 3: Evolution of inequality in equivalized household consumption (CEX). Panel (A) reports the Attanasio-Battistin-Padula (ABP) estimates obtained combining Diary and Interview Survey with their Interview Survey estimate and the Interview Survey in Heathcote-Perri-Violante (HPV). Panel (B) plots the APV series against the series computed by Aguiar and Bills from disposable income minus reported savings.

expenditures more accurately. Panel (A) in Figure 3 plots the IS-based and the IS-DS combined estimates of the variance of log consumption from Attanasio, Battistin, and Padula (2010). The latter series displays an increase almost twice as large over the period 1982-2003, with most of the discrepancy occurring after 1990. In the same figure we also plot the series from Heathcote, Perri, and Violante (2010) that we use in all our baseline calculations. The increase in consumption inequality in this series is comparable to the IS-based series of Attanasio, Battistin, and Padula (2010).

Several authors (e.g., Fisher and Johnson (2006); Blundell, Pistaferri, and Preston (2008); Guvenen and Smith (2010)) have imputed total consumption for households in the PSID based on the expenditure items common in the PSID and the CEX (e.g., food and rent), income and household demographics. Table 3 of Fisher and Johnson (2006) reports a rise in the Gini coefficient of their imputed PSID measure which is almost twice as large as the CEX counterpart between 1984 and 1999.

Aguiar and Bills (2010) exploit the reported amount of active savings and disposable income in the CEX to construct, under a number of assumptions, the measure of consumption residually implied by the household budget constraint. Under this methodology, consumption inequality tracks income inequality closely between 1980 and 2007 showing, once again, a significantly greater increase (roughly twice as large) than what is obtained from the IS-based household expenditure data. Panel (B) of Figure 3 compares this series to the baseline.

While more research is necessary to carefully establish the true dynamics of consumption inequality, these alternative measures seem all to indicate a sharper increase than what initially found from the IS of the CEX. In the meantime, it is useful to exploit alternative data which can be informative of changes in well being, e.g. data on hours worked.

Market hours vs leisure Hours worked in the market are correlated with well being, albeit only imperfectly. Leisure is, theoretically, a better indicator for welfare since it nets out from the time endowment hours spent in the production of market goods as well as home goods. However, leisure is much more difficult to measure properly than market hours because of the lack of detailed data on home production in surveys such as CPS and the CEX, which collect data on income, market hours and consumption expenditures. From time use surveys, Aguiar and Hurst (2009) exploit some limited information on leisure inequality by measuring the difference in leisure across education groups. They study how this measure evolves between 1985 and 2005 and find that over this period, less educated men have increased leisure by 2% and more educated men have decreased their leisure by a similar amount (see their Tables 2, 4A and 5A). In conclusion, the distribution of leisure for men has not changed dramatically, mirroring our finding for market hours, and hence using market hours for welfare calculations is reasonable. For women, instead, the story is quite different. Knowles (2008) shows that from 1975 to 2003 women increased their hours in the market and reduced their hours worked at home, without changing the fraction of the time endowment devoted to leisure. This finding suggests that using female market hours in the welfare calculation would artificially dampen the welfare costs of rising inequality. Because of this fact and the lack of comprehensive information on home work in CEX, we use only male hours in our welfare calculations.

Inflation inequality All the conventional measures of inequality deflate wages, income and consumption across individuals by the same price index –a choice akin to assuming that the bundle of goods consumed is not too different across households at any point in time and that all households pay the same price for the same good or service. A number of recent papers has challenged this view and showed that inflation rates are unequal across income groups.

A vital source of information on inflation heterogeneity is the Homescan database collected by AC Nielsen that records prices and quantities of purchases of several nondurable goods for thousands of US households, and contains detailed demographic information for these same households. From this data set, Broda and Romalis (2009) compute inflation rates by income percentile and find that the annual inflation rate at the 10th percentile has been 0.7% lower

than at the 90th percentile over 1994-2005 (see also Broda, Leibtag, and Weinstein (2009)).¹⁰. On the one hand, these findings could suggest that the adverse welfare effects of rising earnings inequality may have been mitigated by equalizing movements in consumption-good prices. Indeed, Broda and Weinstein (2008) argue that the stark, and commonly reported, finding that real wages at the bottom 10 percentile of the distribution have not risen in real terms since the late 1970s (e.g., Figure 2 in Acemoglu (2002)) is inaccurate: when properly deflated, wages at the 10th percentile rose by 30% from 1979-2005. On the other hand, these findings could reflect the fact that the rich purchase their goods in high-quality shops with better service, or that they devote less time to searching for low price stores. Such shopping pattern should in our view be interpreted as a component of consumption, which complicates the picture even further.

The AC Nielsen data set has limited coverage. In particular, it misses housing services, an important category of household expenditures contributing to a quarter of total nondurable consumption and services. Moretti (2010) shows that from 1980 to 2000, more educated individuals have experienced relatively larger rise in cost of living because they have increasingly concentrated in metropolitan areas characterized by a high housing prices. Deflating nominal wages using a location-specific CPI, over 1/5 of the documented increase in the college premium vanishes.¹¹ Hence, this correction goes in the same direction as for Broda and Romalis (2009). Again, the same caveat of quality differences applies: perhaps higher housing prices in cities simply reflects a growth in the quality of city living (e.g., due to lower crime and larger selection of goods and services than in the 1980s), in which case the housing prices reflect the real consumption value.

Taken together, this body of evidence suggests that some caution should be applied when interpreting the changes in real consumption dispersion.

¹⁰The differential inflation rates across income groups stems both from the different basket composition and the different price paid for the same item. Using the BLS item-specific price indexes (hence common across individuals) and individual-specific expenditure shares from CEX, Hobijn, Mayer, Stennis, and Topa (2009) construct price-indexes for different education groups and estimate that, over the period 1985-2005, the inflation rate for college graduates was 7% lower than for high school graduates. This finding seems to suggest that the lower inflation rates for the income poor documented by Broda and Romalis might be due to the rich paying higher prices for the same items.

¹¹See Van Nieuwerburgh and Weill (2010) for an equilibrium model that delivers this relationship between wages and house prices.

3 The empirical approach

The most direct approach to quantifying the welfare effects of rising wage inequality is that of plainly using observations from survey data on the empirical distribution of consumption and hours worked, the two key arguments of households' utility.¹² Recently, Jones and Klenow (2010) have used a very similar strategy to assess the historical change in welfare across countries and contrast it to the change in GDP, a more traditional measure of growth in well being.

This approach makes the implicit assumption that all the empirical changes in the dispersion of consumption and hours were driven by the shift in the wage structure. In Heathcote, Storesletten, and Violante (2010a), we build a structural dynamic model of the US economy and calibrate it based on household survey data. The calibrated model, with the observed shift in the wage structure as the only input, is able to reproduce the salient trends in the empirical cross-sectional distributions of individual hours worked, household earnings, and household consumption – all endogenous outcomes of the model. Therefore, we conclude that the assumption underlying the empirical approach – that the changes in wage structure explains the dispersion in consumption and leisure – is, to a great extent, vindicated.

In the rest of this section, we describe how different authors have implemented this approach and then we report some findings based on our own calculations.

3.1 Implementation

Comparing distributions of allocations, the thrust of this empirical strategy, requires only a minimal set of assumptions. To fix ideas, consider an overlapping generations economy where, every period, a measure $(1 - \pi)$ of agents is born and a corresponding measure dies. The total population is stationary with measure one. Let $\{\mathbf{c}_i, \mathbf{h}_i\}$ be the lifetime sequence of consumption and hours worked faced by household i and let $U_{t+k}(\mathbf{c}_i, \mathbf{h}_i)$ be the preferences used by birth cohort $t + k$ to evaluate this allocation, i.e.,

$$U_{t+k}(\mathbf{c}_i, \mathbf{h}_i) = \sum_{j=0}^{\infty} (\beta\pi)^j u(c_{i,j,t+k+j}, h_{i,j,t+k+j}) \quad (2)$$

where β is the discount factor, π is the survival rate, and $(c_{i,j,t+k+j}, h_{i,j,t+k+j})$ denotes individual realizations of consumption and hours worked for household i of age j at date $t + k + j$. Define the following Benthamite social welfare function to aggregate utilities across all cohorts alive

¹²See Slesnick (1998) for a survey on the empirical approach to the measurement of welfare.

and still unborn at date t :

$$\mathcal{W}(\mathbf{c}, \mathbf{h}) = \sum_{k=-\infty}^{\infty} \mu_k \int U_{t+k}(\mathbf{c}_i, \mathbf{h}_i) di \quad (3)$$

where $\{\mathbf{c}, \mathbf{h}\}$ is the distribution of lifetime sequences of consumption and hours, and μ_k is the weight on cohort $t+k$ (all the households within a cohort are equally weighted). Let “*” denote the distribution before the shift of the wage structure, and “**” denote the one post shift of the wage structure. We are interested comparing an economy with the “*” allocation forever to an economy with the the “*” allocation until date t and with the “**” allocation forever after. Then, the average welfare effect of rising inequality is defined as the scalar ω that solves

$$\mathcal{W}_t((1 + \omega) \mathbf{c}^*, \mathbf{h}^*) = \mathcal{W}_t(\mathbf{c}^{**}, \mathbf{h}^{**}), \quad (4)$$

where the subscript t on the welfare function \mathcal{W} indicates that the only relevant arguments of \mathcal{W} are the utility terms from t onward, since those indexed before t are common in right-hand side and left-hand side of equation (4) and drop out of the welfare calculation. A negative value for ω represents the fraction of consumption an individual would be willing to give up, in each state at each date, in order to avoid the shift in the distribution of consumption and hours induced by the new wage structure.

The empirical approach must deal with three distinct issues. First, in directly comparing two distributions of consumption and hours at two different points in time one has to deal with the fact that average consumption growth makes the final allocation a better one and this first-order effect is likely to dominate changes in second moments which occurred during the same period. Authors have dealt with this issue by demeaning the data (or equivalently, rescaling the final distribution so that it has the same mean as the initial one). However, demeaning also purges potential level effects (i.e., growth effects on the aggregate level of consumption and leisure) induced by all those same forces that shape wage dispersion. Only through the lens of a model can one identify and measure these level effects. We return on this point in Section 4, when we discuss the “structural approach”.

Second, one must estimate the distribution of lifetime sequences of consumption and hours $\{\mathbf{c}, \mathbf{h}\}$ at two distinct points in time, before and after the shift in the wage structure. Krueger and Perri (2003) exploit the short panel dimension of CEX (one year) and estimate a finite state Markov chain for log consumption and log hours where the transition probabilities across quantiles are time-invariant, but quantiles are allowed to vary over time to reflect the movements in cross-sectional dispersion. As emphasized by Davis (2003) and Storesletten (2003),

a shortcoming of this approach is that the estimated persistence of consumption and hours worked –and hence the estimate of the welfare cost– is likely to be mismeasured because of the extremely short panel dimension in the CEX and because of the large measurement error, known to plague reports of hours worked and expenditures in household surveys.¹³ Even though in theory, welfare calculations are correct as long as the size of measurement error is time invariant, in practice, reporting error, together with the small sample size, makes it very hard to discern precise trends in the data. In sum, the lack of high quality longitudinal data on consumption in the CEX undermines the estimation of a household-level stochastic process.

Attanasio and Davis (1996) chose to circumvent this problem by focusing on the relative movements of wages and consumption across observationally distinct groups. This choice allows the simultaneous use of the best survey data for consumption (CEX) and the best survey data for income (CPS). Their key finding is that persistent changes in relative wages among birth cohort-education groups lead to roughly equal-size changes in the distribution of consumption expenditures. Put differently, the rise in the skill premium translated one for one into consumption differentials between more and less educated households. A drawback of this methodology is that it abstracts from changes in the within-group component of wage dispersion that, as shown in Figure 1, is non negligible.

There is a third way to deal with this issue that allows avoiding the estimation of a stochastic process while, at the same time, retaining within-group variation. It requires a particular choice for the cohort-specific weights μ_k in the social welfare function (3). If the weights are chosen such that $\mu_k = (1 - \beta)(1 - \pi)\beta^{-k}$, then it is easy to see that the social welfare function simplifies to

$$\mathcal{W}_t(\mathbf{c}, \mathbf{h}) = \sum_{k=-\infty}^0 s^k \int u(c_{i,k,t}, h_{i,k,t}) di, \quad (5)$$

where $s^k = (1 - \pi)\pi^k$ is the share of population of age k . The expression in (5) is therefore simply average period utility among all the living individuals at date t .¹⁴ Then all that is needed is the cross-sectional joint distribution of consumption and hours, without any information on individual dynamics – a much less demanding data requirement.

¹³For example, Cogley (2002) suggests that measurement error in CEX consumption biases upward the true variance in individual consumption growth by one order of magnitude. Similarly, Heathcote, Perri, and Violante (2010) find that measurement error accounts for as much as 1/4 of the total variance of log consumption and, clearly, a much bigger share of the within-group component.

¹⁴As explained above, in the welfare function there are additional period-utility terms (deriving from components U_{t+k} for $k < 0$) which are all indexed by dates before t so they do not have any impact on our welfare comparisons and we ignore them.

Finally, a specification for period utility $u(\cdot)$ must be chosen to rank sequences of consumption and hours. This is the only model ingredient needed for the welfare calculation. In particular, since this approach does not try to draw a mapping between wages on the one hand and consumption and hours on the other, no assumptions have to be made on market structure, risk-sharing possibilities, technology or agent’s choice sets. In what follows, we assume the intra-period utility function

$$u(c, h) = \frac{c^{1-\gamma}}{1-\gamma} - \varphi \frac{h^{1+\sigma}}{1+\sigma}, \quad (6)$$

which has the advantage of being defined over consumption and hours and, as such, it avoids the problems in the measurement of leisure discussed above. The parameter γ is the inverse of the intertemporal elasticity of substitution for consumption. The parameter σ captures aversion towards hours fluctuations and $1/\sigma$ measures the Frisch elasticity of labor supply. The preference weight φ captures the strength of an individual’s distaste for work relative to his preference for consumption.¹⁵

3.2 Results

We now put the empirical approach to work in order to quantify the welfare change associated with the shift in the US wage structure. We choose the first and last five years (1980-1984 and 2001-2005) available in our CEX data to represent the joint distribution of consumption and hours worked before and after the shift (i.e., the “*” and the “**” allocations), respectively. We rescale the distributions of consumption and hours in 2001-2005 so that they have the same mean as in 1980-1984.¹⁶ We present three alternative implementations of the empirical approach.

An Atkinson-style calculation In the spirit of Atkinson (1970) and Storesletten (2003), in our first calculation we use the actual realizations of the joint distributions of consumption and hours observed in the CEX. Let I^* be the number of individuals in the 1980-1984 surveys and I^{**} be the number in the 2001-2005 surveys. Then, given the utility specification (6), and

¹⁵In previous work (Heathcote, Storesletten, and Violante (2008)) we have also used a Cobb-Douglas specification for some similar welfare calculations. The advantage of the separable specification in (6) over Cobb-Douglas is that two distinct parameters (γ, σ) regulate the two key elasticities. The disadvantage is that the calibration of the weight φ , in a model with heterogeneity like ours, is not straightforward.

¹⁶To minimize the effect of outliers, we trim the top and bottom 0.5% of the consumption and hours distributions.

Table 1: Atkinson-style Welfare Calculation

	Consumption Equivalent Variation (ω)				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-0.017	-0.037	-0.053	-0.062	-0.064
$\sigma = 2$	-0.018	-0.038	-0.054	-0.063	-0.064
$\sigma = 3$	-0.019	-0.038	-0.054	-0.063	-0.064
$\sigma = 4$	-0.019	-0.039	-0.055	-0.064	-0.065
$\sigma = 5$	-0.019	-0.039	-0.055	-0.064	-0.065

the social welfare function (5), equation (4) becomes

$$\frac{1}{I^*} \sum_{i=1}^{I^*} \left\{ \frac{[(1+\omega)c_i^*]^{1-\gamma}}{1-\gamma} - \varphi \frac{(h_i^*)^{1+\sigma}}{1+\sigma} \right\} = \frac{1}{I^{**}} \sum_{i=1}^{I^{**}} \left\{ \frac{(c_i^{**})^{1-\gamma}}{1-\gamma} - \varphi \frac{(h_i^{**})^{1+\sigma}}{1+\sigma} \right\}. \quad (7)$$

Table 1 reports the values of ω that solve equation (7) for different levels of risk aversion (γ) and Frisch elasticity ($1/\sigma$).¹⁷ In the range $\gamma = 1, \dots, 5$ and $\sigma = 1, \dots, 5$ the welfare losses from the shift in the wage structure vary between 1.7% and 6.5% of lifetime consumption, in line with the findings of Krueger and Perri (2003). As expected, welfare losses increase steeply in γ . The slope with respect to σ is much flatter because, as displayed in Figure 1, the variance of male log hours is basically constant over time.

The upshot of these calculations is that, according to the empirical approach, welfare losses are large. To put these estimates in perspective, recall that the Lucas (1987) seminal calculation of welfare gains from eliminating business cycles in a representative agent economy with log utility is 0.008%. More recently, Krusell and Smith (1999) and Krusell, Mukoyama, Sahin, and Smith (2009) revisited this calculation in an incomplete-markets model with idiosyncratic income risk correlated with aggregate risk and report that the average welfare gain from eliminating cycles is around 0.1% of consumption. Taken together, these calculations reveal that, under the veil of ignorance, US households would be willing to pay at least 10 times more to

¹⁷To calibrate φ , we proceed as follows. Our structural model of Section 4.1 implies the budget constraint

$$c_{it} + a_{i,t+1} - a_{it} = \lambda_t (w_{it} h_{it} + r a_{it})^{1-\tau_t}$$

where the right hand side is disposable income $\tilde{y}_{it} = \lambda_t y_{it}^{1-\tau_t}$ at date t and y_{it} is pre government income. From the utility function specification in (6), the intra-temporal first-order condition for hours worked yields

$$c_{it}^{-\gamma} \lambda_t (1-\tau) y_{it}^{-\tau_t} w_{it} = \varphi_i h_{it}^\sigma.$$

Given a pair (γ, σ) , some externally calibrated values of λ_t and τ_t , and data on (c_{it}, h_{it}, w_{it}) from each individual CEX record, we obtain residually the value for φ_i which allows this condition to hold with equality. From the implied distribution of φ_i , we estimate the median and use it in the welfare calculations in this section.

Table 2: Atkinson-style Welfare Calculation

	CEV (ω) adjusted for inflation differentials				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-0.005	-0.011	-0.017	-0.018	-0.014
$\sigma = 2$	-0.006	-0.013	-0.018	-0.019	-0.015
$\sigma = 3$	-0.006	-0.013	-0.019	-0.020	-0.015
$\sigma = 4$	-0.006	-0.013	-0.019	-0.020	-0.015
$\sigma = 5$	-0.006	-0.013	-0.019	-0.020	-0.015

avoid another rise in wage inequality similar to the one witnessed over the last 30 years than to avoid another recession.

As we discussed in Section 2.2.1, Broda and Romalis (2009) argue that inflation may have evolved differentially across income groups, with the households at the top experiencing the highest price increase. Such consideration would affect our welfare calculations. Figure 7c in Broda and Romalis (2009) plots inflation rates by income percentile between 1994 and 2005 and documents that, when abstracting from the higher shopping-quality services enjoyed by the rich, the inflation for the rich households has been around 0.07% per year higher than for poor households. When we deflate household consumption expenditures in CEX with the Broda and Romalis percentile-specific inflation rates, we find significantly smaller welfare losses, as seen in Table 2.¹⁸ For example, for $\gamma = 1$ and $\sigma = 2$, the welfare loss falls from 1.8% to 0.6% once the differential inflation adjustment is taken into account.

An Attanasio and Davis-style calculation In the spirit of Attanasio and Davis (1996), we perform an alternative exercise. We group individuals by education level (with and without a college degree) and age (25-34, 35-44, 45-54, 55+). In Table 3, we repeat the welfare calculations of Table 1 by using the eight groups, appropriately weighted, instead of the individuals as the unit of analysis: the implied welfare effect reported in Table 3 is determined only by the shift in the distribution of consumption and hours between groups, but it abstracts from the change in the within-group component.

We find that welfare losses are roughly half of those in Table 1. For example, in the parameterization $\gamma = 1$ and $\sigma = 2$, $\omega = 0.9\%$ instead of 1.8%. This finding is consistent with

¹⁸Since the Broda and Romalis data start in 1994, we assume that inflation grew at the same rate across all income groups in the period 1980-1994. Under the alternative assumption that the inflation differentials measured after 1994 took place even during 1980-1994, the welfare losses would turn to welfare gains since the variance of log consumption shrinks, in real terms. For example, in the case $\gamma = 1$ and $\sigma = 2$, the gain would be 1.1 percent.

Table 3: Attanasio & Davis-style Welfare Calculation

	Consumption Equivalent Variation (ω)				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-0.008	-0.017	-0.024	-0.031	-0.036
$\sigma = 2$	-0.009	-0.017	-0.025	-0.031	-0.037
$\sigma = 3$	-0.009	-0.017	-0.025	-0.032	-0.038
$\sigma = 4$	-0.009	-0.018	-0.026	-0.033	-0.039
$\sigma = 5$	-0.009	-0.018	-0.026	-0.033	-0.039

the fact, documented in Figure 1, that the rise in between-group consumption inequality is approximately half of the total.

A Lucas-style calculation Lucas (1987) showed that by assuming lognormality of the stochastic process for consumption, and holding the mean consumption constant across experiments, one can arrive at an analytical and intuitive expression for the welfare cost of business cycles. Here, we follow this approach. Let v_x denote the variance of the random variable x and suppose that log consumption and log hours are distributed as

$$\begin{aligned}\ln c_i^* &\sim N(\mu_c^* - v_c^*/2, v_c^*) \\ \ln h_i^* &\sim N(\mu_h^* - v_h^*/2, v_h^*)\end{aligned}$$

before the shift in the wage structure and

$$\begin{aligned}\ln c_i^{**} &\sim N(\mu_c^{**} - v_c^{**}/2, v_c^{**}) \\ \ln h_i^{**} &\sim N(\mu_h^{**} - v_h^{**}/2, v_h^{**})\end{aligned}$$

after the shift. Then, it is easy to show that, with the utility specification (6) and the social welfare function (5), equation (4) yields

$$\begin{aligned}&(1 + \omega)^{1-\gamma} \frac{1}{1-\gamma} \exp\left((1-\gamma)\mu_c^* - \gamma(1-\gamma)\frac{v_c^*}{2}\right) - \frac{\varphi}{1+\sigma} \exp\left((1+\sigma)\mu_h^* + \sigma(1+\sigma)\frac{v_h^*}{2}\right) \\ &= \frac{1}{1-\gamma} \exp\left((1-\gamma)\mu_c^* - \gamma(1-\gamma)\frac{v_c^*}{2}\right) - \frac{\varphi}{1+\sigma} \exp\left((1+\sigma)\mu_h^* + \sigma(1+\sigma)\frac{v_h^*}{2}\right)\end{aligned}\quad (8)$$

once we normalized the mean consumption and hours in the post shift period to be the same as in the pre-shift period (i.e., $\exp(\mu_x^{**}) = \exp(\mu_x^*)$). Table 4 reports estimates of ω based on this Lucas-style approach. Reassuringly, the estimates of ω are very similar to those in Table 1, especially for moderate utility curvature on consumption.

Table 4: Lucas-style Welfare Calculation

	Consumption Equivalent Variation (ω)				
	$\gamma = 1$	$\gamma = 2$	$\gamma = 3$	$\gamma = 4$	$\gamma = 5$
$\sigma = 1$	-0.018	-0.037	-0.056	-0.076	-0.095
$\sigma = 2$	-0.018	-0.037	-0.056	-0.076	-0.095
$\sigma = 3$	-0.018	-0.037	-0.056	-0.076	-0.095
$\sigma = 4$	-0.018	-0.037	-0.056	-0.076	-0.095
$\sigma = 5$	-0.018	-0.037	-0.056	-0.076	-0.095

As mentioned in Section 2.2.1, the Interview Survey of the CEX may underestimate the rise in consumption inequality. According to Attanasio, Battistin, and Padula (2010) and Aguiar and Bils (2010), the increase in the variance of log consumption obtained combining IS and DS of the CEX is larger than our baseline series computed from the IS (and plotted in Figure 1) by a factor of 2. Using Lucas' approach, one can easily analyze how this reassessment would affect the calculations in Table 4. Since the hours component has a trivial effect under the lognormality assumption, in what follows we ignore it. Abstracting from the hours component, and using a first-order Taylor expansion of type $\exp(x) \approx 1 + x$, from (8) we arrive at the cross-sectional counterpart of Lucas' expression:

$$\omega \approx -\frac{\gamma}{2}\Delta v_c, \quad (9)$$

where $\Delta v_c = v_c^{**} - v_c^*$ is the change in the cross-sectional dispersion of consumption between 1980-1984 and 2001-2005. Equation (9) reveals that, if the true rise in the variance of log consumption is twice as large as in our baseline case, the welfare loss would increase proportionately. For example in the log case ($\gamma = 1$) the consumption equivalent variation ω would be -3.6% of lifetime consumption instead of -1.8% .

3.3 What do we learn from the empirical approach?

The greatest advantage of the empirical approach is that it requires only a minimal set of assumptions on preferences and aggregation of individuals into a welfare function. In particular, no assumption on behavior or market structure is required.¹⁹ Its main drawback is that it is unable, by design, to assess the impact on aggregate consumption and leisure (i.e., the level effects) of those same forces which triggered the shift in wage dispersion.

¹⁹Strictly speaking, this is true only conditional on knowing the value of the relative disutility of hours φ . The calibration of φ requires assuming individuals are on their intratemporal optimality condition.

Different implementations of this approach led to fairly consistent results: in the baseline parameterization ($\gamma = 1$), the average welfare loss from the rise in wage inequality is just below 2 percent of lifetime consumption. Two important caveats are in order. First, inflation dynamics have been quite different across income groups over this period, with the income-rich suffering higher price increases. Second, various alternative calculations of the rise in consumption inequality point towards a larger increase than that measured in the Interview Survey of the CEX. The two corrections have different signs, but they are of the same order of magnitude –between 1 and 2 percent in the case $\gamma = 1$, $\sigma = 2$. Therefore, the 2 percent baseline estimate for the welfare loss can be considered robust.

4 The structural approach

The main reason to adopt a structural approach to welfare calculations is that we wish to quantify the effects that changes in the distribution of relative wages have on aggregate output and labor supply (the level effects): skill-biased demand shifts influence output through human-capital accumulation and larger wage volatility increases productivity through modified labor supply decisions.

To make progress, we need to select a model that offers a mapping between changes in the distribution of wages and changes in the distribution of consumption and labor supply. We choose a simplified version of the “partial insurance” framework described in Heathcote, Storesletten, and Violante (2009a). It is an incomplete-market model featuring the four key channels through which individuals can respond to shifts in the wage distribution: education, labor supply, participation in government redistribution schemes, and private risk sharing. In particular, beyond saving and borrowing through a risk-free asset, agents in our economy are able to perfectly insure a subset of idiosyncratic wage risk. This additional insurance is designed to capture a number of other adjustment mechanisms and institutions that spread risks across individuals or over time. Examples include spousal labor supply, hedging through financial instruments and risk-sharing within a range of networks including families, firms and unions. In the spirit of Deaton (1997), we do not model these mechanisms in detail, but we bundle them together and quantify their overall importance by looking at the residual gap between wage and consumption dispersion, once all the other smoothing channels explicitly modelled are taken into account.²⁰

²⁰Deaton (1997), pages 372–374, writes: “Saving is only one of the ways people can protect their consumption

A significant advantage in following this strategy is a degree of tractability lacking in standard incomplete-market models. Equilibrium allocations for consumption and hours can be obtained in closed form and, as a result, one can solve analytically for the welfare change from the shift in the wage structure. These analytical expressions for welfare reveal all the sources of gains and losses as a function of the structural parameters and can be easily related to the Lucas-style welfare expression (9).

Below, we describe the model formally and discuss in detail its key assumptions. We will focus on a comparison across two steady states. Thus, in describing the environment, we invoke the steady state assumption and drop time-subscripts.

4.1 A model economy with partial insurance

Demographics Time is discrete and continues forever. We adopt the Yaari perpetual youth structure: agents are born at age zero and survive from age j to age $j + 1$ with constant probability $\pi < 1$. A new generation with measure $(1 - \pi)$ enters the economy each period. Thus, the measure of agents of age j is $(1 - \pi)\pi^j$ and the total population size is unity.

Preferences Expected lifetime utility for agent i is given by

$$\mathbb{E}_0 \sum_{j=0}^{\infty} (\beta\pi)^j u_i(c_{ij}, h_{ij}), \quad (10)$$

where the expectation is taken over sequences of shocks defined below. Here c_{ij} denotes consumption at age j , and h_{ij} hours worked. Agents discount the future at rate $\beta\pi$, where $\beta < 1$ is the discount factor. Period utility for individual i is

$$u_i(c_{ij}, h_{ij}) = \ln c_{ij} - \exp(\bar{\varphi} + (1 + \sigma)\varphi_i) \frac{h_{ij}^{1+\sigma}}{1 + \sigma}, \quad (11)$$

a specification which is consistent with balanced long-run growth. The Frisch elasticity of labor supply is equal to $1/\sigma$. The disutility weight on hours worked has a common component $\bar{\varphi}$ and an idiosyncratic fixed component φ_i that is drawn once at the start of an agent's lifetime from the Normal distribution $\varphi_i \sim N(-(1 + \sigma)\frac{v_\varphi}{2}, v_\varphi)$.²¹ Note that this is the same preference specification assumed in Section 3.1 for all our computations based on the empirical approach,

against fluctuations in their income. An alternative is to rely on other people, to share risk with friends and kin, with neighbors, or with other anonymous participants through private or government insurance schemes, or through participation in financial markets." And: "Although it is possible to examine the mechanisms, the insurance contracts, tithes and transfers, their multiplicity makes it attractive to look directly at the magnitude that is supposed to be smoothed, namely consumption."

²¹In the utility function φ is multiplied by $(1 + \sigma)$. This is an innocuous normalization which simplifies the

except for the presence of dispersion in φ_i across agents. Preference dispersion is intended to capture various sources of heterogeneity which generate cross-sectional variation in hours worked and consumption that is independent of variation in productivity. While preference heterogeneity is important for the estimation of the key model's parameters, it plays no role in the welfare calculations.²²

Education There are two possible schooling levels attainable by individuals: high ($s = H$), corresponding to college degree or more, and low ($s = L$), corresponding to high-school degree or less. We keep the model for acquisition of education simple, in line with Heathcote, Storesletten, and Violante (2010a).²³ When they first enter the economy, before drawing their disutility of work φ_i , agents draw a utility cost of attending college χ_i from a lognormal distribution with mean μ_χ and variance v_χ . Taking as given prices, individuals attend college if the expected lifetime utility upon entry in the labor market as college graduates \mathcal{W}_H^0 net of the education cost χ_i , exceeds expected lifetime utility as high-school graduates, \mathcal{W}_L^0 .

Production The final good can be used for private consumption (C) and government consumption (G). Therefore, the aggregate resource constraint of this closed economy is, simply, $C + G = Y$. Aggregate technology is constant returns to scale with college and high school labor as the only inputs. Following a large literature (e.g., Katz and Murphy (1992); Heckman, Lochner, and Taber (1998)), we assume that this technology has the constant elasticity of substitution form

$$Y = \exp(z) \left[\zeta N_H^{\frac{\theta-1}{\theta}} + (1 - \zeta) N_L^{\frac{\theta-1}{\theta}} \right]^{\frac{\theta}{\theta-1}}$$

where N_H and N_L are aggregate effective hours worked by college and high-school-educated workers given by

$$N_s = \int \exp(\alpha_i + \varepsilon_i) h_i di,$$

and where θ is the elasticity of substitution between the two labor inputs. The weight parameter ζ is allowed to vary across steady states, reflecting skill-biased demand shifts, while the elasticity

expressions for equilibrium allocations. Given the distributional assumption, the expected weight on hours in preferences is given by

$$\mathbb{E}[\exp(\bar{\varphi} + (1 + \sigma)\varphi_i)] = \exp(\bar{\varphi}).$$

²²As shown in Heathcote, Storesletten, and Violante (2009a), matching both the wage-hour covariance and the consumption-hour covariance requires substantial cross-sectional preference dispersion in the relative weight on leisure. There, we also allow for stochastic shocks to the relative preference weight on leisure, but we found that on an annual frequency such shocks are negligible.

²³See also Guvenen and Kuruscu (2009) for an analysis of the trends in inequality based on a model with endogenous acquisition of human capital

of substitution parameter θ is fixed. Since labor markets are competitive, the price of an effective hour worked by an individual with schooling level s is its marginal product from the technology described above, or, in logs:

$$p_s = z + \frac{1}{\theta - 1} \log \left[\zeta N_H^{\frac{\theta-1}{\theta}} + (1 - \zeta) N_L^{\frac{\theta-1}{\theta}} \right] + \log \zeta - \frac{1}{\theta} \log (N_s),$$

and the equilibrium log skill premium is

$$p_H - p_L = \log \left(\frac{\zeta}{1 - \zeta} \right) - \frac{1}{\theta} \log \left(\frac{N_H}{N_L} \right). \quad (12)$$

Therefore, what drives the dynamics of the skill premium in the model are skill-biased demand shifts, i.e., an increase in ζ across steady-states.

Individual productivity shocks Individual hourly wages are equal to individual labor productivity (units of effective labor input per hour worked) times the price per effective unit of the schooling-type of labor the individual supplies, i.e. $w_{isj} = \exp(p_s) \exp(\alpha_{ij} + \varepsilon_{ij})$. The terms α_{ij} and ε_{ij} are stochastic components of the wage that are additive in log, orthogonal to each other, and orthogonal to the education component. We assume that the variance of wages, and its dynamics, are the same within each education group.²⁴ The component α_{ij} follows the random walk process

$$\alpha_{ij} = \alpha_{i,j-1} + \omega_{ij},$$

where the innovation ω_{ij} is drawn from the time-invariant (within a particular steady state) Normal distribution with variance v_ω . Agents entering the labor market at age $j = 0$, after the education decision, draw initial realizations α_{i0} from a Normal distribution with cohort-specific variance v_{α_0} . We assume that ε_{ij} is a purely transitory shock, i.e., i.i.d. over time with variance v_ε .²⁵ The statistical process for wages described above (unit root plus i.i.d. shocks) is quite standard in the literature and is consistent with the key features of individual wage dynamics as well as with trends in wage dispersion across the life cycle.²⁶

Finally, we normalize the means of the distributions for α_0 , ε and ω to negative one half their respective variances, which ensures that the average wage of type s workers is given by

²⁴Meghir and Pistaferri (2004) estimate earnings dynamics separately for three educational groups and do not find large differences among them.

²⁵In Heathcote, Storesletten, and Violante (2009a), we assume a richer process for ε , i.e. we let ε be comprised of a permanent shock and a transitory shock. Given our social welfare function, the dynamic properties of ε do not matter for the welfare calculations, so we make this simplifying assumption.

²⁶For example, the empirical autocovariance function for individual wages displays a sharp decline at the first lag, indicating the presence of a transitory component in wages. At the same time, within-cohort wage dispersion increases approximately linearly with age, suggesting the presence of permanent shocks.

$\exp(p_s)$. For the initial steady state we make an innocuous normalization which amounts to assuming that the average wage in the population is equal to one:

$$e^* \exp(p_H^*) + (1 - e^*) \exp(p_L^*) = 1. \quad (13)$$

For the final steady state we assume that, absent any change in enrollment, there would have been no change in average wages. Thus,

$$e^* \exp(p_H^{**}) + (1 - e^*) \exp(p_L^{**}) = 1. \quad (14)$$

This normalization is symmetric to the conditions imposed on the stochastic components for wages, where we always assumed that changes in the wage structure leave the average wage invariant and equal to one.²⁷ With this set of normalizations, we do not hardwire any level effect into the model: all changes in average output between steady states arise because of behavioral responses in the form of modified labor supply or education choices. Note that assumptions (13) and (14) can be re-interpreted as restrictions on the parameter z^* and z^{**} in the aggregate production technology.

Government The government consumes a fraction of output g at every date and finances these expenditures through a progressive tax system. If we let pre-government earnings of individual i be $y_i = w_i h_i$, then disposable earnings are

$$\tilde{y}_i = \lambda (w_i h_i)^{1-\tau} \quad (15)$$

and the government budget constraint is

$$gY = \int \lambda (w_i h_i)^{1-\tau} di. \quad (16)$$

This class of progressive tax functions, discussed in detail by Benabou (2002), is indexed by two parameters. The parameter $\tau \geq 0$ measures the degree of progressivity (with $\tau = 0$ representing a linear tax system with tax rate $1 - \lambda$), while λ captures the overall level of taxation and is determined residually from equation (16) assuming that the government balances its budget every period in equilibrium. Note that the system features marginal tax rates which are increasing in earnings and always generates a transfer (or a negative average income tax) for earnings below the threshold $\lambda^{\frac{1}{\tau}}$. As we show in Section 4.3 below, this class of tax functions provides a remarkable approximation to the actual US system.

²⁷Of course, in the model the distribution across education levels is endogenous, and the actual average wage in the final steady state is given by $e^{**} \exp(p_H^{**}) + (1 - e^{**}) \exp(p_L^{**})$.

Market structure Perfect annuity markets are available to insure against the risk of survival. Moreover agents can trade a risk-free bond and claims contingent on ε , both in zero net supply.²⁸ Finally, we assume that at birth each agent is endowed with zero financial wealth.

4.2 Stationary equilibrium

Nature of the equilibrium The distinctive property of the equilibrium for this economy is that agents choose not to trade the risk-free bond: in equilibrium, the expected marginal rate of substitution between consumption at dates t and $t + 1$ is the same for all households. From this observation, it is immediate to derive that the equilibrium risk-free rate r^* supporting this allocation satisfies the condition:

$$\rho - r^* = (1 - \tau)(2 - \tau) \frac{v_\omega}{2}, \quad (17)$$

where $\rho = \log(\beta)$ and $r = \log(R)$.²⁹ Recall that individuals in our model have two saving motives: an intertemporal motive given by the gap between the degree of patience and the interest rate and a precautionary motive reflecting the variance of permanent shocks to the α component (recall that the ε shock is insurable). In equilibrium, each saving motive applies with equal strength. The intertemporal motive is the same across agents because all agents share common discount factor and face the same economy-wide risk-free interest rate r^* . The precautionary motive is also identical, as a result of the assumptions that shocks are multiplicative, permanent, and drawn from common distributions, that preferences are in the power utility class with respect to both consumption and hours worked, and that all individuals start out with zero wealth. Because the strength of these two saving motives is identical, there exists an economy-wide interest rate r^* at which, in equilibrium, the (negative) intertemporal motive exactly offsets the (positive) precautionary motive, and no agent wants to either borrow or lend.

The closed-form expression for the interest rate helps us visualize this argument. The left-hand side of (17) measures the intertemporal motive to dissave, equal to the difference between the discount rate and the equilibrium interest rate. The right-hand side in (17) captures the precautionary motive for saving, which is proportional to the variance of the permanent shock

²⁸The set of contingent claims is meant to capture the additional private risk sharing discussed earlier. Instead of implementing full insurance through markets, we could have chosen to allow full consumption smoothing with respect to ε within a “large family”, for example, with no impact on equilibrium allocations.

²⁹The survival rate π does not appear in this equation because of the perfect annuity market assumption.

v_ω . Note that, when $\tau = 1$, $r^* = \rho$ as in any complete markets economy because the tax system provides full redistribution and shuts off any need for precautionary savings.

Standard incomplete-market economies do not admit an analytical solution, and numerical methods are required to solve for the equilibrium wealth distribution.³⁰ In our environment, wealth is not a state and the vector $(\varphi, s, \alpha, \varepsilon)$ contains sufficient information to fully describe equilibrium choices. The power of this result lies in the fact that these are all exogenous states and hence allocations –as well as welfare, which is particularly relevant for our application– can be derived in closed form.

This equilibrium, derived in a more general framework and discussed in detail in Heathcote, Storesletten, and Violante (2009a), has its roots in Constantinides and Duffie (1996). Their model is an endowment economy where agents face exogenous unit root shocks to individual log endowments. Since the risk-free asset is not traded, the equilibrium is “autarkic” with consumption equal to the endowment. In our model, instead, it is efficiency units of labor which are exogenous and contain the unit root component. Then, the equilibrium of the model endogenously maps efficiency units into wages (through education decisions and the aggregate technology), wages into earnings (through flexible labor supply), earnings into pre-government income (through private risk sharing), and pre-government income into disposable income (through progressive taxes and transfers). Disposable income retains a unit root, and hence the bond is not traded and consumption equals disposable income.

The generalization of Constantinides and Duffie (1996) is substantial because, while maintaining closed form solutions, our equilibrium is far from autarkic as it contains all the main channels of consumption smoothing.³¹ In particular, fluctuations in the ε component of wages can be fully smoothed thanks to participation to private risk sharing, whereas fluctuations in the α component can only be smoothed through participation in the government redistribution scheme.³² In what follows, we use the term “insurable shock” for ε and “uninsurable shock”

³⁰We use the term standard incomplete markets models to identify the framework developed by Huggett (1993), Aiyagari (1994), Rios-Rull (1995), and Krusell and Smith (1998), among others. See Heathcote, Storesletten, and Violante (2009b) for a survey of the most recent contributions.

³¹In Heathcote, Storesletten, and Violante (2009a) we explain that if individual productivity profiles are allowed to vary, in the style of Lillard and Weiss (1979) and Guvenen (2009) and it is assumed that this variation is known ex-ante by agents, then in equilibrium agents will save and borrow through the risk-free bond in zero net supply. This version of the economy is isomorphic to one where insurable shocks feature a unit root component and remains tractable.

³²In Heathcote, Storesletten, and Violante (2009a), we allow for risk aversion to be different from one. When it is larger than one, labor supply also acts akin to an insurance device against α shocks. The income effect dominates the substitution effect and agents increase hours worked in the wake of a negative shock to α thereby dampening the impact on earnings and consumption.

for α .

Equilibrium allocations The equilibrium allocations for consumption and hours in logs are:

$$\begin{aligned}\log c(\varphi, s, \alpha) &= \kappa_c + (1 - \tau)(p_s + \alpha - \varphi) \\ \log h(\varphi, \varepsilon) &= \kappa_h - \varphi + \frac{1 - \tau}{\sigma + \tau} \varepsilon\end{aligned}\tag{18}$$

where the constant terms in the allocations are given by

$$\kappa_c = \log \lambda - (1 - \tau) \bar{\varphi} + \frac{1 - \tau}{1 + \sigma} \log(1 - \tau) + (1 - \tau) \left(\frac{1 - \tau - \tau(1 + \sigma)}{\sigma + \tau} \right) \frac{v_\varepsilon}{2}\tag{19}$$

$$\kappa_h = -\bar{\varphi} + \frac{1}{1 + \sigma} \log(1 - \tau) - \frac{1 - \tau}{\sigma + \tau} \left(\frac{1 - \tau - \tau(1 + \sigma)}{\sigma + \tau} \right) \frac{v_\varepsilon}{2},\tag{20}$$

and the equilibrium value of λ that balances the government budget is

$$\begin{aligned}\log \lambda \approx & \log(1 - g) + \frac{\tau}{1 + \sigma} \log(1 - \tau) - \tau \bar{\varphi} + (3\tau - \tau^2 + \sigma\tau) \frac{v_\varphi}{2} \\ & + (\bar{p} - \bar{p}_\tau) + \tau(1 - \tau) \frac{v_\alpha}{2} + \frac{\tau(1 - \tau)}{\sigma + \tau} \left(\frac{1 + \sigma}{\sigma + \tau} + 2 + \sigma \right) \frac{v_\varepsilon}{2}\end{aligned}\tag{21}$$

with

$$\begin{aligned}\bar{p} &= \log [e \exp(p_H) + (1 - e) \exp(p_L)], \\ \bar{p}_\tau &= \log [e \exp((1 - \tau)p_H) + (1 - e) \exp((1 - \tau)p_L)].\end{aligned}$$

The consumption allocation reveals that, because of the redistribution imposed by the tax system, only a fraction $1 - \tau$ of uninsurable shocks and education wage differentials transmits into consumption. Heterogeneity in the taste for work effort φ translate into different hours worked and earnings one for one and into consumption proportionately to $1 - \tau$. The insurable productivity shock ε does not affect consumption because of the separability between consumption and hours in preferences.

Turning to hours worked, uninsurable shocks α and education wage differentials are not a source of hours dispersion because income and substitution effects in labor supply exactly offset each other with our preference specification. Insurable productivity shocks affect hours worked proportionately to the tax-modified Frisch elasticity $(1 - \tau)/(\sigma + \tau)$ which is lower than the pure preference-based Frisch elasticity $(1/\sigma)$ exactly because of the progressive tax system.

Note that, through the constant terms κ_c and κ_h , consumption and leisure are both increasing in the variance of the insurable shock. This force is behind some of the welfare gains

from the rise in wage inequality, as we explain formally in Section 4.4. Equation (21) makes it clear that our equilibrium model takes fully into account that changes in the wage structure $(p_H, p_L, v_\alpha, v_\varepsilon)$ affect the general level of taxation necessary to finance government consumption.³³

Turning to the education choice, it is immediate to derive that the difference in expected utility between being a college graduate and a high school graduate is

$$\mathcal{W}_H^0 - \mathcal{W}_L^0 = \frac{(1 - \tau)(p_H - p_L)}{1 - \beta\pi}. \quad (22)$$

The differential expected utility has a very simple form: it is the after-tax discounted present value of the log skill premium. Within-group components do not show up since they affect equally the value of both education levels. If we let $\chi^0 = \mathcal{W}_H^0 - \mathcal{W}_L^0$ be the threshold utility cost at which an individual is just indifferent about attending college, since χ is lognormally distributed, the equilibrium fraction of college educated workers is given by

$$e = \Phi\left(\frac{\ln \chi^0 - \mu_\chi}{\sqrt{v_\chi}}\right) \quad (23)$$

where Φ is the standard Normal CDF.

4.3 Parameterization

We now describe how we choose values for all the structural parameters of the model. To be consistent with the welfare calculations based on the empirical approach, we take 1980-1984 to be the pre-shift steady state (indexed by “**”) and 2001-2005 to be the post-shift steady state (indexed by “***”). The model’s period is set to one year.

To estimate the progressivity parameter τ , we exploit the fact that equation (15) implies a log linear relationship between post-government and pre-government earnings, or

$$\log \tilde{y}_{it} = \lambda_t + (1 - \tau_\tau) \ln y_{it}. \quad (24)$$

We estimate this relationship on the CPS separately for 1980-1984 and 2001-2005 for households with positive labor income.³⁴ We obtain $\tau^* = 0.36$ in the first period and $\tau^{**} = 0.28$ in the second period. In what follows, we use the mean value $\tau = 0.31$ and return to this estimated

³³Equation (21) holds with approximate equality. The approximation involves using the Taylor expansion $\exp(x) \approx 1 + x$. All our numerical examples showed that it is extremely accurate.

³⁴Pre-government income y_{it} is defined as the sum of earnings, asset income and private transfers. Disposable income \tilde{y}_{it} is pre-government income plus public transfers (reported) minus taxes (imputed). See the CPS Appendix in Heathcote, Perri, and Violante (2010) for details.

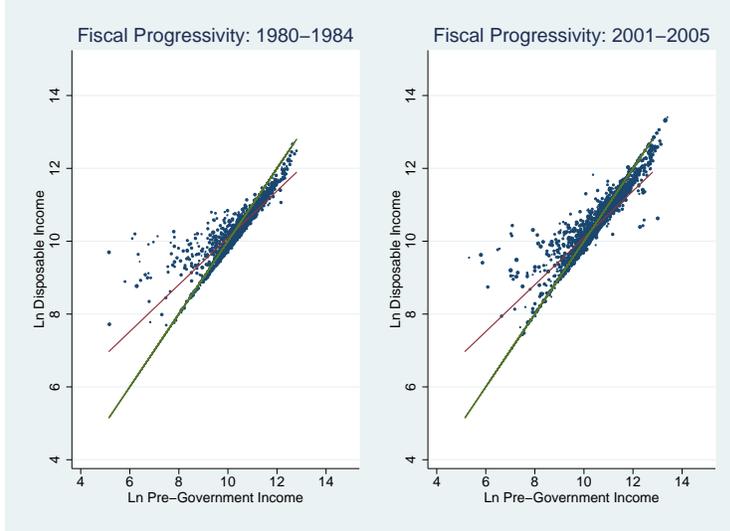


Figure 4: Results of the estimation of equation (24) on the CPS data. The scatterplot represents pairs of log post and pre government income for each household in the sample. The plot also displays the 45 degree line and the regression line with slope $(10 - \tau)$.

decline in progressivity in Section 4.4. Figure 4 reports the fit of equation (24) in the two sub-periods against the CPS pre- and post-government income data. The goodness of fit measured by the R^2 statistic is high, exceeding 0.80 in both samples. Both panels in Figure 4 reveal that our statistical model provides an excellent approximation of the actual tax/transfer system in the US economy in the income range between, roughly \$8,000 and \$400,000. See Heathcote, Storesletten, and Violante (2010b) for a detailed discussion.

We set $\sigma = 2$ which, jointly with our estimate for τ , implies a tax-modified Frisch elasticity of 0.30, in line with the microeconomic estimates for male workers. We set $\beta = 0.972$ to reproduce a risk-free rate of 3% per year from equation (17) when $v_\omega = 0.004$, as estimated by Heathcote, Storesletten, and Violante (2009a). A value of $\pi = (1 - 1/35)$ is consistent with our age range 25-60 in the survey data.

In Heathcote, Storesletten, and Violante (2010a), we report that the fraction of college graduates of age 25-29 increased from 22% in 1980-1984 to 28.7% in 2001-2005.³⁵ We use these numbers for the shares e^* and e^{**} in the model.

The average log wage differential between the two education groups –i.e., the log skill

³⁵The original source is Table A.2 of the Educational Attainment section on the US Census Bureau website, www.census.gov/population/www/socdemo.

premium– in the two steady-states is, respectively,

$$\begin{aligned} sp^* &= p_H^* - p_L^*, \\ sp^{**} &= p_H^{**} - p_L^{**}. \end{aligned}$$

From the CPS (see Figure 1), we compute that $sp^* = 0.31$ and $sp^{**} = 0.52$. These two equations, together with normalizations (13) and (14) yield four equations in the four parameters $(p_H^*, p_L^*, p_H^{**}, p_L^{**})$ with solution $p_H^* = 0.23$, $p_L^* = -0.08$, $p_H^{**} = 0.38$, and $p_L^{**} = -0.14$. The skill-biased technological change parameter rises from $\zeta^* = 0.36$ to $\zeta^{**} = 0.47$. Equation (23) evaluated at both steady states yields the equilibrium enrollment rates in the model. Values $\mu_\chi = 3.26$ and $v_\chi = 6.20$ reproduce observed graduation rates $e^* = 0.220$ and $e^{**} = 0.287$.

To estimate the change in the insurable and the uninsurable components of wage inequality, we use the equilibrium cross-sectional variances and the covariances of the joint distribution of wages, hours and consumption. First, to confront the survey data one must incorporate measurement error in earnings, hours and consumption in the allocations. If we let m_{it}^x be measurement error in the variable x_{it} , we can rewrite the equilibrium allocations after a first stage regression on time, age and education dummies as:

$$\begin{aligned} \log w_{it} &= \alpha_{it} + \varepsilon_{it} + m_{it}^y - m_{it}^h, \\ \log c_{it} &= -(1 - \tau)\varphi_i + (1 - \tau)\alpha_{it} + m_{it}^c, \\ \log h_{it} &= -\varphi_i + \frac{1 - \tau}{\sigma + \tau}\varepsilon_{it} + m_{it}^h, \end{aligned} \tag{25}$$

where the first line recognizes that wages are measured as annual earnings divided by annual hours and, as such, they inherit error from both variables. We take the standard view that the variance of measurement error is i.i.d across agents and over time. Let v_α be the variance of the uninsurable component in the cross-section, i.e. $v_\alpha = v_{\alpha 0} + (1 - \pi)v_\omega \sum_{j=0}^{\infty} j\pi^j$. Then, we have eight parameters to identify and estimate: $\{v_\alpha^*, v_\alpha^{**}, v_\varepsilon^*, v_\varepsilon^{**}, v_\varphi, v_{m^y}, v_{m^h}, v_{m^c}\}$. Consider the following within-group cross-sectional moments implied by the allocations in (25) corrected

for measurement error

$$\begin{aligned}
 \text{var}(\log w_{it}) &= v_{\alpha t} + v_{\varepsilon t} + v_{m^y} + v_{m^h} & (26) \\
 \text{var}(\log h_{it}) &= v_{\varphi} + \left(\frac{1-\tau}{\sigma+\tau}\right)^2 v_{\varepsilon t} + v_{m^h} \\
 \text{cov}(\log w_{it}, \log h_{it}) &= \left(\frac{1-\tau}{\sigma+\tau}\right) v_{\varepsilon t} - v_{m^h} \\
 \text{var}(\log c_{it}) &= (1-\tau)^2 v_{\varphi} + (1-\tau)^2 v_{\alpha t} + v_{m^c} \\
 \text{cov}(\log c_{it}, \log h_{it}) &= (1-\tau) v_{\varphi} \\
 \text{cov}(\log c_{it}, \log w_{it}) &= (1-\tau) v_{\alpha t}
 \end{aligned}$$

and note that we have a separate set of moments for each time period $t = \text{“*”}$ and $t = \text{“**”}$.

From the last three moment conditions, one can identify the variance of the uninsurable shocks $v_{\alpha t}$, preference heterogeneity v_{φ} and the size of measurement error in consumption v_{m^c} . Given an external estimate of the variance of log hours v_{m^h} , the first three moment conditions identify the measurement error in earnings v_{m^y} and the variance of the insurable component $v_{\varepsilon t}$. Table 5 below summarizes the fit of the model and the point estimates under two assumptions about the size of measurement error in hours. The first value ($v_{m^h} = 0.056$) is the one that allows the model to match the level of the cross-sectional variance of log hours and the second value ($v_{m^h} = 0.034$) is a more conservative estimate.

The main result is that approximately 0.05 points of the 0.08 point increase in the residual variance of log wages are due to the uninsurable component and 0.03 points to the insurable component. The changes in the empirical moments between 1980-1984 and 2001-2005, together with the theoretical cross-sectional moments in (26), provide the intuition for our findings. The increases in the variance of log consumption and in the covariance between log consumption and log wages call for a rise in uninsurable risk. The rise in the wage-hour covariance and the gap between the rise in wage inequality and the rise in consumption inequality calls, instead, for an increase in insurable risk. It is important to note that our estimates of the change in v_{α} and v_{ε} –the key inputs of the welfare calculation– are not sensitive to the identifying assumption on the size of measurement error in hours worked v_{m^h} .

Finally, our estimates of preference heterogeneity is in line with the findings in Heathcote, Storesletten, and Violante (2009a). In particular, note that with dispersion in the disutility of work, individuals with the highest effort cost of working, work less and produce less which allows the model to replicate the positive hours-consumption covariance observed in the data.³⁶

³⁶Measurement error in consumption is estimated to be roughly 40% of total cross-sectional consumption

Table 5: Model's Fit and Estimation Results

(A) Variance of Measurement Error in Hours is 50% of Variance of Hours ($v_{mh} = 0.056$)							
Data and Model's Fit							
	$var(\ln w)$	$var(\ln h)$	$cov(\ln w, \ln h)$	$var(\ln c)$	$cov(\ln c, \ln h)$	$cov(\ln c, \ln w)$	
1980-1984	0.318 [0.319]	0.121 [0.111]	-0.031 [-0.032]	0.191 [0.190]	0.034 [0.033]	0.078 [0.077]	
2001-2005	0.400 [0.399]	0.107 [0.114]	-0.025 [-0.024]	0.216 [0.213]	0.027 [0.033]	0.112 [0.112]	
Parameter Estimates							
	v_α	v_ε	v_φ	v_{m^y}	v_{m^c}		
1980-1984	0.112	0.079	0.048	0.072	0.113		
2001-2005	0.162	0.108	0.048	0.072	0.113		
(B) Variance of Measurement Error in Hours is 30% of Variance of Hours ($v_{mh} = 0.034$)							
Data and Model's Fit							
	$var(\ln w)$	$var(\ln h)$	$cov(\ln w, \ln h)$	$var(\ln c)$	$cov(\ln c, \ln h)$	$cov(\ln c, \ln w)$	
1980-1984	0.318 [0.319]	0.121 [0.102]	-0.031 [-0.029]	0.191 [0.189]	0.034 [0.046]	0.078 [0.078]	
2001-2005	0.400 [0.399]	0.107 [0.105]	-0.025 [-0.020]	0.216 [0.213]	0.027 [0.046]	0.112 [0.112]	
Parameter Estimates							
	v_α	v_ε	v_φ	v_{m^y}	v_{m^c}		
1980-1984	0.112	0.018	0.067	0.155	0.104		
2001-2005	0.163	0.047	0.067	0.155	0.104		

4.4 Welfare analysis

We now move to computing the welfare effects of changes in the wage structure. We do so by comparing social welfare across two steady states which differ in the triplet $(\zeta, v_\alpha, v_\varepsilon)$. Recall that because of the skill-biased demand shift and the endogenous education decisions, the two steady states also differ in the skill prices (p_L, p_H) , the enrollment level e and in the average utility cost χ paid by college graduates.

From the social welfare function (5) and the utility function specification in (11), we obtain the counterpart of equation (4) which defines ω implicitly:

$$\begin{aligned} & \sum_{j=0}^{\infty} \left\{ (1-\pi)\pi^j \int \left[\log((1+\omega)c(\varphi, s, \alpha)) - \frac{\exp(\bar{\varphi} + (1+\sigma)\varphi)}{1+\sigma} h(\varphi, \varepsilon)^{1+\sigma} \right] dF_{\varphi, s, \alpha, \varepsilon}^* \right\} - (1-\pi)\bar{\chi}^* e^* \\ & = \sum_{j=0}^{\infty} \left\{ (1-\pi)\pi^j \int \left[\log c(\varphi, s, \alpha) - \frac{\exp(\bar{\varphi} + (1+\sigma)\varphi)}{1+\sigma} h(\varphi, \varepsilon)^{1+\sigma} \right] dF_{\varphi, s, \alpha, \varepsilon}^{**} \right\} - (1-\pi)\bar{\chi}^{**} e^{**}, \end{aligned}$$

where $F_{\varphi, s, \alpha, \varepsilon}$ is the joint distribution of $(\varphi, s, \alpha, \varepsilon)$ in the population. Note the additional term dispersion, and measurement error in earnings between 15% and 30% of total earnings dispersion, depending on the assumption on measurement error in hours.

in the social welfare function which captures the utility costs of attending college: $\bar{\chi}$ is the average utility cost paid by college graduates and $(1 - \pi)e$ the fraction of newborn with college education at every date.

Substituting the allocations in (18) in this equation, we arrive at:

$$\begin{aligned} \log(1 + \omega) = & - (1 - \tau)^2 \frac{\Delta v_\alpha}{2} + \left[2 \left(\frac{1 - \tau}{\sigma + \tau} \right) - \sigma \left(\frac{1 - \tau}{\sigma + \tau} \right)^2 \right] \frac{\Delta v_\varepsilon}{2} \\ & + (1 - \tau) \Delta \mathbb{E}[p_s] + \Delta \log \frac{\mathbb{E}[\exp(p_s)]}{\mathbb{E}[\exp((1 - \tau)p_s)]} - (1 - \pi) \Delta(\bar{\chi} \cdot e). \end{aligned} \quad (27)$$

where Δ denotes the change across steady states. Before discussing this expression, it is convenient to make two minor approximations

$$\begin{aligned} \log(1 + \omega) & \approx \omega \\ \log \left(\frac{\mathbb{E}[\exp(p_s)]}{\mathbb{E}[\exp((1 - \tau)p_s)]} \right) & \approx \tau \mathbb{E}[p_s] \end{aligned}$$

which yield

$$\omega \approx -\frac{1}{2} \underbrace{(1 - \tau)^2 \Delta v_\alpha}_{\Delta \text{var}(\log c_{it}) \text{ w.g.}} + \left[\underbrace{\left(\frac{1 - \tau}{\sigma + \tau} \right) \Delta v_\varepsilon}_{\frac{\partial(\log Y/H)}{\partial v_\varepsilon}} - \underbrace{\frac{\sigma}{2} \left(\frac{1 - \tau}{\sigma + \tau} \right)^2 \Delta v_\varepsilon}_{\Delta \text{var}(\log h_{it})} \right] + \underbrace{\Delta \mathbb{E}[p_s] - (1 - \pi) \Delta(\bar{\chi} \cdot e)}_{\text{Gain from human capital investment}} \quad (28)$$

This expression offers a clean decomposition of the welfare effects of a change in the wage structure into three factors.³⁷ First, an increase in the variance of uninsurable wage risk (a positive value for Δv_α) is welfare reducing. Progressive taxation reduces this welfare loss through ex-post redistribution. Note that the term $(1 - \tau)^2 \Delta v_\alpha$ is the change in the variance of log consumption within educational groups.

Second, an increase in the variance of insurable shocks (a positive value for Δv_ε) is welfare-improving since the coefficient on the term in Δv_ε is always positive. This term is composed of two pieces. The first piece is the increase in average labor productivity associated with an increase in insurable risk. Insurable risk raises average productivity per hour worked because it induces high productivity agents to work relatively long hours. We discuss this effect in detail in Heathcote, Storesletten, and Violante (2008). The negative piece is the increase in the variance of log hours induced by more idiosyncratic insurable risk multiplied by (one half times) the curvature parameter σ which defines the aversion towards hours fluctuations.

³⁷We use expression (28) only to describe the forces at work. In all our calculations, we use the exact expression in (27).

The last component of expression (28) represents the effects of a skill biased demand shift which changes the marginal productivity of college and high-school labor (and hence their prices p_s), the quantity of skilled labor e , and the average utility cost of attending college $\bar{\chi}$. More precisely, this last factor can be expressed as follows

$$\begin{aligned} \Delta \mathbb{E}[p_s] - (1 - \pi) \Delta(\bar{\chi} \cdot e) &= e^* (p_H^{**} - p_H^*) + (1 - e^{**}) (p_L^{**} - p_L^*) \\ &+ (e^{**} - e^*) (p_H^{**} - p_L^*) - (1 - \pi) \int_{\bar{\chi}^*}^{\bar{\chi}^{**}} \chi dF_\chi. \end{aligned}$$

This indicates that if $p_H^{**} > p_H^*$ while $p_L^{**} < p_L^*$, which corresponds to the case in our calibration, then a widening skill premium will benefit low χ agents who choose to graduate from college under both wage structures (a fraction e^* of them), will hurt those agents who choose high school under both wage structures because their utility cost of attending college is too high (a fraction $1 - e^{**}$ of them), and may benefit or hurt intermediate χ agents who choose high school when the college premium is low, but switch to college when the college premium is high. Overall, this factor captures the output gains that accrue to society through investments in human capital that allow workers to take advantage of the higher returns to education. In particular, because of our normalizations, without change in college enrollment ($e^* = e^{**}$) this factor is zero.

4.4.1 Results

We now report the welfare calculations based on the model. These are the structural counterpart of the calculations in Section 3.2. When we substitute the parameter values estimated in Section 4.3 into the welfare expression (27), we obtain that $\omega = 0.014$, i.e., the observed rise in U.S. wage inequality between 1980 and 2005 led to a welfare gain of 1.4%. Looking at the three components, one by one, tells us that the uninsurable risk component contributes for a loss of -0.012 , the insurable risk component contributes for a gain of 0.006 and the human capital accumulation component accounts for a gain of 0.020.

Decomposition The welfare expression (28) shows that there are four channels of adjustment which mediate the transmission of the shift in the wage structure to welfare: labor supply, private risk-sharing, redistributive taxation and human capital investment. We now do a simple series of counterfactuals where we shut down these mechanisms one by one and examine how the welfare calculation is affected. In all these counterfactuals, we first keep the skill prices fixed, i.e. we do not take into account the general equilibrium effects that different individual

choices have on prices. Next, we endogenize skill prices and compute the general-equilibrium version of these same counterfactuals.

To analyze the role of labor supply we set $\sigma = \infty$. By doing so, we remove from welfare the gain originating from the improved labor productivity. From our baseline decomposition it is immediate to see that ω drops to 0.8%. Next, we exclude private risk sharing from the model by assuming that all the increase in wage inequality is uninsurable. Beyond losing the productivity gain, as for the case $\sigma = \infty$, in this counterfactual the losses arising from market incompleteness increase and, overall, we obtain a value for ω close to zero.

When we set $\tau = 0$, we omit any scope for social insurance –the only way in which permanent wage shocks can be smoothed in the model. Moreover, we exclude the possibility that increases in output associated to the shift in the wage structure lead to a gain in terms of reduced average taxation. At the same time, lower progressivity reduces labor supply distortions and increases labor productivity. The first two forces dominate the third, and we obtain a welfare loss of 1%.

Finally, we freeze college enrollment at the level of 1980-1984 and prevent agents from taking advantage of the higher returns to education. This counterfactual leads to a welfare loss over 2%. We conclude that participation in the progressive taxation scheme run by the government and, especially, investment in human capital are the main channels through which households in the US economy were able to respond to the shift in the wage structure.³⁸

Time varying progressivity Our baseline results assume a fixed τ across steady states and, implicitly, attribute all the changes in the cross-sectional moments of the consumption and hours distribution –the source of our estimates of Δv_α and Δv_ε – to changes in the wage structure. However, our estimate of the degree of progressivity τ uncovered a non-negligible change between the two time periods: progressivity of the whole tax and transfer system fell from 1980 to 2005, mostly due to the Tax Reform Act of 1996.³⁹

Another key advantage of the structural approach is that we can incorporate into the model both sources of change, wages and progressivity, and do not need to assume that the entire dynamics in consumption and hours inequality are due to wages. The results of this estimation, where τ is allowed to vary from $\tau^* = 0.35$ to $\tau^{**} = 0.26$, are reported in Table 6.

The main finding is that the estimated increase in the variance of uninsurable risk is smaller,

³⁸In the next draft of the paper, we will perform the general equilibrium experiments, where we keep χ constant, but let skill prices (p_L, p_H) change.

³⁹Our estimates of τ refer to the progressivity of the entire tax and transfer system. Other sources report a sizeable decline in marginal tax rates of federal and state taxes in the same period. See, for example, <http://www.nber.org/~taxsim/marginal-tax-rates>.

Table 6: Model's Fit and Estimation Results: Time Varying Progressivity

(A) Variance of Measurement Error in Hours is 50% of Variance of Hours ($v_{m^h} = 0.056$)						
Data and Model's Fit						
	$var(\ln w)$	$var(\ln h)$	$cov(\ln w, \ln h)$	$var(\ln c)$	$cov(\ln c, \ln h)$	$cov(\ln c, \ln w)$
1980-1984	0.318 [0.321]	0.121 [0.109]	-0.031 [-0.037]	0.191 [0.184]	0.034 [0.031]	0.078 [0.078]
2001-2005	0.400 [0.397]	0.107 [0.115]	-0.025 [-0.020]	0.216 [0.219]	0.027 [0.035]	0.112 [0.111]
Parameter Estimates						
	v_α	v_ε	v_φ	v_{m^y}	v_{m^c}	
1980-1984	0.121	0.067	0.048	0.077	0.113	
2001-2005	0.154	0.113	0.048	0.077	0.113	
(B) Variance of Measurement Error in Hours is 30% of Variance of Hours ($v_{m^h} = 0.034$)						
Data and Model's Fit						
	$var(\ln w)$	$var(\ln h)$	$cov(\ln w, \ln h)$	$var(\ln c)$	$cov(\ln c, \ln h)$	$cov(\ln c, \ln w)$
1980-1984	0.318 [0.320]	0.121 [0.100]	-0.031 [-0.033]	0.191 [0.183]	0.034 [0.043]	0.078 [0.079]
2001-2005	0.400 [0.397]	0.107 [0.106]	-0.025 [-0.017]	0.216 [0.219]	0.027 [0.048]	0.112 [0.110]
Parameter Estimates						
	v_α	v_ε	v_φ	v_{m^y}	v_{m^c}	
1980-1984	0.122	0.003	0.066	0.161	0.104	
2001-2005	0.150	0.052	0.066	0.161	0.104	

and the increase in the variance of insurable risk larger compared to the case with τ constant (see Table 5). This is not surprising, because some of the increase in the variance of log consumption is accounted for by the fact that the tax system has become less redistributive. With these new parameter values, we compute even larger welfare gains from rising wage inequality, ie. $\omega = 0.022$. While the education component is virtually unchanged, now the uninsurable risk component contributes a loss of -0.007 only and the insurable risk component contributes a gain of almost 1 percentage point.

A longer time horizon Even though the rise in wage inequality started in the early 1970s, our calculations compared the early 1980s to the early 2000 because of the lack of comparable consumption data going back that far. However, a close look at our structural moments (26) reveals that one can identify the two key parameters needed in the welfare calculation ($\Delta v_\alpha, \Delta v_\varepsilon$) only using data on individual wages and hours. The change in the variance of log

wages, the variance of log hours and the wage-hour covariance

$$\begin{aligned}
 \Delta var(\log w_{it}) &= \Delta v_{\alpha t} + \Delta v_{\varepsilon t} \\
 \Delta var(\log h_{it}) &= \left(\frac{1-\tau}{\sigma+\tau}\right)^2 \Delta v_{\varepsilon t} \\
 \Delta cov(\log w_{it}, \log h_{it}) &= \left(\frac{1-\tau}{\sigma+\tau}\right) \Delta v_{\varepsilon}
 \end{aligned}
 \tag{29}$$

over-identify the pair $(\Delta v_{\alpha}, \Delta v_{\varepsilon})$, given external estimates of τ and σ . The advantage of being able to study a longer time horizon must be traded off with the inability of using consumption data to identify and estimate the shift in insurable and uninsurable risk. However, given the uncertainty surrounding the true increase in consumption inequality, restricting attention to wage and hours data alone can be considered a strength.

We now repeat our estimation and welfare calculation using the period 1970-1974 as initial steady state. From 1970-1974 to 2001-2005, the skill premium increased somewhat less than from 1980-1984 to 2001-2005, +0.18 instead of +0.21, but college graduation rates of 25-29 increased more, 10.5% instead of +7%. Residual variance of log wages increased more, 0.12 instead of 0.08, residual variance of log hours remained unchanged, and the wage-hour covariance rose by +0.022 instead of 0.006.

The estimation based on these new set of empirical moments yields the result that half of the rise in wage dispersion since 1970 was insurable and half uninsurable. The welfare gain associated with the shift in the wage structure from 1970-1974 to 2001-2005 is equal to 3.5% of lifetime consumption. The uninsurable risk component contributes for a loss of -0.014 , the insurable risk component contributes for a gain of 0.013 and the human capital accumulation component accounts for a gain of 0.036. The larger insurable component of wage risk, appearing in the strong rise of the wage-hour covariance, and the sharp rise in college enrollment are behind this substantial welfare gain.

5 Concluding remarks

What are the welfare consequences of the recent shift in the wage structure in the United States? In this paper, we have summarized the literature on the topic, a body of work that begun almost two decades ago by extending the analysis of trends in inequality to variables more directly correlated with well being than wages, such as consumption and hours of work.

The early welfare calculations were based on changes in the the empirical distribution of consumption and hours, analyzed through the lenses of a social welfare function. We have

repeated this type of calculation here and obtained, similarly to the original studies, welfare losses of the order of 2 percent of lifetime consumption when comparing the early 1980s to the early 2000s. These welfare losses capture the extent to which wage inequality transmits to consumption inequality and they are only marginally affected by the inclusion of hours worked.

We have pointed out that, by demeaning the empirical cross-sectional distribution, these welfare calculations ignore that the same sources behind the shift in the wage structure – the rise in the skill premium and the rise in residual wage volatility– can lead to welfare improvement through their effects on average output and productivity. The picture is therefore less pessimistic than what, at first sight, it appears to be. Individuals have, correctly, interpreted the rise in the skill premium as a rise in the return to human capital and have acquired more education. At the same time, part of the rise in wage volatility was easily insurable –through private risk-sharing or simply by reallocating time within the household– and modified labor supply decisions led to an increase in labor productivity. Overall, we have emphasized that the new wage structure offered novel opportunities that households have largely exploited by investing in human capital and reallocating time efficiently.

Quantifying the welfare gains associated to these channels requires a structural model. We laid out a rich model, based on Heathcote, Storesletten, and Violante (2009a), which has the virtue of delivering a closed-form, hence very transparent, solution for the welfare change (expressed as consumption equivalent variation) of the shift in the wage structure. Under a plausible calibration of this model, our calculations yield welfare gains over 1 percent of lifetime consumption.

The sharp rise in US economic inequality has featured prominently in the public policy debate. Our results suggest that policies that offset the rise in earnings inequality would not necessarily be welfare improving. For example, more progressive taxation dissuades individuals from acquiring additional education in response to a widening skill premium and dampens the productivity gains which emanate from human capital investment (see also Guvenen, Kuruscu, and Ozkan (2009) for a similar argument). Offering additional insurance against increasingly volatile persistent shocks would be a far more effective policy. However, in practice, welfare-augmenting and welfare-reducing drivers of inequality cannot be easily unbundled. The policy challenge is to design institutions and tax-transfer schemes that deliver insurance against misfortune at birth and later in life, while preserving incentives for agents to make efficient investments in education and labor supply decisions.

There are still plenty of open questions to be investigated and partial answers to be refined.

Here, we only list a few. First, the recent re-examination of household consumption expenditure data has led some authors to the conclusion that the relatively small rise in consumption inequality emerging from the CEX may be underestimated by as much as a factor of two. This evidence would point towards larger welfare losses because of imperfect consumption insurance, but our conjecture is that it would not affect the result that the rising skill premium, paired with higher enrollment, is a source of welfare gains.

Second, throughout the paper, we have abstracted from changes in the distribution of hours worked for women, mainly because of the lack of comprehensive data on hours worked at home. In particular, we brushed aside an important trend in the wage distribution: the narrowing gender gap. In Heathcote, Storesletten, and Violante (2009a) we show that the rise in relative wages of women, together with the increase in their participation rates, is an additional source of welfare gains.

Third, the rise in wage inequality started well before 1980, the first year available in the CEX. Exploiting our model, we showed that one can perform a welfare calculation even in absence of consumption data for the period before 1980, by solely relying on labor supply data. Along the lines of Blundell, Pistaferri, and Preston (2008) and Guvenen and Smith (2010), one could improve on our effort by imputing consumption expenditures into the PSID to construct a dataset with a joint distribution of consumption, hours and wages and repeat both the empirical and the structural approach over a longer time period.

Finally, the rise in inequality is not a phenomenon confined to the United States. As documented by Krueger, Perri, Pistaferri, and Violante (2010), in the last 30 years the income distribution widened in a number of other countries (e.g., Canada, UK, Sweden, and Germany), but not every country experience was like that of the US in terms of dynamics of inequality in consumption and labor supply, and hence welfare. A cross-country comparative analysis would shed light on how diverse labor market institutions, family structures, and tax/transfer system affect the welfare implications of a shift in the wage distribution.⁴⁰

⁴⁰The special issue of the Review of Economic Dynamics published in January 2010 (volume 13, issue 1) contains articles detailing the specific experience of twelve different countries and collects the underlying survey data.

References

- ACEMOGLU, D. (2002): “Technical Change, Inequality, and the Labor Market,” *Journal of Economic Literature*, 40(1), 7–72.
- ACEMOGLU, D., AND D. H. AUTOR (2010): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, forthcoming. Elsevier.
- AGUIAR, M., AND M. BILS (2010): “Has Consumption Inequality Mirrored Income Inequality?,” *mimeo*.
- AGUIAR, M., AND E. HURST (2009): “Deconstructing Lifecycle Expenditure,” *mimeo*.
- AIYAGARI, S. R. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *Quarterly Journal of Economics*, 109(3), 659–684.
- ATKINSON, A. B. (1970): “On the Measurement of Inequality,” *Journal of Economic Theory*, 2, 244–263.
- ATTANASIO, O., E. BATTISTIN, AND H. ICHIMURA (2007): “What Really Happened to Consumption Inequality in the US?,” in *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, ed. by E. R. Berndt, and C. Hulten. University of Chicago Press.
- ATTANASIO, O., E. BATTISTIN, AND M. PADULA (2010): “The Evolution of Consumption and Wage Inequality: Evidence and Implications for Consumption Smoothing,” *mimeo*.
- ATTANASIO, O., AND S. J. DAVIS (1996): “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104(6), 1227–1262.
- AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): “The Polarization of the U.S. Labor Market,” *The American Economic Review*, 96(2), 189–194.
- BENABOU, R. (2002): “Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?,” *Econometrica*, 70(2), 481–517.
- BLUNDELL, R. (2010): “From Income to Consumption: Partial Insurance and the Transmission of Inequality,” *Econometric Society Presidential Lecture*.
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98:5, 1887–1921.
- BLUNDELL, R., AND I. PRESTON (1998): “Consumption Inequality and Income Uncertainty,” *The Quarterly Journal of Economics*, 113(2), 603–640.
- BRODA, C., E. LEIBTAG, AND D. E. WEINSTEIN (2009): “The Role of Prices in Measuring the Poors Living Standards,” *Journal of Economic Perspectives*, 23(2), 77–97.
- BRODA, C., AND J. ROMALIS (2009): “The Welfare Implications of Rising Price Dispersion,” *mimeo*.
- BRODA, C., AND D. E. WEINSTEIN (2006): “Globalization and the Gains from Variety,” *The Quarterly Journal of Economics*, 121(2), 541–585.

- COGLEY, T. (2002): “Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey,” *Journal of Monetary Economics*, 49(2), 309–334.
- CONSTANTINIDES, G. M., AND D. DUFFIE (1996): “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 104(2), 219–240.
- CUMMINS, J. G., AND G. L. VIOLANTE (2002): “Investment-Specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences,” *Review of Economic Dynamics*, 5(2), 243–284.
- CUTLER, D. M., AND L. F. KATZ (1991): “Macroeconomic Performance and the Disadvantaged,” *Brookings Papers on Economic Activity*, 1991(2), 1–74.
- (1992): “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980’s,” *The American Economic Review*, 82(2), 546–551.
- DAVIS, S. J. (2003): “Comment on Krueger and Perri: On the Welfare Consequences of the Increase in Inequality in the United States,” *NBER Macroeconomics Annuals*, 18, 132–137.
- DEATON, A. (1997): *The Analysis of Household Surveys*. Johns Hopkins University Press.
- DI NARDO, J., N. M. FORTIN, AND T. LEMIEUX (1996): “Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach,” *Econometrica*, 64(5), 1001–1044.
- FEENSTRA, R. C., AND G. H. HANSON (1996): “Globalization, Outsourcing, and Wage Inequality,” *The American Economic Review*, 86(2), 240–245.
- FISHER, J. D., AND D. S. JOHNSON (2006): “Consumption Mobility in the United States: Evidence from Two Panel Data Sets,” *BE Press: Topics in Economic Analysis & Policy*, 6(1).
- GARNER, T., G. JANINI, W. PASSERO, L. PASZKIEWICZ, AND M. VENDEMIA (2006): “The CE and the PCE: A Comparison,” *Monthly Labor Review*, pp. 20–46.
- GOTTSCHALK, P., AND R. A. MOFFITT (1994): “The Growth of Earnings Instability in the U.S. Labor Market,” *Brookings Papers on Economic Activity*, 25(2), 217–272.
- (2009): “The Rising Instability of U.S. Earnings,” *The Journal of Economic Perspectives*, 23(4), 3–24.
- GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (2000): “The Role of Investment-Specific Technological Change in the Business Cycle,” *European Economic Review*, 44(1), 91–115.
- GUVENEN, F. (2009): “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, 12(1), 58–79.
- GUVENEN, F., AND B. KURUSCU (2009): “A Quantitative Analysis of the Evolution of the US Wage Distribution, 1970–2000,” *NBER Macroeconomics Annual*, 24, 231–276.
- GUVENEN, F., B. KURUSCU, AND S. OZKAN (2009): “Taxation of Human Capital and Cross-Country Trends in Wage Inequality,” *mimeo*.
- GUVENEN, F., AND A. SMITH (2010): “Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach,” *mimeo*.

- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967-2006,” *Review of Economic Dynamics*, 13(1), 15–51.
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2008): “Insurance and Opportunities: A Welfare Analysis of Labor Market Risk,” *Journal of Monetary Economics*, 55(3), 501–525.
- (2009a): “Consumption Insurance and Labor Supply with Partial Insurance: An Analytical Framework,” *mimeo*.
- (2009b): “Quantitative Macroeconomics with Heterogeneous Households,” *Annual Review of Economics*, 1, 319:354.
- (2010a): “The Macroeconomic Implications of Rising Wage Inequality in the United States,” *Journal of Political Economy*.
- (2010b): “Redistributive Taxation in a Partial-Insurance Economy,” *mimeo*.
- HECKMAN, J., L. LOCHNER, AND C. TABER (1998): “Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents,” *Review of Economic Dynamics*, 1(1), 1–58.
- HOBIIJN, B., K. MAYER, C. STENNIS, AND G. TOPA (2009): “Household Inflation Experiences in the U.S.: A Comprehensive Approach,” *SF Fed Working Paper*, 19.
- HUGGETT, M. (1993): “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 17(5-6), 953–970.
- JONES, C. I., AND P. J. KLENOW (2010): “Beyond GDP? Welfare across Countries and Time,” *mimeo*.
- JORGENSON, D. W., AND K. J. STIROH (1999): “Information Technology and Growth,” *The American Economic Review, Papers and Proceedings*, 89(2), 109–115.
- JUHN, C., K. M. MURPHY, AND B. PIERCE (1993): “Wage Inequality and the Rise in Returns to Skill,” *The Journal of Political Economy*, 101(3), 410–442.
- KAPLAN, G., AND G. L. VIOLANTE (2010): “How Much Consumption Insurance Beyond Self-Insurance?,” *American Economic Journal: Macroeconomics*.
- KATZ, L. F., AND D. H. AUTOR (1999): “Changes in the Wage Structure and Earnings Inequality,” in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and D. Card, vol. 3 of *Handbooks of Labor Economics*, chap. 26, pp. 1463–1555. Elsevier.
- KATZ, L. F., AND K. M. MURPHY (1992): “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *Quarterly Journal of Economics*, 107(1), 35–78.
- KNOWLES, J. (2008): “Why Are Married Men Working So Much?,” *mimeo*.
- KRUEGER, D., AND F. PERRI (2003): “On the Welfare Consequences of the Increase in Inequality in the United States,” *NBER Macroeconomics Annual*, 18, 83–121.

- (2006): “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *The Review of Economic Studies*, 73(1), 163–193.
- KRUEGER, D., F. PERRI, L. PISTAFERRI, AND G. L. VIOLANTE (2010): “Cross Sectional Facts for Macroeconomists,” *Review of Economic Dynamics*, 13(1), 1–14.
- KRUGMAN, P. (2005): “The Debt-Peonage Society,” <http://www.nytimes.com/2005/03/08/opinion/08krugman.htm>
- (2007): *The Conscience of a Liberal*. W. W. Norton & Company.
- KRUSELL, P., T. MUKOYAMA, A. SAHIN, AND A. A. SMITH (2009): “Revisiting the Welfare Effects of Eliminating Business Cycles,” *Review of Economic Dynamics*, 12(3), 393 – 404.
- KRUSELL, P., L. E. OHANIAN, J.-V. RIOS-RULL, AND G. L. VIOLANTE (2000): “Capital-skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68(5), 1029–1053.
- KRUSELL, P., AND A. A. SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- (1999): “On the Welfare Effects of Eliminating Business Cycles,” *Review of Economic Dynamics*, 2, 245–272.
- LAZEAR (2006): “Lazear vs Krugman,” <http://gregmankiw.blogspot.com/2006/07/lazear-vs-krugman.html>.
- LEMIEUX, T. (2008): “The Changing Nature of Wage Inequality,” *Journal of Population Economics*, 21(1), 1432–1475.
- LEMIEUX, T., W. B. MACLEOD, AND D. PARENT (2009): “Performance Pay and Wage Inequality,” *Quarterly Journal of Economics*, 124(1), 1–49.
- LILLARD, L. A., AND Y. WEISS (1979): “Components of Variation in Panel Earnings Data: American Scientists 1960-1970,” *Econometrica*, 47(2), 437–454.
- LUCAS, R. E. (1987): *Models of Business Cycles*. Blackwell.
- (2003): “Macroeconomic Priorities,” *American Economic Review*, 37, 1–14.
- MEGHIR, C., AND L. PISTAFERRI (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72(1), 1–32.
- (2010): “Earnings, Consumption and Lifecycle Choices,” *Handbook of Labor Economics*.
- MORETTI, E. (2010): “Real Wage Inequality,” *mimeo*.
- RAJAN, R. (2010): “How Inequality Fueled the Crisis,” *Project Syndicate*.
- RIOS-RULL, J.-V. (1995): *Frontiers of Business Cycles Research* chap. Models with Heterogeneous Agents, pp. 98–125. Princeton University Press.
- SLESNICK, D. T. (1994): “Consumption, Needs and Inequality,” *International Economic Review*, 35(3), 677–703.

- (1998): “Empirical Approaches to the Measurement of Welfare,” *Journal of Economic Literature*, 36, 2108–2165.
- (2001): *Consumption and Social Welfare: Living Standards and their Distribution in the United States*. Cambridge University Press.
- STORESLETTEN, K. (2003): “Comment on Krueger and Perri: On the Welfare Consequences of the Increase in Inequality in the United States,” *NBER Macroeconomics Annuals*, 18, 132–137.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): “Consumption and Risk Sharing Over the Life Cycle,” *Journal of Monetary Economics*, 51(3), 609–633.
- VAN NIEUWERBURGH, S., AND P. O. WEILL (2010): “Why Has House Price Dispersion Gone Up?,” *Review of Economic Studies*, 77(4).
- VIOLANTE, G. L. (2002): “Technological Acceleration, Skill Transferability and the Rise of Residual Inequality,” *Quarterly Journal of Economics*, 117(1), 297–338.