

JONATHAN A. PARKER

Northwestern University

ANNETTE VISSING-JORGENSEN

Northwestern University

The Increase in Income Cyclicity of High-Income Households and Its Relation to the Rise in Top Income Shares

ABSTRACT We document a large increase in the cyclicity of the incomes of high-income households, coinciding with the rise in their share of aggregate income. In the United States, since top income shares began to rise rapidly in the early 1980s, incomes of those in the top 1 percent of the income distribution have averaged 14 times average income and been 2.4 times more cyclical. Before the early 1980s, incomes of the top 1 percent were slightly less cyclical than average. The increase in cyclicity at the top is to a large extent due to increases in the share and the cyclicity of their earned income. The high cyclicity among top incomes is found for households without stock options; following the same households over time; for post-tax, post-transfer income; and for consumption. We study cyclicity throughout the income distribution and reconcile our findings with earlier work. Furthermore, greater top income share is associated with greater top income cyclicity across recent decades, across subgroups of top income households, and, in changes, across countries. This suggests a common cause. We show theoretically that increases in the production scale of the most talented can raise both top incomes and their cyclicity.

Since the early 1970s, economic inequality in the United States—as measured by the distribution of wages and salaries, or of income more broadly, or of consumption expenditure—has been steadily increasing.¹ The consensus explanation for the general increase in inequality is that skill-

1. For wages and salaries this change was first documented by Bound and Johnson (1992) and Katz and Murphy (1992). The increase that began in the 1970s and 1980s continued

biased technological change has raised the earnings of individuals with more skills, as measured, for example, by education. However, accompanying this steady rise in inequality has been a much larger and more rapid increase in the income share of those at the very top of the income distribution. The share of (non-capital gains) income accruing to those in the top 1 percent of the income distribution increased from 8 percent in the early 1980s to 18 percent in 2008; the income share for those in the top 0.01 percent increased from around 0.7 percent to 3.3 percent over the same period (Piketty and Saez 2003, Saez 2010). Both the suddenness and the magnitude of these increases have shifted perceptions about the importance of technological change as the cause of increased income inequality generally and raised the possibility of an important role for other factors, such as “changes in labor market institutions, fiscal policy, or more generally social norms regarding pay inequality” (Piketty and Saez 2003, p. 3).

In this paper we bring together evidence from a variety of datasets to show that, as first argued in Parker and Vissing-Jorgensen (2009), another fundamental shift has occurred across the U.S. income distribution. During the past quarter century the incomes of high-income households have become much more sensitive to aggregate income fluctuations than previously. Before the early 1980s, the incomes of high-income households were more often than not less cyclical than the income of the average household. But since around 1982 the incomes of the top 1 percent have become more than twice as sensitive to aggregate income fluctuations as the income of the average household.

The fact that this increase in the cyclicalness of income of the top 1 percent coincides with the increase in their income share suggests that a common cause underlies both phenomena. We provide further evidence for a link between increased income inequality and increased income cyclicalness at the top by documenting, first, that across *income groups* within the top 1 percent, higher average income is associated with higher income cyclicalness in the 1982–2008 period; second, that across *decades* since the 1970s, cyclicalness of the top 1 percent increases decade by decade as that group’s income share increases; and third, across *countries*, increases in income

through the 1990s and into the 2000s in the top half of the wage distribution (Autor, Katz, and Kearney 2008). On increasing inequality in consumption, see Cutler and Katz (1991), Attanasio and Davis (1996), and Heathcote, Perri, and Violante (2010). Although the survey information on households suggests that the increase in the overall distribution of inequality in expenditure has been significantly less than that observed for income, this may partially be an issue of measurement of expenditure (see, for example, Aguiar and Bils 2010).

cyclicality of the top 1 percent are highly correlated with increases in their income share.

We argue that these facts are not inconsistent with the hypothesis that the increase in top income shares was caused by rapid technological progress in information and communications technologies (ICT) since the early 1980s. If improvements in ICT have increased the ability of the most talented workers to handle more work or to scale their ideas by working with more production inputs, then the ICT revolution could have caused the incomes of the highest paid both to rise and to become more sensitive to economic fluctuations. The intuition is that individuals who have less decreasing returns to scale will operate at a greater scale (that is, with more production inputs) and have lower ratios of gross revenue to production costs, and therefore have greater sensitivity of earnings to business cycles.

Expanding on these contributions, we begin in section I by focusing on the details of the change in income cyclicality of top income groups in the United States. We use the Statistics of Income (SOI) data of Thomas Piketty and Emmanuel Saez (Piketty and Saez 2003, Saez 2010), which are based on tax records, to show that the average income (before taxes and transfers and excluding capital gains) accruing to those in the very top of the income distribution has moved substantially more (in percentage terms) than the overall average in each boom and each recession since 1982, on average rising 5.0 percentage points more per year in each boom and falling 3.7 percentage points more per year in each recession. Before 1982, however, this was not the case.

This high cyclicality is not simply due to capital or entrepreneurial income. High-income tax units (one or more individuals filing a single return) tend to have a significant share of income from wages and salaries (including bonuses), and this type of income has roughly the same exposure to fluctuations as their nonwage income. Wage and salary income is also a major source of the change in cyclicality of top incomes. Before 1982 the wage and salary income of high-income tax units was roughly acyclical, but since 1982 it has been highly cyclical. Also, we show that the top 1 percent of earners come from a broad range of industries and occupations, and we argue that no one industry's or occupation's pay structure is driving our finding.

Further, we provide three pieces of evidence that although high-income households are more likely to have stock options, our main finding is not driven by the potentially endogenous timing of the exercise of stock options. First, in the period since 1997 for which we have data, only about 22 percent of households in the top 1 percent have stock options (that is,

were given stock options during the preceding year or owned stock options when surveyed), and income cyclicality of households in the top 1 percent is roughly similar if one leaves out households with stock options. Second, for a sample of top corporate executives for whom we have information about the value of options granted, we find that income calculated by including options only when granted, rather than when exercised, is highly cyclical. To be clear, this evidence in no way rules out a causal explanation that involves a general rise in pay for performance—indeed, options income is highly cyclical for those who have options, and bonus income may serve a similar purpose for those in the top 1 percent without options income. Our point is simply that the high cyclicality of the wage and salary (and overall) income of the top 1 percent is not spuriously generated by a correlation between the timing of options exercise and aggregate fluctuations. Third, as a further piece of evidence that the high cyclicality is neither due to endogenous timing of income without economic significance nor due to other measurement problems in income data, we show that the cyclicality of the *consumption* of households in the top of the consumption expenditure distribution—specifically, the top 5 percent by initial consumption—is also more than twice that of the average household.

Additional evidence confirming the high cyclicality of top incomes comes from verifying the out-of-sample forecasts made in Parker and Vissing-Jorgensen (2009) based on cyclicality estimates that excluded the recent recession. Income data for 2008 and consumption expenditure data through February 2009 show sharp declines for the top 1 percent during the recent recession, consistent with these predictions.

How does this new fact relate to the prior literature that concludes that low-income households bear the brunt of recessions and benefit the most from expansions? In section II, using data from the Current Population Survey (CPS), we show that the incomes of low-education households are more cyclical than those of high-education households and that the greater cyclicality of the top 1 percent does not appear in the CPS before 1982. Further, looking at the whole distribution using a dataset from the Congressional Budget Office that merges the CPS with the SOI tax data on high incomes, we find that the sensitivity of the wage and salary income of households in the bottom two quintiles to fluctuations in aggregate income is slightly higher than that of households in the third and fourth quintiles and than that of households from the 80th to the 99th percentiles.

However, in the public CPS data for the period since 1982, when one ranks by percentile in the income distribution, the top 1 percent have a higher cyclicality than even the lowest education group (those with less than a high

school diploma). The cyclicity of the top 1 percent is even higher when measured using the CPS top 1 percent income series constructed by Richard Burkhauser and coauthors (2008, 2009) from underlying CPS data not subject to the top coding applied to the public files. Thus, top incomes are highly cyclical, but it is harder to observe this high cyclicity in the publicly available CPS data alone because of top coding, and because cyclicity is high only for very high income households. We conclude that across the distribution of incomes, cyclicity is asymmetrically U-shaped: it is higher for the bottom quintiles than for the middle and the upper-middle class, but much higher for the top 1 percent, and especially for the very highest incomes.

Different cyclicalities of taxes and transfers at different points in the income distribution can lead to differences in cyclicity between pre-tax, pre-transfer cash income and disposable (post-tax, post-transfer) income. We show that taxes and especially transfers significantly reduce the cyclicity at the bottom of the income distribution while making less difference to the cyclicity of the very top. Thus, the cyclicity of top 1 percent incomes relative to the rest of the population is even greater for disposable income than it is for pre-tax, pre-transfer income.

Having established and explored our main finding for the United States, in section III we present evidence from Canada, which has a different tax system, slightly different culture, and better available information on top incomes from tax records. In the Canadian tax data, top income cyclicity is quite similar to that in the United States during the past quarter century. Further, in the Canadian data we are able to follow families across years (that is, we use panel data). Families in the top 1 percent of the income distribution in one year have income changes to the next year that are almost twice as cyclical as for the average. This higher cyclicity for the top 1 percent is similar in repeated cross-sectional data and in panel data, suggesting that the availability of only repeated cross-sectional data in the U.S. tax data is unlikely to substantially affect the estimated U.S. cyclicalities.

Section IV presents evidence of a strong link between increased income inequality and increased income cyclicity at the top by exploiting variation across groups, decades, and countries. We split the top 1 percent into three groups (percentiles 99–99.9, 99.9–99.99, and 99.99–100) and document for the period since 1982 that across these groups, the higher the average income, the higher the income cyclicity. Furthermore, calculating cyclicalities by decade since 1970, we show that for a given top group, as its income share increases, the cyclicity of its income increases. Finally, comparing the period 1970–82 with the period 1982–2007 using data for 10 countries, we find that those with larger increases in the income

share of the top 1 percent also have larger increases in the income cyclical-ity of the top 1 percent.

The link between increased inequality and increased cyclical-ity suggests a common cause of the two phenomena. In section V we argue that the increase in cyclical-ity is not inconsistent with an explanation of the increase in top income shares based on market-driven changes in incomes rather than, for example, changes in social norms. Specifically, we outline an explanation for both phenomena based on the rapid improvements in ICT in recent decades. Skill-biased technological progress that takes the form of lowering the degree of decreasing returns to scale for the highest-skill individuals naturally leads to increases in both the incomes and the income cyclical-ity of these individuals.

We emphasize that our results do not imply that the utility or happiness of high-income households is more cyclical than that of the average household. In fact, if risk aversion is lower at high expenditure levels, the utility of high-income households may be less cyclical than that of lower-income households, even with higher income cyclical-ity. Instead, our main finding establishes a new fact that is informative about changes in incomes and the labor market for high earners and of particular relevance for theories of the recent rise in income shares of high-income households.

I. The Changing Cyclical-ity of High Incomes

In this section we document the changing cyclical-ity of the income that accrues to top percentile groups in the income distribution, using the Statistics of Income data compiled by Piketty and Saez (2003) and extended by Saez (2010). In doing so, we study the timing of the change in cyclical-ity documented by Parker and Vissing-Jorgensen (2009). We show that the dramatic increase in the cyclical-ity of high incomes started in the early 1980s, and that this increase is significantly due to earned income and not just due to the (potentially endogenous) timing of executive stock option compensation.

I.A. The Main Facts

The main advantage of the Piketty-Saez data is that since they are based on administrative data from the Internal Revenue Service (IRS) on individual income tax returns, they provide extensive and accurate measurement of the very top of the income distribution. However, since some low-income households do not file tax returns (and even fewer did in the earlier years covered by the data), there is little detail on the low end of the income distribution. Piketty and Saez use aggregate personal income data

from the national accounts to calculate aggregate taxable income up to 1944; after 1944 they use the available tax return data plus an assumption about the incomes of nonfilers. Using these data, Piketty and Saez track the trend in the income share of the top 1 percent, 0.1 percent, and 0.01 percent of the income distribution, information simply not available in survey-based datasets on wages and incomes. The detail available on tax returns allows the measurement of pre-tax, pre-transfer cash income excluding realized capital gains. We exclude capital gains because our focus is on the timing of income, and the data contain only measures of realized capital gains, not capital gains as they accrue.

The data have some shortcomings, however. First, income excludes income paid as benefits (such as employer-paid health benefits and contributions to pensions) and excludes the employer share of payroll taxes (Social Security, Medicare, and unemployment taxes). Second, the unit of observation in these data is a tax unit, not an individual or a household. There has been a steady downward trend in the number of individuals per tax unit over time. This is a concern for measurement of trends if this ratio changes unevenly across income groups, but it poses less of a concern for our measurement of business cycle exposure. Third, the data are repeated cross sections and contain little information on demographics or other information that could allow one to track income changes for a constant population of households. Thus, the changes in income we report are based on income and income rank for groups of households that overlap but are not completely identical across years.²

Finally, incomes as reported to the IRS may be affected by tax reforms and by a variety of tax avoidance and tax evasion activities such as non-reporting of income, sheltering of income in 401(k)s, and changes in the reporting of income between closely held business profits and personal income. Tax reforms pose a particular concern since they cause changes in total reported taxable income that are potentially different across different filers. To the extent that such changes disproportionately affect high-income filers, this creates an artificially high correlation between changes in aggregate reported taxable income and changes in the reported taxable income of top income filers. To avoid this problem, we do not measure cyclical-ity from correlations with tax return-based aggregates, but instead use

2. We address each of these issues in our analysis of the Canadian data below and argue that focusing on a constant set of households does not lead to materially different results for the income cyclical-ity of the top 1 percent.

aggregates from the national income and product accounts (NIPA; see the online data appendix for details).³ Given this solution, tax reforms as well as the other data issues likely pose larger problems for measuring long-term trends than for measuring cyclicalities (see Reynolds 2007 and Piketty and Saez 2007).

We begin our analysis of these data by reporting the percent growth in income across each boom and recession since 1917, where “boom” and “recession” are defined, respectively, as periods during which NIPA real income per tax unit, before taxes and transfers and excluding capital gains, was increasing, and periods during which it was decreasing. Generally, these periods line up with recessions and expansions as identified by the Business Cycle Dating Committee of the National Bureau of Economic Research.

The dramatic increase in the exposure of high-income tax units to economic fluctuations began in the early 1980s. Table 1 shows the annualized percent change in average income per tax unit for all tax units, for the top 1 percent of the distribution, and for fractional percentiles within the top 1 percent. The final column reports the difference (in percentage points) between this annualized change for the top 1 percent and that for all tax units. Since 1982 the incomes of high-income households have risen more in booms and fallen less in recessions than the average income. According to the final column, since the end of the 1981–82 recession, the average income accruing to the top 1 percent of the income distribution has moved substantially more (in percentage terms) than the overall average in every boom and every recession, on average rising 5.0 percentage points more per year in each boom and falling 3.7 percentage points more per year in each recession.

Further, although one might think it natural for high incomes to be more cyclical, this was not so in the past. In the postwar period before 1982, the incomes of high-income households more often than not moved less (again

3. In our analysis this seems to be an important issue only for the 1986 tax reform (top group cyclicalities are higher in the 1980s if a tax-based measure of aggregate income is used). For the 1993 tax reform, Goolsbee (2000) provides evidence that executives timed the exercise of their options to take advantage of lower tax rates in 1992, thus seemingly raising aggregate income in 1992 at the expense of income in 1993. In the NIPA data, aggregate income growth was marginally negative from 1992 to 1993. To avoid artificially overstating our claim about extreme growth rates for top groups, we include 1993 as a boom year in table 1. Note, however, that Hall and Liebman (2000) argue that the high incomes in 1992 may not have been tax motivated, and they show that income shifting is not evident in response to two tax reforms of the 1980s.

Table 1. Changes in Real Income per Tax Unit by Income Group in Expansions and Recessions, 1917–2008Percent per year except where stated otherwise^a

<i>Period</i>	<i>All tax units</i>	<i>Top 1 percent</i>	<i>99.0th–99.9th percentile</i>	<i>99.9th–99.99th percentile</i>	<i>Top 0.01 percent</i>	<i>Change for top 1 percent minus change for all tax units (percentage points)</i>
<i>Expansions (periods with increasing aggregate personal income per tax unit)</i>						
2003–07	1.8	7.8	5.6	8.7	13.9	6.0
1991–2000	2.6	5.8	4.4	7.5	9.0	3.2
1982–89	2.2	7.9	6.0	10.7	14.3	5.7
1980–81	0.8	-2.7	-3.3	-1.3	-0.7	-3.5
1975–79	1.6	1.4	0.9	2.4	3.7	-0.2
1958–73	2.6	1.9	2.0	1.6	1.0	-0.8
1954–57	3.7	2.6	3.1	1.0	2.0	-1.1
1949–53	5.0	-0.1	0.9	-2.0	-4.1	-5.1
1947–48	1.4	4.7	3.3	8.4	7.5	3.3
1938–44	11.0	3.6	4.5	3.0	-0.7	-7.4
1933–37	8.3	9.3	9.7	9.1	7.8	1.0
1924–29	1.8	4.3	3.0	4.1	10.4	2.5
1921–23	12.1	10.3	9.9	9.7	14.1	-1.8
<i>Recessions (periods with decreasing aggregate personal income per tax unit)</i>						
2007–08	-2.6	-8.4	-6.7	-8.9	-12.7	-5.8
2000–03	-2.3	-5.8	-4.3	-7.7	-8.3	-3.5
1989–91	-1.7	-3.5	-2.2	-6.0	-5.6	-1.8
1981–82	-1.4	2.4	0.3	4.6	15.7	3.9
1979–80	-2.7	-0.9	-1.5	-0.5	3.6	1.8
1973–75	-4.5	-2.5	-3.2	-1.2	1.9	2.0
1957–58	-1.9	-4.7	-4.3	-5.7	-6.1	-2.8
1953–54	-1.1	2.2	2.5	0.2	3.7	3.2
1948–49	-2.3	-4.1	-4.1	-5.3	-1.2	-1.8
1944–47	-5.5	-0.4	0.6	-2.6	-2.4	5.1
1937–38	-8.0	-17.7	-14.4	-22.6	-24.0	-9.7
1929–33	-9.5	-12.8	-11.8	-12.5	-17.7	-3.4
1923–24	-1.2	7.5	6.0	8.8	13.3	8.7
1917–21	-7.6	-10.5	-6.1	-13.2	-22.0	-2.9

Sources: National Income and Product Accounts data, Piketty and Saez (2003), and Saez (2010). See the online appendix (www.brookings.edu/economics/bpea, under “Conferences and Papers”) for details.

a. Geometric annual averages calculated over the indicated period. Income is real pre-tax, pre-transfer income excluding capital gains and per tax unit; the same measure is used to define income groups.

in percentage terms) than the income of the average household. In the post-war period (1947 on) up to 1982, the incomes accruing to the top 1 percent co-moved less with the business cycle than did the income of the average household in 9 of the 12 booms and recessions. Relative to total income per tax unit, income accruing to the top 1 percent of tax units on average rose by 1.2 percentage points per year *less* in each boom and fell by 1.1 percentage points per year *less* in each recession. The difference between this period and the post-1982 period is economically large. Finally, in the pre-1947 period, for which the data are of poorer quality and, after 1941, influenced by wartime policies, the income accruing to the top 1 percent does not appear systematically more or less cyclical than that of the average household.

A striking feature of this change, to which we later return, is that it coincides almost exactly with the acceleration in the share of income accruing to the highest earners documented by Piketty and Saez (2003). In their data the income share of the top 1 percent reached its minimum at 7.7 percent in 1973, grew slightly to equal 8.0 in 1981, and then started rising rapidly to reach 17.7 percent in 2008. The coincident timing of the increase in top income shares and the increase in top income exposure to fluctuations suggests a common cause, as we discuss in sections IV and V.⁴

Notice from table 1 that, consistent with an out-of-sample forecast in Parker and Vissing-Jorgensen (2009), incomes of the top 1 percent fell substantially more than the average income in the recent recession—at least based on 2007–08 growth rates—with an 8.4 percent fall (again in real per-tax-unit terms) for the top 1 percent compared with a 2.6 percent fall for the average tax unit. The fall for the top 0.01 percent is even larger, at 12.7 percent. We emphasize that these numbers exclude capital gains and thus to a large extent are driven by wage and salary income, which fell by 3.3 percent from 2007 to 2008 for the average tax unit, by 6.0 percent for the top 1 percent, and by 17.5 percent for the top 0.01 percent. (We elaborate on the role of earned income for the top income groups below.)

Hereafter we will characterize the cyclical exposure of any income group i by a measure of its income cyclicity we call beta, which is the coefficient

4. Top income shares were also large in the prewar period, a period in which we do not find evidence for higher cyclicity of the incomes of the top 1 percent. Piketty and Saez (2003) argue that different factors drove the income shares of the top 1 percent during the period of declining inequality and during the period of increasing inequality; see our discussion in section IV. See also Kuznets (1953).

on the logarithmic change in income per member in the total population (Y) in a regression where the dependent variable is the log change in income per member of income group i (Y_i):

$$(1) \quad \Delta \ln Y_{i,t+1} = \alpha_i + \beta_i \Delta \ln Y_{t+1} + \varepsilon_{i,t+1}.$$

Beta is thus the elasticity of the income per member of group i with respect to average income, so that if average income growth is 1 percent, we expect the income of group i to grow by β_i percent.

The top panel of table 2 presents our main findings on the change in cyclicalities in terms of beta for the top 1 percent of the distribution and within subgroups of the top 1 percent across periods. The betas of the top 1 percent and the top 0.01 percent of tax units are 2.39 and 3.96, respectively, for the post-1982 period.⁵ These levels of cyclicity represent very large increases relative to prior periods: in the periods before 1982, the betas of all top income groups are less than 1, except for the top 0.01 percent for the period 1917–47.

The second panel of table 2 shows how much more income those in the top 1 percent and its subgroups received relative to the average household. These ratios are calculated from the group income shares (group income share/group size). Income per tax unit in the top groups was relatively high in 1917–47 (income per tax unit for the top 0.01 percent was 194 times the average income), was relatively lower in 1948–82 (65 times the average for the top 0.01 percent), and has been relatively high again since 1982 (207 times the average for the top 0.01 percent). In 2008 the top 1 percent included all tax units with incomes above \$342,000; the threshold for the top 0.01 percent was \$6.4 million. Average income for these two groups was \$906,000 and \$17.1 million, respectively, in that year.

The different betas and the larger share of income earned by top groups together translate into a disproportionate fraction of aggregate income changes falling on high-income households. To estimate the average fraction of aggregate income changes borne by a group, we regress (dollar change in real group income per tax unit) \times (group share of population)/(lagged aggregate real income per tax unit) on the growth rate of aggregate income per tax unit. Across all groups, the numerators sum to the total real dollar

5. It is worth clarifying that there is no mechanical tendency for a group to become more exposed to the cycle as its income share increases, but in fact the opposite. In the limit, as a group's income becomes a larger and larger share of all income, its exposure to the aggregate tends toward 1.

Table 2. Cyclicity of Real Income per Tax Unit, by Income Group, 1917–2008^a

<i>Period</i>	<i>All tax units</i>	<i>Top 1 percent</i>	<i>99.0th–99.9th percentile</i>	<i>99.9th–99.99th percentile</i>	<i>Top 0.01 percent</i>
<i>Income cyclicity (beta)^b</i>					
1982–2008	1.00	2.39 (0.57)	1.75 (0.38)	3.08 (0.80)	3.96 (1.11)
1947–82	1.00	0.72 (0.20)	0.81 (0.16)	0.63 (0.36)	0.02 (0.36)
1917–47	1.00	0.90 (0.17)	0.82 (0.14)	0.94 (0.20)	1.12 (0.31)
<i>Ratio of group average income to average for all tax units</i>					
1982–2008	1.0	13.6	9.2	36.2	206.6
1947–82	1.0	8.7	7.1	18.7	64.6
1917–47	1.0	15.4	10.7	42.6	194.4
<i>Fraction of aggregate income change borne by group^c</i>					
1982–2008	1.00	0.266 (0.059)	0.117 (0.024)	0.082 (0.019)	0.067 (0.018)
1947–82	1.00	0.056 (0.016)	0.046 (0.010)	0.010 (0.007)	–0.000 (0.002)
<i>Alternative measures of beta^b</i>					
Regressing group income growth on median income growth					
1982–2008	0.98 (0.14)	2.27 (0.77)	1.78 (0.51)	2.73 (1.10)	3.43 (1.49)
1967–82	0.93 (0.13)	0.52 (0.25)	0.64 (0.19)	0.32 (0.44)	–0.19 (0.58)
Regressing group income growth on unemployment rate					
1982–2008	–0.023 (0.004)	–0.058 (0.018)	–0.043 (0.012)	–0.076 (0.025)	–0.091 (0.035)
1948–82	–0.021 (0.002)	–0.015 (0.005)	–0.017 (0.004)	–0.013 (0.009)	–0.006 (0.009)

Sources: Authors' regressions using data in table 1, with additional data for median income growth and the unemployment rate. See the online appendix for details.

a. Standard errors are in parentheses.

b. Coefficient on the log growth rate of average income per tax unit for all tax units (top panel) or on the log growth rate in median household income or on the change in the unemployment rate (bottom panels), in a regression where the dependent variable is the log growth rate of average income per tax unit in the indicated group.

c. Coefficient on the growth rate of average aggregate income per tax unit in a regression where the dependent variable is $(\text{change in group average income per tax unit}) \times (\text{group share of population}) / (\text{lagged aggregate average income per tax unit})$.

change in income per tax unit, so the regression coefficients across a complete set of nonintersecting groups would sum to 1. Since 1982 the fractions of income changes borne by the top 1 percent and the top 0.01 percent are 26.6 percent and 6.7 percent—27 times and 670 times their shares in the population—respectively (third panel of table 2).

We emphasize that the increase in top income cyclicality is robust to using other measures of aggregate fluctuations. The fourth panel of table 2 measures cyclicality by beta with respect to changes in median household income (as calculated by the Census Bureau using the CPS) and with respect to changes in the aggregate unemployment rate. In both cases, measured cyclicality of the top 1 percent is lower than that for all tax units in the early period; from there it more than triples, reaching more than double that of the average tax unit in the recent period.

Furthermore, these changes in cyclical exposure represent actual increases in the cyclical volatility of high incomes. That is, the rise in the cyclical exposure of the top 1 percent is much greater than the decline in total income volatility that occurred in the Great Moderation. In the Piketty-Saez data, the standard deviation of the log change in the average income of the top 1 percent rose significantly, from 0.039 during 1947–82 to 0.085 during 1982–2008; the corresponding numbers for the top 0.01 percent are 0.059 and 0.155, respectively. In terms of cyclicality, the standard deviation of the cyclical component $\beta_i \Delta \ln Y_{t+1}$, rose also for all top income groups, as the standard deviation of $\Delta \ln Y_{t+1}$ fell only from 0.029 to 0.023, a much smaller (percentage) fall than the rise in the β_i s in table 2. Thus, for the top 1 percent, the standard deviation of the cyclical component $\beta_i \Delta \ln Y_{t+1}$ rose from 0.021 during 1947–82 to 0.055 during 1982–2008.

1.B. Wages and Salaries

To reiterate, in all of these results, the incomes of high-income groups are measured as cash income before government transfers and taxes, and the income changes are not contaminated by any endogenous timing of realizations of income reported as capital gains. That said, our results so far include income from all other taxable sources: wage and salary income (including bonuses and most stock options), entrepreneurial income, dividends, interest, and rental incomes. We now show that our main findings are driven to a large extent by the changing cyclicality of wage and salary income. We also document that they are not driven by potentially endogenous timing of stock options (more exercising of stock options in booms) or solely due to people with stock options.

Table 3 shows, for the postwar period up to 1982 and the period since, the average share of each group's income that is from each source as defined by the IRS (top panel) and the cyclicality of each type of income (bottom panel). This table documents three main points. First, in the period since 1982, wage and salary income accounts for only a slightly lower share of total income (60 percent) for the top 1 percent than for the average

Table 3. Composition of Income and Cyclicity of Income Growth, by Top Income Group and Income Source, 1947–82 and 1982–2008^a

<i>Income source</i>	1947–82				1982–2008				
	<i>All tax units</i>	<i>Top 1 percent</i>	<i>99.0th–99.9th percentile</i>	<i>99.9th–99.99th percentile</i>	<i>All tax units</i>	<i>Top 1 percent</i>	<i>99.0th–99.9th percentile</i>	<i>99.9th–99.99th percentile</i>	<i>Top 0.01 percent</i>
	<i>Average share of income from indicated source</i>								
Wages and salaries	71.9	45.2	49.4	38.8	67.3	60.3	67.4	53.5	40.0
Entrepreneurial	13.1	28.3	31.2	23.7	10.2	22.8	19.5	25.8	32.0
Dividends	3.5	17.5	11.1	27.1	5.0	6.8	5.1	8.4	12.3
Interest	8.2	5.3	5.0	5.8	15.3	7.7	6.2	8.9	12.2
Rent	3.4	3.8	3.4	4.6	2.1	2.4	1.9	3.5	3.5
	<i>Beta of group's income from indicated source</i>								
Total income	1.00	0.72	0.81	0.63	1.00	2.39	1.75	3.08	3.96
		(0.20)	(0.16)	(0.36)		(0.57)	(0.38)	(0.80)	(1.11)
Wages and salaries	1.12	0.36	0.44	0.20	0.87	2.38	1.32	3.61	6.20
	(0.05)	(0.14)	(0.13)	(0.27)	(0.06)	(0.58)	(0.31)	(1.08)	(1.93)
Entrepreneurial	1.39	1.87	2.08	1.82	1.33	2.07	2.29	0.76	1.53
	(0.25)	(0.68)	(0.59)	(0.99)	(0.33)	(1.31)	(1.13)	(2.91)	(1.78)
Dividends	1.16	0.85	0.96	0.83	1.24	2.65	3.37	2.33	1.64
	(0.29)	(0.38)	(0.39)	(0.68)	(0.57)	(1.26)	(0.97)	(1.62)	(1.93)
Interest	0.00	-0.10	-0.14	-0.04	1.54	4.52	4.41	5.24	3.84
	(0.19)	(0.48)	(0.44)	(0.66)	(0.39)	(1.28)	(1.18)	(1.22)	(1.71)
Rent	0.62	-0.44	-0.17	-0.73	-1.36	-0.26	-0.49	-0.37	-0.54
	(0.41)	(0.87)	(0.98)	(0.93)	(1.29)	(1.61)	(3.61)	(2.07)	(1.54)

Sources: See table 1. See the online appendix for details.

a. Income is total pre-tax, pre-transfer income excluding capital gains. Standard errors are in parentheses.

household (two-thirds). Wages and salaries are a smaller but still significant share of income for the top 0.01 percent (40 percent).

Second, and more important, since 1982 the wage and salary income of high-income groups is much more cyclical than that for all tax units. To maintain comparability across types of income and in the definition of an economic fluctuation, for all types we define cyclicity with respect to fluctuations in NIPA aggregate pre-tax, pre-transfer income excluding capital gains per tax unit. Since 1982 the wage and salary income of the top 1 percent has a cyclicity of 2.4, and that of the top 0.01 percent a cyclicity of 6.2, compared with a cyclicity of less than 1 for all tax units. The cyclicity of wage and salary income of the top 1 percent is about the same as that of their overall income (and thus as the average cyclicity of their other types of income), whereas the cyclicity of wage and salary income of the top 0.01 percent exceeds that of all their other types of income.

Third, the change in cyclicity of the top 1 percent since 1982 is to a large extent driven by the rise in the share of wages and salaries in their total income and the change in its cyclicity, with a smaller role for increased cyclicity of dividend and interest income. The top panel of table 3 shows that the share of wage and salary income in the incomes of the top 1 percent rose by 15 percentage points across periods. The bottom panel shows a dramatic increase in the cyclicity of the wages and salaries of the top 1 percent, from 0.4 in the 1947–82 period to 2.4 in the 1982–2008 period. Across periods there is also a substantial increase in the cyclicity of dividend and interest income for the top 1 percent, but these two sources are smaller shares of income. The cyclicity of entrepreneurial income for the top 1 percent is relatively stable, at around 2 for both 1947–82 and 1982–2008. For the top 0.01 percent, the change in cyclicity is more widespread across categories, but again the largest role is played by wage and salary income.

We next investigate the role of stock options in our findings. The rise of stock options coincides with the rise of income inequality, and the vast majority of stock options are nonqualified options, which are treated for tax purposes as wage and salary income when exercised.⁶ Because our

6. Qualified stock options are taxed as capital gains when exercised and the stocks received are sold, provided that they are held for a year and that the stocks purchased with them are held for another year. The gain resulting from the difference between the strike price and the market price, however, can count toward income for purposes of the alternative minimum tax. We do not deal here with the accounting treatment of stock options for financial reporting, which differs from the tax treatment for the individual; for example, it allows corporations to deduct more on their tax returns than they expense on their financial statements.

analysis so far is based on tax return data, it includes income from non-qualified options in wage and salary income. We are concerned that either endogenous timing of the exercise of stock options (if more are exercised in booms) or a correlation between stock market performance and aggregate income might make our measure of realized top incomes excessively procyclical even if actual economic earnings were not. Thus, we address two questions concerning options. First, is income from options sufficiently prevalent in the top 1 percent to be the main driver of high wage and salary cyclicality? Second, do we still find high cyclicality of top incomes if we include options in income when granted (at values determined by the Black-Scholes model) instead of when exercised (as in the tax data)?

To address the first question, we use the Survey of Consumer Finances (SCF) for 1998, 2001, 2004, and 2007, which contains information on wealth and income (for the preceding calendar year) for a stratified random sample of households that oversamples rich households. These years of the SCF also include the responses to two survey questions about stock options. The first asks whether the household received stock options during the past year, and the second asks whether the household has a valuable asset not otherwise recorded in the interview and then asks the household to state what it is, with stock options being one possible response. SCF data are not top coded, with the exception that a household is dropped if it has a net worth greater than the least wealthy person in the Forbes list of the wealthiest 400 people in the United States.⁷ On average across the four survey years, only 22 percent of households in the top 1 percent of the income distribution had stock options. Furthermore, the cyclicality of income growth (of non-capital gains income, based on aggregate income calculated from SCF data and using 3-year real log growth rates) is around 1.8 both for all households in the top 1 percent and for households in the top 1 percent without stock options. This indicates that income from stock options is not driving our main findings.

To answer the second question, we use data on executive compensation from ExecuComp, which are available for 1992 to 2009. Our sample definition is described in the online data appendix (at www.brookings.edu/economics/bpea, under “Conferences and Papers”). The average number of executives covered in our sample is 6,216 per year. The top panel of table 4 shows that in these data the average total executive compensation (in real 2008 dollars) was \$1.6 million in 1992 based on the value

7. This should not affect our results substantially, since the top 400 families correspond to only a small fraction of even the top 0.01 percent.

Table 4. Cyclicality of Income of Corporate Executives, 1992–2009

	1992	2009
	<i>Millions of 2008 dollars</i>	
<i>Average real total compensation</i>		
Based on value of options granted ^a	1.45	2.43
Based on value of options exercised ^b	1.63	2.39
	<i>Percent</i>	
<i>Average share of total compensation by component, based on value of options granted^c</i>		
Salary	32.6	20.2
Bonus	18.6	5.6
Stock grants	7.0	29.3
Option grants	29.6	19.4
Other ^d	12.2	25.6
	<i>Beta</i>	<i>Standard error</i>
<i>Cyclicality of component income growth^e</i>		
Based on value of options granted		
Total compensation	2.89	0.86
Salary	-0.12	0.13
Bonus	1.01	0.93
Stock grants	2.82	1.02
Option grants	5.36	1.70
Other ^d	0.97	1.57
Based on value of options exercised		
Total compensation	4.39	1.15
Option grants	10.86	2.24
Excluding options		
Total compensation	1.01	0.62
	<i>2007–08</i>	<i>2008–09</i>
	<i>Percent</i>	
<i>Growth rate of total real compensation</i>		
Based on value of options granted	-8.3	-5.3
Based on value of options exercised	-20.1	-18.2

Sources: Authors' calculations using ExecuComp data. See the online data appendix for details.

a. ExecuComp series tdc1.

b. ExecuComp series tdc2.

c. Average compensation from the indicated component divided by average total compensation. Numbers may not sum to 100 because of rounding.

d. For example, nonequity incentive plan compensation.

e. Estimation based on log growth and excluding the 2005–06 growth rate, which may be affected by changes in reporting requirements in 2006.

of options exercised. Using the group income cutoffs in the Piketty-Saez data, on average across 1992–2009, 81 percent of the ExecuComp executives were in the top 1 percent, and 7 percent were in the top 0.01 percent.⁸ The second panel of the table shows that the executives received a substantial fraction of their income in the form of options. The table also reports betas for each income component (calculated from annual averages of each type of income across executives). The beta of overall compensation is 2.9 based on the value of options granted, and 4.4 based on the value of options exercised. Given that only a small fraction of those in the top 1 percent have stock options income (according to the SCF data) and that the beta of executive compensation based on the value of options granted is about two-thirds that based on the value of options exercised (as calculated from the ExecuComp data), we conclude that endogenous timing of options is not likely to have substantially affected our beta estimates for wages and salaries using the Piketty-Saez data.

Interestingly, these findings do not imply that options are not critical for the income cyclicity of top earners who do receive stock options. For executives in the ExecuComp data, options income does drive the high cyclicity of their wage and salary income: their beta of compensation excluding options is around 1. That is, the cyclical component of their income is (granted) options. For these results to be consistent with our results from the SCF, however, it must therefore be that nonoptions wage and salary income is highly cyclical for top earners without options. Bonuses or other incentive pay may play a central role for these households, but our data sources (aside from ExecuComp) do not separately break out bonuses.

A final observation can be made from the ExecuComp data. Table 4 also shows the growth rates of real compensation for executives in this sample for 2007–08 and 2008–09. The negative growth rates for 2007–08 of –8.3 percent and –20.1 percent (depending on which options data are used) confirm the finding based on the data for all top 1 percent tax units in table 1 that top income groups were hit harder by the recent recession than the average household. For 2008–09 the executives in the ExecuComp data did much worse than the average tax unit (for which we estimate, using NIPA data, that wage and salary income fell by 5.3 percent

8. With an average of 137 million tax units across 1992–2009, the top 1 percent consists of, on average, 1,370,000 households, and the top 0.01 percent of, on average, 13,700 households. Households headed by executives represented in ExecuComp thus make up a tiny fraction of both the top 1 percent and the top 0.01 percent.

in real per-tax-unit terms) when we measure income including the value of options exercised, but similar to the average tax unit when we use the value of options granted.⁹

1.C. Who Is in the Top 1 Percent?

To further understand what drives the higher cyclicity of income of the top 1 percent, it is useful to document the characteristics of families in that group and how these have changed across periods. Since this is not feasible in the Piketty-Saez data, we use the March CPS public use microdata files. We study the characteristics of families and their heads for the entire population and for the top 1 percent using pre-tax, post-transfer family income excluding capital gains.¹⁰ Table 5 reports statistics averaged across the 5 years ending in 1982 and across the 5 years ending in 2008.

Heads of families in the top 1 percent tend to be slightly older than the average, are more likely to be married, and are less likely to be retired. They are more likely to be white, self-employed, and more educated. Perhaps surprisingly, the top 1 percent are widely dispersed across industries and occupations. This makes it less likely that a particular industry or occupation is driving most of the high cyclicity of incomes among this group. For example, it is unlikely that the increased cyclicity of the top 1 percent is due only to more of them being employed in finance today than earlier, or to incomes in financial occupations having become more cyclical (although finance may be more important for the top 0.01 percent), for two reasons. First, the share of the top 1 percent in finance (and related industries) is only 16 percent even at the end of our sample period, up by about 4.4 percentage points from the early 1980s. Therefore, whether one assumes that the beta of incomes in the finance industry is constant but that more of the top 1 percent are now employed in finance, or one allows the beta of finance to increase, the beta for finance in the post-1982 period would have to be at least 11 in order for finance to explain

9. The more meaningful comparison here is probably the one based on value of options exercised, since NIPA wages and salaries are based on that concept (see Moylan 2008). The treatment of options in the NIPA is unlikely to materially affect our results, since options income is only a tiny fraction of overall NIPA income. Furthermore, as shown in the bottom panel of table 2, our main results are very similar when we use unemployment or median income to measure aggregate fluctuations.

10. We use this definition of income to match with previous work using the CPS, since comparability is important for our analysis in section II.A.

Table 5. Demographic, Educational, and Occupational Characteristics of Heads of Families in the Top 1 Percent of the Income Distribution, 1978–82 and 2004–08^a

<i>Characteristic</i>	<i>Top 1 percent^b</i>		<i>All families</i>	
	<i>1978–82</i>	<i>2004–08</i>	<i>1978–82</i>	<i>2004–08</i>
	<i>Units as indicated</i>			
<i>Demographics</i>				
Average age	50.7	47.8	45.1	46.9
Percent with children under 18	37.9	50.6	51.5	46.4
Average no. of children under 18	0.7	1.0	1.0	0.9
Percent married	97.8	97.0	87.3	84.6
Percent retired	7.0	12.3	14.8	29.6
Percent white	95.9	88.3	87.6	81.7
Percent self-employed	39.4	27.8	11.6	9.1
	<i>Percent of all family heads</i>			
<i>Education</i>				
Less than high school	5.3	1.3	30.2	12.1
High school diploma	15.6	9.8	33.2	31.3
Some college	13.7	13.0	18.0	27.5
College degree	31.6	33.1	12.3	18.6
Post-college education	33.7	42.8	6.3	10.5
<i>Industry</i>				
Manufacturing and construction	22.0	11.8	28.3	14.9
Finance, insurance, and real estate	11.6	16.0	3.9	5.2
Professional services	24.7	41.8	11.4	23.0
Wholesale and retail trade	13.3	9.2	12.8	9.7
Other	28.4	21.3	43.5	47.1
	<i>1982–85</i>	<i>1998–2001</i>	<i>1982–85</i>	<i>1998–2001</i>
<i>Occupation^c</i>				
Executive, administrative, or managerial	34.7	35.5	10.8	12.3
Professional specialty	29.6	32.1	9.4	11.6
Sales	16.0	13.1	8.3	8.4
Other	19.7	19.3	71.6	67.7

Sources: Authors' calculations using Census public use data from the March CPS files from 1979 to 2009, referring to the previous year's income and labor force characteristics. See the online data appendix for details.

a. "Families" excludes people not living with someone related to them by blood or marriage. This definition includes about 90 percent of households in the top 1 percent of the income distribution and 76 percent of households in the general population (as determined from the 1995 Survey of Consumer Finances). Reported percentages and averages are averaged across years in the indicated period.

b. As defined by CPS family income (pre-tax, post-transfer income excluding capital gains).

c. We use a common occupation coding for income years 1982–2001.

the increase in beta of the entire top 1 percent from 0.7 to 2.4 (top panel of table 2).¹¹

Second, to the extent we can estimate betas of the top 1 percent at the industry or occupation level, we find no evidence that the beta for those in finance is dramatically larger than the betas of other top 1 percent households.¹² Jon Bakija, Adam Cole, and Bradley Heim (2010) provide data for the top 1 percent and the top 0.1 percent that are comparable to the data from Piketty and Saez but contain information about occupations (coded from taxpayer responses to the occupation question on Form 1040). We use their data for 1993, 1997, 1999, and 2001–05 to calculate log growth rates (annualized in the case of 4- or 2-year periods) and regress these on aggregate log growth rates (using the same aggregate variable we used earlier). Four occupations account for more than 5 percent of tax units in the top 1 percent and the top 0.1 percent. These are “executives, managers, and supervisors (non-finance),” “financial professions, including management,” “lawyers,” and “medical.” Using these data, we estimate a beta of 1.99 for the full top 1 percent, and betas for the four subgroups listed of 1.96, 2.34, 1.67, and 0.71. For the top 0.1 percent we estimate a beta of 2.82 for the full group and betas of 2.27, 3.08, 3.60, and 2.34 for the four main subgroups (all estimates listed have associated t statistics of 2 or more). With the exception of medical occupations within the top 1 percent, this suggests that betas are high across all the largest subgroups of the top 1 percent and the top 0.1 percent.

1.D. Consumption

We next turn to the question of whether the high cyclicality of income for high-income households leads to a high cyclicality of consumption

11. We calculate this as follows. Let a' and a denote the share of the top 1 percent employed in finance during 1982–2008 and 1947–82, respectively, and let β' and β denote the income cyclicality of the top 1 percent in these two periods. If the beta of those in the top 1 percent not employed in finance was constant at 1, then $a'\beta' + (1 - a') = 2.4$ and $a\beta + (1 - a) = 0.7$, and thus $a'\beta' - a\beta - (a' - a) = 1.7$. Suppose (based on table 5) that $a' = 0.16$, and assume (to give finance its best chance at being the explanation) that $a = 0.06$ (lower than the pre-1982 value from table 5). Then $0.16 \times \beta' - 0.06\beta = 1.8$. Consider two possible cases: If $\beta = \beta'$, then $\beta' = 18$. Alternatively, if $\beta = 1$, then $\beta' = 11.6$.

12. The CPS data (described in more detail in section II) are problematic for this purpose because values assigned to top-coded observations are not industry specific, implying that betas for top 1 percent households across industries could spuriously look similar. With that important reservation, we find that within each of the four industries listed in table 5, betas for families in the top 1 percent are much larger than for the average family, and the top 1 percent finance industry beta is roughly similar to that for the nonfinance industries taken as a group.

spending. Evidence on this question constitutes a further test of our main finding, as well as of the extent to which consumption is smoothed across these income changes, as would be the case for insurable changes in income or endogenous timing of income. Unfortunately, high-income groups are generally thought to be underrepresented in the Consumer Expenditure Survey (CE), and some CE consumption categories are top coded.¹³ Furthermore, in order to have a sufficient number of households, our analysis here focuses on the top 5 percent of CE households rather than the top 1 percent. Nonetheless, our analysis shows higher cyclicality for high-consumption households.

We use the CE data to construct measures of household-level spending from January 1982 to February 2009 for different groups ranked by their expenditure level in the quarter before the interview. Our consumption measure is nondurables plus some services; the main categories of excluded services are health care, education, and housing (except for the nondurable and service components of household operations). We deflate the reported consumption values using the Bureau of Labor Statistics (BLS) price index for nondurables. For each household we calculate log-consumption growth rates from one quarter to the next and average these across households in a given group (using survey weights). We then calculate annual log growth rates by summing four quarterly log growth rates. For each group we run a time-series regression of the four-quarter log growth rates in consumption per household on the log growth rate of one of four different series (in separate regressions): NIPA pre-tax, pre-transfer income; NIPA disposable (that is, post-tax, post-transfer) income; NIPA personal consumption expenditures (PCE) on nondurables and services; and CE average consumption for all households (using our consumption definition). For comparability across regressions in table 6 and for comparability with the earlier tables, the first three regressions all use the same price deflator, the CPI series from Piketty and Saez, whereas the regression with CE average consumption as the explanatory variable uses the BLS deflator (since both the left- and the right-hand-side variable are based on the same consumption measure).

Table 6 shows that the sensitivity of the consumption of households in the top 5 percent of the distribution (ranked by initial consumption) to

13. Because of the way the CE is structured, the respondent's burden rises with expenditure: more time is required to report more expenditure. Further, there is evidence that underreporting rises with expenditure. See, for example, Aguiar and Bils (2010).

Table 6. Cyclicity of Real Consumption among All Households and the Top 5 Percent, January 1982–February 2009^a

<i>Measure</i>	<i>All households</i>	<i>Top 5 percent</i>
Ratio of group average consumption to average consumption of all households	1.00	2.52
Beta from regression of consumption on on: ^b		
NIPA pre-tax, pre-transfer personal income	0.58 (0.14)	1.94 (0.50)
NIPA post-tax, post-transfer personal income	0.61 (0.23)	2.60 (0.61)
NIPA nondurables and services consumption	1.17 (0.27)	4.80 (0.97)
CE consumption for all households	1.00	2.38 (0.30)
Fraction of total CE consumption fluctuations borne by group ^c	1.00	0.32 (0.04)

Sources: Authors' calculations and regressions using data from the Consumer Expenditure Survey (CE).

a. Consumption includes expenditure on nondurable goods and some services. Groups are defined based on their consumption in the previous survey interview. Changes for all variables in all regressions are measured as 4-quarter log growth rates. Standard errors are in parentheses.

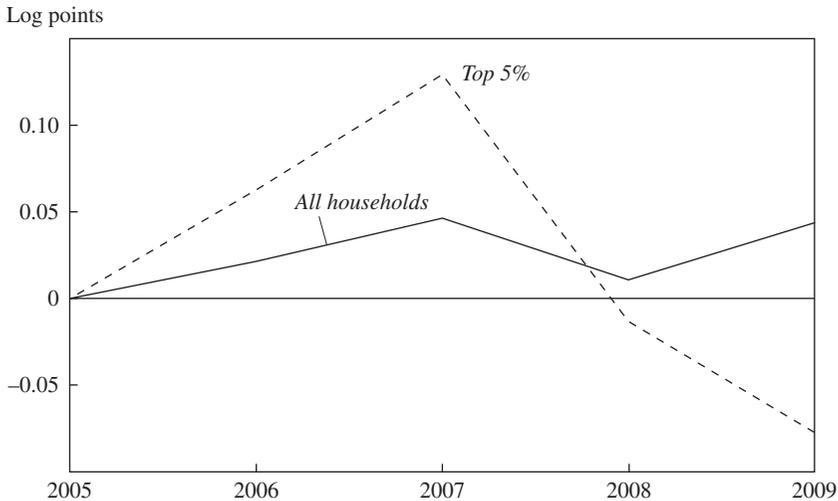
b. Each beta is the coefficient on the log change in the indicated aggregate in a regression where the dependent variable is the log change in consumption per household in the indicated group.

c. Coefficient on the growth rate of aggregate CE consumption per household in a regression where the dependent variable is (change in group average consumption per household) \times (group share of population)/(lagged aggregate average consumption per household).

aggregate income fluctuations is between 1.9 and 2.6, depending on the income measure used, whereas the sensitivity to aggregate consumption fluctuations is almost 5.¹⁴ This compares with a sensitivity of the consumption of the full set of CE households that is substantially less than 1 with respect to NIPA incomes.

The implications of this higher cyclicity are borne out in the expenditure response of high-consumption households to the recent deep recession. Figure 1 shows that CE consumption in the recent recession fell substantially more for high-expenditure households—more than 10 percent from 2007 to 2008—than the average for all households. This finding is

14. The sensitivities of top household consumption to NIPA consumption are a bit lower than similar statistics for a shorter sample reported in Parker and Vissing-Jorgensen (2009). The difference is due not to differences in the samples, but rather to the price index used: Parker and Vissing-Jorgensen (2009) used a PCE deflator to deflate NIPA consumption, whereas the results reported here use the CPI series from Piketty and Saez.

Figure 1. Cumulative Change in Log Real Expenditure per Household, 2005–09^a

Source: Authors' calculations using Consumer Expenditure Survey data for 2005 to February 2009.

a. Annual values are within-year averages (using survey weights) of quarterly consumption (multiplied by 4). The consumption measure is nondurables and some services. Households are sorted on the basis of current consumption.

consistent with the out-of-sample forecast in Parker and Vissing-Jorgensen (2009).¹⁵

These results provide additional evidence that the high cyclicality of top incomes is not due to the endogenous timing of compensation but instead affects the standard of living for top income households. We emphasize, however, that a given percent decline in expenditure presumably has greater welfare implications for a low-expenditure household than for a high-expenditure household. This point, along with the lack of foundation for interpersonal welfare comparisons, suggests that one should not conclude that high-income households suffer more from recessions than do low-income households.

15. For better readability, the figure focuses on annual data, calculated as within-year averages (using survey weights) of quarterly consumption values across households in a given group (multiplied by 4). Furthermore, because we are interested in showing levels of growth rates, not only their betas with respect to aggregates, we sort households according to current consumption rather than consumption in the previous quarter. The latter method is theoretically more meaningful but, with measurement error in consumption, leads to a mechanical negative bias in growth rates for top groups. As discussed in Parker and Vissing-Jorgensen (2009), this does not affect the estimation of betas when log growth rates are used, but it would bias this figure.

To conclude, we find a dramatic increase in the cyclicity of top incomes. This increase occurs for total (non-capital gains) income as well as for wage and salary income alone, and top groups' expenditures are also highly cyclical during the post-1982 period. Furthermore, the top 1 percent are active in a wide range of industries and occupations, making it less likely that a particular industry or occupation is driving most of the high cyclicity of top groups' incomes.

II. Cyclicalities across the Full Income Distribution and the Impact of Transfers, Taxes, and Capital Gains

In this section we use data on the entire distribution of incomes across households to reconcile our findings with the conventional wisdom that low-income households are the most affected by booms and recessions and that this greater sensitivity is due to higher cyclicity of hours worked among this group. Further, in studying the entire distribution, we also characterize how the tax-and-transfer system changes the cyclicity of take-home income. Finally, we track individual families rather than the income distribution across years as an alternative to using repeated cross-sectional data.

II.A. Relating Our Findings to the Conventional Wisdom

Previous studies have shown that the incomes of low-income households are more cyclical because unemployment falls primarily on low-wage workers (Clark and Summers 1981, Kydland 1984), whereas the wages of low-wage households have approximately the same exposure to the business cycle as those of high-wage households (Solon, Barsky, and Parker 1994). The flip side is that economic booms raise the standard of living of low-income households by more than they do high-income households (Card and Blank 1993, Hines, Hoynes, and Krueger 2001). Rebecca Blank (1989, p. 142), for example, concludes that "the income distribution narrows in times of economic expansion." There are three reasons why the conventional wisdom might not have detected the high cyclicity of top incomes: first, the time period, since high-income cyclicity began to rise only in the 1980s; second, the focus on broad groups, since cyclicity is high only for the very top of the distribution; and third, the top coding of incomes in conventional survey datasets, since this masks changes in income at the top end of the income distribution.

To begin, we track income groups using the March CPS public use microdata files for 1968–2008. The definition of income is the standard Census definition, namely, pre-tax, post-transfer income excluding capital

gains, and the unit of observation is a Census-defined family.¹⁶ We drop changes across years with major top code changes, and we note that after 1996 the data report the mean income for families above the top-coded amount, whereas before they simply report the income top code amount in place of actual income when top coded.

Following some of the earlier literature, the top panel of table 7 shows the cyclicity of incomes of low-education families (which are typically also lower-income families) and high-education families (typically higher-income families). Families are categorized according to the characteristics of the head, and we examine cyclicity with respect to average CPS income and the NIPA pre-tax, pre-transfer income (excluding capital gains) measure used in the earlier tables. Even during the period from 1982 on, there is some evidence that the conventional facts about cyclicity hold in that low-education households are more exposed to economic fluctuations than high-education households.

Turning to the top 1 percent, we show in the sixth column of table 7 that from 1968 to 1982, incomes in the top 1 percent of the distribution in the CPS were less cyclical than the average, with a beta of roughly 0.6. Thus, the previous literature is entirely correct about the early part of the period it studies. But when one focuses on the top 1 percent in the period since 1982 (next column), the cyclicity of that group's income is estimated to be 1.97 or 1.00, depending on the measure of income used—estimates as large as those for families with less than a high school education. The higher top 1 percent cyclicity after 1982 is presumably due to the later period, to increases in the top code level at several points over the years, and, after 1996, to the increased variability in amounts reported for top-coded observations.

Finally, we use a measure of the top 1 percent income share in the CPS constructed by Burkhauser and others (2008, 2009) using internal Census Bureau data for the CPS. These data measure top 1 percent income shares more accurately than is possible with the top-coded, publicly released microdata because the internal data are subject only to high-end censoring due to the number of digits allocated to the internally recorded income variable. That said, there are a number of additional issues with the accuracy of internal CPS top income data, and the series of Burkhauser and coauthors does

16. This definition has the benefit of dropping households consisting of unrelated individuals, but the disadvantage of dropping single individuals without children. The latter account for only about 10 percent of households in the top 1 percent, but 24 percent of the general population, according to the 1995 SCF.

Table 7. Cyclicalty of Income by Education and of Top 1 Percent, Using CPS Data^a

<i>Beta</i>	<i>Cyclicalty by level of education, using public use data on families, 1982–2008</i>						<i>Cyclicalty of top 1 percent</i>			
	<i>Less than high school</i>	<i>Some high school</i>	<i>Some college</i>	<i>College graduate</i>	<i>More than college</i>		<i>Using public use data on families</i>	<i>1982–2006, series 1</i>	<i>Using public use and internal data^b</i>	<i>1982–2006, series 2</i>
With respect to CPS income	1.52 (0.19)	0.94 (0.13)	0.91 (0.14)	1.08 (0.14)	0.73 (0.21)		0.55 (0.30)	1.97 (0.75)	2.57 (0.58)	2.44 (0.35)
With respect to NIPA income	0.85 (0.34)	0.67 (0.19)	0.67 (0.20)	0.78 (0.22)	0.63 (0.23)		0.64 (0.46)	1.00 (0.83)	1.81 (0.84)	2.18 (0.62)

Source: Authors' regressions using Census CPS data and data from Burkhauser and others (2010).

a. Public use data are March CPS public use data files from 1968 to 2009; see the online data appendix for further details. When using these data, we drop changes across the following years with major top code changes: 1980–81, 1983–84, 1994–95, and 2001–02. In addition, when using internal CPS data for households (series 1) we drop changes for 1992–93, when the Census allocated more digits to the internal record that stores the income series; series 2, covering tax units, is adjusted for this jump. We do not use income data for 1973–75 because the 1975 data are missing weights, and the 1973 and 1974 data have more than 1 percent of income data top coded. Income for defining the top 1 percent and in the regressions is CPS family income, which is pre-tax, post-transfer income excluding capital gains. Standard errors are in parentheses.

b. Data are based on internal CPS files and are from Burkhauser and others (2010), provided (in updated version) by the authors. Series 1 from Burkhauser and others (2010) measures the top 1 percent income share in the distribution of households, and series 2 measures top 1 percent income share in the distribution of tax units.

not show as significant an increase in top income shares as the tax data (see Burkhauser and others 2008, 2009, and Atkinson and others 2010). Despite these caveats, as displayed in the last two columns of table 7, these internal CPS survey data show an even higher cyclicity of the top 1 percent than the public data, and one that is very similar to that of the Piketty-Saez data from table 3.¹⁷

The previous literature, furthermore, shows that the cyclicity of the incomes of low-income families is largely due to the cyclicity of their hours worked. We now show that hours cyclicity plays only a minor role for the cyclicity of the top 1 percent. First, using all families in the CPS, we calculate average usual hours worked per week in each year for different income groups. For each group we regress the change in log average hours on the change in log real NIPA income (pre-tax, pre-transfer income, excluding capital gains), using data for 1982–2008. The cyclicity of hours for the top 1 percent is 0.26 (but with a standard error of 0.30), which is similar to the cyclicity of hours for all families, which we estimate to be 0.22 (with a standard error of 0.06). Thus, although the results are weak statistically, there is no evidence of a different cyclicity in hours for the top 1 percent. Second, we use the CPS hours data to adjust our measure of the wage and salary income of the top 1 percent from the Piketty-Saez data. We regress log growth in wages and salaries on log growth in hours and use the residual in place of the original wage and salaries series in our analysis of cyclicity. The cyclicity of “hours-adjusted” wages and salaries is estimated to be 2.2 for the top 1 percent, only slightly lower than the cyclicity of the unadjusted series, which is 2.4 (table 3). A similar exercise for the bottom quintile (using merged SOI-CPS data on the bottom quintile’s income, as described in the next subsection) finds that most of the cyclicity for that group is in fact due to the cyclicity of hours worked, consistent with the previous literature.¹⁸ Our analysis implies that, in contrast with the bottom end of the distribution, most of the cyclicity of the

17. To use these data, which report “true” (internal CPS) income shares, we infer the missing amount of income of the top 1 percent (denoted x) and thus also the missing amount from total income in the public CPS (denoted Y) from the relationship $(Y + x) \times$ (internal CPS income share) = $(Y^{1\%} + x)$, where $Y^{1\%}$ is the total income of the top 1 percent of families in the public data. Burkhauser and others (2010) provide two series for the internal CPS top 1 percent income share, one based on households and one based on tax units. Since neither matches our choice of CPS unit of analysis perfectly, we show results based on both.

18. Castro and Coen-Pirani (2008) provide further analysis of the cyclicity of hours, focusing on a comparison of college-educated with non-college-educated individuals.

top 1 percent is due to fluctuating payments for work rather than fluctuating hours worked.

We conclude that our results on income cyclicity both by education group in the CPS and for the top 1 percent before 1982 support the conventional view that low incomes are more cyclical. However, after 1982, even in this conventional survey dataset which has top-coded incomes, high income cyclicity is observable for the top 1 percent, and even higher cyclicity can be measured from versions of the data not subject to the top coding imposed on the public release files. Furthermore, this high cyclicity does not appear to be driven by cyclicity in hours worked, as it is for the bottom income groups. We now turn to a dataset from the Congressional Budget Office (CBO) that combines information from the CPS and the SOI data and allows us to study the entire distribution of income without the confounding issues of top coding.

II.B. Cyclicalities across the Full Income Distribution

To study the complete income distribution, this subsection employs a dataset from the Congressional Budget Office (2008) that merges data from the IRS SOI and data from the CPS to estimate average household income for different groups of households in different years. The two most important differences between the SOI-CPS data from the Congressional Budget Office and the SOI data used in tables 1, 2, and 3 are the unit of analysis and the definition of income used to sort households. The unit of analysis is the household in the SOI-CPS data and a tax unit in the SOI data. In terms of income, in the SOI-CPS data, households are sorted on pre-tax income per effective householder *including* transfers and capital gains, whereas in the Piketty-Saez SOI data, tax units are sorted on pre-tax income *excluding* transfers and capital gains. Our online data appendix provides further details.

The SOI-CPS data confirm our earlier findings for top income groups for this different set of choices about income measurement and unit of analysis. Table 8 shows statistics on the income distribution and cyclicity across the first four quintiles, in detail for the top quintile, and then in further detail for the top 1 percent. Focusing first on wages and salaries and on pre-tax, pre-transfer income excluding capital gains, as in all our analysis up to this point, we find (top panel) that the top 1 percent in the SOI-CPS data earn about 11 times, and the top 0.01 percent about 150 times, the average income; both these results are fairly similar to those reported in the second panel of table 2. The second panel of table 8 shows that all household groupings except the top 1 percent get 60 to 70 percent of their

Cash transfers	7.27	40.43	18.23	10.35	5.87	3.72	3.02	2.38	0.87	1.59	0.37	0.11
Capital gains	5.13	0.36	0.37	0.49	0.77	1.32	2.35	5.90	31.39	23.34	41.52	92.48
Taxes	24.22	10.16	15.55	18.71	21.21	23.51	25.25	28.08	41.22	43.37	45.90	62.78
Post-tax, post-transfer income including capital gains	88.19	130.63	103.04	92.14	85.43	81.53	80.12	80.19	91.05	109.97	95.99	129.81

Beta of indicated income source or measure with respect to aggregate (NIPA) pre-tax, pre-transfer income excluding capital gains

Wages and salaries	0.83	0.79	1.00	0.63	0.78	0.70	0.44	0.67	2.40	1.40	4.15	5.88
Pre-tax, pre-transfer income excluding capital gains	(0.08)	(0.38)	(0.34)	(0.18)	(0.11)	(0.16)	(0.14)	(0.15)	(0.84)	(0.59)	(1.27)	(2.34)
Pre-tax, post-transfer income excluding capital gains	0.90	0.76	0.90	0.66	0.69	0.67	0.67	1.01	2.16	0.77	3.07	3.33
Pre-tax, post-transfer income including capital gains	(0.12)	(0.32)	(0.27)	(0.13)	(0.08)	(0.12)	(0.10)	(0.16)	(0.73)	(0.77)	(1.00)	(1.71)
Post-tax, post-transfer income excluding capital gains	0.78	0.41	0.61	0.48	0.59	0.61	0.66	1.01	2.16	0.78	3.06	3.34
Pre-tax, post-transfer income including capital gains	(0.11)	(0.24)	(0.20)	(0.11)	(0.08)	(0.10)	(0.09)	(0.16)	(0.72)	(0.74)	(0.99)	(1.70)
Post-tax, post-transfer income including capital gains	1.07	0.45	0.64	0.51	0.63	0.68	0.79	1.24	3.28	1.82	4.16	5.47
Post-tax, post-transfer income including capital gains	(0.22)	(0.24)	(0.20)	(0.11)	(0.08)	(0.12)	(0.12)	(0.26)	(1.03)	(0.48)	(1.62)	(2.09)
Post-tax, post-transfer income including capital gains	0.91	0.38	0.50	0.40	0.49	0.51	0.65	1.11	3.48	2.35	4.56	6.09
Post-tax, post-transfer income including capital gains	(0.21)	(0.24)	(0.19)	(0.12)	(0.10)	(0.12)	(0.14)	(0.29)	(1.16)	(0.77)	(1.81)	(2.30)

Fraction of aggregate (merged IRS SOI and CPS) income change borne by indicated group

Pre-tax, pre-transfer income excluding capital gains	1.00	0.034	0.097	0.115	0.158	0.107	0.064	0.132	0.288	0.048	0.098	0.065
Post-tax, post-transfer income including capital gains	1.00	0.016	0.041	0.056	0.100	0.075	0.063	0.140	0.482	0.185	0.159	0.161
Post-tax, post-transfer income including capital gains		(0.013)	(0.020)	(0.018)	(0.018)	(0.011)	(0.009)	(0.012)	(0.052)	(0.020)	(0.018)	(0.022)

Sources: Authors' calculations and regressions using IRS SOI and CPS data merged by the Congressional Budget Office. See the online data appendix for details.

a. Income measures are average income per household. The distribution of income is measured across individuals and is based on household pre-tax, post-transfer income including capital gains, with income adjusted for household size by dividing by the square root of the number of people in the household. Standard errors are in parentheses.

income from wages and salaries. This number drops to 44 percent for the top 1 percent, and 27 percent for the top 0.01 percent.¹⁹ The first two rows of the third panel confirm our main findings on the post-1982 cyclicity of top income groups (compare this panel with the second panel of table 3). For the top 1 percent, both wages and salaries and overall pre-tax, pre-transfer income (excluding capital gains) per householder are more than twice as cyclical as the average income of all households, and for households in the top 0.01 percent, both wages and salaries and overall income are more than three times as cyclical as the average.²⁰

Second, the first two rows of the third panel of table 8 show that the incomes of households in the bottom two quintiles are a bit more cyclical than those of households from the middle quintile up to the 90th to 95th percentile. Thus, even in this period of high exposure of very high income groups, households in the lowest income quintile still have a slightly higher cyclicity of income than households in the middle and upper-middle parts of the distribution, but a much lower cyclicity than those at the top end.²¹

In sum, the recent cyclicity of wages and salaries and pre-tax, pre-transfer income is asymmetrically U-shaped, higher for the bottom two quintiles than for the middle and upper-middle part of the income distribution, and dramatically higher for the top 1 percent and the top 0.01 percent.

II.C. Cyclicity and Transfers, Taxes, and Capital Gains

The different levels and cyclicality of the incomes of different groups in the income distribution lead to different levels and cyclicality of taxes and transfers, and therefore different cyclicality of disposable income and ultimately of consumption. In this section we document that taxes and transfers reduce the cyclicity of income except at the very top. We also investigate the role of capital gains.

19. The somewhat smaller role of wages for the top groups than in the top panel of table 3 is probably due to the fact that households in the SOI-CPS data are sorted by an income measure that includes capital gains.

20. For comparability with earlier tables, the right-hand-side variable in this panel is (as in tables 1, 2, and 3) the log growth rate in real NIPA pre-tax, pre-transfer income per tax unit, excluding capital gains. The alternative would be to use the aggregate income from the SOI-CPS data. This leads to similar results.

21. The cyclicity of the bottom quintile in the SOI-CPS data is not as high as one might have expected from the cyclicity of low-education households in the CPS. When using CPS data with families sorted on income rather than education, we find a cyclicity for the bottom quintile in the CPS that is similar to that found for the bottom quintile in the SOI-CPS data.

First, the top two panels of table 8 show that adding transfers to our definition of income raises the incomes of the lowest quintile substantially but makes only a small difference to the incomes further up the distribution; the ratio of top income to average income falls slightly, since aggregate income is higher when transfers are included. Next, adding capital gains to income works the same way at the other end of the distribution, increasing the incomes of the top groups and so raising their relative incomes, while lowering the relative incomes of the bottom groups. Finally, subtracting taxes lowers the incomes of the top groups the most and so raises the relative incomes of the bottom quintiles.

Second, the third panel of table 8 shows that the income cyclicality of the bottom income groups are significantly reduced by transfers, which are large for the bottom quintile (about 40 percent of pre-tax, pre-transfer, pre-capital gains income) and countercyclical. The cyclicity of income for the bottom quintile falls from 0.76 to 0.41 as a result of transfer income alone, and that of the second quintile falls from 0.90 to 0.61. Third, capital gains increase cyclicity for all groups, and the importance of capital gains rises steadily with income, corresponding to the larger fraction of income coming from capital gains for higher-income groups. Including capital gains raises the income cyclicity of the top 1 percent from 2.2 to 3.3.²² Finally, taxes modestly lower the cyclicity of income for groups below the 99th percentile, while increasing cyclicity for the top 1 percent.

The fourth panel of table 8 summarizes the impact of different income levels and cyclicality by calculating the fraction of aggregate income changes borne by each group. On average, the top 1 percent bears 29 percent of changes in aggregate pre-tax, pre-transfer income excluding capital gains, and as much as 48 percent of changes in aggregate post-tax, post-transfer income including capital gains.²³

Overall, the cyclicity of the middle income groups is more stable across different income measures than that of the top and bottom of the income distribution. The cyclicity of the lowest income groups is significantly reduced by transfers, and that of the top income groups is significantly raised by including realized capital gains.

22. Notice that the betas reported in the third panel of table 8 are all with respect to aggregate pre-tax, pre-transfer income excluding capital gains.

23. In order for these fractions to sum to 1 across groups, we base aggregate income changes on aggregates from the SOI-CPS data themselves. To avoid potential biases in our estimates of betas from having SOI-based data on both the left- and the right-hand side of the regression, we omit the growth rates for the years around the 1986 tax reform (1985–86, 1986–87, and 1987–88).

II.D. The Cyclicalities of Same-Family Income

So far, because we use datasets that have good coverage of the top end of the income distribution, our analysis measures the cyclicalities of the average income of the top 1 percent of the income distribution rather than the cyclicalities of a given set of tax units or households. The top 1 percent of the distribution contains somewhat different people from year to year. Could the cyclicalities of the change in incomes of the group of people that start in the top 1 percent be different from the cyclicalities of the distribution that we have estimated so far? Such a difference could arise, for example, from a correlation between individual income risks and aggregate fluctuations. We have already provided, in our consumption analysis in section I.D, some evidence of high cyclicalities in data covering a constant set of households from one period to the next. Here we further investigate the cyclicalities of same-family income in two ways.

First, we link families across our March CPS extracts (which we also used in section II.A) for 1982–2009. In each year we categorize families into percentiles based on the entire distribution of families, and then we take the subsample of those that can be tracked to the following survey year and calculate the change in average income for each income group from this set of families. Thus, we calculate the annual log change in average income for groups of families that, in the first year of the change, are all within a certain part of the income distribution. Because of the small number of families in the top 1 percent that can be linked across years, the standard errors of the cyclicalities estimated for the top 1 percent in regressions parallel to those in table 7 are very large, around 1.6. For the top 5 percent, the sample is larger and the standard errors are somewhat smaller. The cyclicalities of same-family incomes for the top 5 percent are estimated to be 1.46, with a standard error of 0.80, with respect to average CPS income and 0.80, with a standard error of 0.86, with respect to NIPA income.

Second, in the next section we turn to tax data in which we can track the same families over time. Doing so requires using tax data from another country, but one that has also had an increase in top income inequality.

We can summarize the main results of sections II.A through II.D as follows: First, it is harder to observe high income cyclicalities in the top 1 percent in the public use CPS data, because of top coding and the fact that cyclicalities are high only for very high income families. Second, in looking at the entire distribution of incomes, the cyclicalities of pre-tax, pre-transfer incomes excluding capital gains are asymmetrically U-shaped: they are slightly higher for the bottom two quintiles of the income distribution

than for the next groups up to around the 95th percentile (and even up to the 99th percentile when focusing on wages and salaries), and very high for the top 1 percent and especially the top 0.1 and 0.01 percent. Third, transfers significantly reduce cyclical-ity at the bottom of the income distribution, essentially equating cyclical-ity across the distribution except for the top. The realization of capital gains raises the cyclical-ity of incomes at the very top even higher; taxes generate a small additional increase in cyclical-ity at the top.

III. Canada

Saez and Michael Veall (2007) show that Canada has also had a large increase in income inequality at the high end of the income distribution that roughly coincides temporally with the U.S. increase but is slightly less extreme. Canada has a slightly different tax system and culture but presumably is affected by the same changes in economic factors, such as technology and trade, as the United States. Thus, to provide another observation on the cyclical-ity of top incomes and to provide information about possible causes, we analyze the cyclical-ity of Canadian top incomes. There are also a number of ways in which the Canadian data are better than the U.S. data, most notably in that we can track the same families across years.

Our data come from the Longitudinal Administrative Databank, which contains records for 20 percent of all tax returns filed in Canada from 1982 to 2007. Working with Statistics Canada, we extracted information on the average incomes of families in different groups in the income distribution, both as repeated cross sections and tracking the families in different groups in the income distribution in a given year into the following year, as we were able to do with a subset of the CPS.²⁴ Further, we obtained data on income by source, as in the SOI data from Piketty and Saez, and on taxes and transfers, as in the SOI-CPS data from the CBO. We asked Statistics Canada to rank households and construct groups based on income calculated from pre-tax, pre-transfer income excluding capital gains.

Table 9 summarizes our results on the cyclical-ity of pre-tax, pre-transfer income excluding capital gains for different income groups in Canada with respect to aggregate Canadian income fluctuations for both sampling

24. The Canadian tax system is based on the individual, so tracking families involves summing income across family members (legal and common law spouses and children) if more than one live at the same address.

Table 9. Cyclicity of Income by Measure of Income and Income Group in Canada, 1982–2007^a

<i>Measure or source of income</i>	<i>All families</i>	<i>Ratio of average income in group to average for all families</i>										
		<i>Lowest quintile</i>	<i>Second quintile</i>	<i>Middle quintile</i>	<i>Fourth quintile</i>	<i>80th–90th percentile</i>	<i>90th–95th percentile</i>	<i>95th–99.0th percentile</i>	<i>Top 1 percent</i>	<i>99.0th–99.9th percentile</i>	<i>99.9th–Top 0.01 percent</i>	
Pre-tax, pre-transfer excluding capital gains	1.00	0.04	0.32	0.73	1.26	1.85	2.43	3.41	8.81	6.99	20.12	70.59
Pre-tax, post-transfer excluding capital gains	1.00	0.21	0.43	0.76	1.19	1.70	2.20	3.07	7.86	6.25	17.90	62.66
Pre-tax, post-transfer including capital gains	1.00	0.22	0.43	0.75	1.18	1.68	2.19	3.10	8.34	6.59	19.38	65.82
Post-tax, post-transfer including capital gains	1.00	0.26	0.49	0.80	1.19	1.65	2.09	2.87	6.89	5.57	15.27	50.77
<i>Average share of indicated source in pre-tax, pre-transfer income excluding capital gains (percent)</i>												
Wages and salaries	78.4	65.7	59.5	74.2	83.2	86.2	85.5	77.9	60.5	58.1	65.0	70.0
Pensions	5.5	18.0	14.5	9.8	5.5	3.5	2.9	2.9	2.2	2.5	1.7	0.7
Business and professional income	5.6	-20.1	6.2	4.7	3.6	3.4	4.0	8.2	17.4	21.0	10.6	3.0
Interest and dividends	6.6	27.9	11.6	6.6	4.5	4.1	4.6	7.0	14.6	13.1	17.4	21.2
Other investment income	3.9	8.5	8.2	4.7	3.2	2.8	2.9	4.0	5.3	5.2	5.4	5.1
Pre-tax, pre-transfer income excluding capital gains	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Cash transfers	12.8	571.3	51.3	16.8	7.0	3.5	2.2	1.5	0.7	0.8	0.4	0.1
Capital gains	3.5	27.1	4.0	2.6	2.1	2.1	2.5	4.2	9.4	8.8	11.6	8.3
Taxes	21.5	9.0	10.7	16.1	19.6	21.5	23.0	26.0	35.9	34.2	40.1	40.3
Post-tax, post-transfer income including capital gains	137.8	707.3	166.0	135.5	128.7	127.2	127.7	131.6	145.9	143.9	152.1	148.6

	<i>Beta with respect to average pre-tax, pre-transfer income excluding capital gains for all families</i>											
Pre-tax, pre-transfer income excluding capital gains	0.94 (0.13)	6.21 (1.27)	1.86 (0.31)	1.06 (0.16)	0.73 (0.10)	0.64 (0.09)	0.64 (0.09)	0.75 (0.11)	1.58 (0.29)	1.26 (0.21)	2.17 (0.45)	2.98 (0.85)
Pre-tax, post-transfer income excluding capital gains	0.71 (0.11)	0.36 (0.33)	0.80 (0.21)	0.67 (0.15)	0.59 (0.10)	0.62 (0.10)	0.74 (0.11)	1.57 (0.29)	1.25 (0.21)	2.16 (0.45)	2.97 (0.85)	
Pre-tax, post-transfer income including capital gains	0.79 (0.11)	0.43 (0.31)	0.85 (0.19)	0.71 (0.13)	0.64 (0.09)	0.67 (0.10)	0.84 (0.11)	1.84 (0.39)	1.50 (0.29)	2.63 (0.62)	3.02 (0.91)	
Post-tax, post-transfer income including capital gains	0.71 (0.12)	0.37 (0.32)	0.72 (0.20)	0.63 (0.15)	0.61 (0.12)	0.63 (0.10)	0.77 (0.13)	1.64 (0.48)	1.33 (0.40)	2.51 (0.77)	2.24 (0.95)	
	<i>Same-household beta with respect to average pre-tax, pre-transfer income excluding capital gains</i>											
Pre-tax, pre-transfer income excluding capital gains	0.90 (0.13)	7.40 (3.29)	1.73 (0.36)	1.01 (0.15)	0.77 (0.11)	0.70 (0.12)	0.72 (0.16)	0.93 (0.23)	1.58 (0.35)	1.48 (0.35)	2.20 (0.46)	1.60 (0.93)
Pre-tax, post-transfer income excluding capital gains	0.68 (0.12)	0.51 (0.71)	0.86 (0.26)	0.65 (0.15)	0.60 (0.11)	0.64 (0.15)	0.87 (0.23)	1.56 (0.34)	1.45 (0.34)	2.18 (0.46)	1.59 (0.93)	
Pre-tax, post-transfer income including capital gains	0.76 (0.12)	0.57 (0.72)	0.90 (0.26)	0.71 (0.15)	0.65 (0.10)	0.70 (0.15)	1.00 (0.23)	1.82 (0.41)	1.68 (0.38)	2.54 (0.54)	1.85 (1.07)	
Post-tax, post-transfer income including capital gains	0.69 (0.12)	0.50 (2.53)	0.83 (0.94)	0.68 (0.44)	0.61 (0.64)	0.62 (0.25)	0.88 (0.12)	1.58 (0.13)	1.48 (0.17)	2.38 (0.25)	1.02 (0.49)	

Source: Authors' calculations and regressions using data extracts from the Longitudinal Administrative Databank at Statistics Canada.

a. Individuals are summed within families, and families are ranked by pre-tax, pre-transfer income excluding capital gains in each year. Aggregate income is market income (personal income less transfers) per family. All betas for income measures that include capital gains exclude changes to and from 1994, because that year is an outlier due to a change in tax law (see Saez and Veall 2007). Standard errors are in parentheses.

procedures (same households from year to year, and not). First, focusing on wages and salaries and pre-tax, pre-transfer income excluding capital gains, the top panel of table 9 shows (comparing with table 2) that the ratio of income of the top 1 percent to average income is somewhat lower in Canada than in the United States, although this point should be qualified by possible differences in tax laws and tax avoidance by high-income households between the two countries.²⁵

Second, comparison of the second panel of table 9 with table 3 shows that the top 1 percent in Canada and in the United States get similar shares (about 60 percent) of their income from wages and salaries. However, in Canada the top 0.01 percent get about 70 percent of their income from wages and salaries, compared with only 40 percent in the United States (from table 3).

Third, turning to our main point of interest, the third panel of table 9 shows that top incomes in Canada, as in the United States, are highly cyclical in the period since 1982. In Canada the top 1 percent and the top 0.01 percent have cyclicalities of 1.6 and 3.0 in the recent period, slightly lower than the corresponding cyclicalities in the United States (top panel of table 2), which are 2.4 and 4.0, respectively. The next two sections argue that this pattern across the two countries—higher cyclicity for those at the top of the income distribution—is representative of a close relationship and potentially a common cause of both high income shares and high cyclicity at the top in the period since the early 1980s.

Fourth, table 9 also shows the effect of capital gains, taxes, and transfers in Canada. Looking across rows in the third panel reveals that in Canada the government has little effect on the cyclicity of incomes at the top of the distribution. At the bottom, however, the effect of transfers is far larger in Canada than in the United States (table 8). The beta for the lowest income quintile before taxes and transfers is over 6, compared with 0.76 for the United States, whereas that after transfers is 0.36, quite similar to the 0.41 in the U.S. data. Although one might be tempted to credit the Canadian welfare state, it seems unlikely that the United States and Canada are truly so different in the exposure of pre- versus post-transfer incomes. Instead, the large impact of transfers on the cyclicity of the bottom group in Canada is likely due to very low average pre-tax, pre-transfer

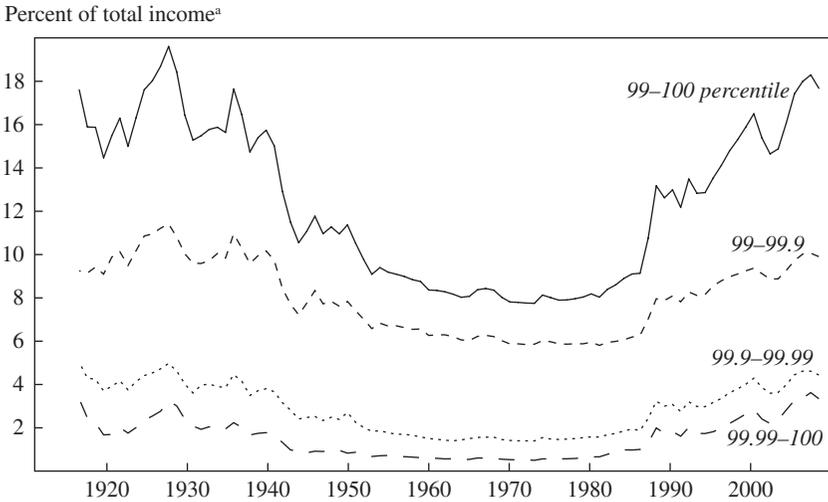
25. We compare the Canadian data with our results in tables 2 and 3 rather than table 8 because tables 2 and 3 (like the Canadian data) are based on sorting households using pre-tax, pre-transfer income excluding capital gains, whereas the data underlying table 8 are available only sorting households using pre-tax, post-transfer incomes including capital gains.

incomes for this group (with very low average incomes, even moderate transfers can change the cyclicity substantially). Lower pre-tax, pre-transfer incomes for the bottom group in Canada are due in large measure to the Canadian groups being defined in terms of an income measure that excludes transfers, and to the SOI-CPS data in table 8 excluding households with negative income from the bottom income category.

Finally, the bottom panel of table 9 shows that in Canada the income changes from one year to the next that occur for those households who are in the top 1 percent in the first year also have a high cyclicity with respect to changes in aggregate Canadian income, roughly similar to that found in the third panel using repeated cross-sectional data. This is something we could not observe in the U.S. tax data. Thus, the cyclical exposure from one year to the next for families that start in the top 1 percent of the income distribution (but who may fall elsewhere in the distribution in subsequent years) is similar to the cyclical exposure of the annually reported top 1 percent of the income distribution (a group that contains somewhat different families from year to year). This is less so, however, for the top 0.01 percent. The three groups of families that start in the various income groups within the top 1 percent (bottom panel) have similar cyclicality, whereas for the same three groups in the annually reported top 1 percent of the distribution (third panel), the top 0.01 percent have (economically) significantly higher cyclicity than the other two groups. Nonetheless, the fact that we estimate high cyclicality for the top 1 percent in both cross-sectional data and panel data is evidence against the hypothesis that the cyclicity of top incomes in panel data in the United States would be quite different from what we have estimated from repeated cross-sectional data. To conclude our discussion of the bottom panel of table 9, we note that the roles of taxes, transfers, and capital gains are broadly similar to those in the third panel.

IV. The Empirical Link between Income Cyclicity and Income Shares at the Top

Having explored in detail the rise in the cyclicity of high incomes in the last three decades, we now show that this increase is closely related to the rise in the share of income accruing to the top of the income distribution. Specifically, we present three pieces of evidence that *the higher the level of income inequality, the higher the income cyclicity of the rich*. We exploit variation across groups, time, and countries. First, in the post-1982 period, the higher a group is in the income distribution (within the top 1 percent),

Figure 2. Income Share of the Top 1 Percent, 1917–2008

Source: Piketty and Saez (2003), updated by Saez (2010).

a. Income is defined as pre-tax, pre-transfer income excluding capital gains.

the higher is that group's income cyclicality. Second, across decades, as the income share of a given top group increases, the cyclicality of its income increases, consistent with the fact that the increase in the income share of the top 1 percent starts at almost exactly the same time as the increase in the income cyclicality of that group. Third, across countries, those with larger increases in the income share of the top 1 percent have experienced larger increases in the income cyclicality of the top 1 percent. This tight empirical link between inequality and cyclicality at the top end of the income distribution in the past quarter century strongly suggests that these two phenomena share a common cause.

Before we turn to this evidence, figure 2 complements the basic facts displayed in tables 1 and 2 by plotting the income shares from the Piketty-Saez data.²⁶ These data show both that the dramatic rise in top income shares started in the early 1980s, when cyclicality also increased, and that cyclicality and the income share of the top income groups are not linked in the first half of the 20th century. Top income shares were large in the pre-war period, a period in which we do not find evidence of higher cyclicality of the incomes of the top 1 percent. Piketty and Saez (2003) argue that

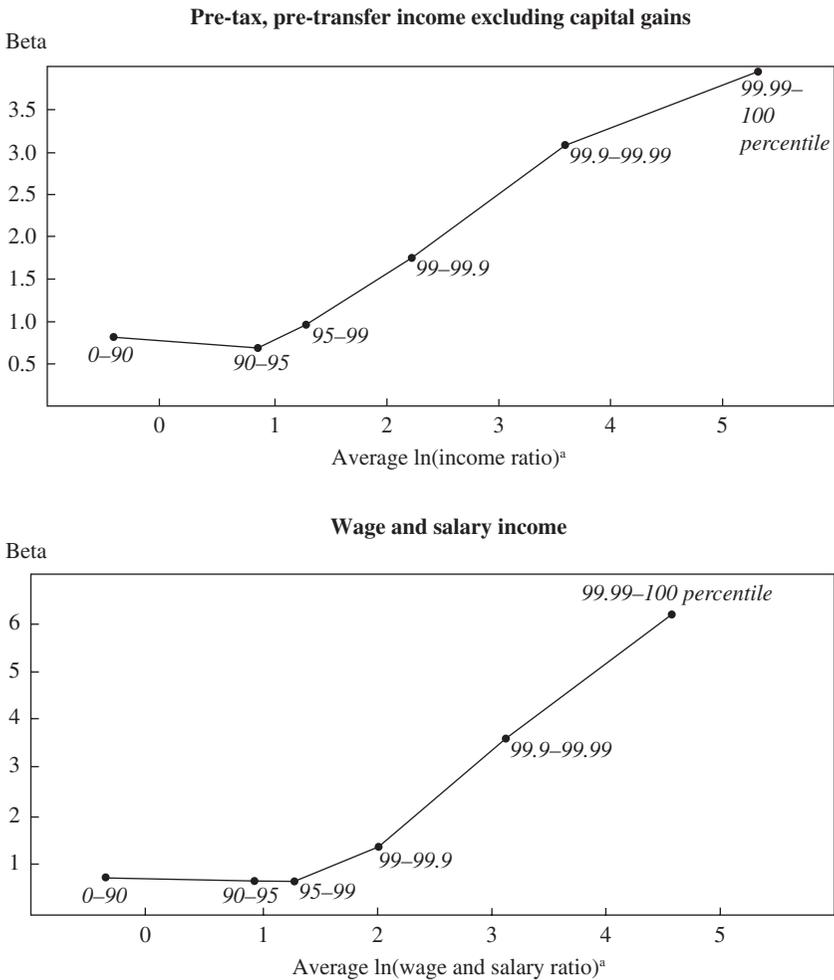
26. As in tables 1, 2, and 3, these shares are for income excluding capital gains, and the data come from Piketty and Saez (2003) as updated in Saez (2010).

different factors drove the income shares of the top 1 percent during the period of declining inequality than during the later period of increasing inequality. They argue that the decline in the income share of the top 1 percent, and of the highest-income groups within the top 1 percent, from around 1928 to around 1970 was driven in large part by declines in capital income (income from dividends and interest), which were in turn due to a combination of the Great Depression and the large tax increases enacted to finance the war; these included large increases in corporate income taxes that almost mechanically reduced distributions to stockholders. In contrast, an increase in wage and salary income is the key driver of the more recent increase in the income share of the top 1 percent. The lack of correspondence between top 1 percent income share and income cyclicality together with the different income composition in the earlier period suggests that the decline in top income shares from 1928 to 1970 was not driven by the same factors as the more recent increases. This is consistent with our explanation for the recent changes: the ICT revolution did not happen in reverse in the early to middle part of the 20th century.

Our first piece of additional evidence of a link between the cyclicality and the income shares of the top 1 percent is that, for groups further up the income distribution within the top 1 percent, there is both a larger income share (relative to the size of the group) and a larger income cyclicality during the period since 1982. Figure 3 graphs the cyclicality of income over the period 1982–2008 for each income group (using data from tables 2 and 3 and the same calculations for other income groups) against the time-series average of the log ratio of that group's average income to the average income of all tax units.²⁷ The first panel of figure 3 focuses on pre-tax, pre-transfer income excluding capital gains, and since we argue that the high cyclicality of wage and salary income is a key driver of the high overall cyclicality of the incomes of the top 1 percent, the second panel focuses on wage and salary income alone. It is apparent from both graphs that groups higher up in the income distribution within the top 1 percent have both higher ratios of income to average income and higher income cyclicality. Inequality at the top is extreme: the incomes of the top 0.01 percent are on average 212 times the average income (see the second panel of table 2). Similarly, cyclicality at the top is extreme: that of the top 0.01 percent is about four times that of the average (six times when one focuses on

27. The betas depicted are from table 3 and are based on growth rates for 1983–2008. The average log income ratios are calculated as the time-series average of the log income ratio (average group income for the year to average income for all tax units for that year), using income ratios for the initial year of each growth rate used (1982–2007).

Figure 3. Betas and Log Ratios of Group Income to Average Income, by Income Group, 1982–2008^a



Source: Authors' calculations based on data from Piketty and Saez (2003), extended by Saez (2010).
 a. Average across years of the log of the group's average income divided by aggregate average income.

wages and salaries only). This again suggests a link between the level of income inequality and income cyclicalities.

A simple statistical description of the relationship is that β in equation 1 is a function of the average log income ratio: $\beta_i = \lambda_0 + \lambda_1(1/T)\sum \ln(Y_{it}/Y_t)$, where the summation is across the T years for each income group i , so that equation 1 becomes

$$(2) \quad \Delta \ln Y_{i,t+1} = \alpha_i + \beta_i \Delta \ln Y_{i,t+1} + \varepsilon_{i,t+1} = \\ \alpha_i + \lambda_0 \Delta \ln Y_{i,t+1} + \lambda_1 \left[\Delta \ln Y_{i,t+1} (1/T) \Sigma \ln(Y_{i,t}/Y_t) \right] + \varepsilon_{i,t+1}.$$

We estimate this relationship by stacking data for percentiles 0–90, 90–95, 95–99, 99–99.9, 99.9–99.99, and 99.99–100 using the growth rates for 1983–2008, for a total of 156 observations. Using pre-tax, pre-transfer income excluding capital gains results in an estimate of λ_1 of 0.65. Using wage and salary data results in an estimate of λ_1 of 1.61. Both estimates are significant at the 1 percent level, showing that cyclicity increases with income share across groups.

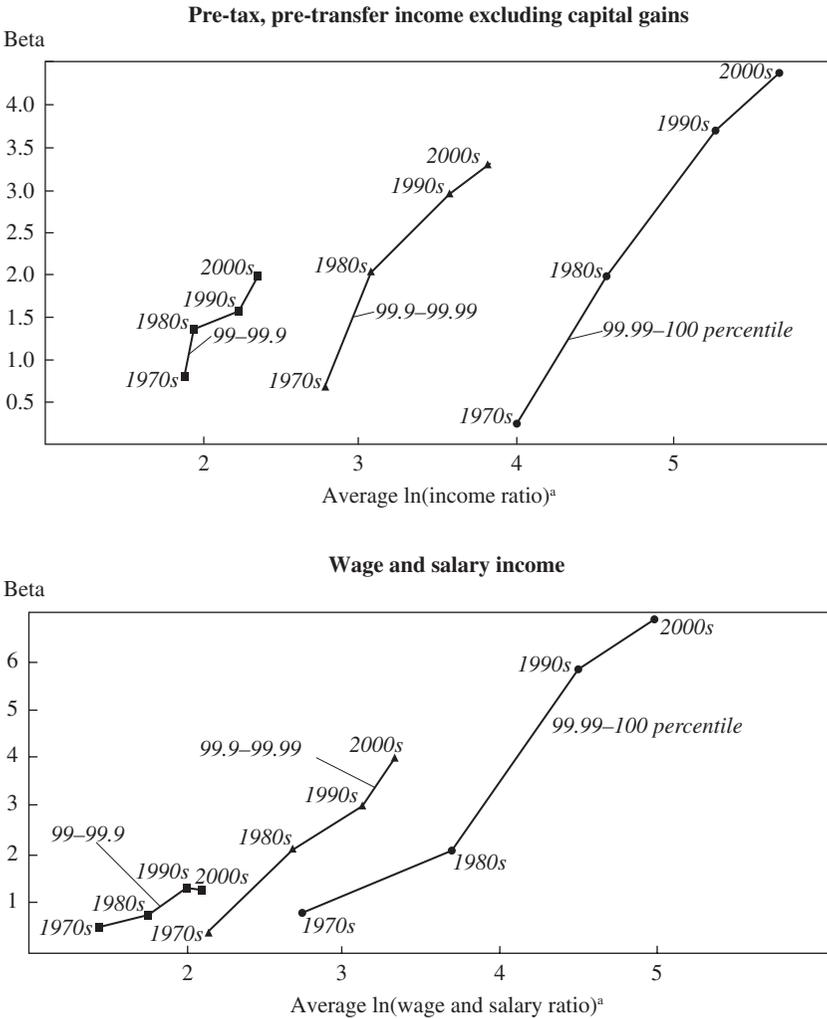
Second, over time since top income shares first began to rise, as a group's income share has increased, so has its cyclicity. To show this, we estimate betas for each high-income group and decade since the 1970s, and the time-series average of the log ratio of that group's average income to the average income of all tax units for each group and decade. Figure 4 plots decadal betas against decadal average log income ratios. Again the top panel focuses on pre-tax, pre-transfer income excluding capital gains, and the second panel on wage and salary income. For each group, both cyclicality and average log income ratios increase over time, leading to a positive association between a group's cyclicity and its average income ratio. This pattern is present both in overall income and in wage and salary income. Notice that when one connects the points by decade, as is done in figure 4, rather than by group, it becomes clear that the relationship between average log incomes and cyclicality is strengthening over time: no relationship was apparent in the 1970s, whereas a strong relationship is observed in the 2000s.

A statistical description of the relationship underlying figure 4 is that β in equation 1 is a function of the log income ratio, now allowing for time-series variation in the ratio, so $\beta_i = \lambda_0 + \lambda_1 \ln(Y_{i,t}/Y_t)$. Equation 1 then becomes

$$(3) \quad \Delta \ln Y_{i,t+1} = \alpha_i + \beta_i \Delta \ln Y_{i,t+1} + \varepsilon_{i,t+1} = \\ \alpha_i + \lambda_0 \Delta \ln Y_{i,t+1} + \lambda_1 \left[\Delta \ln Y_{i,t+1} \ln(Y_{i,t}/Y_t) \right] + \varepsilon_{i,t+1}.$$

We estimate this relationship separately for each of the three subgroups of the top 1 percent—percentiles 99–99.9, 99.9–99.99, and 99.99–100—using the growth rates for 1970–2008. We include $\ln(Y_{i,t}/Y_t)$ itself as an additional regressor to avoid potentially spurious significance of the variable of interest, $\Delta \ln Y_{i,t+1} \ln(Y_{i,t}/Y_t)$. Using pre-tax, pre-transfer income excluding

Figure 4. Betas and Log Ratios of Group Income to Average Income across Decades^a



Source: Authors' calculations based on data from Piketty and Saez (2003), extended by Saez (2010).
 a. Average across years of the group's average income divided by aggregate average income.

capital gains results in estimates of λ_1 of 2.09 ($t = 1.75$) for percentiles 99-99.9, 2.37 ($t = 2.01$) for percentiles 99.9-99.99, and 2.48 ($t = 2.35$) for percentiles 99.99-100. Using wage and salary data results in estimates of λ_1 of 1.46 ($t = 2.25$), 2.86 ($t = 2.31$), and 3.03 ($t = 2.60$) for the three groups, respectively. This confirms, using time-series variation in income shares

(log income ratios) within groups, that higher income shares are associated with higher cyclicalities.

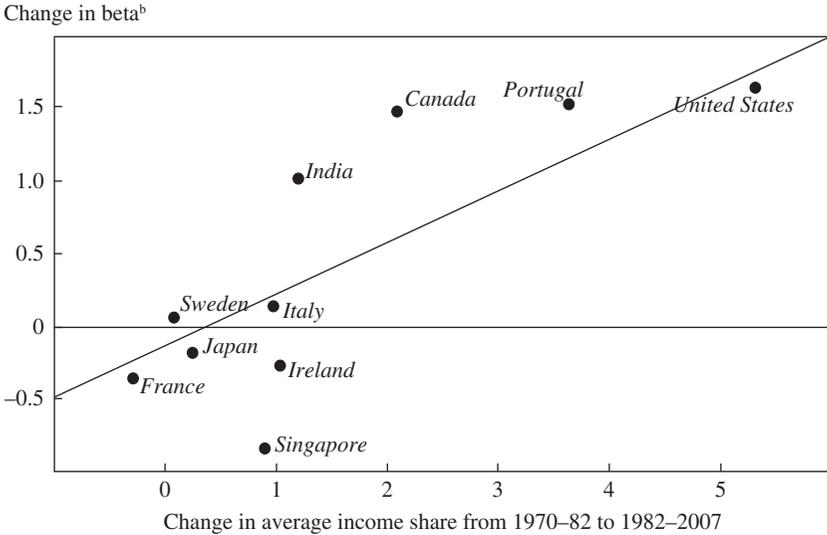
Third, we show that the countries with the largest increases in the income shares of the top 1 percent also have experienced the largest increases in the cyclicalities of incomes of that group. We use the dataset constructed from tax records from Atkinson, Piketty, and Saez (2010), which contains annual time-series data for the incomes of the top 1 percent for 22 countries. We focus on relating changes in top income shares to changes in top income cyclicalities rather than on post-1982 levels of each variable, because of the differences in tax systems across countries and the consequent differences in measurement of top income shares, as well as the host of other differences that exist across countries.²⁸ We estimate income cyclicalities for the top 1 percent in the period from 1982 onward (the period for which we found higher top 1 percent income cyclicalities for the United States) and for the period 1970–82 (as a benchmark period). Of the 22 countries, we exclude 6 (Australia, Finland, Germany, New Zealand, Norway, and the United Kingdom) for which income measures include capital gains and 1 (Switzerland) for which incomes are not available at an annual frequency.²⁹ Furthermore, we require countries to have at least five observations of growth rates in the 1970–82 period and five in the 1982–2007 period, leading us to drop another 5 countries (Argentina, China, Indonesia, Netherlands, and Spain). This leaves 10 countries (Canada, France, India, Ireland, Italy, Japan, Portugal, Singapore, Sweden, and the United States). The original data for Canada extend only to 2000, but we obtained data up to 2007 from Michael Veall.³⁰ As shown in figure 5, there is a positive relationship between the increase in top 1 percent beta and the increase in top 1 percent income shares. The fitted value is from an ordinary least

28. An example is whether the unit of analysis is the family or the individual. See Atkinson, Piketty, and Saez (forthcoming), table 3 and related text.

29. We include Finland in the set for which incomes include capital gains. Although it is possible to calculate top 1 percent income shares from the original article (Jantti and others 2010), it appears infeasible to calculate aggregate totals that fully exclude capital gains.

30. The updated data from Veall start in 1982. We use the original data for earlier years and the updated data from 1982 on, with growth rates always calculated using income data from the same dataset. In the dataset from Atkinson, Piketty, and Saez (2010), U.S. aggregate income is based on tax records. This may bias upward the beta of the top 1 percent in the period that includes the 1986 tax reform (if incomes are distorted by tax reform effects more for the top 1 percent than for other groups). We therefore drop the growth rates for 1987 and 1988 for the United States but could alternatively use NIPA aggregate income as in our earlier analysis (the 1986 growth rate is not an outlier in this dataset).

Figure 5. Changes in Betas and Changes in Income Shares of the Top 1 Percent in 10 Countries, 1970–2007^a



Source: Authors' calculations based on data from Atkinson, Piketty, and Saez (2010).

a. Line represents fitted value.

b. Beta for the top 1 percent for 1982–2007 minus beta for the top 1 percent for 1970–82.

squares regression relating the change in top 1 percent beta to the change in average top 1 percent income shares. The slope coefficient in this regression is 0.42 (the heteroscedasticity-robust standard error is 0.07) and the R^2 is 0.64.

Overall, these three different approaches all suggest that in recent decades, the greater is the top 1 percent income share, the higher is income cyclicity for those in the top 1 percent.

V. Technological Change and Changes in High-Income Shares and Cyclicity

This section argues using a simple example that increases in the scale at which top earners operate naturally lead to increases in both income and income cyclicity at the top of the distribution. We do not provide additional tests to support this interpretation of the facts. Instead we intend in this section to put forward an additional theory, to be considered in future work, about the underlying economic drivers behind these two phenomena.

V.A. Existing Theories for Increasing Top Income Shares

The leading explanation for the broad increase in wage and income inequality that started in the 1970s is that technological change over this period has complemented the skills of highly skilled workers (see, for example, Autor, Katz, and Kearney 2008, Acemoglu and Autor 2010). There is also evidence that changes in economic institutions or regulation (such as minimum wages and unionization) have increased income inequality at the lower end of the distribution. At the very top of the distribution, Piketty and Saez (2003) argue that the speed and size of the increases in relative earnings are inconsistent with the main existing theories based on skill-biased technological change, and that the evidence from top income shares may instead suggest an important role for changing social norms with respect to high earnings. Finally, there is a well-developed literature on the rise in relative compensation for a subset of top earners, namely, corporate chief executive officers (CEOs). Several explanations have been proposed for the rise in relative CEO pay, including a shift in social norms regarding compensation, an increase in managerial power (rent extraction, captured boards), a shift in demand from specific to general skills, an increase in the size of firms, and skill-biased technological change (Kaplan and Rauh 2010, Bertrand 2009).

Top executives are, however, a minority of highly compensated individuals. Steven Kaplan and Joshua Rauh (2010) document that only about 5 percent of earners in the top 0.01 percent are executives of nonfinancial firms. They also show that investment bankers, other financial asset managers (at hedge, venture capital, private equity, and mutual funds), lawyers, and to a smaller extent athletes and celebrities all make up significant fractions of the top income groups.³¹ Kaplan and Rauh argue that the fact that pay has increased dramatically at the top in each of these sectors is evidence against the first three explanations above. Neither social norms nor increased managerial power seems relevant for the pay of many occupations among top earners, such as hedge fund managers, and specific rather than general skills seem more important for lawyers, hedge fund managers, investment bankers, and professional athletes. Expanding on the argument of Xavier Gabaix and Augustin Landier (2008) that increased CEO pay can be explained by increased firm size, Kaplan and Rauh further show

31. Bakija and others (2010) provide much more detailed occupational information, indicating that another large subgroup within the highest income groups is people in medical occupations.

that the leading financial services firms, law firms, and hedge, venture capital, and private equity funds have grown larger over time (measured by inputs or output). This does not fully explain the increase in CEO pay (and the top 1 percent income share), however, since average firm size was increasing before 1980, too. What is needed to explain these facts is that the impact of firm size on top 1 percent pay is higher than it was before, as might arise from skill-biased technological change favoring those at the top. This would amount to a mix of the last two theories listed above.

Of these existing theories, which also predict an increase in the income cyclicality of top earners? The canonical theories of skill-biased technological change require a separate assumption that the technology that complements skill has a very cyclical impact on those at the top of the income distribution.³² Other theories of rising pay at the top similarly require additional assumptions—that the ability of CEOs to “steal” is cyclical or that norms about pay are highly cyclical (for example, because high pay or conspicuous consumption is more stigmatized during recessions).

Although changing institutions and regulations, power structures, or norms may have a role in the changes we have observed, we argue that these changes are not inconsistent with a theory of skill-biased technological change—specifically, changes in ICT—in which these changes have increased the scale at which the top earners operate. We show theoretically that if advances in ICT have increased the ability to scale the application of high skills, this naturally implies both that top incomes will rise and that fluctuations in demand over the business cycle will affect the incomes of the highest-skill individuals disproportionately. The next subsection describes this mechanism, leaving empirical tests or calibrations for future work.

V.B. A Theory of Why Very High Income Individuals Have Higher and More Cyclical Incomes than in the Past

The rise of ICT has allowed the most skilled in any given occupation to apply their talents more broadly, for example, to manage more workers and capital, to entertain more people, or to write more papers. Thus, ICT has lowered the extent to which quality declines when more output is produced; in other words, it has made marginal revenue curves decline more

32. One approach would be to assume that the latest technologies that complement the skills of the most highly paid are tied to new investment (in physical capital of higher quality, in equipment and software, or in organizational capital). Then, since investment is highly procyclical, skill-biased technological change could lead to both higher incomes and higher cyclicality of incomes for those with the highest skills.

slowly with output. This change has raised the operational scale and the earnings of the most skilled. The highest earners tend to have larger fluctuations in their earnings than the rest of the population because those who operate at a large scale naturally have lower profit margins and so are more exposed to cyclical fluctuations.

The following simple model formalizes this argument and illustrates how those with higher incomes tend also to have more cyclical incomes.³³ Let each worker produce earnings according to

$$(4) \quad py - ci = Ai^\alpha - ci,$$

where $0 < \alpha < 1$ and $\alpha A > c$. Further assume that workers earn the full net revenue they contribute to the firm, so that earnings are $\pi = Ai^\alpha - ci$. Very highly skilled workers have higher α than they had in the past, whereas changes in α for lower-skilled workers are zero or comparatively small. A higher α means that a worker's marginal product diminishes less rapidly as the input i increases. The assumption $\alpha A > c$ ensures that high- α workers earn more than low- α workers. The key change in our earnings function that generates both increased cyclicity and increased earnings shares for highly skilled workers is our conjecture that growth in ICT has increased α for very highly skilled workers during the period since 1982.

Three different interpretations of our revenue or earnings function are useful. First, the most obvious interpretation is that all workers produce output of identical quality, but the best produce more for given inputs and have less diminishing returns to scale. In our equations this corresponds to i being inputs, $y = i^\alpha$, $p = A$, and c being the price of the inputs. In this interpretation the ICT revolution increases the returns to scale (that is, reduces the degree of decreasing returns to scale) of the best workers and allows them to work with more inputs; for example, a CEO can manage a larger company.

A second interpretation of equation 4 is that ICT has changed markets so that highly skilled workers are more like superstars in the sense of Sherwin Rosen (1981). That is, highly skilled workers produce the same number of units of output for given inputs as other workers, but as they produce more output, the quality of that output declines more slowly than that of other workers does. As in Rosen (1981, p. 849), "superior talent stands out and does not deteriorate so rapidly with market size as inferior talent does." In this interpretation the ICT revolution has lowered how quickly

33. Our earnings function is in the spirit of the equilibrium model of Lucas (1978).

quality declines with output for the best workers. In our equations this corresponds to i being both input and output ($y = i$), the price p being a function of quality that decreases with output as $p(y) = Ay^{\alpha-1}$ (so $py = Ai^\alpha$), and c being the marginal cost of producing another unit of output. For example, the top lawyers (in the post-1982 world) may be able to write more briefs without the quality of their legal advice suffering as much as would be the case for less skilled lawyers (for example, because of the impact of ICT on the ease with which case histories can be researched).

A final and closely related interpretation is that of an asset manager paid based on performance. In this interpretation let i be assets under management, c the expected return investors can earn elsewhere, and Ai^α the (risk-adjusted) trading profits of the fund.³⁴ The earnings equation then captures the idea that the best fund managers are increasingly able to invest more money without the returns on their investments deteriorating as much as for other fund managers.

Given our assumptions, the optimal level of i is

$$(5) \quad i^* = (A\alpha/c)^{1/(1-\alpha)},$$

with associated earnings for the worker of

$$(6) \quad \pi = (A\alpha/c)^{1/(1-\alpha)} c(1-\alpha)/\alpha.$$

Equation 6 delivers our two main results.

First, because workers with higher α earn higher incomes (by assumption), income inequality and top income shares increase when the α of top earners increases. In equation 6, $d\pi/d\alpha > 0$. This occurs because highly skilled workers generate more revenue for given inputs, and they are optimally matched with more inputs because they have less decreasing returns to scale.

Second, an increase in the α of high- α workers increases the cyclical-ity of their earnings. Assume that business cycle fluctuations are driven by fluctuation in A , representing either market demand shocks or technology shocks. The percentage change in profits depends on α as

$$(7) \quad d\ln \pi / d\ln A = 1/(1-\alpha),$$

34. This structure corresponds to a setting with no management fees and a 100 percent carry, but the results should generalize to a more standard contract with a 2 percent fee and 20 percent carry.

which is positive and increasing in α . Thus, the cyclicity of the earnings of a worker increases if the worker's α rises. Note that $d \ln \pi / d \ln A$ does not depend on whether i is adjusted optimally in response to the change in A , since by the envelope theorem, $d\pi/di = 0$ at the initial value of A . Thus, the high cyclicity of earnings is driven not by a higher cyclicity of inputs, but by the increased sensitivity of earnings to demand that comes from working with a higher level of inputs. That said, the input scale of more highly skilled workers is more cyclical in this model; it is just not the cause of greater income cyclicity.

The intuition for the different cyclicalities is that more highly skilled workers optimally are working with more inputs (interpretation 1) or producing more output (interpretations 2 and 3). Their revenue is similarly sensitive to fluctuations as those of lower-skill individuals, since fluctuations in A affect the revenue of each worker proportionately. However, for highly skilled workers, costs are higher relative to revenue, so ci is more substantial relative to Ai^α for these workers. This makes their earnings, $\pi = Ai^\alpha - ci$, more cyclical. For low-skilled workers, inputs are small relative to revenue, so fluctuations in earnings are in percentage terms more similar to fluctuations in revenue.

VI. Conclusion

Coinciding with the increase in the income share of top earners since the early 1980s has been an increase in the cyclicity of the incomes of top earners. The high cyclicity that we document for top incomes, including wages and salaries, appears to be linked empirically to increases in the income shares of top earners, based on variation over time, across groups of top earners, and across countries. This increased cyclicity and its link to increased income shares should contribute to a better understanding of the reasons behind the increase in top income shares.

We propose that the information and communications revolution provides a natural way to think about how technological change may have raised both top income shares and top income cyclicity. The change in technology that we suggest—increased scale or increased “superstar”-type production by top earners—generates a simple connection between income shares and cyclicity, in that the earnings of those operating on a larger scale naturally become more sensitive to the business cycle. Our brief analysis of our posited mechanism leaves open the question of how well it can quantitatively match the documented changes in cyclicity over time and across countries.

ACKNOWLEDGMENTS For helpful comments we thank the editors, our discussants Rebecca Blank and Erik Hurst, as well as Jeff Campbell, Xavier Gabaix, Takashi Yamashita, participants at the Brookings Panel conference, and seminar participants at Kellogg School of Management and the Federal Reserve Banks of Cleveland and Chicago. We thank David Autor and Melanie Wasserman for help with the CPS files. We thank Brian Murphy and Habib Saani of Statistics Canada for help with the Longitudinal Administrative Databank. We thank Jeff Larrimore for providing the data from Burkhauser and others (2008, 2009) and Michael Veall for providing data from Saez and Veall (2007) updated to 2007. We thank the Zell Center and Kellogg School of Management for funding. Nicolas Ziebarth provided excellent research assistance.

The authors report no relevant potential conflicts of interest.

References

- Acemoglu, Daron, and David Autor. 2010. "Skills, Tasks and Technologies: Implications for Employment and Earnings." Working Paper no. 16082. Cambridge, Mass.: National Bureau of Economic Research.
- Aguiar, Mark, and Mark Bilts. 2010. "Has Consumption Inequality Mirrored Income Inequality?" Working paper. University of Rochester.
- Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. Forthcoming. "Top Incomes over a Century or More." *Journal of Economic Literature*.
- Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104, no. 6: 1227–62.
- Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in U.S. Wage Inequality: Revising the Revisionists." 2008. *Review of Economics and Statistics* 90, no. 2: 300–23.
- Bakija, Jon, Adam Cole, and Bradley T. Heim. 2010. "Jobs and Income Growth of Top Earners and the Causes of Changing Income Inequality: Evidence from U.S. Tax Return Data." Working paper. Williams College (November).
- Bertrand, Marianne, 2009. "CEOs." *Annual Review of Economics* 1: 1.1–1.29.
- Blank, Rebecca M. 1989. "Disaggregating the Effect of the Business Cycle on the Distribution of Income." *Economica* 56, no. 222 (May): 141–63.
- Bound, John, and George Johnson. 1992. "Changes in the Structure of Wages in the 1980's: An Evaluation of Alternative Explanations." *American Economic Review* 82, no. 3: 371–92.
- Burkhauser, Richard V., Shuaizhang Feng, Stephen P. Jenkins, and Jeff Larrimore. 2008. "Estimating Trends in US Income Inequality Using the Current Population Survey: The Importance of Controlling for Censoring." Working Paper no. 14247. Cambridge, Mass.: National Bureau of Economic Research (August).
- . 2009. "Recent Trends in Top Income Shares in the USA: Reconciling Estimates from March CPS and IRS Tax Return Data." Working Paper no. 15320. Cambridge, Mass.: National Bureau of Economic Research (September).
- Card, David, and Rebecca M. Blank. 1993. "Poverty, Income, Distribution, and Growth: Are They Still Connected?" *BPEA*, no. 2: 285–339.
- Castro, Rui, and Daniele Coen-Pirani. 2008. "Why Have Aggregate Skilled Hours Become So Cyclical since the Mid-1980s?" *International Economic Review* 49, no. 1: 135–84.
- Clark, Kim B., and Lawrence H. Summers. 1981. "Demographic Differences in Cyclical Employment Variation." *Journal of Human Resources* 16, no. 1: 61–79.
- Congressional Budget Office. 2008. *Historical Effective Tax Rates, 1979 to 2005: Supplement with Additional Data on Sources of Income and High-Income Households, and Methodology*. Washington (December). www.cbo.gov/publications/collections/tax/2009.cfm.
- Cutler, David M., and Lawrence F. Katz. 1991. "Macroeconomic Performance and the Disadvantaged." *BPEA*, no. 2: 1–74.

- Gabaix, Xavier, and Augustin Landier. 2008. "Why Has CEO Pay Increased So Much?" *Quarterly Journal of Economics* 123, no. 1: 49–100.
- Goolsbee, Austan. 2000. "What Happens When You Tax the Rich? Evidence from Executive Compensation." *Journal of Political Economy* 108, no. 2 (April): 352–78.
- Hall, Brian J., and Jeffrey B. Liebman. 2000. *The Taxation of Executive Compensation*. Tax Policy and the Economy series, vol. 14. Cambridge, Mass.: National Bureau of Economic Research.
- Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13, no. 1 (January): 15–51.
- Hines, James R., Hilary Hoynes, and Alan B. Krueger. 2001. "Another Look at Whether a Rising Tide Lifts All Boats." In *The Roaring Nineties: Can Full Employment Be Sustained?* edited by Alan B. Krueger and Robert Solow. New York: Russell Sage Foundation.
- Jäntti, M., M. Riihelä, R. Sullström, and M. Tuomala. 2010. "Trends in Top Income Shares in Finland." In *Top Incomes: A Global Perspective*, edited by A. B. Atkinson and T. Piketty. Oxford University Press.
- Kaplan, Steven N., and Joshua Rauh. 2010. "Wall Street and Main Street: What Contributes to the Rise in the Highest Incomes?" *Review of Financial Studies* 23, no. 3: 1004–50.
- Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–87: Supply and Demand Factors." *Quarterly Journal of Economics* 107, no. 1: 35–78.
- Kuznets, Simon. 1953. *Shares of Upper Income Groups in Income and Savings*. New York: National Bureau of Economic Research.
- Kydland, Finn E. 1984. "Labor-Force Heterogeneity and the Business Cycle." *Carnegie Rochester Conference Series on Public Policy* 32: 173–208.
- Lucas, Robert E., Jr. 1978. "On the Size Distribution of Business Firms." *Bell Journal of Economics* 9, no. 2 (Autumn): 508–23.
- Moylan, Carol E. 2008. "Employee Stock Options and the National Economic Accounts." *Survey of Current Business* 88, no. 2 (February): 7–13.
- Parker, Jonathan A., and Annette Vissing-Jørgensen. 2009. "Who Bears Aggregate Fluctuations and How?" *American Economic Review* 99, no. 2 (May): 399–405.
- Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *Quarterly Journal of Economics* 118, no. 1: 1–39.
- . 2007. "Response by Thomas Piketty and Emmanuel Saez to: The Top 1% . . . of What? By Alan Reynolds." University of California, Berkeley.
- Reynolds, Alan. 2007. "Has U.S. Income Inequality Really Increased?" *Policy Analysis* 586 (January 8): 1–24.
- Rosen, Sherwin. 1981. "The Economics of Superstars." *American Economic Review* 71, no. 5 (December): 845–58.

- Saez, Emmanuel. 2010. "Tables and Figures Updated to 2008." University of California, Berkeley (July). elsa.berkeley.edu/~saez/TabFig2008.xls.
- Saez, E., and M. R. Veall. 2007. "The Evolution of High Incomes in Canada, 1920–2000." In *Top Incomes over the 20th Century: A Contrast between Continental European and English-Speaking Countries*, edited by A. B. Atkinson and T. Piketty. Oxford University Press.
- Solon, Gary, Robert Barsky, and Jonathan A. Parker. 1994. "Measuring the Cyclicity of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics* 109, no. 1 (February): 1–25.

Comments and Discussion

COMMENT BY

REBECCA M. BLANK¹ This paper by Jonathan Parker and Annette Vissing-Jorgensen is highly interesting. Its primary conclusion, that incomes have become markedly more cyclical at the very top of the income distribution in the past 25 years, is surprising and intriguing. The paper presents a new fact about the world that was not previously known, and this makes it likely that the paper will stimulate further research and debate.

For an empirical economist, there is much to like in this paper. The authors do an extremely thorough job of data analysis. They use multiple datasets to confirm and test their results, with substantial attention to proving the robustness of what they find. Any careful reader will come away impressed by the serious data work in the paper and persuaded that the cyclicity of incomes among the top 1 percent of U.S. households has indeed increased. That said, as with most papers that uncover new facts, there is more work to be done to understand and interpret this result, so that it informs the theoretical framework that economists use when thinking about income generation, inequality, and macroeconomic change.

It is important to be clear about what the results in this paper do *not* show. The greater cyclicity that the authors discuss appears to be focused at the very top of the income distribution, particularly among the top 1 percent of households. Hence, this result does not overturn the frequently noted result that incomes are more cyclical among lower-income families than among higher-income families. On average, income in the bottom quintiles is more cyclical than in the middle quintile, as the authors' table 8 demonstrates. Furthermore, the authors reiterate the fact that the cyclical

1. These comments reflect the personal opinion of the author and do not necessarily represent the views of the Department of Commerce or the U.S. Government.

nature of unemployment, in particular, seems to lead to income cyclicalities among lower-income families.

On the question of who is most hurt by cyclical downturns, nothing in this paper refutes the widely held belief, buttressed by substantial evidence, that lower-income families (particularly those headed by someone with less education, working in a lower-wage job) experience greater economic deprivation in a recession than do other families. These families experience a disproportionate share of unemployment and are more likely than other families to need government assistance to survive economically during bad economic times. The fact that income and consumption patterns (as the authors show) are also highly cyclical at the very top of the income distribution is less likely to signal deprivation, although it may well create real stress within these families. Households in the top 1 percent of the income distribution have substantial savings and assets and can smooth their consumption if they wish. This means that the consumption cyclicalities that they experience (matching their income cyclicalities) is best viewed as an economic choice on their part. In contrast, consumption cyclicalities among very poor families who have no savings is much more likely to be an involuntary and unavoidable response to changes in earnings and income.

It would therefore be inaccurate to interpret the results in this paper as saying something about well-being. Parker and Vissing-Jorgensen are clear on this point, but it is worth stressing nonetheless. The results in this paper do, however, inevitably raise the question of why this cyclicalities has increased among households at the very top of the income distribution, particularly given the close relationship between rising cyclicalities and increases in absolute levels of inequality, which the authors document. At the end of the paper, Parker and Vissing-Jorgensen present a theory that focuses on changes in information and communications technologies (ICT) that have increased the ability of highly skilled persons to leverage their skills and expand their income, leading to rising inequality. The authors' model suggests that this exposes them to greater cyclical fluctuations.

I find this model a plausible story, although it is just that at the moment—a possible story, without supporting evidence. To investigate whether the data support this theory, one would want to look at changes in earnings levels and cyclicalities among high-earning workers who might have greater “leverage” due to the ICT revolution, and among those who might be less affected by this phenomenon. Unfortunately, when one is exploring a phenomenon that is primarily visible in only the top 1 percent of the population, such investigations are hard to pursue.

What is happening in ICT may be only part of the change in the economic environment facing top-earning workers. The expanded global markets in which more and more companies are operating also provide scope to utilize the gains from ICT that did not exist before. It might have been useful for the authors to say more about globalization and how it relates to their theory.

My biggest hesitation about the causal hypothesis that Parker and Vissing-Jorgensen present is that it is unclear to me why it would be limited to workers at the extreme top of the income distribution. Both the greater global marketplaces and the expanded possibilities created by new ICT should have benefited many higher-skilled workers. The authors' results suggest that the increased income cyclicality they observe is closely related to cyclicality in wages and salaries among the topmost earners and does not reflect rising cyclicality in hours of work or in other forms of income. At a minimum, this suggests that compensation among the very top earners is more tied to overall economic performance than it is among workers even slightly lower in the earnings distribution. Perhaps additional theoretical structure is needed to explain why compensation practices at the very top differ from those even a little lower on the wage spectrum.

For instance, one question I would be very interested in knowing more about is how compensation packages for top earners differ across industries and occupations. Although Parker and Vissing-Jorgensen indicate that the top 1 percent of earners are spread across industries (table 5 in their paper), my guess is that there may be different compensation practices for (say) those who manage money for large manufacturing firms than for those who manage the firm's operations. And the ability of new ICT to enlarge the possible value generated by these different top managers might also vary. It may take a series of more micro-focused case studies, looking at very highly paid senior people in a selected group of industries and occupations, to better understand and investigate both the authors' theory and their empirical results.

Let me close with a comment about the data. As the authors note, it is extremely difficult to study the phenomenon of income cyclicality at the top because very few of the available datasets are large enough to produce a reasonable-sized sample among the top 1 percent of earners. And very few available datasets are accurate enough to produce informative data about that group, even if their samples were larger. Among survey statisticians there is widespread concern about lower survey response rates among the extremely wealthy. (Of course, sample weighting techniques can adjust for this, but a small number of observations with larger weights will lead to less accuracy.)

In addition, noisy data can lead to a top 1 percent sample that includes households whose actual income would not place them in this category. For all of these reasons, annual cross-sectional datasets based on relatively small samples of the population (such as the Current Population Survey or the Consumer Expenditure Survey) are probably of limited value in addressing the questions raised by this paper. For this reason, I would place less reliance on tables 5 and 6, which use those data, than on other results in the paper. Even if one combines a number of years' data together to produce a larger sample, data reliability questions may still pose problems for the researcher.

This means that there are probably two datasets best suited to look at this small sliver of the population: the Statistics of Income data, which the authors use intensively, and the American Community Survey (ACS). The ACS, which the authors do not use, replaced the old "long form" of the decennial census after the 2000 census. It collects information monthly on a wide variety of indicators (including income, earnings, and family composition) from a random sample of families. In any one year, the ACS samples a little over 1.9 million households. Although the ACS lacks data from before the 2000s, and so cannot be used to investigate long-term trends in cyclicity, it can be used to look in much greater depth at who the families and individuals are at the very top of the income distribution in recent years, and at how different types of households and families responded to the Great Recession. Those who want to explore these issues further should think about the possibilities provided by the ACS for this research question.

Overall, this is a fine paper. In some ways it merely adds to the puzzle of why and how inequality and earnings among very top earners have changed over the past 25 years. But by adding a new fact about income cyclicity, and closely linking that fact with rising incomes among these earners, the paper provides information that will help economists winnow out the various theories that have been proposed to explain widening inequality. The most believable explanations will be those that explain both the rising levels of income and the rising income cyclicity in this group.

COMMENT BY

ERIK HURST This paper by Jonathan Parker and Annette Vissing-Jorgensen documents an interesting, important, and novel set of facts pertaining to the cyclicity of income for very high income individuals. The paper shows that in recent years, households in the top 1 percent of the income distribution have much more cyclical incomes than most other households. Additionally, the paper shows that this high relative cyclicity

is a relatively recent phenomenon, that it moves in lockstep (decade by decade) with the well-documented increase in income inequality driven by the increasing income share for these households, and that it is robust to controlling for stock options, household fixed effects, and taxes and transfers. The facts are very carefully documented, and I have no comments whatsoever on the existing empirical work in the paper.

The second part of the paper lays out a simple theory to explain these facts. In particular, it asks what factors could possibly result both in an increasing share of income earned by very high income individuals and in an increasing cyclicity of income for those individuals. The authors propose a model where information and communications technologies have increased the optimal production scale for the most talented individuals. Nothing in the paper convincingly supports or convincingly refutes this theory. Rather, as the authors note, it is simply one theory that could simultaneously generate increasing income inequality and increasing cyclicity among those with very high incomes.

My comments are structured in two parts. First, I want to emphasize that the authors make no claims about the welfare costs of recessions. They are very clear about this. However, it is a point worth reemphasizing so that the paper's implications are not misconstrued. Second, I will offer some new facts related to the changing nature of compensation that took place for higher-income households during this period. In particular, bonus income increased in importance for high-income households during the 1990s and early 2000s. As I show below, bonus income is much more cyclical than other types of income and is more closely associated with the finance industry than with other industries.

WHAT THIS PAPER IS NOT ABOUT. Upon reading this paper, one is tempted to use the facts that it documents to make statements pertaining to the distributional costs of business cycle fluctuations. The authors caution readers against making such types of calculations. I want to underscore this point.

The authors show (convincingly) that the cyclicity of income is much higher for those with very high incomes than for other income groups and that this cyclicity has been increasing over time. Do these results imply that the cost of business cycles, in terms of standard utility-based measures of welfare, is higher for those with very high incomes than for those at other points of the income distribution? Do the results imply that over the last two and a half decades, those with very high incomes are bearing an increasing brunt of business cycle variation in terms of changes in welfare? The answer to both of these questions is a resounding no. Variations in

income (and, to a lesser extent, in consumption) do not map directly onto variations in standard, utility-based measures of welfare. Households with sufficient wealth can self-insure against income fluctuations by accumulating and then drawing down assets. Households can maintain consumption flows despite variation in consumption outlays by delaying the replacement of durables, and even some goods traditionally defined as nondurable, such as clothing or vacation spending, have aspects of durability. Finally, given standard assumptions about household preferences, concave utility functions imply that a given change in expenditure will have a much smaller effect on utility for individuals with very high expenditure than for individuals with lower expenditure.

Two other facts need to be emphasized. First, households with very high incomes may have anticipated the increase in risk to their incomes that the authors document, and if so, one would expect them to have demanded compensation for bearing that risk. This could explain the fact that those with very high incomes are earning higher returns on their labor and simultaneously facing more variable labor income streams. The story is analogous to the difference between investing in stocks and investing in bonds. If the earnings of those with very high incomes have become more stock-like (taking more of an equity stake in their employing firm through their labor investments), it is not surprising to see them bearing more risk and receiving higher returns. Second, and a related point, the variation in income for these households could be either transitory or permanent. In order to compute standard welfare calculations using income and expenditure data (even if one could measure the service flow of expenditure correctly), one needs to know whether the observed variation in income was perceived as a transitory shock or as a permanent shock. To the extent that business cycle variation implies differences in expectations about the evolution of the permanent component of income for individuals at different points of the income distribution, welfare calculations again become complicated.

Collectively, the results in this paper do not suggest that the brunt of business cycles in terms of changing well-being is being disproportionately borne by those with very high incomes in recent periods. What the paper does show is that the income of those at the top of the income distribution has become more cyclical. I view these results as potentially informative about the changing nature of compensation in the economy over the last few decades, not as an input into how we think about the distributional costs of cyclical variation. Like the authors, I would caution readers against using the paper's results to draw conclusions about how

cyclical variations affect well-being for individuals at different points of the income distribution.

THE INCREASING IMPORTANCE OF BONUS INCOME AT THE TOP. The paper left me with a few lingering questions about which components of earnings are driving the results. First, how important are bonuses for individuals at the top of the income distribution? Second, are bonuses more important for individuals in some professions than in others? Third, has the composition of bonus-receiving professions been changing over time? Fourth, is bonus income more cyclical than other types of income? Finally, can bonuses help explain the correlation between the increased share of income and the increased cyclical of income for very high income households?

Some of these questions are hard to answer with existing datasets. I will try to provide some information on some of these questions using data from the Panel Study of Income Dynamics (PSID), and will then discuss further why bonus income could help explain the facts documented in the paper. I wish to emphasize that these PSID results are meant to be only suggestive. The PSID is not an ideal dataset for analyzing the earnings behavior of very high income households, because of its limited sample size.

The PSID disaggregates labor earnings into the following categories: regular wage and salary income, bonus income, income from commissions, tips, overtime compensation, and business income. For my analysis I use data from the 1995, 1997, 1999, 2001, 2003, 2005, and 2007 waves of the PSID, and I pool bonus and commission income together, because commissions, like bonuses, could be related to work effort and could vary with the state of the aggregate economy. I restrict the sample to male heads of households between the ages of 16 and 70 who were currently employed and had positive earnings during the preceding year. The earnings reports I use are total earnings (from all sources) within a particular category from the preceding year. For example, bonus earnings reported in the 1995 wave of the PSID refer to all bonuses earned during calendar 1994. All earnings data are converted into 2000 dollars, when applicable.

To compute earnings percentiles, I rank all earnings for individuals within the sample separately for each year. Given the sample sizes, I classify households into the top 2.5 percentiles (the richest households), percentiles 2.5–5.0, percentiles 5.0–10.0, and the bottom 90 percentiles. I look at three measures: the share of households receiving either bonus or commission earnings, the share of total earnings that come from either bonus or commission earnings, and the fraction of household heads who work in the finance industry. As it turns out, the inclusion of commissions adds lit-

Table 1. Importance of Bonus Income across the Income Distribution, Pooled Years^a
Percent

<i>Indicator</i>	<i>Labor earnings percentiles</i>			
	<i>Bottom 90</i>	<i>5–10</i>	<i>2.5–5</i>	<i>Top 2.5</i>
Fraction receiving bonus income, all heads of household	9	19	22	29
Share of bonus income in total income, all heads of household	1	3	4	8
Share of bonus income in total income, bonus recipients only	11	17	20	28
Sample size	25,028	1,117	542	533

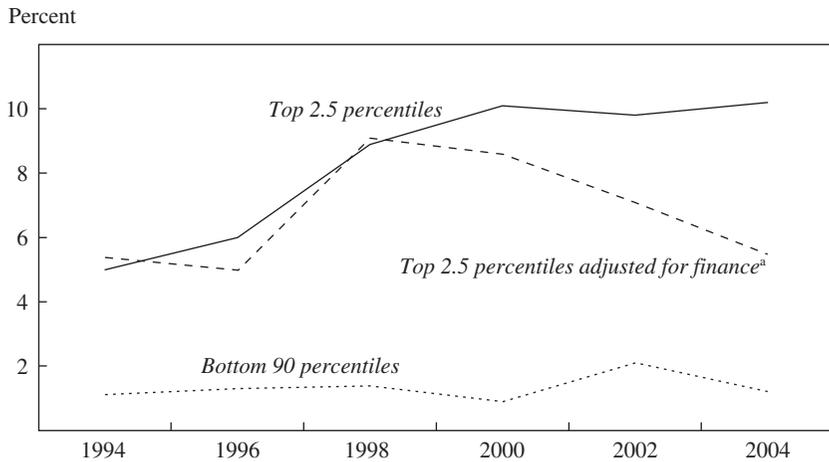
Source: Author’s calculations using data from the 1995, 1997, 1999, 2001, 2003, 2005, and 2007 waves of the Panel Study of Income Dynamics.

a. Sample includes all currently employed male heads of household between ages 16 and 70 who had positive income in the preceding year. Percentiles are defined within each year separately. All differences are statistically significant from each other except for the 5–10 percentile and 2.5–5 percentile comparisons.

tle to the analysis; essentially all the results are driven by bonuses rather than commissions, and therefore in what follows I refer to the sum of bonus and commission income simply as bonus income.

My table 1 shows, first, the fraction of household heads in each of the above percentile ranges who received bonus income. These results pool the data across all years. Only 9 percent of household heads in the bottom 90 percentiles of the income distribution received bonus income. For the other income groups, the fraction receiving bonus income rises with income, reaching 29 percent in the highest income group. The table also shows the average fraction of income that comes from bonuses across all households within the different percentile ranges. This is calculated as the simple average of the bonus share across all individuals within each range. This share likewise increases as one moves up the earnings ladder. For example, the average individual in the top 2.5 percentiles gets about 8 percent of earnings from bonuses, compared with 1 percent for the average individual in the bottom 90 percentiles. Finally, the third line of table 1 shows the average bonus share for those households who reported positive bonus income. The conclusion from table 1 is that bonus income is more important for higher-earning than for lower-earning households.

Figure 1 shows the time-series patterns in the bonus share of earnings for household heads in the bottom 90 percentiles and for those in the top 2.5 percentiles. The figure shows a dramatic increase in the share of income earned from bonuses between 1994 (from the 1995 survey) and

Figure 1. Share of Income from Bonuses, 1995–2007

Source: Author's calculations using data from the 1995, 1997, 1999, 2001, 2003, 2005, and 2007 waves of the Panel Study of Income Dynamics.

a. Plot of the coefficients on the year dummies in the regression reported in the second results column in table 2. Sample includes all currently employed male heads of household between ages 16 and 70 who had positive income in the preceding year.

2004 (from the 2005 survey) for the latter group. For example, whereas in 1994 roughly 5 percent of this group's earnings came from bonuses, in 2004 that figure was roughly 10 percent. In contrast, those in the bottom 90 percentiles show no discernable trend in the share of income earned from bonuses.

To summarize, the PSID results show that the share of income from bonuses among households at the top of the income distribution was increasing at the same time that these households, according to the data that the paper uses, were seeing both an increased share of total income and an increased cyclicity of income. This suggests that the rise in bonus income among these households may relate to the patterns documented by Parker and Vissing-Jorgensen.

Is there a statistical relationship between the receipt of bonus income and working in the finance sector? Steven Kaplan and Joshua Rauh (2010) show that individuals in the finance sector increased their share in the very top of the income distribution during the 1990s and the early 2000s. The same patterns hold in the PSID data. In 1994, 12 percent of individuals in the top 2.5 percent of the income distribution were in the finance industry; by 2004 this figure had risen to nearly 18.5 percent.

Table 2. Regressions Explaining Bonus Income with Finance Industry Employment and Income, Pooled Years^a

<i>Independent variable</i>	<i>Dependent variable</i>		
	<i>Dummy for positive bonus</i>	<i>Bonus share of total income</i>	
		<i>All heads of households</i>	<i>Bonus recipients only</i>
Dummy for employment in finance industry	0.075 (0.026)	0.041 (0.012)	0.159 (0.052)
Dummy for income in top 10 percentiles	0.103 (0.021)	0.024 (0.007)	0.057 (0.030)
Dummy for income in top 2.5 percentiles	0.087 (0.026)	0.043 (0.017)	0.084 (0.047)
Constant	0.091 (0.004)	0.009 (0.001)	0.104 (0.009)
Sample size	27,220	27,220	2,902

Source: Authors' regressions using data from the 1995, 1997, 1999, 2001, 2003, 2005, and 2007 waves of the PSID.

a. Sample includes all currently employed male heads of household between ages 16 and 70 who had positive income in the preceding year. Percentiles are defined within each year separately. Robust standard errors are in parentheses.

Table 2 shows the results of three regressions. Each regresses some measure of the importance of bonus income on a dummy variable indicating whether the individual is in the finance industry, a dummy for whether the individual is in the top 10 percent of the income distribution, and a dummy for whether the individual is in the top 2.5 percent of the income distribution. (If the individual is in the top 2.5 percent, both the top 10 percent dummy and the top 2.5 percent dummy have a value of 1.) I run these regressions on the entire pooled sample. As the table shows, being in the finance industry increases the likelihood of receiving a bonus, the share of income that comes from bonuses, and the share of income coming from bonuses conditional on receiving a bonus.

Given that the finance industry has been increasing in importance over time, a natural question is how much of the increasing share of bonus income for those individuals with very high income during the 1990s and early 2000s (documented above) was simply due to the increasing prominence of individuals in the finance industry in that group. To address this, I run two regressions on a sample that includes only those individuals in the top 2.5 percentiles of the income distribution. The first simply regresses the share of income from bonuses on year dummies. The second regresses the same dependent variable on year dummies, a dummy for whether the

individual was in the finance industry, and an interaction of the finance dummy with the year dummies. Figure 1 also plots the coefficients on the year dummies from these regressions and shows that a substantial part of the increase in the bonus share of earnings for this group, particularly after 1998, was due to the increasing importance of the finance industry.

The PSID data do not go back far enough in time to allow a full analysis of the cyclical nature of bonus income. However, it is not a leap to think that bonus income is more cyclical than other types of income, given that it is usually linked to firm performance or profits. If that is the case, then as bonus income has been a more important component of income for those with very high incomes, this could be a cause of the increased cyclical nature of income for these individuals.

What can the increasing importance of bonus income reveal about the relationship between the rising share of total income accruing to very high income individuals and the increased cyclical nature of income for these individuals? One possibility is that the facts documented in the paper are simply driven by the increasing share of very high income individuals working in the finance industry. On average, individuals employed in finance receive a larger share of their income as bonuses, and they are more likely to be represented among the very rich. Although this is likely to be some of the story, it is not the entire story. As shown in the paper, some evidence suggests that it is unlikely that the compositional switch within the group of very high income individuals toward finance solely explains their results.

The rise in importance of bonus income does reveal that the nature of compensation has been changing. Ex ante, higher-income individuals are relying more on bonus income as a form of compensation. Bonus income is more risky than some other forms of compensation in that it is directly tied to firm profits. To be willing to bear this risk, these high-income individuals need to be compensated for it. As a result, the shift toward bonus income can be consistent with the rising share of income for these households as well as with the increased cyclical nature of income. But why has the compensation structure changed such that those who had very high incomes to begin with are willing to bear this additional income risk? Are such risk-sharing agreements efficient, in that they better align incentives between the high-income workers and the firm? Are the high-income workers becoming synonymous with the firm itself? If these workers are now willing to take on more risk of the firm's profitability, does that imply that the other workers are now facing less risk? Does it imply that other *investors* in the firm are facing less risk? The facts in this paper should be leading

economists to ask a whole new series of questions about the allocation of risk within the economy.

SUMMARY. Overall, this is a very nice paper. The methodology is well executed, and the results are well documented. The question remaining is what is driving those results. The paper proposes one story. But there is nothing in the paper that confirms (or contradicts) this story. It appears that the changing nature of compensation of very high income individuals in the form of the rising importance of bonus income is potentially part of the story. The facts documented in the paper, collectively, should point researchers toward addressing a whole series of interesting questions.

REFERENCE FOR THE HURST COMMENT

Kaplan, Steve, and Joshua Rauh. 2010. "Wall Street and Main Street: What Contributes to the Rise in the Highest Income?" *Review of Financial Studies* 23, no. 3: 1004–50.

GENERAL DISCUSSION George Perry observed that developments in the financial sector can largely explain the sharp rise in economy-wide inequality between 1982 and 2008 that the authors analyze. In the authors' data, wages and salaries of the top 0.1 percent of the income distribution were \$183 billion higher in 2008 than if they had just kept up with the average rise since 1982. Wages and salaries per worker in finance rose nearly twice as much as the economy-wide average over this period, and total wages and salaries in finance in 2008 were \$154 billion higher than if the per worker average had simply kept up with the rest of the economy. Hedge funds, which were in their infancy at the start of the 1980s, managed an estimated \$2.5 trillion in 2008, which would account for roughly \$100 billion of financial wages and salaries that year. Hedge funds are also characterized by high earnings volatility, as are other leveraged financial activities that generate very high incomes and greatly expanded over this period. All this suggests that finance is not just part of the income distribution story but the dominant part. Economies of scale have always existed in finance. What has changed in the financial sector is the increasing application of leverage and risk.

Refet Gürkaynak found one of the most fascinating findings in this paper to be that not just income, but also consumption, has become more volatile at the top of the income distribution. That finding is surprising, because one would think that individuals at the very top also have sufficient wealth to smooth their consumption. He further suggested that the

paper's findings might be explained in terms of the standard risk-return relationship in finance. If a greater proportion of an individual's compensation is in the form of bonuses, which are more volatile than wages and salary, that individual would have to be compensated more on average to be willing to accept that risk.

Benjamin Friedman was likewise fascinated by the finding of higher consumption volatility at the high end. He proposed three potential explanations. First, available statistics other than those from tax records (which do not report consumption directly) are unreliable at the extremes, and so the finding might simply be spurious. Second, even though people at the top of the income distribution also have higher wealth-to-income ratios, much of that wealth is in illiquid form and so might not be available to smooth consumption. Third, consumption by people at the very top may be lumpier. Whatever the explanation, it was a puzzling finding that seemed to go against accepted knowledge. Erik Hurst added that even if consumption is volatile for the really wealthy, their utility is probably not much affected. Happiness data show that happiness is not more volatile for the very rich than for other households.

William Nordhaus thought that what might be going on at the very top end is that some people are able to impose a "tax" on the profits of companies that they control, in the form of bonuses, stock options, or perks. Because profits are cyclical, this income will also be cyclical. To the extent that compensation structures are becoming more incentive-based, moving away from a fixed base pay, this should contribute to making top incomes more cyclical. Nordhaus was also concerned that capital gains are a very large omitted part of income. To the extent that some cyclicalities of the capital gains component is not getting measured, that would be another explanation for the paper's finding.

James Hines noted that the Tax Reform Act of 1986 changed not only tax rates but also the definition of taxable income. Some of the difference in the proportion of income going to the very top depends on this definitional change. Also, because tax rates on the very rich are much lower today than in earlier decades, the rich have less incentive to avoid classifying some income as taxable income. The estate tax, which has seen an extreme reduction recently, also bears on the decisionmaking of top income groups. Assets can now be given to a trust in the name of a child and will not show up as income.

Gary Burtless argued that another important change was in the incentive to hold income within corporations as opposed to organizing the firm so that the income is immediately treated as though distributed to all of the

owners. Before the Tax Reform Act of 1986, there was a strong incentive for corporate income to be held within companies rather than distributed to rich shareholders; after the reform, this incentive changed. Many companies were reorganized so that company income was taxed only once, as personal income to shareholders. Income at the top might be more cyclical today in part because some income was formerly sheltered within the corporation. Under current law, all the cyclicity in that corporate income will be reflected directly in the owner's personal income tax.

Robert Gordon noted that the share of total executive income taking the form of stock options rose dramatically during the 1990s. Also, the two big recent episodes of stock market volatility were synchronous with the business cycle, making it difficult to distinguish between its role and that of the stock market cycle. Gordon suggested that quite a bit of the increase in top income cyclicity might be due to the increased dependence of very top income earners on stock options. He proposed as a possible explanation a general increase in the market power of managers, which could help to explain both the increase in inequality and the increased downward responsiveness of labor hours to the decline in output, as has occurred in the last two recessions. The question was how much of this shift in market power is due to growing strength at the top versus growing weakness at the bottom. It could be that the eroding market power of workers at the bottom created a vacuum, and the top moved in.

Karen Dynan was interested in how the authors' findings related to the so-called Great Moderation. If top income groups accounted for a greater share of total income in this period, and at the same time were experiencing greater income volatility, how does that square with the stylized fact of greater moderation in the macroeconomy, and how might that inform the understanding of that period?

Justin Wolfers found the paper's analysis to be extremely thorough across datasets, and the results as reliable as they could be given certain weaknesses in the data. He also noted that the main finding is not only well supported, but surprising. Two years ago, if someone had surveyed 100 labor economists and asked them whether rich people were more likely than others to get hurt by recessions, the majority, he believed, would have said no. Wolfers also suggested further testing the theory using data from the so-called Great Compression of the 1920s through the 1950s—a period that also saw a large shift in income inequality but in the opposite direction.

Robert Hall observed that the rational thing to do when one's lifetime resources change immediately and dramatically is to change one's

consumption immediately and dramatically. There is no reason to think that high-income households would act any differently in such circumstances. Following up on Gürkaynak's point about the risk-return relationship, Hall noted that there is also a lot of evidence that ordinary wages contain an insurance element, especially among longer-term workers, who are typically insulated from layoffs. Somebody has to stand on the other side of this insurance deal. To the extent that that somebody is the high-income shareholders of the same firms, this could explain the observed volatility of their incomes.

Bruce Meyer reported that a student of Anthony Atkinson had found that high-income shares of total income rise dramatically after financial crises. This result comes from 50 years of data from many countries. He wondered how much of what the paper found to be going on is about the timing of income changes in response to financial shocks.

Laurence Ball thought the facts reported in the paper were basically right, and he agreed with Perry that hedge funds must be a big part of the story. But he also wondered how precisely hedge funds might be driving the observed change in cyclicalities. Possible explanations included regulatory changes, changes in social norms, changes in tax rates, or some combination of those elements.