Volume Title: Improving the Measurement of Consumer Expenditures

Volume Author/Editor: Christopher D. Carroll, Thomas F. Crossley, and
John Sabelhaus, editors

Series: Studies in Income and Wealth, volume 74

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-12665-X, 978-0-226-12665-4

Volume URL: http://www.nber.org/books/carr11-1

Conference Date: December 2-3, 2011

Publication Date: May 2015

Chapter Title: The Evolution of Income, Consumption, and Leisure
Inequality in the United States, 1980–2010

Chapter Author(s): Orazio Attanasio, Erik Hurst, Luigi Pistaferri

Chapter URL: http://www.nber.org/chapters/c12675

4

The Evolution of Income, Consumption, and Leisure Inequality in the United States, 1980–2010

Orazio Attanasio, Erik Hurst, and Luigi Pistaferri

Recent research has documented that income inequality in the United States has increased dramatically over the last three decades. There has been less of a consensus, however, on whether the increase in income inequality was matched by an equally large increase in consumption inequality. Most researchers have studied this question using data from the Consumer Expenditure Survey (CE) and some studies have suggested that the increase in consumption inequality has been modest. Unfortunately, there is now mounting evidence that the CE is plagued by serious nonclassical measurement error, which hinders the extent to which definitive conclusions can be made about the extent to which consumption inequality has evolved over the last three decades.

In this chapter, we use a variety of different techniques to overcome the measurement error problems with the CE. First, we use data from the diary

component of the CE, focusing on categories where measurement error has been found to be less of an issue. Second, we explore inequality measures within the CE using the value of vehicles owned, a consumption component that is considered to be measured well. Third, we try to account directly for the nonclassical measurement error of the CE by comparing the spending on luxuries (entertainment) relative to necessities (food). This is similar to the recent approach taken by Browning and Crossley (2009) and Aguiar and Bils (2011). Finally, we use expenditure data from the Panel Study of Income Dynamics (PSID) to explore the dynamics of alternative measures of consumption inequality. All of our different methods yield similar results. We find that consumption inequality within the United States between 1980 and 2010 has increased by nearly the same amount as income inequality.

4.1 Introduction

This chapter studies the evolution of the distribution of well-being over the last thirty years in the United States. Our study has three distinctive features. First, we look at different measures of well-being (e.g., income, consumption, and leisure) to assess whether they paint similar pictures with regard to trends in inequality. This is important not only because variables such as consumption and leisure are likely to affect well-being directly, but also because the joint characterization of the evolution of the distribution of these variables can be informative about the nature of the shocks that have affected individual incomes, about the ability of individual households to buffer them and, ultimately, about the potential need for government interventions. Second, we measure inequality in well-being using different indexes and looking at different population groups, which helps us understand movements in the entire distribution and, in particular, whether the trends we observe tend to be concentrated in certain groups within the population. Finally, we draw our inference from disparate sources of data that differ by the quality and the type of well-being measures available, which is useful to assess the robustness of our conclusions. In summary, our analysis of a variety of different data sources suggests that the well-documented rise in income inequality during the last thirty years was accompanied by an increase in consumption inequality of nearly the same magnitude.

It is a very well-known fact that, starting in the early 1980s, inequality in wages (and earnings) in the United States has increased dramatically, both in absolute terms and within groups defined by observable characteristics such as education, labor market experience, occupation, gender, and race. The rise in inequality has been attributed to a combination of many forces, including skill-biased technology changes (such as the computerization of the labor force), institutional factors (such as the decline in unionization and the falling real value of the minimum wage), and the impact of international trade. Some authors have argued that the rise in wage and earnings inequal-

ity has been of a structural or permanent nature; others have noticed that structural factors have been accompanied by a rise in transitory factors of similar or even higher magnitude.[1]

The distinction between temporary and persistent shifts in the wage distribution is important because the nature of the policy interventions aimed at reducing the welfare effects of the rise in inequality depends on identifying correctly what caused it. If the increase in wage inequality is mainly due to unskilled individuals losing ground due to technology shocks making their skills obsolete, policies that try to retrain the unskilled may be effective. In contrast, if the rise in wage inequality is primarily due to transitory forces (such as increased turnover in the labor market), then short-run income support policies are more appropriate to reduce the welfare consequences of increasing inequality.

The distinction between temporary and persistent forces also highlights the usefulness of measures of welfare, such as consumption or leisure, that are likely to depend on long-run (or permanent) income. This consideration has spurred a large and growing literature looking at trends in consumption inequality. A first set of contributions, which includes among others Cutler and Katz (1992), Attanasio and Davis (1996), Slesnick (2001), Attanasio, Battistin, and Ichimura (2007), Krueger and Perri (2006), Meyer and Sullivan (2009), Attanasio, Battistin, and Padula (2010), and Aguiar and Bils (2011), had as a primary objective verifying whether the trends in consumption inequality mirror the trends in wage or earnings inequality. Implicitly, the question that these papers try to answer is whether the worries induced by the well-documented increased dispersion in the wage and earnings distributions were confirmed by observing an increase in consumption inequality of similar magnitude. Another set of contributions, such as Deaton and Paxson (1994), Attanasio and Davis (1996), Blundell and Preston (1998), Krueger and Perri (2006), Blundell, Pistaferri, and Preston (2008), Parker, Vissing-Jorgensen, and Ziebarth (2009), Heathcote, Perri, and Violante (2010), and Attanasio and Pavoni (2011), use information on consumption (and sometimes income) inequality to test a number of theoretical predictions, such as the hypothesis of complete markets, the presence of partial insurance against income shocks, or evidence for endogenous incomplete markets due to asymmetric information or limited commitment.

We complement and extend the existing literature in a number of directions. First, and most importantly, we analyze the evolution of consumption inequality with a variety of empirical strategies, using different consumption measures, and using consumption data from many alternative data sets. When exploring the changing nature of consumption inequality within the United States, most of the studies cited above use nondurable expenditure

---

1. For a detailed discussion of these issues see, for example, Autor, Katz, and Kearney (2008) and the citations within.

data from the Consumer Expenditure Survey (CE). It is now well documented that the CE has measurement problems that are nonclassical in a way that will likely bias the estimates of trends in consumption inequality. For example, many papers document the fact that aggregate measures of expenditure from the CE does a poor job at reproducing the level of expenditure in national account data (see Garner and Maki 2004). The most worrying feature is the fact that the large discrepancy between CE aggregate consumption measures and the personal consumption expenditures (PCE) aggregates has been increasing over time. Additionally, Aguiar and Bils (2011) document that higher-income households are increasingly likely to underreport their expenditures relative to lower-income households. If true, such measurement error will mechanically result in trends in consumption inequality to be increasingly biased downward. This can be one reason why authors who have used CE data have concluded that the rise in consumption inequality during the last thirty years was only a small fraction of the rise in income inequality during the same period (see, for example, Krueger and Perri 2006).

We start the empirical analysis of this chapter by replicating the analysis of consumption inequality in the main (interview) CE survey. However, we also perform many exercises to try to overcome the measurement error problems in the CE data. First, we examine consumption inequality in categories of the CE that have been found to be measured well relative to the PCE in all years of the survey. Using the properties of a simple demand system where the consumption categories are measured with an error structure that we specify, we can then scale up the measures of consumption inequality in these categories by the income elasticity for that category to get a measure of overall consumption inequality.[2] Second, we use data from the diary component of the CE where measurement error in some of the categories have been found to be less problematic. Third, we look at the stock of car owning in the CE and use the imputed value of those vehicles to create an alternative measure of consumption inequality. Finally, we can use expenditure data from the PSID—where systematic changes in measurement error has not been documented—to compute trends in overall consumption inequality.

All of the different methods tell a very similar story. During the last thirty years, consumption inequality evolved very similarly to income inequality. In particular, our estimate of the standard deviation of log income increased by roughly 0.2 log points between 1980 and the latter part of the first decade of the twenty-first century. Depending on our sample and measure of expenditure, our preferred estimates of the increase in the standard deviation of log consumption ranged between 0.15 and 0.2 log points during this time period (depending on the sample and the measure of consumption used). All of these estimates are much larger than the estimates obtained using

---

2. The approach we take is related to ideas in Browning and Crossley (2009) and Aguiar and Bils (2011).

the interview CE survey data, without accounting for the changing nature of measurement error within the survey. The striking feature is how robust these estimates are across the different surveys and consumptions measures we explore.

Our second contribution is to document the evolution of leisure inequality within the United States during the last thirty years. We show that despite the fact that consumption and income inequality increased dramatically between high- and low-educated households during this time period, the change in actual utility differences between the two groups was muted by the fact that low-educated households were spending much more time in leisure relative to their highly educated counterparts.

We also look at different aspects of the change in the distribution of income, consumption and leisure inequality. To do this, we look at trends in inequality both at the top of the distribution (as measured by the 90th–50th percentile difference) and at the bottom of the distribution (as measured by the 50th–10th percentile difference). Lastly, we explore the evolution of leisure inequality during this time period. We find that despite the fact that higher-income individuals experienced a rapid rise in consumption relative to lower-income individuals, higher-income individuals experienced a smaller change in leisure relative to lower-income individuals.

Overall, our results suggest that there has been a substantial rise in consumption and leisure inequality within the United States during the last thirty years. The rise in income inequality translated to an increase in actual well-being inequality during this time period because consumption inequality also increased. Some of this increase, however, was offset by the fact that leisure inequality increased as well, in particular with lower-income individuals taking more leisure relative to their highly educated counterparts.

4.2 A Conceptual Framework

In this section, we expand upon some of the conceptual issues we need to address to assess the changing nature of income, consumption, and leisure inequality. Additionally, we will introduce the conceptual framework we will be using to address the measurement error within the CE data.

4.2.1 Consumption versus Income Inequality

Most analyses of inequality focus on income, not consumption. Partly, this is due to data availability. Data sets containing information on measures of household resources (wages, earnings, income, etc.) are more frequently available, have typically larger samples, and have more consistent variable definitions than data sets containing information on consumption.

While an analysis of income inequality is very valuable, one may argue that analyzing trends in consumption inequality may be even more infor-

mative from a welfare point of view. Since individuals' utility is typically defined over consumption of goods rather than income per se, one may argue that measures of consumption inequality get closer to an ideal measure of inequality in household welfare than income inequality. Moreover, large changes in income inequality may reflect transitory variations, and these may have small welfare effects if households can smooth their consumption against transitory shocks. In other words, consumption might be a better proxy of "permanent income." Consumption inequality might therefore provide a more reliable measure of inequality in long-term living standards than income. Finally, a study of consumption inequality allows researchers to study allocation of disposable income to different commodities, which differ in their necessity/luxury characteristics. This analysis may be important insofar as an increase in food-spending inequality is perceived as being more worrying than, say, an increase in the inequality of spending on holidays.

In practice, it may be important to study income and consumption inequality simultaneously. Their joint analysis may be informative about smoothing possibilities available to consumers, as well as distinguishing between external shocks in insurance opportunities as opposed to fundamental changes in the income process caused by, say, labor market reforms, technological changes, and so forth. Moreover, one can distinguish between income- and consumption-based measures of poverty and study their evolution over the business cycle.

While the main focus of this chapter is the analysis of the evolution of consumption inequality, partly because the trends on income inequality are much better known and partly to put the consumption inequality figures into context, we start our result section (4.4) with some discussion on the evolution of income inequality, where income is measured by total household income divided by the number of adult equivalents.

As we discuss in section 4.3, we will be using different data sources, some of which have an established use in the analysis of income and earnings inequality. We will also discuss the fact that different pictures emerge when we consider inequality of consumption measures from different data sources that rely on completely different samples. Comparing income inequality in the same data sources can then be informative about the nature of these differences, with the two main alternatives being the different nature of the consumption information contained in the data sets and the composition of the samples used in the analysis.

4.2.2 Measures of Inequality and Changes in the Distribution

When looking at the evolution of consumption and income inequality, we will start by considering the evolution over time of the standard deviation of the log of both consumption and income in the samples described below. However, the evolution of the standard deviation of log consumption (or

income) is only one way to characterize the changing inequality within the distribution of interest. It maybe that a given change in the standard deviation corresponds to a large change in the difference between the top of the distribution and its middle with nothing much happening in the bottom of the distribution. As a result, to provide a more complete picture of what has happened to consumption and income inequality in the last thirty years, we will also be looking at the difference between the 90th percentile and the median as well as the difference between the median and the 10th percentile of the respective distributions.

4.2.3 Inequality in Different Dimensions: Skill and Year of Birth Groups

The statistics mentioned in the previous subsection will be computed on the whole sample we use. It may be of considerable interest, however, to consider the evolution of the distribution over time in other dimensions, as they might suggest direct economic interpretations to what has happened. An important dimension we will be looking at is that of the difference across skill groups (as proxied by the education achievement of the household head). In particular, we will be looking at inequality both across and within different skill groups. The evolution of differences in income and consumption between skill groups might reflect the evolution of the prices of different skills in the labor market, which in turn have been associated to technological progress and other innovations that are likely to be permanent and difficult to insure and smooth out. Inequality within skill groups will reflect both the evolution of unobserved skill prices and other factors. When considering this decomposition, it is clear that the simultaneous analysis of consumption and income inequality can be particularly informative about the nature of the shocks we observe and household ability to smooth them out.

An issue that potentially affects many measures that have been considered in the literature is the fact that when we follow the evolution of inequality measures over time, they might reflect changes in the composition of the sample we are considering. This concern may be particularly salient when exploring the patterns of inequality over long periods of time. This is true for the overall sample and, even more so, in the case of the skill groups, as the fraction of, for instance, high school dropouts declines monotonically over the sample period. To address this issue, one can consider the evolution of inequality within groups whose membership is (approximately) constant over time. For instance, one can define groups by year of birth (of the household head), and by doing so follow the same group of individuals over time.[3] The evolu-

3. Groups defined by the year of birth of the household head can change in composition, however, for several reasons. First, it is possible that family formation and dissolution is different for individuals of different economic status. Second, there are strong differences in mortality rates between rich and poor individuals that are likely to make observed cohorts progressively "richer." Finally, it is possible that migration patterns are also related to economic status.

tion of the distribution of a given variable, being consumption or income, in the overall sample can mask very different dynamics for a fixed group of individuals. This is particularly the case in the presence of strong cohort effects. Moreover, theoretical models of insurance of income shocks have specific implications for the evolution of the relative distribution of income and consumption. Following the evolution of these distributions over the life cycle can therefore be particularly interesting. For reason of space, we discuss results related to the evolution of inequality over the life cycle in an appendix available on our website.[4]

4.2.4 Measuring Inequality: Accounting for Measurement Error

As discussed in the introduction, one of the main issues we have to deal with when studying the distribution of consumption and its evolution over time is the presence of large measurement error of the nonclassical type in the CE. The CE, however, contains details on hundreds of commodities that, in turn, can be aggregated into different categories, some of which have been documented as providing a good match to PCE data (see Bee, Meyer, and Sullivan, chapter 7, this volume; Garner and Maki [2004]). One possible approach to study the evolution of the inequality of overall consumption is therefore to focus on consumption categories that are well measured. To get an estimate of the changing nature of total consumption inequality, one simple approach is to compute the extent of consumption inequality using the specific consumption category that is measured well and then scale that measure up by the category's income elasticity. We do this below.

Additionally, we can take a stand on the nature of the measurement error in the consumption data. Let's denote with C_{it} the total consumption of household i in period t. Suppose that total consumption is made of K different categories with $C_{it} = \Sigma_{k=1}^{K} q_{it}^k$, where q_{it}^k is the spending on consumption category k by household i in period t. Let's consider two commodities that are known to be measured without systematic error, q_{it}^1 and q_{it}^2, and suppose that commodity 2 is a necessity, while commodity 1 is a luxury. As usual, we define a necessity as a commodity whose elasticity with respect to total expenditure is less than one and a luxury as a commodity whose income elasticity is greater than one.

Suppose that in the case of commodities q_{it}^1 and q_{it}^2 spending on those categories can be expressed by the following equations:

(1) $$q_{it}^1 = C_{it}^{\alpha_1} u_{it}^1 v_t^1 \ , \ \alpha_1 > 1$$

(2) $$q_{it}^2 = C_{it}^{\alpha_2} u_{it}^2 v_t^2 \ , \ \alpha_2 < 1$$

Equations (1) and (2) represents two Engel curves. They relate the expenditure of each of the two commodities to total expenditure (with α_1 and α_2 being

---

4. See http://www.stanford.edu/pista.

the income elasticities), some aggregate factors v_t^j (with $j = 1, 2$), such as relative prices, and some unobserved idiosyncratic taste shocks (u_{it}^j, $j = 1, 2$). We will assume that the idiosyncratic taste shocks are i.i.d across households and that their distribution is constant through time. Taking the ratio between q_{it}^1 and q_{it}^2 one obtains:

$$
(3) \qquad \frac{q_{it}^1}{q_{it}^2} = C_{it}^{\alpha_1 - \alpha_2} \frac{v_t^1 u_{it}^1}{v_t^2 u_{it}^2}.
$$

Taking logs of this expression, one gets:

$$
(4) \qquad \begin{aligned} \log(q_{it}^1) - \log(q_{it}^2) &= (\alpha_1 - \alpha_2)\log(C_{it}) + (\log(v_t^1) - \log(v_t^2)) \\ &\quad + (\log(u_{it}^1) - \log(u_{it}^2)) \end{aligned}.
$$

Computing the cross-sectional variance of both sides of equation (4) and assuming for the time being that the idiosyncratic taste shocks are uncorrelated with total expenditure, one obtains:

$$
(5) \qquad \begin{aligned} \text{Var}(\log(q_{it}^1) - \log(q_{it}^2)) &= (\alpha_1 - \alpha_2)^2 \text{Var}(\log(C_{it})) \\ &\quad + \text{Var}(\log(u_{it}^1) - \log(u_{it}^2)) \end{aligned}.
$$

Expression (5) deserves several comments. First, the aggregate shocks, by their very nature and because they enter additively in equation (4), do not contribute to the variance of the right-hand side. Second, the left-hand side of equation (5) is observed and can be computed in a data set that contains detailed information on consumption. In situations where a reliable measure of total consumption is not available because some of its components are affected by substantial measurement error whose variance is changing over time, the interesting question is the extent to which we can use such a variable as an approximation for the level or the changes in total consumption inequality. Notice that, because of the choice of commodities, $(\alpha_1 - \alpha_2) \neq 0$ so that the left-hand side of equation (5) will be varying with changes in the variance of total expenditure. If one is willing to assume that the variance of the taste shocks is invariant over time, then *changes* in the left-hand side will be driven entirely by *changes* in the variance of total consumption. Indeed, changes in the left-hand-side will be proportional to changes in such a variance, where the factor of proportionality is given by $(\alpha_1 - \alpha_2)^2$. Information on total expenditure elasticities derived from other sources can be used to evaluate the size of such a factor of proportionality.[5]

The approach we propose is similar to the idea discussed in Aguiar and Bils (2011) and even more so to the approach proposed by Browning and Crossley (2009). Browning and Crossley (2009), in particular, consider the

---

5. Aguiar and Bils (2010) use information on demand systems to address the measurement error problems in the CE in the United States. Our approach is, however, different from theirs in that they attempt to address systematic measurement error within the CE. Like them, we also find that consumption inequality tracks income inequality over this time period.

evolution of the covariance of two "noisy measures" of total consumption. Notice that from components, this would imply considering:

(6) $\text{Cov}(\log(q_{it}^1), \log(q_{it}^2)) = (\alpha_1\alpha_2)\text{Var}(\log(C_{it})) + \text{Cov}(\log(e_{it}^1), \log(e_{it}^2)),$

where the e's include both the aggregate and individual shocks in equations (1) and (2). Again, assuming that the second term in equation (6) does not change over time, one can use changes in the covariance on the left-hand side of equation (6) and knowledge of the income elasticities to back out the evolution of total consumption.

In what follows, we will look at the ratio of the changing variance of expenditure on entertainment services relative to the changing variance of expenditure on food at home. The latter is a luxury, while the former is a necessity (see Blow, Lechene, and Levell, chapter 5, this volume). Moreover, as indicated, for instance, in the study by Bee, Meyer, and Sullivan (chapter 7, this volume), these components of the CE are relatively well measured over the sample period we study. Moreover, when aggregated up, the ratio of the resulting aggregates to PCE from the national accounts is relatively constant.

4.3 Data: Surveys and Sample Selections

4.3.1 The Consumer Expenditure Survey (CE)

Survey Overview. Studying income and consumption inequality entails nonnegligible measurement issues. The first is data availability. In the United States there is only one data set with a comprehensive measure of consumption, the Consumer Expenditure (CE) Survey. Other data sets include incomplete consumption information, ranging from just food (the Panel Study of Income Dynamics [PSID] before the 1999 redesign, as well as most proprietary scanner data sets), to spending on only child care and rent (the Survey of Income and Program Participation [SIPP]), to measures that have much more details about expenditure but still fall short of covering the entirety of household budget (the PSID after the 1999 redesign). We begin our analysis with the CE data given that it is designed to provide a comprehensive measure of spending for US households. However, as we discuss below, we are aware that the CE has its own limitations.

The CE survey has a long history, dating back to the beginning of the twentieth century. The main purpose of the survey is to collect information to be used in computing the weights for the Consumer Price Index (the CPI). For this reason, the CE survey contains comprehensive and detailed information about consumption expenditure and its components. Until 1980, the CE was performed roughly every ten years. In 1980, however, it was radically redesigned and became a survey that is run continuously. It is made of two separate and independent samples: the Interview and the

Diary surveys. The former is a rotating panel available on a continuous basis since 1980. Households are interviewed every three months for, at most, five quarters. The first interview is a preparatory one and no data pertaining to it are released. From the second interview, the respondent in each household is asked to report detailed expenditures on hundreds of categories in each of the three months preceding the interviews. These categories are almost exhaustive of total consumption, the only exception being personal care items. Some items, however, are extremely aggregated. The best example is food at home, which is a single category. The information on expenditure is then complemented with information on mortgages, cars (including loans to finance their purchases), credit cards, health and education expenditures, and so on. Finally, the interview survey also includes extensive socioeconomic information on the household, ranging from detailed demographic information to labor supply and earning information on each household member, to some information on assets.

A large part of the income and demographic information are also found in the diary survey. The information on expenditure in this sample, however, is collected with a radically different method. In the diary survey, households are asked to fill in a register (a diary) detailing their spending for two continuing weeks. Until 1986, the diary survey contained information only on "frequently purchased items" such as food and personal care items. The information on food is much more detailed than in the interview survey. Starting in 1986, the diary survey becomes an almost exhaustive expenditure survey, with substantial overlap with the interview survey. Despite this overlap, however, the Bureau of Labor Statistics, which runs the survey, uses the diary for some expenditure components and the interview for others. The presumption is that one survey is better at measuring some components and the other is better for others.

The CE survey is a remarkable data source. Given its richness, it is not surprising that over the last twenty years it has been extensively used by economists for a variety of purposes. However, there are well-known issues with the CE. A particular worry is the lack of correspondence between aggregates derived from the CE survey and the personal consumer expenditure series published in the national accounts. Not only does the CE seem to underestimate substantially the level of PCE consumption, but the ratio of CE aggregates to PCE aggregates has declined substantially over time. Moreover, there is now increasing evidence of nonrandom nonresponses and attrition. In what follows, we will use the CE data without referring explicitly to these issues, although the approach we sketched in section 4.2.4 was designed precisely to deal with the fact that comprehensive measures of consumption derived from the CE might be plagued by substantial measurement error with an increasing variance over time. As we show below, the CE data, without adjusting for potential measurement error issues, provides a very different picture of consumption inequality than does the CE where

measurement error issues are confronted directly or with the PSID where the measurement issues are not as problematic.

Sample Selection. Within the CE data, we select households whose household head is between age twenty-five and sixty-five. With this choice we want to avoid a number of issues that are relevant for very young households and those that are approaching the last part of the life cycle where retirement and health problems become particularly relevant. Family formation and dissolution, binding liquidity constraints for the young group, pressing health problems for the older one, are only some of the issues we want to avoid. Additionally, given that the CEX excludes households living in rural areas from their sampling frame, we drop such households from the sample in all years of our analysis. Finally, we drop from our sample all households with incomplete income responses. The reason for this is that we want to match the sample documenting consumption inequality with the same sample with which we measure income inequality.

Variable Definitions. There are a number of issues one needs to tackle before even starting to analyze trends in consumption inequality. First, which definition of consumption should we focus on? The distinction between durable and nondurable goods is important as it drives a wedge between the concepts of spending and consumption. For nondurable goods the two concepts coincide, but for durable goods (especially large-ticket items) spending is typically done upfront, but the same good provides services over multiple periods. We will be interested in measuring inequality in consumption, rather than inequality in spending per se, and hence will focus our analysis primarily on nondurable spending. When using the CE diary data, it is only possible to construct a comprehensive measure of nondurable consumption starting in 1986. For the CE interview data our nondurable spending data starts in 1980.

While some items are naturally included (e.g., food) or excluded (e.g., furniture) from the definition of nondurable consumption and services, there are a number of arbitrary choices one needs to make. To make our figures comparable with those of other researchers, we decided to include clothing and footwear in our nondurable expenditures measures. On the other hand, we exclude expenditure on health and education, as we see them more as investment in the stock of human capital. On conceptual grounds, we also exclude payments of interest on loans and mortgages (as well as the repayment of the principal). Finally, and somewhat more arbitrarily, we exclude contributions and donations to charities. A complete definition of our measure is reported in appendix A (on website). In addition to nondurable consumption and services, we also consider two additional flow aggregates: the expenditure on food at home and the expenditure on nondurable entertainment. Nondurable entertainment expenditures include items such as cable television subscriptions, DVDs, music, and so forth.

We do explore inequality patterns using the one durable commodity that is measured in a very rich manner in the CE: the amount of vehicles owned by

the household. The CE contains, in a special module, detailed information on the type of cars held by each household. In particular, the make, model, and year is known in addition to a number of car characteristics. Furthermore, if the car has been purchased (new or used) in the twelve months preceding the interview, the purchase price is also reported. We use these data to impute a value for the cars for which no price is reported. Effectively, for the cars for which we have a value, we run an hedonic regression that includes make, model, and year identifiers as well as age and several characteristics. We then use the parameters of this regression to interpolate the value of all the cars in the survey and obtain, for each household, the value of the stock of cars they hold. The procedure is described in detail in appendix B (on website) and is similar to the one used by Padula (1999).[6]

Another relevant issue when we measure inequality in household consumption is that households differ in size and composition, implying important differences in needs as a function of, say, the age and number of children in the households and so forth. To account for these differences we will equivalize household consumption by dividing total consumption by an adult equivalence scale. We use the Organisation for Economic Co-operation and Development (OECD) scale, defined as $S = 1 + 0.7 * (A - 1) + 0.5 * K$ (where A is the number of adults and K the number of children, age eighteen or less, in the household). A final issue is how to deflate monetary variables in our data. One option is to use a global deflator (the CPI), and another is to use commodity-specific deflators, which may be important in the presence of differential trends in relative prices. Here we use the general CPI-Urban deflator (in 1983–1984 USD).[7]

4.3.2 The Panel Study of Income Dynamics

Survey Overview. The PSID is a longitudinal survey of US families that started in 1968 with two subsamples, the Survey Research Center (SRC) sample, which was representative of the US population (60 percent of the initial sample), and the Survey of Economic Opportunity (SEO) sample, which was oversampling poor families (the remaining 40 percent of the initial sample). The main feature of the PSID is that it follows the original survey households as well as households that get formed as a branch of the original ones (e.g., sons or daughters forming their own household unit). Data have been collected yearly from 1968 to 1997, and biannually after that. The latest available survey refers to 2008. The data is primarily geared toward collecting data on labor market items such as labor supply, wages, and so forth. However, the PSID has also collected information on consumption, especially food (at home and away from home) and, in some waves, rent, utilities, and child care. After the 1997 wave the PSID was rede-

---

6. We thank Mario Padula for help with this procedure.
7. For food we use the CPI food deflator.

signed. The survey became biannual and richer in certain interview components (such as expenditure on various commodities, health, wealth, and detailed information on spousal sources of income). For our purposes, the most relevant change was that the PSID started collecting richer and more detailed information on household spending, which now covers 70 percent of total CE spending. See below for a definition.

Sample Selection. Our sample includes all survey households with a head (typically the male) between age twenty-five and sixty-five. We exclude the Latino subsample and keep the SEO subsample, but use sampling weights throughout the analysis. We also exclude observations with outlier records on total household income and food.

Variable Definitions. Similarly to the CE survey, the PSID can also be used to address the dynamics of consumption distributions. To do so, we will look both at direct measures of consumption in the PSID (food, and the post-1997 consumption measure), and imputed consumption measures.

Food consumption is the sum of food at home, food away from home, and the value of food stamps. The post-1997 consumption measure (or "70 percent measure" from now on) includes information on spending on utilities (electricity, heating, water, miscellaneous utilities); home insurance premiums; health (health insurance premiums, nursing care, doctor visits, prescriptions, other health spending); vehicle spending (vehicle insurance premiums, vehicle repairs, gasoline, parking); transportation (bus fares, taxi fares, other transportation expenses); education (tuition, other school expenses); and child care. To match the nondurable consumption definition from the CE, we also consider an alternative measure that excludes spending on education and health.

We adopt two procedures for imputing a measure of total consumption in the PSID. The first measure follows Ziliak (1998) and is based on a simple budget constraint accounting (the "Ziliak measure" from now on). Consumption is defined as the difference between income and the change in assets. Assets are the sum of liquid assets and equity (the difference between the self-reported home value and the remaining principal on the home mortgage). Before 1999, data on asset stocks (with the exception of housing) are reported only every five years (starting in 1984). We thus impute liquid assets by taking the ratio of income from liquid assets (which is available every year) and the return on the T-bill. This imputed measure of consumption is not available for 1981–1982 because no data on equity are available for 1981.[8]

The second measures, based on Blundell, Pistaferri, and Preston (2008), impute consumption using the estimates of a food demand equation from the CEX (the "BPP measure" from now on). In particular, we use the CE data set to estimate (on a sample where the head is age twenty-five to sixty-five

---

8. Note that this imputation procedure provides more correctly a measure of total consumption, rather than total nondurable consumption.

and for each year for which we have data) a regression of log food onto the number of children, a quadratic in the household head's age, a dummy for self-employment, education dummies, log consumption, and the interaction of the latter with education dummies:

$$\ln f_{it} = X_{it}\beta_t + \ln C_{it}\gamma_t(E_{it}) + \varepsilon_{it.}$$

We then use the estimated coefficients in the CEX to impute a measure of consumption in the PSID:

$$\ln \hat{C}_{it} = \frac{\ln f_{it} - X_{it}\hat{\beta}_t}{\hat{\gamma}_t(E_{it})}.$$

We refer the interested reader to Blundell, Pistaferri, and Preston (2008, 2004) for more technical details about this imputation procedure.

Similarly to what done with CE data, in the PSID we also equivalize household consumption using the OECD scale, and deflate nominal values using the general CPI-Urban deflator or the food CPI when we use just food data (both deflators are expressed in 1983–1984 USD).

4.3.3 Time-Use Surveys

Survey Overview. To examine the trends in leisure inequality during this time period, we use data from the 1985 Americans' Use of Time survey and the 2003–2007 American Time Use Survey. The 1985 Americans' Use of Time survey was conducted by the Survey Research Center at the University of Maryland. The sample of 4,939 individuals was nationally representative with respect to adults over the age of eighteen living in homes with at least one telephone. The survey sampled its respondents from January 1985 through December 1985. The 2003 American Time Use Survey (ATUS) was conducted by the US Bureau of Labor Statistics (BLS). Participants in the ATUS, which includes children over the age of fifteen, are drawn from the existing sample of the Current Population Survey (CPS). The individual is sampled approximately three months after completion of the final CPS survey. At the time of the ATUS survey, the BLS updated the respondent's employment and demographic information. During 2003, roughly 1,700 individuals completed the survey each month, yielding an annual sample of over 20,000 individuals. During the 2004–2007 period, roughly 1,160 individuals were surveyed per month yielding an annual sample of just about 14,000 individuals.[9]

Each survey is based on twenty-four-hour time diaries. Survey personnel assign each activity to a category in a set classification scheme. The more refined the classification scheme, the less the survey needs to rely on the judgment of surveyors in correctly coding activities. The ATUS represents the state of the art of time-use surveys for the United States and reports 406

---

9. See Aguiar and Hurst (2007) for a detailed discussion of both surveys.

detailed time-use categories. The 1985 Americans' Use of Time survey used a scheme that included slightly less than 100 categories.

All data in the surveys are weighted so that they are nationally representative using the provided survey weights. Moreover, we also weight the data so that each day of the week is represented equally.

Sample Selection. For both surveys we restrict the sample to those individuals between the age of twenty-five and sixty-five (inclusive). Given that the data is collected at the individual level, we did not restrict the data to include only household heads. We also restricted the data to include only those households that had complete time diaries in that all twenty-four hours were accounted for and were able to be classified into discrete time-use categories.

Variable Definitions. We break the allocation of time into a number of broad time-use categories. As we have constructed the categories, they are mutually exclusive and they sum to the household's entire day. In other words, each person in the survey has twenty-four hours of nonoverlapping activities. Time spent on an activity includes any time spent on transportation associated with that activity.

In terms of our analysis, we use the time-use surveys to construct measures of leisure. Our definition of leisure follows the definition of Aguiar and Hurst (2007). In particular, we think of leisure time as being the time not allocated to market work or to home production (cooking, cleaning, mowing the lawn, etc.). We also exclude time spent taking care of one's children, time spent allocated to health care (going to the doctor), and time spent in educational attainment from our measure of leisure. Our measure of leisure therefore sums together time spent watching television; socializing (relaxing with friends and family, playing games with friends and family, talking on the telephone, attending/hosting social events, etc.); time spent exercising or participating in sports (playing sports, attending sporting events, exercising, running, etc.); reading (reading books and magazines, reading personal mail, reading personal email, etc.); enjoying entertainment events and hobbies (going to the movies or theatre, listening to music, using the computer for leisure, doing arts and crafts, playing a musical instrument, etc.); and all other similar activities.[10]

4.4 The Evolution of Income Inequality

As mentioned above, the main aim of this chapter is the study of the evolution of the *distribution* of welfare, which we will mainly approximate by the distribution of consumption. Before delving in the evidence on con-

---

10. We exclude the following from our measure of leisure: time spent eating, time spent sleeping, and time spent in personal maintenance (grooming, etc.). Aguiar and Hurst (2007) include such activities in some of their leisure measures. Our results are not sensitive to whether or not we include such activities in our leisure measures.
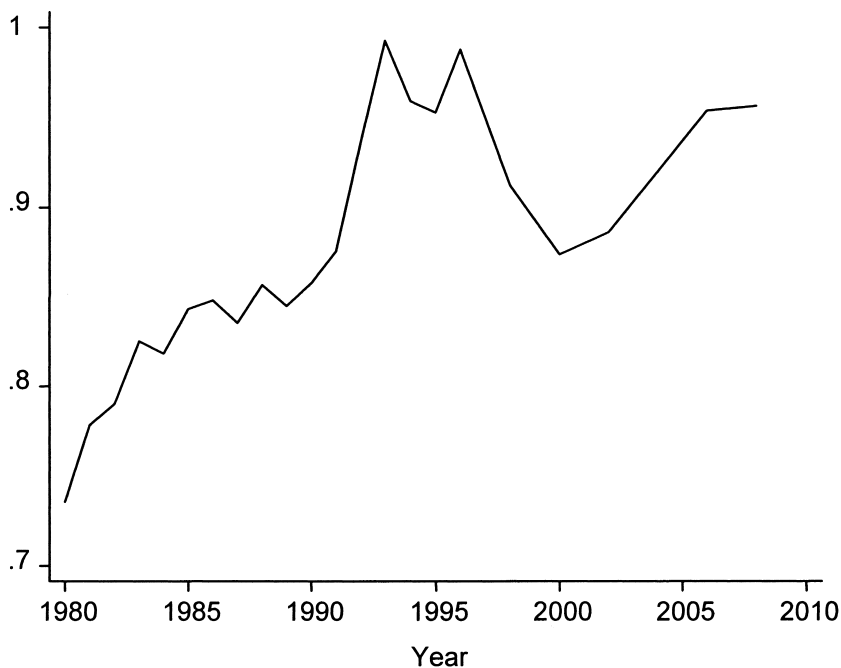
Fig. 4.1 Inequality in (equivalized) family income, PSID

sumption and how its distribution is measured in different data sources and with different definitions of consumption, we provide some evidence on the evolution of the distribution of household income. This piece of evidence is much more familiar and uncontroversial.

Figure 4.1 shows how income inequality, as measured by the standard deviation of logs, has evolved over the 1980–2008 period using PSID data. Our measure of income is before-tax family income, scaled by the OECD equivalence scale. Before-tax family income includes labor earnings, financial income, and public and private transfers received by all household members. All data are deflated using the CPI for urban households (in 1983–1984 USD) and weighted using the PSID longitudinal sampling weights.

The figure summarizes well-known facts. Income inequality, measured by the standard deviation of the logarithms, rises quite rapidly and dramatically over the 1980s until the mid-1990s; it slows down (and it even declines) during the second half of the 1990s before rising again throughout the first decade of the twenty-first century.[11] Between 1980 and 2008, the overall increase in the standard deviation of logs is large, at roughly 0.2.

---

11. We find similar trends using the Gini coefficient, which is less subject to the influence of extreme values.

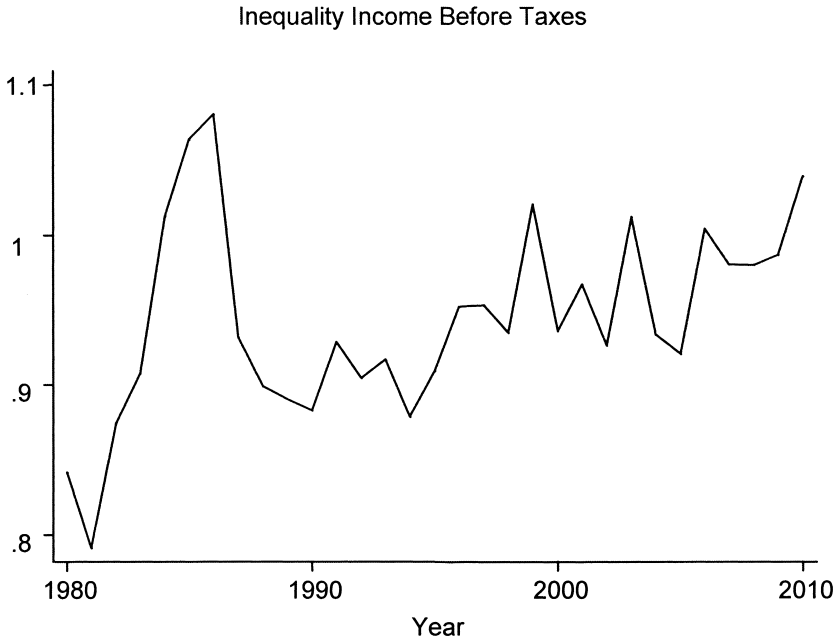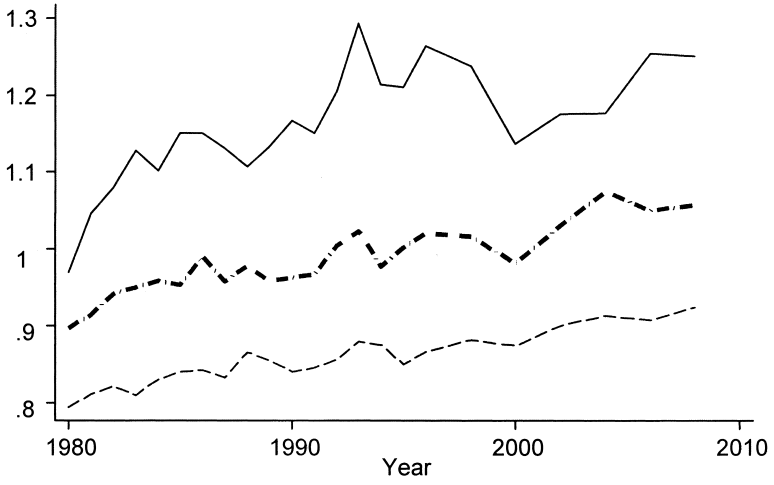Inequality Income Before Taxes

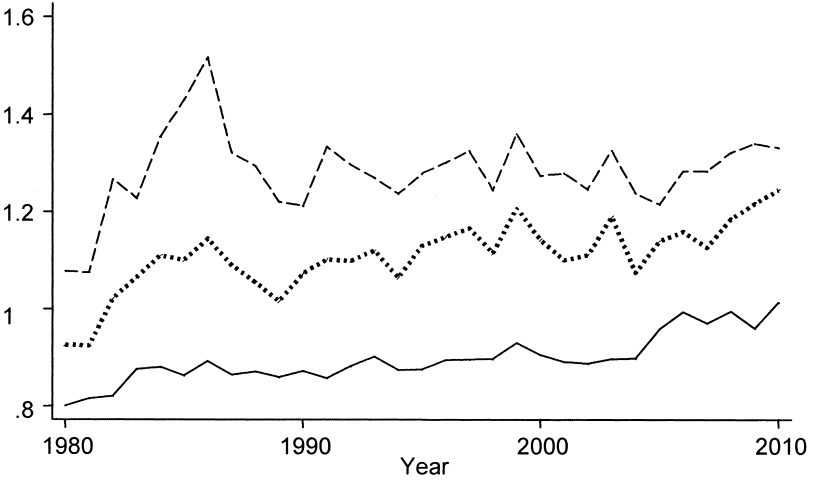**Fig. 4.2 Inequality in (equivalized) family income, CE interview survey**

Figure 4.2 illustrates the change in income inequality using income measures from the CE instead of the PSID. In the CE, we use "before-tax family income" as our measure of income. As with the PSID data, we account for family size differences across households by normalizing income by the OECD equivalence scale. As visible from figure 4.2, the overall increase in the standard deviation of log income within the CE during the last thirty years was nearly identical to the similar measure in the PSID data (roughly 0.2). Notice, however, that there are a few differences between the PSID trends and the CE trends. First, the level of inequality in all years is higher in the CE. This is likely due to the fact that the CE income measure is measured with more error in all years. The CE survey is designed to measure consumption, not income. The PSID data, in contrast, has as its primary goal that of measuring income well. Second, the CE data suggests a sharp rise in income inequality in the early 1980s that is not present in the PSID data. The patterns in the PSID match well the patterns found using data from other surveys such as the Current Population Survey (see Autor, Katz, and Kearney 2008).

As we mentioned above, in addition to the evolution of overall inequality, as measured, for instance, by the standard deviation of logs, we also want to consider the evolution of different parts of the distributions of the variables of interest. In figures 4.3 (PSID data) and 4.4 (CE data), we provide some

Fig. 4.3 (Equivalized) income inequality in different parts of the distribution, PSID data

Fig. 4.4 (Equivalized) income inequality in different parts of the distribution, CE interview survey

Legend:
— 50th-10th percentile difference   ----- 90th-50th percentile difference
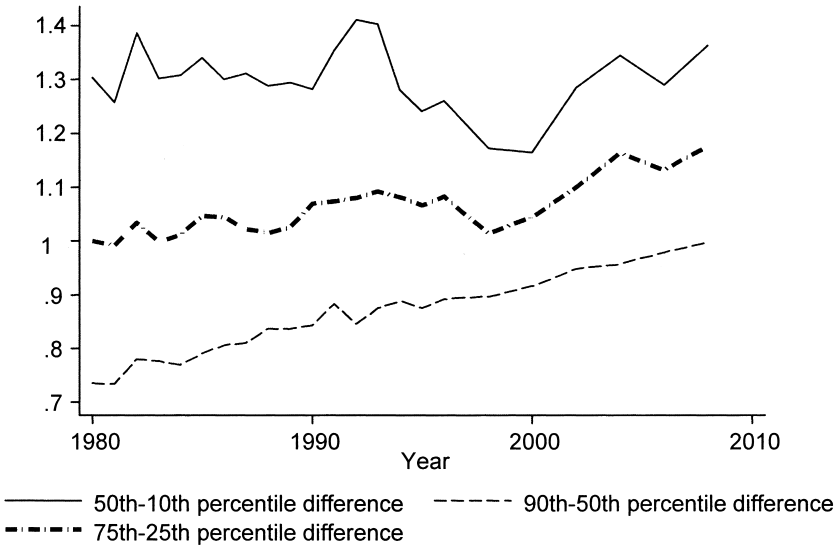·-·-·- 75th-25th percentile difference

Fig. 4.5 Inequality in family earnings, various points of the distribution

information on the evolution of income inequality in different parts of the distribution. To this purpose, we plot the 90th–50th, the 50th–10th, and the 75th–25th percentile difference. Focusing our attention on the PSID data in figure 4.3, we first notice that the difference between the 50th and 10th percentile is considerably larger than the difference between the 90th and the 50th and (to a lesser extent) than the difference between the 75th and 25th percentile. The figure also shows that the decline in income inequality of the 1990s comes primarily from a decline in inequality at the bottom of the distribution, while inequality at the top of the distribution increases almost monotonically throughout our sample period. It should also be noted that the PSID does not sample very rich households, and hence substantially underestimates the rise in inequality that has occurred at the very top of the income distribution. Indeed, over the entire period the rise in inequality at the bottom is larger than at the top (measured in log points). When comparing figures 4.3 and 4.4, notice in particular the steady increase, in both data sets, of the 90th–50th and 75th–25th differentials. In both data sets, the largest increases, especially in the first part of the sample, however, are registered for the 50th–10th differential.

Some of the trends in the above figures are induced by the dynamics of public transfers at the bottom of the distribution. We can assess this fact using the PSID data by comparing figure 4.3 with figure 4.5. Figure 4.5 is similar to figure 4.3, except that the income measure is "family earnings" rather than "family income." The family earnings measure does not include transfer payments (and financial income). In figure 4.5, we find that
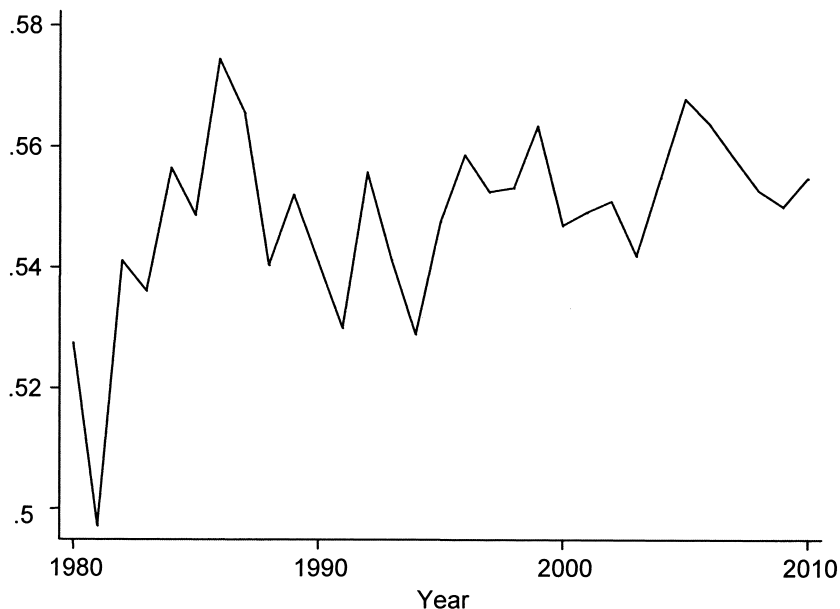
Fig. 4.6 Standard deviation log-equivalized nondurable consumption, CE interview data

inequality in family earnings at the top has risen substantially, while that at the bottom has remained fairly constant (except a drop and a subsequent rise from the mid 1990s to the middle of the first decade of the twenty-first century). Indeed, the main difference between figures 4.3 and 4.5 is the fact that in the former the 50th–10th differential line increases throughout the 1980s and early 1990s, while it is constant in the latter.

4.5 Consumption Inequality: Nondurable Expenditure in the CE Interview Survey

We start our analysis of consumption inequality by looking at data from the interview survey of the CE. We start here because this is often used as a starting point by researchers exploring the evolution of consumption inequality within the United States over the last thirty years. Figure 4.6 reports the standard deviation of log consumption of nondurable commodities and services, as defined in appendix A (on website), and as measured in the interview survey. As with the income data, we adjust the consumption data for differences in family composition using the same OECD scale mentioned above.

The results shown in figure 4.6 are staggering. First of all, we notice that the level of the standard deviation of log consumption is considerably lower than the standard deviation of log household income in all years of the sur-
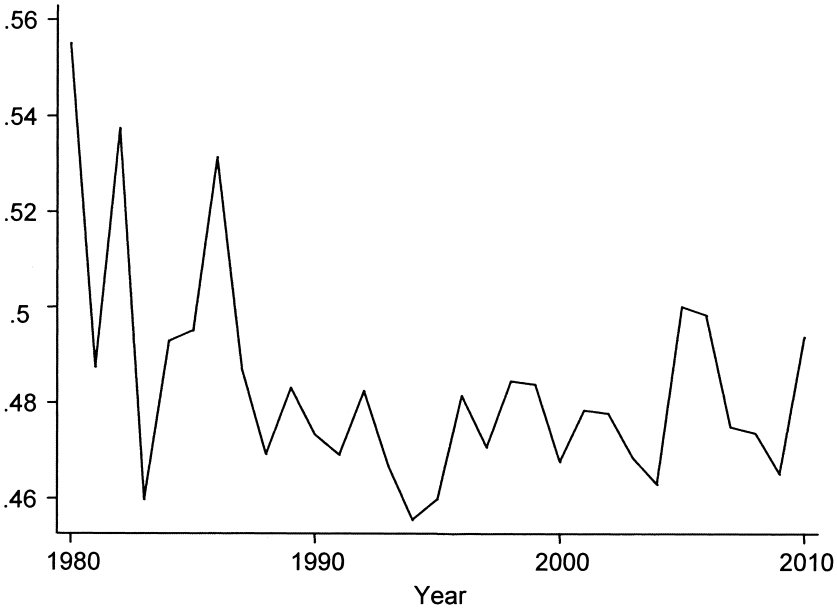
Fig. 4.7 Standard deviation log-equivalized food at home consumption, CE interview data

vey. This, by itself, is not surprising. Some of the variation in income documented in figures 4.1 and 4.2 is transitory. Households are able to smooth out such transitory shocks through borrowing and saving. This implies that the standard deviation in log consumption should be lower, on average, than the variation in log income. The staggering part of figure 4.6 is that the dramatic increase in the standard deviation of log income was not matched by any meaningful increase in the standard deviation of log consumption. Even starting in 1982, where the level of the standard deviation of logs is lowest, the standard deviation of log nondurable consumption as measured in the CE interview survey does not increase by more than 0.06. This is just over a third of the increase witnessed for income. Even this number, however, is likely overstated. Starting in 1983, there was essentially no increase in the standard deviation of log consumption. The bulk of the increase from 1982 through 2010 occurred between 1982 and 1983. This feature of the CE data on consumption inequality has been discussed, among others, by Attanasio, Battistin and Ichimura (2007), Attanasio, Battistin and Padula (2010), Krueger and Perri (2006), Heathcote, Perri, and Violante (2010).

In figure 4.7, we look at the evolution of inequality in (log) consumption of food at home, as measured in the interview survey. Again, we adjust the measure for family size using the OECD equivalence scales. We show these data so as to compare it with the food measures in the CE diary data and
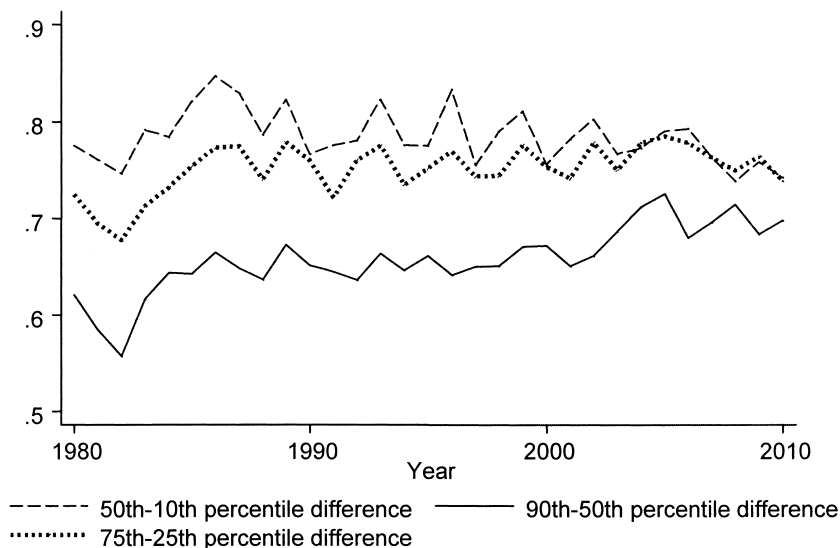
Fig. 4.8 Movements in the distribution of nondurable consumption, CE interview data

with the food measures in the PSID data. Not surprisingly, the standard deviation of this variable is always below that of the standard deviation of total nondurable consumption as shown in figure 4.6. Food at home is more of a necessity good and, as a result, varies less in the cross section. Moreover, the evolution over time indicates that it does not change much until early in the first decade of the twenty-first century, when it exhibits a slight increase. Although there is much more noise in the data, the results in figure 4.7 are broadly consistent with the results from figure 4.6. Given the noise over time in this picture, one could conclude that the inequality of food at home has not increased over time and, relative to the early 1980s, it has, if anything, declined. But if the food measure in the CE is plagued with systematic measurement error over time, then consumption inequality measured using food at home data may also be biased downward.

Having described the evolution of overall inequality in our sample, we now look—as we did for income—at the evolution of different parts of the consumption distribution. As in figures 4.3 and 4.4, we now plot the difference between the 90th and 50th percentile, between the 50th and 10th, and between the 75th and 25th for consumption using the CE interview data. These results are shown in figure 4.8.

As with income, of the three differentials, the largest is the one between the 50th and the 10th, followed by that between the 75th and the 25th and then by the one between the 90th and the 50th. The differences between the three lines, however, are much less pronounced that in the case of family income. This is particularly true for the difference between 50th and 10th and 75th

and 25th percentile in that the second part of the samples are actually very close to each other.

Interestingly, in the case of consumption, as with income, the last two increase steadily for the sample period, while the first (the difference between the 50th and 10th) is flat. However, while the income differential between the 90th and 50th percentile of household (log) income increases by over 0.2 points, the increase over the whole period is half that size in the case of consumption. The increase in the difference between the 75th and 25th is again about 0.2 points in the case of income and less than 0.1 points for nondurable consumption.

4.6 Consumption Inequality Does Track Income Inequality: Beyond the Aggregate CE Interview Measures of Consumption

Given the measurement error in the CE, which has been widely documented in many studies, we are not sure how much faith to put in the results on consumption inequality in figures 4.6, 4.7, and 4.8. In this section, we use other measures from the CE where measurement error may be less of an issue, as well as data from the PSID. We do this to see if the patterns using these other data sets and consumption measures yield a different story relative to the CE interview data, but a consistent story among themselves. The results again are striking. Across all the other measures we consider—where, to reiterate, measurement error is less of an issue—consumption inequality has increased by only slightly less than the increase in income inequality.

4.6.1 CE Diary Data: Total Expenditure

As we mentioned in the data section, the CE survey is made of two components: the interview survey and the diary survey. While the figures we have considered so far are derived from the former, analogous figures can be constructed using the latter, especially after 1986, when the diary survey became comprehensive and includes virtually all consumption categories. In figure 4.9, we plot the standard deviation of log total consumption for the 1986–2010 period. Again, we adjust the data for differences in family size.

When comparing figures 4.6 and 4.9, two features emerge. First, the level of inequality measured in figure 4.9 is considerably larger. This is not particularly surprising because of the structure of the two surveys: the diary survey covers only two weeks and infrequently purchased items can induce a considerable amount of additional inequality in the cross section. What is most surprising, however, is the increase in inequality is considerably larger and more persistent in the diary survey. This second feature has been discussed extensively in Attanasio, Battistin, and Ichimura (2007) and in Attanasio, Battistin, and Padula (2010). Attanasio, Battistin, and Ichimura (2007), in particular, rule out a number of simple explanations for this differ-
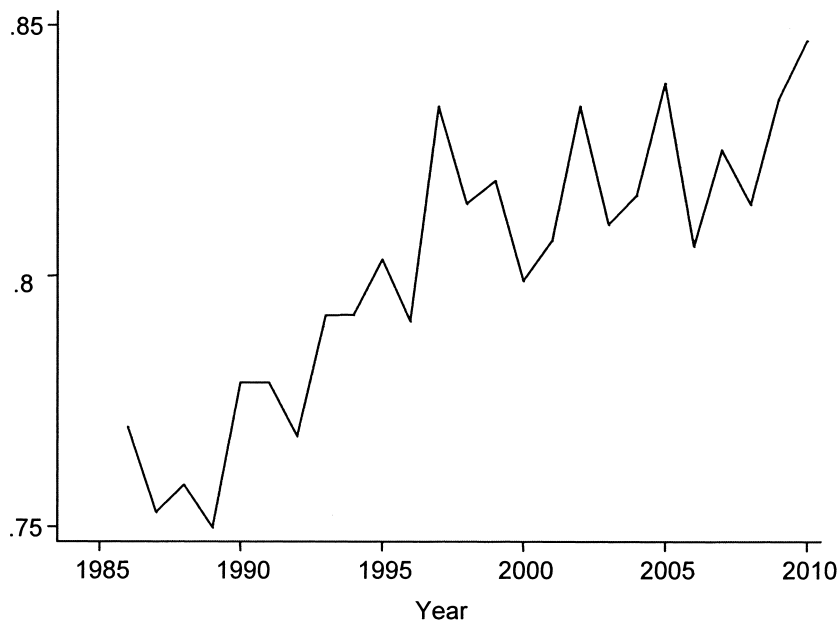
Fig. 4.9 Standard deviation log-equivalized nondurable consumption, CE diary data

ence, including a decrease in the frequency of shopping that could increase the number of zeros in a two-week diary.

The increase in the measure of the standard deviation of log expenditure using the diary data is 0.10, which is about one-half the increase in the measured increase in income inequality. Now, this measure may be understated if the expenditure categories that comprise the diary data are more likely to be necessities (like food). In that case, the increase in the consumption inequality from the diary survey would have to be scaled up by the income elasticity for the goods in the diary survey to get an overall measure of the change in inequality for total consumption. We do not do that here. Instead, we look at specific categories within the diary data, particularly food (and later on, entertainment). We turn to that analysis next. The take away from this section, however, is that within the goods in the diary data, there is a substantial increase in consumption inequality during the last three decades. The timing of the increases in consumption inequality from the diary data also matches closely the timing of the changes in income inequality over this time period.

4.6.2 CE Diary Data: Food

In figure 4.10, we explore the evolution of consumption inequality using food expenditures reported in the diary data. Researchers at the BLS believe that food data in the diary is measured with much less error than in the interview and, indeed, the main motivation for having the diary survey is to mea-
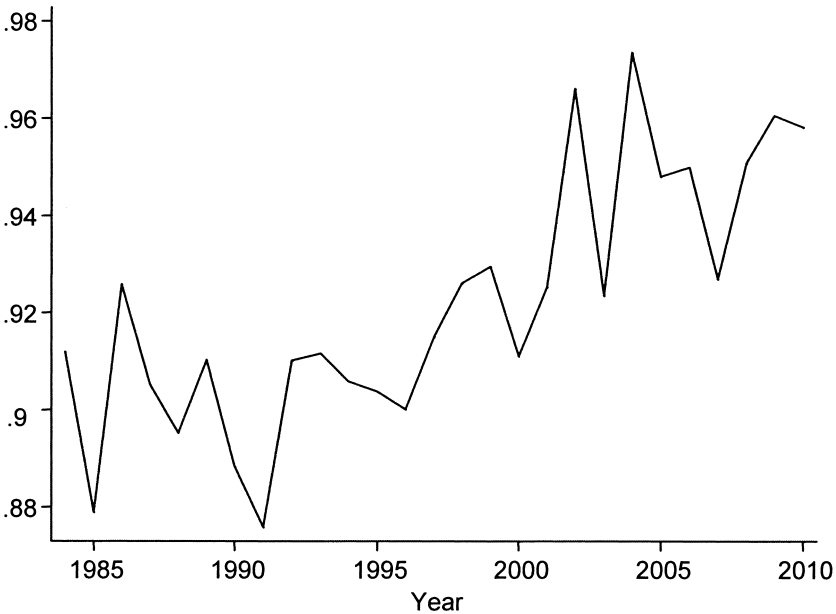
Fig. 4.10 Standard deviation log-equivalized food consumption, diary survey

sure more accurately what the BLS defines "frequently purchased items." As seen in figure 4.10, the standard deviation of log food expenditure at home in the diary data increased by between 6 and 8 percentage points. Estimates from Aguiar and Bils (2011) find that the income elasticity for total food spending is 0.5. Using this estimate, it follows that a simple back-of-the-envelope calculation suggests that total consumption inequality increased by between 12 and 16 percentage points over the period examined.[12]

Two things are of interest from the food results in figure 4.10. First, using the food data, the rise in consumption inequality was roughly 80 percent of the rise in income inequality during the sample period (0.16/0.20). Second, the rise in consumption inequality using the food data in the diary is higher than the rise in consumption using total consumption in the diary (0.16 vs. 0.10). Again, this is likely because the "total" nondurable expenditure measure in the diary is not a complete representation of nondurable expenditures. Given that it likely contains more reports of food expenditures than other nondurable expenditure measures, it may also need to be deflated using the income elasticity. As we show below, the results from the food data in the CE diary matches well the results on consumption inequality from the PSID.

12. This simple back-of-the-envelope calculation assumes that the income elasticity remains constant. If the elasticity is declining over time, our calculation is overestimating the increase in total consumption inequality.

4.6.3 CE Diary Data: Food versus Entertainment Spending

As we discussed in section 4.2.4, one way to deal with measurement error problems that can affect our overall measures of consumption is to focus on components of consumption for which the measurement issues are less severe and, possibly, stable over time. In this section we combine information on two such measures: consumption of food in the home and expenditure on entertainment goods and services (excluding durable goods).

In figure 4.11A, we plot the standard deviation of the log of the ratio of entertainment to food at home expenditure using the measures from the diary data. As argued in section 4.2.4, under certain conditions, this should be proportional to the standard deviation of log nondurable consumption, with the factor of proportionality depending on the difference between the income elasticities of the two commodities. As said, Aguiar and Bils (2011) estimate the income elasticity for total food spending estimates to be 0.5; they also estimate the elasticity for entertainment goods and services to be about 1.9. According to these estimates, one would then adjust the increase in figure 4.11A by a factor of 1.4 (1.9–0.5). The study by Blow, Lechene, and Levell (chapter 5, this volume) reports a similar expenditure elasticity of food at home of 0.5, but their estimate of the expenditure elasticity of entertainment is lower, at 1.5—implying a difference of 1—and hence suggesting that the increase in figure 4.11A would be in no need of adjustment.

As seen from figure 4.11A, the standard deviation of the log ratio of the two commodities increased over the whole sample period by about 0.15.[13] This figure would imply an increase in total nondurable consumption inequality (as measured by the standard deviation of logs) of the same size, if we take the Blow, Lechene, and Levell elasticities (while it would be lower using the Aguiar-Bils estimates). Again, this metric suggests that roughly 75 percent of the increase in income equality has translated to an increase in consumption inequality (0.15/0.20).

The ratio between food at home and entertainment expenditure can also be computed in the interview data. In figure 4.11B, we report the path of the standard deviation of the log of such a ratio, to be compared to figure 4.11A.

We find that, in the case of this ratio, the inequality measure that emerges from the interview CE data is considerably larger than what was obtained with total nondurable consumption expenditure and not inconsistent with the evidence coming from the diary survey. In particular, depending on when one starts counting, the increase in the standard deviation of the log ratio is between 0.15 and 0.25.

These results further suggest that the problems associated with the CE interview survey and the discrepancies between that and the diary survey might be attributed to difficulties in measuring certain specific commodities

---

13. In figure 4.11A, we do not need to adjust our results by changes in family size, as we are considering the (log) of a ratio.
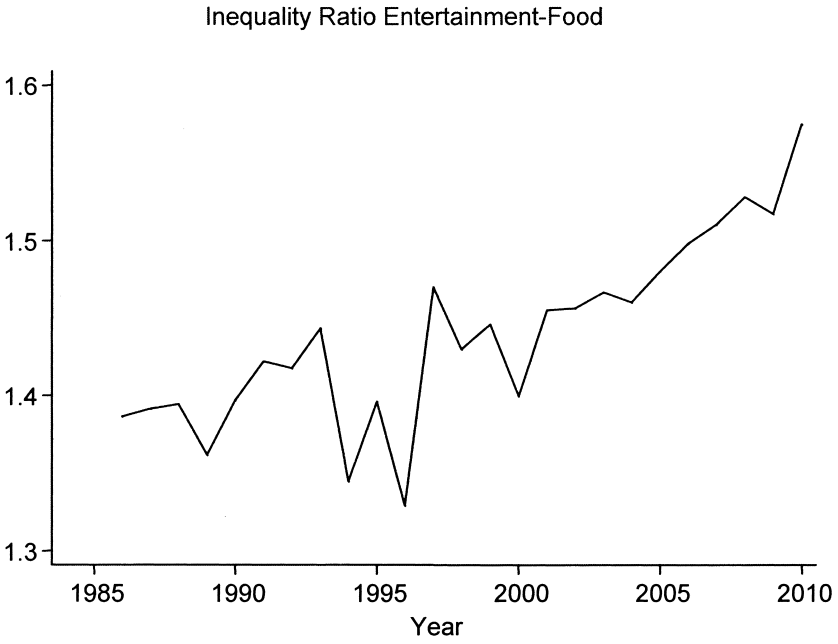
Inequality Ratio Entertainment-Food

Fig. 4.11A SD log ratio entertainment/food spending, CE diary data
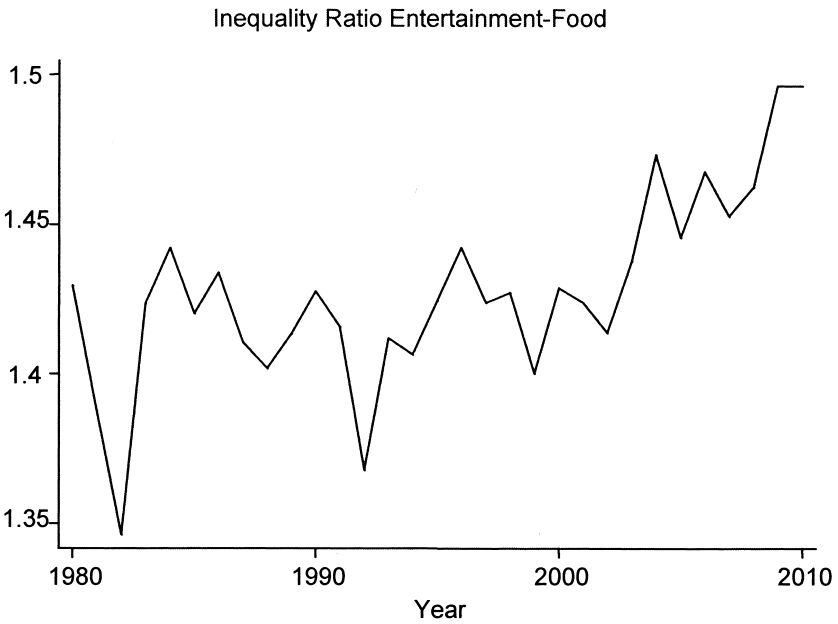
Inequality Ratio Entertainment-Food

Fig. 4.11B SD log ratio entertainment/food spending, CE interview data
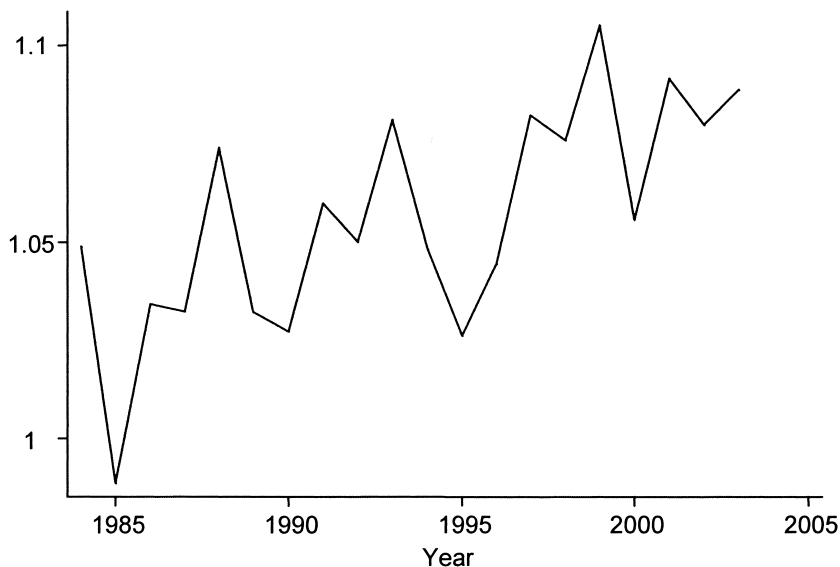
Fig. 4.12 Coefficient of variation stock of cars, CE interview survey

within the interview survey. It may be worth remembering that household income inequality in the CE interview survey increases as much as in other data sets, such as the PSID.

4.6.4 CE Interview Data: Stock of Car Holdings

If it is true that the different conclusions between the CE diary data and the CE interview data arise because some components of the CE interview data are fraught with changing nonclassical measurement error, one ideal exercise would be to find a category within the CE interview data that is measured with less error. To do this, we look at one such measure in the CE: the stock of cars owned by the household. This is an interesting exercise to perform for at least two reasons. First, the data on the value of the car stock in the CE seems to be of excellent quality, both in terms of the expenditure on cars and in terms of the composition of the stock of existing cars. Second, cars are durables and large. Moreover, adjusting the stock of cars is subject to transaction and adjustment costs. One would therefore guess that decisions about cars reflect long-run expectations about permanent income.

When analyzing the stock of cars one has to take into account the fact that there are a number of households that do not own a car. This prevents us from computing the log of the value of cars for these households. To deal with this issue we follow two different approaches. In figure 4.12, we plot the coefficient of variation of the stock of cars, defined as the standard deviation

Fig. 4.13 SD log stock of cars, CE interview data

of car values divided by its mean. Both the mean and the standard deviation are computed including the zeros for households who do not own a car. In figure 4.13, instead, we only use households who own at least one car and plot the standard deviation of the log value for these households.

When looking at the coefficient of variation we see an increase of about 0.1, which happens especially in the first part of the sample period. The increase documented in figure 4.12 for the standard deviation of logs is actually larger, at almost 0.2.

4.6.5 PSID Data: Food and Nondurable Expenditures

Figure 4.14 uses PSID data and plots consumption inequality (as measured by the standard deviation of the logs) against time. Here we use the five different consumption measures we can construct from the PSID: (a) food consumption over the 1980–2010 period; (b) the so-called 70 percent measure available over the 1998–2010 period (with spending on health and education); (c) the 70 percent measure excluding spending on health and education;[14] (d) the imputation based on the CEX estimation of the food demand function over the 1980–2010 period (BPP measure), and (e) the imputation based on the difference between income and the change in assets over the 1980–2010 period (Ziliak measure).

14. Attanasio and Pistaferri (2013) use the consumption measure available in the redesigned PSID 1998–2010 to impute consumption inequality for the years before 1999.
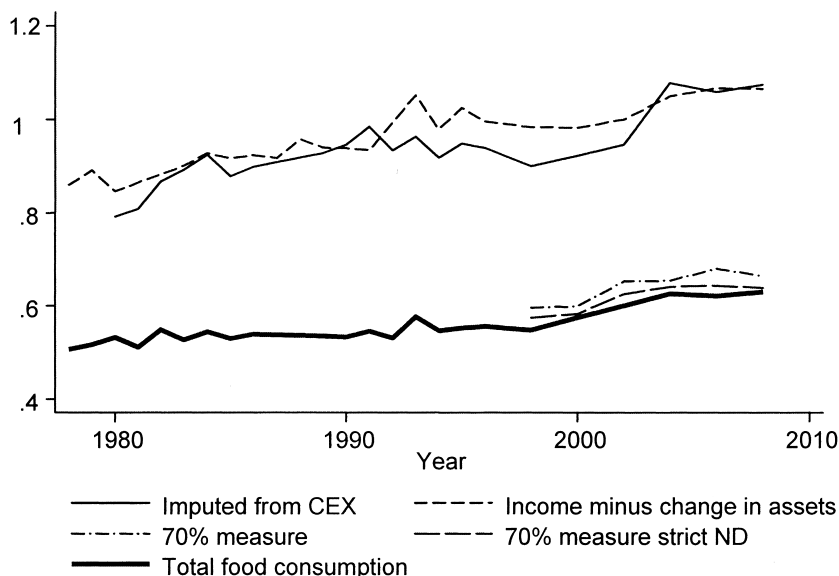
Fig. 4.14 Trends in consumption inequality from the PSID

A number of features of figure 4.14 are worth noting. First, the five measures rank as we might expect: food has less variance than the 70 percent measure, which in turn displays less variance than the imputed (more comprehensive) measures. Second, it is remarkable that the two more comprehensive measures—despite being obtained with two completely different imputation procedures that use completely different modules of the survey—display very similar trends and levels. Third, both imputed measures show a considerable increase in consumption inequality over the 1980s and into the early 1990s, of almost 0.2. In other words, the composite measures from the PSID show that consumption inequality has tracked income inequality nearly exactly.[15] Fourth, the increase in inequality in food consumption in the PSID is nearly identical to the increase in inequality in food consumption from the CE diary survey. This is reassuring given that there is no evidence that systematic measurement error has been changing in the PSID. Fifth, the increase in total consumption inequality based on the PSID food data is 0.2 (0.10/0.5). Again, to get this, we scale up the estimate based upon the total food income elasticity of 0.5. Finally, the consumption inequality measure using the new consumption data available for the 1998–2010 period confirms the general increasing trends.

---

15. Both imputed measures, of course, do not identify the levels of consumption inequality because of imputation errors. However, if such errors are stationary, then the imputed measures are unbiased estimates of the *change* in consumption inequality over the sample period.
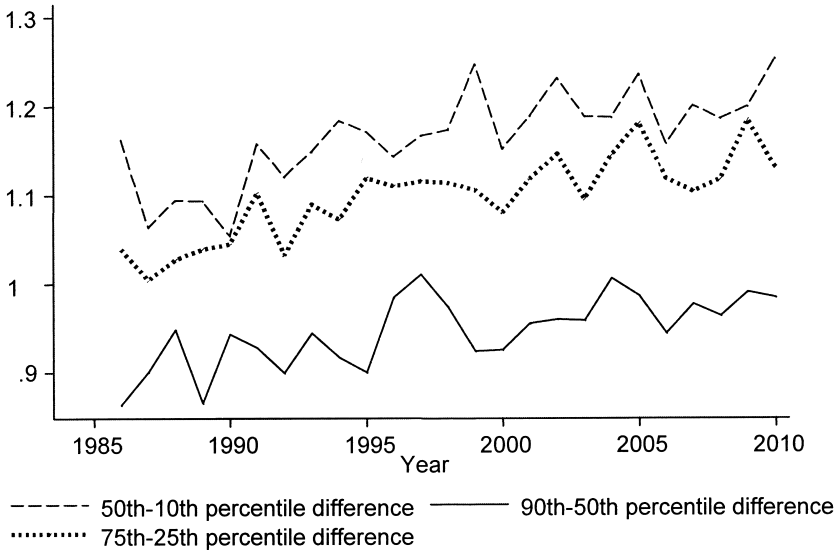
Fig. 4.15 Movements in the distribution of nondurable consumption, CE diary survey

Taken together, the results from the PSID data shown in figure 4.14 is that consumption inequality and income inequality tracked each other nearly identically during this time period. We highlight one more aspect emerging from the PSID data. Given the results in figure 4.1 and the results in figure 4.14, we can compute the simple correlation between the total consumption inequality measures based on the imputation procedures and the income inequality measures from the PSID over this time period. A simple ocular examination of the figures shows that the time-series patterns of the inequality measures are very similar. Both inequality measures (imputed consumption inequality and income inequality) increased during the 1980s, leveled off during the 1990s, and increased further in the first decade of the twenty-first century. This is reflected in the simple correlation measures. For example, the simple time-series correlation between the PSID-imputed consumption inequality measure (based on the CE elasticities) and the PSID income inequality measure was 0.69, and that between the Ziliak measure and the income inequality measure was 0.88.

4.7 Different Dimensions of Inequality

In this section, we use our various other consumption measures within the CE and the PSID to explore the evolution of consumption inequality at different points in the distribution. In figure 4.15, we consider the percentile differentials we have used so far to study the evolution of the distribution of
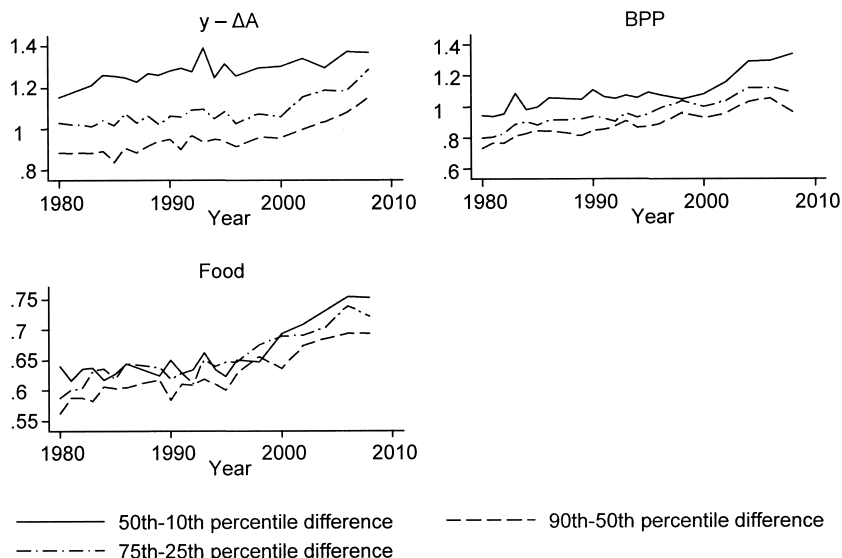
Fig. 4.16 Consumption inequality in different parts of the distribution

consumption. In figure 4.15, we plot the same differentials between percentiles in figure 4.8, but instead use the diary total nondurable consumption measure. Notice the 50/10 differential, which was substantially flatter in figure 4.8 and increases throughout the sample period as does the 75/25 differential. The 90/10 differential, which in figure 4.8 was the one that increases the most, increases but slightly less than the other differentials. Again, the changes in consumption inequality throughout the distribution using the CE diary data better matches the time-series trends in income inequality at similar percentile points.

Figure 4.16 explores whether the changes in consumption inequality within the PSID are coming from the bottom or the top part of the distribution. In the top-left panel, we show the results using the Ziliak measure (defined as the difference between income and the change in assets). First, in the 1980–1995 period the rise in inequality is explained by movements in both tails—the 75th–25th percentile difference is indeed very stable. The rise in inequality at the bottom is, if anything, more pronounced than the rise in inequality at the top. However, in the second half of the sample, the top part of the distribution starts detaching itself more dramatically from the rest—the rise in the 90th–10th percentile difference and in the 75th–25th percentile difference is very pronounced, while the 50th–10th percentile difference remains stable. Interestingly, trends in food consumption inequality (bottom panel) are different—there is less heterogeneity in the movement of different parts of the food consumption distribution. Partly reflecting
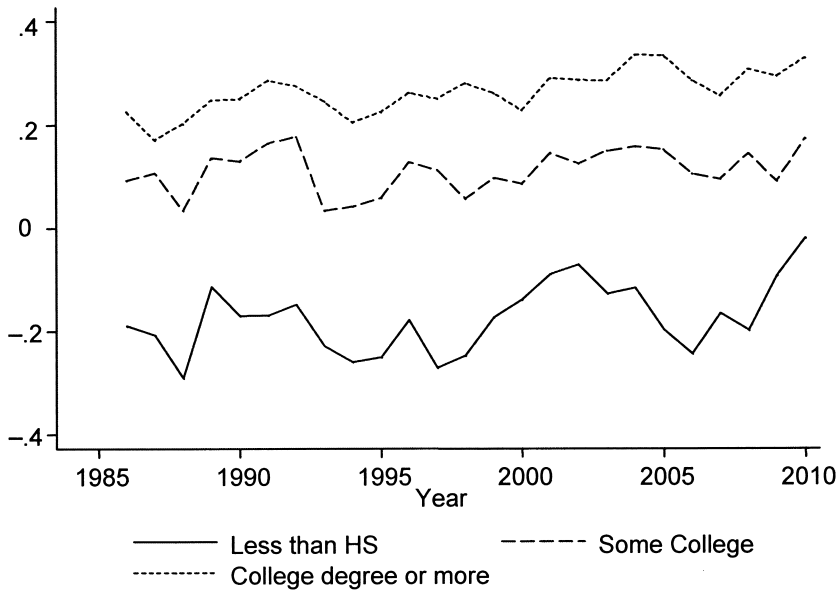
Fig. 4.17 Consumption inequality across skill groups (averages relative to those with a high school degree), CE diary survey

this, the percentile differences computed using the BPP's imputed measure of consumption (top-right panel) display less stark trends.

In addition to the overall sample and its distribution, it is also interesting to cut the sample by skill levels of the household level. We know from a large literature in labor economics that the return to education has increased dramatically over the period we are considering and that, in the case of income and wages, there have been large increases in inequality both across skill level and within skill levels (possibly reflecting changes also in the price of unobservable skills).

In figure 4.17, we divide the CE diary sample into four groups on the basis of the education of the household head: the first group is formed of households headed by an individual with a college degree or more, the second group by households headed by an individual with some college experience, the third by high school graduates, and the fourth by high school dropouts. We then express, for every year, average log consumption as difference from average log consumption of the third group (the high school graduates).

The graphs in figure 4.17 indicate a steady increase in the return to education as measured by the difference between high school graduates and college graduates (the top line). The differential between college graduate and high school graduates increases from about 0.2 to about 0.33 at the end of the sample. These changes in the first part of the sample (across different years of birth cohorts) were studied in Attanasio and Davis (1996) in rela-
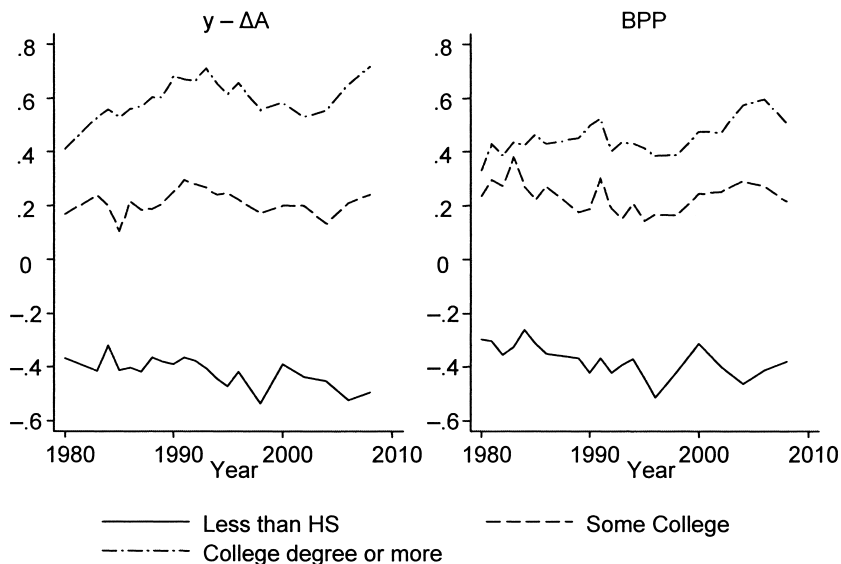
Fig. 4.18 Consumption inequality across skill groups (averages relative to those with a high school degree), PSID consumption measures

tion to similar changes in relative wages. Similar patterns for the PSID are shown in figure 4.18.

Having considered changes across education groups, we now look at changes within education groups. In particular, in figures 4.19 and 4.20, we plot the standard deviation of log consumption within each of the four education groups listed above for the CE diary data and the PSID consumption measures, respectively. Within-skill group consumption inequality increased dramatically across all skill groups in both surveys.

4.8 Leisure Inequality: Time-Use Surveys

The results in the prior section suggest strongly that consumption inequality has tracked income inequality rather closely over the last thirty years. Does this mean that the actual inequality in well-being tracked income inequality over this period? In the standard model, utility is a function of both consumption and leisure time. It is therefore natural to look at the changes in inequality in leisure in conjunction with the changes in inequality in consumption so as to get a better measure of changes in the inequality of total well-being across individuals.[16]

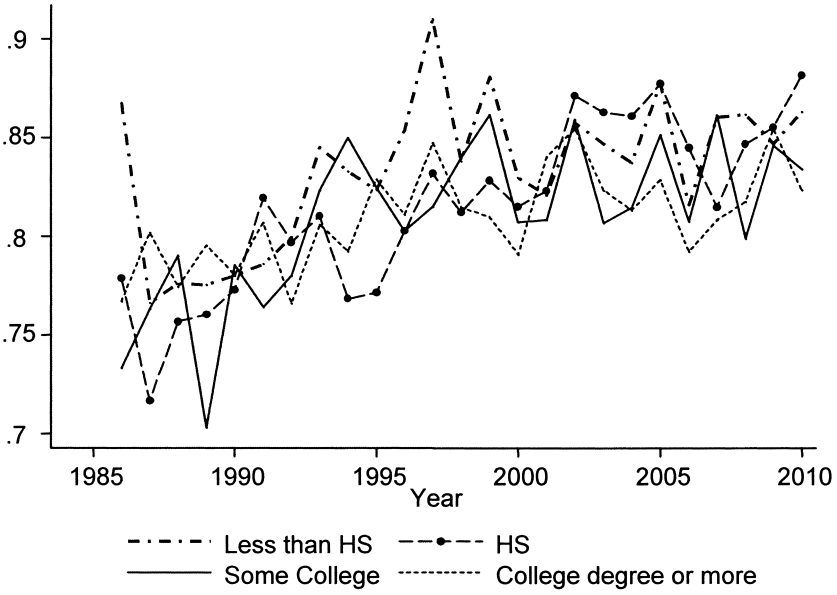16. The inequality in leisure has been explored by Aguiar and Hurst (2007, 2009).

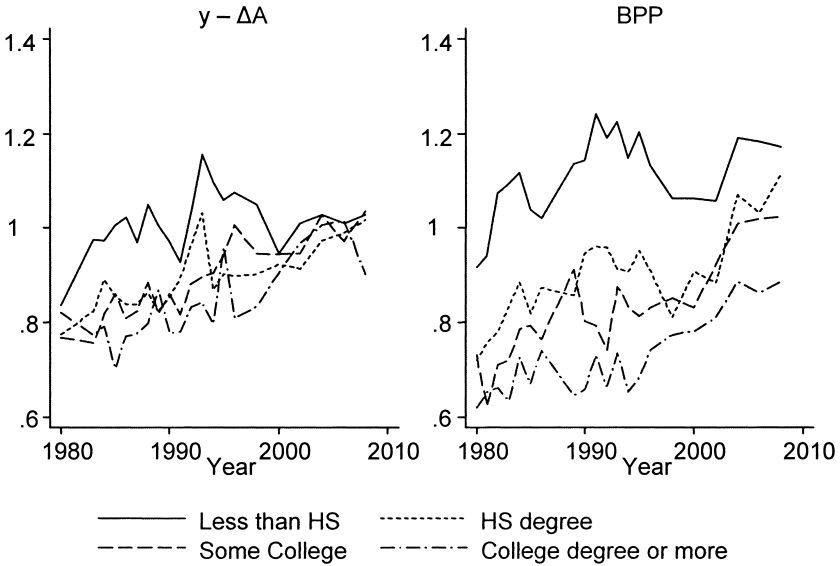**Fig. 4.19** Consumption inequality within skill groups, CE diary survey

Fig. 4.20 Consumption inequality within skill groups, PSID data

Table 4.1   **Leisure measures (hours per week) by group, 1985 and 2003–2007**

	Leisure (hours per week)		
Sample	1985	2003–2007	Change
Men			
Low educated	36.6	39.1	2.5
High educated	34.4	33.2	–1.2
Women			
Low educated	35.0	35.2	0.2
High educated	32.2	30.3	–1.9

To make the results from the time-use surveys comparable with some of the results in the prior sections, we explore the changes in leisure inequality across skill groups. Table 4.1 explores the hours per week spent in leisure for low- and high-educated men and women in both 1985 and 2003–2007. As discussed above, our measure of leisure includes the actual time the individual spends in leisurely activities like watching television, socializing with friends, going to the movies, and so forth.

A few things are of note from table 4.1. First, in 1985, low-educated men took only slightly more hours per week of leisure than high-educated men. As above, we define high educated as those with more than twelve years of schooling. A similar pattern holds for women. However, by 2007, the leisure differences between high- and low-educated men are substantial. Specifically, low-educated men experienced a 2.5 hours per week gain in leisure between 1985 and 2007. High-educated men, during the same time period, experienced a 1.2 hour per week decline in leisure. The new effect is that leisure inequality increased dramatically after 1985. Again, similar patterns are found for women.

Figure 4.21 shows the distribution of leisure for low-educated men in 1985 (dashed line) and 2003–2007 (solid line). Figure 4.22 shows similar patterns for higher-educated men. Most of the increase in leisure occurred as a result of changes in the upper tail of the leisure distribution. A greater share of low-educated men in 2003–2007 are taking more than fifty hours per week of leisure than in 1985. This is not the case for higher-educated men. If anything, there is slightly lower proportion of higher-educated men taking more than fifty hours per week of leisure in 2003–2007 than there was in 1985.

The patterns shown in table 4.1 and figures 4.21 and 4.22 show that the overall inequality measures between high- and low-skilled individuals becomes muddied when one combines the results for consumption and leisure. While it is true that the consumption of the highly educated has grown rapidly relative to the consumption of the low educated, it is also true that leisure time of the low educated has grown rapidly relative to the leisure time of the highly educated. In order to make overall welfare calculations,

Kernel Density of Leisure Time in 1985 and 2003-2007
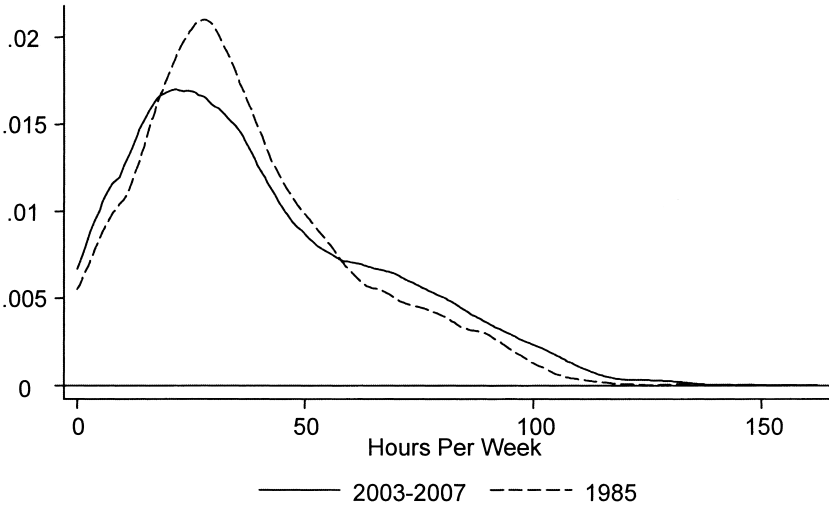Sample: Low Educated Males Between the Ages of 25 and 65

Hours Per Week

——— 2003-2007   – – – – 1985

Fig. 4.21 Kernel density of leisure time for low-educated men (1985 and 2003–2007)

Kernel Density of Leisure Time in 1985 and 2003-2007
Sample: Low Educated Males Between the Ages of 25 and 65

Hours Per Week
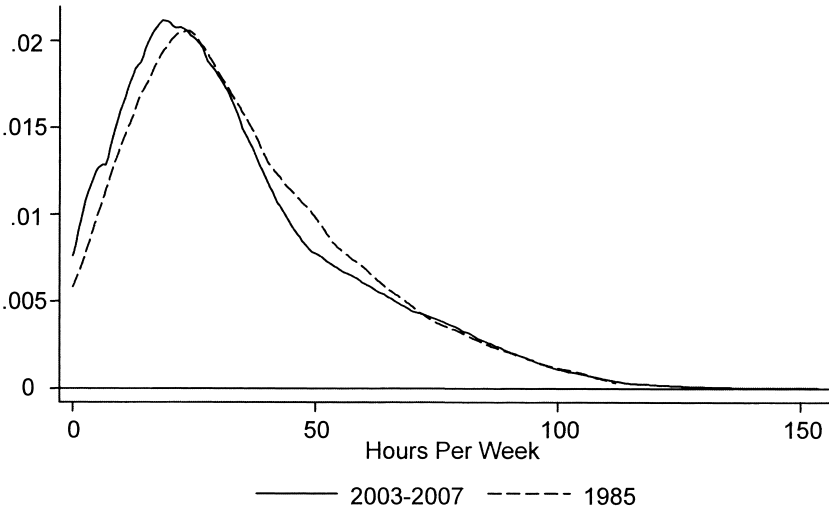
——— 2003-2007   – – – – 1985

Fig. 4.22 Kernel density of leisure time for highly educated men (1985 and 2003–2007)

one needs to take a stance on how the leisure time is valued. But, as long as leisure has some positive value, the increase in consumption inequality between high- and low-educated households during the past few decades will overstate the true inequality in well-being between these groups.

4.9 Conclusions

In this chapter, we have documented thoroughly that the increase in income inequality was matched by an increase in consumption inequality of comparable magnitude. In particular, between 1980 and 2010, the standard deviation of log income increased by roughly 0.2 percentage points. Across our various preferred measures, the standard deviation of log consumption increased by roughly 0.10 to 0.2 percentage points with most of the estimates being in the 0.15 to 0.20 range.

The main innovation of the chapter is to show that the data on the increase in consumption inequality within the United States is very robust to alternative measures of consumption and across alternative data sets. The one outlier in terms of estimates of consumption inequality is total nondurable expenditure from the CE interview survey (where no attempt is made to adjust for measurement error). For this measure, the standard deviation of log consumption increased by only 0.06 percentage points—with most of the increase coming before 1982. As shown by Aguiar and Bils (2011), the CE interview data is plagued with nonclassical measurement error that biases estimates of consumption inequality downward. Given that many researchers used the nondurable consumption measure from the CE interview data as their primary measure of consumption inequality, they have naturally concluded that consumption inequality has not increased much over the last thirty years. However, some researchers have found rising consumption inequality using other measures. These various estimates have lead researchers to conclude that the extent of the increase in consumption inequality is still an open debate among economists.

Our results in this chapter, however, show that such a conclusion is unwarranted. Across every other measure of consumption we analyzed, consumption inequality increased substantially. Some of these measures came from the CE diary survey, like food and entertainment spending (where measurement error in those categories has been found to be less of a problem). Some measures came from the CE interview survey, like the stock of cars (where quality of data appears fairly high). Finally, some of our measures come from the PSID (where systematic measurement has not been found to be a problem). Not only do these other measures of consumption inequality mirror the overall change in income inequality, the timing of the changes also line up very closely.

Within the context of the CE, it is clear that the interview survey is plagued by some serious measurement problems. These can arise from a variety of

sources. In terms of the discussion of inequality, however, the evidence we have presented seems to indicate that these problems are not caused by the specificity of the interview sample and by the fact that it might exclude, for a variety of reasons, households from the extremes of the income distribution. The main evidence we provide in this respect is that we do find evidence of increasing inequality in the interview survey. It is apparent in income, in the value of the car stock, and in the ratio of food at home to entertainment. It is therefore likely that the issue lies with the measurement of specific items or the degree to which expenditures are fully reported by certain groups. This evidence can be valuable in the redesign of the CE survey.

From a methodological point of view, more work is needed to formalize the use of several components of the CE survey, both to do imputation in other data sets and to make inferences about overall (nondurable) consumption. We also believe more work is needed in order to understand the consequences for welfare of changes in relative prices and the consequent shifts in commodity-demand patterns. Finally, if would also be helpful to think about overall measures in well-being by thinking about the overall change in consumption inequality jointly with the overall change in leisure inequality. We think future work should attempt to make a composite change in the inequality of well-being by formally linking the two measures.

References

Aguiar, Mark, and Mark Bils. 2011. "Has Consumption Inequality Mirrored Income Inequality?" NBER Working Paper no. 16807, Cambridge, MA.

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.

———. 2009. *The Increase in Leisure Inequality: 1965–2005*. Washington, DC: American Enterprise Institute for Public Policy Research.

Attanasio, Orazio, Erich Battistin, and Hidehiko Ichimura. 2007. "What Really Happened to Consumption Inequality in the United States?" In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, Studies in Income and Wealth, vol. 67, edited by Ernst R. Berndt and Charles R. Hulten, 515–44. Chicago: University of Chicago Press.

Attanasio, Orazio, Erich Battistin, and Mario Padula. 2010. *Inequality in Living Standards since 1980: Income Tells Only a Small Part of the Story*. Washington, DC: AEI Publishing.

Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104 (6): 1227–62.

Attanasio, Orazio, and Nicola Pavoni. 2011. "Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption." *Econometrica* 79 (4): 1027–68.

Attanasio, Orazio, and Luigi Pistaferri. 2013. "What Do We Know about Consumption Inequality? Some Evidence Using the New PSID Consumption Measure." Working Paper, Stanford University.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008. "Trends in US Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics* 90 (2): 300–23.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2004. "Imputing Consumption in the PSID Using Food Demand Estimates from the CEX." IFS Working Paper no. W04/27, Institute for Fiscal Studies.

———. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5): 1887–921.

Blundell, Richard, and Ian Preston. 1998. "Consumption Inequality and Income Uncertainty." *Quarterly Journal of Economics* 113 (2): 603–40.

Browning, Martin, and Thomas Crossley. 2009. "Are Two Cheap, Noisy Measures Better Than One Expensive, Accurate One?" *American Economic Review* 99 (2): 99–103.

Cutler, David M., and Lawrence F. Katz. 1992. "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s." *American Economic Review* 82 (2): 546–51.

Deaton, Angus, and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *Journal of Political Economy* 102 (3): 437–67.

Garner, Thesia I., and Atsushi Maki. 2004. "The Gap between Macro and Micro Economic Statistics: Estimation of the Misreporting Model Using Micro-Data Sets Derived from the Consumer Expenditure Survey." Working Paper no. 3, Econometric Society.

Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante. 2010. "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006." *Review of Economic Dynamics* 13 (1): 15–51.

Krueger, Dirk, and Fabrizio Perri. 2006. "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory." *Review of Economic Studies* 73 (1): 163–93.

Meyer, Bruce D., and James X. Sullivan. 2009. "Five Decades of Consumption and Income Poverty." NBER Working Paper no. 14827, Cambridge, MA.

Padula, Mario. 1999. "Euler Equations and Durable Goods." CSEF Working Paper no. 30, Centre for Studies in Economics and Finance.

Parker, Jonathan, Annette Vissing-Jorgensen, and Nicolas L. Ziebarth. 2009. "Inequality in Expenditure in the Twentieth Century." Working Paper, Northwestern University.

Slesnick, Daniel T. 2001. *Consumption and Social Welfare: Living Standards and Their Distribution in the United States.* Cambridge: Cambridge University Press.

Ziliak, James. 1998. "Does the Choice of Consumption Measure Matter? An Application to the Permanent-Income Hypothesis." *Journal of Monetary Economics* 41 (1): 201–16.