Uncovering Recent Trends in Consumption and Wealth Inequality

Robin Brace

August 2019

An essay submitted to the Department of Economics
in partial fulfillment of the requirements for
the degree of Masters of Arts

Queen's University
Kingston, Ontario, Canada

August 2019

Acknowledgments

I would like to thank my supervisor, Dr. Brant Abbott, for his continued advice and feedback. He patiently answered all my emails and questions and offered exceptional guidance throughout this process. Thank you for your kindness and support. I would also like to thank my family, friends, and my partner. I could not have succeeded without their love and encouragement. All errors are my own.

Contents

1 Introduction

Much of the existing research on inequality focuses on rising wealth and income inequality. While they are easier to measure, and there is less controversy surrounding their recent trends, income and wealth are decidedly less indicative of well-being than consumption. Utility functions evaluate consumption and leisure as opposed to income or wealth and consumption can provide a more accurate depiction of an individual's standard of living. However, there is little consensus about basic consumption inequality patterns over the last several decades or about how individuals consume in relation to their wealth. In order to evaluate the welfare effects of inequality, more work is needed to understand consumption inequality dynamics.

The debate over trends in consumption inequality stems from inconsistencies in its measurement. Historically, researchers have used the Consumer Expenditure Survey (CE) to measure individual consumption in the United States. Initial research indicated that consumption inequality was not following the same trends as income and wealth inequality, but had instead remained relatively flat since the 1980s (Krueger and Perri 2006; Heathcote, Perri, and Violante 2010; Meyer and Sullivan 2017). Conversely, Garner et al. (2006) and Parker, Vissing-Jorgensen, and Ziebarth (2009) documented that the consumption data published by the CE are inconsistent with equivalent spending data in the NIPA tables. This evidence indicates that consumption microdata is measured with error and thus creates significant difficulties for computing consumption inequality.

A growing discussion addresses the challenges in quantifying consumption inequality due to measurement error. Papers use various imputation methods to compute total consumption either because of incomplete data or to address measurement error. Blundell, Pistaferri, and Preston (2008) maintain that the measurement error is classical in nature and account for it by imputing total consumption in the CE. More recently, Aguiar and Bils (2015) document that the measurement error is systematic in nature and develop a new estimation method that accounts for non-classical error. The method presented by Aguiar and Bils (2015) is novel because it calculates total consumption while allowing for systematic measurement error. Consumption inequality for Aguiar and Bils (2015) is estimated by evaluating how rich vs. poor households allocate spending across goods. Even if consumption is mis-measured at the household level, the ratio of consumption between goods will be consistent for a given household. The estimation follows a two-step process. First, consumption elasticities are estimated from an Engel curve demand system. The second stage inverts the demand system and recovers an estimate of total consumption. This method controls for household and good-level systematic measurement error. Aguiar and Bils (2015) use their approach to conclude that consumption inequality has indeed been increasing since the 1980s.

While these methods have been used to compare trends in consumption inequality to trends in income inequality, little is known about how consumption and wealth inequality interact. This paper evaluates how consumption inequality has evolved with wealth inequality from 2004 to 2017. I apply the method developed by Aguiar and Bils (2015) to produce an estimate of consumption inequality in the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF).

The SCF is the primary source for wealth and income data in the United States. Unlike other surveys, it accurately represents the top end of the distributions, where most of the wealth and income are held. Because the SCF only collects food consumption data, total consumption must be imputed in order to compare wealth and consumption inequality jointly. Only Fisher et al. (2018) compare consumption and wealth inequality directly in the SCF, but their imputation method does not account for systematic measurement error. Until recently, the CE was the dominant source for consumption data in the United States. However, the PSID expanded their consumption data in 2005 so that it captures all consumption categories at aggregate levels; it has since become an increasingly popular resource for consumption analysis. This paper uses expenditure elasticities calculated in the PSID to calculate total consumption and consumption inequality in the SCF by exploiting the relative differences in the demand for food at home, away, and delivered in the SCF. Because these goods are relative necessities and luxuries, there is enough variation in the elasticities to identify the demand system. This allows me to compare consumption inequality to the most compelling measure of wealth inequality available in the United States.

This paper contributes to the debate over consumption inequality as it relates to wealth inequality by providing a rigorous measure of consumption in the SCF. It is the only paper to jointly present trends in consumption and wealth inequality that are robust to systematic measurement error. I correct for systematic mis-measurement in consumption data by applying the method presented by Aguiar and Bils (2015) to the data in the SCF. The remainder of the paper is organized as follows: section II explores the related literature and compares previous imputation methods; section III describes the data used; section IV details the methodology presented by Aguiar and Bils (2015) and its extension in this paper; section V presents the main results and Section VI concludes.

2 Related Literature

A large portion of the related literature considers consumption inequality in contrast to income inequality. This research measures whether individuals are able to insure against income shocks and smooth consumption, and points to the permanence or temporality of income inequality. Krueger and Perri (2006), Heathcote, Perri, and Violante (2010) and Meyer and Sullivan (2017) find that consumption inequality has not kept pace with income inequality and conclude that this is due to the transitory nature of income shocks. This evidence was used to argue that the implications of rising income inequality have not induced large welfare effects because individuals are able to insure against transitory income shocks (Attanasio and Pistaferri 2014). Meanwhile, Blundell, Pistaferri, and Preston (2008) argue instead that this trend occurs because individuals do not react enough to permanent income shocks. More recently, Aguiar and Bils (2015), Attanasio, Hurst, and Pistaferri (2012), Attanasio and Pistaferri (2014) and Fisher et al. (2018) have found that consumption inequality has indeed mirrored trends in income inequality and attribute the previous contradictory findings to a lack of quality data.

While the effects of income inequality on consumption inequality inform us about the permanent income hypothesis and the life-cycle hypothesis, consumption inequality is uniquely important when researching the welfare effects of inequality. Income measures can be quite volatile from one year to the next, even when people are living under very similar circumstances. By contrast, individual consumption is generally more stable and, thus, may better reflect an individual's living standards. Accurate consumption measurement is essential in evaluating the level of deprivation at the bottom end of the income distribution. Meyer and Sullivan (2009) claim that poverty measures related to consumption provide a more holistic picture of destitution and economic health. Attanasio and Pistaferri (2016) argue that since economists construct utility curves that evaluate consumption and leisure, we should evaluate inequality with respect to these measures as opposed to with income and the budget constraint. Studying wealth inequality in conjunction with consumption inequality is important because a person's stock of wealth gives them the power to self-insure and transfer wealth across generations; wealth is what links income and consumption inequality. Fisher et al. (2018) prioritize a multidimensional analysis of inequality and evaluate trends in the joint distributions of income, consumption and wealth. They point out that an increase in income inequality need only result in an increase in either consumption or wealth inequality, but not necessarily both.[1] This emphasizes the importance of contrasting the trends of wealth and consumption inequality.

Wolff (2014) studies wealth inequality from 1983 to 2010 and finds that wealth inequality in the United States was relatively stable until 2007 as median wealth grew substantially in the 1990s and early 2000s. However, wealth inequality rose steeply during the Great Recession as median wealth fell by 47 percent. Saez and Zucman (2016) find that the share of wealth held by the top 1% has been growing strongly since the 1970s. Their estimates show that the top 0.1% share of wealth grew from 7% in 1978 to 22% in 2012, which highlights the importance of accurate measurement at the top end of the distribution. Wealth inequality is compounded by high savings rates and rising capital income at the top end of the income distribution (Saez and Zucman 2016). The recent plunge in median wealth indicated by Wolff (2014) and Saez and Zucman (2016) can be explained by slow growth in pensions coupled with high levels of consumer credit, mortgage and student debt. Kopczuk (2015) notes that the nature of wealth accumulation has changed over the last few decades; wealth accumulation now relies more on income related to labor rather than capital income. Relating wealth accumulation to consumption, Straub (2018) challenges the traditional view that consumption is a linear function of permanent income. He estimates a permanent income elasticity of consumption of 0.7, which indicates that the rich save a larger portion of their income than the poor. He incorporates non-homothetic preferences into a precautionary savings model and, when calibrated, the model matches the high levels of wealth inequality observed in 2014.

A branch of research that informs this paper more closely is that which examines the quality of consumption data available in the United States. The most extensive

1. Fisher et al. (2018) find that joint measures of inequality exceed any univariate measure and express concern that analysis relying on any single measure of inequality may understate recent patterns in inequality. Heathcote, Perri, and Violante (2010) also directly contrast income, wealth and consumption inequality.

source of consumption data is available through the Consumer Expenditure (CE) Survey. This survey reports individual expenditure on hundreds of different items and is available annually since 1980. However, Garner et al. (2006), Parker, Vissing-Jorgensen, and Ziebarth (2009), and others, document a gap between the CE data and consumption reported in the NIPA tables. This gap indicates that measurement error exists in the CE data. Meanwhile, Heathcote, Perri, and Violante (2010) report that this gap is expanding over time. The discrepancy between the CE data and the NIPA tables are the likely culprits for why it appears that consumption inequality has not followed the trends in income inequality. A number of papers discussed below attempt to address this bias through various imputation methods.

Prior to 1999, the PSID only reported expenditure on food and housing. The survey expanded its consumption questions in 1999 and 2005 and now the PSID reports consumption measures that closely match equivalent aggregates in the CE (Blundell, Pistaferri, and Preston 2008; Li et al. 2010; Heathcote, Perri, and Violante 2010). Before the PSID expansion, no survey contained comprehensive data on both income and expenditures, so imputation methods were also used to impute consumption across surveys. Previous work comparing consumption data in the CE to that in the PSID after the 1999 expansion indicate that both likely contain systematic measurement error. There exist multiple gaps between measures in the two surveys: expenditure on education in the PSID only accounts for 86% of that in the CE, expenditure on food is 8% higher in the PSID while spending on transportation is documented to be 5% lower in the PSID than in the CE (Andreski et al. 2014; Li et al. 2010). A small part of these discrepancies can be accounted for through definitional differences. After correcting for measurement error, Aguiar and Bils (2015), Attanasio and Pistaferri (2014) and Fisher et al. (2018) find that consumption inequality has followed the same increasing trend as income inequality. These comparisons illustrate the importance of taking measurement error into account when analyzing consumption data.

Imputing consumption was common practice before the PSID contained detailed consumption data. Skinner (1987) imputes total consumption in the PSID by regressing total consumption in the CE on the consumption items present in both the PSID and the CE. The coefficients from the CE regression are used to calculate a linear prediction of total consumption in the PSID. Since the issue of measurement error in CE data has been highlighted, Battistin (2003) and Orazio, Battistin, and Ichimura (2007) have used the diary components of the CE to correct for measurement error. Blundell, Pistaferri, and Preston (2008) impute total consumption in the PSID using the CE by estimating the demand for food in the CE, conditional on prices, non-durable expenditure, and various demographic variables. This function is then inverted to obtain a measure of non-durable consumption in the PSID. This method diverts from Skinner (1987) because it allows demand to vary with prices, non-durable expenditure and demographics (Blundell, Pistaferri, and Preston 2008). They assume that consumption is measured with classical measurement error and argue that a valid instrument for total expenditure will produce unbiased estimates of total expenditure in the PSID. Attanasio and Pistaferri (2014) also use the method presented by Blundell, Pistaferri, and Preston (2008) to impute total consumption in the PSID. They use the expanded categories presented in 1999 to impute total consumption backwards before all consumption categories were presented. Since the PSID consumption data in 1999 appear to match the NIPA tables, they argue that this method should present an accurate depiction of consumption inequality throughout the history of the PSID. They use in-sample verification to test that their imputation method replicates the trends observed in the more recent PSID data.

Browning and Crossley (2009) target measurement error in consumption data through an Engel curve approach. They argue that using two noisy measures can produce an accurate measure of total consumption and can dominate one expensive and relatively accurate measure. They suggest employing two categories of spending (ideally one luxury and one necessity) so that the covariance of the two spending variables will reveal the variance of total expenditure and present an estimator that is centered on the true value. Aguiar and Bils (2015) identify that the under-reporting of consumption data in the CE is not uniform across income groups and is thus subject to non-classical measurement error. To account for this, they employ an Engel curve approach, as in Browning and Crossley (2009), but exploit differences across goods within a demand system rather than extracting a common source of variation. Aguiar and Bils (2015) estimate consumption inequality by comparing how households allocate spending on relative luxuries vs. relative necessities. Since this method relies on the ratio of spending across goods to identify total consumption, it is robust to household-level measurement error. The method presented by Aguiar and Bils (2015) acts as the foundation for this paper and I will therefore go into more detail in the methodology section of this paper. Aguiar and Bils (2015) are able to reproduce results consistent with those presented by Blundell, Pistaferri, and Preston (2008) when using raw CE data, but find that their corrected measure of income inequality shows that individuals are less able to insure against income shocks than Blundell, Pistaferri, and Preston (2008) originally claim.

Fisher et al. (2018) is the only research to impute consumption and compare it to wealth inequality in the Survey of Consumer Finances (SCF). They use the CE to impute consumption for items not reported in the SCF. Their imputation method involves calculating the ratio of reported consumption to total consumption in the CE for the variables included in the SCF. They use the coefficients from the CE to predict the share of total consumption reported in the SCF. They arrive at the total imputed level of consumption in the SCF by dividing the reported consumption in the SCF by the imputed share. This method relies on the shares of consumption being unbiased in the CE data and the authors do not account for any measurement error in the CE data. The SCF is a rich source of wealth data because it is uniquely able to capture behaviour at the top ends of the income and wealth distributions. Both the CE and the PSID report low survey response rates from high-income and high-wealth individuals so the data presented are missing representation from the top of the income and wealth distributions; the full distributions are not represented in either survey (Sabelhaus and Groen 2000). Using the SCF data is important because it provides a more complete view of inequality.

As outlined above, many papers study consumption inequality as it relates to income inequality to evaluate the extent to which individuals can insure against income shocks, both permanent and temporary. Few compare consumption inequality with wealth inequality and none do so while taking into account the well-documented non-classical measurement error in consumption data. This paper addresses that issue by applying the estimation method developed by Aguiar and Bils (2015) to the Survey of Consumer Finances. This paper fills a gap in existing literature by accounting for non-classical measurement error in consumption data and imputing unbiased levels of total consumption for respondents in the SCF. In doing so, this paper presents a more complete view of consumption and wealth inequality throughout the entire distribution.

3 Data

The data for this paper comes from two main sources, the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF). The PSID is a longitudinal survey that collects detailed information on income, transfers, housing and consumption in the United States. The SCF is a cross-sectional survey that specializes in collecting information on income and wealth, particularly at the top end of the distributions. In order to use a consistent measure of consumption, data are used from 2005-2017 in the PSID and from 2004-2016 in the SCF.

The PSID began in 1968 and sampled over 18,000 individuals in 5,000 different families in the United States. These families, and their offspring, are surveyed every two years so that now over 11,000 families and multiple generations are represented. Originally created to study income and poverty dynamics, the PSID collects information on a range of topics including employment, income, wealth, expenditures, health , education, family structure etc. The PSID contains a representative sample, a supplementary low-income subsample, and a Latino subsample, which was added in 1990. Only the representative sample and their descendants are included in this study. The unit of observation in the PSID is a family unit which represents a group of economically-dependent people living together.

The SCF is largely considered the gold standard of wealth data in the United States (Pfeffer et al. 2016). The survey releases data every three years on over 6,000 families. The goal of the survey is to learn about American families' finances and thus, in order to produce accurate data on wealth and income, the questionnaire includes many detailed questions about different kinds of assets, incomes, and liabilities. Because such a large share of wealth is held at the very top of the wealth distribution, the SCF oversamples wealthy households in the hopes of better understanding how wealth is distributed. Weights are used to make calculations in reference to a representative sample. These qualities ensure that the SCF produces the most reliable wealth data for American families. The unit of observation in the SCF is the primary economic unit, which is defined as an economically dominant individual and all people who are dependent on them.

Wealth in the PSID is defined as the sum of farm and business assets, the value in a family's checking and savings accounts, their stocks, real estate holdings, vehicles, annuities and other assets net of debt value plus the value of home equity. The SCF assesses wealth through a large variety of questions and eventually publishes a family's net worth as the difference between their assets and debts.

Figure 1: Share of Wealth Owned by Top 10% in the PSID and the SCF

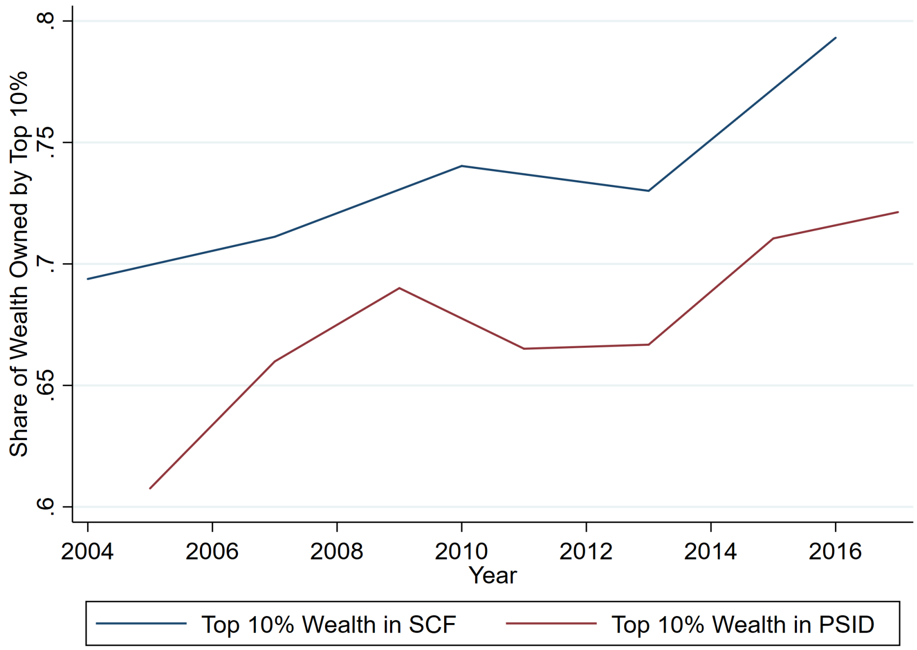

Notes: This figure depicts the share of wealth held by the top 10% of the wealth distribution in the PSID and the SCF. Definitions of wealth for each series are described in the main text.

Figure 1 shows the share of wealth held by the top 10% of the wealth distribution. Both surveys show increasing levels of wealth concentration at the top end of the distribution. In 2004, the SCF reports that the top 10% hold 69% of the wealth in the United States compared to a reported level of only 60% in 2005 by the PSID. By the end of the sample, the share has risen to 79% in 2016 and 72% in 2017 according to the SCF and PSID respectively. Although they follow similar trends, the figure clearly shows that the SCF captures higher levels of wealth concentration than the PSID. We know from the survey sampling methods that the SCF depicts a more accurate picture of wealth at the top of the distribution.

Accordingly, Figure 2 shows wealth inequality calculated in the SCF and PSID. Wealth inequality here is calculated as the ratio of reported wealth in the highest income quintile compared to the lowest income quintile. Inequality calculated from

Figure 2: Comparing Wealth Inequality in the PSID and the SCF

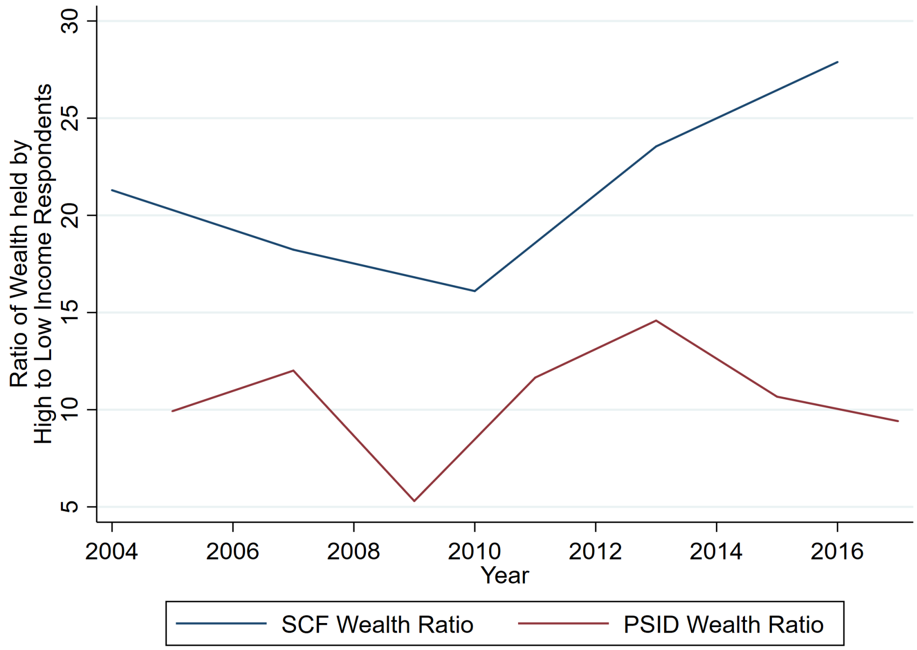

Notes: This figure depicts the ratio of wealth held by the top income quintile to the wealth held by the bottom income quintile. It represents wealth inequality throughout the sample in each survey, as measured by the raw data. Definitions of wealth for each survey are described in the main text.

the PSID has fluctuated throughout the sample, but ultimately has not increased at all. Comparably, when calculated with the SCF data, wealth inequality dropped in 2010 but has steadily increased since. Higher levels of inequality are captured with the SCF than the PSID throughout the entire sample. Both Figures show that the different surveys produce very different results. Pfeffer et al. (2016) attribute the difference between the PSID and SCF wealth data to the fact that the SCF accurately represents wealthy households. As such, Figures 1 and 2 exemplify the importance of using SCF data when evaluating wealth inequality.

The PSID expanded its consumption questions significantly in 1999 and 2005. Since 2005, the PSID presents consumption in the categories of food consumed at home, away from home and food delivered, the cost of housing, transportation, clothing, housing repairs and furnishings, education, childcare, health care, trips and recreation.

Table 1: Average Share of Total Expenditure

	Share PSID	Share (A&B)
Total Food	15.69	16.3
Food at Home	10.88	11.7
Food Out	4.58	4.6
Food Delivered	.24	.
Housing	37.48	32.5
Transportation	19.32	20.6
Clothing	2.89	5.1
Home Repairs	3.82	.
Furnishings	2.24	1.5
Education	3.53	1.3
Childcare	1.22	1.5
Health Care	7.97	5
Trips	3.81	.
Recreation	1.94	2.2

Notes: This table presents the average share of expenditure for each good for 2011-2013. The first column presents the calculations from the PSID and the second column presents the comparable estimates calculated by Aguiar and Bils (2015), where applicable. The categories used by Aguiar and Bils (2015) do not map directly to all aggregates in the PSID. As such, these are rough comparisons. The blank entries exist because no comparable category is available.

The analysis presented in this paper begins with data in 2005 from the PSID in order to ensure a consistent representation of consumption. These categories closely match the aggregated consumption categories presented by Aguiar and Bils (2015), except the PSID does not have separate consumption categories for tobacco and alcohol. Table 1 presents each good's average share of total expenditure for 2011-2013. The first column presents the share of total expenditure in the PSID and the second column provides the share for a similar aggregate compiled from Aguiar and Bils (2015). The blank entries in the second column appear because there are no comparable good categories. The Consumer Expenditure Survey (CE) asks very detailed consumption questions about specific goods and are aggregated by Aguiar and Bils (2015) where the

PSID asks more general consumption questions and presents the data in an aggregated form. Both sources present relatively similar calculated shares; housing comprises the largest share of expenditure, followed by transportation, food at home, and health care in both surveys.

The SCF contains expenditure data only on food consumed at home, away from home and food delivered. The goal of this paper is to impute total consumption in the SCF to evaluate consumption and wealth inequality together. This is achieved by using food expenditure elasticities in the PSID. Table 2 compares raw food expenditure data in the PSID and the SCF in each year the data are available. The first two columns present average spending on food at home, the second two present average spending on food out, and the last two columns present average spending on food delivered. Generally speaking, the average expenditure on food is similar between the two surveys. Average expenditure on food at home and food delivered is slightly higher in the SCF while food away from home is reported to be slightly higher in the PSID compared to the SCF. Interestingly, food expenditure has decreased in all three categories over the sample in the SCF while average expenditure on food at home and food away from home has increased in the PSID.

The demographic control variables used to calculate the PSID elasticities are the age of the reference person, the number of individuals in the family unit, and the number of labor earners in the family unit. The SCF does not provide labor income on the rest of the family, so that demographic control is omitted when calculating the deviation from the mean expenditure in the SCF calculations. Post-tax income in the PSID is calculated using Kimberlin, Kim, and Shaefer (2014) TAXSIM method.

To help motivate the paper, Figure 3 shows the share of wealth, pre and post-tax income, and consumption held by the top 10% of each distribution in the PSID. The share of wealth is much higher than the rest and both wealth and income shares are

Table 2: Comparing Average Food Expenditure

	Food at Home		Food Out		Food Delivered	
	SCF	PSID	SCF	PSID	SCF	PSID
2004	2485.29	.	941.64	.	91.94	.
2005	.	2320.00	.	1087.79	.	69.56
2007	2451.54	2310.44	926.29	1069.58	69.54	65.73
2009	.	2311.08	.	957.70	.	47.56
2010	2461.74	.	859.69	.	54.89	.
2011	.	2286.56	.	940.26	.	52.35
2013	2426.80	2317.73	876.76	963.74	51.87	47.23
2015	.	2394.04	.	1038.29	.	52.94
2016	2246.15	.	933.12	.	55.90	.
2017	.	2503.33	.	1094.32	.	56.51

Notes: This table compares the food expenditure data in the PSID to that presented in the SCF. Expenditure in each food category is compared for the years available.

increasing over the sample. By contrast, the share of consumption by the top 10% was higher than income at the beginning of the sample, but is decreasing steadily over time. Figure 4 presents equivalent trends in the SCF. Specifically, the share of wealth and income held by the top 10% of the distribution increase from 69% to 79% and 43% to 53% respectively from 2004-2016. As in the PSID, the SCF consumption share also decreases slightly over the sample from 26% to 24%. Consumption in Figure 4 is calculated from the raw SCF data and represents aggregate food consumption. It is the sum of food consumed at home, out, and delivered. The contrasting trends of wealth and income to consumption are what provoked much of the research in consumption and income inequality. The more recent literature concludes that consumption data is published with systematic measurement error. These Figures illustrate the importance of considering measurement error when contrasting trends in wealth, income, and consumption inequality. It is clear that failing to account for measurement error when analyzing consumption data will lead to significantly biased results.

The main motivation for this paper is to correct for non-classical measurement er-

Figure 3: Comparing Trends in Shares held by Top 10% in PSID

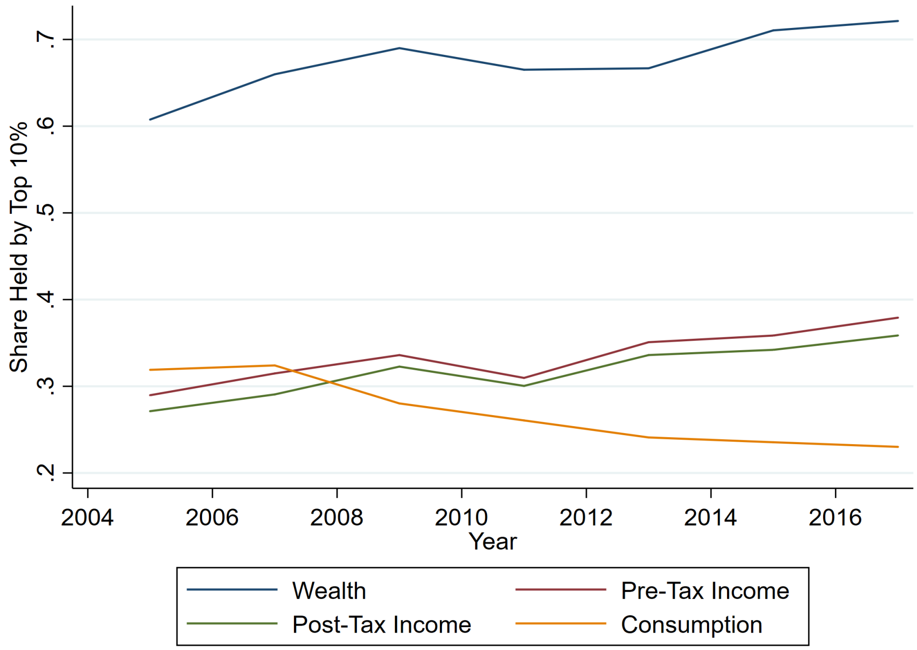

Notes: This figure depicts share of consumption, wealth, and pre and post-tax income held by the top 10% of each distribution. Consumption here is the aggregate of all consumption categories in the PSID. Wealth and pre-tax income are defined in the main body of the text. Post-tax income is calculated using the Kimberlin, Kim, and Shaefer (2014) TAXSIM method.

ror in consumption data in order to contrast corrected consumption inequality with wealth inequality. The extensive consumption data across all good categories in the PSID permits unbiased calculations of total consumption in the PSID while imputing an unbiased measure of total consumption in the SCF enables a comparison of consumption inequality to wealth inequality throughout a well-represented sample. This paper therefore combines the strengths of the PSID and the SCF to re-evaluate trends in consumption and wealth inequality with corrected data.

Figure 4: Comparing Trends in Shares held by Top 10% in SCF

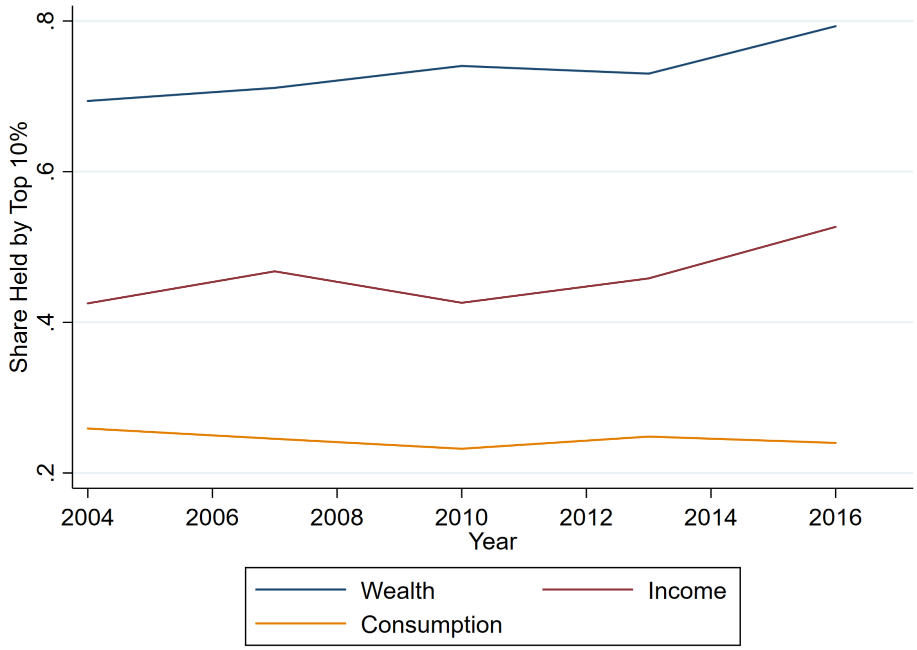

Notes: This figure depicts the share of food consumption, wealth, and income held by the top 10% of each distribution in the SCF. Food consumption is the aggregate of food consumed at home, out, and delivered. The other data categories are defined in the main body of the Data section.

4 Methodology

The methodology section of this paper includes two subsections. The first outlines the method developed by Aguiar and Bils (2015) while the second outlines the extension in this paper.

4.1 Aguiar and Bils (2015)

This section presents the model developed by Aguiar and Bils (2015). The goal of this method is to estimate consumption inequality free of non-classical measurement error. It accomplishes this by using a demand system to calculate total consumption, accounting for systematic measurement error. The key idea is that consumption inequality for Aguiar and Bils (2015) is estimated by using the ratio of consumption across good categories rather than the level of consumption. Since the ratio of spending on relative luxuries and necessities identifies total expenditure, the method does not require that overall spending is measured perfectly. As it will be shown below, this method accounts for both household-specific and good-specific measurement error.

Let us first classify the measurement error. Let $h = 1, ..., H$ indicate households, let $i = 1, ..., I$ be income groups where $I = 5$, let $j = 1, ..., J$ be categories of goods, let t index the year, and let x_{hjt} denote observed spending by household h on good j at time t. Note that $X_{ht} = \sum_{j=1}^{J} x_{hjt}$ indicates the total expenditure by household h at time t. Let ζ_{hjt} denote the measurement error. We can allow observed consumption to be measured with error where x^*_{hjt} is the true level of consumption:

$$x_{hjt} = x^*_{hjt} e^{\zeta_{hjt}} \tag{1}$$

The error in equation (1) can be decomposed into 3 components. Let ψ^j_t denote the error that is common across all individuals for the consumption of good j at time t. Let ϕ^i_t be the error that is common within a specific income group i at time t. Finally, let v_{hjt} be the residual error that is household-good specific at time t. Aguiar and Bils (2015) assume that v_{hjt} is classical measurement error, and is therefore independent of j and h at t. Both ψ^j_t and ϕ^i_t are systematic and need to be eliminated through the estimation process. Equation (2) below illustrates the error term, ζ_{hjt}, decomposed into its three components:

$$\zeta_{hjt} = \psi^j_t + \phi^i_t + v_{hjt} \tag{2}$$

There are two main steps in the estimation method presented in Aguiar and Bils (2015). The first step consists of estimating total expenditure elasticities for each good from a log-linear Engel curve. This step eliminates good-specific measurement error by taking differences across households. The second stage inverts the demand system to recover how consumption inequality has evolved throughout the observed time period. Income-specific measurement error is accounted for in this stage through income-time dummmies. To estimate the expenditure elasticity for each good, Aguiar and Bils (2015) estimate a log-linear approximation to the Engel curves:

$$lnx^*_{hjt} - ln\bar{x}^*_{jt} = \alpha^*_{jt} + \beta_j lnX^*_{ht} + \Gamma_j Z_h + \varphi_{hjt} \tag{3}$$

Here, $\bar{x}^*_{jt}$ represents the average spending on good j at time t for all individuals, Z_h is a collection of demographic dummies for each household that include the age of the head of the household, the number of income earners, and the size of the household. The coefficient on this vector of demographic dummies, Γ_j, varies across goods. Any changes in demand over time that are driven by relative prices are captured by the good-time intercepts, α^*_{jt}. The error term represents individual taste shocks and the second-order error from the log-linear approximation. The coefficients of interest from equation (3) are the expenditure elasticities, β_j, for each good j. Aguiar and Bils (2015) assume that the elasticity for each good is stable across time and test this assumption by estimating the elasticities in various time periods. In terms of observables, the equation becomes:

$$lnx_{hjt} - ln\bar{x}_{jt} = \alpha_{jt} + \beta_j lnX_{ht} + \Gamma_j Z_h + u_{hjt} \tag{4}$$

The observed residual u_{hjt} in equation (4) is composed of the income-specific systematic measurement error ϕ^i_t, the mis-measurement term v_{hjt}, and the individual taste shocks φ_{hjt}:

$$u_{hjt} = \phi^i_t + v_{hjt} + \varphi_{hjt} \tag{5}$$

There are several things to note when estimating equation (4). First, not every individual has positive expenditure in each expenditure category. As such, instead of using the log specification, the authors use the percentage deviation from the average expenditure on good j at time t. This will later be represented by $\tilde{x}_{hjt} = \frac{x_{hjt} - \bar{x}_{jt}}{\bar{x}_{jt}}$. Second, because the mean observed expenditure on good j at time t is included on the left-hand side of equation (4), the good-time specific systematic measurement error denoted above as ψ_t^i has been differenced out. Finally, υ_{hjt} in the error term will be correlated with observed total expenditure. To address this, an instrumental variables approach is used to instrument total expenditure.

Aguiar and Bils (2015) use two instruments for total expenditure. The first method instruments total expenditure with a dummy for the household's income group and a continuous variable for log after-tax income while the second instrument exploits specific timing in the CE survey questionnaires. The second method is not applicable in this paper but both instruments report similar results in Aguiar and Bils (2015). The IV methods used by Aguiar and Bils (2015) are constructed to solve only the classical measurement error[2]. Because the IVs only account for the classical measurement error, the instruments may still be correlated with the systematic measurement error that is common within an income group i at time t, denoted by ϕ_t^i in equation (5). This would lead to biased estimates of the expenditure elasticities estimated in equation (4). In this case, the elasticities may vary from one year to the next due to differences in income-specific measurement error. The estimated inequality calculated from the second-stage is therefore conditional on the initial level of observed of consumption inequality. Aguiar and Bils (2015) explore the stability of the expenditure elasticities, and their effect on calculated inequality, by estimating the first-stage using various years. They find that the elasticities are very stable and other years produce very similar inequality results. Their reported measures for the log change in consumption inequality from 1980-2010 vary by at most 0.05 when using alternate years

2. This IV method is similar to that used by Blundell, Pistaferri, and Preston (2008) to account for classical measurement error when imputing consumption in the PSID from the CE.

to estimate their first stage. The IV addresses the classical measurement error in the observed expenditure data. The next stage will address the systematic measurement error.

Once the expenditure elasticities are calculated in the first step, the second step inverts the demand system to estimate changes in consumption inequality over time. The dependent variable, $\hat{x}_{hjt}$, is the deviation from the mean expenditure on good j by household h at time t, after adjusting for demographics:

$$\hat{x}_{hjt} \equiv \tilde{x}_{hjt} - \hat{\Gamma}_j Z_h \tag{6}$$

Substituting the demographic adjustments into equation (3) and then substituting in the average log expenditure for income group i provides us with our estimating equation by income quintile. Beginning with equation (3), we have:

$$\begin{aligned} ln x^*_{hjt} - ln \bar{x}^*_{jt} &= \alpha^*_{jt} + \beta_j ln X^*_{ht} + \Gamma_j Z_h + \varphi_{hjt} \\ \hat{x}_{hjt} &= \alpha^*_{jt} + \phi^i_t + \beta_j ln X^*_{ht} + \varphi_{hjt} + v_{hjt} \\ &= \alpha^*_{jt} + \phi^i_t + \beta_j ln X^*_{it} + \beta_j (ln X^*_{ht} - ln X^*_{it}) + \varphi_{hjt} + v_{hjt} \\ &= \alpha^*_{jt} + \phi^i_t + \beta_j ln X^*_{it} + \varepsilon_{hjt} \end{aligned} \tag{7}$$

where the residual term has been simplified to include $\beta_j(lnX^*_{ht} - lnX^*_{it}) + \varphi_{hjt} + v_{hjt}$.

Treating the calculated elasticities from the first step as data, consumption inequality can be calculated by regressing $\hat{x}_{hjt}$ on a vector of good-time dummies and income-time dummies, denoted below by D^j_t and D^i_t respectively, and the estimated elasticities from the first stage interacted with income-time dummies:

$$\hat{x}_{hjt} = \alpha^*_{jt} D^j_t + \phi^i_t D^i_t + \hat{\beta}_j D^i_t ln X^*_{it} + \varepsilon_{hjt} \tag{8}$$

The coefficient on the term interacting the elasticities with the income-time dummies represents the estimate of total expenditure by individuals in income group i at time t, lnX^*_{it}. In order obtain an unbiased and consistent estimate of total expenditure, we require that the error term be uncorrelated with $\hat{\beta}_j D^i_t$; the individual component of mis-measurement v_{hjt} and the individual taste shock φ_{hjt} must be orthogonal to the expenditure elasticities conditional on income group. The systematic error is accounted for through the income-time dummies and the good-time dummies[3]. The standard errors are computed by bootstrapping. The measure of consumption inequality presented is then the difference between total expenditure for income group 5 and that of group 1 in a given time period: $lnX^*_{5t} - lnX^*_{1t}$. For ease of interpretation, the lowest income group is kept as the base level so all expenditure estimates presented are relative to $i = 1$. In this way, Aguiar and Bils (2015) present an estimate of trends in consumption inequality that account for non-classical measurement error.

4.2 Application and Extension to SCF

This paper is primarily interested in how trends in consumption inequality compare to trends in wealth inequality. The richest and most reliable source for wealth data in the United States is found in the Survey of Consumer Finances (SCF). However, the SCF only contains expenditure data on food consumed at home, food consumed away from home and food delivered. Therefore, in order to calculate consumption inequality in the SCF, I impute total consumption using expenditure elasticities calculated from the PSID. This method accounts for systematic measurement error as in Aguiar and Bils (2015) and exploits the variation between the expenditure elasticities of a relative necessity (food consumed at home) and a relative luxury (food consumed away from home) to impute total consumption in the SCF. The variation in the β_j for each

3. Aguiar and Bils (2015) outline possible scenarios where their identification system may fail and execute several robustness checks to validate their methodology. I encourage curious readers to see their original work for the details of all exceptions.

food category is enough to identify an estimate of total consumption for each income and wealth group. An additional modification is that I also include information for respondents' wealth in the first stage of the estimation process. The error shown in equation (1) can be further decomposed to include the error that is common within a specific wealth group w at time t, represented by ϕ_t^w. Using the same notation as above, equation (2) therefore becomes:

$$\zeta_{hjt} = \psi_t^j + \phi_t^i + \phi_t^w + \upsilon_{hjt} \tag{9}$$

The first stage is estimated with PSID data to recover the expenditure elasticities. Using the same notation as above and where $\tilde{x}_{hjt} = \frac{x_{hjt} - \bar{x}_{jt}}{\bar{x}_{jt}}$ is the percentage deviation from the average expenditure on good j in year t by household h. The first stage is estimated by:

$$\tilde{x}_{hjt} = \alpha_{jt} + \beta_j lnX_{ht} + \Gamma_j Z_h + u_{hjt} \tag{10}$$

where the residual u_{hjt} is composed of income-specific systematic measurement error ϕ_t^i, wealth-specific systematic measurement error ϕ_t^w, mis-measurement υ_{hjt} and taste shocks φ_{hjt}:

$$u_{hjt} = \phi_t^i + \phi_t^w + \upsilon_{hjt} + \varphi_{hjt} \tag{11}$$

The matrix Z_h contains dummy variables for the age of the head (16-24, 25-37, 38-50, 51-64, 65-80, and 81-104), the number of individuals working in a family unit (0, 1, 2, 2+), and the size of the family unit (less than or equal to 2, 3 or 4, and 5 or more). The demographic coefficient, Γ_j, varies with each good j and α_{jt} is a good-time dummy. As above, the correlation between the mis-measurement term embedded in the observed residual, υ_{hjt}, and total expenditure will result in classical measurement error. I account for this by instrumenting total expenditure with dummy variables for both income and wealth and continuous variables for after-tax income and wealth. It is important to incorporate wealth in this IV specification because, as previously noted, academics have observed that the percentage of wealth consumed is a decreasing function of wealth; wealthier people consume increasingly smaller portions of their wealth. I present the recovered elasticities where only income is used as an IV, only wealth is used, and where both income and wealth are used to instrument total expenditure in Table 3. This IV method addresses the classical measurement error from mis-measurement in the expenditure data. As in Aguiar and Bils (2015), expenditure elasticities are assumed to be stable over time in order to produce unbiased first-stage estimates. Since the left hand side of the estimated equation, $\tilde{x}_{hjt}$, is the percentage deviation from mean expenditure, good-time specific measurement error is differenced out in this stage. The variables of interest from the first stage are the expenditure elasticities for each good, β_j.

The second stage is calculated in both the PSID and the SCF and uses both income and wealth dummies to identify total consumption. The demand system is inverted to recover total expenditure by income and wealth group. First, the deviation from mean expenditure on good j by household h at time t after taking demographics into account must be calculated:

$$\hat{x}_{hjt} \equiv \tilde{x}_{hjt} - \hat{\Gamma}_j Z_h \tag{12}$$

Adapting equation (8) to include the wealth data, consumption inequality by wealth and income groups are recovered by regressing $\hat{x}_{hjt}$ on a vector of good-time dummies, denoted by D_t^j, wealth-time and/or income-time dummies, which are denoted by D_t^i and D_t^w respectively, and the estimated elasticities from the first stage interacted with income-time dummies and/or wealth-time dummies. I estimate the following equation:

$$\hat{x}_{hjt} = \alpha^*_{jt} D_t^j + \phi_t^i D_t^i + \phi_t^w D_t^w + \hat{\beta}_j D_t^i D_t^w ln X^*_{iwt} + \varepsilon_{hjt} \tag{13}$$

The coefficient on the final term represents the estimate of total consumption by income group i and/or wealth group w at time t, lnX^*_{iwt}. I estimate three specifications of equation (13). The first includes only the dummy variables for each income quintile; ie. D^w_t does not exist. In this case, inequality is the difference between the total log consumption by income quintile 5 and income quintile 1. Mathematically, it is represented by $lnX^*_{it} - lnX^*_{it}$, where i are 5 and 1, respectively. This is the exact estimation method used in Aguiar and Bils (2015). The second specification includes only the wealth quintile dummies. Since D^i_t are not included, inequality in the second specification is defined as the difference in total log consumption by wealth quintile 5 and wealth quintile 1. It is $lnX^*_{wt} - lnX^*_{wt}$, where w are 5 and 1, respectively. The final specification incorporates both income and wealth dummies, as presented in equation (13). In this case, inequality is the difference between total log consumption by those in both income quintile 5 and wealth quintile 5 to those in both the bottom income and wealth quintiles. This estimate of inequality is seen as $lnX^*_{55t} - lnX^*_{11t}$ where the numbers 5 and 1 represent the income and wealth quintiles. When both income and wealth quintile dummies are used, the final term produces 25 coefficients (one for each wealth and income quintile interaction). All three specifications are presented in the results section. The tables specify which specification the columns refer to by stating whether the elasticities are interacted with income dummies, wealth dummies, or both.

Equation (13) is first estimated in the PSID using the variation uncovered by consumption across all goods. It is then estimated using only the variation between food consumed at home, away, and delivered in both the PSID and the SCF. Because these categories are relative luxuries and necessities, and their resulting expenditure elasticities are different, we are still able to identify lnX^*_{it}, lnX^*_{wt}, and lnX^*_{iwt} through the variation within the demand system. We assume that the individual mis-measurement, v_{hjt}, and the individual taste shocks, φ_{hjt}, are orthogonal to the expenditure elasticities conditional on income or wealth group. The systematic measurement error is accounted for through the good-time dummies and the income-time or wealth-time dummies in both specifications.

This application of the Aguiar and Bils (2015) methodology accounts for systematic measurement error in consumption data and uncovers how consumption inequality has evolved over time with respect to wealth inequality. The expenditure elasticities are calculated with PSID data, which represents all categories of consumption, and differences out good-time specific measurement error. Next, the elasticities from the PSID are used with the SCF data to impute total expenditure for each income and wealth quintile. This step accounts for income/wealth-specific measurement error by including income/wealth-time dummies. The results show unbiased estimates of consumption inequality.

5 Results

Table 3 presents the results from the first stage of the estimation process. The first column presents the estimated elasticity from each good's Engel curve using both wealth and income as an IV for total expenditure, the third and fifth columns present the results using only wealth or income as instruments, respectively. The standard errors are reported next to each estimate. The second stage estimations employ the estimates from the first column. The elasticities are fairly consistent across all IV methods but are less similar when the share of expenditure is smaller. As expected, the results show that clothing, home repairs, furnishings, education, childcare, trips and recreation are relative luxuries while food at home and transportation are relative necessities. Looking at the first column, we can see that food at home has a relatively low expenditure elasticity of 0.52 while food out and food delivered have higher elasticites of 1.05 and 0.96 respectively. This variation is important because it

Table 3: Expenditure Elasticities

	All		Wealth		Income	
	Beta	SE	Beta	SE	Beta	SE
Total Food	0.68	0.01	0.65	0.02	0.69	0.01
Food at Home	0.52	0.01	0.49	0.02	0.51	0.02
Food Out	1.05	0.03	1.01	0.04	1.10	0.03
Food Delivered	0.96	0.12	1.07	0.15	1.01	0.13
Housing	1.00	0.03	1.02	0.04	1.10	0.03
Transportation	0.76	0.02	0.73	0.03	0.77	0.03
Clothing	1.45	0.11	1.74	0.15	1.77	0.13
Home Repairs	1.63	0.20	2.49	0.28	2.22	0.24
Furnishings	1.60	0.13	1.87	0.17	1.90	0.15
Education	0.97	0.10	1.82	0.12	1.56	0.10
Childcare	1.54	0.10	1.09	0.13	1.41	0.11
Health Care	0.91	0.04	1.10	0.05	1.07	0.04
Trips	1.79	0.05	2.11	0.06	2.11	0.05
Recreation	1.53	0.08	1.70	0.11	1.74	0.10

Notes: This table presents the calculated expenditure elasticities from the first stage of the estimation process. The first two columns present the elasticities and standard errors when both continuous wealth and income and the income and wealth quintiles are used to instrument log total expenditure. The next columns present the elasticities and standard errors when only continuous wealth and the wealth quintiles are used or only continuous income and income categories are used, respectively.

is what permits identification when imputing total consumption in the SCF.

The estimated trends in consumption inequality in the PSID are presented in Table 4. The first three columns use all elasticities when estimating inequality while the last three use only the food elasticities. Columns (I) and (IV) incorporate only the income dummies from equation (13), columns (II) and (V) incorporate only the wealth dummies from equation (13) and columns (III) and (VI) use both income and wealth dummies in equation (13). The row labeled "Interaction" summarizes which specification is used. The first row shows the estimated consumption inequality in 2005 and each subsequent row shows the log change in consumption inequality between

Table 4: Trends in Consumption Inequality in the PSID

	(I)	(II)	(III)	(IV)	(V)	(VI)
2005	1.42	0.83	1.37	1.39	0.53	0.89
2007	-0.26	-0.02	-0.34	-0.18	-0.03	-0.25
2009	-0.27	-0.14	-0.58	-0.08	-0.05	-0.40
2011	-0.16	0.06	-0.22	0.05	0.13	-0.64
2013	-0.37	-0.24	-0.60	-0.04	0.07	-0.46
2015	-0.40	-0.09	-0.51	-0.00	-0.19	-0.63
2017	-0.46	-0.19	-0.67	-0.27	-0.03	-0.53
Interaction	Income	Wealth	Both	Income	Wealth	Both
Elastities Used	All	All	All	Food	Food	Food

Notes: This table presents the estimated change in consumption inequality in the PSID. The specifications differ in which interaction term is used on the total consumption coefficient. Columns (I) and (IV) interact the elasticities with only the income quintile dummies; columns (II) and (V) use only the wealth dummies; columns (III) and (VI) use both income and wealth, as it is explicitly written in equation (13). The first three columns show the estimated inequality using all elasticities and the last three use only the food elasticities. The first row shows the initial level of log inequality in 2005 and each subsequent row presents the change in log inequality with respect to the base year.

the highest and lowest income and wealth quintiles relative to the base year (2005). Looking first at consumption inequality estimated with the income quintiles, the first row reports an estimated log inequality of 1.42 in 2005. However, the results show decreasing consumption inequality throughout the sample. There is a 27% decrease by 2009 and a 46% decrease through to 2017. The second column estimates much smaller levels of inequality in the base year. In 2005, the second specification, which estimates inequality with the wealth quintiles, estimates a log inequality of 0.83. By 2009, inequality decreases by 14% using the wealth specification and 19% by 2017. The difference in inequality estimated using income and wealth illustrates the phenomena that wealthy individuals do not consume as much as they should. As such, consumption inequality is smaller when wealth quintiles are used. The final specification using both income and wealth dummies reports inequality of 1.37 in the base year which decreases 58% by 2009 and 67% by 2017.

Figure 5: Comparing Estimates using all Elasticities vs. only Food Elasticities

(a) Income Interaction

(b) Wealth Interaction

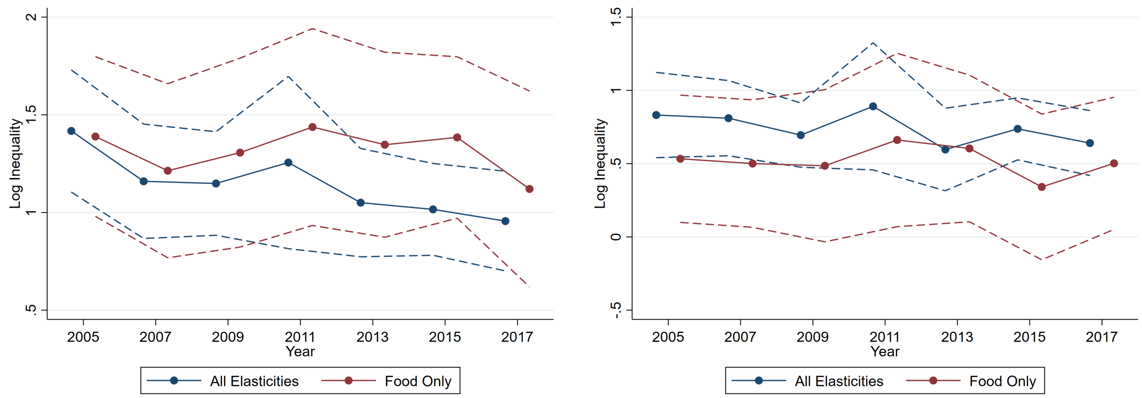

Notes: These figures compare the estimates of log inequality using all elasticities to log inequality calculated using only the food elasticities in the PSID. The solid lines represent the estimates of log inequality in each year while the dashed lines represent the 95% confidence intervals. The Figure on the left, 5a, reports the inequality calculated with respect to the income quintile dummies. They are comparable to the columns (I) and (IV) of Table 4. The Figure on the right, 5b, reports the inequality calculated with respect to the wealth quintile dummies. These results are comparable to the estimates presented in columns (II) and (V) of Table 4.

Comparing the inequality estimates using only food elasticities to those using all elasticities, we can see that all series follow the same trend. Estimated inequality is highest when income quintiles are used in both specifications. However, the series presented in column (IV) shows a significantly smaller decrease in consumption inequality than column (I); by 2017, inequality falls by only 27% compared to 46% when all elasticities are used. Initial inequality using the wealth quintiles and only food elasticities is only 0.53 compared to 0.83 when all elasticities are used. Both series using the wealth specification remain relatively stable over the sample.

To justify only using the food elasticities, Figures 5a and 5b show the inequality calculated with all elasticites and with only the food elasticities in the PSID. Each series is accompanied by a 95% confidence interval, shown in dashed lines. The figures show that both series largely fall within the same confidence intervals. The errors are larger when only food elasticities are used. Nevertheless, the figures support our ability to impute total consumption in the SCF using only the food elasticities.

Finally, Table 5 presents inequality estimated in the SCF. Recall that these estimates are obtained by using the food elasticities calculated in the PSID to estimate total consumption in the SCF. As above, the first row presents the estimated inequality in 2004, which is the base year for the SCF. The remaining rows show the log change in consumption inequality between the year listed and the base year. Again, all three specifications are used and are identified in the table by the row labelled "interaction". Similar to the results using PSID data, the inequality estimate is significantly higher when it is estimated with respect to income quintiles rather than wealth quintiles. The SCF results show increasing inequality throughout the sample in all specifications. Consumption inequality between the highest and lowest income quintiles increases increases 36% throughout the sample. When wealth is used, inequality increases by 2% over the sample. When inequality is calculated by the difference in total log consumption by individuals in both the top income and wealth quintiles to those in the bottom quintile for both income and wealth, inequality increases by 7%. Because the SCF captures the top end of the wealth and income distributions more accurately, we are able to obtain higher levels of imputed consumption for the wealthy individuals. These higher levels of consumption at the top end of the distribution are why inequality appears to increase when estimated in the SCF as opposed to with the PSID data.

The SCF has a significantly better representation of the top end of the income distribution. Roughly speaking, the PSID sample should produce similar data to the bottom 90 percent of the SCF. To illustrate that the increased levels of consumption inequality observed in the SCF data are due to its superior sampling methods, Table 6 estimates consumption inequality in the SCF with a restricted sample. Table 6 presents results comparable to Table 5, but they are calculated without the top 10% of the income and wealth distributions in each year. As in Table 5, the first column

Table 5: Trends in Consumption Inequality in the SCF

	(I)	(II)	(III)
2004	1.22	0.98	1.13
2007	0.16	0.09	0.18
2010	0.45	0.39	0.17
2013	0.55	0.37	0.39
2016	0.36	0.02	0.07
Interaction	Income	Wealth	Both

Notes: This table presents the estimated change in consumption inequality in the SCF. The first column uses the income specification, the second column uses the wealth specification; the third column uses dummy variables for both income and wealth quintiles when estimating total consumption. The first row shows the base level of inequality in 2004 and each subsequent row shows the change in inequality relative to the base.

represents inequality with respect to income quintiles, the second column presents inequality with respect to wealth quintiles and the final column uses both the income and wealth quintiles to calculate inequality. The overall levels of inequality are significantly smaller with the restricted sample across all three specifications. As predicted, the restricted sample also indicates that consumption inequality has decreased over from 2004-2016, as in the PSID. The first specification shows a decrease in log consumption inequality by 14% from 2004 to 2016. Although the levels of falling consumption inequality are not as large in the restricted SCF sample as in the PSID estimations, it is clear that the increasing inequality measured in the full SCF is due to its inclusion of the top ends of the income and wealth distributions. The PSID produces biased estimates because it fails to capture the top end of the distribution. This significant difference highlights the importance of imputing consumption and SCF data.

Table 6: Trends in Consumption Inequality in the SCF from the bottom 90% of the distribution

	(I)	(II)	(III)
2004	1.08	0.46	0.47
2007	-0.20	-0.45	-0.26
2010	-0.04	0.19	-0.04
2013	0.17	0.05	0.25
2016	-0.14	-0.05	0.07
Interaction	Income	Wealth	Both

Notes: The SCF data used for this table has been restricted to omit the top 10% of the income and wealth distribution in each year. This table presents the estimated change in consumption inequality in the SCF, once the data has been restricted. The first column uses the income specification, the second column uses the wealth specification; the third column uses dummy variables for both income and wealth quintiles when estimating total consumption. The first row presents the base level of inequality. Each subsequent row shows the change in log inequality relative to the base year.

The impact of imputing consumption in the SCF and the PSID is clearly seen when we compare estimated consumption inequality to the level of consumption inequality calculated in the raw data for both the SCF and the PSID. These comparisons are presented in Figure 6 and Figure 7. Figure 6 shows that consumption inequality in the PSID raw data has remained relatively constant throughout the sample. By comparison, the PSID estimated consumption inequality shows a significant downward trend in consumption inequality. Ignoring the systematic measurement error present in the data produces biased estimates of consumption inequality while the estimated measure captures much more variation. Figure 7 represents the same comparison of raw inequality estimates to calculated consumption inequality in the SCF. The SCF

Figure 6: Comparing Calculated Consumption Inequality to Raw Inequality in the PSID

Notes: This figure compares consumption inequality observed in the raw PSID data to the consumption inequality calculated with the estimation method in the PSID. Inequality in the raw dataset is defined as the ratio of consumption by the top income quintile to the consumption by the bottom income quintile. The derivation of the calculated measure of inequality is described in the text, with the income specification.

shows very low and stable levels of consumption inequality throughout the sample of raw data. Consumption inequality begins to increase a bit in 2013, but the levels are still much smaller than in the PSID. The imputed levels of consumption inequality in the SCF are, however, increasing steadily over the sample. The ratio of consumption by the top income quintile to the bottom income quintile increases from 1.25 to 1.79 in 2013, before falling to 1.59 in 2016. As in the PSID, the imputed estimates of consumption inequality in the SCF show much more variation and higher levels of consumption inequality than the raw estimates. These graphs illustrates how misleading the biased consumption data can be. They help summarizes the importance of the process presented in this paper by exemplifying the bias found in the raw data.

Figure 7: Comparing Calculated Consumption Inequality to Raw Inequality in the SCF

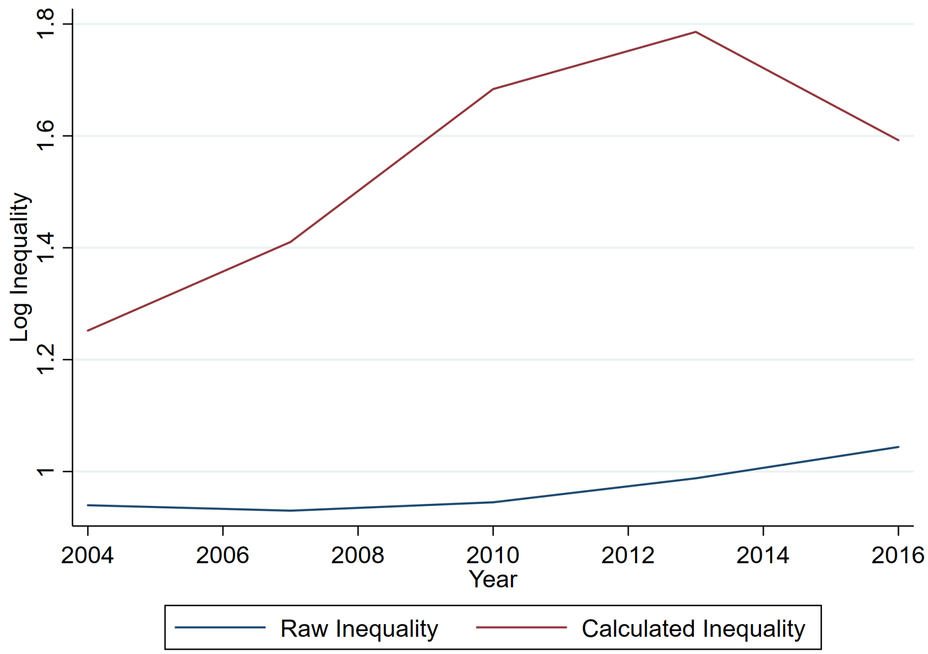

Notes: This figure compares consumption inequality observed in the raw SCF data to the consumption inequality calculated with the imputation method in the SCF. Inequality in the raw dataset is defined as the ratio of consumption by the top income quintile to the consumption by the bottom income quintile. The derivation of the calculated measure of inequality is described in the text, with the income specification.

Recall that the initial wealth inequality trends presented in Figures 1, 2, and 3 showed that wealth inequality has been increasing throughout the sample. Consumption inequality that did not correct for mis-measurement and failed to capture consumption throughout the entire wealth distribution initially showed that consumption inequality was decreasing in the sample presented. Once consumption has been corrected for measurement error, the results from the SCF show that consumption inequality has been increasing since 2004. The results show that wealth and consumption inequality are therefore exhibiting the same trends.

6 Conclusion

The results presented in this paper show that consumption inequality has continued to rise in recent years. These results contradict initial analysis with raw consumption data, but do follow the trends in wealth inequality. The basis of this finding lies in the fact that raw consumption data is measured with systematic error and fails to capture accurate trends. The method used in this paper to correct for systematic mis-measurement follows that presented by Aguiar and Bils (2015), but extends it to SCF data. First, expenditure elasticities are estimated for various good categories using an Engel curve approach. Next, the elasticities are used to obtain an estimate for total consumption. Inequality is calculated by comparing total consumption at the top end of the income/wealth distributions to that at the bottom of the distributions. This approach relies on the ratio of goods consumed, and thus eliminates systematic measurement error at the household level. Error is also allowed to vary across a good-year intercept. The model assumes that good elasticities are constant across time, that the demand system is correctly specified, and that there is enough variation within the demand for food at home, food away, and food delivered to identify total consumption in the SCF. The results present unbiased estimates of consumption inequality.

This paper applies the estimation method to the SCF because it is the most extensive survey for wealth data in the United States. Other surveys fail to represent the top end of the income and wealth distributions. The estimation process depends on accurate measures of income and wealth. Therefore, when the top end of the income and wealth distributions are not captured, an estimated measure of total consumption is biased downward. To capture the full extent of consumption inequality, it was necessary to apply this method to data presented in the SCF.

Much of the research on inequality examines trends and causes of income and wealth inequality, but comparably little evaluates consumption inequality. Consumption is more informative about deprivation and relates more directly to utility; it is therefore important to evaluate consumption inequality when researching the impact of increasing inequality. This paper shows that, once accounting for measurement error, consumption inequality does follow the increasing trend of wealth inequality.

References

Aguiar, Mark, and Mark Bils. 2015. "Has consumption inequality mirrored income inequality?" *American Economic Review* 105 (9): 2725–2756.

Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni. 2014. "Estimates of annual consumption expenditures and its major components in the PSID in comparison to the CE." *American Economic Review* 104 (5): 132–135.

Attanasio, Orazio, Erik Hurst, and Luigi Pistaferri. 2012. *The evolution of income, consumption, and leisure inequality in the US, 1980-2010.*

Attanasio, Orazio, and Luigi Pistaferri. 2014. "Consumption inequality over the last half century: some evidence using the new PSID consumption measure." *American Economic Review* 104 (5): 122–126.

———. 2016. "Consumption inequality." *Journal of Economic Perspectives* 30 (2): 3–28.

Battistin, Erich. 2003. *Errors in survey reports of consumption expenditures.*

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption inequality and partial insurance." *American Economic Review* 98 (5): 1887–1921.

Browning, Martin, and Thomas Crossley. 2009. "Are two cheap, noisy measures better than one expensive, accurate one?" *American Economic Review* 99 (2): 99–103.

Fisher, Jonathan, David Johnson, Timothy M. Smeeding, and Jeffrey P. Thompson. 2018. "Inequality in 3-D: Income, Consumption, and Wealth."

Garner, Thesai, George Janini, William Passero, Laura Paszkiewicz, and Mark Vendemia. 2006. "The Consumer Expenditure Survey: A Comparison with Personal Expenditures." *Bureau of Labor Statistics.*

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.

Kimberlin, Sara, Jiyoun Kim, and Luke Shaefer. 2014. "An updated method for calculating income and payroll taxes from PSID data using the NBER's TAXSIM, for PSID survey years 1999 through 2011." *Unpublished manuscript, University of Michigan. Accessed May* 6:2016.

Kopczuk, Wojciech. 2015. "What do we know about the evolution of top wealth shares in the United States?" *Journal of Economic Perspectives* 29 (1): 47–66.

Krueger, Dirk, and Fabrizio Perri. 2006. "Does income inequality lead to consumption inequality? Evidence and theory." *The Review of Economic Studies* 73 (1): 163–193.

Li, Geng, Robert F. Schoeni, Sheldon Danziger, and Kerwin Kofi Charles. 2010. "New Expenditure Data in the PSID: Comparisons with the CE." *Monthly Labor Review* 133 (2): 29–39.

Meyer, Bruce D., and James X. Sullivan. 2009. *Five decades of consumption and income poverty.*

———. 2017. *Consumption and Income Inequality in the US Since the 1960s.*

Orazio, Attanasio, Erich Battistin, and Hidehiko Ichimura. 2007. "What Really Happened to Consumption Inequality in the US?" *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches, University of Chicago Press, Chicago.*

Parker, Jonathan, Annette Vissing-Jorgensen, and Nicholas Ziebarth. 2009. "Inequality in Expenditure in the Twentieth Century." *presentation slides, NBER Summer Institute.*

Pfeffer, Fabian T., Robert F. Schoeni, Arthur Kennickell, and Patricia Andreski. 2016. "Measuring wealth and wealth inequality: Comparing two US surveys." *Journal of economic and social measurement* 41 (2): 103–120.

Sabelhaus, John, and Jeffrey A. Groen. 2000. "Can permanent-income theory explain cross-sectional consumption patterns?" *Review of Economics and Statistics* 82 (3): 431–438.

Saez, Emmanuel, and Gabriel Zucman. 2016. "Wealth inequality in the United States since 1913: Evidence from capitalized income tax data." *The Quarterly Journal of Economics* 131 (2): 519–578.

Skinner, Jonathan. 1987. "A superior measure of consumption from the panel study of income dynamics." *Economics Letters* 23 (2): 213–216.

Straub, Ludwig. 2018. "Consumption, Savings, and the Distribution of Permanent Income." *Unpublished manuscript, Harvard University.*

Wolff, Edward N. 2014. "Household wealth trends in the United States, 1983–2010." *Oxford Review of Economic Policy* 30 (1): 21–43.