

Consumption Inequality

Orazio P. Attanasio and Luigi Pistaferri

Much of the debate over the rising levels of inequality in the United States and other developed countries is phrased in terms of income, or in terms of components of income like wages and earnings. But for economists, a basic utility function of individuals typically refers to consumption and leisure, not income.

The distinction between income and consumption could make a meaningful difference in thinking about inequality if the distribution of consumption at a given point in time is less wide than that of income, or if its evolution over time is smoother than that of income. Consumption can differ from income if consumers borrow or save, or if they receive transfers from other family members or the government in response to income shocks. The joint analysis of consumption and income inequality can be informative in several ways. It can show the presence (or lack) of such consumption-smoothing mechanisms. It can shed light on the nature of income shocks, and in particular the extent to which they should be understood as temporary (which may be easier to smooth out for consumption

■ *Orazio P. Attanasio is Jeremy Bentham Professor of Economics, University College London, London, United Kingdom, and Research Fellow, Institute for Fiscal Studies, both in London, United Kingdom. Luigi Pistaferri is Professor of Economics and Ralph Landau Senior Fellow, Stanford Institute for Economic Policy Research (SIEPR), both at Stanford University, Stanford, California. Both authors are Research Associates, National Bureau of Economic Research, Cambridge, Massachusetts, and also Research Fellows, Centre for Economic Policy Research (CEPR), London, United Kingdom. Their email addresses are o.attanasio@ucl.ac.uk and pista@stanford.edu.*

[†] For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <http://dx.doi.org/10.1257/jep.30.2.3>

doi=10.1257/jep.30.2.3

purposes) or permanent. If one is interested in the effects of inequality on those in the poorest segments of society, consumption might reveal different insights than income—for example, because of different dynamics in the relative prices of goods consumed by rich and poor households. Finally, higher consumption of leisure could partly offset lower consumption of goods when it comes to overall welfare measurement.

In this essay, we begin with a discussion of the sources of consumption data, and some of the issues that arise when looking at data on consumer spending while trying to infer the economically relevant concept of consumption. We then offer our interpretation of the research that has compared trends in income inequality with trends in consumption inequality. The narrative has evolved very sharply from arguing that trends in consumption inequality are quite different from those of income inequality to concluding that they track each other closely. This change in findings has been shaped in substantial part by the adoption of strategies aimed at dealing with measurement problems in consumption data, as well as by some reinterpretation of the underlying economic forces.

We then discuss some additional aspects of consumption. We look at specific data on inequality in consumption of food, ownership of major household appliances, leisure, and persistence in consumption across generations. These comparisons suggest ways in which aggregate consumption inequality fails to tell the entire story. In the concluding section, we take stock of the evidence and summarize challenges for future work. Our main conclusion is that researchers interested in measuring inequality in well-being need to go beyond the fact that consumption is unequally distributed and realize that a full picture of the evolution in welfare requires taking a stand on quality concerns and on the value that people attach to leisure, among other things.

Consumption Data: Sources and Concepts

What Consumption Data Do We Have?

US researchers who want to study income inequality have a considerable array of data sources from which to choose. If they want household survey data on incomes, the Current Population Survey, Panel Study of Income Dynamics, Survey of Income and Program Participation, National Longitudinal Survey of Youth, and even the decennial Census (for studying long-run trends) all offer large, consistent samples and detailed information on income and its components. Researchers may also have access to administrative-level data (where measurement error issues may be less important with regard to income data), such as data from the Internal Revenue Service and data from the W-2 forms that employers use to report income paid. Overall, datasets with measures of household income resources (such as wages, earnings, and income) are more frequently available, typically have larger samples, and have more consistent variable definitions than datasets containing information on consumption (Pistaferri 2015).

In contrast, household surveys on household expenditure are rare, small, and lack a consistent longitudinal component. The Consumer Expenditure Survey (CE), the only dataset with comprehensive and detailed information on household expenditure and its components, is available on a continuous basis since 1980. It is used by the Bureau of Labor Statistics primarily to form weights placed on price changes of goods in the computation of the overall Consumer Price Index. The CE is composed of two distinct surveys. In the Interview survey, respondents are sampled every three months for a total of four quarters. In the Diary survey, respondents fill a two-week diary of their expenditures and are sampled only once. In producing aggregate means, the Bureau of Labor Statistics routinely uses and tabulates certain items from the Interview survey and others from the Diary survey.

The other dataset that is widely used by academic researchers to study consumption behavior is the Panel Study of Income Dynamics (PSID), which is available on an annual basis from 1968 to 1997, and on a biannual basis since then. The initial goal of the PSID was to study income dynamics (and poverty) between and across generations. For this reason, information on consumption was considered ancillary, and until 1997, the PSID collected information only on a few consumption items: food (at home and away from home), home rent, and (occasionally) utility payments. Starting with the 1999 wave, however, the PSID began to collect information on a larger range of items, covering about 70–90 percent of the spending collected in the Consumer Expenditure Survey. Respondents typically report spending for broad categories, with the reference period being (with some exceptions) the previous calendar year. Blundell, Pistaferri, and Saporta-Eksten (2016) show that these data track national accounts aggregates well.

Comprehensive administrative data on consumption spending are not available, but some partial sources do exist. For example, some researchers have used spending data from credit card expenditure records (as in Gross and Souleles 2002; Aydin 2015). Others have used data on spending, income, and assets for consumers using online financial aggregators such as Mint.com or Check.com (Baker 2014; Gelman et al. 2014). Finally, there is spending data from checkout scanners from the Nielsen Homescan datasets (Handbury 2014; Broda and Weinstein 2008), which refer primarily to grocery store items. While these new sources of administrative data on consumption constitute remarkable steps ahead, they are either not representative of all households or not representative of all the goods that people buy.¹

Looking at Spending, Thinking about Consumption

Consumption can be harder to measure accurately than income, and measurement errors may be differently severe in different parts of the income distribution.

¹For some Nordic countries, researchers have proposed using longitudinal administrative tax record information on income and wealth to create consumption using the intertemporal budget constraint, so that expenditure can be derived as income minus the change in wealth; see Browning and Leth-Petersen (2003) and De Giorgi, Frederiksen, and Pistaferri (2015) for Denmark; Koijen, Van Nieuwerburgh, and Vestmanz (2015) for Sweden; and Autor, Kostøl, and Mogstad (2015) for Norway.

Survey data like the Consumer Expenditure Survey typically report consumer spending, which may not coincide with consumption for at least four reasons. First, consumption is overstated relative to spending for those who have made durable purchases in the current period, and understated for those who made durable purchases in the past. Most surveys of household consumption have no information on the stock of existing durables owned by the household. In the Consumer Expenditure Survey, the only exception is cars, as consumers report the year, type, and make of the cars they own. However, there is no information on the resale values, which must instead be estimated. For other durables, there is some information on ownership and number of items owned, but no information on current values. As for housing, the survey contains information on imputed services, as homeowners report how much their house would rent for. Second, some consumption is received in kind, through transfers from friends, relatives, private institutions (charities or churches), or the government (in the form of in-kind or voucher-provided benefits like food stamps, school lunches, health care services through Medicaid or Medicare, rent subsidies, and so forth). Third, some consumption is produced at home using time and good inputs, like child care provided by parents or siblings. Finally, the conversion from spending to consumption requires knowledge of prices paid for the goods people consume, a requirement that is usually solved by assuming that households face the same prices. This assumption is violated in practice for many nontradeable goods (such as housing), and even for tradeable, homogenous goods due to shop-specific effects, bulk purchases, or loyalty cards. Ignoring the distinction between consumption and spending can either understate or overstate differences in well-being across individuals with different levels of spending.

Survey nonresponse and measurement errors create a different set of issues. Sabelhaus et al. (2015) study nonresponse rates in the Consumer Expenditure Survey and conclude that they have risen over time, especially among the high-income population. Sabelhaus and Groen (2000), using a variety of techniques, argue that the ratio of consumption to income for richer households is downward biased. This may affect the measurement of trends in aggregate consumption and consumption inequality, respectively. In this journal, Meyer, Mok, and Sullivan (2015) conclude that measurement error in survey data has also increased over time. In principle, there is no obvious reason to expect errors in reporting spending to be worse than errors in reporting income. On the one hand, the changing nature of spending modes and patterns may heighten reporting errors; on the other hand, the move from cash to e-commerce may facilitate the collection of administrative data on spending. Several papers discuss strategies to elicit consumption information in general-purpose surveys (for instance, Crossley and Winter 2015).

The combination of these issues makes any measure of consumption naturally problematic, no matter how much effort researchers put into making it accurate. In contrast, income measured in surveys is arguably closer to the relevant economic concept (except in cases involving businesses). However, we should also note that income is not easier to measure than consumption for every household. For low-income individuals (in both developed and developing countries) income can

be complex to measure, because it includes a myriad of different sources including wages, interpersonal and government transfers, and so on. In comparison, consumption for the poorest households may be fairly straightforward. The situation is probably reversed for well-off households for which administrative income data might be accessible and reliable, while their consumption can be complex and varied and difficult to measure (Deaton and Grosh 2000).

Survey Data versus National Accounts Data on Consumption

Consumer spending data available in the Consumer Expenditure Survey appears increasingly detached from the Personal Consumption Expenditure (PCE) data collected by the Bureau of Economic Analysis (which forms the basis for the national income and product accounts data). For example, Passero, Garner, and McCully (2015) report that the ratio of total expenditures in the CE data compared to the PCE data has declined from 0.70 in 1992 to 0.58 in 2010. Of course, whenever two different methods of measuring a similar economic concept give different answers, it's a matter for concern.

Some of the discrepancy is due to the fact that the two series measure different concepts and cover different entities. The Personal Consumption Expenditure data includes the value of goods and services purchased by US resident households (including imputed rents for owner-occupied housing) and by nonprofit organizations on behalf of households (typically, employer-paid health insurance and medical care, and expenses associated with life insurance and pension plans). It also includes purchases by US government civilian and military personnel stationed abroad and US residents traveling or working abroad for one year or less. For most consumption categories, the PCE is estimated using a “commodity-flow” method. This approach computes the value of domestic output based on data from the Census of Manufactures, which looks at the value of manufacturers’ shipments and inventories. Next, domestic consumption (denominated in producers’ prices) is estimated by adding imports and subtracting exports and changes in inventory. Finally, the value of consumer purchases is converted from producers’ prices to purchasers’ prices by adding wholesale margins and taxes, transportation costs, and retail margins, and taxes. Clearly, many steps in the calculation of the PCE are also likely to contain sizeable measurement error.

Both the population coverage and the methodology for obtaining total spending is different in the Consumer Expenditure Survey (as discussed in Slesnick 1991). The Personal Consumption Expenditure data includes institutionalized individuals; the Consumer Expenditure Survey does not. The Consumer Expenditure Survey excludes spending made by US residents abroad and by nonprofit institutions on behalf of households (with the most obvious difference being the value of Medicare and Medicaid spending, which tripled in real terms between 1990 and 2014). The PCE concept includes imputed rents on owner-occupied housing, while the Consumer Expenditure Survey aggregates typically exclude them. Indeed, the discrepancy between the two measures is less dramatic when comparing items that are conceptually comparable and definitionally similar. Passero, Garner, and

McCully (2015) compare different components of consumption and conclude that “non-durables are most alike for the CE and PCE with about 93 percent of total non-durable expenditures identified as comparable within the CE and within the PCE.” Their conclusion is that “focusing on comparable goods and services only, CE to PCE ratios have steadily decreased,” but slightly less than when comparing unadjusted statistics. For example, the CE–PCE ratio for total comparable goods and services decreased from 84 percent for 1992 to 74 percent for 2010 (as opposed to 70 and 58 percent, respectively, in unadjusted figures). They write: “The greatest decline in CE to PCE ratios is for durables, with a decrease of 24 percentage points,” from 0.82 to 0.62. The decline is smaller for services (0.95 to 0.86) and for nondurables (0.70 to 0.63).

Bee, Meyer, and Sullivan (2015) assess the performance of the Consumer Expenditure Survey on a good-by-good basis and report three findings. First, in general the Interview survey performs better than the Diary survey in matching numbers from the Personal Consumption Expenditure data for some categories. Second, the coverage ratios are excellent for some goods (food at home, rent, and utilities) and, in those areas, have not changed appreciably over time; on the other hand, the coverage ratios for other items (such as clothing or alcoholic beverages) are low and declining. Finally, some durable stocks and durable purchases appear to be reported sufficiently well (new vehicles), while the quality of others has worsened considerably (furniture). Overall, they conclude that the consumption categories that tend to be reported poorly are those that involve small and infrequent purchases, while large and regular purchases are reported sufficiently well. Moreover, they write that “based on observable characteristics, the [Consumer Expenditure Survey] appears to be fairly representative, although there is strong evidence of under-representation at the top of the income distribution and under-reporting of income and expenditures at the top.” This is another reason for the growing discrepancy between the CE aggregates and the PCE from the National Accounts. If consumption growth is higher at the top of the distribution, declining survey response among the rich may easily explain the deterioration of the match between CE aggregates and the PCE.

Quantity versus Prices

Survey data like the Consumer Expenditure Survey measure expenditure, which is the product of prices and quantities. To make comparison of consumption across periods meaningful, researchers deflate expenditure using an overall measure of the cost of living, typically the Consumer Price Index (CPI). Indeed, as mentioned above, the CE is collected primarily to compute the weights for the CPI. However, average weights may not be relevant for all households. It is possible that the composition of the consumption basket varies substantially and systematically across different households, as a direct consequence of differences in access to resources as well as needs and tastes. Luxuries will be more prominent in the expenditure basket of the rich, while necessities will account for a larger share of poor households’ expenditure. Health costs may be more relevant for older individuals

and certain types of entertainment more relevant for younger ones. Therefore changes in relative prices can have distributional consequences.

Moreover, there may be price differences across space (or stores within a given geographical location), or across time within space (because of high-frequency sales) even for relatively homogenous goods. Because of differences in prices over space, consumers might have an incentive to search for the best deals and these incentives may vary for individuals with different values of time.

The presence of differences in prices for homogeneous goods and changes in the relative prices of goods that are more or less relevant for different groups of individuals might lead researchers to overstate or understate the level and trends of inequality. If the poor live in relatively cheaper areas (or if they shop in relatively less-expensive stores), or if the prices for the goods that they typically purchase grow less than the prices of the goods typically purchased by rich households, then inequality in consumption (that is, spending deflated by an index that accounts for household-specific price differences) will be less (and grow more slowly) than inequality in spending deflated with a common price index. There are two reasons why this may be the case. First, increasing trade with low-wage countries lowers the price of imported goods. This may reduce consumption inequality even in the absence of any change in income inequality if goods imported from low-wage countries are relatively more important in the consumption basket of low-income individuals than in the basket of high-income individuals. Moreover, the diffusion of mega-stores (such as Walmart) has likely benefited low-income individuals more than high-income individuals. These differences are partly attenuated by the consideration that better quality and the experience of shopping in certain stores have an amenity value.

Datasets where researchers can disentangle the two components of expenditure are hard to come by. The Nielsen Homescan data is one exception, but it is limited to groceries and a few other items. Other data sources (such as the ACCRA Cost of Living Index produced by the Council for Community and Economic Research) provide city-specific indexes on a few categories of interest. In the Consumer Expenditure Survey, the geographical detail is very limited due to confidentiality concerns (and in any case, it would miss information on the type of store where goods are purchased). Later in the paper, we discuss some recent work on the implications of price inequality.

Does Consumption Inequality Track Income Inequality?

Consumption Smoothing: Why and How

Is consumption inequality a better measure of changes in welfare than income inequality? For economists, using consumption inequality has theoretical appeal. The life-cycle hypothesis of Modigliani and Brumberg (1955) and the permanent income hypothesis of Friedman (1957), which constitute the workhorse theory of how people make their consumption decisions, suggest that risk-averse households

prefer a smooth to a variable consumption flow. Hence, households would choose consumption to be a constant fraction of their permanent or lifetime income, not current income. Because current income can be highly volatile from one year to the next, it may give a partial snapshot of people's living standards. The extent to which households can achieve a smooth consumption flow depends on the tools they have to move resources over time and states of nature. Savings can be used to absorb certain income shocks and can be accumulated for such a purpose. Other tools for consumption smoothing may include access to credit and insurance markets, and interpersonal and government transfers.

The ability to move resources across time and states explains why consumption may not track income. Consumption may exceed current income because a consumer is borrowing (permanent income is above current income, as in the case of a medical student taking out loans in the expectation of higher future earnings) or it may be below current income because the consumer is saving (and the doctor is now repaying medical school loans). Large wealth effects can also have a considerable influence on consumption independently of income. It is then possible for the income distribution to reveal no changes in well-being even though the underlying consumption distribution is shifting in response to wealth effects. Consumption may vary from income for other reasons as well. For example, consumption falls below current earnings and wages because of taxes paid, and above them because of government transfers—a different form of consumption smoothing especially relevant for households at the bottom of the distribution. Even if *full* smoothing is not feasible, perhaps because of borrowing restrictions and imperfections in insurance and credit markets, some consumption smoothing would still occur. Recent surveys on consumption (or marginal utility) smoothing are Browning and Crossley (2001) and Attanasio and Weber (2010).

These considerations imply that how consumption (and welfare) react to changes depends on the tools available for consumption smoothing and on the nature of income changes. In general, we can think of current income as including two components: one reflects long-term or permanent factors such as the level of skills and human capital, while the other reflects temporary or transitory factors (like being out of the labor force due to human capital investments, fertility, job loss, and the like). For any given household, permanent income changes slowly over time. When it does, it is because of unpredictable events, like the way in which new technologies affect the price and quantity demanded of one's skills (for better or worse), along with unexpected promotions or demotions, a change in the local economy that affects wage levels, and other factors. MaCurdy (1982) and Abowd and Card (1989) are two representative papers in a vast labor literature that tries to decompose income (or earnings, or wages) into transitory and permanent components.

The distinction between temporary and persistent shifts in the wage or income distribution has important policy and welfare implications. Policies aimed at reducing inequality under the two scenarios are very different. In the first case, it is probably necessary to reduce inequality in the endowments of

human capital, whilst in the latter it may be sufficient to improve the access to smoothing mechanisms. In theory, permanent shocks are harder to absorb and insure and are thus more likely to be reflected in substantial changes in consumption and welfare. In contrast, temporary shocks are easier to smooth through borrowing or running down accumulated assets. Hence, if all changes in income inequality were of a transitory nature, we could expect no large changes in consumption inequality.

The Evidence on Consumption and Income Inequality

The mainstream narrative on consumption inequality has evolved considerably over time, from earlier uncertainty over whether consumption was rising less than income inequality to the current belief that it has been rising just about as much as income inequality.

The first few papers in the earlier literature that looked at different dimensions of inequality in consumption were Cutler and Katz (1991, 1992) and Slesnick (1994). These papers used the Consumer Expenditure Survey data (in an era in which the measurement error issues discussed above were less well-known than they have since become). Cutler and Katz (1991) found that “changes in the distribution of consumption parallel changes in the distribution of income.” In contrast, Slesnick (1994) found that consumption inequality had grown more modestly than income inequality.² A number of later studies (for example, Krueger and Perri 2006) found evidence similar to Slesnick (1994). In the terminology of consumption smoothing and the permanent income hypothesis, this finding can be interpreted as implying that a sizeable proportion of shocks to income were both temporary and insurable. Evidence on this comes from two sources: direct evidence on the income process, and indirect evidence coming from the response of consumption to income changes of different natures.

On the first front, researchers working on income inequality were finding strong evidence that the rise in income inequality was partly of transitory nature—which in turn implies that some portion of the rise in inequality could be more smoothed in consumption. Gottschalk and Moffitt (1994) argued that the rise in the variance of the transitory component of income (what they called “income instability”) represented about one-third to one-half of the overall rise in income inequality observed in the 1980s and 1990s. (Income instability is often measured using the variance of income changes, or growth. This is because the change in income across two periods is approximately equal to the change in the transitory component if permanent income evolves along the expected path.)

The fact that a good chunk of the rise in income inequality was of a transitory nature would not matter much (in terms of separating income from consumption inequality) if consumers were unable to smooth transitory shocks. However, several

²The differences between the two studies arise partly from their different consumption definitions. Cutler and Katz (1991) include spending on durables (other than housing and vehicles), while Slesnick (1994) imputes service flows. Moreover, Slesnick adjusts for topcoding in some spending categories.

papers show that consumers are able to smooth short-run shocks (Dynarski and Gruber 1997), although less so if they have low assets or low education (Blundell, Pistaferri, and Preston 2008). Attanasio and Davis (1996) focused on the relationship between relative wages and relative consumption across different groups in the US population, where groups were defined on the basis of the year of birth of the household head and on their educational achievement. They found that long-run relative movements in wages across these groups were mirrored in relative movements in consumption. This correlation was driven by the relative movements across education groups: the increases in the return to education in terms of wages and earnings seemed to be reflected in increases in the return to education in terms of consumption. The online appendix available with this paper at <http://e-jep.org> contains an update of Attanasio and Davis (1996), extending the data to 2012. As in the original paper, we find that when we consider the impact on consumption of one-year wage changes, the variability of which is probably dominated by temporary fluctuations in wages that can be smoothed in some way, we do not find a significant relationship between relative changes in wages and consumption. However, when considering longer (five- and eight-year) horizons, where instead persistent wage factors are more likely to be at play, the relationship between changes in consumption inequality and income inequality becomes strongly significant.

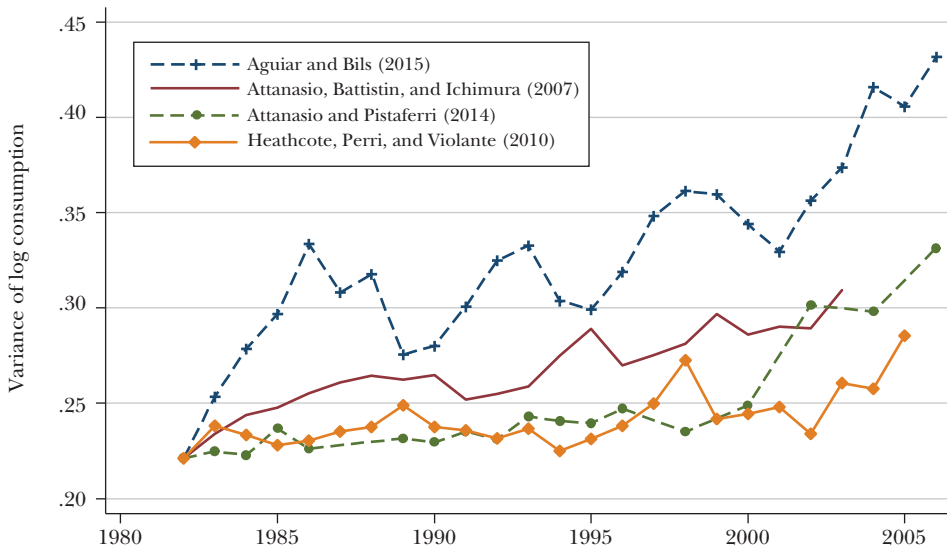
In keeping with these two pieces of evidence, Krueger and Perri (2006) show that while in the 1980–2003 period the variance of log income increases from 0.35 to 0.57, the variance of log consumption increases only from 0.18 to 0.24. In other words, inequality rises in both income and consumption, but the rise in income inequality is much larger.

More recently, however, researchers have started to question the evidence about consumption inequality, rethinking the measurement issues that arise from considering measures of expenditure in the Consumer Expenditure Survey. The survey seems to be affected by serious nonclassical measurement error whose importance is increasing over time. One possible strategy is to focus on the components of the Consumer Expenditure Survey that appear to be measured most accurately and to use alternative datasets for other components of consumption. The challenge of course is that one would like to make statements about inequality in overall consumption, not necessarily about inequality in some components. Once one corrects for the measurement problems afflicting the Consumer Expenditure Survey, uses alternative data sources when these seem preferable, or measures consumption in alternative ways, consumption inequality seems to rise by more than previously believed, and to track income inequality closely.

A number of papers develop this view using a variety of data sources and empirical approaches (Attanasio, Battistin, and Ichimura 2007; Aguiar and Bils 2011; Attanasio and Pistaferri 2014). Figure 1 gives an overall view of the evolution of consumption inequality over time and across papers (and empirical strategies). In this figure, consumption inequality is measured by the variance of log consumption (deflated by the Consumer Price Index and expressed in per capita terms).

Figure 1

The Evolution of Consumption Inequality over Time as Measured by Different Papers



Note: Heathcote, Perri, and Violante (2010) used the Interview survey of the Consumer Expenditure Survey. Attanasio, Battistin, and Ichimura (2007) combined consumption items from the Interview survey and the Diary survey (in the attempt of picking the survey component that best measures each item). Aguiar and Bills (2015) used the Consumer Expenditure Survey but computed consumption as the difference between disposable income and active savings. In Attanasio and Pistaferri (2014), we used Panel Study of Income Dynamics consumption data available from 1999 onward, estimated an inverse demand function for food for the 1999–2009 period, and then used the estimated coefficients to predict consumption for the period before 1999 (when only food data were available). Consumption inequality is measured by the variance of log consumption (deflated by the Consumer Price Index and expressed in per capita terms).

Heathcote, Perri, and Violante (2010) used the Interview survey of the Consumer Expenditure Survey, and their findings reproduce the flat profile of consumption inequality shown by Krueger and Perri (2006). Attanasio, Battistin, and Ichimura (2007) combined consumption items from the Interview survey and the Diary survey (attempting to pick the survey component that best measures each item), with results showing a more marked increase in inequality. Aguiar and Bills (2015) used the Consumer Expenditure Survey but computed consumption as the difference between disposable income and active savings, and they find an even larger increase in inequality. Finally, in Attanasio and Pistaferri (2014), we used Panel Study of Income Dynamics consumption data available from 1999 onward, estimated an inverse demand function for food for the 1999–2009 period, and then used the estimated coefficients to predict consumption for the period before 1999 (when only food data were available). We found that inequality also increases more

than the Heathcote, Perri, and Violante measure, especially in the last years of the sample period.³

To obtain a sense of how much consumption inequality grows relative to income inequality, and how the response depends on the methodology used to measure consumption, consider the following calculation. Over the period considered in the figure, the variance of the log of family income from the PSID (deflated by the Consumer Price Index and expressed in per capita terms) increases by 27 points (or about 20 points when using an after-tax measure available in the Consumer Expenditure Survey, as reported by Heathcote, Perri, and Violante 2010). If we take the Aguiar and Bils measure of consumption inequality shown in Figure 1 as the most credible, the variance of log consumption increases by about 18 points over the same time period. In contrast, the Heathcote, Perri, and Violante measure would suggest an increase of only about 10 points. Meyer and Sullivan (2013) show that the tracking between income and consumption inequality is stronger at the top of the distribution (as measured by the 90th–50th percentile difference) and in the 1980s and 1990s than in subsequent years. Aguiar and Bils’s (2015) core exercise is actually to measure consumption inequality by looking at how high- and low-income households allocate spending to luxuries and necessities. In particular, inequality in the luxury/necessity spending ratio (scaled by the difference in demand elasticities, which can be obtained from estimation of a demand system) is shown to provide a measure of consumption inequality that is robust to measurement error in overall spending, as well as to household-specific measurement errors (for example, more severe underreporting by high-income households) and good-specific measurement errors (more severe underreporting for some goods than others). Using this alternative metric, Aguiar and Bils confirm that over the 1980–2007 period, inequality in consumption grows as much as income inequality.

The common element of the papers above is that once one makes an attempt to move away from the traditional measurement of consumption inequality using the Interview component of the Consumer Expenditure Survey, and tries to correct for the measurement problems, then the trends in consumption inequality appear much steeper than initially believed. Of course, the conclusion reached by these papers may also be premature, because the strategies adopted, while ingenious, are based on data that may have different types of measurement problems.

One aspect that seems to militate in their favor, however, is that the change in the consensus about the trends in consumption inequality has been accompanied by changes in the consensus on the evolution of income inequality, based on

³An important caveat is that the consumption series used in the four papers are not identical. For example, the Attanasio and Pistaferri (2014) measure is limited by the fact that the PSID collects a limited amount of information on expenditure, while the Aguiar and Bils (2015) measure, by definition, does not use any consumption information. The Heathcote, Perri, and Violante (2010) and Attanasio and Pistaferri (2014) measures include out-of-pocket spending on health and education, while the Attanasio, Battistin, and Ichimura (2007) measure excludes them. The differences between the series should thus be seen as illustrative. We normalize all series to equal the Heathcote, Perri, and Violante (2010) value in 1982 (the first year in which we observe all four series).

improved income data. Recent work using administrative data about income—which is less prone to measurement error issues than survey data—finds that most of the increase in wage and earnings variance has been structural, or of a more permanent nature (for example, DeBacker et al. 2013; Kopczuk, Saez, and Song 2010; Guvenen, Ozkam, and Song 2014). Kopczuk, Saez, and Song (2010) use Social Security data over a very long time horizon and present a formal decomposition between total, transitory, and persistent earnings variances. Using this decomposition, the rise in total variance during the period of interest is primarily driven by a rise in structural factors. In contrast, there is very little evidence of a rise in the variance of the transitory component.

These recent findings about the nature and dynamics of income inequality are consistent with the revised thinking in the dynamics of consumption inequality. If income volatility is stable, it means the variance of the transitory component has not increased. Hence, the bulk of the change in income inequality has occurred because of a rise in the variance of the permanent component. It is possible for consumption inequality to rise less than income inequality even in a setting in which income volatility is stable. This is because consumers may be able to insure even some shocks to their permanent income, at least partially. For example, the Disability Insurance program seeks to attenuate the economic cost of permanent shocks to health that result in permanent inability to work. But it is clearly more difficult to smooth changes in permanent income, and as a consequence, it is not surprising that consumption inequality rises by roughly as much as income inequality. Indeed, their rise is explained by the same forces (absent strong insurance mechanisms).

Inequality in Prices

As mentioned earlier, recent research has started to look at data on individual purchases and has documented the existence of important heterogeneity in prices of even very homogenous goods, both in different stores and within a store over short periods of time, through the use of sales and discounts. One of the first papers to document the existence of substantial heterogeneity in the prices of very homogeneous goods is Aguiar and Hurst (2007). They correlate observed prices and shopping behavior with consumer characteristics. They show that older consumers are more likely to shop longer and more frequently and, probably as a consequence, pay lower prices for similar goods. Griffith, Leibtag, Leicester, and Nevo (2008) use British scanner data to show that low-income households realize considerable savings by buying in bulk and by buying economy brands, while savings from sales, coupons, and the like are nonlinear in income—specifically, higher at the top and bottom of the income distribution. More recently, Nevo and Wong (2015) show that during the Great Recession, consumers switched to buying more on sale, using more coupons, buying more generics and larger pack sizes and these changes were larger in states that suffered larger increases in unemployment rates.

Kaplan and Menzio (2015) use scanner data on a large sample of US households covering grocery stores purchases in 54 geographical markets over the 2004–2009 period. They find that the distribution of prices is symmetric and with fatter tails

than the normal distribution, and its average standard deviation is between 19 and 36 percent. They also show that, when decomposing the variability of prices of homogeneous goods into a store component, a store-good component, and a transaction component (the dispersion of prices within a store), most of the variability of prices in their sample is explained by the latter two factors. They suggest that price dispersion is more likely to be driven by intertemporal price discrimination and search frictions than differences in amenities or marginal costs across stores. This hypothesis is explored more fully in Kaplan, Menzio, Rudanko, and Trachter (2016).

The extent to which differences in prices actually paid affect the dynamics of consumption inequality, either through differences in consumer baskets or through price heterogeneity induced and sustained by frictions and retailer behavior, is an open question and one of considerable interest. The availability and use of scanner data on individual transactions can be very useful in this respect, as is the development of models that incorporate price discrimination and frictions in price-setting behavior. More broadly, the measurement of consumption and income inequality is a lively area of research. The existing work is undoubtedly subject to improvement as better data or more creative approaches to overcome measurement issues come along.

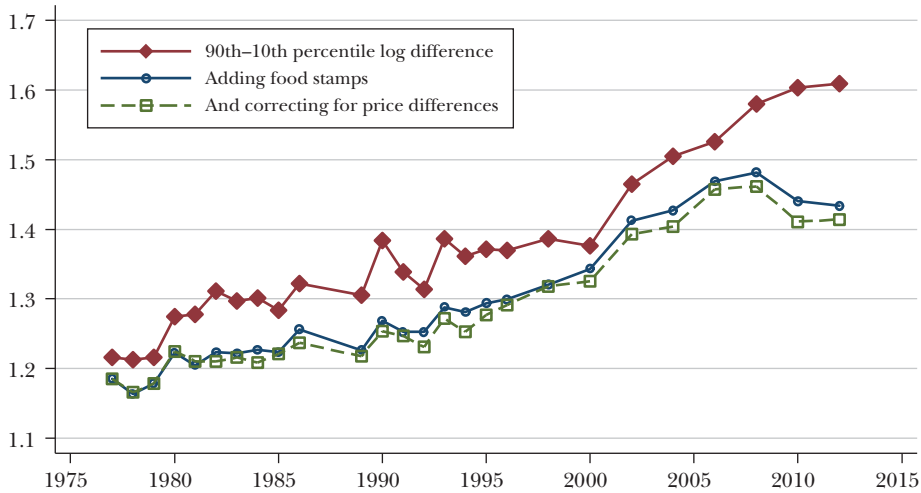
Inequality in Components of Consumption

Looking at inequality of consumption across specific components of consumption may be interesting for several reasons. First, the measurement of some components of consumption is of better quality than others, thus alleviating concerns about whether results are affected by measurement error. Second, the analysis of different groups of commodities with different income elasticities can be informative about the nature of shocks and about mechanisms for smoothing consumption. Third, changes in the patterns of expenditure on durables can be informative about the perception of future shocks, because individuals know that commodities such as furniture or cars provide services for long periods and can be sold only subject to large transaction costs. Finally, disparities in consumption necessities such as food may be more worrying from a welfare point of view than disparities in the consumption of luxuries, such as exotic vacations.

Food

In Figure 2, we use data from the Panel Study of Income Dynamics and plot the difference between the 90th and the 10th percentile of the logarithm of food consumption distribution over the 1977–2012 period. Food consumption is defined as the sum of spending on food at home, food away from home, and the monetary value of food stamps. Data are in real terms and adjusted for family composition by dividing by an OECD scale (defined as $\$1 + 0.7(n - 1) + 0.3k$, where n is the number of adults and k the number of kids). PSID food data exist before 1977, but it is only in 1977 that the Food Stamps Act established national standards of eligibility.

Figure 2

The 90th–10th Percentile Log Food Difference

Source: Authors using data from the Panel Study of Income Dynamics.

Note: This figure plots the difference between the 90th and the 10th percentiles of the logarithm of food consumption distribution from 1977 to 2012.

Our sample includes all households whose head is aged 25–85. The sample includes the poverty subsample of the Panel Study of Income Dynamics and hence sampling weights are used throughout. To emphasize the distinction between consumption and spending (which in the case of food may be particularly relevant due to government transfers), we plot the 90th–10th percentile difference both including and excluding food stamps from our definition of food consumption.

Clearly, inequality in food consumption is rising. Most of the rise is coming from a decline in spending at the bottom (not shown separately here). The difference between the top and intermediate lines shows clearly the insurance value of government transfers. In particular, during the Great Recession the availability of food stamps allowed poor households to maintain their food consumption, while spending declined substantially.

We should also note that some of the lower spending on food by the poorest households may be due to a decline in the prices of the food items that they purchase (Broda, Leibtag, and Weinstein 2009). To have some sense about the importance of price differentials, we also plot the 90th–10th percentile difference allowing the price deflator to be good-specific (that is, food at home plus food stamps, and food away from home). Correcting for price differentials has a small effect, although it is more pronounced in recent years. Because the price of food at home (a necessity consumed in large fractions by households at the bottom of the distribution) has been steadily declining relative to the price of food away from home (a luxury

consumed in large fractions by households at the top of the distribution), inequality in consumption is lower when adjusting for these price differences.

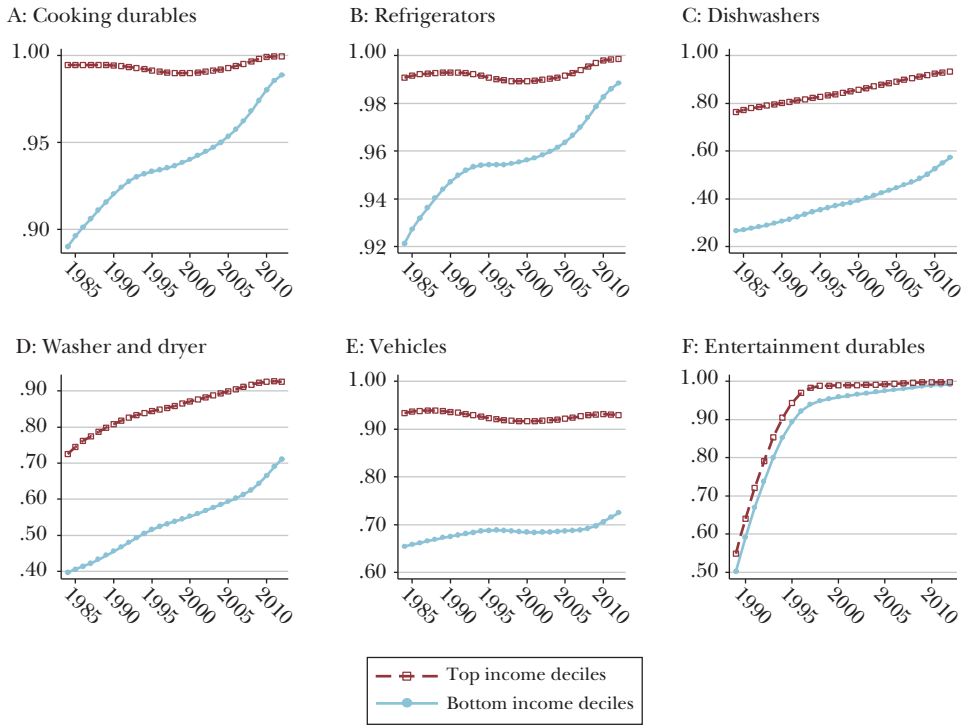
The decline in spending on food consumption at the bottom of the distribution may not indicate a decline in caloric intakes. Households may spend less on food without modifying the caloric intake of the food purchases they make (as argued by Aguiar and Hurst 2005). Indeed, Singh et al. (2009) report that energy intake is not statistically different between US adults with income below the poverty line and those with income above 500 percent of the poverty line. After all, sugars and fats, which are high in calories, can be considerably less expensive than diets based on vegetables, fruits, whole grains, and lean meat.

The different qualities of food raise a question: Should an assessment of inequality in food consumption be based on its monetary cost, its energy content, its healthfulness, or some other measure of quality? The US Department of Health and Human Services and the US Department of Agriculture have proposed to measure diet quality with an index known as the Healthy Eating Index (HEI). The index gives a 0–10 score to 12 food components (like Total Fruits, Whole Fruits, and so on). For some “good” components (like Total Vegetables) a higher intake means a higher score, while for some “bad” components (like Saturated Fat) the opposite is true. Wang et al. (2014) use data from the 1999–2010 National Health and Nutrition Examination Survey, and compare the HEI index for people of different socioeconomic background and education. They find that in the population at large the quality of food consumed increases monotonically over the sample period. However, individuals with low socioeconomic status (defined by those with less than high school and income below 130 percent of the poverty line, the eligibility threshold for food stamps) make no progress in terms of HEI from 1999 to 2010, while most of the improvements are concentrated among medium and high socioeconomic status groups. We want to stress that while these differences reflect changes in the “quality” of food consumed between rich and poor individuals, they are silent regarding the reasons. One possibility is that tastes for healthy food changed differently for rich and poor individuals (or that the rich were more receptive or attentive to “eating healthy” campaigns). Another possibility is that salience was similar but the higher price of healthy food or its lower availability in poor neighborhoods represent significant “barriers to entry” in healthier eating habits for poor individuals.

Durable Goods

Consumer durables reflect an element of standard of living that may not be captured by current spending (as people buy them infrequently). Hence, another way to look at consumption inequality is to see how many and what type of households own certain home appliances and durable goods. Also in this case—as we did in the discussion of food consumption—we stress that the quality of what is being consumed or purchased can matter substantially in thinking about the welfare consequences of inequality.

Figure 3

Share Owning Durables in the Top and Bottom Income Deciles

Source: Authors using data from the Consumer Expenditure Survey

Note: For different categories of durables, the figure compares ownership rates of the bottom and the top after-tax income deciles.

Given that the Panel Study of Income Dynamics has no information on durables except cars, we use the Consumer Expenditure Survey, which contains consistent series on durable goods ownership over a long period time. For some appliances like refrigerators, washing machines, and others, “availability” is probably a more appropriate term than “ownership” if such items are attached to the housing unit. Figure 3 offers a comparison of ownership rates for the bottom and the top after-tax income deciles. For some categories, we have a long series from 1984 to 2012; for others, the series starts in 1989. For most categories, there is evidence of catching up. For example, at the beginning of the time period, ownership of cooking durables (stoves or microwaves) and refrigerators is almost universal among the top 10 percent households, while the proportion of households in the bottom income decile owning such appliances is below 90 percent. For refrigerators at the start of the period, the difference is less but still noticeable. By 2012, these differences have largely disappeared. While there is convergence for these categories, the catch-up rates for dishwashers and for washers and dryers are much slower. Ownership of cars has also converged, albeit at a slower rate than food-related appliances. Finally,

there is only a small difference in the fractions owning entertainment durables (TVs, sound systems, DVD players, PCs, and so on) throughout the entire period.

There are two caveats to these findings. First, a convergence in ownership does not imply convergence in the number of appliances owned. Indeed, for the durables for which this information is available (vehicles and entertainment), there is no evidence of convergence. Moreover, there may be a large quality difference between high-end and low-end appliances, but the existing data are not rich enough to measure the quality of the durables owned by socioeconomic status.

Inequality in the Consumption of Leisure

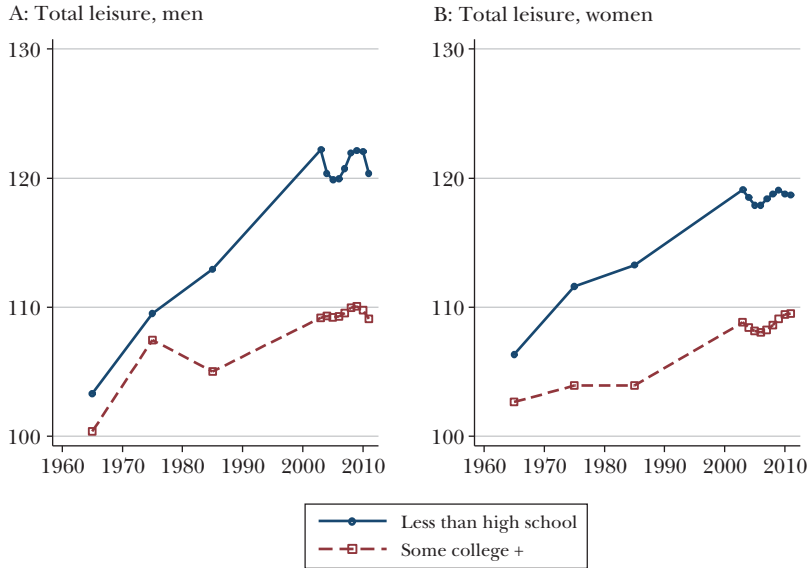
Economists traditionally write the utility function of individuals as comprising consumption and leisure. Perhaps greater inequality in the consumption of goods and services is being partially offset by greater equality in leisure time?

Measuring leisure is complex. Aguiar and Hurst (2009) have looked at trends in time use as a way of measuring leisure time. We follow a similar strategy here. In particular, we use surveys collecting information on time use over the last 50 years: the 1965–66 Americans' Use of Time, the 1975–76 Time Use in Economics and Social Accounts, the 1985 American Use of Time, and the 2003–11 integrated American Time Use Survey. The datasets do not include detailed or consistent information on income or on other measures of economic resources. Thus, we will use education as a rough measure of socioeconomic status.

To display the sharpest differences, we consider only the top and bottom education categories: individuals with less than a high school degree, and those with at least some college. In these datasets, people report the number of minutes they spend in various activities in the previous 24 hours. The main time use categories are: “work” (including time spent searching for jobs), “chores” (all household activities such as cooking, cleaning, and others), “child care,” “social” (watching sports, going to movies, partying, and so on), “organizational” (for example, volunteering, religion), “personal care” (sleeping, eating, and so on), “shopping,” “education,” “active leisure” (sport activities, playing games, and others), and “passive leisure” (watching television, listening to radio, relaxing, and others). All figures are weighted with sampling weights (except 1965–75 where no weights were released) and expressed in hours per week.

In Figure 4, we plot trends in total leisure time—the sum of social activities, active and passive leisure, and time devoted to personal care—controlling for the day of the week the diary was filled in. This measure of leisure may not be without problems. For example, it includes time spent assisting or helping adult household members (which for some people may represent “chores”); it excludes gardening or cooking (which for some individuals may represent a form of leisure). However, excluding personal care does not affect the main trends. The less-than-high-school education group has more leisure, and in the last few decades most of the growth of leisure has happened for this group, too. Moreover, growth has been stronger for men than women.

Figure 4

Trends in Total Leisure Time

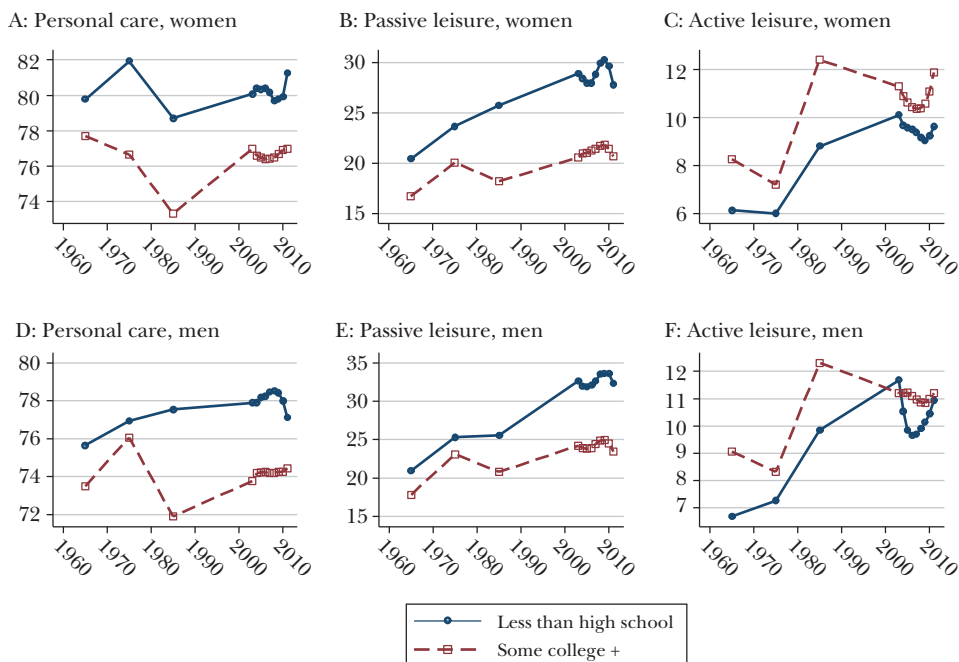
Source: Authors using data from the 1965–66 Americans’ Use of Time, the 1975–76 Time Use in Economics and Social Accounts, the 1985 American Use of Time, and the 2003–11 integrated American Time Use Survey.

Note: In the figure, we plot trends in total leisure time—the sum of social activities, active and passive leisure, and time devoted to personal care.

How do we interpret these trends? It is tempting to argue that the rise in consumption inequality has been to some extent counterbalanced with increasing leisure time among the poor, implying that the increase in the inequality of well-being is less severe than what consumption data alone may suggest. But any such conclusion needs to be hedged around with cautions. Some of the increase in leisure is involuntary, due to lack of job market opportunities. Excluding recession years from the analysis gives the same broad picture of Figure 4. Moreover, if we repeat the analysis only for the employed, we find that the differences are smaller but still significant (especially in the 2000s). Of course, an analysis that conditions on employment does not solve the problem of making welfare comparisons across income groups that include leisure, as employment itself results from and is affected by a combination of supply and demand factors. It is also possible that there is substantial heterogeneity in preferences for leisure—which may also help to explain different educational choices in the first place.

What component of leisure time is driving these trends? In Figure 5, we decompose total leisure into three components: personal care, active leisure (plus social activities), and passive leisure. There are some notable trends. First, for both low-educated women and (especially) men, the increase in total leisure time visible from

Figure 5
Decomposing Total Leisure
(in hours per week)



Source: Authors using data from the 1965–66 Americans' Use of Time, the 1975–76 Time Use in Economics and Social Accounts, the 1985 American Use of Time, and the 2003–11 integrated American Time Use Survey

Note: Here we decompose total leisure into three components: personal care, active leisure (plus social activities), and passive leisure.

Figure 4 is coming primarily from an increase in time devoted to passive leisure activities. Second, time spent on active leisure and social activities also increases, resulting in greater similarity between high- and low-educated individuals. Finally, time spent on personal care is stable. It is possible that these changes represent an evolution of preferences for leisure across education (income) groups. It is also possible that increasing availability of durable goods in the lower-income groups (documented earlier) frees up time previously devoted to housework.

Consumption Mobility

While there is a popular image of the United States as a land where high rates of mobility across the income distribution are possible, in practice intergenerational income mobility has not changed much over the last 40 years (Chetty, Hendren,

Kline, and Saez 2014). In fact, some European countries display more intergenerational income mobility than does the United States (Black and Deveraux 2011). There are also vast geographical differences in mobility across US regions.

Do the trends in intergenerational mobility in consumption mirror those found for income? To study this topic, we need longitudinal information on consumption that follows multiple generations. The Panel Study of Income Dynamics offers such data. In particular, for each household, it follows the children, the “splitoff” households, when they leave the parental home. A few authors have looked at the intergenerational dimension of the PSID data in the context of risk-sharing within the family (for example, Hayashi, Altonji, and Kotlikoff 1996; Attanasio, Hurst, and Pistaferri 2015).

We first construct a measure of household consumption using the Panel Study of Income Dynamics data. Specifically, we define consumption in this data as the sum of spending on food, rent, health, home insurance, utilities, car insurance, car repair, gasoline, parking and transportation, education, child care, clothing, vacation, and entertainment (with the last three categories only available since 2005). We add the monetary value of food stamps and imputed rents for homeowners and free-rent households. To obtain a measure of consumption and income, we deflate by the Consumer Price Index and as before use the OECD adult equivalence scale. This measure of consumption is available for the survey years 1999–2013. To reduce the impact of measurement error, we take moving averages across three subsequent surveys for both consumption and income.

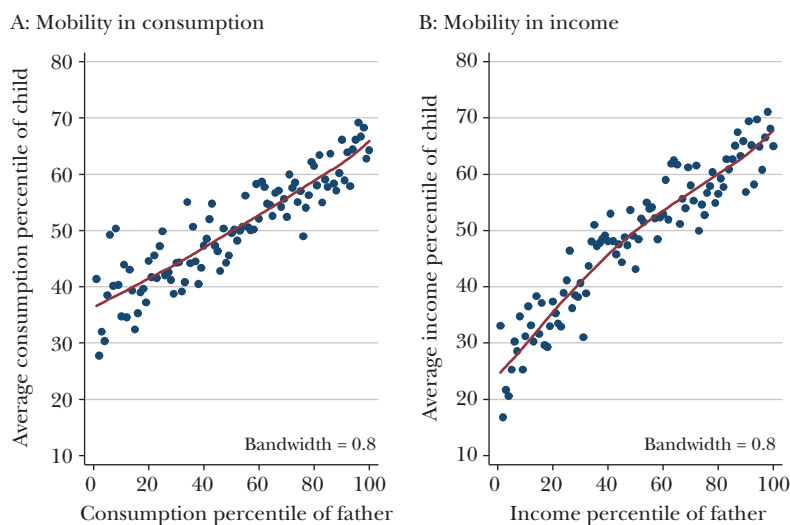
For each year in which the household is observed, we compute the percentile occupied by the household relative to the head’s reference birth cohort (born in the 1900s, 1910s, and so on). We do this for the father and his children. Next, we look graphically at the relationship between the average percentile occupied by the children and the percentile occupied by the father. If there is no relationship between the ranks of parents and children, the (local) regression lines we plot separately for consumption and income in Figure 6 should be flat; on the other hand, perfect correlation between the ranks of parents and children would give a 45-degree line.⁴

We find that the slope of the local regression line for income gradient is higher than that for consumption, implying greater intergenerational mobility in consumption than income. This finding is especially true at the bottom of the distribution. Hence, as consumption is more equally distributed than income, there is also more intergenerational mobility when looking at consumption than income.⁵

⁴If we follow the suggestion of Chetty, Hendren, Kline, and Saez (2014) of conditioning on the age of parent and child, we get similar, though less-precise, results. Interestingly, the relationship we plot in the right panel of Figure 6 is remarkably similar to that reported by Chetty et al., despite the enormous differences in sample sizes.

⁵Wodon and Yitzhaki (2002) extend the traditional Sen (1973) welfare function to the dynamic case. Sen’s welfare function increases with aggregate income and declines with inequality. Wodon and Yitzhaki’s social welfare function increases with intergenerational mobility. Hence, social welfare is higher when considering consumption than when considering income not only because of less unequal distribution of consumption (relative to income), but also because of higher intergenerational mobility.

Figure 6

Intergenerational Mobility in Consumption and Income

Source: Authors using Panel Study of Income Dynamics data.

Note: If there is no relationship between the ranks of parents and children, the regression lines we plot separately for consumption and income should be flat; on the other hand, perfect correlation between the ranks of parents and children would give a 45-degree line. See text for details.

The explanation of the former phenomenon is in all likelihood the tendency to smooth-out income shocks whenever possible (through saving and borrowing, public programs, or informal mechanisms). As for intergenerational mobility, one can conjecture that parents transfer genetic endowments of ability (which will be reflected in both consumption and income) as well as preferences (which will be reflected primarily in consumption). However, the extent of similarity between the consumption of parents and the consumption of children also depend on the credit and insurance market frictions faced by the two generations. (Similarities can also depend on the point of the life-cycle we are observing father and child, but we neglect this complication here.) In the end, whether there is more or less intergenerational mobility in income or consumption is an empirical matter, and the data we present constitute one of the first pieces of relevant evidence in this regard.

Conclusion

The goal of this paper has been to discuss what we do and do not know about the evolution of consumption inequality in the United States, while contrasting it with trends in income inequality. There is now some cumulating evidence

showing that increasing disparities in income are approximately replicated by increasing disparities in consumption. These findings suggest that a substantial portion of the nature of the shocks to income and wages that have generated the observed and well-documented increase in income inequality over the last 35 years should be viewed as permanent rather than temporary, and that households have only a limited ability to absorb such shocks for more than a short period.

While much attention in discussions of inequality has been given to the top of the income and consumption distribution, the left tail is also of considerable interest, both from a scientific and policy point of view. We have considered different components of consumption, inequality in leisure, and also the intergenerational transmission of consumption inequality. When looking at individual components, some of them show greater equality, and some raise difficult questions of how to adjust for quality changes. Ownership of major durables, which in principle raise living standards, has also been converging rapidly between low- and high-permanent-income households. While inequality in food consumption has increased, there is little evidence of growing inequality in caloric intakes—partly as a result of assistance provided by government programs (like food stamps) supplementing private spending, partly from price declines of some food items, and probably in part because low-income people spend more time searching for lower prices. The latter trend involving time use is actually more general: the consumption of leisure has increased among low-socioeconomic status individuals at a faster pace than among the higher educated.

What do we conclude about whether disparities in well-being have increased? Our opinion is that, despite the fact that some studies have suggested the opposite, inequality in the consumption of nondurables and services has increased substantially over the last few decades and has paralleled the increase in inequality in income and earnings. A consequence of this is that the increase in income inequality is reflected in an increase in inequality in welfare and well-being.

Some important caveats, however, are in order. We have provided evidence on specific goods, leisure, durable ownership, and even mobility across generations that points to relative utility gains realized by the lower-income groups. These relative gains arise as a consequence of lower prices for the goods they typically purchase, increasing availability of leisure time, increasing durable ownership, and improved consumption opportunities for their children. Obviously, assigning a value to these utility gains is hard, and we do not attempt to do this. It should also be pointed out that most of these gains have been in quantity terms, not quality terms.

This discussion suggests that if using consumption is in principle a better way than income to measure the well-being of households, a complete welfare analysis will need to go beyond looking at aggregate categories of household expenditure, and consider in addition the value that people assign to time and the quality of goods they consume, among other factors.

References

- Abowd, John M., and David Card.** 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57(2): 411–45.
- Aguiar, Mark A., and Mark Bils.** 2011. "Has Consumption Inequality Mirrored Income Inequality?" NBER Working Papers 16807.
- Aguiar, Mark, and Mark Bils.** 2015. "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review* 105(9): 2725–56.
- Aguiar, Mark, and Erik Hurst.** 2005. "Consumption vs. Expenditure." *Journal of Political Economy* 113(5): 919–48.
- Aguiar, Mark, and Erik Hurst.** 2007. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *Quarterly Journal of Economics* 122(3): 969–1006.
- Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura.** 2007. Chap. 17 in "What Really Happened to Consumption Inequality in the United States?" In: *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, edited by Ernst E. Berndt and Charles R. Hulten. National Bureau of Economic Research.
- Attanasio, Orazio P., and Steven J. Davis.** 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104(6): 1227–62.
- Attanasio, Orazio P., Erik Hurst, and Luigi Pistaferri.** 2015. "The Evolution of Income, Consumption, and Leisure Inequality in the US, 1980–2010." Chap. 4 in *Improving the Measurement of Consumer Expenditures*, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus. University of Chicago Press.
- Attanasio Orazio, and Luigi Pistaferri.** 2014. "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure." *American Economic Review* 104(5): 122–26.
- Attanasio, Orazio, and Guglielmo Weber.** 2010. "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy." *Journal of Economic Literature* 48(3): 693–751.
- Autor, David, Andreas Ravndal Kostøl, and Magne Mogstad.** 2015. "Disability Benefits, Consumption Insurance, and Household Labor Supply." <https://bfi.uchicago.edu/research/working-paper/disability-benefits-consumption-insurance-and-household-labor-supply-0>.
- Aydin, Deniz.** 2015. "The Marginal Propensity to Consume Out of Liquidity." http://stanford.edu/~daydin/DAYdin_MPCL.pdf.
- Baker, Scott R.** 2014. "Debt and the Consumption Response to Household Income Shocks." April. http://web.stanford.edu/~srbaker/Papers/Baker_DebtConsumption.pdf.
- Bee, Adam, Bruce D. Meyer, and James X. Sullivan.** 2015. "The Validity of Consumption Data: Are the Consumer Expenditure Interview and Diary Surveys Informative?" In *Improving the Measurement of Consumer Expenditures*, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus. University of Chicago Press.
- Black, Sandra, and Paul Deveraux.** 2011. "Recent Developments in Intergenerational Mobility." Chap. 16 in *Handbook of Labor Economics*, vol. 4, Part B, edited by Orley Ashenfelter and David Card. North Holland Press, Elsevier.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston.** 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98(5): 1887–1921.
- Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten.** 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review* 106(2): 387–435.
- Broda, Christian, Ephraim Leibtag, and David E. Weinstein.** 2009. "The Role of Prices in Measuring the Poor's Living Standards." *Journal of Economic Perspectives* 23(2): 77–97.
- Broda, Christian, and David E. Weinstein.** 2008. *Prices, Poverty and Inequality: Why Americans Are Better Off Than You Think*. Washington, DC: AEI Press.
- Browning, Martin, and Thomas F. Crossley.** 2001. "The Life-Cycle Model of Consumption and Saving." *Journal of Economic Perspectives* 15(3): 3–22.
- Browning, Martin, and Søren Leth-Petersen.** 2003. "Imputing Consumption from Income and Wealth Information." *Economic Journal* 113(488): F282–F301.
- Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez.** 2014. "Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States." NBER Working Paper 19843.
- Crossley, Thomas F., and Joachim K. Winter.** 2015. "Asking Households about Expenditures: What Have We Learned?" Chap. 1 in *Improving the Measurement of Consumer Expenditures*, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus. University of Chicago Press.
- Cutler, David M., and Lawrence F. Katz.** 1991. "Macroeconomic Performance and the Disadvantaged." *Brookings Papers on Economic Activity* no. 2, pp. 1–74.

- Cutler, David M., and Lawrence F. Katz.** 1992. "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s." NBER Working Paper 3964.
- Deaton, Angus, and Margaret Grosh.** 2000. "Consumption." Chap. 5 in *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study*, Vol. 1, edited by Margaret Grosh and Paul Glewwe. Washington, DC: World Bank.
- DeBacker, Jason, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos.** 2013. "Rising Inequality: Transitory or Persistent? New Evidence from a Panel of U.S. Tax Returns." *Brookings Papers on Economic Activity*, Spring, pp. 67–142.
- De Giorgi, Giacomo, Anders Frederiksen, and Luigi Pistaferri.** 2015. "Consumption Network Effects." November 23. <https://www.economics.utoronto.ca/index.php/index/research/downloadSeminarPaper/60199>.
- Dynarski, Susan, and Jonathan Gruber.** 1997. "Can Families Smooth Variable Earnings?" *Brookings Papers on Economic Activity* 1: 229–305.
- Friedman, Milton.** 1957. *A Theory of the Consumption Function*. Princeton University Press.
- Gelman, Michael, Shachar Kariv, Matthew Shapiro, Dan Silverman, and Steven Tadelis.** 2014. "Harnessing Naturally Occurring Data to Measure the Response of Spending to Income." *Science*, July 11, 345(6193): 212–15.
- Gottschalk, Peter, and Robert Moffitt.** 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *Brookings Papers on Economic Activity* 25(2): 217–72.
- Griffith, Rachel, Ephraim Leibtag, Andrew Leicester, and Aviv Nevo.** 2008. "Timing and Quantity of Consumer Purchases and the Consumer Price Index." NBER Working Paper 14433.
- Gross, David B., and Nicholas S. Souleles.** 2002. "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data." *Quarterly Journal of Economics* 117(1): 149–85.
- Guvenen, Fatih, Serdar Ozkan, and Jae Song.** 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy* 122(3): 621–60.
- Handbury, Jessie.** 2014. "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities." Unpublished paper.
- Hayashi, Fumio, Joseph Altonji, and Lawrence Kotlikoff.** 1996. "Risk-Sharing between and within Families." *Econometrica* 64(2): 261–94.
- Heathcote, Jonathan, Fabrizio Perri, and Gianluca L. Viuolante.** 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006." *Review of Economic Dynamics* 13(1): 15–51.
- Kaplan, Greg, and Guido Menzio.** 2015. "The Morphology of Price Dispersion." *International Economic Review* 56(4): 1165–1206.
- Kaplan, Greg, Guido Menzio, Leena Rudanko, and Nicholas Trachter.** 2016. "Relative Price Dispersion: Evidence and Theory." NBER Working Paper 21931.
- Koijen, Ralph, Stijn Van Nieuwerburgh, and Roine Vestmanz.** 2015. "Judging the Quality of Survey Data by Comparison with 'Truth' as Measured By Administrative Records: Evidence from Sweden." Chap. 11 in *Improving the Measurement of Consumer Expenditures*, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus. University of Chicago Press.
- Kopczuk, Wojciech, Emmanuel Saez, and Jae Song.** 2010. "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937." *Quarterly Journal of Economics* 125(1): 91–128.
- Krueger, Dirk, and Fabrizio Perri.** 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies* 73(1): 163–93.
- MacCurdy, Thomas E.** 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis." *Journal of Econometrics* 18(1): 82–114.
- Meyer, Bruce D., Wallace K. C. Mok, and James X. Sullivan.** 2015. "Household Surveys in Crisis." *Journal of Economic Perspectives* 29(4): 199–226.
- Meyer, Bruce D., and James X. Sullivan.** 2013. "Consumption and Income Inequality in the U.S. since the 1960s." <https://www3.nd.edu/~jsullivan4/Inequality3.6.pdf>.
- Modigliani, Franco, and Richard Brumberg.** 1955. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." Chap. 15 in *Post Keynesians Economics*, edited by Kenneth K. Kurihara. Rutgers University Press.
- Nevo, Aviv, and Arlene Wong.** 2015. "The Elasticity of Substitution between Time and Market Goods: Evidence from the Great Recession." http://sites.northwestern.edu/awo760/files/2015/10/Paper_Aug19_2015-y5zgke.pdf.
- Passero, William, Thesia I. Garner, and Clinton McCully.** 2015. "Understanding the Relationship: CE Survey and PCE." Chap. 6 in *Improving the Measurement of Consumer Expenditures*, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus. University of Chicago Press.
- Pistaferri, Luigi.** 2015. "Household Consumption: Research Questions, Measurement Issues, and Data Collection Strategies." *Journal of Economic and Social Measurement* 40(1–4): 123–49.

- Sabelhaus, John, and Jeffrey A. Groen.** 2000. "Can Permanent-Income Theory Explain Cross-Sectional Consumption Patterns?" *Review of Economics and Statistics* 82(3): 431–38.
- Sabelhaus, John, David Johnson, Stephen Ash, David Swanson, Thesia I. Garner, John Greenlees, Steve Henderson.** 2015. "Is the Consumer Expenditure Survey Representative by Income?" Chap. 8 in *Improving the Measurement of Consumer Expenditures*, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus. University of Chicago Press.
- Sen, Amartya.** 1973. *On Economic Inequality*. Clarendon Press.
- Singh, Gopal K., Mohammad Siahpush, Robert A. Hiatt, and Lava R. Timsina.** 2011. "Dramatic Increases in Obesity and Overweight Prevalence and Body Mass Index among Ethnic-Immigrant and Social Class Groups in the United States, 1976–2008." *Journal of Community Health* 36(1): 94–110.
- Singh, Rajni, Berdine R. Martin, Yvonne Hickey, Dorothy Teegarden, Wayne W. Campbell, Bruce A. Craig, Dale A. Schoeller, Deborah Anne Kerr, and Connie M. Weaver.** 2009. "Comparison of Self-Reported, Measured, Metabolizable Energy Intake with Total Energy Expenditure in Overweight Teens." *American Journal of Clinical Nutrition* 89(6): 1744–50.
- Slesnick, Daniel T.** 1991. "The Standard of Living in the United States." *Review of Income and Wealth* 37(4): 363–86.
- Slesnick, Daniel T.** 1994. "Consumption, Needs and Inequality." *International Economic Review* 35(3): 677–703.
- Slesnick, Daniel T.** 2001. *Consumption and Social Welfare*. Cambridge Books.
- Wang, Dong D., Cindy W. Leung, Yanping Li, Eric L. Ding, Stephanie E. Chiuve, Frank B. Hu, and Walter C. Willett.** 2014. "Trends in Dietary Quality among Adults in the United States, 1999 through 2010." *JAMA Internal Medicine* 174(10): 1587–95.
- Wodon, Quentin, and Shlomo Yitzhaki.** 2002. "Inequality and Social Welfare." *A Sourcebook for Poverty Reduction Strategies*, Vol. 1: *Core Techniques and Cross-Cutting Issues*, edited by Jeni Klugman, 75–104.