



Inequality in 3-D: Income, Consumption, and Wealth

Jonathan Fisher (Stanford University, United States), David S. Johnson (University of Michigan, United States), Timothy M. Smeeding (University of Wisconsin, United States), and Jeffrey P. Thompson (United States Federal Reserve Board of Governors)

Paper prepared for the 34th IARIW General Conference

Dresden, Germany, August 21-27, 2016

Session 6D: Household Wealth II

Time: Thursday, August 25, 2016 [Afternoon]

August 15, 2016
IARIW Version

Inequality in 3-D: Income, Consumption, and Wealth

Jonathan Fisher (Stanford University)
David Johnson (University of Michigan)
Timothy Smeeding (University of Wisconsin)
Jeffrey Thompson (Federal Reserve Board of Governors)

I. Introduction

Economic inequality is multi-dimensional. Income, consumption, and net wealth, independently and jointly, inform the perception and reality of inequality. Yet most studies of inequality limit analysis to one dimension. Even those using more than one measure present the measures separately and often from different datasets. Because the three measures are not perfectly correlated, our understanding of inequality deepens, broadens, and becomes more nuanced by studying inequality in two and three dimensions. We are the first to study inequality in three conjoint dimensions, and to do so we use income, consumption, and net wealth from the 1989-2013 Surveys of Consumer Finance (SCF).

We begin with a more conventional analysis by presenting inequality results for the three measures individually, demonstrating that our sample replicates the one-dimensional understanding of inequality for each measure. Moving beyond the conventional analysis, we present the conjoint distribution of income, consumption, and wealth using share analysis, transition matrices, and summary measures of cross-sectional mobility. Our focus is on two questions. What does inequality in two and three dimensions look like? And, has inequality in multiple dimensions increased by less, by more, or by about the same as inequality in one dimension?

A further contribution moves beyond focusing on the top of the distribution. Much of the recent research focuses on the share of the resource measure held by the top percentiles, motivated in large part by the seminal work of Piketty and Saez (2003). While the top is important and is the driver of much of the increase in inequality, multidimensional inequality may look different at the bottom and middle of the distribution. Here we show results for the top of the distribution and for the complete distribution, again allowing for a more complete understanding of inequality.

The results from this paper are not only relevant to the literature on the level and trend in inequality. Our uni-dimensional inequality results can help calibrate macroeconomic models such as the ones found in Krusell and Smith (1998), Castenada, Diaz-Gimenez, and Rios-Rull (2003), Benhabib, Bisin, and Zhu (2011),

Kaplan and Violante (2014), and, Krueger, Mitman, and Perri (2016). These types of models can also be extended to include the joint distributions we provide. Until our findings, macroeconomic research did not have results to calibrate their models to the joint distribution. These joint results will allow future macroeconomic models to better understand the underlying dynamics and heterogeneity across households. For instance, we build on the results in Kaplan, Violante, and Weidner (2014) by identifying that households are more than just low wealth and high wealth households. They find the wealthy hand-to-mouth households that they identify as having high illiquid wealth but little liquid savings, such as elders living in costly cities in a house they own outright. In general, understanding the distribution of wealth income *and* consumption greatly expand our analytical power, especially when one component that is not highly covariant with another is left out of the analyses. For instance, the relationship between income and consumption is often studied, but without knowledge of the way wealth and wealth change affects consumption (Fisher, Johnson, and Smeeding, 2014).

The existing inequality literature typically studies one measure at a time. For example, Piketty and Saez (2003) use tax records and Burkhauser et al. (2012) use survey data, but both study income inequality alone using different datasets, with much discussion on how to combine the top income end of the tax files with the more generally available household survey literature. Those studying consumption inequality often compare the trend in consumption inequality to the trend in income inequality, but the focus is on the univariate distributions and not the joint distribution (e.g., Attanasio and Pistaferri, 2014; Aguiar and Bils, 2015; Fisher, Johnson, and Smeeding, 2015; Meyer and Sullivan, 2016). Similarly, wealth inequality is often studied alone or is compared to income inequality, but rarely do we consider the effects of wealth accumulation on consumption (e.g., Wolff, 2014; Saez and Zucman, 2015).

A few wealth inequality studies present information on the joint distribution of income and wealth, such as Saez and Zucman (2015) who report the share of income held by the top 1% of wealth holders using tax data. While an important step, Saez and Zucman (2015) do not have data on consumption, only have pre-tax pre-transfer taxable income, use tax-filing units instead of households, and focus on the very top for the joint distribution. Jäntti, Sierminska, and Smeeding (2008) focus on the middle and on the bottom of the distribution by studying the net worth of the low and middle-income populations in a cross-national context using the same datasets. Smeeding and Thompson (2011) and Armour et al. (2014) capitalize wealth holdings into income to study the level and trend in income inequality, including in their income definition the imputed flows of income from capital, whether realized or not, but they do not account for the effects of wealth accumulation on well-being. Hence, an important component of multi-dimensional inequality, the stock of wealth is more than just an annuitized income flow, as it represents the power to consume and/or the power to transfer across generations, regardless if the imputed or real value of wealth changes on incomes per se.

Heathcote, Perri, and Violante (2010) and Krueger, Mittman, and Perri (2016) come closest to our approach. Heathcote et al. (2010) present income, consumption, and wealth inequality together, but they use a different survey for each measure, which necessitates focusing on the univariate distributions. Krueger, Mittman, and Perri (2016) use the Panel Study of Income Dynamics (PSID) and present the shares of income and consumption by wealth quintile, with the goal of building a real business cycle model to help explain how the cross-sectional distribution of wealth shapes business cycle dynamics, similar to our recent PSID work (Fisher, et al, 2016a) Krueger et al. (2016) also make an important step in understanding two-dimensional inequality by showing the share of income and the share of consumption held by each wealth quintile, but we go further in exploring multi-dimensional inequality and moreover, we are able to capture the top of the distributions on income and wealth by using the Survey of Consumer Finances, as that data is not available in the PSID.

Auerbach, Kotlikoff, and Koehler (2016) use a micro-simulation program with inputs from the 2013 distributions of earnings, income, and wealth from the SCF to project lifetime resources and lifetime spending for each household. They use these measures to compare the distribution of projected lifetime resources and projected lifetime spending in order to understand the redistribution built into fiscal policy. They find that the top 1 percent own 24 percent of net worth but only 10 percent of remaining expected lifetime spending, a ratio of 2.4 to 1. Their wealth share to lifetime spending share is comparable to our wealth-to-consumption share for the top 5% in 2013, which is 2.8 to 1.¹ In comparison to our work here, Auerbach et al. (2016) impute consumption to the SCF in a different manner but arrive at the same basic conclusion that consumption inequality is considerably smaller than wealth inequality. We instead impute current consumption rather than lifetime consumption, a much less heroic task, and we include trends in multi-dimensional inequality rather than just levels.

The common thread through all of this inequality research is increasing economic inequality, whether using income, consumption, or wealth. Given the consensus of increasing inequality, the necessity of studying the three measures begs for attention. Income, consumption, and wealth holdings and positions are not perfectly correlated. The life-cycle pattern of the measures best demonstrates this imperfect correlation. Younger adults often have consumption exceeding income along with low or negative wealth, while older adults often have relatively high consumption and high wealth but low income (Fisher, Johnson, Smeeding, and Thompson, 2015). The *Report by the Commission on the Measurement of Economic Performance and Social Progress* (Stiglitz, Sen, and Fitoussi, 2009) also argues for the joint study of the three measures, stating that, “the most pertinent measures of the distribution of material living standards are probably based on *jointly* considering the income, consumption, and wealth position of households or individuals.” Blundell (2014), in

¹ The top 5% of the wealth distribution hold 62.6% of the wealth and the top 5% of the consumption distribution have 22.4% of consumption.

his address to the Royal Statistical Society, also highlights the importance of all three measures, stating that: “...the results of the research presented here provide a strong motivation for collecting consumption data, along with asset and earnings data...”

The only way to understand the conjoint distribution is to have income, consumption, and wealth in the same survey. The Panel Study of Income Dynamics (PSID) asks about income, consumption, and wealth in every wave since 1999. Because the PSID contains all three measures, it represents a ready-made source, as evidenced by Krueger et al. (2016) and Fisher, Johnson, Latner, Smeeding and Thompson (2016a). The PSID, however, misses the top of these distributions, which is of particular interest because the top drives much of the increase in inequality over the last four decades (e.g., Pfeffer, Schoeni, Kennickell and Andreski, 2014). Another drawback of the PSID is that it only has all three measures since 1999. Before 1999, the PSID asked a limited set of consumption questions and did not ask wealth in every wave.²

The Survey of Consumer Finances (SCF) captures the top of the income and wealth distributions better than any other survey extant and has contained a consistent sample and consistent measures since 1989 (Bricker, et. al. , 2016; Bricker, et al, 2014; Dettling, et al, 2015) . But, the SCF has an incomplete measure of consumption, including just spending on food, mortgage or rent, and the stock of vehicles. We impute the residual consumption components to the SCF using the Consumer Expenditure (CE) Survey. This represents the first time consumption is imputed to the SCF to study inequality.³

Our results depend crucially on the imputation. One concern is the quality of the source data. The CE Survey reports lower aggregate expenditures than those reported in the Personal Consumption Expenditures (PCE) (Bee, Meyer, and Sullivan, 2014). One source of the under-reporting is that the CE Survey receives a lower response rate from high-income zip codes (Sabelhaus et al., 2014). The SCF, on the other hand, oversamples from the high-income zip codes that the CE misses, which will help better match PCE totals.

The SCF oversampling high-income zip codes creates a separate issue – the CE may lack support to impute consumption for the highest income SCF households. We must assume the estimated functional form between the share of income consumed and income observed in the bottom 95% of the distribution holds for the highest income SCF households. We test the quality of our imputation by comparing our imputed SCF aggregates to the PCE and find that the imputed SCF consumption

² Fisher and Johnson (2006), Fisher, et. al.(2016a), and Blundell, Pistaferri, and Preston (2012), among others, impute total consumption to the PSID using the CE Survey, but this only solves one of our problems with the PSID. The PSID still misses the top of the distributions and only asks wealth in 1984, 1989, 1994, and from 1999 on.

³ Bostic, Gabriel, and Painter (2009) impute consumption to the SCF to study housing and financial wealth effects.

measure matches the level and trend in the PCE from 1989-2013. We also compare the level and trend in inequality to the CE and find higher consumption inequality in the SCF, as expected, and we find higher growth in consumption inequality in the SCF.⁴ Lastly, Fisher et al. (2016a) show that the SCF and PSID joint distributions look similar, at least for the part of the distributions that are available in both data sets.

To preview our results, we find that there is a large overlap between the top and bottom of the two-dimensional distributions, identical to the twin peaks phenomenon seen in the mobility literature (e.g., Fisher and Johnson, 2006). In three dimensions, one in three households that are in the bottom wealth quintile are also in the bottom quintile of income and consumption. At the other extreme, one in two households that are in the top wealth quintile are also in the top quintile of income and consumption. Finally, the results indicate that inequality in two dimensions is growing faster than inequality in one dimension.

The results improve our understanding of inequality in the United States since 1989 by exploring the joint relationship between various measures of economic inequality. The results indicate that the correlation between the three measures is high but not perfect. Therefore, the individuals who comprise the top in any one measure do not necessarily comprise the top in another. While there is stickiness at both the top and the bottom of individual measures, fluidity also exists, especially as the ranking of one measure relates to the ranking of another.

The picture of inequality drawn here both aligns with previous research in that it is rising in all three dimensions, but the picture provided here also improves the clarity by incorporating the relationship between various key measures of economic well-being that comprise inequality. In so doing our analysis points to wealth as the most important source of inequality since it can be deployed to stabilize both consumption and income, while offering self-insurance and the option of transfers to younger generations to ensure their future economic position (see also Fisher, et al., 2016).

II. Inequality and the Budget Constraint

To frame our understanding of inequality in three dimensions, we start with the budget constraint.

$$Y_t = C_t + \Delta W_t \tag{1a}$$

$$W_t = W_{t-1}*(1+r_t) + \Delta W_t \tag{1b}$$

⁴ We have also compared results of the joint distribution of our three measures between the PSID and SCF and find similar patterns between the two surveys (see also Fisher, Johnson, Latner, Smeeding, and Thompson, 2016a).

Only two of our three dimensions, income and consumption flows, appear directly in the budget constraint in equation 1a. The change in wealth includes the interest earned or lost on last period's wealth, other changes in the value of assets, and any saving or dissaving. Wealth, the stock represented by equation 1b, is therefore a function of last period's wealth plus interest earned, minus interest paid, and any saving or dissaving.⁵

We start from the observation that income inequality is increasing, and we want to understand how this increase in income inequality could affect consumption inequality and wealth inequality. To frame a basic understanding, assume a world with no income inequality – everyone has the same amount of income in t , and everyone makes the same consumption and savings decisions such that there is no consumption or wealth inequality. Now suppose one person's income doubles while everyone else's income stays the same in year $t+1$. The person with double income must increase consumption or savings, meaning inequality must increase in consumption or wealth, but it is not guaranteed that inequality must increase in consumption and wealth. A priori, a rise in income inequality does not have to lead to an increase in consumption inequality and consumption inequality.

Blundell, Pistaferri, and Preston (2008) present a formal model for how changes in income inequality translate to changes in consumption inequality. Real log income contains a permanent component and a mean-reverting transitory component. The change in log unpredictable consumption contains three terms: the effect of a permanent change in income with a corresponding marginal propensity to consume (MPC); the effect of a transitory change in income with its MPC; and, a random component that represents innovations to consumption independent of changes in income.

If households can completely self-insure against income shocks, the MPC out of permanent shocks and the MPC out of transitory shocks is zero, suggesting that an increase in income inequality generated by changes in permanent or transitory shocks does not affect consumption inequality. Instead wealth inequality increases. On the other extreme if households have zero ability to self-insure and the MPCs instead equal zero, then an increase in income inequality completely passes through to consumption inequality, with no change in wealth inequality. Anything between the two extreme MPCs leads to an increase in consumption inequality and an increase in wealth inequality when income inequality increases.

Some research on consumption inequality finds that consumption inequality increased much less than income inequality, arguing that households were

⁵ Attanasio and Pistaferri (2016) argue for another dimension to inequality, leisure. They view inequality through the utility function, rather than the budget constraint, and argue that inequality in consumption and inequality in leisure are most important. They still look at inequality one dimension at a time rather than the bivariate distribution of consumption and inequality. See also Aguiar and Hurst (2007) for a study on leisure inequality.

experiencing more transitory income shocks, which has an empirically lower marginal propensity to consume than permanent shocks, and these transitory shocks allowed households to smooth consumption (e.g., Krueger and Perri, 2006; Blundell, Pistaferri, and Preston, 2008; and, Meyer and Sullivan, 2016). If income inequality is increasing because of larger, randomly distributed transitory income shocks, then neither consumption inequality nor wealth inequality need increase even as income inequality increases. Permanent income has not changed so households do not change consumption in the face of the transitory shocks. The positive transitory shock is saved, and wealth is drawn down in the face of a negative transitory shock, leaving overall wealth inequality (relatively) unchanged.

More recent research finds that consumption inequality increased by about the same amount as income inequality (Attanasio and Pistaferri, 2014; Aguiar and Bils, 2015; Fisher, Johnson, and Smeeding, 2015). In the model of Blundell, Pistaferri, and Preston (2008), the observation that income inequality and consumption inequality increased by about the same amount would indicate that households are sensitive to transitory shocks and these reactions depend on the level of wealth, as low wealth households cannot adjust to shocks. Fisher, Johnson, Latner, Smeeding, and Thompson (2016a, 2016b) use the PSID and show that the marginal propensity to consume out of predictable income shocks is higher for low wealth households.

Another possible scenario is that wealth inequality could increase independent of a change in income inequality. Fagerang, Guiso, Malacrino, and Pistaferri, (2016) find that returns to assets vary substantially across households. If high wealth households receive a higher rate of return than low wealth households, wealth inequality would increase with no change in income inequality. As those high wealth households consume out of the extra wealth (e.g., Bostic, Gabriel, and Painter, 2009; Carroll, Otsuka, and Slacalek, 2011), consumption inequality would increase as well, independent of a change in income inequality. Wealth effects could help explain why consumption inequality and income inequality do not always move in tandem, and wealth effects could help explain why consumption inequality fell during the Great Recession while income inequality was flat or increased slightly. High wealth households may have experienced larger negative wealth shocks, which led high wealth households to cut back consumption more than lower wealth households (Fisher, Johnson, and Smeeding, 2014).

While theoretically it is not pre-determined that an increase in income inequality would lead to an increase in consumption inequality and/or an increase in wealth inequality, the empirical record suggests that the increase in income inequality led to an increase in consumption inequality and an increase in wealth inequality, although both consumption inequality and wealth inequality could have increased absent an increase in income inequality. Thus, we expect to see that inequality in two dimensions and inequality in three dimensions should also increase. We now turn to how we measure income, consumption, and wealth before turning to results showing inequality in one, two, and three dimensions.

III. Data and Imputation

In this section we describe the data we use, the definitions of income, consumption, and wealth we use, and our approach to imputing residual consumption to the SCF. Then we present results judging the quality of the imputation.

The Survey of Consumer Finances

We use data from the nine waves of the Federal Reserve Board's triennial Survey of Consumer Finances (SCF) conducted between 1989 and 2013. The survey collects detailed information about households' financial assets and liabilities, and has employed a consistent instrument and sample frame since 1989. To support estimates of a variety of financial characteristics as well as the overall distribution of wealth, the survey employs a dual-frame sample design. The national area-probability sample provides good coverage of widely spread characteristics. Because of the concentration of assets and non-random survey response by wealth, the SCF also employs a list sample that is developed from statistical records derived from tax returns under an agreement with the Internal Revenue Service's Statistics of Income Branch. The list sample consists of households with a high probability of having high net worth. The SCF joins the observations from the national sample and the list sample through weighting. The weighting design adjusts each sample separately, and final weights are adjusted so that the combined sample is nationally representative of the population and assets. These weights are used in all results.

The unit of analysis in the SCF is the "primary economic unit", which refers to a financially-dependent related group living together. This concept is distinct from either the household or family units employed by the Census Bureau, but is conceptually close to the consumer unit used in the Consumer Expenditure (CE) Survey data. We make no restrictions regarding the age of the household head. We do not restrict to prime-age working household heads. Our interest lies in economy-wide inequality, not inequality among a restricted age group.

We use after-tax income in all results, which includes capital gains income. Wealth, or net worth, is assets less liabilities. Assets include all financial assets (e.g., stocks, bonds, savings and retirement accounts) and non-financial assets (e.g., real estate, businesses, and vehicles). Liabilities include mortgages, credit card balances, student and automobile loans, and other miscellaneous forms of debt.

The SCF includes some consumption questions. Since 2004, the SCF has asked about food spending. Expenditures on automobiles are asked every wave, and the consumption value of automobiles is estimated based on the stock of automobiles. Renters are asked the dollar value of rent paid, and homeowners report payments for mortgage interest and principal along with property taxes. Because the SCF does not include full consumption, we impute the remaining components of consumption.

To impute consumption, we use the Consumer Expenditure (CE) Survey, which we describe next. Then we describe our imputation methodology.

We use data from the nine waves of the Federal Reserve Board's triennial Survey of Consumer Finances (SCF) conducted between 1989 and 2013. Several features of the SCF make it appropriate for addressing the questions of interest. The survey collects very detailed information about households' financial assets and liabilities, and has employed a consistent instrument and sample frame since 1989. As a survey of household finances and wealth, the SCF includes some assets that are broadly shared across the population (bank savings accounts) as well some that are held more narrowly and that are concentrated in the tails of the distribution (direct ownership of bonds).⁶ To support estimates of a variety of financial characteristics as well as the overall distribution of wealth, the survey employs a dual-frame sample design.

A national area-probability (AP) sample provides good coverage of widely spread characteristics. The AP sample selects household units with equal probability from primary sampling units that are selected through a multistage selection procedure, which includes stratification by a variety of characteristics, and selection proportional to their population. Because of the concentration of assets and non-random survey response by wealth, the SCF also employs a list sample which is developed from statistical records derived from tax returns under an agreement with SOI.⁷ (See Bricker et al (2014) for additional details on the SCF list sample.) This list sample consists of households with a high probability of having high net worth.⁸ The SCF joins the observations from the AP and list sample through weighting.⁹ The weighting design adjusts each sample separately using all the useful information that can be brought to bear in creating post-strata. The final weights are adjusted so that the combined sample is nationally representative of the population and assets.¹⁰

Recent work by Dettling et al, 2015 shows that the income and wealth data in the SCF match very well, both in level and in trend over time, with the National Income and Product Accounts (NIPA), the Financial Accounts of the United States, and other macro aggregates. Moreover, Bricker et al, 2016, show that the SCF captures high

⁶ It is important to emphasize that the publicly-released SCF data are cleaned of any identifying information about the responding family, including any geographic information about the family. The Federal Reserve does release summary information by Census region, though (see Bricker et al, 2014). The empirical analysis in this paper uses the internal SCF data in order to identify the household's state, MSA, and county of residence.

⁷ See Wilson and William J. Smith (1983) and Internal Revenue Service (1992) for a description of the SOI file. The file used for each survey largely contains data from tax returns filed for the tax year two years before the year the survey takes place.

⁸ For reasons related to cost control on the survey, the geographic distribution of the list sample is constrained to that of the area-probability sample.

⁹ The evolution of the SCF weighting design is summarized in Kennickell (2000).

¹⁰ The SCF weights were revised in 1998 to incorporate home ownership rates by race (Kennickel, 1999). Weights for earlier years were updated to reflect the revised methodology.

and rising top income and wealth shares, as in the administrative tax data used by others as well as the NIPA, on a conceptually comparable basis.

The Consumer Expenditure Survey

The CE Survey interviews households four times over one year, with the consumption questions covering the previous three months. We aggregate the four quarters of consumption to arrive at annual consumption. In the last interview, the CE asks about income over the previous twelve months, which provides us with a measure of consumption and income that cover the same twelve months. As the households who remain in the sample for four quarters are more likely to be higher income, homeowners, and older, we follow the procedures in Fisher and Johnson (2006) to re-weight the sample to represent the quarterly sample.

Consumption is total spending on all goods and services for current consumption (excluding life insurance, pensions, and cash contributions) less the purchase price of vehicles and the expenditures for home-ownership plus the service flow from vehicles, the reported rental equivalence of home-ownership and the value of federal government rental assistance. As with other research on consumption, we do not include goods obtained through barter, home production, or in-kind gifts from other households or organizations. In contrast to other research, however, our measure of consumption includes all other components of consumption, and includes education, health care expenses, and other durable goods. Excluding these components of consumption would break the explicit relationship between income, consumption, and wealth explicit in equations 1a and 1b.

The CE Survey began imputing income in 2004 but did not impute previous years. We use the imputed income variables from Fisher, Johnson, and Smeeding (2015). Another issue is that the CE income variables underwent a major redesign in 2013. The concern is that published mean income in the CE decreased by \$1,800 between 2012 and 2013 in nominal dollars (U.S. Bureau of Labor Statistics, 2015), while real mean income in the American Community Survey and the Current Population Survey increased between the same years. The CPS shows that mean income increased by \$1,400 between 2012 and 2013, which suggests that mean income in the CE is underreported at the mean in 2013 by approximately \$3,200. To address this concern, in the imputations we use the household's rank in the after-tax income distribution instead of the level of income, under the assumption that the under-reporting in 2013 is proportional to income.

Imputation Methodology

The CE data are appended to the SCF data to impute consumption from the CE to the SCF. Consumption is only imputed for the components that are not already asked in the SCF. Known consumption items account for approximately 15 percent of consumption. The remaining items are imputed.

The independent variables include a suite of demographic characteristics (e.g., age, race, and education) along with urban status and Census division. A quadratic in the

rank of income is included, along with indicators for whether the household reported: negative income, receiving government transfer income, receiving wage or salary income, receiving positive capital income (e.g., interest and dividends), and receiving negative capital income.¹¹ The consumption components available in the SCF are also included as independent variables.

A unique aspect of the SCF we utilize in our imputation is the series of questions regarding spending relative to income. Specifically the SCF asks: “over the past year, would you say that your spending exceeded your income, that it was about the same as your income, or that you spent less than your income?” Respondents are asked to exclude any investments made and to treat the net pay down of debt as spending less than income. For those that purchased a home or automobile in the previous year, they are asked to leave aside those expenses in answering the question.

We take the answer to this question as a noisy indicator of spending relative to income. Respondents may interpret what it means to spend “about the same as your income” differently. The income concept is not defined in the question.

The CE does not ask a similar question. Instead we create the variable in the CE to match the SCF weighted totals. We take the observed percentage of households in each group from the SCF and assign the same percentage of CE households. In practice, this means that any household reporting consumption being less than 80% of after-tax income is classified as spending less than income. Those that spend more than income are those reporting consumption at least 120% of income. Those spending about the same as income are those spending between 80% and 120% of income.

Separate imputation equations are estimated for the three spending-to-income groups. Imputed consumption in the SCF is not constrained to be within these specific bands. In this way, we are just using the spending-to-income groups as noisy indicators of consumption rather than as strict rules limiting our imputation. Mean and median spending relative to income is higher in the SCF group that reported spending is larger than income, but there is overlap in the distribution of spending-to-income among the three groups.

Note that we impute total consumption, not individual components as Aguiar and Bils (2015) do. Aguiar and Bils worry about differential underreporting. We correct for potential differential underreporting by estimating total consumption with separate income groups and by imputing consumption as a share of income.

Judging the Quality of the Imputation

¹¹ Because capital gains and losses are not reported in the CE, capital gains and losses are excluded from income in the imputations. However in the results presented below, income includes capital gains and losses.

To judge the quality of the imputation, there must be a proper benchmark. One simple comparison is back to the original CE data. Before comparing imputed consumption in the SCF to the CE, it is important to remember that there are expected to be differences between the two surveys in consumption. The SCF captures high-income households that are missed by the CE (see Sabelhaus et al., 2014; Bricker, et al, 2016).¹² To this point, Figure 1 shows the ratio of SCF before-tax income to CE before-tax income.¹³ The CE matches the SCF up to at least the 75th percentile. At the 95th percentile and above, the SCF income is higher than the CE income. Given that the CE misses the top of the income distribution, we expect that SCF consumption will also be higher at the top.

Figures 2A and 2B confirm the expectation that SCF consumption exceeds CE consumption at the top. These two figures compare mean consumption by income quintile in the SCF and CE, with Figure 1A presenting the components of consumption reported in the SCF and Figure 1B presenting the imputed components. Mean reported consumption in the SCF and CE overlap each other until about the 70th percentile of before-tax income. Mean imputed consumption in the SCF exceeds mean consumption in the CE starting around the 60th percentile. The difference is particularly large for the top 10 percent of the SCF income distribution. But these differences are expected as Figure 1 shows that before-tax income in the CE is considerably lower than before-tax income in the SCF at and above the 90th percentile in 2013.

Another test of the quality of consumption data is to compare aggregate consumption to the national aggregates reported in the Personal Consumption Expenditures (PCE). The CE is known to underestimate some PCE categories and the overall PCE (Bee, Meyer, and Sullivan, 2014; Garner et al., 2006). Figure 3 shows the difference between the CE and PCE aggregates, and the figure shows that the SCF aggregates close the gap. The SCF is still below the PCE, but the gap is smaller using the SCF than the CE, and the gap is 40-60 percent lower in the most recent years. That the SCF better matches the PCE suggests that the CE fails in matching the PCE in large part because the CE misses consumption by the high-income households that are not responding to the CE.¹⁴

Combined, the results in Figures 2 and 3 provide confidence that the imputation of consumption to the SCF is behaving as we would expect. Consumption is higher in the SCF than the CE, and it is higher at higher income levels. With higher consumption especially at the top of the distribution, the SCF consumption aggregates approach spending from the PCE.

¹² Note that high-income households are missed in all surveys, such as the Current Population Survey and American Community Survey. Among household surveys, only the SCF with its oversample of high-income zip codes is fully representative of the top 5-10% of the income distribution.

¹³ The before-tax income measure in Figure 2 excludes capital gains because capital gains are not in scope in the CE income measure.

¹⁴ The PCE itself is also an estimate of aggregate consumption with its own errors and should not be treated as the gold standard for the measure of aggregate consumption.

Lastly, Figure 4 displays the Gini coefficient for consumption in the SCF and CE. As expected, the consumption Gini is higher in the SCF than in the CE in every year because the SCF capture the top of the income distribution, which is largely missed in the CE. When removing households in the top 5% of the income distribution, the consumption Gini in the CE and SCF line up almost perfectly at the two end points, with some differences in the years between but telling the same overall story.

Overall, the results indicate that the SCF differs from the CE in the predictable ways but matches the CE over the part of the income distribution where the two should line up. With that fact established, we move to studying inequality in two and three dimensions. There remains a concern that we built an artificial correlation between income and consumption by imputing consumption using income as an independent variable. However, we find that most of our results on the joint distributions are similar to those found in Fisher et al. (2016b) who use the PSID, where consumption is reported instead of imputed.

IV. Inequality in Two Dimensions

Two dimensional comparisons are much easier than are three dimensional comparisons, and here we make a set of them in multiple combinations and using the Shorrocks index. In general we find that the top quintile of any distribution gained in own and cross shares, and the bottom quintiles lost in own and cross shares over the period of our study. The top quintile has a higher share of income, consumption, and wealth in 2013 than 1989, and there is a stronger correlation between the three measures in 2013 as well. Inequality in any two dimensions combined grew faster than did inequality in any one dimension alone. Further, we find that the gains at the top of all three distributions came from all four lower quintiles. The top quintile has a higher share of income, consumption, and wealth in 2013 than 1989.

Top 5% Shares

We start with a share analysis in the spirit of Piketty and Saez (2003). Before presenting inequality in two dimensions using share analysis, Figure 5 shows the one-dimensional shares of after-tax income, consumption, and wealth held by the top 5% of the respective distribution. One difference between our work with SCF and PSID is that the SCF includes top income earners and wealth holders. While the PSID allows us to add intra-generational mobility and to use all three inequality dimensions from the same dataset, the SCF gives us perspective on the top end of the distribution and macro-micro matches that are not possible with the PSID. For comparison, Figure 5 also includes the share of wealth held by the top 5% of wealth holders from Saez and Zucman (2015), the share of taxable income held by the top 5% of taxable income tax units from Piketty and Saez (2013), and the share of consumption of the top 5% of consumption households from the CE.

The level and trend in wealth shares are remarkably similar considering they originate from different sources. Both series indicate that the share of wealth held by the top 5% of the wealth distribution was around 55% in 1989 and increased to around 60% by 2010. The share of income held by the top 5% is almost identical between the SCF and Piketty and Saez (2003). Lastly, Figure 5 shows the share of consumption for the top 5% of the consumption distribution from the SCF and CE, with similar conclusions from Figure 4 – the SCF shows a higher level of inequality and a larger increase in inequality when compared to the CE.

For summary measures between 1989 and 2013, the growth in the shares is 15 percent for wealth, 14 percent for income, and 41 percent for consumption. The 41 percent growth in consumption inequality is outside the typical range presented in previous research. If not for the results in Figure 4 showing that the growth in the Gini coefficient for consumption in the CE matches the SCF, excluding the top 5% from the SCF, we would be more concerned that the observed growth in the consumption share was a by-product of our imputation.

Now we move to bivariate distribution by examining the percent of one measure held by the top 5% of a separate measure. For example, what percent of wealth is held by the top 5% of the income distribution or the top 5% of the consumption distribution? By definition the top 5% of the income distribution holds less wealth than the top 5% of the wealth distribution, unless the two groups are identical in which case the share would also be identical. But we can learn two important things by presenting these cross-shares. First, how big of a difference is there in the share of wealth held by the top 5% of wealth and the share of wealth held by the top 5% of income or consumption? The difference tells us how much the two distributions overlap. Second, has the share of wealth held by the top 5% of income or consumption increased over time and has it increased by a larger percentage than the uni-dimensional shares? The trend in the cross-shares indicates whether two-dimensional inequality has increased faster than uni-dimensional inequality.

We start with the percent of households in the top 5% of each pair of our three measures, shown in the first panel of Figure 6. In 1989, 2.5 percent of households were in the top 5% of the income and wealth distributions, meaning half of the households that were in the top 5% of the income distribution were also in the top 5% of the wealth distribution. Over half, 2.6 percent, were in the top 5% of wealth and consumption, and 2.3 percent were in the top 5% of income and consumption.

All three series also increase between 1989 and 2013, indicating a growth in two-dimensional inequality as more households are in the top 5% of at least two measures in 2013 than in 1989. The bottom panel of Figure 6 shows the growth in the percent of households in the top 5% of each pair of measures. The highest growth is observed for the income and consumption series, increasing by over 20 percent between 1989 and 2013, to 2.8% of households in the top 5% of the income and consumption distributions. By 2013, 3.1 percent of households were in the top 5% of wealth and consumption.

Figure 7A extends the two-dimensional inequality analysis in Figure 6 and focuses on the own and cross-shares for all measures. In other words, instead of focusing on the percent of households in the top 5%, as in Figure 6, we focus on the share of resources held by those in the top 5%. For wealth in 2013, the top 5% of the wealth distribution held 63 percent of wealth, and the top 5% of the consumption distribution held 53 percent of wealth. The top 5% of income held 51 percent of wealth. The fact that the top of the consumption distribution held more wealth than the top of the income distribution indicates that there is a stronger relationship between consumption and wealth than income and wealth, which may be driven by retired households that have little current income and consume out of wealth and its accretion.

We can also look at the same set of concepts using income and then consumption. The share of income held by the top 5% of income households is 29% in 2013, while the top 5% of consumption and the top 5% of wealth hold 25 percent of income. Similarly with consumption, the share of consumption held by the top 5% of consumption households is 22 percent, while the top 5% of wealth hold 20 percent of consumption and the top 5% of income hold 19 percent of consumption.

We now turn to whether the cross correlation is stronger in 2013 than in earlier years. Figures 7B-7D show the growth in the own and cross-shares, indexed to equal 100 in 1989. In each two-way comparison, the cross-share increases by a higher percent than the own share, indicating a faster increase in inequality in two dimensions. For example with wealth, the percent of wealth held by the top 5% of wealth households increased by 15 percent between 1989 and 2013, while the percent of wealth held by the top 5% of consumption households increased by almost 30 percent (Fisher et al. (2016b) find similar results using the PSID and the top 20% as a cutoff). The pattern of a faster increase in the cross-shares occurs across all three measures. The larger growth in the cross-shares than in the own shares indicates an increase in inequality in two dimensions because there is a greater correlation between the two dimensions in both directions as rich gain and poor lose out.

The top 5% share results represent a detailed look at the top of the distributions and have a long history in economics back to Kuznets, but focusing on top shares misses a deeper understanding of the rest of the distribution. One way to illuminate the dynamics throughout the distribution is by applying the top 5 share results to the entire distribution. The next sub-section turns to the own and cross-shares by quintile to highlight what is happening across the full distribution.

Shares by Quintile

Here we present share results using quintiles of each distribution. We can now see if the top 20% is experiencing the same type of growth in shares as the top 5%, and we can see what quintiles are losing share to the top. While these results shed light on the entire distribution, which has often been ignored in previous research, it still

presents inequality in one dimension. To understand the dynamics of inequality in two dimensions, we then move to the cross-share analysis to determine if inequality in two dimensions is increasing using quintiles.

Figure 8 presents the share of each measure held by each quintile of the respective distribution. The first panel shows the share of income held by each quintile of the income distribution. In 2013, the top quintile of the income distribution received almost 60 percent of income, while the second quintile received just 20 percent. The bottom quintile received 3.5 percent of all income. Between 1989 and 2013, the share held by the top quintile grew by 6.6 percent, while the middle three quintiles all experienced a decrease in share by 8 to 10 percent. The bottom quintile had a slight increase in the share of income, increasing from 3.0 percent to 3.5 percent.

The second panel in Figure 8 shows the shares of consumption, by consumption quintile. The top consumption quintile had almost half of consumption in 2013, up from 40 percent in 1989. Similar to income, the top consumption quintile experienced an increase in its share while the remaining quintiles all lost share, anywhere from a 7 to 17 percent decrease in share.

Wealth shares by quintile tell the same story, but with even more extreme shares and share changes. The top wealth quintile held 87 percent of wealth in 2013, up from 81 percent in 1989. The middle three quintiles of the wealth distribution all lost shares, falling by 50 percent in the second quintile, and as others have reported much of this share took place during the Great Recession (Krueger, Mitman, and Perri, 2016). But the actual loss is relatively small, given that the second quintile only held 1.2 percent of wealth in 1989 to begin with. The bottom wealth quintile has a negative share of wealth in both 1989 and 2013 because most households have zero or negative wealth in the bottom wealth quintile. What we can take from the own shares analyses by quintile is that the top quintile of each distribution gained share between 1989 and 2013 while the other four quintiles lost share or remained flat.

Now we move to our measures of inequality in two dimensions using the cross-shares by quintile, analogous to the cross-shares in Figure 7. The first panel in Figure 9 presents the indexed own and cross-shares for quintile 1 of the three distributions. Remember that the indexed values show the trend since 1989. A number above 100 means the quintile experienced an increase in the share since 1989, while a number below 100 means a decrease in the share. As shown in Figure 8, the bottom 20 percent of income held more income in 2013 than in 1989, from 3.0 percent to 3.5 percent. However, the bottom quintiles of wealth and consumption had a smaller increase in the share of income, with the bottom quintile of consumption actually experiencing a decrease in the share of income, from 5.7 percent to 5.3 percent. Because the bottom quintiles of wealth and consumption had a smaller increase in the share of income than the bottom quintiles of income experienced, or a decrease in the share in the case of consumption, we take this as

another indicator that inequality in two dimensions increased by more than inequality in one dimension.

The wealth story for the bottom quintiles matches what is seen for income among the bottom quintiles. The bottom quintile of wealth had an increase in the share of wealth held, while still holding a negative share of wealth, while the bottom quintiles of income and consumption held a smaller share of wealth in 2013 than in 1989. Again, this suggests that inequality in two dimensions increased by more than inequality in one dimension across the entire distribution.

Being in the bottom quintile of any of the three resource measures means a decrease in the share of consumption between 1989 and 2013 (Figure 9), but the biggest decrease is for consumption quintiles. This is not consistent with a faster increase in inequality in two dimensions than in one dimension, but it still indicates an overall increase in inequality two dimensions.

A similar pattern emerges among the middle quintile in the second panel of Figure 9. All nine series in this panel of Figure 9 are below 100 in 2013, indicating a smaller share in 2013 than in 1989 for all own and cross-shares among those in the middle quintile, which means that inequality in one and two dimensions is higher in 2013. Hence if the middle quintile is the middle class, they lose in all three dimensions. There is, however, no clear pattern of a greater increase in inequality in two dimensions for the middle quintile. Generally the own shares fall by the same percent or by slightly more than the cross-shares.

The last panel of Figure 9 shows the top quintile. These results match what we see for the top 5% in Figure 7, even though the trends are less steep than with the top 5 percent. The cross-shares increase by more than the own shares, for all combinations of resource measures, again providing evidence that inequality in two dimensions increases by more than inequality in one dimension, at least in the top quintile of the distribution, and more so in the top ventile.

Overall, Figures 8 and 9 tell a compelling story. The top quintile of the distribution gained in own and cross shares, and the bottom quintiles lost own and cross shares. The top quintile has a higher share of income, consumption, and wealth in 2013 than 1989, and there is a stronger correlation between the three measures in 2013 as well. We also see that the gains at the top came from all four lower quintiles.

Transition Matrices

Another way to illuminate the dynamics of two-dimensional inequality throughout the distribution is with transition matrices and summary measures such as the Shorrocks measure of mobility, which is based on the transition matrix. The transition matrix displays the percent of households in the various quintiles of each distribution and is more often used to study inter or intra-generational mobility over time within the same resource measure.

In our context, the transition matrix and Shorrocks measure provide additional indicators of the degree of correlation between the two chosen variables. Figure 10 shows the transition matrices between income and consumption, income and wealth, and wealth and consumption for 2013. As an example of how to interpret a transition matrix, approximately 13.1 percent of households were in the bottom quintile of the income distribution and the bottom quintile of the consumption distribution, and 13.8 percent of households were in the top income and consumption quintiles (Figure 10). Note that only 20 percent of households can be in any pair of quintiles. Thus another way to interpret the observation that 13.8 percent of households were in the top quintile of the income and consumption quintiles is that 69% of households in the top quintile of income were also in the top quintile of consumption. This is a stronger correlation than we saw in the top 5% from Figure 6, where 56% of households in the top 5% of income were also in the top 5% of consumption. The transition matrix allows us to see where the 31% of households that are in the top quintile of income but not the top quintile of consumption. The vast majority (24% of the 31%) is in the fourth quintile of consumption.

The mobility literature finds a twin peaks phenomenon (e.g., Fisher and Johnson, 2006 and Fisher et al. (2016b)), where there is more stickiness at the top and bottom of the distributions than in the middle of the distributions. We find the same twin peaks for each pair of resource measures (Figure 10). We see a stronger relationship between income and consumption than between income and wealth. As mentioned above, 13.8 percent of households are in the top quintiles of income and consumption, and 13.1 percent are in the bottom quintile of both. Only 8.8 percent of households are in the bottom quintiles of income and wealth, and 12.1 percent of households are in both top quintiles. The transition matrix for consumption and wealth looks more like the matrix for income and consumption.¹⁵

Another way to utilize the transition matrix is by presenting a summary measure of mobility such as the Shorrocks index. The Shorrocks index is calculated as five minus the sum of the main diagonal, all divided by four. This summary measure can be interpreted as the percent of individuals that change quintiles. Table 1 presents the Shorrocks index for all three pairs of resource measures and for 1989 and 2013. In 1989, 62.5 percent of households were in different quintiles of the consumption and income distributions, while in 2013 61.9 percent of households were in different consumption and income quintiles.

As Figure 10 shows, Table 1 also confirms that there is more mobility (i.e., a smaller correlation) in the income and wealth distributions than in the other two distributions, while income and consumption are most highly correlated among the three pairs. These patterns hold in 1989 and in 2013.

¹⁵ Fisher et al. (2016b) also find a larger concentration at the top of the joint distributions.

In contrast to the share analysis above, the Shorrocks measure of mobility provides no evidence for a clear increase in inequality in two dimensions between 1989 and 2013. There is a small decrease in mobility between 1989 and 2013 for consumption and income, from 62.5 percent of households in different quintiles in 1989, while 61.9 percent are in different quintiles in 2013, a slight increase in the correlation between the two distributions. On the other hand, there is a small increase in mobility for consumption and wealth between 1989 and 2013. Because a significant portion of the growth in inequality occurred in the top 5%, expanding the bins to include the top 20% instead of just the top 5% may change our view slightly on inequality in two dimensions.

While the transition matrices do not obviously show a greater increase in inequality in two dimensions than one dimension, they allow us to examine the entire distribution and not just the top 5%. To understand inequality and its dynamics, we need to view the entire distribution, which is even more important for trying to understand inequality in more than one dimension.

V. Inequality in Three Dimensions

As there are no existing measures of inequality in three dimensions, we continue to use the shares concept from measuring inequality in two dimensions and a new ordinal measure based on ventile ranks later in this section.

To analyze the share of income, we show the share of income held by the top 5% of the income distribution along with the share of income held by households in the top 5% of the income distribution *and* the top 5% of the consumption distribution. It is then the share of income held by those households in the top 5% of two dimensions, not just one. Similar to our analysis of inequality in two dimensions, inequality in three dimensions is increasing faster than inequality in one dimension if the share of income held by the top 5% of the income and consumption distributions grows faster than the share of income held by the top 5% of the income distribution.

We return to Figure 6 first, which shows the percent of households in the top 5% of all three resource measures. In 1989, 1.6 percent of households were in the top 5% of income, consumption, and wealth, indicating that almost one in three households in the top 5% of income were also in the top 5% of consumption and wealth. By 2013, 2.2 percent of households were in the top 5% of all three measures. The bottom panel of Figure 6 shows that this represents a 42 percent increase in the percent of households in the top 5% of all three measures. By this measure of inequality in two and three dimensions, inequality in three dimensions increased faster than inequality in two dimensions.

Figure 12 presents analogous three dimensional inequality results to the two dimensional inequality results in Figure 7. Here we focus on the growth in the shares, not the levels, but we will reference levels in the text. The first panel in

Figure 12 uses income as the resource measure. Between 1989 and 2013, we see the result shown above that the share of income held by the top 5% of income increased by 14.8 percent. The share of income held by the top 5% of the consumption and wealth distributions increased by 48.5 percent between 1989 and 2013, from 15.9 percent to 23.6 percent. The same pattern exists for the top 5% of income and consumption and the top 5% of income and wealth, relative to the growth in the share of income held by the top 5% of income. Once again inequality in three dimensions increased by more than inequality in one dimension.

The pattern follows when using consumption or wealth as the resource measure as well. The share of consumption by the top 5% of the consumption distribution increased by 41.6 percent, from 15.8 to 22.4 percent. The share of consumption increased by 87.0 percent for those households in the top 5% of the income and wealth distributions, from 7.6 to 14.2 percent. For wealth, the largest increase in the share was for the top 5% of the income and consumption distributions, increasing by 42.6 percent, from holding 29.3 of wealth to holding 41.8 percent of wealth. In contrast, the top 5% of the wealth distribution increased the share of wealth by 15.7 percent.

It is worth reemphasizing that this increase in inequality in three dimensions is much greater than the increase in one dimension. Thus viewing inequality through one dimension greatly understates the growth in inequality in two and three dimensions. The conclusion is that the U.S. is becoming more unequal than is generally understood in the income or consumption or wealth inequality literature once all three are considered together.

We now return to transition matrices to measure inequality, but now moving to three dimensional inequality. Figure 9 shows the transition matrix between income and consumption in 2013 by wealth quintile. The figure shows the strength of the correlation between income, consumption, and wealth. For example, 6.8 percent of all households are in the bottom quintiles of income, consumption, and wealth. Another way to interpret this result is that more than one in three households that are in the bottom of the wealth quintile are also in the bottom quintile of income and consumption. Less than 1 percent of households who are in the bottom quintile of wealth reach the top quintile of income *or* consumption. Once again multiple measures reinforce one another

At the other extreme, 10.0 percent of households are in the top quintiles of income, consumption, and wealth, or one in two households that are in the top quintile of wealth are also in the top quintiles of income and consumption. Almost 98 percent of households that are in the top quintile of wealth are in the top three quintiles of the consumption distribution, suggesting importantly that high wealth provides a minimum level of consumption, regardless of income. Given that median wealth in the top wealth quintile is \$1.9 million in the SCF, it should not be surprising that being in the top wealth quintile also means households are in the top three quintiles of consumption.

In the middle wealth quintile, households are relatively evenly spread throughout the income and consumption distributions, though fewer reach the top quintile. Almost 13 percent of households are in the middle wealth quintile along with the middle three quintiles of the income and wealth distributions, meaning 64 percent of those households in the middle wealth quintile are also in the middle three quintiles of the income and consumption distributions.

Taken together, Figure 9 provides evidence of a twin peaks phenomenon in three dimensions. Being in the bottom wealth quintile means you are also overrepresented in the bottom income and consumption quintiles, and being in the top wealth quintile means you are overrepresented in the top of the income and consumption quintiles. Future work will also look for trends in 3D inequality from 1989 to 2013.

Lastly, we extend the idea behind transition matrices to create another rank only measure of inequality in two and three dimensions. We assign everyone to a ventile of the three individual distributions, and then we add their ventile rank for two measures and for three measures. Thus, a household in the second ventile of income and the fourth ventile tile of consumption will have a score of six when measuring inequality in two dimensions with income and consumption. Adding a wealth ranking in the 3rd ventile gives the household a score of nine. The range of scores in two dimensions is [2, 40] and in three dimensions is [3, 60]. We then calculate the Gini coefficient of the scores and look at the trends over time, which is shown in Figure 13.

The results show that inequality in three dimensions increased by about 1% between 1989 and 2013 including a decline in the early 1990s, a steep rise from below .99 to above a 1.01 from 1995 to 2004 and dip during the Great Recession (GR) followed by a rebound through 2013. Inequality in income and consumption taken together increased by about the same percent, but with a peak in 2001 and without a large drop during the GR. In contrast, inequality of two dimensional ranks in income and wealth and consumption and wealth followed the same ups and downs as the three dimensional measure, but with the GR depressing the income and wealth combination the most.

In future work, we plan to use more of these types of comparisons, essentially simplifying the distributions into a combination relative positions and then looking at the medians with across each ranking. *Some feedback on this method is much appreciated.*

VI. Conclusions

We do not live in a world where we need to choose between income, consumption, or wealth as the superior measure of well-being. All three individually and jointly determine the well-being of households, and we have access to data that contain all

three measures, including the values for the highest strata of income and wealth. Individually, we show that inequality in one dimension has increased since 1989 for each measure. Further, our results point toward an increase in inequality in two dimensions and in three dimensions, and with this increase in inequality in two and three dimensions being faster than the increase in inequality in one dimension. The top 5 percent increased its concentration in all three dimensions. The top quintile of the distribution gained in own and cross shares, and the remaining four quintiles lost own and cross shares. The top quintile has a higher share of income, consumption, and wealth in 2013 than 1989, and there is a stronger correlation between the three measures in 2013 as well. The gains at the top came from all four lower quintiles. And finally, inequality in three dimensions increases the fastest.

By presenting results using income, consumption, and wealth for the same households, we firmly believe that we improve our understanding of inequality in the U.S. The three measures are not perfectly correlated, but there is much higher correlation in the tails of the distribution, suggesting that uni-dimensional inequality understates the true level of overall economic inequality. Most research on inequality focuses on income alone, but that may well not be the best single measure of inequality. Economists prefer consumption as a measure of permanent income, and we argue that wealth incorporates both the ability to increase income and the ability to consume directly, allowing top quintile and top ventile households to smooth shocks in income or consumption at their own whim.

In future work, we plan to devote some added energy to analysis of the bottom quintile and the way that each dimension affects upward mobility. The role of debt in holding back families at the bottom will be examined in the same spirit that we look at the way that top shares of wealth and income convey advantage.

Wealth is a stock that can be used to stabilize consumption in times of misfortune, or to increase realized income flows. Because wealth is so highly skewed and becoming more skewed over time, it also allows the wealthy to ensure the economic success of their children. That is, personal and business wealth provides one an ever-increasing cushion against economic misfortune and a dynastic advantage to maintaining one's own social position over time. This advantage can also extend across generations providing the ability to save, purchase durables, finance education, borrow more cheaply, and the ability to do more for inter-vivos transfers, especially for the following generations. Accumulated wealth clearly leads to benefits for the children of high-wealth households over other children, and thus may compromise equality of opportunity and diminish intergenerational mobility (Fisher, Johnson, Latner, Smeeding, and Thompson, 2016a). And this might well be the most important consequence of the tripartite increases in the cross-nationally high levels of inequality observed in the United States over this period.

References

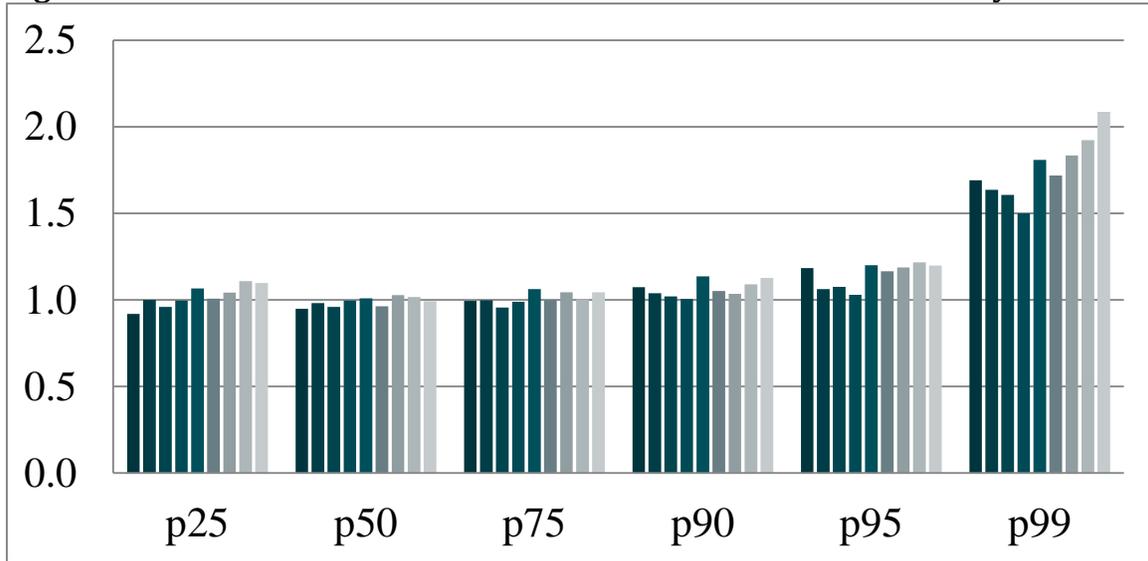
- Aguiar, M., and Bils, M. "Has Consumption Inequality Mirrored Income Inequality," *American Economic Review* 105(9), 2015.
- Aguiar, M., and Hurst, E., "Measuring Trends in Leisure: The Allocation of Time over Five Decades," *Quarterly Journal of Economics* 122(3), 2007.
- Armour, P., Burkhauser, R., and Larrimore, J.. "Levels and Trends in United States Income and its Distribution: A Crosswalk from Market Income Towards a Comprehensive Haig-Simons Income Approach," *Southern Economic Journal*, 81(2) 2014.
- Attanasio, O., and Pistaferri, L., "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure," *American Economic Review* 104(5), 2014.
- Attanasio, O., and Pistaferri, L., "Consumption Inequality," *Journal of Economic Perspectives* 30(2), 2016.
- Auerbach, A., Kotlikoff, L., and Koehler, D., "U.S. Inequality, Fiscal Progressivity, and Work Disincentives: An Intragenerational Accounting," *NBER Working Paper* 22032 (February 2016).
- Bee, C. A., B. Meyer, and J. Sullivan, "Micro and Macro Validation of the Consumer Expenditure Survey", in C. Carroll, T. Crossley, and J. Sabelhaus (eds), *Improving the Measurement of Consumer Expenditures*, University of Chicago Press, Chicago, IL, 2014.
- Benhabib, J., Bisin, A. and Zhu, S., "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents," *Econometrica* 79(1), 2011.
- Blundell, R. 2014. "Income Dynamics And Life-Cycle Inequality: Mechanisms and Controversies," *The Economic Journal*, 124, 289–318.
- Blundell, Richard, Pistaferri, Luigi and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, 98:5, pp. 1887–192, 2008.
- Bostic, R., Gabriel, S., and Painter, G., "Housing Wealth, Financial Wealth, and Consumption: New Evidence from Micro Data," *Regional Science and Urban Economics* 39 2009.
- Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus. 2016. "Measuring income and wealth at the top using administrative and survey data". *Brookings Papers on Economic Activity*, in press

- Bricker, Jesse, Lisa J. Dettling, Alice Henriques, Joanne W. Hsu, Kevin B. Moore, John Sabelhaus, Jeffrey Thompson, and Richard A. Windle, 2014. "Changes in U.S. Family Finances from 2010 to 2013: Evidence from the Survey of Consumer Finances," *Federal Reserve Bulletin*, Vol. 100, no. 4.
- Burkhauser, Richard V., Shuaizhang Feng, Stephen Jenkins and Jeff Larrimore. "Recent Trends in Top Income Shares in the USA: Reconciling Estimates from March CPS and IRS Tax Return Data." *Review of Economics and Statistics*, 94(2), 2012.
- Carroll, C., Otsuka, M., and Slacalek, J. "How Large are Housing and Financial Wealth Effects? A New Approach," *Journal of Money, Credit, and Banking* 43(1), 2011.
- Castenada, A., Diaz-Gimenez, J., and Rios-Rull, J., "Accounting for the U.S. Earnings and Wealth Inequality," *Journal of Political Economy*, 111(4), 2003.
- Dettling, Lisa J., Sebastian J. Devlin-Foltz, Jacob Krimmel, Sarah J. Pack, and Jeffrey P. Thompson (2015). "Comparing Micro and Macro Sources for Household Accounts in the United States: Evidence from the Survey of Consumer Finances," Finance and Economics Discussion Series 2015-086. Washington: Board of Governors of the Federal Reserve System.
- Fagerang, A., Guiso, L., Malacrino, D., and Pistaferri L., "Wealth Returns Dynamics and Heterogeneity," Paper presented at the 2016 ASSA Meetings.
- Fisher, J. and Johnson, D. "Consumption Mobility in the U.S.: Evidence from two panel data sets," *Topics in Economic Analysis and Policy*, Vol 6:1, 2006.
- Fisher, J., Johnson, D., and Smeeding, T., "Inequality of Income and Consumption in the U.S.: Measuring the Trends in Inequality from 1984 to 2011 for the Same Individuals" *Review of Income and Wealth* 61(4), 2015.
- Fisher, J., Johnson, D., and Smeeding, T., "Exploring the Divergence of Consumption and Income Inequality during the Great Recession," 2014; Paper presented at AEA 2014.
- Fisher, J., Johnson, D., Latner, J., Smeeding, T., and Thompson, J., "Inequality and Mobility using Income, Consumption, and Wealth for the Same Individuals," 2016a; *RSF Journal of the Social Sciences* 2(6).
- Fisher, J., Johnson, D., Latner, J., Smeeding, T., and Thompson, J., "The Synergistic Distributions of Income, Consumption, and Wealth," 2016b, Working Paper.
- Fisher, J., Johnson, D., Smeeding, T., and Thompson, J., "The Demography of Inequality: Income, Consumption, and Wealth, 1989-2013" 2015. Paper presented at PAA 2015.

- Garner, Thesia I., George Janini, William Passero, Laura Paszkiewicz, and Mark Vendemia. 2006. "The CE and the PCE: A Comparison." *Monthly Labor Review* 66 (September), 20-46.
- Heathcote J., Perri, F., and Violante, G., "Unequal We Stand: An Empirical Analysis of Economic Inequality in the US, 1967-2006", *Review of Economic Dynamics* 13 2010
- Internal Revenue Service, 1992, *Individual Income Tax Returns*, 1990 Internal Revenue Service, United States .
- Jäntti, M., E. Sierminska, and T. M. Smeeding. 2008. "How is Household Wealth Distributed? Evidence from the Luxembourg Wealth Study." In *Growing Unequal*, Paris: OECD, pp. 253–271.
- Kaplan, G., and Violante, G., "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica* 82(4), 2014.
- Kaplan, G., Violante, G., and Weidner, J., "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity* (Spring), 2014.
- Kennickell, A. B., 1998, "List Sample Design for the 1998 Survey of Consumer Finances," FRB Working Paper, Washington: Board of Governors of the Federal Reserve System.
- _____, 1999. "Revisions to the SCF Weighting Methodology: Accounting for Race/Ethnicity and Homeownership," FRB Working Paper, Washington: Board of Governors of the Federal Reserve System.
- _____, 2000, "Wealth Measurement in the Survey of Consumer Finances: Methodology and Directions for Future Research," FRB Working Paper, Washington: Board of Governors of the Federal Reserve System.
- Krueger, D., Mitman, K., and Perri, F., "Macroeconomics and Household Heterogeneity," *NBER Working Paper 22319*. (June 2016).
- Krueger, D. and Perri, F. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *Review of Economic Studies*, 73, pp. 163-193.
- Krusell, P., and Smith, A., "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy* 106(5), 1998.
- Meyer, B., and Sullivan, J., "Consumption and Income Inequality in the U.S. Since the 1960s," Working Paper, July 2016.
- Pfeffer, F., Schoeni, R., Kennickell, A., and Andreski, P. "Measuring Wealth and Wealth Inequality: Comparing Two U.S. Surveys." Working paper. 2014.

- Piketty, T., and Saez, E. 2003. "Income Inequality in the United States: 1913-1998," *Quarterly Journal of Economics* 118(1).
- Sabelhaus, J., Johnson, D., Ash, S., Swanson, D., Garner, T., Greenlees, J., and Henderson, S. "Is the Consumer Expenditure Survey Representative by Income?" in C. Carroll, T. Crossley, and J. Sabelhaus (eds), *Improving the Measurement of Consumer Expenditures*, University of Chicago Press, Chicago, IL, 2014.
- Saez, E., and Zucman, G. 2015. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *Quarterly Journal of Economics* forthcoming.
- Stiglitz, J.E., Sen, A., Fitoussi, J. 2009. *Report by the Commission on the Measurement of Economic Performance and Social Progress*. United Nations Press
- U.S. Bureau of Labor Statistics, "Consumer Expenditures in 2013," *BLS Reports*, Report 1053 (Feb 2015).
- Wilson, O. H., and J. William J. Smith, 1983, "Access to Tax Records for Statistical Purposes," *Proceedings of the Section on Survey Research Methods*, American Statistical Association, pp. 595-601.
- Wolff, Ed. "Household wealth trends in the United States, 1983-2010," *Oxford Review of Economic Policy*, 30(1) 2014.

Figure 1: Ratio of SCF Before-Tax Income to CE Before-Tax Income by Year



Sources: Consumer Expenditure Survey and Survey of Consumer Finances.
Notes: The figure shows the ratio of SCF before-tax income to CE before-tax income. SCF before-tax income excludes capital gains, as capital gains are not reported in the CE. Were capital gains included, the differences would be larger at the top of the distribution. There is one bar per year for each percentile shown in the figure. The years correspond to the years in the SCF: left-to-right 1989, 1992, 1995, 1998, 2001, 2004, 2007, 2010, and 2013.

Figure 2A: Mean Reported Consumption by Before-Tax Income Vingtile in the SCF and CE Survey, 2013

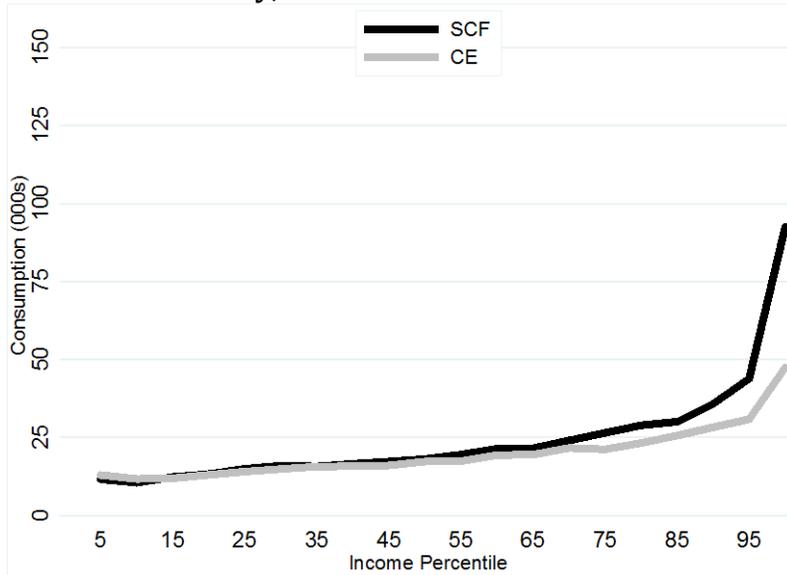
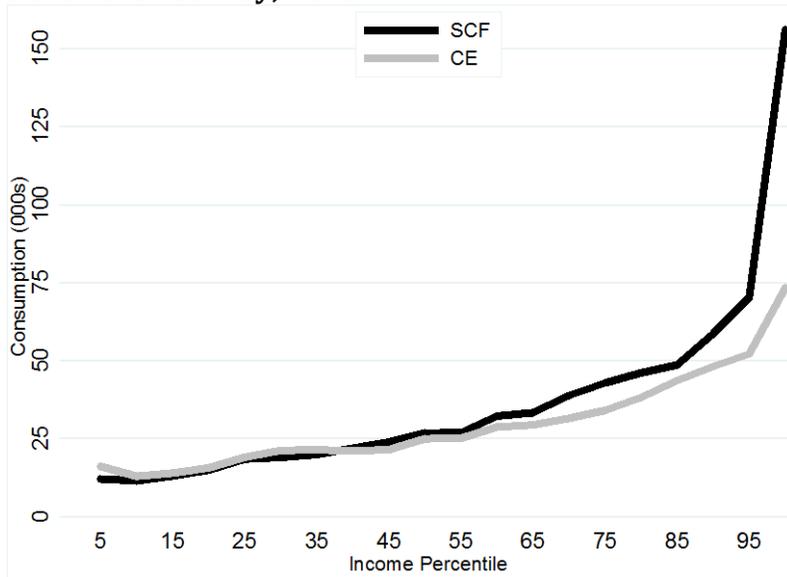


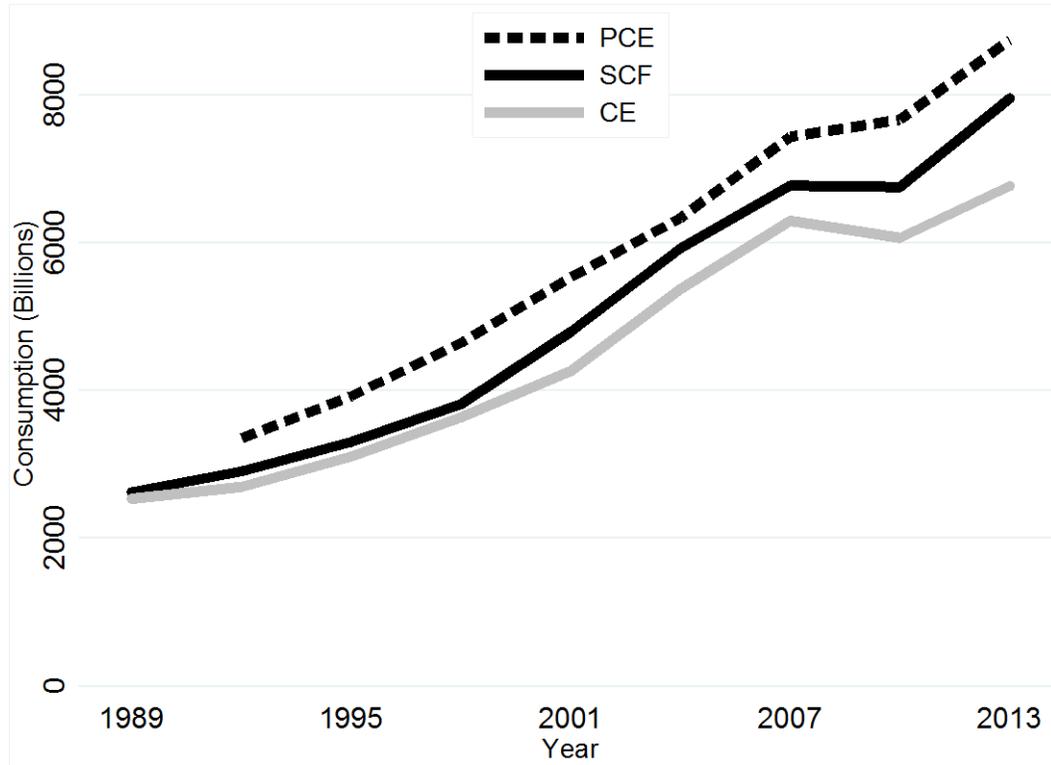
Figure 2B: Mean Imputed Consumption by Before-Tax Income Vingtile in the SCF and CE Survey, 2013



Sources: Consumer Expenditure Survey and Survey of Consumer Finances.

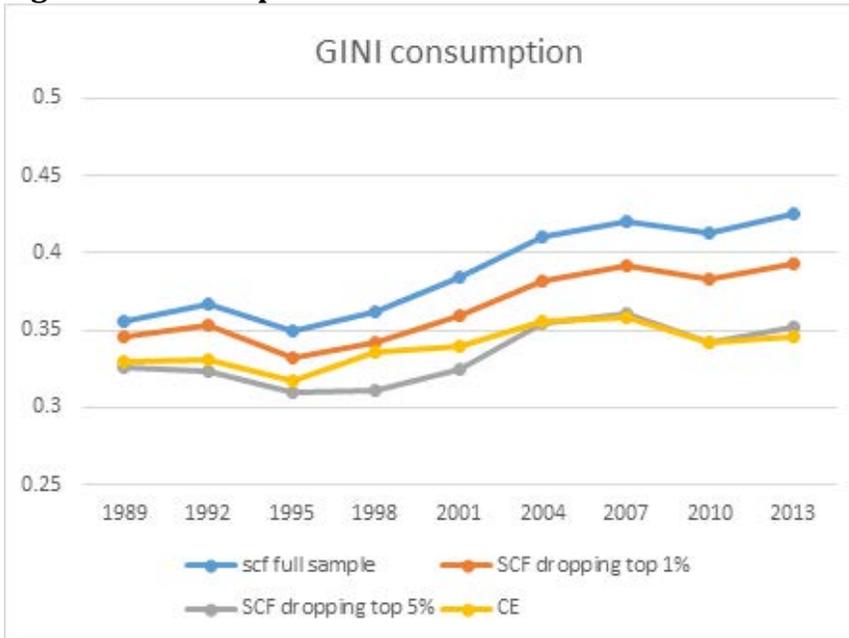
Notes: The top figure shows the consumption components reported in both the CE and SCF. It shows mean reported consumption by income vingtile. The bottom figure shows mean imputed consumption by income vingtile.

Figure 3: Comparing Consumption Aggregates Between the PCE, CE, and SCF: 1989-2013



Notes: PCE aggregates come from Garner et al. (2006) and subsequent updates. The PCE aggregates only include comparable items and definitions.

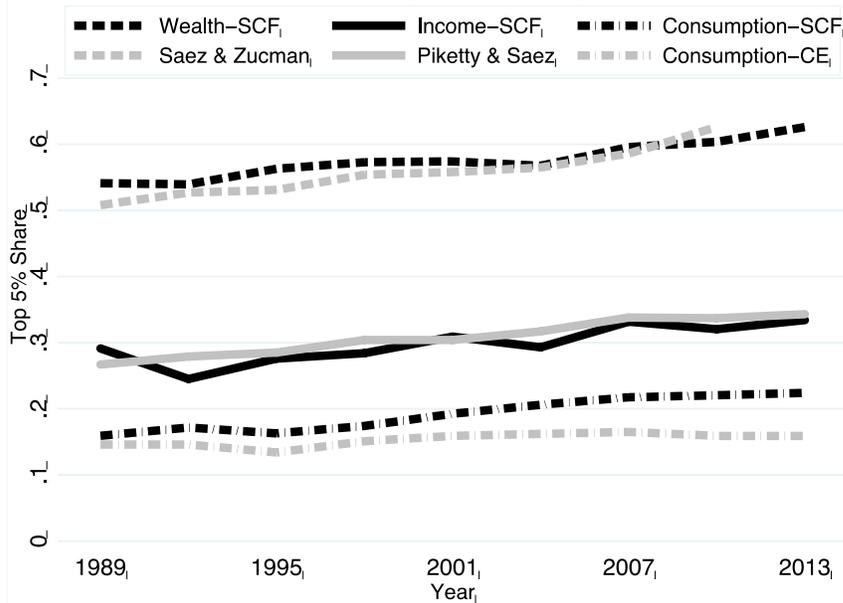
Figure 4: Consumption Gini in the SCF and CE: 1989-2013



Sources: Survey of Consumer Finances and Consumer Expenditure Survey

Notes: The line dropping the top 5% removes the top 5% of the income distribution from the SCF and then calculates the Gini coefficient for consumption in the SCF. This sample attempts to mimic the sample in the CE.

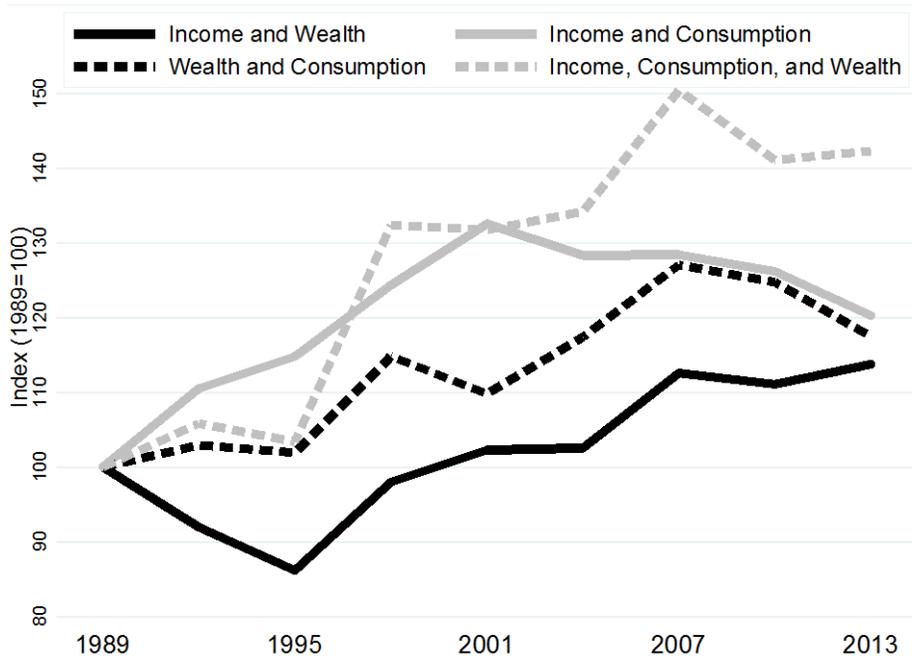
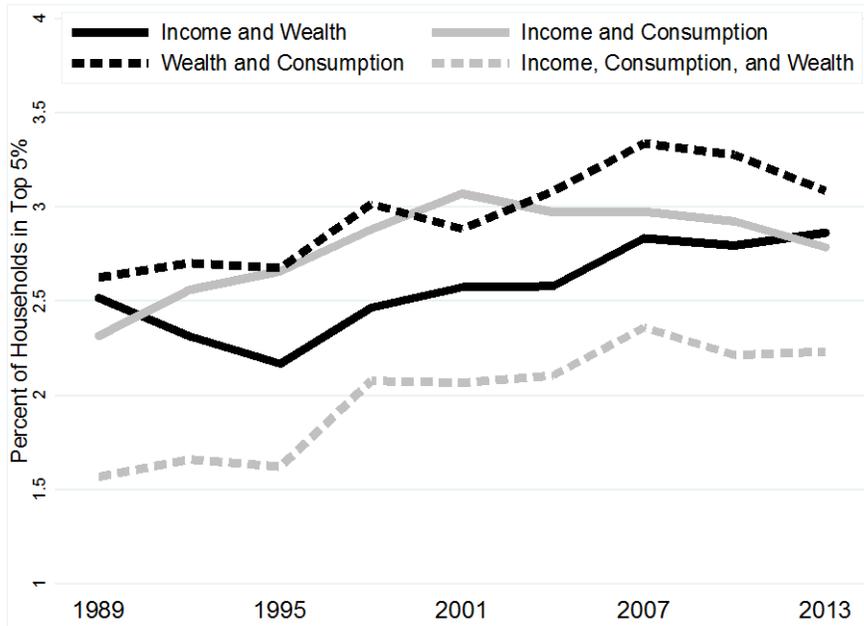
Figure 5: Shares Held by Top 5% of Respective Distribution, 1989-2013



Sources: Survey of Consumer Finances and Consumer Expenditure Survey

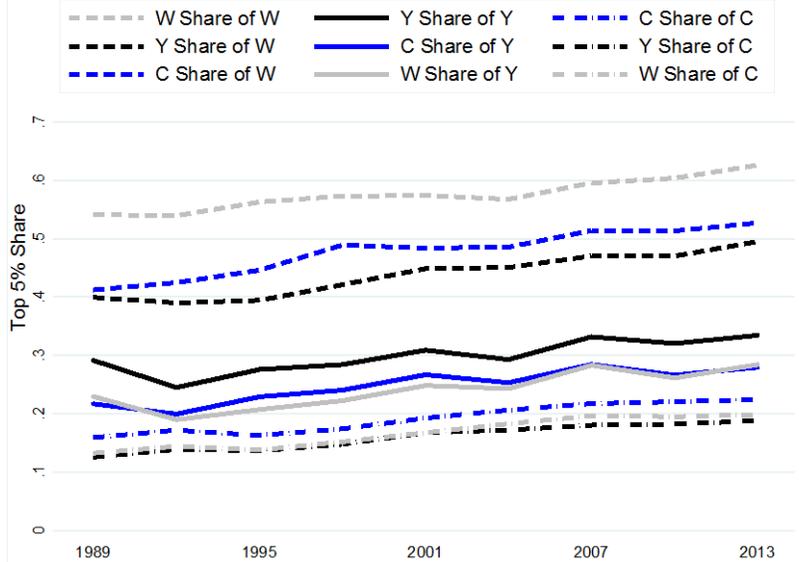
Notes: The non-SCF wealth shares come from Saez and Zucman (2015). The non-SCF income shares come from Piketty and Saez (2003) and from updates on the World Wealth and Income Database <<http://www.wid.world/>>.

Figure 6: Percent of Households in Top 5% of Two or Three Measures (1989-2013)



Sources: Survey of Consumer Finances, 1989-2013

Figure 7A: Top 5% Own and Cross-Shares (1989-2013)



Notes: Three solid lines represent share of wealth held by top 5% of Y, C, and W. Dashed lines represent share of income earned by top 5% of the three distributions. Dotted lines represent share of consumption consumed by top 5% of the three distributions.

Figure 7B: Top 5% Own and Cross-Shares Indexed to 1989=100

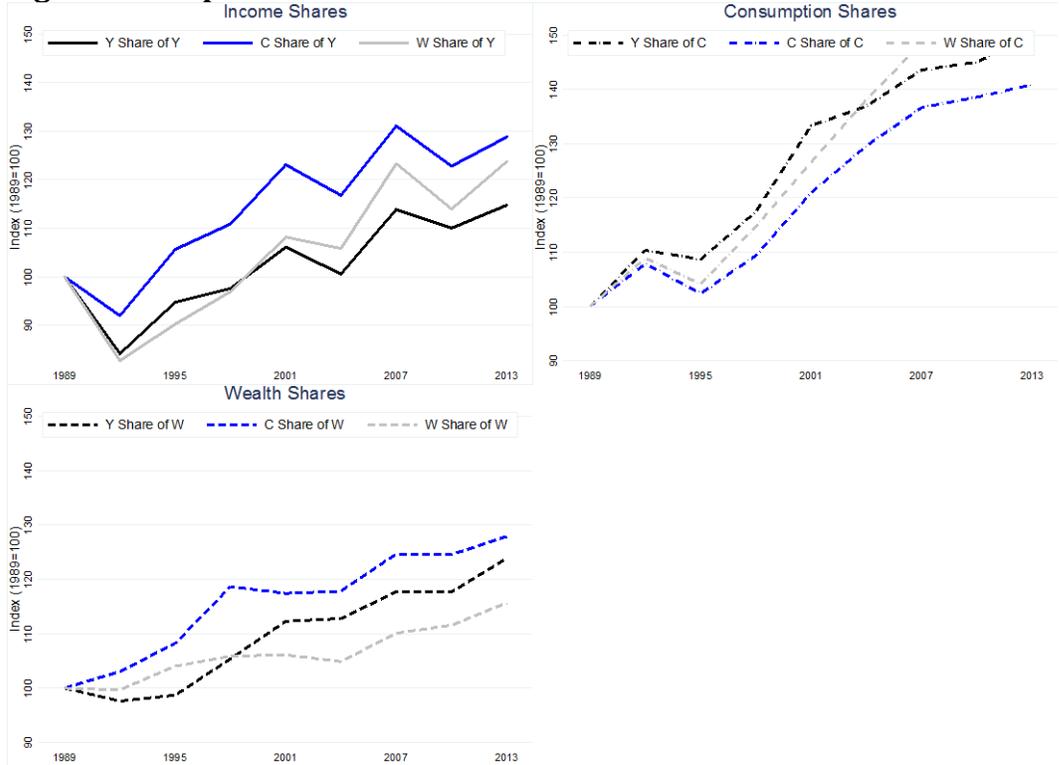


Figure 8: Own Shares by Quintile: Income, Consumption, and Wealth (1989-2013)

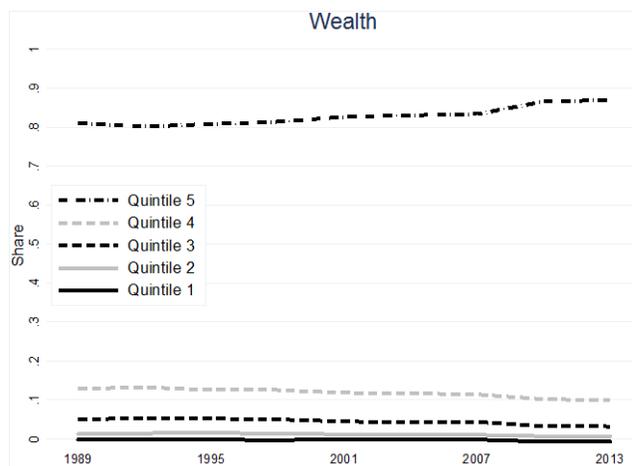
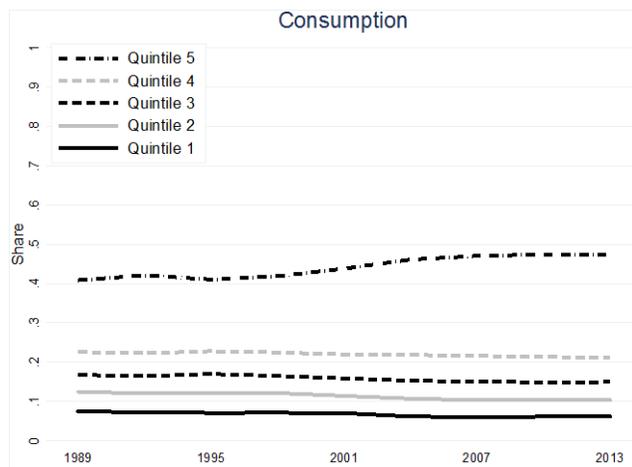
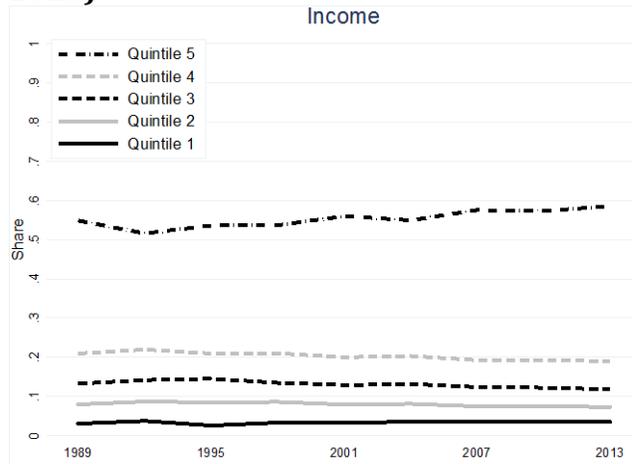
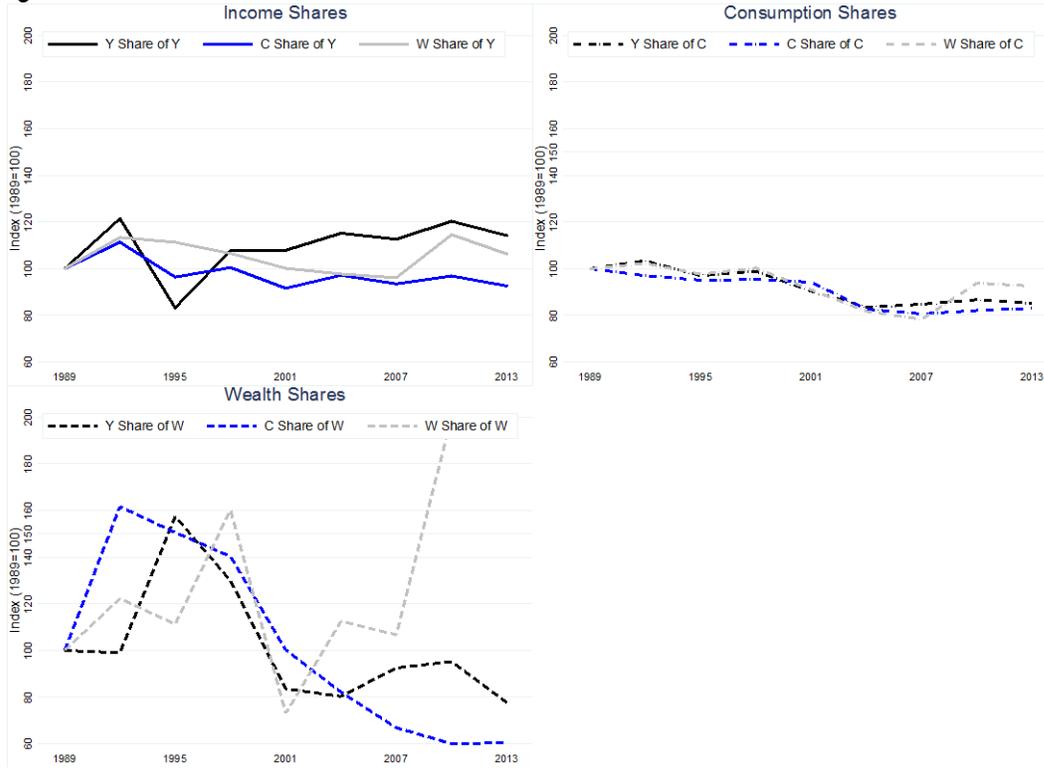


Figure 9: Indexed Own and Cross-Shares by Quintile (1989-2013)

Quintile 1



Quintile 3

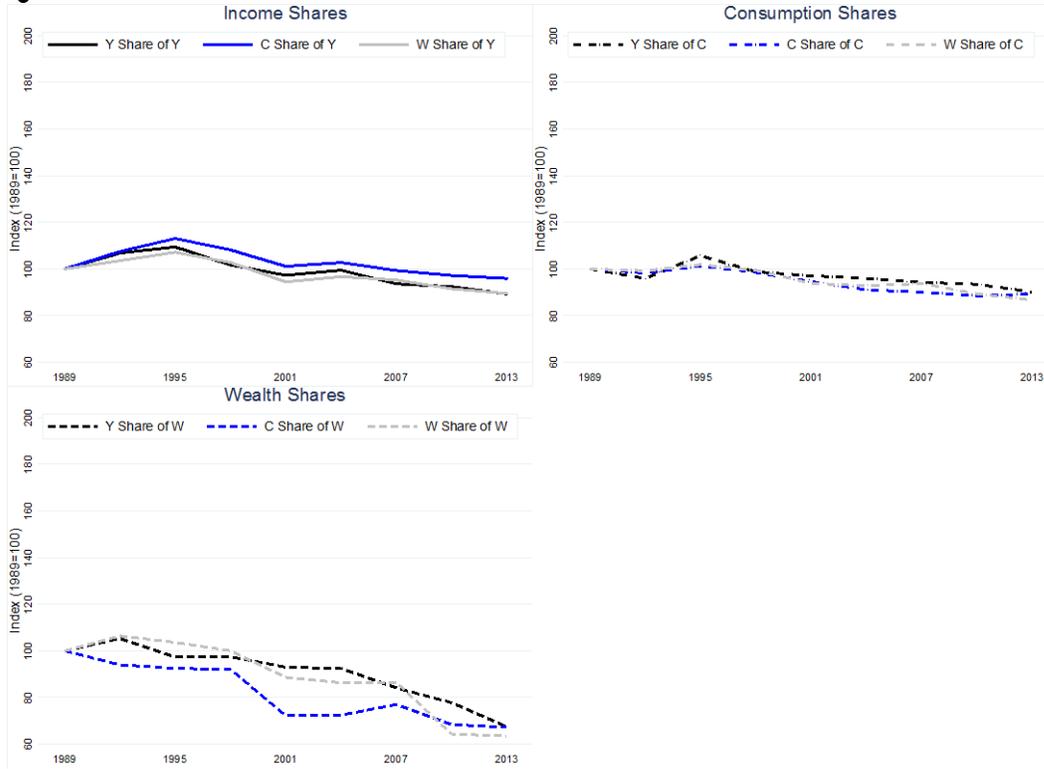


Figure 9 continued: Indexed Own and Cross-Shares by Quintile (1989-2013)
Quintile 5

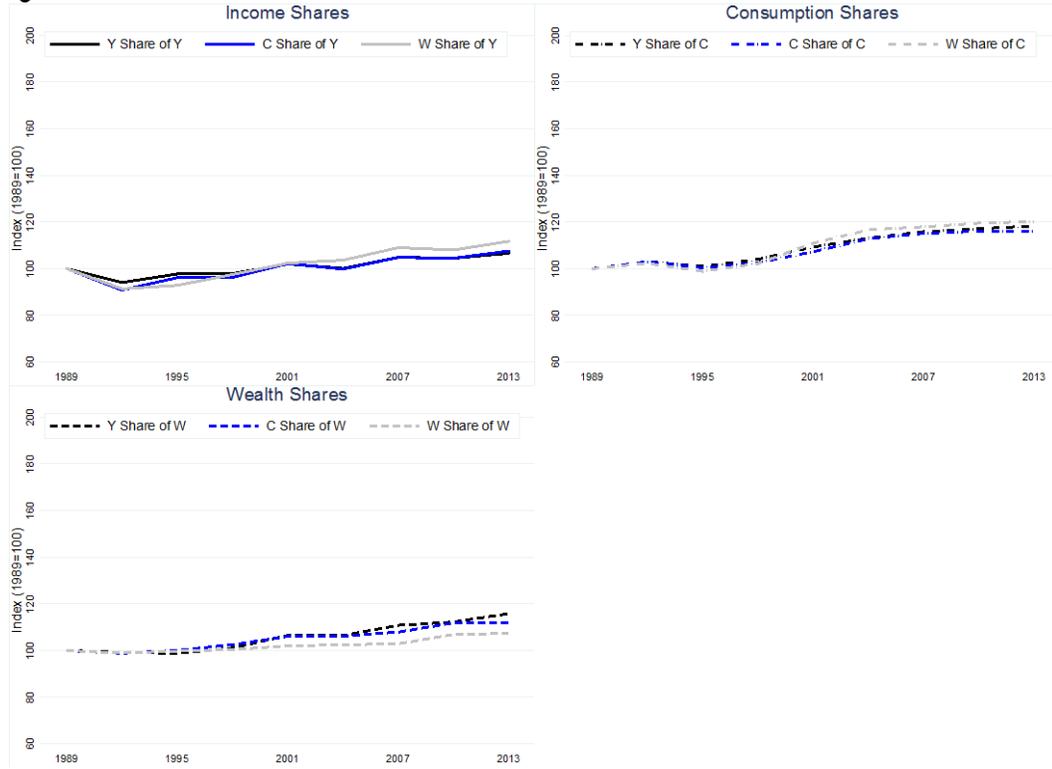
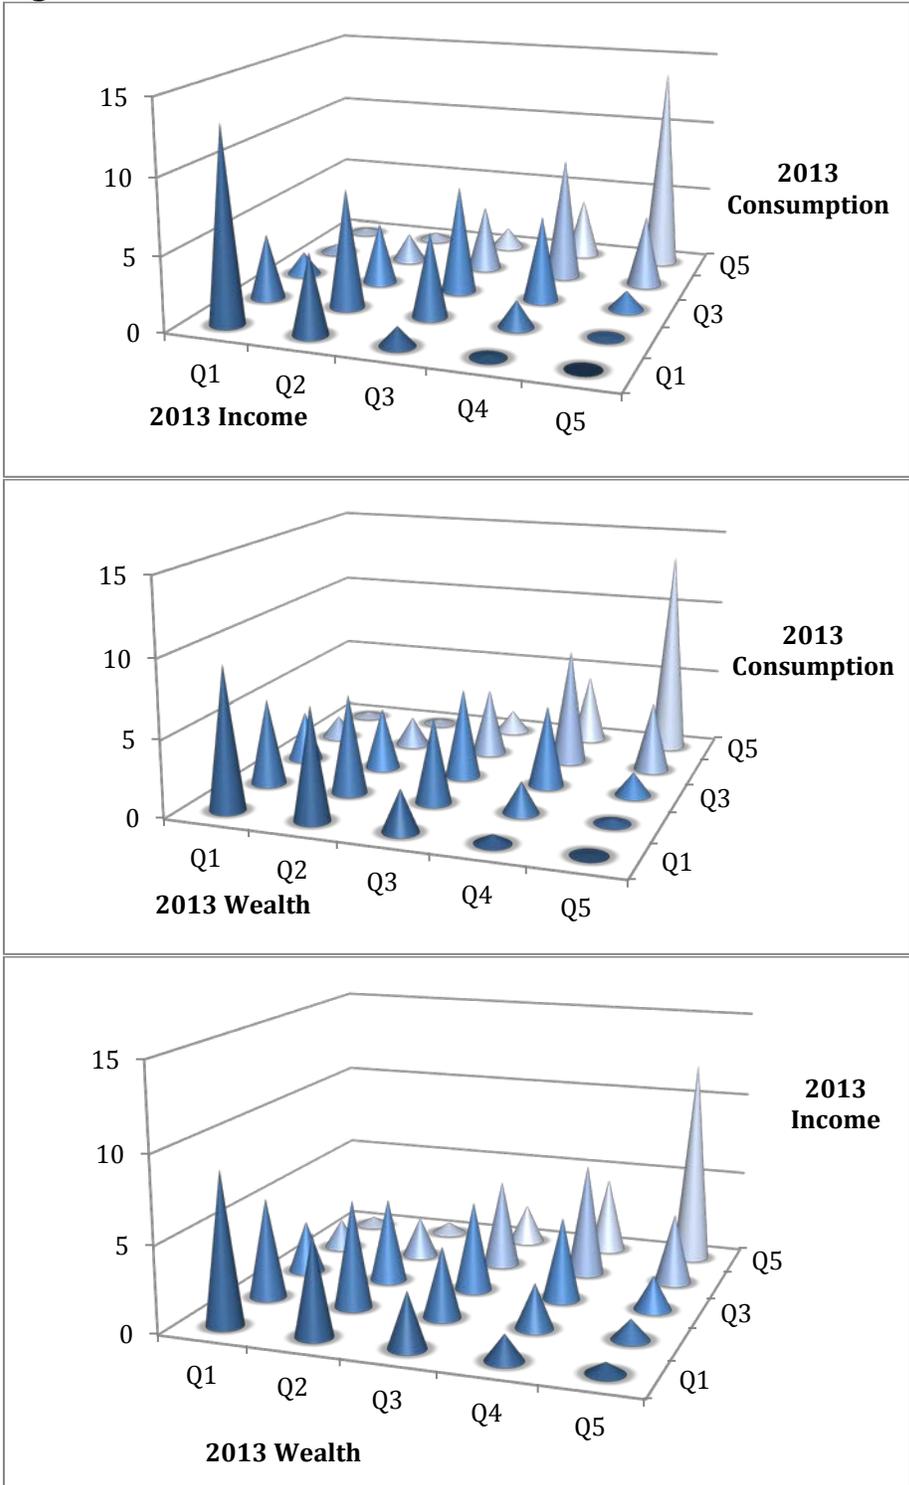
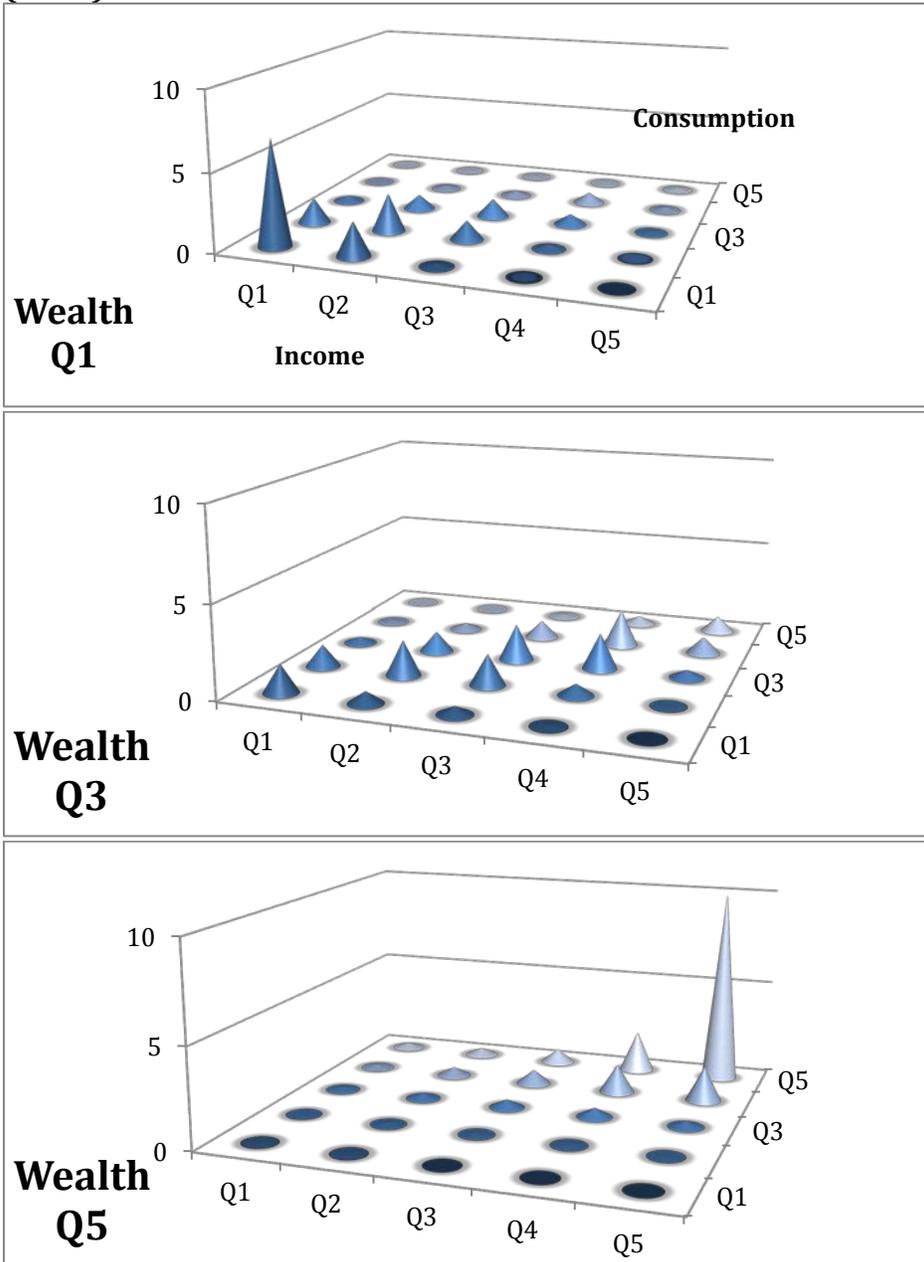


Figure 10: Transition Matrices Between Pairs of Resource Measures (2013)



Source: 2013 Survey of Consumer Finances

Figure 11: Consumption and Income Transition Matrix by Wealth Quintile (2013)



Source: Survey of Consumer Finances

Figure 12: Indexed Own and 3-D Cross-Shares for the Top 5% (1989-2013)

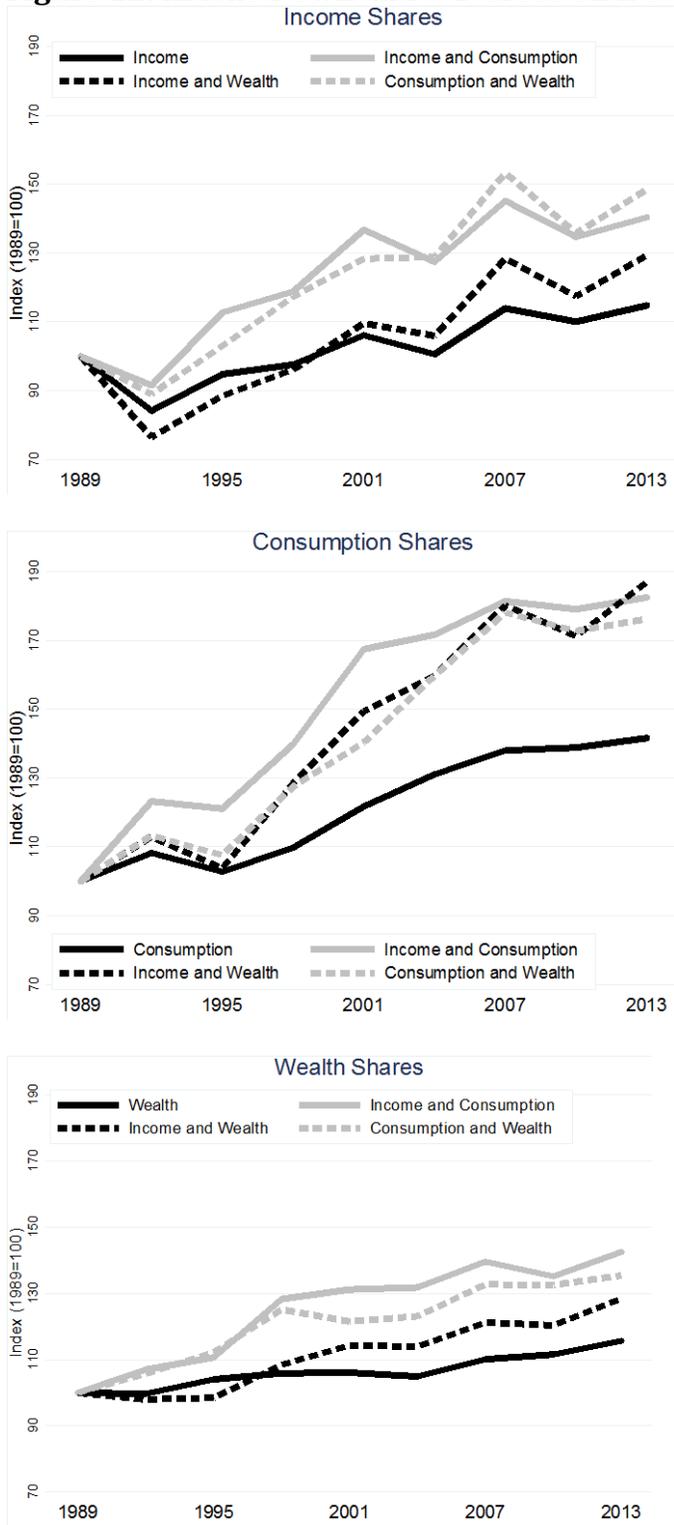


Figure 13: 3-D Measure using the Gini of Additive Ranks (1989-2013)

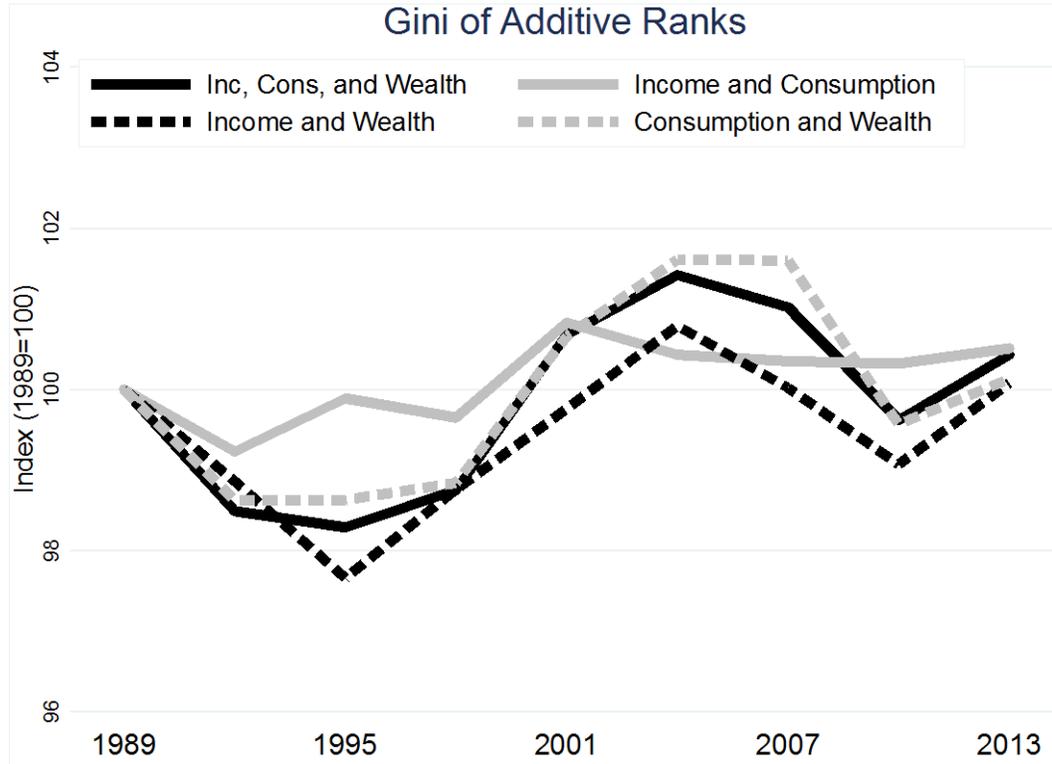


Table 1: Shorrocks Index of Mobility (1989 and 2013)

	1989	2013
Consumption x Income	0.625	0.619
Consumption x Wealth	0.702	0.716
Income x Wealth	0.766	0.762

Source: Survey of Consumer Finances

Notes: The Shorrocks Index can be interpreted as the percent of households that are in different quintiles of the two distributions. For example, in 1989 62.5 percent of households were in different quintiles of the consumption and income distributions.