

Who Bears Aggregate Fluctuations and How?

By JONATHAN A. PARKER AND ANNETTE VISSING-JØRGENSEN*

As we write in late December 2008, the economy is mired in a year-long recession, the US stock market is down 40 percent for the year, and real per capita consumption of nondurables and services has fallen roughly 1 percent over the last year. The welfare costs of these declines depend significantly on their allocation across households.

In this paper we study differences in exposure to aggregate fluctuations across households, focusing on high-consumption and high-income households. In doing so, we bring together two somewhat disparate literatures. One line of research has documented increases in income and consumption inequality over the past 25 years (e.g., David Cutler and Lawrence Katz 1991; Dirk Krueger and Fabrizio Perri 2003; Thomas Piketty and Emmanuel Saez 2003; Georgio E. Primiceri and Thijs van Rens, forthcoming). This work, typically framed within a basic permanent income model, focuses on the extent to which income shocks are insured and pays less attention to the extent to which insurance of aggregate shocks differs across households. This contrasts with the literature in asset pricing that has documented that equity risk is born disproportionately by households with large stock market wealth (e.g., N. Gregory Mankiw and Stephen P. Zeldes 1991; Parker 2001; Christopher Malloy, Tobias J. Moskowitz, and Vissing-Jørgensen, forthcoming). This work studies differences in the covariation of consumption growth only with equity returns, and not with aggregate fluctuations more generally.¹ For our purposes, this is significant since

the share of aggregate income that comes from labor is roughly double the share coming from capital. As with the literature on consumption inequality, this research is limited by underrepresentation of households with very high consumption in standard consumption datasets.

We have five main results. First, the consumption growth of high-consumption households is significantly more exposed to aggregate fluctuations than that of the typical household in the Consumer Expenditure (CEX) Survey. The exposure to aggregate consumption growth of the consumption growth of households in the top 10 percent of the consumption distribution in the CEX is about five times that of the average household. Second, this pattern predicts a significant decline in consumption inequality over the past year. With real aggregate per capita consumption growth about 3 percentage points less than its historical mean (of 2 percent) during the past year, the ratio of consumption of the top 20 percent to the bottom 80 percent is expected to fall by about 9 percentage points, relative to its evolution under trend growth.

Third, we provide evidence on the channels that lead to higher exposure for high-income households using income data from the tax return data set assembled by Piketty and Saez (2003). In the period covered by the CEX, we show that a higher exposure of the *income* of rich households to aggregate consumption and income fluctuations is a likely contributor to their higher *consumption* exposure. High-income households (top 1 percent) earn more than half of their non-capital gains income from wage income, and their wage income is far more exposed to aggregate fluctuations than that of lower-income households. Fourth, we find even higher income exposure to aggregate fluctuations for very high-income households (top 0.01 percent) than for high-income

*Parker: Finance Department, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208-2001 (e-mail: Jonathan-Parker@Kellogg.Northwestern.edu); Vissing-Jørgensen: Finance Department, Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, IL 60208-2001 (e-mail: a-vissing@Kellogg.Northwestern.edu). We thank Francisco Gomes for helpful comments and Emmanuel Saez and Thomas Piketty for posting their IRS data.

¹ Malloy et al. (forthcoming) show higher sensitivity of the consumption growth of wealthy stockholders to both the stock market and to aggregate consumption growth.

Here, we instead sort households by income and consumption levels, and analyze aggregate income and consumption fluctuations, exposure by type of income, and inequality.

households, suggesting that our consumption estimates may understate the exposure of high-consumption households thought to be omitted from the CEX.

Finally, we find a striking change in the exposure of the incomes of high-income households: prior to the early 1980s, the incomes of high-income households were *not* more exposed to aggregate fluctuations. Thus, while high-income households currently bear an inordinately large share of aggregate fluctuations, this is a recent occurrence.

I. The Exposure of the Consumption of High-Consumption Households to Aggregate Fluctuations and Its Implications for Inequality

We use the CEX surveys from 1982 to 2004 to study whether the consumption of high-consumption households is more exposed to aggregate fluctuations. We construct average log consumption growth rates for each group and period using expenditures on nondurable goods and a subset of services deflated by the CPI for nondurables.² To reduce the impact of households exiting and entering the CEX, we construct average consumption growth for a group by taking a CEX-weighted average across household-level, quarterly log growth rates within a group. A household's percentile in the (weighted) cross-sectional distribution of consumption is defined based on its expenditures in the first period of the change.³ We sum four quarterly average log changes to obtain annual growth rates (available at monthly frequency), defined as the log growth rates from a three-month period to the same three-month period the following year.

² The expenditure definition and CEX sample period follow Malloy et al. (forthcoming). Since we cannot adjust for family size or changes in family size in the income data used in the section below on income, we do not adjust consumption for family size effects in the CEX. All CEX levels and growth rates are thus to be interpreted as per household.

³ Sorting on initial consumption leads to nonstandard measurement error when analyzing consumption growth rates of percentile groups. We show in an online Appendix (available at <http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.2.399>) that using log growth rates gives consistent estimates of the sensitivity of group consumption growth to aggregate consumption growth if the primitive measurement error is classical. For consistency across datasets we use log growth rates in all tables.

Panel A in Table 1 shows the extent of consumption inequality across percentile groups in the CEX. Panel B shows the results of regressing a group's average annual change in log consumption on month dummies and the contemporaneous change in log aggregate real per capita consumption of nondurable goods and services, deflated as described in Parker (2001) and constructed from monthly data to match the timing of the CEX series. The growth rate of those in the top 10 percent of households in the distribution of consumption typically changes by about 5 percentage points when the growth rate of aggregate consumption per capita changes by 1 percent, while the change in the growth rate of the bottom 80 percent is only about a half percent.⁴ Exposure to aggregated CEX consumption is somewhat less concentrated on high-consumption households than National Income and Product Accounts (NIPA) consumption (panel D). The top 5 percent of households are estimated to be about 4.5 (=2.51/0.56) times more exposed to changes in CEX consumption than those in the bottom 80 percent.

If the sensitivities in panel D were all equal to one, the share of aggregate consumption fluctuations borne by a group would be its share of initial consumption: the fraction the group constitutes of the population times the average ratio of group consumption to average consumption (e.g., $0.10 \times 2.15 = 0.215$ for the top 10 percent). We estimate the actual fraction borne by a group by regressing (Change in real group consumption per household) \times (Group share of population) / (Lagged aggregate real consumption per household) on the growth rate in aggregate real consumption per household. Across subgroups of households, the numerators sum to the total real dollar change in consumption per household, so the regression coefficients sum to one. As shown in panel E, the fraction borne by the top 10 percent of household is 45 percent—driven by both higher average consumption and higher sensitivity to aggregate shocks.

⁴ At a quarterly frequency there is a smaller difference across groups; at lower frequencies, there is a larger difference across groups. The results are similar if the left-hand side is calculated from the change in the log of mean group consumption (as opposed to the mean of the log changes) among households present in both periods.

TABLE 1—EXPOSURE OF CONSUMPTION GROWTH TO AGGREGATE CONSUMPTION GROWTH

All CEX households	Bottom 80 percent	Top 20 percent	Top 10 percent	Top 5 percent
<i>Panel A: Average consumption to total average consumption</i>				
1	0.79	1.83	2.15	2.52
<i>Panel B: Sensitivity to NIPA consumption growth</i>				
1.10 [3.66]	0.54 [1.47]	3.36 [5.62]	5.29 [4.02]	5.33 [3.55]
<i>Panel C: Biased sensitivity of unbalanced panel to NIPA consumption growth</i>				
1.33 [3.28]	1.15 [3.42]	1.59 [2.54]	1.86 [2.38]	2.25 [2.21]
<i>Panel D: Sensitivity to total CEX consumption growth</i>				
1	0.56 [7.92]	1.69 [13.50]	2.01 [8.29]	2.51 [6.02]
<i>Panel E: Fraction of total CEX fluctuations borne by group</i>				
1	0.39 [7.40]	0.62 [11.53]	0.45 [7.97]	0.34 [6.09]

Notes: *t*-statistics based on Newey-West standard errors in brackets. All regressions use annual changes from the same three-month period one year ago.

Given the large exposure of high-income and high-consumption households to movements in aggregate income or consumption, we expect recent poor aggregate economic performance to reduce inequality. To match recent figures, we consider the effect of a decline in the growth rate in aggregate real per capita consumption of nondurables and services from 2 percent to minus 1 percent, i.e., a 3 percentage point growth rate decline (for one year). Based on panel B in Table 1, a 3 percentage point decline in the growth rate of aggregate real per capita consumption of nondurables and services will lead to a decline of about 1.5 percentage points in the growth rate of real per household consumption for those in the bottom 80 percent, of around 10 percentage points for those in the top 20 percent, and of around 16 percentage points for those in the top 10 percent—all relative to trend. Thus, the consumption of the typical household in the top 20 percent (top 10 percent) will decline by about 9 percent (15 percent) relative to the consumption of the average household in the bottom 80 percent.

We conclude that in the CEX data, the consumption of high-consumption households is more exposed to aggregate booms and busts than that of the typical household. But we also see the exposure rising significantly only quite high in the distribution, and, according to

Bureau of Labor Statistics (BLS) statisticians, households in the top 5 to 10 percent of the distribution of expenditures are underrepresented in the CEX. What we denote the top 10 percent may thus represent the eighty-fifth to the ninety-fifth percentile. To investigate this further, and to better understand the channels behind higher consumption exposure, we turn to information on high income households.

II. The Exposure of Incomes of High-Income Households to Aggregate Fluctuations

We use the Piketty and Saez (2003) data on taxable income by type and income level, which are based on large samples of very high-income households. We initially study the period from 1982 (to match our CEX data) to 2006. For our purposes, these data have two disadvantages: they do not track the same households over time and they do not have information on low-income households. We subsequently show that the first disadvantage likely biases down the extent to which exposure to aggregate fluctuations differs by income, and we account for the second by using national totals from NIPA Table 2.1 (as do Piketty and Saez 2003).

For two reasons, we focus on income excluding capital gains. First, the IRS data measure

TABLE 2—EXPOSURE OF INCOME GROWTH BY INCOME PERCENTILE, 1982–2006

Type of income	All tax units	Top 10 percent	Top 1 percent	Top 0.1 percent	Top 0.01 percent
<i>Panel A: Average income in group to average for all tax units</i>					
Total	1	3.2	10.7	41.3	157.9
<i>Panel B: Average percent of income from source</i>					
Wage	68.0	77.5	60.7	49.0	40.3
Nonwage	32.0	22.5	39.3	51.0	59.7
<i>Panel C: Aggregate consumption growth beta</i>					
Total	1.98 [5.14]	2.60 [3.32]	4.69 [2.62]	7.30 [2.64]	8.62 [2.59]
Wage	1.86 [6.08]	2.53 [4.08]	5.44 [3.08]	9.86 [2.55]	15.22 [2.71]
Nonwage	2.25 [3.09]	2.03 [2.30]	2.80 [1.44]	3.51 [1.42]	2.71 [0.87]
<i>Panel D: Aggregate total income growth beta</i>					
Total	1	1.26 [5.34]	2.22 [3.70]	3.23 [3.36]	3.71 [3.16]
Wage	0.82 [12.67]	1.07 [5.28]	2.28 [3.69]	4.37 [3.24]	5.96 [2.90]
Nonwage	1.38 [10.04]	1.92 [3.01]	2.11 [2.48]	1.95 [1.87]	1.52 [1.34]
<i>Panel E: Fraction of aggregate income change borne by group</i>					
Total	100	40.30 [5.34]	23.90 [3.86]	13.40 [3.52]	5.80 [3.17]

Notes: *t*-statistics in brackets. Percentiles refer to tax units. Total income excludes capital gains and transfers and does not subtract taxes. Panel E is similar to Table 1 panel E, except for being based on income rather than consumption.

only realized capital gains and the timing of capital gains realizations is an endogenous choice of a household. Second, since we study only the cash flow from human capital (wages), for comparability we study only the cash flow from nonhuman capital (dividends, interest, rental income and proprietors' income). This said, income including true capital gains on human and financial capital would of course be preferable.

Panel A of Table 2 shows the extent of income inequality in the data. Panel B shows that wage income is still a substantial source of income for high-income households. Even the top 0.1 percent have nearly half their income from wages, with the remainder constituted by 28 percent from proprietors' income (roughly triple the population average), 10 percent from dividends (roughly double the population average), and 14 percent from interest and rental income.

Turning to the stochastic properties of income, we regress the log growth rate in real income per

tax unit onto the log growth rate of either aggregate real consumption per tax unit or aggregate real income per tax unit (calculated from NIPA consumption data, NIPA total income across the five subcategories of income, and the IRS number of tax units). Panels C and D show that the incomes of very high-income groups have dramatically larger sensitivities to aggregate growth rates: the incomes of the top 0.1 percent of tax units have sensitivities of about 7 to aggregate consumption growth and about 3 to aggregate income growth.⁵ Figure 1 displays this strikingly different cyclicity across groups.

What drives these differences in exposure? One might expect that the labor income of high-income households is more insulated from

⁵Results are similar if we omit the two years after the 1986 tax reform which have unusually large income growth.

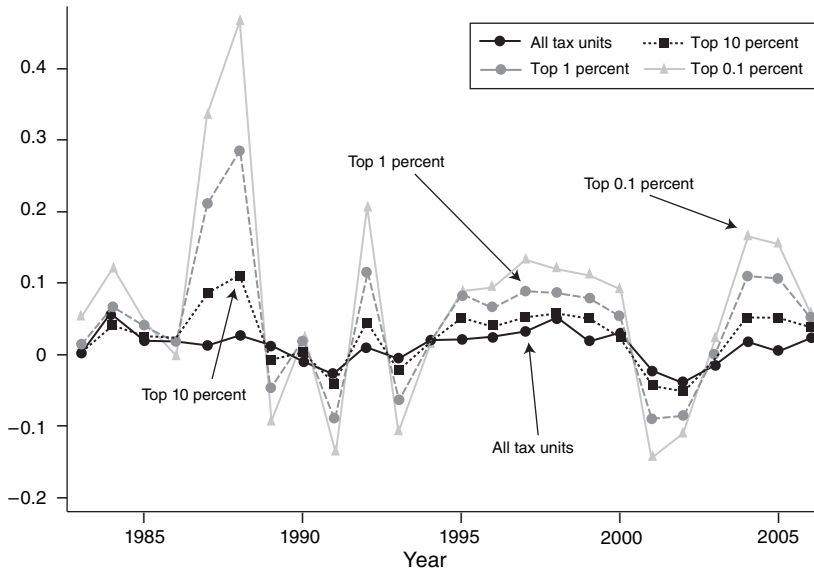


FIGURE 1. GROWTH RATES OF NON-CAPITAL GAIN INCOME BY GROUP, 1982–2006

aggregate fluctuations than that of low-income household, but in fact it is more exposed. The high sensitivities for the rich are mostly due to the higher sensitivity of their wage income to changes in aggregate consumption and income.

A concern in our use of this income data is that the percentile group into which a household is allocated in a given year is based on income in that year. Thus, the incomes used to calculate a given growth rate do not represent the same households in each period (unlike our analysis in the CEX). If high-income households are more exposed to aggregate fluctuations, some of them will fall into lower percentile groups when aggregates fall and will rise up the distribution when aggregates rise. This composition bias actually biases *down* the relative exposure of high-income groups. A high-income group's measured decline in bad times is reduced by initially lower-income, less-exposed households entering the high-income group, and the group's measured rise in good times is reduced by these households leaving the high-income group. The converse occurs for lower-income groups, biasing upward their measured exposure to fluctuations.

To judge the empirical significance of this bias, we introduce it artificially into our analysis of consumption in the CEX. Panel C of Table 1 reports the exposure to aggregate consumption

of the consumption of a changing population of households in the CEX constructed analogously to the tax data. Relative to the unbiased (fixed-group) sensitivities in panel B, panel C shows greater exposure to aggregate growth rates for the low-consumption group and lower exposure for the high-consumption groups. Therefore, the estimates of the exposure of high-income households to aggregate fluctuations in Table 2 are likely downward biased.⁶

The finding in Table 2 that exposure of *income* to aggregate consumption (and income) fluctuations increases dramatically from the top 10 to the top 1 or top 0.01 percent suggests that the exposure of *consumption* to aggregate consumption fluctuations of very rich households is likely to be larger than documented in Table 1, given underrepresentation of the very rich in the CEX. However, we cannot conclude this with certainty since differential consumption exposure across groups may be driven not only by differential income exposure but also by differences in capital gains exposure or in the relation between wealth and consumption across groups.

⁶ The different approaches to constructing group level growth rates also affect average growth rates for the groups.

TABLE 3—EXPOSURE OF INCOME GROWTH BY INCOME PERCENTILE, 1929–1982

Type of income	All tax units	Top 10 percent	Top 1 percent	Top 0.1 percent	Top 0.01 percent
<i>Panel A: Average income in group to average for all tax units</i>					
Total	1	2.9	8.9	27.6	83.6
<i>Panel B: Average percent of income from source</i>					
Wage	69.8	66.0	42.0	30.1	18.0
Nonwage	30.2	34.0	58.0	69.9	82.0
<i>Panel C: Aggregate consumption growth beta</i>					
Total	1.62 [7.22]	1.55 [10.87]	2.01 [7.25]	1.85 [4.74]	1.98 [4.31]
Wage	1.49 [5.66]	0.71 [4.55]	0.65 [2.97]	0.42 [1.30]	0.32 [0.45]
Nonwage	1.83 [8.65]	2.61 [8.72]	2.86 [7.42]	2.29 [4.76]	2.22 [4.38]
<i>Panel D: Aggregate total income growth beta</i>					
Total	1	0.62 [8.45]	0.84 [6.60]	0.82 [4.83]	0.75 [3.54]
Wage	1.03 [31.10]	0.33 [4.91]	0.07 [0.65]	−0.04 [−0.29]	−0.37 [−1.20]
Nonwage	0.91 [12.71]	1.02 [6.77]	1.30 [8.02]	1.10 [5.45]	0.93 [4.09]

III. Changes in the Exposure of the Incomes of High-Income Households

The larger exposure of the incomes of high-income households is only a recent phenomenon: prior to the last 25 years, the incomes of high-income households were *not* more exposed to aggregate fluctuations. Table 3 shows a subset of the statistics from Table 2, but over the period from 1929 to the date of availability of the CEX data sample (with roughly similar conclusions if we focus on the postwar period up to 1982).

High-income households have less of their income from wages and more from dividends, relative to the more recent period, suggesting higher exposure of the very high-income households to stock market fluctuations prior to 1982. More importantly, panel C and D show that in the earlier period the incomes of high-income households have about the same sensitivity to aggregate consumption as the income of all households, and a lower sensitivity to aggregate income. This is due mainly to lower exposure of the wage income of the rich to changes in aggregate fluctuations in

the earlier period and to lower exposure of nonwage income (disproportionately earned by the rich) to changes in aggregate income in the earlier period.

We find this fact—high-income households becoming more exposed to aggregate changes in income and consumption—tantalizing. It begs further study both in terms of measurement and in terms of understanding the underlying labor and capital market mechanisms.

REFERENCES

- Cutler, David M., and Lawrence F. Katz. 1991. "Macroeconomic Performance and the Disadvantaged." *Brookings Papers on Economic Activity*, 2: 1–61.
- Krueger, Dirk, and Fabrizio Perri. 2003. "On the Welfare Consequences of the Increase in Inequality in the United States." In *NBER Macroeconomics Annual*, ed. Mark Gertler and Kenneth Rogoff, 83–121. Cambridge, MA: MIT Press.
- Malloy, Christopher, Tobias J. Moskowitz, and Annette Vissing-Jørgensen. Forthcoming. "Long-Run Stockholder Consumption Risk

and Asset Returns.” *Journal of Finance*.

Mankiw, N. Gregory, and Stephen P. Zeldes. 1991. “The Consumption of Stockholders and Non-stockholders.” *Journal of Financial Economics*, 29(1): 97–112.

Parker, Jonathan A. 2001. “The Consumption Risk of the Stock Market.” *Brookings Papers on Economic Activity*, 2: 279–333.

Piketty, Thomas, and Emmanuel Saez. 2003. “Income Inequality in the United States, 1913-1998.” *Quarterly Journal of Economics*, 118(1): 1–39.

Primiceri, Giorgio E., and Thijs van Rens. Forthcoming. “Heterogeneous Life-Cycle Profiles, Income Risk and Consumption Inequality.” *Journal of Monetary Economics*.