Expenditure Cascades

By

Robert H. Frank and Adam Seth Levine

Evaluative judgments are known to depend heavily on context. For example, the same car that would have been experienced as having brisk acceleration in 1950 would seem sluggish to most drivers today. Similarly, a house of given size is more likely to be viewed as adequate the larger it is relative to other houses in the same local environment. And an effective interview suit is one that compares favorably with those worn by other applicants for the same job.

Although the link between context and evaluation is uncontroversial among behavioral scientists, the reigning economic models of consumer behavior completely ignore it. These models assume that each person's consumption spending is completely independent of the spending of others.

In contrast, James Duesenberry's relative income hypothesis explicitly acknowledges the link between context and evaluation.[1] In this paper we employ a variant of his model to explore the relationship between context and spending patterns. In this effort, we exploit data that allow us to quantify the effects of substantial increases in income inequality that have occurred in recent decades. According to the life-cycle and permanent income hypotheses, these increases should have no effect on individual spending decisions. In contrast, the relative income hypothesis predicts a substantial change in spending patterns in response to these changes. From statistical analysis of

[1] Duesenberry, 1949.

U.S. Census data for the 50 states and 100 most populous counties, we find evidence that rapid income growth concentrated among top earners in recent decades has stimulated a cascade of additional expenditure by those with lower earnings.

I. Expenditure Cascades

Milton Friedman's permanent income hypothesis continues to provide the foundation that underlies modern economic analysis of spending and savings.[2] According to this model, a family spends a constant proportion of its permanent income, rich or poor. The model thus predicts that savings rates should be independent of household income and should remain stable over time.

Both predictions are at odds with experience. It has long been shown, for example, that savings rates rise sharply with permanent income in cross-section data.[3] Savings rates have also shown substantial variation over time. According to U.S. Department of Commerce estimates shown in Figure 1, the aggregate personal savings rate has fallen from an average of roughly 10 percent in the mid-1970s to almost zero today.

---

[2] Friedman, 1957.

[3] See, for example, Mayer, 1972. Mayer rejects Friedman's original conjecture that this pattern is explained by the unresponsiveness of spending to transitory income changes, arguing that it cannot explain why people in high-income occupations save at higher rates than people in low-income occupations.

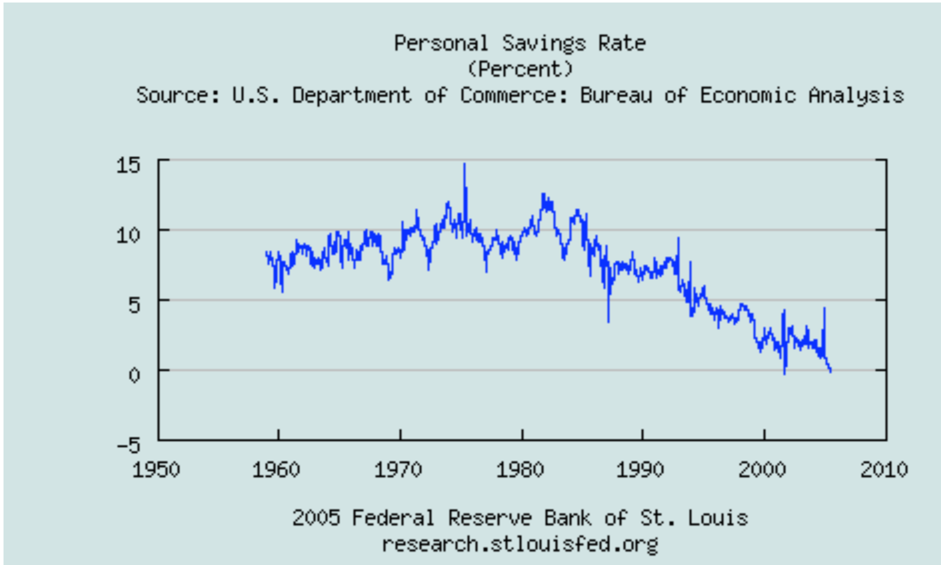

Figure 1. The Personal Savings Rate in the United States
Source: Federal Reserve Bank of St. Louis

The recent experience of middle-income families also casts doubt on Friedman's portrayal of the relationship between household income and spending. In 1980, the median size of a newly constructed house in the United States was approximately 1,600 square feet. By 2001, however, the corresponding figure had grown to over 2,100 square feet—more than twice the corresponding growth in median family earnings.[4] During the same period, the median household experienced substantial growth in consumer debt. One in five American households currently has zero or negative net worth.[5]

Why has consumption expenditure grown so much more rapidly than predicted by traditional economic models? We use the term *expenditure cascade* to describe a process whereby increased expenditure by some people leads others just below them on the income scale to spend more as well, in turn leading others just below the second group to

[4] Median house size growth: http://www.census.gov/prod/2003pubs/02statab/construct.pdf; http://www.census.gov/hhes/income/histinc/f03.html. Income growth rates: Center on Budget and Policy Priorities, 2003.

[5] Wolff, 2002.

spend more, and so on. Our expenditure cascade hypothesis is that a pervasive pattern of growing income inequality in the United States has led to the observed decline in savings rates.

II. An Illustrative Model

Consider an economy with N consumers arranged in descending order with respect to their permanent incomes. According to the permanent income hypothesis, individual i's current consumption, C_i, is proportional to his permanent income, Y_i:

$$C_i = kY_i, \qquad i = 1, \ldots, N, \tag{1}$$

where **k** is a parameter unrelated to permanent income level or rank. According to this model, each consumer's spending is independent of all income levels other than his own:

$$dC_i/dY_j = 0, \qquad \forall\ i \neq j\ . \tag{2}$$

Thus, according to the permanent income hypothesis, changes in the distribution of income should have no effect on individual spending levels. If someone's income does not change, his spending will remain the same, even if the income and spending levels of others change substantially.

In contrast to this baseline model, we consider the following variant of the relative income hypothesis:

$$C_i = k(1-a)\, Y_i \; + a\, C_{i+1}, \qquad i = 1, \ldots, N-1 \tag{3}$$

and

$$C_N = kY_N, \tag{4}$$

where C_i and Y_i again denote current consumption and permanent income levels of the i^{th} consumer, and where C_{i+1} denotes the current consumption level of the individual whose permanent income ranks just ahead of i's own. The parameter **k** is defined as before, and the parameter **a** (where $0 \leq \mathbf{a} \leq 1$) represents the extent to which each individual's spending is influenced by the spending of those with higher incomes. For **a** = 0, the spending of others has no influence at all, and the model collapses to the permanent income hypothesis. For **a** = 1, an individual's spending level is determined entirely by the spending level of the individual whose income just outranks his own. As indicated in equation 4, the highest ranking member of a group consumes according to the relationship assumed in the permanent income hypothesis. In a crude way, this model captures what are perhaps the two most robust findings from the behavioral literature on demonstration effects: 1) the comparisons that matter most are highly localized in time and space; and 2) people generally look to others above them on the income scale rather than to those below.[6]

A more realistic model would allow explicitly for the possibility that a consumer is also influenced by others more distant on the income scale. But even in our simple model, the influence of such others is captured indirectly through a chain of step-by-step comparisons. For example, if a given consumer were to spend an additional $100, the

[6] For a survey of the relevant literature, see Frank, 1985, chapter 2.

spending levels of the four individuals ranked just below him would go up by 100a, $100a^2$, $100a^3$, and $100a^4$ dollars, respectively.

For illustrative purposes, we consider a hypothetical 11-member reference group with **k** = 0.8 and **a** = 0.5. If the highest ranked member in this group consumes 80 percent of his income, lesser-ranked members will consume according to equation (3), which, for the assumed parameter values, simplifies to

$$C_i = 0.4\, Y_i + 0.5\, C_{i+1}\,, \qquad i = 1, \ldots, N\text{-}1. \tag{5}$$

For the initial income distribution shown in the left panel of Figure 2, the corresponding savings rates are shown in the right panel. They range from a high of 20 percent for the highest ranked member (the savings rate that we would see for everyone if the parameter **a** were equal to zero, as under the permanent income hypothesis), to a low of 12 percent for the lowest-ranked member. The average savings rate for the group is 15.6 percent, or 4.4 percentage points lower than it would have been in the absence of income inequality.

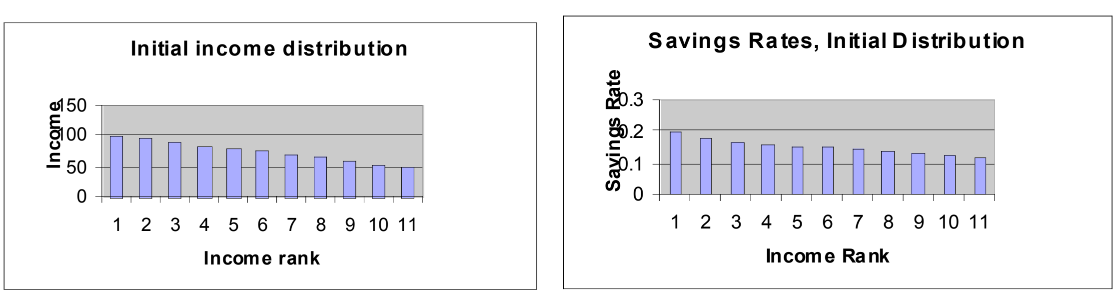

Figure 2. Income Rank and Savings Rates, Initial Income Distribution

We now alter the initial distribution by increasing the incomes of only the two highest-ranked members. In the new distribution, the highest-ranked member earns not 100, but 150; and the second-ranked member earns not 95 but 120. The incomes of the remaining members are the same as in the original distribution. The resulting is an expenditure cascade that lowers the savings rates of all the remaining members. The median earner, with an income of 75 in both distributions, saves at a rate of almost 15 percent under the original distribution, but only 12.3 percent under the new distribution. The savings rate for the group as a whole is now only 11.6 percent, a full 4 percentage points lower than it was under the original distribution.

Some economists object that concerns about relative consumption can affect savings rates in the manner described only if consumers are myopic. After all, if a consumer is induced to spend more today because of higher current spending by others, she will have even lower relative consumption in the future. Perhaps so. Yet it may still be rational to be responsive to community consumption standards.

Consider, for example, the fact that in most communities, the median family on the earnings scale now pays much more for housing, in real terms, than its counterpart in 1980. This family would find it easier to live within its means if it simply spent less on housing than others in the same income bracket. But because the quality of public schools in the United States is closely linked to local property taxes, which in turn depend on local real estate prices, this family would then end up having to send its children to below-average schools.[7] In the same vein, a job seeker could live more comfortably for the time being by refusing to match the increased expenditures of others on interview

---

[7] In the light of evidence that any given student's achievement level rises with the average socio-economic status of his or her classmates, property values and school quality will be positively linked even in jurisdictions in which school budgets are largely independent of local property values.

suits. Yet doing so would entail a reduced likelihood of landing the best job for which he was qualified. It is thus clear that being influenced by community consumption standards need not imply myopia. On the contrary, it may be a perfectly rational response on the part of consumers in pursuit of widely recognized goals.

On the other hand, there is considerable evidence that myopia is a salient feature of human psychology.[8] The pain of enduring lower relative living standards today can be experienced directly. In contrast, the pain of enduring lower relative standards in the future can only be imagined. So even though expenditure cascades can exist in the absence of myopia, they are undoubtedly strengthened by it.

In any event, if individual spending is influenced by the spending of others in the manner assumed in our simple model, an increase in income inequality will give rise to a reduction in savings rates. In the next section we examine how the increase in inequality assumed in our illustration compares with the actual recent growth in inequality.

III. Changing Patterns of Income Growth

In the United States, income growth from 1945 until the end of 1970s was well-described by the famous picket fence chart shown in Figure 3. Incomes grew at about the same rate for all income classes during that period, a little under three percent per year.

[8] Pigou, 1929, and more recently, Ainslie, 1992; Laibson, 1998; and O'Donoghue and Rabin, 1999.

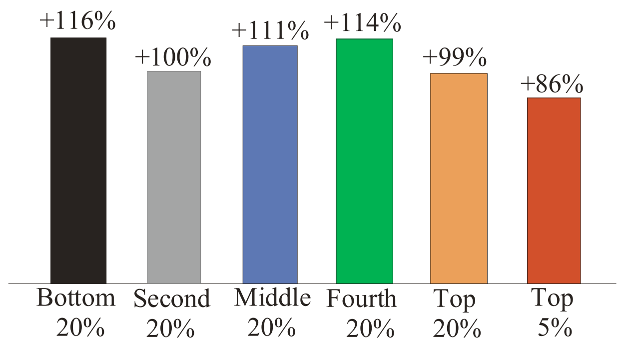

Figure 3. Changes in Before-Tax Household Incomes, 1949-1979.
Source: http://www.census.gov/hhes/income/histinc/f03.html

That pattern began to change at some point during the 1970s. During the 24-year period shown in Figure 4, the real pre-tax income of people at the bottom income distribution remained essentially unchanged, and gains throughout the middle of the income distribution were extremely small. For example, median family earnings were only 12.6 percent higher at the end of that period than at the beginning. Income gains for families in the top quintile were substantially larger, and were larger still for those in the top five percent. Yet even for these groups, income growth was not as great as during the earlier period. The later period was thus a period of both slower growth and much more uneven growth.

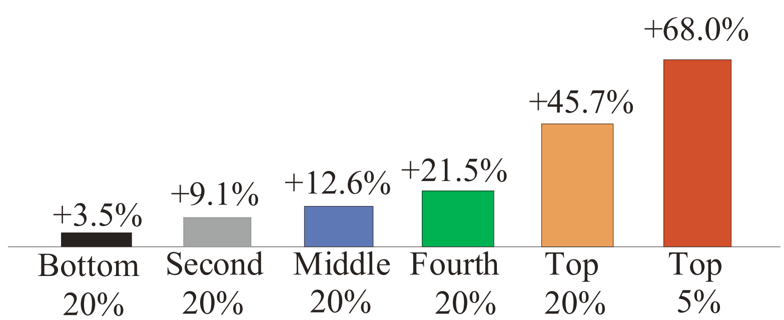

Figure 4. Changes in Before-tax Incomes, 1979-2003.
http://www.census.gov/hhes/income/histinc/h03ar.html

Income inequality has also increased in two important ways not portrayed in Figures 3 and 4. One is that changes in the income-tax structure during the Ronald Reagan presidency significantly shifted real after-tax purchasing power in favor of those atop the economic ladder, a change that was reinforced by additional tax cuts targeted toward high-income families during the first term of George W. Bush. A second change not reflected in Figures 1 and 2 is the magnitude of the earnings gains recorded by those at the very top of the income ladder.

Figure 5 portrays some of the results of these two additional effects. Note that the bottom 20 percent of earners (net of both tax and transfer payments) gained slightly more ground than in Figure 4, which showed pre-tax incomes (net of transfer payments). Note also that the gains accruing to the top one percent in Figure 5 are almost three times as large the corresponding pre-tax gains experienced by the top five percent. For people in the middle quintile, however, growth in after-tax incomes occurred at essentially the same modest pace as growth in pre-tax incomes.

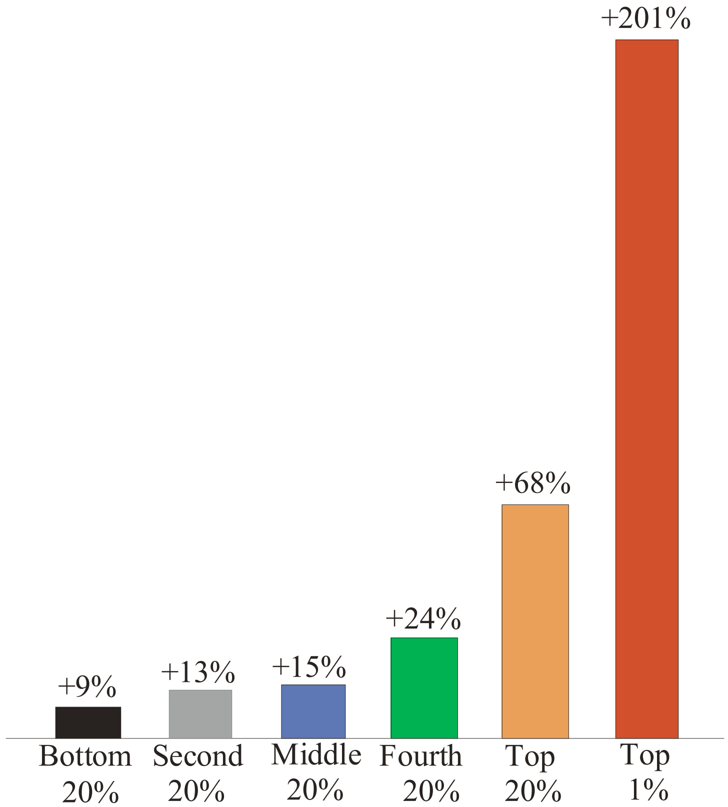

Figure 5. Change in After-Tax Household Income, 1979-2000
Source: Center on Budget and Policy Priorities, "The New, Definitive CBO Data on Income and Tax Trends," Sept. 23, 2003

For present purposes, an important feature of recent experience is that the aggregate pattern of income changes repeats itself in virtually every income subgroup. Thus, if we look at the top quintile of the earnings distribution, earnings growth has been relatively small near the bottom of that group and only slightly larger in the middle, but much larger among the top one percent. We see the same pattern again among the top one percent. In this group, the lion's share of income gains have accrued to the top tenth of one percent.

Only fragmentary data exist for people that high up in the income distribution. But a few snapshots are available. For more than 25 years, for example, *Business Week* has conducted an annual survey of the earnings of CEOs of the largest U.S. corporations.

In 1980, these executives earned 42 times as much as the average American worker, a ratio that is larger than the corresponding ratios in countries like Japan and Germany even today. But by 2001, the American CEOs were earning 531 times the average worker's salary. There is evidence that the gains have been even more pronounced for those who stand higher than CEOs on the income ladder.[9]

A similar pattern of inequality growth is observed when we look within occupations and educational groups. It shows up, for example, among college graduates, dentists, real estate agents and high school graduates.[10] *The upshot is that almost irrespective of the identities of the members of a person's personal reference group, income inequality within that group is likely to have grown sharply in recent decades.* Even for the wealthiest reference groups, for which average incomes have risen most sharply, most members are thus likely to have seen their incomes decline relative to those of their most prosperous associates.

IV. Three Specific Hypotheses

In its simplest form, the expenditure cascade hypothesis is that increasing income inequality within any reference group leads to a reduction in the average savings rate for that group. Our attempts to test this hypothesis are grounded on the observation that income growth patterns for most population subgroups in the United States in recent decades are roughly like the one shown for the population as a whole in Figure 5. Within most groups, people at the top have enjoyed robust earnings growth, while others have seen their incomes grow much more slowly. Our claim is that the new context created by

---

[9] See, for example, Krugman, 2002. As Wolff, 2002, has shown, the distribution of household net worth has also be come more right-skewed in recent decades.

[10] See Frank and Cook, 1995, chapter 5.

higher spending at the top of each group has caused others within the group to save a smaller proportion of their incomes.

An ideal test of this claim would examine how an individual's spending responds when other members of his or her personal reference group alter their spending. But because we cannot identify the specific persons who constitute any individual's personal reference group, we are forced to rely on crude proxies.

We begin by assuming that the amount of income inequality within a person's personal reference group varies directly with the amount of inequality in the geographic area in which that group is embedded. This assumption is more palatable for narrowly defined geographic areas than for broad ones. Thus, for example, the within-reference-group level of inequality for an individual is likely to correspond more closely to the degree of inequality in the city in which he lives than to the degree of inequality in his home country. In one version of our study, we employ samples of persons segregated by state of residence. In another, we employ samples from the 100 most densely populated counties. Our inequality measures for both sets of jurisdictions come from the 1990 and 2000 installments of the United States Census.

Do people who live in high-inequality jurisdictions in fact save at lower rates than those who live in low-inequality jurisdictions? Unfortunately, the Census does not record information that would enable us to construct reliable estimates of household savings rates by state or county. We are thus forced to examine alternative restatements of the hypothesis that are amenable to testing with available data.

A more general statement of the hypothesis is that families living in high-inequality areas will find it harder to live within their means than their counterparts in low-inequality areas. This observation suggests that the expenditure cascade hypothesis can be tested by examining the relationships between various measures of financial distress and measures of income inequality.

Families respond to financial distress in multiple ways, some of which leave clear footprints in data available from the Census or other sources. Beyond saving at lower rates, for example, they tend carry higher levels of consumer debt, which increases their likelihood of filing for bankruptcy. In addition, families who cannot afford to carry the mortgage payments for houses in conveniently located neighborhoods with good schools often respond by moving to cheaper, more remote neighborhoods, thus increasing their average commute times. And like other forms of distress, financial distress may increase the level of stress in personal relationships, thus increasing the likelihood of marriages ending in divorce. We have found that for both state and county data, growth in inequality between 1990 and 2000 is positively linked with growth in each of these three measures of financial distress. But because the narrower county level data are preferable from the perspective of our theory, we report only the results of our analyses of those data. Our decision to focus on the most populous counties was driven in part by Thorstein Veblen's observation that "…consumption claims a relatively larger portion of the income of the urban than of the rural population… [because] the serviceability of consumption as a means of repute is at its best…where the human contact of the individual is widest and the mobility of the population is greatest."[11]

[11] Veblen, 1899, p. 66.

V. Empirical Results

In this section, we present the results of empirical studies of the link between inequality and the likelihood of filing for bankruptcy, between inequality and the likelihood of filing for divorce, and between inequality and commute times.

We calculated two measures of income inequality in household incomes. The first was the ratio of the 90th percentile household income to 50th percentile household income (P9050).[12] The second is the Gini coefficient, a number between zero and one that indicates the level of inequality across the entire income distribution of an area.[13] For present purposes, the Gini coefficient is the preferred inequality measure, because it is Lorenz consistent[14] and accounts for the real income loss experienced by those in the lower reaches of the income distribution between 1990 and 2000, the specific time frame covered by our data. In the results we report below, we thus confine our attention to regressions in which our inequality measure was based on the Gini coefficient. (Results for regressions using the P9050 measures were qualitatively similar.)

To control for unobserved heterogeneity across states and counties, we ran all our regressions in first-difference form. In our bankruptcy regressions, for example, the value of the dependent variable for each area is the difference between that area's bankruptcy filings in 2000 and the corresponding number for 1990. Similarly, the area inequality variable we used was the difference between its Gini coefficient in 2000 and the

---

[12] P9050 ratios for states were calculated using 1-percent microdata samples provided by the Decennial U.S. Census. The ratios for counties were estimated using income brackets. For 1990, these brackets came from 1990 Census Summary File 3, tables P80 and P80A. For 2000, see 2000 Census Summary File 3, tables P52 and P53.

[13] We used a program provided by the U.S. Bureau of the Census to calculate Gini coefficients.

[14] An inequality measure is Lorenz consistent if and only if it is simultaneously consistent with the anonymity principle (permutations among people do not matter for inequality judgments), population principle (cloning the entire population and their incomes does not alter inequality), relative income principle (only relative, and not absolute, income matters), and Dalton principle (regressive transfers from poor to rich increase inequality).

corresponding measure in 1990.[15] Because both years were at approximately the same point in the business cycle, we do not expect this external influence to bias our results.[16]

Our first-difference regression models thus take the following general form:

$$\Delta dep_i = a + b\Delta ineq_i + c\Delta x_i + \Delta u_i \,, \tag{5}$$

where $\Delta dep_i = dep_{2000i}-dep_{1990i}$, the change in the dependent variable for area *i*, $\Delta ineq_i = ineq_{2000}-ineq_{1990}$, the change in the Gini coefficient for area *i*, Δx_i is a vector of the corresponding changes in other possible exogenous influences on the dependent variable (with **c** its vector of response coefficients), and Δu_i is an error term, assumed i.i.d.[17] The list of exogenous variables is recorded separately for each regression.

A. Bankruptcy

Individuals and married couples may file for non-business bankruptcy under Chapters 7, 11, or 13. To assess whether increases in inequality increase the likelihood of such filings, we use the total number of non-business bankruptcies under any of these three chapters as the basis for constructing our dependent variable.[18]

In addition to ΔGini, exogenous variables for our bankruptcy regressions include a mix of economic and socio-demographic characteristics measures employed by authors

---

[15] Some Decennial Census data, such as income, are for the year prior to the year of the census. In order to match income data with financial distress, we use non-business bankruptcies for 1989 and 1999. Welfare data used in the divorce rate regressions are from 1990 and 2000.
[16] Nor do we expect legislation to skew the results, because most bankruptcy law occurs at the federal level.
[17] To test for heteroskedasticity, we used a special form of White's test that regresses the squared residuals of the original regression on the predicted values and the squares of the predicted values. We reject the null hypothesis of homoskedasticity if the F-test on the two independent variables is significant. Instead of reporting the results of this test in every regression, homoskedasticity is assumed unless otherwise stated.
[18] All bankruptcy data come from the American Bankruptcy Institute website <http://www.abiworld.org/stats/stats.html>.

in the bankruptcy literature, all translated into first-difference form.[19] Economic factors include the change in the twentieth percentile household's nominal income (ΔNomP20),[20] the change in the proportion of total households in which both husband and wife work (ΔTwoWorker), and the change in the unemployment rate (ΔUnemploy). Socio-demographic characteristics include the change in average household size (ΔHHsize), the change in the proportion of total population black (ΔBlack), the change in the proportion of total population Asian and Pacific Islander (ΔAsian), the change in the proportion of total population ages 18-29 (ΔAge1829), and the change in the proportion of total population ages 15 and older divorced (ΔDivorce). In addition, since the number of people filing for bankruptcy in a county is population-sensitive, we include the change in the total county population aged 18 and over as an independent variable (ΔAdultPopulation). Finally, we include the change in population per square mile (ΔDensity). Only the last of these variables, ΔDensity, does not appear in standard bankruptcy studies. We added it to control for the possibility that it might be correlated with social forces that influence the likelihood of filing for bankruptcy.

At the outset, we had no prior views about what functional form would best capture the relationship between income inequality and financial distress. Simple linear regressions of the change in non-business bankruptcies on the change in income inequality revealed a positive, significant relationship in both our state and most populous county samples. But the goodness of fit was generally better in regressions involving the logarithms of the changes in bankruptcy and inequality measures. Also, this specification

[19] See, for example, Lawless, 2002; Summers and Carroll, 1987; and Hermann, 1966.
[20] Lacking price index data at the county level, we were forced to nominal income. But since the 1990s was a period of relatively low inflation, the change in nominal income for a county ought to be a good approximation for the corresponding change in real income.

was robust across our state and county samples and facilitated easily-interpretable results in terms of elasticities. In Table 1, we report the results for the ΔlnGini measures for the 100 most populous counties.

The coefficient for ΔlnGini suggests that, as hypothesized, changes in income inequality are positively and significantly associated with changes in the number of non-business bankruptcy filings in our sample of the 100 most populous counties. A one percent increase in the Gini coefficient is associated with an 8.73 percent rise in the number of non-business bankruptcies. This is a remarkably strong effect. For our sample of the 100 most populous counties, the Gini coefficients increased by an average of 4.41 percent between 1990 and 2000. Our estimate thus implies that increased inequality in these counties was associated with an almost 40 percent increase in bankruptcy filings between 1990 and 2000. This estimate seems reasonable given that, on average, non-business bankruptcies increased 148 percent in our sample.

Dependent Variable: Change in the natural logarithm of the number of non-business bankruptcies
Sample: 100 Most Populous Counties

Independent Variable	Coefficient	Standard Error	T-statistic	P
Constant	0.724	0.183	3.95	0.0002
ΔlnGini	8.732	2.338	3.73	0.0003
ΔNomp20	-0.00008686	0.00002439	-3.56	0.0006
ΔDensity	0.00003352	0.00002577	1.30	0.1968
ΔlnAdultPopulation	1.431	0.434	3.30	0.0014
ΔBlack	-0.883	1.672	-0.53	0.5987
ΔAsian	-2.595	2.745	-0.95	0.3472
ΔTwoWorker	11.339	4.167	2.72	0.0078
ΔUnemploy	2.771	3.585	0.77	0.4416
ΔAge1829	-8.311	3.170	-2.62	0.0103
ΔDivorce	11.172	7.577	1.47	0.1440
ΔHHSize	-1.490	0.752	-1.98	0.0505

$R^2 = 0.5173$, Adj $R^2 = 0.4570$

Table 1. The Relationship between Inequality and the Likelihood of Bankruptcy

Note also in Table 1 that changes in the absolute income of the 20th percentile household are negatively and significantly associated with changes in bankruptcy filings. This finding is consistent with the traditional view that households with more money should be better able to meet their financial obligations. But the effect is small, and does not rule out the notion that a household's desired consumption may increase hand in hand with income. Although the ΔDensity variable is not statistically significant at conventional levels, this may reflect the existence of threshold effects, since density is extremely high in most of the 100 most populous counties.

B. Divorce Rates

The dependent variable in our divorce regressions is the change in the proportion of the total area population aged 15 and over that is divorced. In these regressions, too, we include the standard economic and socio-demographic factors discussed by other authors in the relevant literature.[21] The main economic factor is the change in the log of the maximum state welfare benefit for a family of three, which captures the impact of the 1996 welfare reform that gave states greater latitude in distributing welfare benefits (ΔlnWelfare). The socio-demographic factors include the change in the proportion of total population aged 25 and over with at least a bachelor's degree (ΔEdu), the change in the proportion of women aged 16 and over in the labor force (ΔWomenLF), the change in the proportion of total households receiving retirement income (ΔRetInc), and the change in the average household size (ΔHHSize).

Table 2 reports our results for the ΔlnGini specification for the 100 most populous counties.

Dependent Variable: Change in the natural logarithm of the proportion of total population ages 15 and over divorced
Sample: 100 Most Populous U.S. Counties

Independent Variable	Coefficient	Standard Error	T-statistic	P
Constant	0.080	0.018	4.50	<0.0001
ΔlnGini	1.207	0.277	4.35	<0.0001
ΔlnWelfare	0.049	0.066	0.73	0.4671
ΔEdu	-0.700	0.289	-2.42	0.0173
ΔWomenLF	1.283	0.378	3.40	0.0010
ΔRetInc	1.322	0.694	1.90	0.0600
ΔHHSize	-0.502	0.104	-4.84	<0.0001

$R^2 = 0.5423$, Adj $R^2 = 0.5128$

Table 2. The Relationship between Inequality and Divorce Rates

[21] See, for example, Americans for Divorce Reform, 2003; and Nakonezny et al., 1995

Note in Table 2 that a one percent rise in the Gini coefficient is associated with a 1.21 percent increase in the proportion of divorced persons in highly populated counties. Given that the average change in the Gini coefficient between 1990 and 2000 was 4.41 percent for counties in our sample, the estimate implies that increased inequality was associated with a 5.34 percent increase in the number of divorces during this period.

C. Travel Time to Work

In these regressions, our dependent variable is the change in the proportion of all workers aged 16 and over whose daily commute is one hour or more. Here again we include a variety of economic and demographic characteristics that are known to affect our dependent variable.[22] We include changes in the median household income (ΔNomP50) . Because of studies finding a positive relationship between race and commute time, particularly for African-Americans,[23] we control for racial characteristics by including the change in the proportion of total population white (ΔWhite) and the change in the proportion of total population black (ΔBlack). We also include the change in the density of the population (ΔDensity), this time to control for changes in congestion on the roads and in the public transit systems. Finally, we include the change in the proportion of total population receiving retirement income (ΔRetInc), to control for the portion of the population that is probably not commuting.

Again the results for the state and county regressions were broadly similar. Unlike the earlier regressions, however, we found that ΔGini provided a somewhat tighter fit than ΔlnGini in these regressions and was more robust across our state and county

[22] See, for example, Khattak et al., 1999.
[23] Ibid.

samples. Table 3 reports our results for that specification for the 100 most populous counties.

Dependent Variable: Change in the proportion of total workers ages 16 and over with one hour or longer daily commute
Sample: 100 Most Populous Counties

Independent Variable	Coefficient	Standard Error	T-statistic	P
Constant	-2.261e-5	0.00576	-0.00	0.9969
ΔGini	0.403	0.160	2.51	0.0137
ΔnomP50	8.920e-7	3.727e-7	2.39	0.0187
Δwhite	-0.041	0.049	-0.84	0.4018
Δblack	0.156	0.066	2.35	0.0207
Δdensity	-1.917e-7	1.010e-6	-0.19	0.8492
Δretinc	-0.340	0.172	-1.98	0.0505

$R^2 = 0.2344$, Adj $R^2 = 0.1850$

Table 3. The Relationship between Inequality and Commuting Time

The estimated coefficient for ΔGini suggests that, as hypothesized, increases in income inequality are positively associated with changes in financial distress, as manifested in this instance by decisions to buy cheaper, but less conveniently located, housing. For counties in our sample, the Gini coefficient went up by an average of 0.018 between 1990 and 2000. Our estimate thus implies that increased inequality is on average associated with an increase of 0.0073 in the proportion of adults with commutes longer than one hour. For a county that began with the average value of that proportion in 2000 (0.09), increased inequality is thus associated with a rise of almost 8 percent in the number of adults with long commutes. For Fairfax County, Virginia, in which the proportion of adults with long commutes in 2000 was 0.097, and which had the largest growth in inequality during the decade (ΔGini= 0.038), our estimate suggests that approximately 16 percent more adults in the county had long commutes in 2000 than if inequality had not grown.

VI. Our Findings in Context

For our three specific measures of financial distress, our findings are consistent with the expenditure cascade hypothesis and at odds with the permanent income hypothesis. Economists seldom change their views about the efficacy of conventional models on the basis of isolated regression findings, nor should they. It is important to recognize, however, that our findings are part of a broader fabric of theoretical and empirical research that conveys a consistent message.

On the theoretical side, our best current understanding of the conditions that molded human nervous systems lends no support to models in which individuals care only about absolute resource holdings. No serious scientist disputes the Darwinian view that animal drives were selected for their capacity to motivate behaviors that contribute to reproductive success. In the Darwinian framework, reproductive success is all about relative resource holdings.

For example, frequent famines were an important challenge in early human societies, but even in the most severe famines, there was always some food. Those with relatively high rank got fed, while others often starved. On the plausible assumption that individuals with the strongest concerns about relative resource holdings were most inclined to expend the effort necessary to achieve high rank, such individuals would have been more likely than others to survive food shortages.

Relative resource holdings were also important in implicit markets for marriage partners. In most early human societies, high-ranking males took multiple wives, leaving many low-ranking males with none.[24] So here, too, theory predicts that natural selection will favor individuals with the strongest concerns about relative resource holdings. The motivational structure expected on the basis of theoretical considerations is thus consistent with the expenditure cascade hypothesis but inconsistent with models in which only absolute consumption matters.

On the empirical side, our findings on the link between inequality and various measures of financial distress complement similar findings by other researchers. Using OECD data across countries and over time, for example, Bowles and Park found that total hours worked were positively associated with higher inequality, both as measured by the 90/50 ratio and the Gini coefficient.[25] Using specially constructed 2000 Census data for a sample of 200 school districts in the United States, Ostvik-White found that median house prices were substantially higher in school districts with higher levels of income inequality, as measured by the 95/50 ratio, even after controlling for median income.[26]

The expenditure cascade hypothesis is also consistent with detailed patterns in cross-section data that are not predicted by the permanent income or life-cycle hypotheses. For example, as James Duesenberry observed in his 1949 book, a black family with a given absolute income would have had higher relative income in the segregated neighborhoods of the era than a white family with the same absolute income. And as Duesenberry predicted, the savings rates of black families with a given income level were higher than those of white families with the same income. The permanent

---

[24] Konner, 1982.
[25] Bowles and Park, 2002.
[26] Ostvik-White, 2003.

income hypothesis and the life cycle hypothesis, both of which disavow any role for context in consumption decisions, predict that families will save at the same rate irrespective of where they stand in their respective local distributions of income.

The expenditure cascade hypothesis is also consistent with observed patterns in international savings rates that are not predicted by traditional consumption theories. The aggregate savings rate, for example, was lower in the United States than in Europe in 1980, and the gap has grown larger during the ensuing years. One could invoke cultural differences to explain the initial gap, but the prevailing view is that cultures have grown more similar to each other with globalization, which leaves growth in the savings gap unexplained. The expenditure cascade hypothesis suggests, more parsimoniously, that the observed patterns in the savings data should mirror the corresponding patterns in the inequality data. It thus suggests that Americans saved less than the Europeans in 1980 because inequality was much higher in the United States than it was in Europe. And it suggests that the savings gap has grown wider because income inequality has been growing faster in the United States than in Europe in the years since then.[27]

Finally, the expenditure cascade hypothesis suggests a plausible answer to the question of why aggregate savings rates have fallen even though income gains have been largely concentrated in the hands of consumers with the highest incomes. As noted earlier, formal versions of the permanent income and life-cycle hypotheses predict no link between aggregate savings rates and differential rates of income growth across income classes. As a practical matter, however, modern specifications of these models have been forced to accommodate the fact that savings rates rise sharply with permanent incomes in cross-section data. If we take that fact as given, the observed pattern of

[27] Smeeding, 2001.

income growth in recent decades would seem to imply a secular upward trend in aggregate savings rates. After all, the lion's share of all recent income gains have accrued to prosperous families with the highest savings rates. And yet, as noted, aggregate savings rates have fallen sharply.

The expenditure cascade hypothesis suggests that the apparent contradiction may stem from the fact that the patterns of income change within wealthy groups have mimicked those we observe for the population as a whole. As noted earlier, available evidence suggests that no matter how we partition the population, income gains are highly concentrated among top earners within each group. Again, the expenditure cascade hypothesis stresses that local comparisons matter most. So even though more income is now flowing to members of prosperous groups, most members of such groups have been losing ground relative to their most prosperous peers. If it is relative income that drives the bequest motive and if local context is what really matters, the observed decline in aggregate savings rates is not anomalous.

VI. Concluding Remarks

Although persuasive theoretical and empirical evidence suggests that evaluations of consumption goods depend on context, prevailing economic models of consumption disavow any link between spending and context. This disavowal has become increasingly difficult to justify. Prevailing models predict that savings rates will not vary with permanent income; and that savings rates at all levels—individual, local, or national—should be insensitive to changes in the distribution of income. Prevailing models also predict that changes in income inequality should not influence either the number of hours people choose to work or the median price of housing where they live. Each of these predictions is contradicted by experience.

Economists have generally responded by incorporating ad hoc modifications into traditional theories—as, for example, by positing a bequest motive for wealthy consumers to accommodate the fact that savings rates rise sharply with permanent income in cross-section data. Such moves, however, generally raise more questions than they answer. Why, for example, should only the wealthy wish to leave bequests to their children?

Our claim is that existing fact patterns and theoretical constraints can be accommodated parsimoniously by simple variants of James Duesenberry's relative income hypothesis. We have argued that a simple model incorporating context-dependence predicts a clear link between income inequality and observed savings rates. Such a model predicts, for example, that the savings rate of any reference group will decline when income inequality within that group rises. This prediction is consistent with observed patterns in U.S. Census data for the 50 states and the 100 most populous counties between 1990 and 2000, a period during which income inequality was rising rapidly. It is also consistent with links found by other authors between inequality and hours worked. It is consistent as well with links found by other authors between inequality and median house prices. Finally, it is consistent with numerous observed patterns in cross-national savings data.

On the strength of available theoretical and empirical evidence, Mr. Duesenberry's relative income hypothesis clearly merits a closer look.

References

Ainslie, George. *Picoeconomics*, New York: Cambridge University Press, 1992.

Americans for Divorce Reform. *Correlations of Divorce Rates with Other Factors*, 25 April 2003 http://www.divorcereform.org/cor.html.

Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. "What Accounts for the Variation in Retirement Wealth Among U.S. Households?" *American Economic Review*, 91, Sept. 2001: 832-857.

Bowles, Samuel and Yongjin Park. "Emulation, Inequality, and Work Hours: Was Thorstein Veblen Right?" Santa Fe Institute mimeograph, 2002.

Bodkin, Ronald G. "Windfall Income and Consumption," *American Economic Review*, Sept. 1959, 49: 602-614.

Carroll, Christopher D. "Why Do the Rich Save So Much?" in Joel Slemrod, ed., *Does Atlas Shrug: The Economic Consequences of Taxing the Rich*, New York: Oxford University Press, 1998.

Diamond, Peter A. and Jerry A. Hausman. "Individual Retirement and Savings Behavior," *Journal of Public Economics,* 23, 1-2 (February-March 1984): 81-114.

Duesenberry, James, *Income, Saving, and the Theory of Consumer Behavior,* Cambridge, Mass.: Harvard University Press, 1949.

Frank, Robert H. *Choosing the Right Pond.* New York: Oxford University Press, 1985.

Friedman, Milton. *A Theory of the Consumption Function,* Princeton, NJ: Princeton University Press, 1957.

Herrmann, Robert O. "Families in Bankruptcy, A Survey of Recent Studies." *Journal of Marriage and the Family, v28*, n3 (1966): 324-330.

Khattak, Asad J., Virginie Amerlynck, and Roberto G. Quercia. "Are Travel Times and Distances to Work Greater for Residents of Poor Urban Neighborhoods?,"July 1999, accessed on 25 April 2003 <http://www.unc.edu/~khattak/wktt.html>.

Konner, Melvin. *The Tangled Wing,* NY: Holt, Rinehart, and Winston, 1982.

Krugman, Paul. "For Richer: How the Permissive Capitalism of the Boom Destroyed American Equality." *New York Times Magazine*, October 20, 2002: 62-142.

Laibson, David. "Life Cycle Consumption and Hyperbolic Discounting Functions," *European Economic Review Papers and Proceedings, 42,* 1998: 861-871.

Lawless, Robert M. “The Relationship between Nonbusiness Bankruptcy Filings and Various Basic Measures of Consumer Debt,” August 12, 2002, accessed on April 25, 2003, http://www.law.unlv.edu/faculty/rlawless/busbkr/body_filings.htm.

Liviatan, Nissan. "Tests of the Permanent-Income Hypothesis Based on a Reinterview Savings Survey," in *Measurement in Economics: Studies in Mathematical Economics and Econometrics in Memory of Yehuda Grunfeld*, edited by Carl F. Christ and others. Stanford, Calif.: Stanford University Press, 1963, pp. 29-59.

Mayer, Thomas. “The Propensity to Consume Permanent Income,” *American Economic Review*, 56, (1966): 1158-1177.

Mayer, Thomas. *Permanent Income, Wealth, and Consumption.* Berkeley: University of California Press, 1972.

Modigliani, Franco and R. Brumberg. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data," in *Post-Keynesian Economics,* edited by K. Kurihara, London: Allen and Unwin, 1955.

Nakonezny, Paul A, Robert D. Shull, and Joseph Lee Rodgers. “The Effect of No-Fault Divorce Law on the Divorce Rate Across the 50 States and its Relation to Income, Education, and Religiosity.” *Journal of Marriage and the Family. v57*, n2, 1995: 477-488.

O’Donoghue, Ted and Matthew Rabin. “Doing It Now or Later,” *American Economic Review, 89,* March, 1999: 103-124.

Ostvik-White, Bjornulf. “Income Inequality and Median House Prices in 200 School Districts,” Cornell University Institute of Public Policy Masters Thesis, Ithaca, NY, 2003.

Pigou, A. C. *The Economics of Welfare*, London: Macmillan, 1929.

Smeeding, Timothy M. “The Gap Between Rich and Poor: A Cross-national Perspective on Why Inequality Matters and What Policy Can Do To Alleviate It,” prepared for the National Institute of Social Security Research, Tokyo, Japan, March 21, 2001.

Summers, Lawrence H. and Chris Carroll, “Why Is U.S. National Saving So Low?” *Brookings Papers on Economic Activity* 1987:2 (1987).

Thaler, Richard and H. Shefrin, "An Economic Theory of Self-Control," *Journal of Political Economy, 89,* April, 1981: 392-405.

Veblen, Thorstein. *The Theory of the Leisure Class*, 1899, Alan Wolfe, ed., New York: The Modern Library, 2001.

Wolff, Edward N. *Top Heavy: A Study of Increasing Inequality of Wealth in America,* NY: The New Press, 2002.