

NBER WORKING PAPER SERIES

HAS CONSUMPTION INEQUALITY MIRRORED INCOME INEQUALITY?

Mark A. Aguiar
Mark Bills

Working Paper 16807
<http://www.nber.org/papers/w16807>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2011

We thank Yu Liu for outstanding research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Mark A. Aguiar and Mark Bills. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Has Consumption Inequality Mirrored Income Inequality?

Mark A. Aguiar and Mark Bills
NBER Working Paper No. 16807
February 2011
JEL No. E2,E21,J31

ABSTRACT

We revisit to what extent the increase in income inequality over the last 30 years has been mirrored by consumption inequality. We do so by constructing two alternative measures of consumption expenditure, using data from the Consumer Expenditure Survey (CE). We first use reports of active savings and after tax income to construct the measure of consumption implied by the budget constraint. We find that the consumption inequality implied by savings behavior largely tracks income inequality between 1980 and 2007. Second, we use a demand system to correct for systematic measurement error in the CE's expenditure data. Specifically, we consider trends in the relative expenditure of high income and low income households for different goods with different income (total expenditure) elasticities. Our estimation exploits the difference in the growth rate of luxury consumption inequality versus necessity consumption inequality. This "double-differencing," which we implement in a regression framework, corrects for mis-measurement that can systematically vary over time by good and income group. This second exercise indicates that consumption inequality has closely tracked income inequality over the period 1980-2007. Both of our measures show a significantly greater increase in consumption inequality than what is obtained from the CE's total household expenditure data directly.

Mark A. Aguiar
Department of Economics
University of Rochester
P.O. Box 270156
Rochester, NY 14627-0156
and NBER
mark@markaguiar.com

Mark Bills
Department of Economics
University of Rochester
Rochester, NY 14627
and NBER
bils@troi.cc.rochester.edu

1 Introduction

We revisit the issue of whether the increase in income inequality over the last 30 years has translated into a quantitatively similar increase in consumption inequality. Contrary to several influential studies discussed below, we find that consumption inequality has closely tracked income inequality over the period 1980-2007. Like most of the previous literature that argues the opposite, we base our conclusions on the Consumer Expenditure Survey's (CE) interview survey. However, we focus on two new measures of consumption that under our stated assumptions adjust for the systematic measurement error in the CE. The first measure is the CE's data on savings, from which we calculate consumption via the budget constraint. The second is a demand system, from which we estimate relative consumption growth using relative expenditures on luxuries and necessities. Both measures show a substantial increase in consumption inequality, similar in magnitude to the increase in income inequality.

An influential paper by [Krueger and Perri \(2006\)](#), building on related work by [Slesnick \(2001\)](#), uses the CE to argue that consumption inequality has not kept pace with income inequality.¹ In an exercise similar to Krueger and Perri's, we show that relative after-tax income inequality increased by 33 percent (.33 log points) between 1980 and 2007, where our conservative measure of income inequality is the ratio of those in the 80-95th percentiles to those in the 5-20th percentiles.² The corresponding increase in consumption inequality for the same two groups is 17 percent.

We reassess these facts using two alternative measures. Our first exercise is simply budget constraint accounting. The mirror image of the differential trends between income and consumption inequality in the CE is a growing gap in savings favoring high income households. Based on reported consumption expenditures, the high income group increased their savings rate from 25 percent to 38 percent between 1980 and 2007, while the low income group maintained a savings rate of roughly -30 percent over this period. The implied savings rates using CE income and consumption are implausible. For the overall mean, the implied savings rate in the CE increases from 10 percent in 1980 to over 20 percent in 2007. This contrasts with savings out of disposable income reported in the flow of funds accounts, which falls from 10 percent to almost zero, as well as is inconsistent with other micro data sets (see [Bosworth and Anders, 2008](#) and [Bosworth and Smart, 2009](#)). This discrepancy is in line with the well documented decline in aggregate consumption reported in the CE relative to

¹For other contributions to this literature, see [Blundell and Preston \(1998\)](#), [Blundell et al. \(2008\)](#), and [Heathcote et al. \(2010\)](#).

²For the period 1980-2004, [Krueger and Perri \(2006\)](#) report a log change in the 90/10 income ratio of approximately 0.36 for income, and 0.16 for consumption.

NIPA (e.g., [Garner et al., 2006.](#))

In addition to expenditures and income, the CE asks detailed questions on savings flows directly. These questions include net payments of loans, changes in deposit balances, purchases of stocks, etc. The average reported savings rate in the CE declines over time, consistent with flow of funds and NIPA, but in contrast to the savings rate implied by the CE's consumption data. Calculating implied expenditure as income minus savings, we obtain an increase in relative consumption inequality of 28 percent, which is substantially closer to the relative change in disposable income of 33 percent than that implied by reported expenditure. The CE's savings measures are noisy (particularly regarding new mortgages), and so we view them primarily as a consistency check on the reported consumption data, and only secondarily as an independent measure of consumption itself.

Our preferred measure of consumption inequality uses the CE's expenditure data, but corrects for systematic measurement error. Our modeling of measurement error is fairly general. In particular, we allow for time-dependent multiplicative measurement error that is good specific as well as income-group specific. The former allows for the mis-measurement of particular goods to vary over time, such as the possibility that the under-reporting of luxuries has increased relative to the under-reporting of necessities. The latter allows for the measurement to be income-group specific, such as the possibility that the under-reporting of expenditure of high income households across all goods has increased relative to the under-reporting by low income households. This modeling of measurement error captures systematic mis-measurement that is correlated with the characteristics of the good and the income-characteristics of the households. We also allow for mis-measurement at the level of good-income group interaction (clothing of the rich versus clothing of the poor), but restrict this joint mis-measurement to be independent of the characteristics of the goods, in particular the good's income (total expenditure) elasticity.

Our estimation procedure consists of two steps. First, we estimate good-specific total expenditure elasticities using a simple log-linear demand system. To do this, we use the 1972-73 CE, separating our first stage sample from the post-1980 period of focus. In the second stage, we consider the difference in expenditure growth across goods and across income groups. To see how this approach addresses mis-measurement, take expenditures on food at home (a necessity) versus nondurable entertainment (a luxury) as an example. The relative expenditure on food at home across income groups remained essentially constant between 1980 and 2007. Given a non-zero estimated expenditure elasticity of 0.5 for food at home, this suggests zero change in relative total expenditures. While comparing the same good across income groups controls for (multiplicative) mis-measurement of food in

each period, it does not control for the possible mis-measurement correlated with income. For this, we can add a second good, nondurable entertainment. Over the same period, the high income-low income ratio of expenditure on nondurable entertainment increased by 0.75 log points. Given an estimated elasticity of 2.0 for entertainment, this implies a change in relative expenditure of 38 percent. Again, this controls for good-specific measurement error, but not mis-measurement correlated with income. However, any mis-measurement that is specific to income groups, but that is uniform across goods, can be eliminated by differencing across goods. That is, the difference in relative expenditure growth rates will equal the difference in expenditure elasticities times the change in total expenditure inequality (plus an idiosyncratic error term). Solving this equation, the relative growth in these two goods implies a change in consumption inequality of 48 percent. Our procedure is thus a difference-in-difference estimate, where one difference eliminates good-specific mis-measurement and the second difference eliminates income group-specific mis-measurement.

While food and entertainment are interesting due to their extreme income elasticities, the CE data contains expenditure on many goods. We therefore implement this procedure using all goods in a regression framework. Our estimates suggest that consumption inequality increased by 30 percent between 1980 and 2007, approximately the same as the change in income inequality, and even larger than that obtained from the budget constraint accounting. We find this estimate is stable across different subsets of goods, different weighting schemes across goods, and alternative first-stage elasticity estimates.

We also consider trends in inequality in different sub-periods. We find that income inequality increased by 20 percent between 1980 and the mid-1990s, and then by an additional 13 percent between 1995 and 2007. The inequality in reported CE expenditure increased by 13 percent in the first sub-period, and then by 4 percent in the latter half of the sample. Reported consumption inequality does not keep pace with income inequality in either sub-period. Using our demand system estimates, we find that consumption inequality increased by roughly 19 percent between 1980 and the mid-1990s, and then by additional 13 percent through 2007, for a total increase of 32 percent. These estimates more closely track the profile of income inequality, with a larger increase in the 1980s, and a smaller but still significant increase thereafter.

We are not the first to reassess trends in consumption inequality, particularly with a focus on mis-measurement of CE interview expenditures. [Attanasio et al. \(2005\)](#) use the diary component of the CE to correct for mis-measurement in the interview survey, and document a large increase in consumption inequality. Their analysis does not extend back to the 1980s due to data limitations. Our analysis uses interview survey data, but brings in

data on savings as well as the differential trends across goods and income groups to address mis-measurement. Our paper is complementary to an independent effort by [Parker et al. \(2009\)](#), who focus on the gap between CE expenditures and those reported by NIPA to obtain a corrected estimate of consumption inequality.

There is a large literature concerning consumption inequality that precedes or is not focused on the issues raised by Slesnick and Krueger and Perri. An important paper by [Attanasio and Davis \(1996\)](#) documents that the increase in the college premium observed for wages in the 1980s is mirrored by similar increases in consumption inequality. However, [Attanasio and Davis \(1996\)](#) do not address the relative trends within education groups, which is where [Krueger and Perri \(2006\)](#) show the conflict between income and consumption inequality trends is starkest. Other important papers in this earlier literature include [Cutler and Katz \(1992\)](#) and [Blundell and Preston \(1998\)](#). [Sabelhaus and Groen \(2000\)](#) also discuss mis-measurement in the context of the relationship of consumption and income. For trends in inequality for a number of countries and time periods, see the papers collected in [Krueger et al. \(2010\)](#). There is also a large literature on consumption versus income inequality over the life cycle, starting with [Deaton and Paxson \(1994\)](#).³ These papers often use the CE for consumption data, and are therefore subject to the measurement error problems addressed in this paper. We leave the question of whether our approach has implications for trends in life cycle inequality to future research.

The remainder of the paper is organized as follows. Section 2 describes the data set; section 3 analyzes the CE’s savings data; section 4 performs our main demand-system analysis; section 5 contains a discussion of potential mis-specification and concludes.

2 Data

In this section we describe our data set, leaving to the data appendix a more detailed discussion of variable construction and our sample. Our data is from the Consumer Expenditure Survey’s interview sample. This is a well known consumption survey that has been conducted continuously since 1980. We also use the earlier waves of 1972 and 1973 for estimation of the demand system. (The survey was not conducted between 1973 and 1980.)

The survey is large, consisting of over 5,000 households in most waves. Each household is assigned a “replicate” weight designed to map the CE sample into the national population,

³See also, [Storesletten et al. \(2004\)](#), [Heathcote et al. \(2005\)](#), [Guvenen \(2007\)](#), [Huggett et al. \(2009\)](#), and [Aguiar and Hurst \(2009\)](#).

which we use in all calculations. An initial interview collects information about household characteristics as well as other baseline information. Each household is then re-interviewed once a quarter for up to four consecutive quarters, which we refer to as interviews 1 through 4. Each interview records expenditures on detailed categories over the preceding three months. The final interview also updates income and demographic information for the preceding 12 months. For income and demographics, we use the responses from the last interview. Income, expenditure, and savings variables are all recorded at the household level. Demographics such as age, sex, and education, are those reported by the “reference person,” who is identified by the response to the question who owns or rents the house. We define income and expenditures at the household level, rather than creating adult equivalence scales. However, when estimating household demand equations we control for demographic dummy variables that reflect the number of household members, number of household earners, and the reference member’s age.

On the income side, we use the CE measures of total household labor earnings, total household income before tax, and total household income after tax. These variables are reported in the last interview and cover the previous 12 months. Before-tax income in the CE includes labor earnings, non-farm or farm business income, social security and retirement benefits, social security insurance, unemployment benefits, workers’ compensation, welfare (including food stamps), financial income, rental income, alimony and child support, and scholarships. Our measure of before-tax income is that reported in the CE, but we add in food as pay and other money receipts (lump sum receipts from estates, gambling winnings, proceeds from the sale of personal items, etc.). For consistency, as we count receipts of alimony and child support as income, we subtract off payments of alimony and child support. Finally, as rental equivalence is a consumption expenditure for home owners, but not out of pocket housing expenses (mortgage interest, property taxes, insurance, repairs and maintenance, etc.), we include rental equivalence minus out-of-pocket housing costs as part of before tax income as well. Our measure of after tax income deducts personal taxes from our measure of before tax income. These taxes are federal income taxes, state and local taxes, and payroll taxes. (We include government retirement and railroad retirement contributions as payroll taxes.) Note that federal income taxes can be negative, especially as they capture earned income credits. We consider an alternative measure of after tax income by replacing self-reported federal income taxes with taxes calculated from the NBER’s TAXSIM program. We discuss those results as a robustness check in section 3.

The CE asks respondents a number of questions on active savings. For example, they record net flows to savings accounts, purchases of assets (including houses and business),

payments of mortgages, payments of loans, purchases and sales of vehicles, etc. The detailed components of savings are reported in the data appendix. The CE records the total outstanding credit balances in the 1st and 4th interviews, which are 9 months apart. From this, we calculate net payments of credit by taking the difference and scaling up to an annual measure by $4/3$. The other net worth items are reported as flows and do not require differencing across interviews.

While the CE contains fairly rich data on savings, it is designed to measure consumption and not savings. We use the savings data primarily as a consistency check, via the budget constraint, on reported consumption. As we show in section 3, the average saving rate reported in the CE appears to be broadly consistent with that obtained from the flow of funds or national income accounts, although there are marked differences. In particular, the data on new mortgages in the CE raise the question of whether the CE accurately records the net effect of refinancing on savings. The CE data implies sharp up-ticks in new mortgages around 1993 and in the 2000s, which is consistent with published statistics on refinancing. However, there are a number of reported new mortgages without a corresponding purchase of a house or a significant paying down of an existing mortgage. New mortgages for households who do not purchase a home are on average nearly 14 times the reported reductions to existing mortgages. In particular, the CE data imply an average “cash out” percentage of 73 percent from new mortgages not associated with a house purchase. This high rate is not supported by other studies of refinancing, which suggest that roughly 13 percent of the new mortgage is taken out as cash and the remainder is used to pay off existing mortgages and related costs (see [Greenspan and Kennedy, 2007](#)).

These questionable transactions, while not many in number, nevertheless affect the mean savings rate due to their size. To address this potential measurement error, we identify questionable new mortgages as those that are greater than 1.5 times the sum of the purchase price of a new house plus any lump sum payments or reductions to existing mortgages. Only 7 percent of the sample has a questionable new mortgage, but roughly three quarters of the new mortgages fall into this category. For these mortgages, we top code the amount of the mortgage as the sum of the full amount of any house purchase plus the payment on existing mortgages plus one third the reported mortgage amount. This implies that at most one third of the new mortgage amount is taken out as cash. This reduces the average implied “cash out” ratio of refinanced mortgages to 14 percent, consistent with the number reported by [Greenspan and Kennedy \(2007\)](#). In what follows, we typically present two savings series, the raw series using reported mortgages (labeled “unadjusted”) and the alternative series which uses the adjusted mortgages (labeled “adjusted”). As documented in section 3, it turns out

that the adjustment affects mean savings rates, but does not have a significant impact on the consumption inequality implied by the budget constraint.

The CE reports expenditure on hundreds of separate items. We aggregate these into 20 groups, which are listed in table 2. The division of expenditures into groups is governed by several criteria. The first is to respect BLS categorization of similar goods. The second is to define groups broadly enough to ensure consistency across the various waves of the survey. The third is to define groups narrowly enough that they span a wide range of expenditure elasticities. We adhere to the groupings created by the BLS in published statistics with a couple of exceptions. We have grouped telephone equipment and services with appliances, computers, and related services rather than with utilities, based on our prior regarding expenditure elasticities. We combine expenditure on reading materials with other nondurable entertainment expenditures because alone it represents a trivial expenditure share (about 0.2%).

For expenditure on housing services, we use rent paid for renters and self-reported rental equivalence for home owners. For the eight quarterly surveys conducted in 1980 and 1981 households were not asked about rental equivalence. We impute the rental equivalence for homeowners in these early waves based on non-housing expenditures as well as demographics. In particular, we use the the two years of surveys conducted in 1982 and 1983 and regress reported rental equivalence on total expenditures minus out of pocket housing expenditure, after tax income, and a set of dummies for age, marital status, family size, and number of earners. We then fit this regression for the earlier waves that do not report a housing service measure.

For durables other than housing we use direct expenditure, and do not impute service flows. This is motivated by our use of income groups as the unit of analysis (described below), and the assumption that aggregating over many households provides a good proxy for the consumption of durable services at a point in time. We show in section 4 that our estimates are not sensitive to the exclusion of housing and other durables.

Reported expenditures on food at home are notably lower for the 1982 to 1987 CE waves. This disparity appears to reflect different wording in the questionnaire for those years. To adjust for this drop, we increase food at home expenditure by 11% for these years. This 11% adjustment is derived from a regression for surveys 1980 to 1989 of log food at home expenditures on log after-tax income, log total expenditure, quadratic time trends, and a zero/one dummy variable that equals one for years 1982 to 1987. This adjustment is similar to that employed by [Krueger and Perri \(2006\)](#).

Income, saving, and household total expenditures are expressed in constant 1983 dollars using the CPI-U. Note that we use the aggregate CPI to deflate total expenditures, and do not deflate separately by expenditure category. This keeps all elements of the budget constraint in the same units. All results based on individual expenditure categories are expressed for one set of households relative to others (e.g., high versus low income) at a point in time, so price deflation is not an issue.

We aggregate expenditures for each household across the four interviews, so each household appears once in the sample. We assign households to years based on the month of the first interview, with households starting the survey in the fourth calendar quarter assigned to the next year.

CE survey waves from 1981 through 1983 include only urban households, and so for consistency we restrict our analysis to urban residents for the entire sample period. Our analysis employs the following further restrictions on the CE urban samples, both for the 1980-2007 and 1972-1973 samples. First, we restrict households to those with reference persons between the ages of 25 and 64. Second, we only use households who participate in all four interviews, as our income measure and most savings questions are only asked in the final interview. Third, we restrict the sample to those which the CE labels as “complete income reporters,” which corresponds to households with at least one non-zero response to any of the income and benefits questions. Fourth, we eliminate households that report unusually large expenditures on our smaller categories. In particular, we exclude any household that records spending more than half of after tax income on any category, with the exception of housing, food, and vehicle purchases. Finally, to eliminate outliers and mitigate any time-varying impact of top-coding, we exclude households in the top and bottom five percent of the before tax income distribution. (The extent of top coding dictates the five percent trimming.) We are left with 9,845 households for the 1972-73 sample, and 55,003 households for the period 1980-2007. The data appendix details how many households are eliminated at each step.

From this sample, we divide households into 5 bins based on before tax income, with the respective bins containing the 5-20, 20-40, 40-60, 60-80, and 80-95 percentile groups, respectively. For each income group and each year, we average expenditure, income, and savings variables across the member households, using the household sampling weights. Our primary measure of inequality is the ratio of the mean of the top income group to the mean of the bottom income group. When estimating the expenditure elasticities, reported in table 2 below, we control for demographics. To do this, we further divide each income group into 18 demographic cells, based on age range (25-37, 38-50, 51-64), number of earners (<2, 2+),

and household size (≤ 2 , 3-4, 5+). The analysis is done by averaging within each cell and weighting the cell by the sum of the underlying household sampling weights.

3 Budget Constraint Accounting

In this section, we review the trends in income and consumption inequality using our CE sample. We then discuss the CE savings rates, and introduce our first alternative measure of consumption based on the budget constraint.

3.1 Trends in Income and Consumption Inequality

We begin with labor earnings. The top line in figure 1 depicts the trend in labor earnings inequality. As discussed in section 2, inequality is the ratio of the mean for the top income cells to the mean for the bottom income cells. Keep in mind that the allocation of respondents into the high and low income groups is based on before tax income, and so the cells are the same for all lines in figure 1. There is substantial year-to-year movement, reflecting in large part sampling error, so we average over multiple years in table 1. In particular, we look at the three multi-year periods 1980-82, 1992-1995, and 2005-2007. For the 1980-1982 period, average household labor earnings in 1983 dollars was \$44,995 for our top income group and \$7,002 for our bottom income group, for a ratio of 6.43. Labor earnings for the top income group grew by 30 percent (in log points) through 2007, while labor earnings for the low income grew by 10 percent, resulting in a ratio of 7.88 in 2005-2007. This implies an increase in earnings inequality of 20 log points over the full period. The increase in inequality from 1980-82 to the mid 1992-1995 period is even larger at 29 percent. But this is largely driven by years 1992-1993 which, from figure 1 appear as outliers for earnings.

The next line in figure 1 is before tax income. Inequality in this broader measure of income is lower at each point in time, but also shows a steady increase over time. In particular, this ratio increases from 4.76 in 1980-82 to 6.40 in 2005-2007, for an increase of 30 percent over this period. This change is reported in the second row of table 1. Inequality in total household income, after deducting taxes, grew by slightly more than before tax income, with an increase of 33 percent over the sample period (Row 2 of table 1). As a robustness check on the CE measure of after tax income, we computed federal income taxes using the NBER's TAXSIM program, and used this in place of the CE's self-reported income tax to calculate after tax income from before tax income. This alternative measure of after tax income

inequality increased from a ratio of 3.8 in 1980-82 to a ratio of 5.01 in 2005-2007, for a log change of 28 points. This exercise suggests that respondents in the CE are under-reporting the progressivity of federal income taxes relative to TAXSIM, and this gap is increasing over time. Nevertheless, the differences do not substantially change the conclusion that income inequality increased significantly over this period, on the order of 30 percent.⁴

Figure 1 also depicts consumption inequality between the top income group and the bottom income group. The increase is much less than that of earnings or income, the feature highlighted in Krueger and Perri (2006). In table 1, we see that consumption inequality increased by only 17 percent over the full period, with 13 percent of that change occurring in the first half of the sample. The decline of CE measured consumption relative to NIPA varies across goods. If under-measurement becomes more problematic for luxuries over time, this could bias downward any measured rise in consumption inequality. However, this does not appear to be the case. In figure 1 we adjust consumption expenditures to ensure that good-by-good CE aggregates track those from the national income accounts, using proportional adjustment factors from Meyer and Sullivan (2009). This has little impact, as the two measures of consumption inequality are nearly identical. As stated in the introduction, our finding that inequality in expenditures has risen much less than income inequality is consistent with a number of papers in the literature.⁵

We have also computed inequality relative to the middle income group, which represents the 40th to 60th percentiles. For simplicity, we will refer to this as the 50th percentile. The 30 percent increase in before tax income inequality reported in table 1 can be broken down into an increase of 19 percent for the 90-50 ratio, and 11 percent for the 50-10 ratio. Similarly, the 33 percent increase in after tax income inequality is composed of a 20 percent increase for the 90-50 ratio and 14 percent increase for the 50-10 ratio. For consumption, the 17 percent increase is due to a 13 percent increase in the 90-50 ratio and a 4 percent increase in the 50-10 ratio. That is, there is very little reported increase in consumption inequality in the bottom half of the sample.

⁴The rise in income inequality we observe in the CE is broadly consistent with patterns in other data. Meyer and Sullivan (2009) measure income inequality using income information in the Current Population Surveys (CPS). There are differences in methodology from our approach; for instance, their statistics adjust for family size using equivalence scales. Nevertheless, they show for 1980-2007 an increase in the 90-10 differential in after tax income of 27 percent. Heathcote et al. (2010) also examine after tax income based on CPS data, but report a larger increase in the 90-10 differential for 1980-2005 of a little over 50 percent.

⁵ For instance, Meyer and Sullivan (2009) show an increase in the 90-10 differential in consumption expenditure for their CE sample of only about 10 percent for 1980 to 2007, or about a third of the rise they see in the 90-10 differential for after tax income over those years in the CPS. Heathcote et al. (2010) show an increase in the 90-10 differential in consumption expenditure for their CE sample for 1980 to 2006 of about 17 percent, similar to our number. But that is also only about a third as large as the rise they see in the 90-10 differential for after tax income over those years in their CPS sample.

3.2 Saving Rates

We now turn to implied and observed saving rates, beginning with mean saving rates. Figure 2 depicts the personal saving rate reported in the flow of funds accounts.⁶ There is a clear downward trend in this series, starting from 12.2 percent for 1980-82 and falling to 1.7 percent for 2005-07. This downward trend in the personal saving rate is well known, and is similar to that implied by the national income accounts.

The implied savings rate in the CE data can be computed as one minus the mean ratio of consumption expenditures to mean after tax income. This series is also depicted in figure 2. The implied saving rate has a dramatically different trend, increasing from 12.5 percent for 1980-82 to 23.2 percent for 2005-07. This sharp increase in implied savings is at odds with the flow of funds or national income accounts, and is the counterpart to the previously discussed increasing gap between CE and NIPA expenditure. Because inequality in CE expenditures rises only half as much as income inequality, the data imply a large rise in the dispersion in implied savings rates between high income and low income groups: For the top income group the implied savings rates increases dramatically from 27 percent for 1980-82 to 38 percent for 2005-07, while for the bottom group is nearly unchanged, going from -26 to -25 percent.

Although the trend in CE consumption expenditure is far below that seen in national accounts, this does not necessarily imply that the CE fails to depict accurately trends in consumption inequality. For example, if the CE's ability to capture expenditures has declined by the same multiplicative factor for all income groups, then the mis-measurement will not bias our ratio-based measures of consumption inequality. However, such a scenario has somewhat extreme implications for relative saving rates. That is, suppose we uniformly increase expenditure across groups in 2005-07 to generate a decline of 10.5 points in the aggregate CE savings rate, which is the decline observed in the flow of funds. This implies an adjustment factor of 28 percent.⁷ Given that $\frac{Savings}{Income} = 1 - \frac{Consumption}{Income}$, this implies a downward adjustment of saving rates that is 28 percent times the consumption to income ratio. Because the consumption-income ratio is much higher for low income groups, it

⁶Specifically, the saving rate is personal saving without consumer durables divided by disposable income. A similar pattern is obtained using the national income and product accounts, where savings is disposable personal income minus personal outlays.

⁷Specifically, let γ denote our adjustment factor, so we increase consumption by a factor of $(1 + \gamma)$ uniformly across households. The adjustment to the saving rate is: $\Delta \frac{S}{Y} = -\gamma \frac{C}{Y}$. To match the 10.5 point decline in the saving rate observed in the flow of funds, the aggregate CE saving must be adjusted down by 21.2 points in 2005-2007. As the ratio of aggregate CE consumption to income in 2005-2007 is 0.767, an adjustment factor of 27.6 percent is required: $(-0.276)(0.767) = -0.212$.

requires an extreme decline in their savings rate. In particular, the implied savings rate for top income group must decline modestly from 27 percent for 1980-82 to 21 percent for 2005-07, while for the bottom group is must go from -26 all the way down to -59 percent.⁸ We would suggest that such a trend decline in savings rate for the bottom group is extreme, especially given that income is defined to include transfers and given that the very lowest income households are trimmed from the sample.

Figure 2 also reports the saving rate constructed from the CE’s savings data. The series labeled “unadjusted” is the sample mean of reported savings divided by mean after tax income for each year. The mean savings rate falls from 3 percent in 1980 to -20 percent at the end of the sample. This decline is nearly a mirror image of the increase implied by consumption data, implying an inconsistency between the CE’s consumption, income, and savings data that is increasing over time. The decline, while of the right sign, is much larger than that obtained from flow of funds accounts. As mentioned in section 2, there is a measurement issue concerning new mortgages, which underlies the large decline generally, and the sharp swings around 1993 and 2003 in particular. As described in section 2, we construct an alternative savings series designed to address the mis-reporting of new mortgages. This series is the “adjusted” series in figure 2, which more closely tracks flow of funds savings and eliminates part of the sharp downward spikes in savings in the mid-1990s and 2000s.

While mean savings rates are a useful check on the data, we are primarily interested in relative consumption. Using the budget constraint, we can use reported income and savings at the household level to construct an alternative measure of consumption. Specifically, we define implied consumption as after tax income minus reported savings, which is denoted $Y - S$ in figure 1 and table 1. Starting with table 1, the unadjusted savings data implies an increase in consumption inequality of 19 percent between 1980 and 1995, and 28 percent for the full sample. The adjusted savings implies similar trends in consumption inequality, with increases of 13 and 26 percent, respectively. Both series suggest a greater increase in consumption inequality than that implied by reported expenditure. In particular, the 16 point gap between the increase in after tax income and consumption inequality (33 versus 17) is reduced to about a 5 point gap.⁹

⁸These numbers are obtained as follows. The ratio of consumption to income for the high income group averages 62 percent for 2005-2007, the adjustment lowers the saving rate of the high income group by $(-0.276)(0.618)=-0.171$. That is, the saving rate of the high income group is adjusted down from 38 percent to 21 percent. The consumption-income ratio for the low income group is 1.25, and the adjustment to the saving rate is therefore -0.34, requiring an adjustment of the saving rate from -25 percent to -59 percent.

⁹For reference, we also computed $Y - S$ using the after tax income calculated using TAXSIM’s federal income tax. The increase in implied consumption inequality is 25 percent using unadjusted savings and 20 percent using adjusted savings. The modestly smaller increase in inequality reflects that this alternative after tax income measure also displays a smaller increase in inequality.

Figure 1 depicts the ratio of high income to low income consumption implied by reported savings for each year. Given that the unadjusted and adjusted savings series yield nearly identical trends in inequality, we only plot the adjusted series to avoid clutter. The savings-implied consumption inequality tracks income inequality quite closely, and is different in both level and trend from the reported consumption data.

Moreover, recall from the previous sub-section that the increase in consumption inequality was particularly small for the 50-10 ratio. For $Y - S$, there is more balance across the top and bottom of the income distribution. In particular, the overall increase of 26 percent for the adjusted $Y - S$ is composed of a 12 percent increase in the 90-50 ratio and a 14 percent increase in the 50-10 ratio.

As previously emphasized, reported savings is not a focus of the CE, and one may reasonably question conclusions drawn solely from reported savings. Our primary focus is to use the savings data as a consistency check on the CE's consumption data. It turns out that the savings data tell a much different story regarding consumption inequality than do the expenditure data. This inconsistency raises the question of whether the expenditure data is subject to systematic measurement error that biases our estimates of consumption inequality. Addressing this potential measurement error is the focus of the next section.

4 Demand System Estimates of Consumption Inequality

In this section we present our main results. We first discuss how our econometric methodology corrects for several classes of mis-measurement. We then estimate a simple demand system which we use to generate our estimates of consumption inequality growth.

4.1 Econometric Approach

To set notation, let the index $h = 1, \dots, H$, represent cells defined by income and demographics as described in section 2. Let $i = 1, \dots, I$ denote the $I = 5$ income groups. With 18 demographic groups for each of the 5 income group, we have $H = 90$. Let $j = 1, \dots, J$ index our $J = 20$ goods; and let t index time (year). With this notation, let x_{hjt} denote reported expenditure on good j at time t by income-demographic group h , where we average over households in each cell using the CE replication weights. Let X_{ht} denote total expenditure

at time t by group h ; that is, $X_{ht} = \sum_{j=1}^J x_{hjt}$.

We assume that x_{hjt} is measured with error, with the degree of mis-measurement depending on time, income group, and good. Note that this is the systematic measurement error that survives averaging across households within each income-demographic group. In particular, let x_{hjt}^* denote the true expenditure, and

$$x_{hjt} = x_{hjt}^* e^{\psi_t^j + \phi_t^i + v_{hjt}}. \quad (1)$$

Here, ψ_t^j reflects mis-measurement of consumption good j at time t that is common across respondents (e.g., food may be under-reported for all households); ϕ_t^i represents mis-measurement specific to i at time t that is common across goods (e.g., the rich may under-report all expenditures); and v_{hjt} is good-group specific measurement error (e.g., food expenditures of the rich are under-reported). We assume that v_{hjt} is classical measurement error; in particular, it is independent of the characteristics of good j and group h at each date t . Without loss of generality (given the presence of ψ_t^j and ϕ_t^i), we normalize the mean of v_{hjt} to be zero for all t .

Our estimation consists of two steps. First, we estimate the total expenditure elasticities for each good. We assume that Engel curves are log-linear and so expenditure elasticities are constant. Of course, this can only be true locally, unless all elasticities are one. Nevertheless, it provides a tractable framework to address the mis-measurement of expenditure in the CE. We discuss mis-specification further in Section 5. We assume that true expenditure satisfies:

$$\ln x_{hjt}^* = \alpha_{jt}^* + \beta_j \ln X_{ht}^* + \mathbf{\Gamma}_j \mathbf{Z}_h + \varphi_{hjt}. \quad (2)$$

The term \mathbf{Z}_h is a vector of demographic dummies corresponding to age, number of earners per household, and family size, reflecting the categories used to construct the demographic cells. We allow the coefficient vector on demographics $\mathbf{\Gamma}_j$ to vary across goods. The error term φ_{hjt} represents idiosyncratic relative taste shocks which we assume are independent of total expenditure and independent of expenditure elasticities β_j .

We estimate expenditure elasticities using the 1972-73 Consumer Expenditure Survey. Using the early sample allows us to separate estimation of the expenditure elasticities from the estimation of the implied growth in total consumption expenditures post-1980. Specifically, let $t = 0$ denote observations from 1972 and 1973. We estimate expenditure elasticities using observed expenditures:

$$\ln x_{hj0} = \alpha_{j0} + \beta_j \ln X_{h0} + \mathbf{\Gamma}_j \mathbf{Z}_h + u_{hj0}, \quad (3)$$

where $\alpha_{jt} \equiv \alpha_{jt}^* + \psi_t^j$ subsumes the good-time specific measurement error into the intercept, and

$$u_{hj0} = \phi_0^i + v_{hj0} + \varphi_{hj0}. \quad (4)$$

We pool the years 1972 and 1973, but allow for a different good-specific intercept for each year.

A concern with estimating a demand system like (3) is that mis-measurement of individual goods is cumulated into total expenditure, inducing correlation between the measurement error captured in the residual and observed total expenditure. A standard technique is to instrument total expenditure with income and other proxies for total expenditure. However, we are already using income-category averages, which eliminates measurement error uncorrelated with income. Nevertheless, as modeled above, there may be measurement error that is common across households within an income group. This issue is mitigated by the fact that, at least using NIPA expenditure as our metric, measurement error is less of an issue in the 1972-73 survey than in the later waves. (See [Meyer and Sullivan, 2009](#).) Moreover, we can use NIPA to adjust expenditure for each good j . Specifically, we inflate or deflate expenditure on good j with a good specific multiplicative constant so that aggregate expenditure for each good j in the 1972-73 CE sample equals the corresponding NIPA expenditure for 1972-73. If this adjustment is not sufficient, our estimated elasticities may be biased. To the extent our elasticities are subject to attenuation bias, we will over-predict expenditure differences across income groups. In other words, when we invert the demand system in the second stage, we will predict expenditure inclusive of the measurement error characteristic of the 1972-73 survey.¹⁰ We revisit the possibility that our estimates are driven by first-stage measurement error in Section 5 below.

The second stage of our estimation is to invert the demand system (2) to recover an estimate of consumption growth post-1980. We first adjust expenditure for demographics and pool by income group. Specifically, let

$$\ln \hat{x}_{ijt} \equiv \frac{I}{H} \sum_{h \in i} \omega_{ht} \left(\ln x_{hjt} - \hat{\Gamma}_j \mathbf{Z}_h \right),$$

where ω_{ht} is the normalized sum of the CE sample weights for demographic group h in year t , and $\hat{\Gamma}_j$ is the estimate of Γ_j from (3). That is, $\ln \hat{x}_{ijt}$ is the average expenditure of income

¹⁰As our unit of observation is an income-demographic cell, we can instrument for expenditure with any variables that predict expenditure across the income-demographic cells, other than the demographic controls. The obvious instrument is a cell's (after-tax) income, which is equivalently to averaging over all demographic cells within an income group. We have explored this specification and found that it generates essentially the same estimates for the expenditure elasticities.

group i in year t on good j adjusted for demographics. Using (2), we have

$$\ln \hat{x}_{ijt} = \alpha_{jt} + \beta_j \ln X_{it}^* + \phi_t^i + \varepsilon_{ijt}, \quad (5)$$

where $\varepsilon_{ijt} = \varphi_{ijt} + v_{ijt} = \frac{I}{H} \sum_{h \in i} \omega_{ht} (\varphi_{hjt} + v_{hjt})$.¹¹ Given our assumptions, ε_{ijt} is independent of β_j . Therefore, we can obtain a consistent estimate of $\ln X_{it}^*$, up to a normalization, by least squares. We only have identification up to a normalization given the presence of α_{jt} .¹² Note that changes in systematic measurement error over time are captured by good-time and income group-time dummies. Identification comes from the fact that if the income of group i increases relative to that of group i' , it will increase its relative expenditure, but the increase will fall disproportionately on luxuries.

To implement (5), we regress $\ln \hat{x}_{ijt}$ on a vector of good-time dummies (whose coefficients correspond to α_{jt}), a vector of income-time dummies (whose coefficients correspond to ϕ_t^i), and the interaction of income-time dummies and β_j . The coefficients of the last group of variables will be the estimate of $\ln X_{it}^*$. To address the issue of normalization, we estimate expenditure relative to the lowest income group. That is, we have a consistent estimate of consumption inequality: $\delta_{it} \equiv \ln X_{it}^* - \ln X_{1t}^*$. To estimate trends over time, we restrict ϕ_t^i and δ_{it} to be constant within the windows 1980–1982, 1992–1995, and 2005–2007, but allow the good-time intercept terms α_{jt} to vary year by year. Our two-step procedure requires adjusting the second stage standard errors, which we do by bootstrapping.¹³

4.2 Results

Table 2 reports the results of our first stage estimates of each good’s total expenditure elasticity. The table also includes the average share of each good out of total expenditure for our 1972-73 CE sample. The first column of estimates uses the reported expenditure in the CE. The second column of estimates adjusts each good by a constant proportion to bring the aggregate CE expenditure for each good in line with its share in the national income and product accounts. We base this adjustment on conversion factors reported in Meyer and Sullivan (2009). This only affects the estimated elasticities through any effect on the distribution of total household expenditures, our right hand side variable. As can be seen from the table, this adjustment does not have a substantial impact on the estimated

¹¹The residual term will also contain estimation error related to $\hat{\Gamma}_j$, which we suppress in the notation.

¹²That is, the mean of $\ln X_{it}^*$ is not identified as $\alpha_{jt} + \beta_j \ln X_{it}^* = \alpha_{jt} - \beta_j \delta + \beta_j (\ln X_{it}^* + \delta)$.

¹³Specifically, we draw with replacement from the micro data for all years and re-estimate both stages. In a previous version, we adjusted the standard errors following Murphy and Topel (1985). Neither methodology implies a substantial adjustment to the standard errors.

elasticities. The standard errors reported next to each estimate suggest that our first stage has a fair degree of precision, particularly for the goods with large expenditure shares.

The estimated elasticities range from 0.25 for tobacco to about 2.0 for non-durable entertainment, education, and charitable contributions. Consistent with other studies, food at home has a fairly low expenditure elasticity (0.48), while food away from home has a high elasticity (1.45). Vehicle purchases is also a large category with a fairly high income elasticity. Housing services, our largest expenditure category, has an expenditure elasticity of 1.

To provide a sense of how these expenditure elasticities are informative about relative consumption inequality, we first consider two goods – food at home and non-durable entertainment. These goods have reasonably large shares and very different expenditure elasticities. We plot the relative expenditure (high income to low income) for each good in figure 3, along with the ratio for total expenditures. Food at home shows essentially no change between 1980 and 2007, implying equal growth rates for high income and low income households. Conversely, over this period high income households substantially increased entertainment expenditure relative to low income households. Specifically, the log ratio of expenditure on entertainment increased 0.75 log points.

On the one hand, the stable ratio of food at home expenditure suggests little change in consumption inequality. On the other hand, the dramatic rise in relative entertainment expenditures suggests the opposite – even with the large expenditure elasticity of 2.0, this change in entertainment expenditure still implies an increase in consumption inequality of 38 percent. However, in the terminology of the previous subsection, both these series are contaminated with systematic measurement error.

To see what we can learn from these two series, recall from equation (5) that $\ln x_{ijt} = \alpha_{jt} + \phi_t^i + \beta_j \ln X_{it}^* + \varepsilon_{ijt}$, for each good j and income group i .¹⁴ The series depicted in figure 3 are ratios of high income spending to low income spending for the same good, so they are already cleaned of the good specific intercepts. Differencing across the two goods eliminates the income specific measurement error ϕ^i . An unbiased estimate of the change in consumption inequality is therefore the change in relative entertainment expenditure minus the change in food at home, all divided by the difference in expenditure elasticities. We perform this calculation for 2005-2007 versus 1980-82. The change in consumption inequality implied by relative spending on entertainment and food is 0.48 log points, nearly three times the 0.17 log point change implied by reported expenditure, and even considerably larger than

¹⁴We are omitting the demographic shifters for this two-good exercise.

the 33 percent change in disposable income inequality.

While food at home and non-durable entertainment are informative due to their large difference in expenditure elasticities, they represent only two of our expenditure categories. Using all goods reduces the impact of the idiosyncratic error terms and provides more precise estimates. Table 3 reports our second stage regression estimates of the log change in consumption inequality from (5).¹⁵ We focus on the change in consumption inequality between the highest income and lowest income groups relative to 1980-82, and discuss other inter-group comparisons below. The first row of table 3 reports the estimated inequality in the pooled base period 1980–1982. This is the estimate of $\ln X_5^* - \ln X_1^*$ for the first three years of our sample. The row labeled “Log Change 1980–1995” is the estimated change in inequality between 1980-82 and 1992-95. Similarly, the row labelled “Log Change 1980–2007” corresponds to the estimated change in consumption inequality between 1980-82 and 2005-07.

Column (1) reports the second stage estimates using ordinary least squares. The first row reports the estimated log inequality in the pooled period 1980–1982, which is 0.95. For comparison, table 1 reports a log ratio of $\ln(2.47)=0.90$ for 1980-82, which differs from our second stage point estimate for that period by only 0.05 points. This implies that the level of consumption inequality estimated with our two-step procedure is similar to that obtained from reported expenditure for the beginning of our sample. This similarity, however, does not persist over time. The next two rows of estimates in column (1) report that the estimated change in consumption inequality is 19 percent for the early period and 32 percent for the full sample. These numbers are close to those for after tax income reported in table 1, and differ from changes in reported consumption inequality. One issue with OLS is that it weights all goods equally in the second stage. This raises the question of whether goods with small shares or greater heteroscedasticity are driving the results. Column (2) estimates the second stage using two-step feasible generalized least squares. Specifically, we allow heteroscedasticity across goods to capture that the size of taste shocks or idiosyncratic measurement error may differ across goods. To estimate good-specific residual variances, we use residuals from the OLS specification of column (1). We use these to weight the final estimation. GLS implies an initial log inequality of 0.94, and an increase in inequality in the first half of the sample of 23 percent, and a full-sample increase of 29 percent. The GLS estimates suggest a little more increase in the first half of the sample, but a similar full-sample increase as the OLS estimate.

¹⁵We only report estimates using our benchmark elasticities (column (1) of table 2). As can be surmised from the fact that the two columns of table 2 are nearly identical, the second stage estimates using column (2) elasticities are also similar to the benchmark estimates, and we therefore omit them.

Column (3) performs the same GLS regression but excludes categories that contain durables.¹⁶ Non-durable consumption avoids the issue of imputed service flow that complicates measures of durable consumption. But, because we maintain the same first stage, these estimates are still of total consumption inequality, not just non-durable consumption inequality. We find that the estimated increase in inequality is stable to this alternative sample. Specifically, we find a 23 percent increase in inequality in the first half of the sample, and 31 percent over the full sample.

The final column of table 3 implements weighted least squares, where the weights reflect the share of each good in personal consumption expenditures (PCE) from the national income accounts. Specifically, we calculate the share of each good out of total PCE for each year, then average the shares over the sample period 1980-2007 and use these shares to weight the goods in the second stage regression. For health expenditures we downweight its share for each year to a factor equal to the share of private expenditures, out of pocket and private insurance, out of total national health expenditures; this factor averages 49 percent for 1980-2007.¹⁷ The baseline log inequality is slightly lower (0.90) in this specification, and the corresponding increase over time slightly lower as well. Specifically, we estimate a change in consumption inequality of 18 percent for the early period, and 26 percent for the full sample.

The second stage estimation uses all five income categories, and therefore produces an estimate of inequality across any two income groups. For example, the OLS point estimates suggest that consumption inequality between the middle income group and the low income group increased by 0.03 points between 1980-82 and 2005-07. The GLS estimate is higher, at 0.09. These estimates are not markedly different than those obtained using the CE expenditure data directly. As discussed at the end of section 3.1, the reported increase in inequality for the 50-10 ratio is 0.04 over this period. This implies that the two-step estimation has the greatest impact at the top of the income distribution. Specifically, reported CE expenditure implies an increase in consumption inequality between the top and middle income groups of 0.13 log points between 1980-82 and 2005-2007, while our two-step estimation suggests an increase of 0.29 log points for OLS and 0.20 for GLS.

Table 4 report the estimates for income-specific measurement error, ϕ_t^i . In particular, we report the difference between the highest income group and the lowest income group: $\phi^5 - \phi^1$. These are the coefficients on the income group dummies on their own and not interacted with β_j . The rows and columns are arranged in the same manner as in table 3. The first row is the

¹⁶Specifically, from the goods listed in table 2, we exclude vehicles, appliances, furniture, and entertainment equipment.

¹⁷The data source is Centers for Medicare & Medicaid Services, Office of the Actuary, National Health Statistics Group

estimated level in 1980-82, and the remaining two rows of estimates report the changes over time. The first row point estimates suggest that (relative) income-specific mis-measurement is not substantial early in the sample. The point estimates are small and not statistically different from zero. This is consistent with the fact that log consumption inequality reported in the first row of table 3 is similar to that calculated from direct responses and reported in levels in table 1. However, the relative measurement error increases in magnitude over time, representing a substantial problem by the end of the sample. For example the OLS specification, column (1), suggests that relative mis-measurement increases by 27 log points by 2005-2007 in the direction of under-reporting of high income expenditure or over-reporting of low income expenditure. A similar pattern is reported for the other specifications. (The weighted least squares reports the smallest change, -0.14, which is consistent with the fact that it also reports a smaller implied increase in inequality in table 3.) The results imply that relative under-measurement of high income expenditure is growing over time, with an increase of 10 log points in the first half of the sample, and 27 log points for the entire sample, using the OLS estimates.¹⁸ These measurement error estimates are the flip side of the difference between table 3 and 1. The fact that our estimation suggests a greater increase in inequality over time than is suggested by reported total expenditures implies that income-group mis-measurement is growing larger over time (in the direction of under-reporting inequality).

5 Discussion and Conclusion

The results of the previous section suggest that increases in consumption inequality may mirror that of income inequality to a greater extent than implied by reported total expenditure. The Engel curve approach has a number of advantages as well as potential limitations. On the positive side, it allows us to use the detailed expenditure reports on different classes of goods to correct for income-specific measurement error. The attraction of the CE is that it is a comprehensive survey of expenditure across many goods, and this richness can be exploited using a simple demand system. Moreover, the results of the budget constraint accounting suggest that some care is warranted in interpreting the overall expenditure numbers. The demand system analysis assesses the extent to which systematic measurement error is behind

¹⁸This 27 point increase can be broken down into 6 points for the relative mis-measurement between the middle and bottom income groups, and 21 for the high-middle comparison. For GLS, the 22 point increase can be broken down into 11 points for the middle-bottom comparison, and 11 points for the high-middle comparison. This is consistent with the discussion above regarding that our two stage correction has the greatest impact on the top half of the income distribution using OLS, while GLS yields a more balanced decomposition.

the failure of the budget constraint.

A key assumption is the validity of our Engel curve estimates. We use the 1972-73 CE in order to avoid the contamination of our first stage estimates with the suspected measurement error in the later waves of the CE. This raises the question of whether the earlier CE is also subject to measurement error that may potentially bias our first-stage elasticities. One way to assess the implications of this potential bias is to assume that our second stage results are *only* due to mis-measurement in 1972-73. That is, the reported change in consumption inequality using total expenditure is accurate, and our two-stage results differ only due to mis-measured income elasticities reflecting measurement error in 1972/73.

To see what this implies, consider that our OLS estimate of the change in inequality is 1.9 times that obtained from reported expenditure (0.32 versus 0.17). We can generate this as a spurious result by assuming reported total expenditure in the 1972-73 CE overstates the log consumption gap between high income and low income households in those years by a factor of 1.9. That is, $\ln X_i = 1.9 \ln X_i^* - b$, where X_i is reported expenditure, X_i^* is true expenditure, and b is an arbitrary adjustment to the mean (to reflect, for example, that the 1972-73 CE aggregate expenditure is not very different from that implied by NIPA). Such mis-measurement would bias down our first stage elasticities, and bias up our second stage estimates of inequality by the same factor of 1.9, explaining the gap between tables 1 and 3. However, such an interpretation has fairly extreme predictions for consumption and savings behavior in 1972-73. For example, using reported expenditures, the implied savings rate in 1972-73 was 24 percent for high income households and -12 percent for low income households. Picking our b parameter to keep aggregate consumption unchanged, our hypothetical mis-measurement in 1972-73 implies that we should adjust consumption of the high income group downward by 14 percent and consumption of the low income group upwards by 37 percent. In turn, this implies increasing the savings rate for high income households from 24 to 34 percent and decreasing the savings rates of low income households from -12 all the way down to -54 percent. That is, such a re-interpretation of our results requires an extreme (and arguably implausible) savings gap in 1972-73.

Yet another alternative interpretation is that our first stage is mis-specified. As noted above, the log-linear specification is necessarily a local approximation, unless all elasticities are one. However, if we assume that expenditure elasticities revert towards one along a growth path, then we are overstating the heterogeneity in expenditure elasticities in later years. This biases down our measure of consumption inequality over time, leading to an understatement of the increase in inequality. It would take a fanning out of elasticities to generate an over-estimation of the increase in inequality. More generally, our estimation

states that high income households are shifting their expenditure towards certain goods relative to low income households. These shifts are not random, but are concentrated toward goods that have higher expenditure elasticities in our 1972-73 sample. One could construct scenarios in which these shifts are driven by instability in expenditure elasticities, but such scenarios would not resolve the budget constraint violations documented in section 3, nor directly address the fact that the micro data is increasingly at odds with aggregate measures of consumption. The income-specific measurement error approach that we have explored in this paper is a parsimonious explanation that speaks to both the observed relative shifts in consumption composition as well as the measurement issues suggested by budget constraint accounting. The results suggest that the increase in consumption inequality has been large and of a similar magnitude as the observed change in income inequality.

Data Appendix

In this appendix we describe construction of the variables in our data set and the impact of sample restrictions on our samples for 1972-73 and for 1980-2007. Unless noted, the descriptions apply to all years of the data. All data is available from the authors' web page.

Construction of variables from CE

The income variables we examine are total household labor earnings, total household income before tax, and total household income after tax. These variables are principally based on responses in the last quarterly interview that cover income from the previous 12 months. Household labor income sums all household member earnings, before deductions, over the past 12 months. The before-tax income in the CE (FINCBTAX) includes labor earnings, business (including farm), and professional income, interest, dividend, rental, and royalty income, income from social security and railroad retirement benefits, income from pensions and annuities, scholarships or stipends, workers' compensation and veterans' benefits, and alimony and child support received. It also includes the following transfer payments: public assistance (welfare) payments including those related to job training, food stamps, supplemental security income, and unemployment benefits.

We adjust this measure of before tax income in the following ways to be consistent with budget accounting. We add in food as pay and other money receipts. The latter includes lump-sum receipts of alimony and child support, lump-sum receipts from estates,

selling household items, prizes or gambling winnings, and refunds of insurance payments, property taxes, or employer over withholding on social security taxes. We subtract alimony and child support payments, to be consistent with those receipts being treated as income. We also subtract expenditures that we do not treat as consumption. These include life insurance premiums, occupational expenses, fees for financial services, finance charges, legal fees, funeral expenses, moving expenses, and support for college students. We treat the implicit rental from owner-occupied housing both as a component of expenditures and a part of income. So we add home owner's estimate of rental equivalence to before tax income. At the same time we subtract expenses of home ownership for mortgage interest, property taxes, expenditures for capital repairs and replacements, home insurance, security systems, pest control, and other maintenance expenses both from income and expenditures.

We subtract personal taxes from our measure of before tax income to arrive at a measure of after tax income. These taxes include federal, state and local income taxes. We also subtract the income contributed to social security by all household members during the year, as well as contributions for government or railroad retirement programs. The CE measure of social security contributions is estimated by the BLS. Contributions for government and railroad retirement programs is annualized by the BLS based on deductions for household members in the most recent pay period. Our measure of after tax income differs from the CE measure (FINCATAX) due to all the adjustments listed above to before tax income, and because we subtract contributions to social security, government, and railroad retirement programs. We consider an alternative measure of after tax income by replacing self-reported federal income taxes with taxes calculated from the NBER's TAXSIM program. We do not adjust for state and local taxes, as we do not know the state of residence for many households in the CE. We also considered replacing social security taxes with TAXSIM values, but this has little effect on the results. This is not surprising, as the social security contributions in the CE are estimated by the BLS as well.

We aggregate CE expenditure items into 20 groups, as described in the text. Our definitions of expenditures by good closely follow definitions in the CE with a few exceptions, most notably for housing services. As in the CE, for renters we define housing by rent paid. But for home owners we use self-reported rental equivalence rather than out of pocket expenditures. This adjustment was described above in discussing adjustment to income. For the eight quarterly surveys conducted in 1980 and 1981 households were not asked about rental equivalence. We impute the rental equivalence for homeowners in these early waves as follows. We use the the two years of surveys conducted in 1982 and 1983 and regress reported rental equivalence on total expenditures minus out of pocket housing expenditure,

after tax income, and a set of dummies for age, marital status, family size, and number of earners. We then fit this regression for the earlier waves that do not report a housing service measure. For vacation homes there is no measure of self-reported rental equivalence before 1999. So, for all years, we measure expenditures on vacation homes, like the CE, based on expenditures for mortgage interest, taxes, and maintenance.

We differ from the CE measure of expenditures on vehicles in that we subtract the value of used vehicles that are sold by a household, even when this is separate from any vehicle purchase. (Both our measure and the CE expenditure on vehicles, by using net payments for vehicle purchases, implicitly deducts the value of vehicles traded in as part of purchases.) We also adjust the reported expenditures on food at home in the CE for the 1982 to 1987 waves. Spending on food at home shows a distinct drop for these waves, apparently reflecting a difference in the questionnaire wording from other waves. To adjust for this drop, we increase food at home expenditure by 11% for these waves. This 11% adjustment is derived from a regression for surveys 1980 to 1989 of log food at home expenditures on log after-tax income, log total expenditure, quadratic time trends, and a zero/one dummy variable that equals one for the waves from 1982 to 1987. This adjustment is similar to that in [Krueger and Perri \(2006\)](#).

Our measure of total expenditure will differ from the BLS measure of total expenditure in the CE (TOTEXP) due to these adjustments. It also differs because we treat a set of expenditures (e.g., alimony payments, life insurance, financial fees, social security contributions) as deductions from income, rather than as consumption expenditures. We also treat payments to private pensions as a component of savings, whereas the CE includes these as part of total expenditure.

The CE asks respondents a number of questions on active savings. The BLS employs these responses to publish statistics on net changes in assets and liabilities (see addenda to Current Expenditure Tables, www.bls.gov/cex/home.htm#tables). In each quarterly interview, households report the net change in savings accounts and purchases and sales of stocks and other financial assets. In addition, households report new loans undertaken, including mortgages and home equity loans, and reports equity payments against mortgages and other loans. Households also report purchases and sales of real assets including houses, businesses, home improvements, and vehicles. They report the net changes in money borrowed or loaned to other households. The CE records the total outstanding credit balances in the first and fourth interviews covering expenditures, which are 9 months apart. We estimate net payments of credit by subtracting the fourth interview's value from that in the first, and annualize by multiplying by 4/3. (Because all other responses for savings already reflect

changes in assets or changes in liabilities, these do not require differencing across interviews.)

Our measure of net changes in assets and liabilities differs in a couple respects from the CE measures reflected in BLS published statistics. The primary difference is that we add payments into private pensions as a form of savings (not as a component of expenditures). Secondly, we do not include net purchases of vehicles, as we treat these as a component of expenditures.

As discussed in the text, the data on new mortgages in the CE raise the question of whether the CE accurately records the net effect of refinancing on savings. We observe a number of reported new mortgages without a corresponding purchase of a house or a significant paying down of an existing mortgage. The CE data imply an average “cash out” percentage of 73 percent from new mortgages not associated with a house purchase, a rate not supported by studies of refinancing. (For instance, [Greenspan and Kennedy, 2007](#), finds that 13 percent of the value of new mortgages is taken in the form of cash, not used to pay off existing mortgages or to pay related fees.) To address this potential measurement error, we construct an alternative measure of household savings that caps the amount of net borrowing (cash out) associated with new mortgages at one third the size of that mortgage. This reduces the average implied cash out ratio of refinanced mortgages to 14 percent, close to the number reported by [Greenspan and Kennedy \(2007\)](#).

Lastly, we create demographic variables for age of the reference person (identified by who owns or rents the residence), the number of household members, and number of household earners, with all variables based on responses in the households final quarterly interval. (For the 1972-73 surveys these variables are defined by the first household interview; so we add one year to the reference person’s age.) These variables are used to divide households within each of five income groups into cells, as described in the text.

The impact of sample restrictions

We impose a set of sample restrictions; the impact of these restrictions is reported Table [A 1](#). We begin with 19,975 households for 1972-73 surveys and about 225,443 households for 1980-2007 surveys. We aggregate expenditures for each household across the four interviews—so each household appears only once in the sample. There is considerable attrition across surveys, especially for the latter surveys. For the 1980-2007 surveys, the BLS responds to attrition by introducing households with the second, or later, survey instrument, so as to keep a balanced panel across interview quarters. Focusing on households that begin with the

first survey instrument reduces the potential sample of households to 165,638 for 1980-2007.

We make the following restrictions on these samples. The 1981 through 1983 surveys include only urban households. For consistency we restrict the samples to urban households for their entirety. This reduces the 1972-73 sample by 16 percent to 16,809 and the 1980-2007 sample by 9 percent to 150,465. We restrict households to those with reference persons between the ages of 25 and 64, reducing the sample for 1972-73 by 27 percent to 12,220 and for 1980-2007 by 28 percent to 108,156. In order to contrast household expenditures with income, it is necessary to have measures of expenditures and income over comparable periods. In turn this requires that households participate in all four interviews in order to be present for the income variables in the final interview. This reduces the 1972-73 sample by 3 percent, to 11,898, and the 1980-2007 sample by 31 percent, to 74,422. We require households to be “complete income reporters,” which the BLS defines as respondents with values for some major source of income, such as wages, self-employment income, or Social Security income. (Even complete income reporters might not have provided full accounting for all household members.) This restriction reduces the 1972-73 sample by 6 percent to 11,182 and the 1980-2007 sample by 14 percent to 63,794. We drop households that report implausibly large spending on smaller goods categories. More exactly, we require that households spend less than half of their after tax income on any category, unless it is housing, food, or vehicle purchases. This restriction reduces the sample for 1972-73 by 2 percent to 10,978 and the 1980-2007 sample by 4 percent to 61,064. (Of those eliminated, 28 households showed negative after tax income for the 1972-73 sample, 857 households did so for 1980-2007.) Lastly, in order to eliminate outliers and to mitigate the impact of time-varying top-coding, we exclude households in the top and bottom five percent of the before tax income distribution. (The fraction of households top coded on income fluctuates from about one to just over four percent across survey waves.) This results in a sample of 9,845 households for 1972-73 and 55,003 households for 1980-2007.

References

Aguiar, Mark and Erik Hurst, “Deconstructing Life Cycle Expenditure,” *working paper*, 2009.

Attanasio, Orazio and Steven Davis, “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 1996, *104*, 1227–1262.

—, **Erich Battistin, and Hidehiko Ichimura**, *What Really Happened to Consumption Inequality in the US?*, Chicago: Chicago University Press, 2005.

- Blundell, Richard and Ian Preston**, “Consumption Inequality and Income Uncertainty,” *Quarterly Journal of Economics*, 1998, *113*, 603–640.
- , **Luigi Pistaferri, and Ian Preston**, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 2008, *98*, 1887–1921.
- Bosworth, Barry and Rosanna Smart**, “Evaluating Micro-Survey Estimates of Wealth and Saving,” *CRR WP 2009-4*, 2009.
- and **Sarah Anders**, “Saving and Wealth Accumulation in the PSID, 1984-2005,” *CR-RWP 2008-2*, 2008.
- Cutler, David and Lawrence Katz**, “Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s,” *American Economic Review*, 1992, *82*, 546–551.
- Deaton, Angus and Christina Paxson**, “Intertemporal Choice and Inequality,” *Journal of Political Economy*, 1994, *102*, 437–467.
- Garner, Thesai, George Janini, William Passero, Laura Paszkiewicz, and Mark Vendemia**, “The Consumer Expenditure Survey: A Comparison with Personal Expenditures,” *Bureau of Labor Statistics*, 2006.
- Greenspan, Alan and James Kennedy**, “Sources and Uses of Equity Extracted from Homes,” *Federal Reserve Board Finance and Economics Discussion Series 2007-20*, 2007.
- Guvenen, Fatih**, “Learning your Earning: Are Income Shocks Really Very Persistent,” *American Economic Review*, 2007, *97*, 687–712.
- Heathcote, Jonathan, Fabrizio Perri, and Gianluca Violante**, “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006,” *Review of Economic Dynamics*, 2010, *13*, 15–51.
- , **Kjetil Storesletten, and Giovanni Violante**, “Two Views of Inequality over the Life Cycle,” *Journal of the European Economic Association*, 2005, *3*, 765–775.
- Huggett, Mark, Gustavo Ventura, and Amir Yaron**, “Sources of Lifetime Inequality,” *working paper*, 2009.
- Krueger, Dirk and Fabrizio Perri**, “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *The Review of Economic Studies*, 2006, *73*, 163–193.
- , – , **Luigi Pistaferri, and Gianluca Violante**, “Cross-Sectional Facts for Macroeconomists,” *Review of Economic Dynamics*, 2010, *13*.
- Meyer, Bruce and James Sullivan**, “Five Decades of Consumption and Income Poverty,” *NBER Working Paper 14827*, 2009.
- Murphy, Kevin and Robert Topel**, “Estimation and Inference in Two-Step Econometric Models,” *Journal of Business and Economic Statistics*, 1985, *3*, 370–379.

Parker, Jonathan, Annette Vissing-Jorgensen, and Nicholas Ziebarth, “Inequality in Expenditure in the Twentieth Century,” *working paper*, 2009.

Sabelhaus, John and Jeffrey Groen, “Can Permanent-Income Theory Explain Cross-Sectional Consumption Patterns,” *The Review of Economics and Statistics*, 2000, *82*, 431–438.

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

Storesletten, Kjetil, Chris Telmer, and Ameir Yaron, “Consumption and Risk Sharing over the Life Cycle,” *Journal of Monetary Economics*, 2004, *51*, 609–633.

Table 1: Trends in Inequality – Ratio of High Income to Low Income Respondents

	1980–82	1992–1995	2005–2007	Log Change 1980–1995	Log Change 1980–2007
Labor Earnings	6.43	8.57	7.88	0.29	0.20
Before Tax Income	4.76	5.87	6.40	0.21	0.30
After Tax Income	4.21	5.13	5.87	0.20	0.33
Consumption Expenditures	2.47	2.80	2.93	0.13	0.17
Income minus Savings	4.20	5.09	5.54	0.19	0.28
Income minus Savings (adjusted)	4.22	4.83	5.45	0.13	0.26

Note: High income refers to respondents who report before tax household income in the 80th through 95th percentiles. Low income refers to respondents in the 5th through 25th percentiles. The elements of the first three columns are the ratio of the average of high income respondents to the average for low income respondents, where the averages are taken over the pooled years indicated at the head of the respective column. The last two columns are the log change in the first two columns and the first and third columns, respectively. All variables are converted into constant dollars before averaging. The row labeled Income minus Savings is reported after tax income minus reported savings. The final row, labeled Income minus Saving (adjusted) limits the amount of cash taken out of refinanced mortgages, as described in the text. Definitions of each series and sample construction are given in the data section.

Table 2: Engel Curves from 1972/1973 Expenditure Survey

Good Category	1972/1973		
	CE Share	(1)	(2)
Housing	20.8	0.99 (0.03)	0.95 (0.03)
Food at home	15.6	0.48 (0.03)	0.45 (0.03)
Vehicle purchasing, leasing, insurance	10.7	1.48 (0.07)	1.40 (0.07)
All other transportation	9.8	0.97 (0.03)	0.93 (0.03)
Food away from home	5.2	1.45 (0.05)	1.41 (0.04)
Health expenditures including insurance	5.2	0.99 (0.06)	0.94 (0.06)
Utilities	4.5	0.68 (0.05)	0.63 (0.05)
Applicances, phones, computers, with services	4.4	0.70 (0.03)	0.68 (0.03)
Men's and women's clothing	4.3	1.42 (0.04)	1.38 (0.04)
Furniture and fixtures	2.9	1.74 (0.07)	1.67 (0.07)
Shoes and other apparel	2.8	0.89 (0.04)	0.88 (0.04)
Entertainment equipment and subscription television	2.8	1.19 (0.06)	1.15 (0.06)
Entertainment fees, admissions, reading	2.1	2.00 (0.06)	1.94 (0.06)
Domestic services and childcare	1.8	1.39 (0.16)	1.37 (0.15)
Tobacco, other smoking	1.6	0.25 (0.05)	0.24 (0.05)
Children's clothing (up to age 15)	1.4	0.36 (0.09)	0.34 (0.09)
Education	1.3	2.10 (0.24)	2.02 (0.23)
Personal care	1.2	1.08 (0.06)	1.04 (0.06)
Alcoholic beverages	1.0	1.49 (0.09)	1.49 (0.08)
Cash contributions, not for alimony/support	0.5	2.07 (0.15)	2.01 (0.14)
Total expenditure adjusted based on NIPA?	No	No	Yes

Note: The first column presents each goods average share of total expenditure for 1972-73, calculated as the cross-sectional average expenditure on each good for each year divided by average total expenditure for that year and then averaged across years. The remaining columns report two alternative estimates of each goods expenditure elasticity, with associated standard errors in parentheses. Column (1) uses the reported expenditure data, while column (2) adjusts expenditure on each good by a constant multiple to ensure aggregate expenditure equals NIPA expenditure. The data come from the 1972 and 1973 CE surveys. See text for details of sample construction and regression specification. All specifications include demographic control dummies for age, household size, and number of earners.

Table 3: Trends in Consumption Inequality Based on Relative Expenditure Patterns

	(1)	(2)	(3)	(4)
Log Inequality 1980–1982	0.95 (0.09)	0.94 (0.08)	0.97 (0.08)	0.90 (0.07)
Log Change 1980–1995	0.19 (0.09)	0.23 (0.08)	0.23 (0.08)	0.18 (0.07)
Log Change 1980–2007	0.32 (0.09)	0.29 (0.08)	0.31 (0.08)	0.26 (0.07)
Categories Included	All	All	Those without durables	All
Specification	OLS	GLS	GLS	WLS (NIPA Shares)

Note: This table reports the estimated change in income inequality obtained from the second stage regressions. The estimated parameters represent the relative growth in total expenditure for high income households relative to low income households. See the specification in the text for full details. The first column implements the second stage by OLS; the second column implements the second stage using two-step generalized least squares, where we allow expenditure on each good to have its own residual variance; the third column implements GLS while omitting all good categories containing durables; the last column implements weighted least squares, using the average NIPA for 1980–2007 shares as weights. Standard errors are calculated using a bootstrap with 500 replications.

Table 4: Change in Relative Income-Specific Measurement Error

	(1)	(2)	(3)	(4)
Relative Mis-Measurment 1980–1982				
High Income – Low Income	0.01 (0.08)	0.01 (0.08)	-0.04 (0.08)	0.04 (0.06)
Change 1980–1995	-0.10 (0.09)	-0.14 (0.08)	-0.13 (0.08)	-0.07 (0.07)
Change 1980–2007	-0.27 (0.09)	-0.22 (0.08)	-0.22 (0.08)	-0.14 (0.07)
Categories Included	All	All	Those without durables	All
Specification	OLS	GLS	GLS	WLS (NIPA Shares)


Note: This table reports the change in the estimated income-specific measurement error for high income respondents relative to low income respondents: $\phi^5 - \phi^1$ from equation (5). The specification for each column is the same as in table 3. The first row is the level for the period 1980-82, and the next two rows report the change in the levels between 1980-82 and 1992-95 and 1980-82 and 2005-07, respectively. Standard errors are calculated using a bootstrap with 500 replications.

Table A 1: Sample Construction

	1972-73	1980-2007
Total Number of Households	19,975	225,443
Households who enter at “first” interview”	NA	165,683
After Sample Restriction:		
Urban	16,809	150,465
Ages 25 to 64	12,220	108,156
Full-year of Interview Coverage	11,898	74,422
Complete Income Reporter	11,182	63,794
No Expenditure Outliers	10,978	61,064
Truncate Before-Tax Income: 5-95 ptile (Final Sample)	9,845	55,003


Note: This table reports the sample size after each restriction. The first row reports the original CE sample obtained from the BLS. Each sample restriction is discussed in the data appendix. The final row represents the sample used in the analysis.

Figure 1: Trends in Inequality




Note: This figure depicts the ratio of high income to low income respondents reported labor earnings, before tax income, after tax income, income minus savings, NIPA adjusted consumption expenditures, and consumption expenditures. High income refers to respondents who report before tax household income in the 80th through 95th percentiles. Low income refers to respondents in the 5th through 25th percentiles. The income minus saving (Y-S) sample uses the adjusted mortgage series as described in the text. Definitions of each series and sample construction are given in the data section.

Figure 2: Mean Saving Rates



Note: This figure depicts the mean savings rates. The line labeled $1 - C/Y$ refers to implied savings computed as after tax income minus reported consumption expenditures. The line labeled “Flow of Funds” is the flow of funds aggregate private savings rate out of disposable income. The lines labeled S/Y refer to average reported savings divided by average reported after tax income. Adjusted and unadjusted refer to whether we adjust reported new mortgages, as described in the data section of the text. Definitions of each series and sample construction are given in the data section of the text.

Figure 3: Food and Entertainment Expenditure – Ratio of High Income Households to Low Income Households



Note: This figure depicts the relative spending of high and low income households on nondurable entertainment and food at home, as well as total expenditures. Points are averaged over 3 year spans centered on the corresponding horizontal axis label. The year 1993 corresponds to the four year average 1992–1995.