Conspicuous Consumption and Inequality

by

Iuliia Nesterova

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Applied Economics

Waterloo, Ontario, Canada, 2023

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	Francisco Alvarez-Cuadrado Associate Professor, Department of Economics, McGill University
Supervisor(s):	Francisco M. Gonzalez Professor, Department of Economics, University of Waterloo
Internal Member:	Ana M. Ferrer Professor, Department of Economics, University of Waterloo Pierre Chaussé Associate Professor, Department of Economics, University of Waterloo
Internal-External Member:	Juan Moreno-Cruz Associate Professor and Canada Research Chair, School of Environment, Enterprise and Development, University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

While Chapter 2 and 3 are not sole-authored (co-authored with my supervisor Professor Francisco M. Gonzalez and Professor Ana M. Ferrer both from the Department of Economics at the University of Waterloo), I have made the major contribution to the work at all stages including the development of the research question, data selection, method choice, implementation of analysis and interpreting results.

Abstract

My research is centered around understanding consumption behavior and its relationship with inequality.

In Chapter 1, I study how consumption inequality in the United States has evolved over time, with a particular focus on distinguishing two major expenditure components: services and goods. I argue that such distinction is important to understand inequality between high and low income groups. I show that increases in consumption inequality over the period 1984-2018 were driven mostly by rising inequality in expenditure on services rather than goods. I further show that most of it was driven by increased inequality in young households expenditure in services, whereas older households have experienced no change in inequality of either good or service expenditures. As modern societies undergo the transformation into service societies, this research contributes to our understanding of the diverse effects of inequality and informs policy decisions to ensure that this transformation benefits all.

Chapter 2 proposes a canonical model of intertemporal choice in which both current and future conspicuous consumption can distort household consumption behavior. What makes our model tractable is that we assume that each consumer cares about the expected comparison of relative consumption, which provides a parsimonious characterization of positional concerns. We show that equilibrium consumption behavior is a function of the distribution of conspicuous consumption in an individual's reference group as well as her own permanent income. In turn, the distribution of conspicuous consumption is a function of the distribution of permanent income. The relevant empirical implication is that an individual's consumption, by itself, is no longer a valid proxy for the individual's permanent income if relative consumption matters.

In Chapter 3, we document a robust effect of visible inequality on household expenditures in the United States over the period 2010–2018. To that end, we exploit variation in the cross-sectional distribution of visible consumption — expenditures in clothing, personal care, food away from home and vehicles — for younger and older households across regions of the United States and over time. We find that rising inequality in expenditure on visible goods within the different groups is associated with an increase in average spending on those same goods as well as an increase in total expenditures by the average household in the group. Our main findings are not likely to be a symptom of correlated differences in preferences across generations, selection effects across geographical locations, alternative sources of state-level variation over time, or measurement error. Rather, they most likely reflect actual distortions associated with consumption externalities. We conjecture that historically low interest rates and the rise of social media underlie our findings.

Acknowledgements

I would like to acknowledge Professors Francisco M. Gonzalez, Ana M. Ferrer, Pierre Chaussé, Thomas Parker, and Jean Guillaume Forand, for their guidance and support, without which this thesis would not have been possible. I want to thank them for always being generous with their time and fostering an open door collaborative environment in the Economics Department at the University of Waterloo. I will cherish the numerous courses, seminars, and conversations that contributed to the stimulating and inspiring academic atmosphere I was fortunate to be a part of.

I am particularly grateful to my supervisor, Francisco M. Gonzalez, for shaping my PhD journey into a meaningful experience. He has provided me with a strong foundation in rigorous and critical thinking and has always being a source of inspiration through his passion to research and original ideas. I would also like to express my gratitude to Ana M. Ferrer for her generosity with her time and her willingness to share her extensive knowledge in empirical analysis. She has consistently offered support and encouragement to me, even with her many commitments.

I'd also like to thank all the faculty members, staff, and my colleagues for their assistance and valuable insights during my PhD. They have contributed to making my experience at the University of Waterloo memorable and meaningful.

Table of Contents

List of Figures

List of Tables

Chapter 1

Consumption and Income Inequality: Goods vs Services

1.1 Introduction

In this paper, I study how consumption inequality evolved over time, with a particular focus on distinguishing two major expenditure components: services and goods. I argue that such distinction is important to understand inequality in consumption, given the increasing importance of service transactions in the economy and the importance of some services, such as health and education, on well-being. Using US household-level data for expenditure and income over four decades, I show that increases in consumption inequality – measured as the ratio of average expenditure of top income quintile to the bottom income quintile – are driven mostly by inequality in service expenditures. Further, the growing inequality in services is more pronounced among younger households.

It is well recognized that consumption is generally a better measure of welfare than income. The main reason being that consumption better reflects permanent income, while current income is subject to considerable year-to-year variation and therefore a less reliable measure of overall well-being. Measures of income also may not reflect all available resources, including wealth or credit, which can be used to maintain consumption in the presence of income shocks.

I argue that understanding the distinctive role of inequality of expenditure on services in overall consumption inequality is important since spending on services is a major component of household budgets, and inequality in spending on services can contribute significantly to overall economic inequality. Spending on services is also critical to human development and social welfare. Access to high-quality healthcare, education and transportation can improve people's health, education outcomes, and mobility, and in turn increase their economic opportunities. In this regard, public policies that aim to reduce income inequality may need to address inequality in access to and utilization of services. Further, policies targeting specific services such as healthcare or education may need to consider how they affect different income groups.

The changing demographics that affect most Western economies are likely to affect consumption in the future. In the US, there were 34.4 million people aged sixty-five or older in 1998, which is predicted to double by 2030. This will likely increase the importance of understanding consumption patterns in more detail.

For this analysis, I use a sample of Consumer Expenditure Survey (CEX) from 1980 to 2018 and consider expenditure in services as household's reported spending on health, utilities, food away from home, domestic services and childcare, education, entertainment fees, vehicle maintenance, public transport and personal care services. Expenditure in goods includes reported spending on vehicle purchases, leasing, food at home, gas, appliances, entertainment equipment, men's and women's clothing, furniture and fixtures, shoes and other apparel, alcoholic beverages, tobacco, children's clothing, personal care goods (Cravino et al.,2019). I define inequality as the ratio of average income (or expenditure) of the top income quantile group to the bottom income quantile, where these are defined using reported before-tax income (Aguiar and Bils, 2015). A cursory examination of the data shows that before-tax income inequality has grown through three decades before the great recession of 2008-2009. Further, while inequality in total expenditure remained constant through the period, inequality in service's expenditure increased by 11 log points and inequality in good's expenditures *decreased* by 21 log points. In particular, if we look at changes in total expenditure after the great recession, inequality in consumption of services has increased by as much as total consumption inequality.

The literature on consumption inequality discusses extensively the challenges associated with identifying trends in consumption inequality over time due to the presence of non-classical measurement error (Pistaferri and Attanasio, 2016). Particularly when the importance of measurement increases over time (Bee at al., 2015). In what follows, I use the methodology developed in Aguiar and Bils (2015) to examine changes in inequality in expenditure on goods and services free of bias due to non-classical measurement error and under-reporting that is present in CEX. Aguiar and Bils (2015) suggested to tackle these issues by looking at how high- and low-income households allocate spending to luxuries and necessities. As a first step, they estimate expenditure elasticities for different goods and services. Then they use the ratio of spending on different categories scaled by the difference in demand elasticities (estimated in a first step) for these categories. This approach allows them to recover the change in distribution of consumption inequality (measured as a difference in spending relative to bottom income group) that is robust to (a) measurement error in total spending, (b) household-specific measurement errors - to correct for more severe underreporting by high-income households - and (c) good-specific measurement errors to account for more severe underreporting for some goods than others. I apply this methodology to the examination of changes in the inequality of expenditure on goods and services separately to assess the influence of each of these components on changes in inequality in total expenditure.

This chapter answers three questions. 1) Do the trends in increasing inequality in consumption observed by Aguiar and Bils (2015) persisted after the great recession? 2) Is inequality in consumption driven by consumption in goods or consumption in services? 3)

Do expenditure patterns between goods and services change along the life cycle?

I find that robust estimates of the separate trends in goods and service expenditure show that inequality in services expenditure has significantly increased over the 1984-2018 period. In contrast, inequality in good expenditures has remained relatively stable over the period. I further examine whether life cycle patterns of expenditure may affect these trends and conclude that most inequality changes in consumption expenditure seem to arise from increased inequality in young households' expenditure in services, whereas older households have experienced no change in inequality of either good or service consumption.

1.2 Consumption in Goods and Services

It is well established that income and consumption inequality do not necessarily track each other. Krueger and Perri (2006) examined the relationship between income and consumption inequality using CEX data from 1980 to 2003 and conclude that over this time period, changes in income inequality do not necessarily lead to changes in consumption inequality. They proposed a theoretical framework that incorporates differences in borrowing constraints and preferences to explain why the relationship between income and consumption inequality may not be straightforward.

One such factor is the existence of potential differences in borrowing constraints. Individuals who face more stringent borrowing constraints might be unable to borrow enough to maintain their consumption levels despite an increase in their income. As a result, an increase in income inequality may not necessarily lead to an increase in consumption inequality. Another factor is the difference in household preferences. Individuals may have different tastes for consumption and saving, leading to variations in how changes in income are allocated between the two. For example, a person who places a higher value on saving may choose to save more of their income than someone who places a higher value on consumption. As a result, changes in income inequality may not have a direct impact on consumption inequality if income is being allocated differently between individuals with different preferences.

Russell (1999) and Paulin (2000) offer empirical evidence of differences in consumption patterns between the young and the old. These studies show that older households increased expenditure by a larger fraction than did younger households between 1984 and 1997 and that trends in terms of what they spent on also changed, with older households spending more in entertainment and food away from home during this period.

The theoretical framework proposed by Krueger and Perri (2006) seeks to incorporate these and other factors to better understand the complex relationship between income and consumption inequality. Their analysis highlights the importance of considering the heterogeneity of individuals in understanding the effects of income inequality on consumption patterns.

A few years earlier, Deaton and Paxson (1994) already highlighted the relationship between intertemporal choice and inequality, using data from a survey of US households. They unravel differences in household discount rates, with low-income households having a higher discount rate than high-income households, meaning that they place less value on future consumption relative to current consumption. They also found that the degree of impatience is higher among households headed by young and less-educated individuals. As a result, consumption inequality tends to be highest among the elderly in the United States. This is primarily due to higher medical expenses and greater health care needs among the elderly, which can lead to greater consumption risk and income shocks if they experience a health crisis. In contrast, younger households tend to have more stable consumption patterns and lower consumption inequality.

The literature on the inequality of household expenditure on services also highlights significant disparities between income groups in accessing and consuming various services. Several studies have documented that high-income households tend to consume more of higher-quality services, such as healthcare, education, and transportation, than low-income households. Boppart (2014) also documents empirical evidence that lower income households spend a larger fraction of their budget on goods than do higher income households, with the largest differences in expenditure patterns arising in health, utilities, domestic care and childcare. He also points out an interesting empirical fact that the share of goods in total expenditure and price of goods relative to services decline over time. Healthcare has emerged as a major source of inequality in consumption of services, with low-income households experiencing higher out-of-pocket expenditures and facing greater difficulty in accessing healthcare services.

There is evidence that changes in the prices of services, such as healthcare and education, have contributed to rising inequality over time. For example, healthcare costs have risen faster than inflation and wages, leading to a larger share of income going towards healthcare for low-income households. Cutler and Lleras-Muney (2010) show that high-income households tend to have better access to and consume more high-quality services such as healthcare, education, and transportation than low-income households. They also highlight the important role that income plays in determining access to and utilization of services.

This chapter contributes to the literature on consumption inequality by introducing the distinction between goods and services. It also explores how trends in consumption inequality have differed between older and younger households.

1.3 Inequality Trends in CEX

In this section, I present the trends in income and consumption inequality using Consumer Expenditure Survey's interview sample (CEX). I begin with the 1984 wave, as this is the earliest year for which the survey has been consistently conducted.[1]

[1] CEX began in 1980. Data for 1980, 1981 waves is available and 1982 and 1983 waves are missing. In 1980 and 1981 waves, households who own their homes were not asked about self-reported rental equivalence, which is a key measure for expenditure on housing.

In CEX, each household is assigned a "replicate" weight to enable mapping of the CEX sample into the national population, which is utilized in all calculations. Expenditure data is collected through household interviews conducted over a period of up to four consecutive quarters, with detailed information on expenditures recorded for the preceding three months. The final interview collects information on earnings, income, and taxes from the preceding 12 months, aligning with the period for expenditures. Income, expenditure, and savings are all recorded at the household level, while demographic dummy variables reflecting the number of household members, number of household earners, and the reference member's age are included in estimating household demand equations.

To construct the sample, I closely follow the approach taken in Aguiar and Bils (2015). In terms of income, I utilize CEX measures of total household labor earnings, total household income before tax, and total household income after tax, which are reported in the last interview and cover the previous 12 months. Before-tax income comprises various sources, such as labor earnings, non-farm or farm business income, security and retirement benefits, social security insurance, unemployment workers' compensation, welfare (including food stamps), financial income, rental income, alimony and child support, and scholarships. While the measure of before-tax income is derived from the CEX, I also include food as pay and other money receipts, such as gambling winnings to match Aguiar and Bils' sample. Furthermore, I include rental equivalence minus out-of-pocket housing costs as part of before-tax income, given its status as a consumption expenditure for homeowners.

After-tax income is calculated by deducting personal taxes, including federal income taxes, state and local taxes, and payroll taxes, from before-tax income. Federal income taxes may be negative as they capture earned income credits. CEX adopted the TAXSIM model developed by the National Bureau of Economic Research in 2015 to enhance the quality of tax measures, resulting in substantial improvements in the availability and accuracy of tax data. In particular, it led to a significant increase in the amounts of tax paid by households and a decline in after-tax income and savings levels. These results are observed for both the population as a whole and for demographic groups within the population in

CEX Interview Survey. (Paulin, 2015).

The CEX reports expenditure on hundreds of separate items, which I aggregate into 20 groups to ensure consistency across various waves of the survey and to cover a broad range of expenditure elasticities. I utilize these 20 groups to estimate demand and further aggregate them into goods and services categories. Expenditures on services include household spending on health, utilities, food away from home, domestic services and childcare, education, entertainment fees, vehicle maintenance, public transport, and personal care services. Expenditures on goods include spending on vehicle purchases and leasing, food at home, gas, appliances, entertainment equipment, clothing, furniture and fixtures, alcoholic beverages, tobacco, and personal care goods. We use rent paid for renters and self-reported rental equivalence for homeowners to measure expenditure on housing services. For durables other than housing, I utilize direct expenditure and do not impute service flows (Aguiar and Bils, 2015).

To document trends in inequality, I categorize households into five groups based on their before-tax income. These groups are 5-20, 20-40, 40-60, 60-80, and 80-95 percentile bins. For each group, I calculate the average income and total expenditure for the households within a group.

Figure 1.1 depicts the high-income to low-income ratio for before-tax income (which includes transfers), after-tax income, and consumption expenditures since 1984 to 2018. To offset substantial year-to-year movements, Table 1.1. summarizes trends in inequality by averaging the high-income to low-income ratio over three-year periods.

Income inequality, particularly before-tax income, increases steadily over time from 1984 to 2018 with a 28 log points increase for before-tax income (see Table 1.1). After-tax income increases substantially less (6 log points). However, this contains the change in computing after-tax income mentioned above. Splitting the changes into before and after the great recession shows a slightly larger increase in after tax income inequality than before-tax income inequality before the great recession.[2] Figure 1.1 shows trends in inequality for these variables over the sample period. Even with the substantial drop in after-tax income inequality from 2015 reflecting the introduction of a new tax data measurement in CEX, the upward trend in after-tax income inequality persists. In contrast to income inequality, consumption inequality remains mostly flat during the same period (Table 1.2. Note, however, that inequality of expenditure in services increased by 11 log points while inequality of expenditure on goods declined by 21 log points.

Table 1.1: Trends in inequality

	1984-1986	**2010-2012**	**2016-2018**	**log change 1984-1986/2016-2018**	**log change 1984-1986/2005-2007**	**log change 2010-2012/2016-2018**
Labor Earnings	7.83	9.02	9.74	0.22	-0.01	0.08
Before Tax Income	5.52	6.65	7.33	0.28	0.14	0.10
After Tax Income	4.99	6.01	5.30	0.06	0.16	-0.13
Total Expenditure	2.88	2.79	2.87	0.00	0.01	0.03

Notes: Table presents descriptive statistics using a sample of Consumer Expenditure Survey after the financial crisis. Columns 1-3 report the average ratio of the top income quintile to the bottom income quintile over three-year periods. Column 4 represents a log difference over the sample period. Columns 5 and 6 further break the change into before and after the financial crisis. All variables are converted into 1983 constant dollars before averaging.

[2] I report similar measures in Table A.1 in the appendix using the sample 1984 to 2010, to show the drastic change in inequality before and after the great recession as reported by Aguiar and Bils (2015).

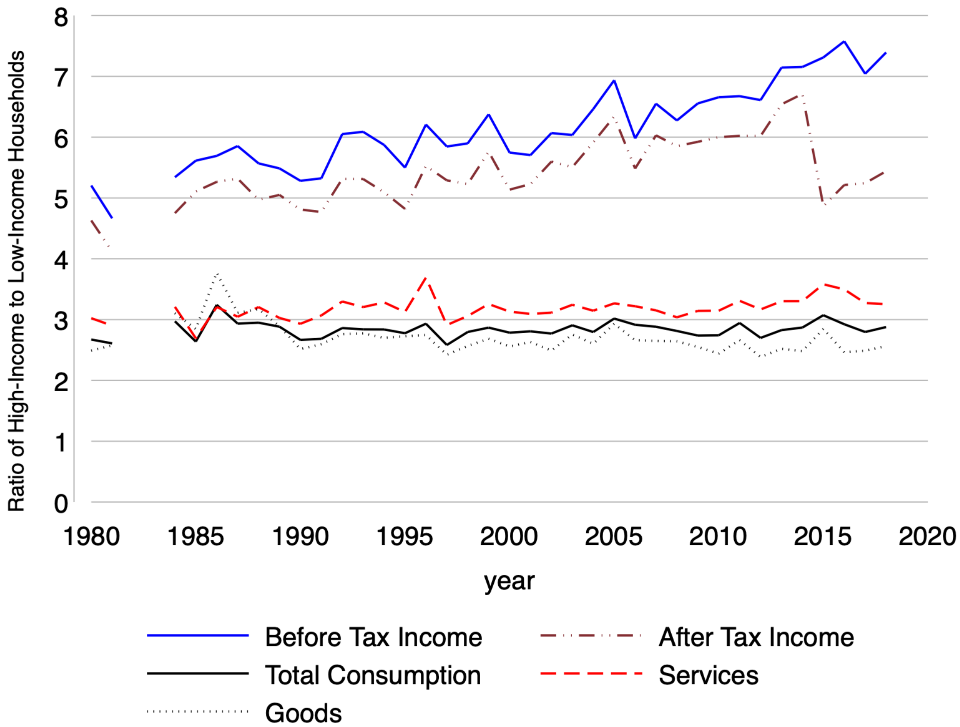

Figure 1.1: Trends in Inequality, 1980–2018

Notes: This figure illustrates trends in income and consumption inequality using CEX sample for period from 1980 to 2018. Inequality is measured as a ratio of high-income to low-income respondents. The high income group includes households with before-tax income larger than 80$^{\text{th}}$ percentile and smaller than 95$^{\text{th}}$ percentile.

Table 1.2: Consumption Inequality: Services vs Goods

	1984-1986	1991-1993	2005-2007	2010-2012	2016-2018	log change 1984-1986/2016-2018	log change 1984-1986/2005-2007	log change 2010-2012/2016-2018
Total Expenditure	2.88	2.77	2.92	2.79	2.87	0.00	0.01	0.03
Expenditure on Services	2.99	3.19	3.2	3.21	3.34	0.11	0.07	0.04
Expenditure on Goods	3.09	2.71	2.71	2.49	2.51	-0.21	-0.13	0.01

Notes: Table presents descriptive statistics using a sample of Consumer Expenditure Survey after the financial crisis. Columns 1-5 reports average ratio of top income quintile to bottom income quintile over three-year periods. All variables are converted into 1983 constant dollars before averaging.

1.4 Methodology

The previous subsection has discussed trends in inequality observed in the Consumer Expenditure Survey. Specifically, how the increase in income inequality is more pronounced when compared to consumption inequality. This discrepancy has also been extensively documented in the existing literature. Heathcote, Perri, and Violante (2010) also show that consumption inequality is significantly lower than income inequality. They argue that this discrepancy arises due to various factors, such as borrowing constraints, precautionary saving, and the role of government transfers and taxes.

Blundell, Pistaferri, and Preston (2008) focus on understanding the factors that contributed to the divergence between consumption and income inequality and argue that consumption inequality may be a more accurate measure of individuals' economic well-being compared to income inequality alone. They suggest that the presence of insurance mechanisms, such as government transfers, private insurance, and familial support contribute to the differences between consumption and income inequality. These mechanisms can help individuals smooth consumption in the face of income fluctuations, reducing consumption inequality relative to income inequality.

However, the observed divergence between consumption and income inequality in CEX does not necessarily imply that income inequality is growing faster compared to consumption inequality. Several studies note so when raising concerns about the validity of consumption data in CEX. They document the presence of nonclassical measurement error and underreporting of income and expenditures by the top income household in CEX. Both of these issues can lead to biased estimates of consumption inequality based on CEX.

In particular, Bee, Meyer, and Sullivan (2015) explore the issue of nonclassical measurement error in the Consumer Expenditure Survey (CEX) and its implications for the validity of consumption data. They observe significant discrepancies in reported expenditure levels and patterns between the Interview Survey and the Diary Survey, suggesting that individuals may not accurately recall or report their expenses. They also find that measurement error varies across different types of consumption expenditures. Certain categories, such as food and housing, exhibit relatively low measurement error, while others, such as clothing and transportation, show higher levels of error. These variations suggest that the accuracy of reported consumption data may be influenced by the nature of the expenditure and the ease of recall.

Attanasio, Hurst, and Pistaferri (2013) demonstrate that the CEX underestimates the expenditure share of high-income households, which can distort measures of inequality relying on CEX data. The authors use data from the Panel Study of Income Dynamics (PSID) where consumption questions were significantly expanded starting from the 1999 survey, and it now provides more detailed and accurate information on expenditures compared to the CEX. They compare the expenditure data from the PSID with the CEX data to identify discrepancies and estimate the degree of underreporting. The paper finds evidence of significant underreporting of expenditure in the CEX, particularly at the top of the income distribution. The authors argue that high-income households may have a stronger incentive to underreport their expenditure due to various factors, such as tax planning strategies, the desire for privacy, or the complexity of their spending patterns. Underreporting is also found to be changing over time, which posses a challenge in capturing trends in consumption inequality in CEX.

These concerns motivate the methodology introduced in Aguiar and Bils (2015). Their paper shows that the growth in consumption inequality in the United States tracks the growth in income inequality after correcting for measurement error in the data. They argue that the apparent divergence between income and consumption inequality in previous studies may have been driven by measurement error in the data used to measure consumption. Once the measurement error is corrected, the authors find that consumption inequality has grown at a similar pace than income inequality over the period 1980-2010.

I adopt a similar methodology to correct for measurement error when looking at consumption in services and consumption in goods separately. The main goal is to avoid measuring consumption inequality directly in CEX, i.e. tracking the total household expenditure of different income groups over time. Instead, Aguiar and Bils (2015) proposed to estimate relative consumption growth across income groups by observing how households in these groups have shifted their expenditures across different goods over time. The key advantage of their approach is that it does not require that overall expenditures of households be well measured.

Aguiar and Bils (2015) methodology consist of two main steps. First, assume that a household's i reported expenditure on good j in the year t x_{ijt} is measured with error, with the degree of mis-measurement depending on the household's income group, the good and time:

$$x_{ijt} = x^*_{ijt} e^{\zeta_{ijt}} \tag{1.1}$$

where x^*_{ijt} is the true expenditure and ζ_{ijt} is the measurement error assumed to be of the following form:

$$\zeta_{ijt} = \psi_{jt} + \phi_{it} + v_{ijt} \tag{1.2}$$

where the first component ψ_t^j represents errors in measuring consumption good j in a year t that do not vary across households, for instance, underreporting of transportation expenses for all households in a given CEX survey wave. The second component ϕ_t^i refers to the measurement error specific to a particular household i in a year t that does not vary across goods, for example, underreporting of all expenditures by wealthy households. The term v_{ijt} refers to the measurement error specific to a particular household and consumption good.

Aguiar and Bils (2015) based their methodology on the identification assumption that v_{ijt} measurement error is of a classical form. Their main idea is to control for systematic mis-measurement at the good-time and income-time level and assume that the residual measurement error at the household-good-time level has zero mean for all households i and all CEX survey years t, i.e. it has to be independent of the good's j expenditure elasticity when we condition on household's income group. To understand why a household's income group is an important factor here, we need to describe the estimation steps that constitute Aguiar and Bils (2015) methodology.

Step I uses log-linear approximation to the Engel curves to estimate total expenditure elasticities for each good.

$$logx_{ijt} - log\bar{x}_{jt} = \alpha_{jt} + \ \beta_j logX_{it} + \Gamma_j Z_i + u_{ijt} \tag{1.3}$$

The dependent variable is the logarithm of the household's i expenditure on good j in year t demeaned using average expenditure on good j across all households observed in the year t. To handle the cases of zero reported expenditure on specific goods, Aguiar and Bils (2015) replace $logx_{ijt} - log\bar{x}_{jt}$ with the percentage deviation from average expenditure on that good in that year: $\tilde{x}_{ijt} = \frac{x_{ijt}-\bar{x}_{jt}}{\bar{x}_{jt}}$. They also consider five types of households, based on households' reported before-tax income.

Subtracting average expenditure on each good across all households differences out errors in measuring expenditure on a good j in a year t that are common across households, the term ψ_{jt} in the equation 1.2 above. On the right-hand side of the equation 1.3 $logX_{it}$ is log total expenditure of a household i in a year t, Z_i is a vector of demographic controls, and u_{ijt} is the error term that includes the following three components: ϕ_t^i measurement error that is constant across goods and specific to a particular household i in a year t, η_{ijt} second-order error of the log-linear approximation to the Engel curves, and v_{ijt} mismeasurement assumed to be of a classical form.

$$u_{ijt} = \phi_{it} + \eta_{ijt} + v_{ijt} \tag{1.4}$$

The main coefficient of interest in Step I is β_j elasticity of household's expenditure on good j with respect to household's total expenditure.

In estimating Engels curves, measurement error in expenditure on individual goods could potentially introduce a correlation between the measurement error captured in the residual and the observed total expenditure. To address this concern, a common approach is to use reported income as an instrument for total expenditure. This approach takes advantage of the fact that total expenditure is determined by the household's permanent income, which is correlated with current income. To instrument total expenditure, we utilize the household's after-tax income along with dummies representing the household's income group.

Step II takes Engel curves estimates obtained in Step I and uses them as regressors to infer total expenditure.

$$\hat{x}_{ijt} = \alpha_{jt} G_{j,t} + \phi_{it} D_{i,t} + (\hat{\beta}_j * D_{i,t}) logX_{it}^* + \epsilon_{ijt} \tag{1.5}$$

The dependent variable is defined as household's i expenditure on a good j in year t after removing the estimated effect of demographic controls $\hat{\Gamma}_j$ obtained in Step I.

$$\hat{x}_{ijt} \equiv \tilde{x}_{ijt} - \hat{\Gamma}_j Z_i \tag{1.6}$$

where $\tilde{x}_{ijt}$ denotes demeaned household's expenditure on each good and Z_i is a vector of demographic controls same as in equation 1.3

On the right-hand side $G_{j,t}$ represents a vector of indicator variables for each good multiplied by year dummies, thus, the term $\alpha_{jt}G_{j,t}$ is meant to capture the effect of errors in measuring consumption good j in year t which do not vary across households. $D_{i,t}$ is a vector of indicator variables for household's income group interacted with year dummies. The term $\phi_{it}D_{i,t}$ isolates the effect of the measurement error specific to a particular household i in the year, t which does not vary across goods. The third term is the main focus. $\hat{\beta}_j * D_{i,t}$ denotes the interaction between Step I elasticity estimates $\hat{\beta}_j$ and the vector of indicators for household's income group interacted with year dummies. The coefficient of this interaction term $logX^*_{it}$ represents estimated total expenditure corrected for the measurement error for each income group. The lowest income group is the baseline income group; therefore, the estimated coefficient represents the total expenditure of households that belong to an income group i relative to the lowest income group in a given year t. To reduce variability in yearly data, I restrict the time dummies to the following three-year windows: 1984-1986, 1991-1993, 2005-2007, 2008-2010, 2016-2018

1.4.1 Services vs Goods

I assume that services are weakly separable from other commodities which together constitute total expenditure and focus on the allocation of expenditure on services among different service components. To that end, I modify Aguiar and Bils(2015) methodology as follows:

In Step I, I estimate the Engel Curves to obtain the elasticity of each component included in expenditure on services with respect to the total household's spending on services:

$$logs_{ijt} - log\bar{s}_{jt} = \alpha_{jt} + \beta_j logS_{it} + \Gamma_j Z_i + u_{ijt} \tag{1.7}$$

where s_{ijt} is household's i expenditure on service j in a year t, e.g. household's expenditure on personal care services, $\bar{s}_{jt}$ is average expenditure on service j across all households observed in year t, and S_{it} is the sum of household's expenditure on all services in year t

Then, in Step II I take the Engel curves estimated elasticities of services obtained in Step I and use them as regressors in a regression of household's expenditure on services with respect to total expenditure in services.

$$\hat{s}_{ijt} = \alpha_{jt} G_{j,t} + \phi_{it} D_{i,t} + (\hat{\beta}_j * D_{i,t}) log S^*_{it} + \epsilon_{ijt} \tag{1.8}$$

The same procedure is replicated to infer trends in a household's expenditure on goods.

In Step I, I use 2SLS to estimate expenditure elasticities of each good, where log total expenditure is deemed to proxy household's permanent income and is instrumented with log after-tax household income and dummy variables for household's income quintile. I use the sample of 1994-1996 CEX waves to estimate expenditure elasticities. The motivation for using this period is to have different sample periods in Steps I and II. Using different samples ensures that residual error terms in Step I and Step II are different, hence, regressors in Step II are orthogonal to the error term. The choice of 1994-1996 period hinges on the assumption that expenditure elasticities are stable over time.[3]. In Step II, I estimate equation 1.8 by ordinary least squares. Note that in Step I equation 1.7 is estimated separately for each good to obtain $\hat{\beta}_j$, while Step II uses a pooled sample of all goods to capture the relative shifts in expenditure across goods with different elasticities.

Note that the validity of Aguiar and Bils's (2015) exercise in this setting relies heavily on the assumption of separability of goods and services when estimating expenditure elasticities. While this approach was employed previously in a study of components of food consumption by Hayashi (1995), its validity in the separation of goods and services is the subject of future work.

[3] Robustness checks conducted in Aguiar and Bils (2015) show that Step II estimates are relatively stable when expenditure elasticities are estimated at different points in the sample

1.4.2 Results

Table 1.3 presents the main results. The estimation procedure is the same in all columns. Column(1) reports estimates of expenditure on all components, while column(2) and column (3) estimate expenditure on services and goods respectively. The first row in Table 1.3 reports total expenditure estimates of the top income households (above 80^{th} percentile before-tax income) relative to the bottom income households (below 20^{th} percentile before-tax income). The next four rows, (2) to (5), estimates changes in total expenditure of top income households relative to bottom income households for the period specified. Standard errors are computed by bootstrap with 400 repetitions to correct for the use of estimated coefficients from Step I. Bootstrap procedure is sampling with replacement from the full sample and estimating Step I and Step II each time.

Results reported in Column(1), rows (1) through (3) are comparable to the estimates reported in Aguiar and Bils (2015). Consumption inequality (ratio of spending by top income households to spending by bottom income household) in the base period 1984-1986 is estimated to be 0.94, i.e. average consumption of top income household was $\exp(0.94) = 2.56$ times larger than average consumption of bottom income households. This differs from the observed total expenditure ratio in row 3 of Table 1.1 (2.88) by only 0.32 ($2.88–2.56 = 0.32$), suggesting that the measurement error did not significantly alter inequality estimates at the beginning of our sample.

The next three rows show the estimated growth in consumption inequality relative to the baseline 1984-1986 period. Inequality was growing through 2005-2007 with a slowdown during the financial crisis period, in agreement with results in Aguiar and Bils (2015). The fifth row shows the change in inequality over the full period. The total increase in consumption inequality between 1984-1986 and 2016–2018 years (fifth row in Table 1.3) continued to slow down beyond the great recession. It is estimated to be 29 percent, which is comparable to the increase in inequality observed in before-tax income and labor earnings.

Column(2) and (3) further decompose this increase into inequality in expenditure on services and goods, these columns summarize the central results of the paper. While the baseline inequality estimates are similar for both expenditure on service and expenditure on goods - top income households on average spend three times more on both services ($\exp(1.152) = 3.16$) and goods ($\exp(1.111) = 3.04$) than their lower income counterparts - the trends in inequality over time between the two differ in a stark manner. Column(2) shows that inequality in expenditure on services has been growing at a similar pace as estimated total expenditure inequality in Column(1). In contrast, inequality in expenditure on goods shows no significant change throughout the sample, with all estimates reported in column (3) being insignificant and small in magnitude.

1.4.3 Age

As mentioned above, preferences are reported to influence consumption behavior, and particularly consumption in goods versus consumption in services. Consumption in services is likely highest among the elderly, who may have already acquired significant amounts of durable goods, but may require higher levels of services due to greater health care needs among the elderly. Older households are also more subject to inequality risk in consumption of services (Deaton and Paxson, 1994). In contrast, younger households tend to have more stable consumption in goods. Partly because the need to accumulate a certain amount of durable goods over time, and also because the lower likelihood of requiring some services, like health services. In this subsection, I explore whether reported inequality trends in expenditure in services and expenditure in goods follow different patterns over the life-cycle and hence may be affected by demographic trends.

To that end, I first look at observed inequality trends in reported expenditure for different age groups. I split the sample into two age groups - 25-44 and 45-64. The age cut-off is chosen to preserve the similar sample size between both groups and also capture potential changes in expenditure patterns related to the households transitioning through

Table 1.3: Estimated change in consumption inequality over time

	Total Expenditure (1)	Expenditure on Services (2)	Expenditure on Goods (3)
Log ratio top to bottom, 1984-1986	0.942***	1.152***	1.111***
	(0.052)	(0.100)	(0.076)
Log change 1984-1986/1991-1993	0.193***	0.301**	-0.036
	(0.072)	(0.136)	(0.098)
Log change 1984-1986/2005-2007	0.396***	0.489***	0.044
	(0.072)	(0.130)	(0.112)
Log change 1984-1986/2008-2010	0.354***	0.385***	-0.058
	(0.072)	(0.128)	(0.097)
Log change 1984-1986/2016-2018	0.288***	0.274*	-0.059
	(0.080)	(0.143)	(0.101)
Categories Included	All	Services	Goods
Observations	592320	296160	266544
R-squared	0.432	0.511	0.264

Notes: This table reports second-stage regression estimates of change in consumption inequality obtained using OLS. Consumption inequality is defined as the average consumption of households in the top income quintile divided by the average consumption of households that fall into the bottom income quintile. Column(1) reports the estimated change in total consumption inequality. Column(2) and (3) break it down to estimated change in inequality of consumption of services and goods, respectively. Row(1) summarizes baseline log-inequality in the period 1984-1986. Rows (2) - (5) show change in consumption inequality in respective periods relative to the baseline years 1984-1986.
All specifications include demographic control dummies for age, household size, and number of earners in the household. The standard errors are bootstrapped with 400 repetitions.

their life cycle.

Table 1.4 reports the ratio of average expenditure of top income households relative to average expenditure of bottom income households, computed using the CEX sample. Overall, households between 45 and 64 year of age experience larger levels of consumption inequality compared to their younger counterparts. This observation is consistent with the existing literature on life-cycle inequality. (Deaton and Paxson, 1994)

The observed movements of inequality in total expenditure are rather small for both age groups. Inequality in expenditure on services grew significant larger for older households relative to younger ones (0.16 vs 0.06 log points) between 1984-1986 and 2016–2018 years. Also, inequality has been falling after 2010-2012 period for total expenditure and both expenditure components for younger households, while it has kept growing for older households.

To provide estimates of inequality trends free of measurement error, I apply the previous methodology to the sample of younger and older households separately. Estimation results for 25-44 and 45-64 age groups are reported in the tables 1.5 and 1.6 respectively. Both tables have the same structure as table 1.3 and only differ in the households' age groups included in the sample.

The estimates show that younger households start with relatively modest levels of inequality both in total expenditure and the goods and services expenditure components in 1984–1986 years (first row in Table 1.5). The top-to-bottom ratio is initially estimated between 2.3 for services and 3.0 for goods. A difference with the main results is that we observe a continuous increase in total consumption inequality for younger households. Inequality grew by 27 log points during the 80s, which corresponds to 31% increase in the ratio of consumption of top to bottom income group, and by 50 log points between the initial period and the great recession (60 log points including the recession). The total growth in inequality is over 50% (fifth row in Table 1.5). This corresponds to an increase in the ratio of 1.7.

Table 1.4: Consumption Inequality over the lifecycle

	1984-1986	**1991-1993**	**2005-2007**	**2010-2012**	**2016-2018**	**log change 1984-1986/2016-2018**	**log change 1984-1986/2005-2007**	**log change 2010-2012/2016-2018**
				Age Group 25-44				
Total Expenditure	2.88	2.79	2.82	2.77	2.72	-0.06	-0.02	-0.02
Expenditure on Services	2.85	3.13	3.09	3.13	3.04	0.06	0.08	-0.03
Expenditure on Goods	3.05	2.69	2.47	2.39	2.26	-0.3	-0.21	-0.06
				Age Group 45-64				
Total Expenditure	2.85	2.72	2.99	2.8	2.98	0.04	0.05	0.06
Expenditure on Services	3.03	3.16	3.25	3.23	3.54	0.16	0.07	0.09
Expenditure on Goods	3.21	2.74	2.95	2.59	2.76	-0.15	-0.08	0.06

Notes: Table presents descriptive statistics using a sample of Consumer Expenditure Survey after the financial crisis. Columns 1-5 reports average ratio of top income quintile to bottom income quintile over three-year periods. All variables are converted into 1983 constant dollars before averaging.

In column (2) I show that this increase in consumption inequality is due to the significant growth in inequality in service expenditures. For the full period, this represents a 4.2 times higher ratio, or almost three times larger (point estimate of 1.43). In contrast, inequality in expenditure in goods has been declining, with a significant decline by 27% between 1984-1986 and 2016-2018.

Older households, on the other hand, initially experienced larger levels of inequality (first row in Table 1.6). Baseline inequality ratio in services expenditure was $\exp(1.385) = 4$, and inequality in consumption $\exp(0.846) = 3.2$, larger than baseline inequality estimated for younger households. However, these estimates do not change significantly over the period, not even for total expenditure.

Table 1.5: Estimated change in consumption inequality over time — Age 25-44

	Total Expenditure (1)	Expenditure on Services (2)	Expenditure on Goods (3)
Log ratio top to bottom, 1984-1986	0.844***	0.846***	1.099***
	(0.054)	(0.102)	(0.131)
Log change 1984-1986/1991-1993	0.271***	0.636***	-0.113
	(0.081)	(0.164)	(0.158)
Log change 1984-1986/2005-2007	0.487***	0.858***	-0.098
	(0.088)	(0.178)	(0.160)
Log change 1984-1986/2008-2010	0.591***	1.247***	-0.107
	(0.096)	(0.195)	(0.168)
Log change 1984-1986/2016-2018	0.541***	1.434***	-0.266*
	(0.120)	(0.259)	(0.160)
Categories Included	All	Services	Goods
Observations	298,160	149,080	134,172
R-squared	0.783	0.838	0.455

Notes: Sample: CEX households with household head being older than 25 but younger than 44. Table reports second-stage regression estimates of change in consumption inequality obtained using 2SLS. Consumption inequality is defined as the average consumption of households in the top income quintile divided by the average consumption of households that fall into the bottom income quintile. Column(1) reports the estimated change in total consumption inequality. Column(2) and (3) further break it down to estimated change in inequality of consumption of services and goods, respectively. Row(1) summarizes baseline log-inequality in the pooled period of 1984-1986. Rows (2) - (5) show change in consumption inequality in respective periods relative to the baseline years 1984-1986.
All specifications include demographic control dummies for age, household size, and number of earners in the household. The standard errors are bootstrapped with 400 repetitions.

Table 1.6: Estimated change in consumption inequality over time — Age 45-64

	Total Expenditure (1)	Expenditure on Services (2)	Expenditure on Goods (3)
Log ratio top to bottom, 1984-1986	1.078***	1.385***	1.172***
	(0.107)	(0.177)	(0.088)
Log change 1984-1986/1991-1993	0.027	-0.059	0.016
	(0.138)	(0.225)	(0.121)
Log change 1984-1986/2005-2007	0.112	0.032	-0.033
	(0.134)	(0.220)	(0.129)
Log change 1984-1986/2008-2010	0.060	-0.151	-0.137
	(0.129)	(0.210)	(0.115)
Log change 1984-1986/2016-2018	0.121	-0.090	-0.079
	(0.142)	(0.230)	(0.118)
Categories Included	All	Services	Goods
Observations	294,160	147,080	132,372
R-squared	0.774	0.786	0.775

Notes: Sample: CEX households with household head being older than 45 but younger than 64. Table reports second-stage regression estimates of change in consumption inequality obtained using 2SLS. Consumption inequality is defined as the average consumption of households in the top income quintile divided by the average consumption of households that fall into the bottom income quintile. Column(1) reports the estimated change in total consumption inequality. Column(2) and (3) further break it down to estimated change in inequality of consumption of services and goods, respectively. Row(1) summarizes baseline log-inequality in the pooled period of 1984-1986. Rows (2) - (5) show change in consumption inequality in respective periods relative to the baseline years 1984-1986.
All specifications include demographic control dummies for age, household size, and number of earners in the household. The standard errors are bootstrapped with 400 repetitions.

1.5 Conclusion

In this paper, I look at changes in consumption inequality over the period 1984 to 2018, using US household-level data for expenditure and income from the CEX. The innovative feature of this analysis is distinguishing between two major expenditure components: services and goods. I argue that such distinction is important to understand inequality in consumption. I follow Aguiar and Bils (2015) empirical methodology to provide estimates that are free of non-classical measurement error and under-reporting present in CEX. The results show that increases in consumption inequality – measured as the ratio of average expenditure of the top income quintile to the bottom income quintile – are driven mostly by inequality in service expenditures.

The results extend our knowledge of the trends in consumption inequality to the year 2018. In this regard, I show that growth of inequality in total consumption has continued to slow down beyond the great recession. More importantly, the estimates show that this trend masks a considerable increase in inequality in the consumption of services paired with a decline in the inequality in consumption of goods. Further, the rising inequality in the consumption of services is fuelled by young households.

Understanding the inequality of expenditure on services is important for understanding economic inequality and designing effective public policies that promote social welfare and reduce inequality. Spending on services is a major component of household budgets, and inequality in spending on services can contribute significantly to overall economic inequality. High-income households tend to spend more on services such as healthcare, education, and transportation, which can exacerbate income inequality if low-income households cannot afford these services.

It is important to note that spending on services is critical for human development and social welfare. Access to high-quality healthcare, education, and transportation can improve people's health, education outcomes, and mobility, which in turn can increase their economic opportunities and well-being. Therefore, the inequality of expenditure on services has important policy implications. For example, public policies that aim to reduce income inequality may need to address inequality in access to and utilization of services.

Rising inequality in the consumption of services among young households, calls for a better understanding of what specific services may be contributing to it. To the extent that this may prevent access to essential services for a significant group of the population, it may signal an increase in the precariousness of these households and call for additional policies targeting specific services.

Chapter 2

A canonical model of conspicuous consumption[1]

2.1 Introduction

Relative consumption seems to matter to people (Luttmer, 2005). However, it does not follow that relative concerns must lead to distorted consumption, let alone insufficient saving (Corneo and Jeanne, 1998; Arrow and Dasgupta, 2009). The reason is that consumption choices are inherently intertemporal and so even if current consumption improves one's relative position today, it will also weaken one's relative position in the future. From an empirical standpoint, the identification problem lies in that the influence of group behavior on individual behavior can be mistaken for the aggregation of similar individual behavior (Manski, 1993). Furthermore, standard consumption theory suggests that credible identification of these separate forces is particularly problematic in this case. The complicating

[1]This chapter is co-authored with Ana Ferrer and Francisco M. Gonzalez. We are very grateful to Curtis Eaton for helpful discussions.

factor is that household expenditures are related to the household's *permanent* income, which is unobservable (Friedman, 1957).

It follows that the identification of distortions associated with consumption externalities necessarily relies on a priori assumptions about consumption behavior. Accordingly, an explicit theoretical framework is needed in order to assess the role of concrete identification assumptions. In this chapter, we propose a tractable canonical framework that it is useful for this purpose. Specifically, we propose a model of intertemporal choice that extends standard consumption theory (Modigliani and Brumberg, 1954; Friedman, 1957) to incorporate a role for relative consumption (Veblen, 1899; Duesenberry, 1949).

Our model has three key features. First, for some goods, households care solely about their own consumption, but for others they care about their consumption level relative to those of other households in their reference group. We refer to the latter class of goods as *conspicuous* goods and to the former as *non-conspicuous* goods. Second, an individual's reference group consists of other individuals in the same location/age group/year. Third, our specification of positional concerns is different from standard specifications, which typically posit individual comparisons against some weighted average of consumption by others (e.g., Abel, 1990). Instead, we assume that the utility an individual enjoys from her conspicuous consumption involves the comparison of her consumption of conspicuous goods with that of every other consumer in her reference group. What makes our model tractable is that we assume that each consumer cares about the expected comparison of relative consumptions, which provides a parsimonious characterization of positional concerns.

Under these assumptions, we show that equilibrium consumption behavior is a function of the distribution of conspicuous consumption in an individual's reference group as well as her own permanent income. In turn, the distribution of conspicuous consumption is a function of the distribution of permanent income. The relevant empirical implication is that an individual's consumption, by itself, is no longer a valid proxy for the individual's permanent income if relative consumption matters. The reason is straightforward: two individuals with the same consumption, but different reference groups, must have different permanent incomes if consumption externalities in fact distort consumption behavior.

In practical terms, our model implies that an individual's consumption is a valid proxy of her permanent income, conditional on the distribution of expenditure in conspicuous goods across individuals in her reference group. This insight is the centerpiece of our empirical strategy in Chapter 3. By contrast, standard consumption regressions either fail to control for permanent incomes (Bertrand and Morse, 2016) or assume that consumption is a valid proxy for permanent income even in the absence of controls for the reference group (Charles et al., 2009).

Our analysis contributes to the literature on consumption inequality (Attanasio and Pistaferri, 2016) by calling attention to the effect of visible consumption inequality on consumption behavior. Our canonical model illustrates why one should expect an individual's consumption expenditures to be a function of the distribution of permanent incomes in her reference group as well as her own permanent income. It also illustrates alternative mechanisms through which visible inequality might matter.

First, provided that positional concerns are strong enough, an increase in visible inequality has a positive effect on total household spending if and only if the interest rate is sufficiently low. Intuitively, that is the case in which the positive effect of conspicuous consumption on a household's current social position outweighs the additional negative effect on the household's future social position.[2] Second, an increase in visible inequality has a positive effect on visible spending, and a negative effect on non-visible spending, if and only if positional concerns are sufficiently strong. Equivalently, a necessary and sufficient condition is that the marginal utility of own conspicuous consumption increases at an increasing rate with the conspicuous consumption of others.

The rest of this chapter is organized as follows. Section 2.2 describes the model in detail. Section 2.3 characterizes the unique subgame perfect equilibrium of the model.

[2]This feature of our model is shared by alternative intertemporal models of conspicuous consumption (Corneo and Jeanne, 1998; Arrow and Dasgupta, 2009; Xia, 2010). Alvarez-Cuadrado and Long (2011) and Ray and Robson (2012) are two other interesting dynamic models of relative consumption.

Section 2.4 concludes.

2.2 The model

In this section, we present a simple model in which conspicuous consumption can lead to distortions in current and future consumption behavior. The model has the following important features. First, it focuses on intertemporal choice, so we can understand the mapping between permanent income and consumption behavior. Second, for some goods, households care solely about their own consumption, but for others they care about their consumption level relative to those of other households in their reference group. We refer to the latter class of goods as *conspicuous* goods and to the former as *non-conspicuous* goods. Third, an individual's reference group consists of other individuals in the same location/year/age group.

Formally, there is a continuum of consumers, with mass one, living for two periods. Each consumer faces a standard intertemporal choice problem. The consumer i is endowed with an amount a_i of a single asset at the beginning of period 1, and she also receives incomes y_i and y_i' in period 1 and 2, respectively. Asset holdings from period 1 to period 2 earn an interest rate r, and each consumer can borrow and lend freely, as long as all debt is repaid at the end of the second period.

There are two types of consumption goods with the same price, which is normalized to one. Let $c_i \geq 0$ and $v_i \geq 0$ be individual i's consumption in period 1 of each of the two goods. The stock of assets held by the consumer i at the beginning of period 2 is given by $a_i' = (1+r)(a_i + y_i - c_i - v_i)$, and her second-period consumption is given by $c_i' + v_i' = a_i' + y_i'$. By convention, primed variables denote second-period values.

Accordingly, consumer i's intertemporal budget constraint is given by

$$c_i + \frac{c_i'}{1+r} + v_i + \frac{v_i'}{1+r} = Y_i, \tag{2.1}$$

where Y_i denotes consumer i's permanent income:

$$Y_i \equiv a_i + y_i + \frac{y_i'}{1+r}$$

Let $F_{Y'}$ and F_Y denote the cumulative distribution functions of second-period incomes and permanent incomes, respectively.

Our specification of preferences reflects the fact that the individual i has positional concerns with respect to consumption of the "positional" good, v_i. Accordingly, we refer to v_i as consumer i's *conspicuous* consumption. Consumer i's preferences are defined over consumption in the two periods according to

$$U_i = \alpha u\left(c_i\right) + E\left(u\left(v_i/v_j^{\delta}\right)\right) + \frac{1}{1+\rho}\left[\alpha u\left(c_i'\right) + E\left(u\left(v_i'/v_j'^{\delta}\right)\right)\right], \tag{2.2}$$

with $\alpha > 0$ and $0 < \delta \leq 1$, where $\rho > 0$ is the common rate of time-preference, E denotes the expectation operator, and the expectation in each of the two periods is taken over the distribution of v_j and v_j', respectively.

Our specification of positional concerns is different from standard specifications, which typically posit individual comparisons against some weighted average of conspicuous consumption by others. Here, instead, we assume that utility from conspicuous consumption v_i involves consumer i's comparison of v_i against the consumption v_j of every other consumer $j \neq i$. To make the problem tractable, we also assume that consumer i's positional concerns are additively separable across all $j \neq i$, so consumer i's utility is affected by the distribution of conspicuous consumption only through the average utility associated with positional concerns.

Since $v_i/v_j^{\delta} = v_i^{1-\delta}\left(v_i/v_j\right)^{\delta}$, positional concerns are such that consumer j's conspicuous consumption v_j affects consumer i' utility only through a geometric average of v_i and v_i/v_j, where δ is the geometric weight on v_i/v_j.[3] The parameter δ reflects the strength

[3]In this sense, our specification of positional concerns is as in Abel (2005).

of positional concerns. If $\delta = 0$, the problem reduces to one with two non-positional consumption goods. When $\delta = 1$, consumer i's utility depends on v_i only through v_i/v_j.

We assume that u is an isoelastic function

$$u(z) = \frac{z^{1-\sigma}}{1-\sigma}, \text{ with } \sigma > 1, \tag{2.3}$$

where the elasticities of consumer i's marginal utility with respect to c_i and v_i, respectively, are both constant and equal to $-\sigma$. Similarly, note that $1/\sigma$ measures the elasticity of intertemporal substitution, both with respect to conspicuous and non-conspicuous consumption.

A subgame perfect equilibrium is characterized by a distribution of conspicuous consumption in period 1, F_v^*, and in period 2, $F_{v'}^*$, and a distribution of second-period assets, $F_{a'}^*$, such that (1) for all i, v_i and a_i' are optimal from the viewpoint of the first period, and (2) for all i, v_i' is optimal from the viewpoint of the second period, given a_i' and $F_{a'}^*$.

Properties of the utility function

The main implications of the model are driven by the properties of the individuals' utility function. Thus, we begin with a discussion of those properties. Since preferences are additively separable in conspicuous and non-conspicuous consumption, the precise manner in which a consumer i is concerned about her own position relative to an arbitrary consumer $j \neq i$ depends on the curvature of the function u, as reflected in the value of σ, and the strength of positional concerns, as reflected in the value of δ.

Lemma 1 *The utility function u has the following properties:*

(i) $\frac{\partial u\left(v_i/v_j^\delta\right)}{\partial v_j} < 0$*, for all* $\sigma > 0$*;*

(ii) $\frac{\partial^2 u\left(v_i/v_j^\delta\right)}{\partial v_j^2} < 0$ *if and only if* $\sigma > 1 + 1/\delta$*;*

(iii) $\frac{\partial^2 u\left(v_i/v_j^\delta\right)}{\partial v_i \partial v_j} > 0$ *if and only if* $\sigma > 1$*;*

(iv) $\frac{\partial^3 u\left(v_i/v_j^\delta\right)}{\partial v_i \partial v_j^2} > 0$ *if and only if* $\sigma > 1 + 1/\delta$*.*

Part (i) of the lemma implies that the utility cost that consumer j inflicts on consumer i increases with v_j, for all $\sigma > 0$.

Part (ii) implies that the marginal cost of v_j to consumer i increases with v_j if and only if $\sigma > 1 + 1/\delta$. If $\delta = 1$, this requires that $\sigma > 2$. Even if $\delta = 1/4$, it just requires that $\sigma > 5$. Of course, as δ falls towards zero, positional concerns disappear altogether. Note that the condition $\sigma > 1 + 1/\delta$ is equivalent to $\delta > 1/(\sigma - 1)$. For a fixed value of $\sigma > 1$, this condition is met if and only if positional concerns are sufficiently strong.

Part (iii) implies that consumer i's incentive to engage in conspicuous consumption increases with v_j if and only if $\sigma > 1$. Thus, $\sigma > 1$ seems the most plausible assumption not only with respect to non-conspicuous consumption, but also with respect to conspicuous consumption. Note that as σ approaches 1 one obtains the log-utility case (up to a constant), in which case the interdependence in consumption behavior disappears, even though positional externalities remain present, and so consumers continue to affect each other's utility.

Part (iv) implies that not only consumer i's marginal utility of v_i increases with v_j (Part (iii)), but it does so at an increasing rate if and only if $\sigma > 1 + 1/\delta$, equivalently, if and only if $\delta > 1/(\sigma - 1)$.

Note two implications of the above properties for overall utility. First, for a fixed level of v_i, $E\left(u\left(v_i'/v_j'^\delta\right)\right)$, and thus U_i, will decrease, for all i, with an increase in the distribution of conspicuous consumption in the sense of first-order stochastic dominance, since $u\left(v_i/v_j^\delta\right)$ is a decreasing function of v_j. Thus, everything else equal, individuals dislike it when conspicuous consumption in their reference group grows.

Second, for a fixed level of v_i, $E\left(u\left(v_i'/v_j'^{\delta}\right)\right)$, and thus U_i, will decrease, for all i, with a mean-preserving spread in the distribution of conspicuous consumption if and only if $\sigma > 1 + 1/\delta$, since $u\left(v_i/v_j^{\delta}\right)$ is a concave function of v_j if and only if $\sigma > 1 + 1/\delta$ (Part (ii) of Lemma 1). Accordingly, $\sigma > 1 + 1/\delta$ is a necessary and sufficient condition for a more unequal distribution of conspicuous consumption to have a negative effect on the utility of individuals, everything else equal. As noted above, for a fixed value of $\sigma > 1$, this is the case whenever positional concerns are sufficiently strong.

2.3 Subgame-perfect equilibrium

In the first period, consumer i's problem is given by

$$\max_{v_i, a_i'} \left\{ \frac{\alpha\left(a_i + y_i - v_i - \frac{a_i'}{1+r}\right)^{1-\sigma}}{1-\sigma} + E\left(\frac{v_i^{1-\sigma} v_j^{\delta(\sigma-1)}}{1-\sigma}\right) \right.$$
$$\left. + \frac{1}{1+\rho} \max_{v_i'} \left\{ \alpha \left[\frac{(a_i' + y_i' - v_i')^{1-\sigma}}{1-\sigma}\right] + E\left(\frac{v_i'^{1-\sigma} v_j'^{\delta(\sigma-1)}}{1-\sigma}\right) \right\} \right\}.$$

The second line shows consumer i's problem in the second period. In the Appendix, we characterize the solution to this problem and then proceed to solve the first-period problem. The following proposition characterizes the unique subgame perfect equilibrium of the model.

Proposition 1 *There is a unique subgame-perfect equilibrium. The levels of conspicuous and non-conspicuous consumption in period* 1 *are given by*

$$v_i^* = H(V^*)\, Y_i \;\; \textit{and} \;\; c_i^* = (V^*/\alpha)^{-1/\sigma} H(V^*)\, Y_i,$$

for all i, where the function H and the equilibrium outcome $V^* = E\left((v_j^*)^{\delta(\sigma-1)}\right)$ are char-

acterized in the Appendix as a function of the distribution of permanent incomes, F_Y. The levels of conspicuous and non-conspicuous consumption in period 2 for all i are given by

$$(c_i')^* = \left(\frac{1+r}{1+\rho}\right)^{\frac{1}{\sigma}} c_i^* \;\; and \;\; (v_i')^* = \left(\frac{1+r}{1+\rho}\right)^{\frac{1}{(1-\delta)\sigma+\delta}} v_i^*.$$

The proposition provides a sharp characterization of the equilibrium, despite the fact that the consumption of every individual is a function of the consumption of others in her reference group, which are determined simultaneously. This property of the model is worth noting, since it is well known that the presence of peer effects can create existence problems and otherwise complicate the characterization of equilibrium behavior.

In equilibrium, consumers spend different shares of their permanent income on conspicuous and non-conspicuous consumption. Each of the two shares is a function of the distribution of conspicuous consumption across individuals in the same age group that period. In the Appendix, we show that v_i^* is increasing and c_i^* is decreasing in $V^* = E\left((v_j^*)^{\delta(\sigma-1)}\right)$, and that V^* is uniquely determined as a function of $E\left(Y_i^{\delta(\sigma-1)}\right)$, where recall that $Y_i \equiv a_i + y_i + \frac{y_i'}{1+r}$ is consumer i's permanent income.

Note that the growth rates of conspicuous and non-conspicuous consumption are different, and they are both independent of the behavior of the group. Accordingly, in the present context one cannot detect the presence, or absence, of peer effects by looking at growth rates of either conspicuous or non-conspicuous consumption.

It is instructive to compare the two growth rates.

Corollary 1 $\frac{(c_i')^*}{c_i^*} > \frac{(v_i')^*}{v_i^*}$ *if and only if* $r < \rho$.

This property reflects the interaction between positional externalities and the consumers' willingness to substitute consumption intertemporally. When $r < \rho$, so consumers are relatively impatient, consumption profiles are tilted towards the present, with both $(c_i')^* < c_i^*$ and $(v_i')^* < v_i^*$. This situation aggravates positional concerns in the present period, raising the incentive to frontload conspicuous consumption relative to non-conspicuous consumption. Similarly, if $r > \rho$, consumers have an incentive to postpone consumption, and so $(c_i')^* > c_i^*$ and $(v_i')^* > v_i^*$. This in turn would aggravate positional concerns in the future, raising the incentive to postpone conspicuous consumption relative to non-conspicuous consumption. Finally, note that saving rates are independent of the distribution of permanent incomes when equilibrium levels of conspicuous and non-conspicuous consumption are both constant over time, which is the case if and only if $r = \rho$. Thus, saving rates are not distorted in this case.

The next two propositions characterize the effect of a general increase in wealth and that of an increase in inequality, respectively, in group permanent incomes.

Proposition 2 *Consider a general increase in permanent incomes, defined as an increase in F_Y in the sense of first-order stochastic dominance. For all i, c_i^*/Y_i falls; v_i^*/Y_i and $\frac{v_i^*}{c_i^*+v_i^*}$ rise; $\frac{c_i^*+v_i^*}{Y_i}$ rises if and only if $r < \rho$.*

Recall that we assume that $\sigma > 1$ throughout. Not surprisingly, the effect of changes in the distribution of wealth on consumption behavior depends on the elasticity of marginal utility, $-\sigma$, common to both v_i and c_i. Intuitively, as reference group wealth is higher, consumer i's incentive to increase v_i/c_i rises whenever the elasticity of marginal utility is less than -1 (i.e., whenever $\sigma > 1$), in which case when others increase their conspicuous consumption, the utility cost of not raising one's own conspicuous consumption increases as well. Recall that consumption behavior becomes independent of the distribution of wealth as σ approaches $\sigma = 1$.

However, note that the effect on total consumption also depends on whether $r < \rho$. This condition matters because it determines how positional concerns in the future compare with positional concerns in the present.

Proposition 3 *Consider an increase in the inequality of permanent incomes, defined as a mean-preserving spread of F_Y. For all i:*

(i) If $\sigma \in (1, 1 + 1/\delta)$, c_i^*/Y_i rises; v_i^*/Y_i and $\frac{v_i^*}{c_i^*+v_i^*}$ fall; $\frac{c_i^*+v_i^*}{Y_i}$ rises if and only if $r > \rho$.

(ii) If $\sigma > 1 + 1/\delta$, c_i^*/Y_i falls; v_i^*/Y_i and $\frac{v_i^*}{c_i^*+v_i^*}$ rise; $\frac{c_i^*+v_i^*}{Y_i}$ rises if and only if $r < \rho$.

Recall from Lemma 1 that $\sigma > 1 + 1/\delta$ is a necessary and sufficient condition for a mean-preserving spread in the distribution of conspicuous consumption to decrease utility (part (ii)) and also for consumer i's marginal utility of v_i to increase with v_j at an increasing rate (part (iv)). The latter property of the utility function drives the response of individual conspicuous consumption to increases in the inequality of permanent incomes. Thus, individuals engage in more conspicuous consumption when permanent incomes are more unequal if and only if positional concerns are strong enough — formally, if and only if $\delta > 1/(\sigma - 1)$.

In the Appendix, we prove that the marginal rate of substitution of conspicuous for non-conspicuous consumption is a monotone function of $V^* = E\left((v_j^*)^{\delta(\sigma-1)}\right)$. Accordingly, when permanent incomes become more unequal, then individual non-conspicuous consumption falls exactly when individual conspicuous consumption rises. Once again, the effect on total consumption depends on whether $r < \rho$, which determines how positional concerns in the future compare with positional concerns in the present.

2.4 Discussion and Conclusion

Our analysis illustrates the conditions under which conspicuous consumption may or may not increase when *permanent* income inequality increases. Specifically, in the context of our framework, the following statements are equivalent: (i) individuals engage in more conspicuous consumption when permanent incomes are more unequal, (ii) positional concerns are strong enough, (iii) the marginal cost of others' conspicuous consumption to a consumer increases with others' conspicuous consumption, (iv) a consumer's marginal utility of own conspicuous consumption increases with that of others at an increasing rate.

Our theoretical framework provides an empirical strategy to control for an individual's permanent income when conspicuous consumption does distort individual behavior. The framework is special in several dimensions. For example, there is no risk, preferences are time-consistent and isoelastic, utility is additively separable across consumption goods, and households are heterogenous only with respect to permanent incomes. It is well known that these assumptions have a number of counterfactual implications, and we do not wish to impose them on the data. However, other broad features of our theoretical framework are general. Thus, building on Proposition 1, the empirical strategy we develop in the next chapter rests on the assumption that if conspicuous consumption distorts consumption behavior, then equilibrium expenditures are a function of both the household's permanent income and the distribution of conspicuous consumption within the reference group.

The proof of Proposition 1 shows that the distribution of conspicuous consumption is in fact a function of the distribution of permanent incomes across households in the same reference group. It should be clear that this is not just a feature of our model, but a general feature of models of relative consumption, as long as they also recognize that consumption decisions are inherently intertemporal. However, despite the theoretical insight of this class of models, empirical analyses of conspicuous consumption tend to take as given that individual behavior responds to changes in the *distribution of current incomes* (e.g., Charles et al., 2009; Bertrand and Morse, 2016).

Previous empirical analyses of consumption rely on textbook consumption theory to argue that consumption expenditure is a valid proxy of permanent income (e.g., Aguiar and Bils, 2015). An important insight from our theoretical framework is that this approach is potentially problematic in the context of conspicuous consumption. As illustrated in Proposition 1, changes in a household's permanent income are reflected into changes in the consumption (conspicuous, non-conspicuous and total) of the household only *conditional* *on the distribution of conspicuous consumption of the reference group*. Failing to control for group behavior implies that two individuals with the same consumption, but exposed to different group behavior, will have different permanent incomes if consumption externalities in fact distort consumption behavior. This implies that the comparison of individual behavior across reference groups is particularly problematic, because one cannot make sure the comparison involves individuals with the same permanent incomes. By contrast, the literature on conspicuous consumption has simply assumed that expenditure is a good proxy of permanent income, even in the absence of controls for group behavior, even though such an assumption is flawed precisely when conspicuous consumption matters.

Chapter 3

An empirical analysis of the effect of visible inequality on household expenditures[1]

3.1 Introduction

In this chapter, we use data from the Consumer Expenditure Survey (CEX) and from the Panel Study of Income Dynamics (PSID) to document an effect of consumption inequality on household expenditures in the United States over the period 2010–2018. There is no comparable effect before the Great Recession. We argue that our findings are driven by neither confounding correlated effects nor a failure to account for permanent income. Rather, they are most likely the symptom of actual distortions associated with consumption externalities.

[1]This chapter is co-authored with Ana Ferrer and Francisco M. Gonzalez. We have benefited from comments by Pierre Chaussé, Herb Emery, David Green and Tom Parker.

Our empirical strategy relies on a model of intertemporal choice that extends standard consumption theory (Modigliani and Brumberg, 1954; Friedman, 1957) to incorporate a role for relative consumption (Veblen, 1899; Duesenberry, 1949). We provided an analysis of the model in Chapter 2 and showed that equilibrium consumption behavior is a function of the distribution of conspicuous consumption in an individual's reference group as well as her own permanent income. To implement this insight empirically, one needs a measure of conspicuous consumption, a definition of reference groups and a proxy for the distribution of conspicuous consumption within each reference group. We address these three issues as follows.

First, we construct measures of conspicuous versus non-conspicuous consumption by breaking down total consumption expenditure into more versus less visible categories (Charles et al., 2009). We follow Bertrand and Morse (2016) by using Heffetz's (2011) visibility scores to rank consumption categories, then aggregate them into two broad categories: visible versus non-visible expenditures. Our main empirical measure of conspicuous consumption consists of annual expenditure on clothing, personal care, food away from home, and vehicles.

Second, we define the reference group of an individual in a given year as consisting of all other households who live in the same region that year and who belong to the same broad age group. Ultimately, we exploit variation in the distribution of log visible expenditure for 2 age groups (the head of the household is aged either 25–44 or 45–64) across the 4 U.S. regions (northeast, south, midwest and west) over 35 years (1984–2018).

Third, we approximate the relevant features of the distribution of log spending on visible goods by its first two moments. This approximation is appropriate in the present context because the distribution of visible consumption expenditure in the CEX data is reasonably well approximated as log normal. Battistin, Blundell and Lewbel (2009) show that the cross-sectional distribution of household consumption expenditure in the CEX is approximately log normal. The distribution of visible consumption expenditure is closer to log normal than the distribution of income, although not as close as that of total expenditure.

It is well understood that not only total expenditure and its component expenditures are jointly determined and therefore endogenous, but also measurement error in total expenditure and each of its components are likely to be correlated. To address these problems, we instrument individual expenditures with measures of current income, which is the usual solution in the consumption literature (Aguiar and Bils, 2015). Furthermore, we use lagged measures of the mean and the standard deviation of visible spending in order to exclude the possibility that our estimates are just capturing contemporaneous correlation between measurement errors in the aggregate and individual variables. Our working assumption is that if individual behavior is influenced by group behavior, such a response comes with some delay.

As explained above, our theoretical framework provides an empirical strategy to control for an individual's permanent income when conspicuous consumption does distort individual behavior. However, identifying a causal effect running from group behavior to individual behavior requires that we untangle this *endogenous* effect from confounding *correlated* effects. The latter arise when agents belonging to the same group exhibit similar behavior because they have similar individual characteristics or face similar institutional environments (Manski, 1993). For example, it is plausible that increasing housing prices in some states drive consumption expenditure in those states, or that variation in redistributive policies drives variation in consumption inequality across states and over time. These are non-trivial concerns for analyses of conspicuous consumption that exploit state-level variation in the CEX data (e.g., Charles et al., 2009; Bertrand and Morse, 2016). To address these concerns, our main econometric specifications include the interaction of state and time fixed effects in addition to controlling for the permanent income of the household. This strategy makes it unlikely that the effect of visible inequality on household expenditures is driven by correlated variation in the distribution of unobserved permanent incomes, including correlated variation in relative prices, and in fiscal policies.

Relatedly, individuals may sort into states according to both their generation (i.e., their broad age group) and their preferences in such a way that their visible spending is spuriously correlated with visible inequality. We use fixed effects to deal with these endogeneity concerns. We find that the effect of dispersion in log visible spending on household expenditure over the period 2010–2018 is robust to controlling for the interaction of generation and time fixed effects, and for reference-group-specific time trends, in addition to controlling for the interaction of state and time fixed effects. It is also robust to controlling for average log visible consumption of the reference group as well as individual permanent income.

A further endogeneity concern arises because of the systematic underreporting of consumption in the CEX data among top income households, which is known to be growing over time and so it may give rise to non-classical measurement error (Aguiar and Bils, 2015; Bee et al., 2015; Sabelhaus et al., 2015). We address this concern in several ways. First, we replace the mean and the standard deviation of the distribution of log visible spending within a reference group with the median and the ratio of the 95th to the 5th percentile, which are less dependent on spending at the top. This strategy serves to limit the impact of non-classical measurement error. We do not have a reason to suspect a bias in our estimates of the effect of median visible spending of the reference group. However, one may still expect some bias in our estimates of the effect of visible inequality. We address this potential bias in two ways. First, we instrument the ratio of the 95th to the 5th percentile of log visible spending in the reference group — our measure of *visible inequality* — with the ratio of the 95th to the 5th percentile of the log of income in the same period using data from the Current Population Survey (CPS). Second, we run our main specification using PSID data for the period 2010-2018.

Overall, our analysis contributes to the literature on consumption inequality (Attanasio and Pistaferri, 2016) by calling attention to the effect of visible consumption inequality on consumption behavior. In Chapter 2, we provided a canonical intertemporal model of conspicuous consumption and visible inequality. In this chapter, we provide evidence of its main implications. Our model illustrates why one should expect an individual's consumption expenditures to be a function of the distribution of permanent incomes in her reference group as well as her own permanent income. Our empirical analysis exploits two implications of this result. First, as explained above, consumption, by itself, is not a good proxy of permanent income. Second, controlling for group behavior in consumption regressions requires controlling for the distribution of *permanent*, as opposed to the distribution of *current*, incomes.

We find that both visible and total spending increase, but non-visible spending declines, when visible inequality within one's reference group increases. Through the lens of our model, these effects can be understood in terms of the intertemporal budget constraint: visible inequality increases individual visible consumption at the expense of both current non-visible consumption and current savings.

In Section 3.2, we describe the data and discuss our empirical methodology in detail. We also report estimates of the effect of visible inequality on visible spending over the period 1984–2018. In Section 3.3 we argue that visible inequality has had a significant effect on household spending behavior over the period 2010–2018, but we do not find a comparable effect before the Great Recession. In Section 3.4 we focus on the period after the Great Recession and confirm the robustness of our main findings to alternative specifications. In Section 3.5 we discuss our main findings and conclude. We conjecture that historically low interest rates and the rise of social media are the two main drivers of the effect of visible inequality on household expenditures that we identify. We also highlight the role of household spending on food away from home.

3.2 Data and empirical methodology

Our data comes from the Consumer Expenditure Survey (CEX), a cross-sectional survey conducted by the U.S. Bureau of Labor Statistics. The CEX contains extensive information about consumers' buying habits along with their demographic characteristics. Households participate in the survey for fifteen months. Each household is interviewed five times over five consecutive calendar quarters before it is replaced by a new household unit. Approximately one-fourth of the households contacted each quarter are new to the survey. Households are asked about their expenditures over the preceding three months in the second to fifth interviews, and about their annual income in the last interview.

For our analysis, we use the years 1984 through 2018. Following previous work also using CEX, we include only households who completed all four interviews and reported at least one source of income or benefits (Bertrand and Morse, 2016), which reduces our sample by approximately half of the observations. We keep only households that reported nonzero expenditures in food and those for whom the sum of food expenditure is less than total expenditure. We also exclude households who reported large expenditure shares - greater than one half of the household after tax income in any category other than food or shelter - and focus on urban households only, forgoing about thirty thousand observations (Aguiar and Bils, 2015). In order to avoid consumption choices being distorted by schooling and retirement, we further restrict the sample to households where the reference person is between 25 and 64 years of age, approximately seventy-three percent of the remaining households. Finally, we eliminate households with before-tax income over the 95th and below the 5th percentile, as these are truncated in the data, and those with missing relevant variables. Our final sample contains 76,545 household/year observations. Expenditure and income data are expressed in constant 1983 dollars. We use CEX survey weights through the analysis to make our results representative of the national population.

To identify the effect of visible (consumption) inequality, which is endogenous, we construct measures of the income distribution in each generation-state-year cell to serve as instruments. We follow Bertrand and Morse (2016) and use the large samples of the March Current Population Survey (CPS) to obtain such measures. The household income is the sum of total income for all adult household members, which includes income from business, farm rent, and government transfers, as well as wage income. We then compute the median and the P95/P05 percentile ratio of the household income distribution for each generation-state-year cell, using the household weights provided in the CPS, and merge it with the corresponding CEX households in that group.

We consider the same expenditure categories as Aguiar and Bils (2015), listed in Table C.1 in the appendix. Like them, we measure expenditure on vehicles as a rental equivalence, i.e. computed as rental service flows of vehicle consumption accounting for depreciation. A similar approach is employed for expenditure on shelter. Specifically, annual housing consumption of homeowners is calculated as the sum of mortgage payments, expenses related to repairs, and property taxes.

Table C.2 in the Appendix summarizes the main variables used in the analysis by age group. Average family income after tax is over \$27,000 and total spending \$22,000. Approximately, 28 percent of total spending is in housing and 12 percent goes to visible consumption goods. Older households have slightly higher family income and total expenditure, and devote less resources to visible consumption than younger households. The majority of the household heads are white, male, and have some post-secondary education. Most households have four members or fewer and typically have two earners. Younger households are smaller, more likely to have a college degree, more concentrated in the West and less likely to be white than older households

We confirm our main results using the Panel Study of Income Dynamics (PSID) for the years 1999-2018. The PSID is a longitudinal survey of a representative sample of US individuals and their families, covering consumption expenditure (biannually) since 1999. Despite covering consumption expenditures for a shorter time period than the CEX, the PSID has reputedly better measures of income, and there is no evidence that consumption measures in the panel suffer from non-classical measurement error. The consumption categories considered since 1999 also align with the measures we consider using the CEX.

Our starting point is the idea that households may not only care about their own consumption, but also about their consumption level relative to those of other households in their reference group. We use standard terminology and refer to the latter class of goods as *conspicuous* goods and to the former as *non-conspicuous* goods. We are interested in understanding whether relative concerns cause distortions in individual consumption behavior.

We need to address three key issues: (i) the empirical specification of conspicuous goods, (ii) the empirical specification of the reference group and (iii) the identification of the effect of group consumption on individual consumption. We discuss each of these in turn.

3.2.1 Conspicuous versus non-conspicuous goods

Consumption data from the Consumption Expenditure Survey (CEX) over the period $1984 - 2018$ suggests that the average level and the dispersion of household consumption expenditure correlates very differently with average consumption in different categories of consumption. To illustrate this, we constructed state-year measures of mean log spending in each of fifteen standard categories, as well as measures of the mean and the standard deviation of log total spending in a state and a year. We then regressed mean log spending in each category on the mean and the standard deviation of log total spending in a state and year, controlling for state fixed effects and a linear time trend. Figure 3.1 plots the estimated coefficients on the mean (left panel) and the standard deviation (right panel) for every category as a function of the visibility index constructed by Heffetz (2011). The effects of both the mean and the standard deviation of the distribution of log consumption across households in a state and year on individual household spending tend to be increasing in visibility, i.e., negative for the least visible categories and positive for the most visible categories.

Although the relationships illustrated in Figure 3.1 cannot be interpreted as causal, they raise a number of issues that are likely relevant when using data on consumption expenditure to understand how conspicuous consumption may distort consumption behavior. First, rather than aggregating consumption expenditure across all categories, it may be important to distinguish conceptually between conspicuous and non-conspicuous consumption. With respect to this, we use visible spending, a standard measure of spending on publicly observable goods, to proxy for spending on conspicuous goods. We define visible spending as the sum of annual consumption on clothing, personal care, vehicles and food away from home. Second, interpreting the above relationships as symptoms of conspicuous consumption raises questions about the relevant consumption externalities. If there is in fact an effect of group behavior on individual behavior, what and to whom do individuals respond to? We address these questions next.

3.2.2 Reference group

A natural hypothesis is that individuals only respond to observable consumption among the members of their reference group. However, the definition of the reference group is problematic. Even if one assumes that observationally identical individuals belong to the same reference group, individuals are heterogeneous in multiple dimensions (e.g., age, gender, race) and so different reference groups are overlapping. Furthermore, individuals with identical observable characteristics surely have heterogeneous attitudes towards conspicuous consumption. In this sense, it is useful to take as given some broad characteristic of reference groups as salient. However, a further problem may arise because some of the naturally salient characteristics of reference groups (e.g., geographical location, education, occupation, workplace) are endogenous.

To the extent that consumers value their own conspicuous consumption relative to that of others, age is perhaps the most salient characteristic of any such reference group. Indeed, it is well known that the distribution of consumption varies systematically with age. For example, the dispersion of expenditures in the CEX is increasing over the life cycle.[2] Here, we argue that conspicuous and non-conspicuous consumption patterns also

[2] In principle, this observation can be explained by the textbook permanent income hypothesis (Deaton and Paxson, 1994). However, Aguiar and Hurst (2008) show that the cross-sectional dispersion of entertainment expenditures in the CEX actually declines over the life cycle, which undermines the view that dispersion is driven by idiosyncratic permanent income shocks.

differ by the age of the household. The left panel in Figure 3.2 shows how the mean of log expenditure in visible, non-visible and total consumption changes with age. Visible consumption remains relatively stable for young households and quickly declines after age forty-five or so, while consumption in other goods increases by 0.5 log points until about the same age, remaining fairly constant after that threshold. The right panel in Figure 3.2 further shows that the dispersion of both visible and non-visible consumption, as measured by the standard deviation, increase after age forty, but it is more pronounced in visible consumption, suggesting a greater heterogeneity in this type of consumption for older households.

Accordingly, we focus on broad categories of age (or generations) as the salient characteristic of reference groups and allocate each household to one of two generations, either young or old, based on the age of the head of the household being above or below 45. While one expects that there are many dimensions defining the group that individuals may choose to compare themselves to — by social class, ethnic group, political affiliation, or wealth level — we consider these as subdivisions of a broader categorization along age that is rooted in differences in life cycle stages. A finer categorization may shed further light on conspicuous consumption. However, we find that comparisons within broad age groups are enough to produce significant evidence of conspicuous consumption behavior.

Our choice of the geography of reference groups is driven by practical considerations. We assign each household to one of the four regions identified in the CEX. We make this choice so we have enough observations in a given generation/region/year cell to construct reliable measures of the reference group's consumption distribution. We do, however, include state fixed effects in our regressions to account for changes at the state level that affect household consumption. Note that a narrower specification of reference groups raises problems of selection that tend to exacerbate the underlying identification problems. By contrast, a broader specification of the reference group mitigates selection problems but introduces within-group heterogeneity that may make it difficult to identify social interactions.

Accordingly, our analysis below exploits variation in the distribution of log visible spending for 280 age/region/year cells, where most cells have between 200 and 400 households. The minimum cell size is 41.

3.2.3 Empirical methodology

From an empirical standpoint, there are two fundamental problems. First, the influence of group behavior on individual behavior can be mistaken for the aggregation of similar individual behavior (Manski, 1993). With respect to this, relying on state-level variation over time seems problematic, because there are so many correlated effects (e.g., those associated with variation in local prices or redistributive policies) that may drive a correlation between individual and group behavior. Second, in the case of consumption externalities, the additional problem is that household expenditures should be related to the household's expectations of future income, which is inherently unobservable. In the presence of these two problems, the credible identification of significant distortions associated with consumption externalities necessarily relies on assumptions about consumption behavior. This brings attention to the importance of spelling out the underlying theoretical framework.

Our identification strategy relies explicitly on the theoretical framework we presented in the previous chapter. Specifically, it rests on the assumption that if conspicuous consumption distorts consumption behavior, then equilibrium expenditures are a function of both the household's permanent income and the distribution of conspicuous consumption within the reference group. Accordingly, to study the effect of group behavior on individual spending one would like to estimate the following regression:

$$\begin{aligned}\ln\left(\text{visible spending}\right)_{hgrt} &= \text{distribution of} \ln\left(\text{group visible spending}\right)_{grt} \\ &+ \ln\left(\text{permanent income}\right)_{hgrt} \\ &+ (\text{X})_{hgrt} + \varepsilon_{hgrt}.\end{aligned} \tag{3.1}$$

where h indexes households, g indexes generations, r indexes regions and t indexes time.

To account for systematic differences in consumption levels across different types of households, we control for the vector X_{hgrt} of household socio-demographic characteristics. These include household head's gender, race and education, a quadratic in household head's age, and indicator variables for the number of earners and household's size. The variable of interest is the distribution of $\ln\left(\text{group visible spending}\right)_{grt}$, a vector of moments of the distribution of log visible spending of the household's reference group.

How should one control for the relevant features of group behavior? As illustrated in the previous chapter, one may need to control for all the moments of the distribution. Assuming that individual consumption is a function of some *ad hoc* moment, or percentile, of the distribution of consumption spending may lead to non-trivial misspecification. A natural strategy is to use the first two moments of the distribution of visible spending in a household reference group. We think this strategy is appropriate in the present case because the cross-sectional distribution of visible spending is approximately log normal. However, Aguiar and Bils (2015) show some systematic under-reporting of consumption in the CE among top income households, which is, furthermore, growing over time. Thus, it is likely that the level and growth of both the mean and the standard deviation of log consumption are biased downwards. In order to address this issue, we control for the median and the ratio of the 95th percentile to the 5th percentile of the distribution of log visible spending, rather than controlling directly for the first two moments of the distribution of visible spending in the reference group. Further, to attenuate the effect of measurement error in consumption, we use lagged values of these two moments. This allows us to exclude the possibility that our estimates are just capturing contemporaneous correlation between measurement errors in the aggregate and individual variables. Intuitively, it seems plausible that if household consumption choices are influenced by consumption of the reference group, such response would likely come with some delay.

There are further issues with the identification of group effects. It is well known that unobservable determinants of household consumption may be correlated across households within the same location and age group, and over time because individuals in the group face a similar institutional environment. For example, it could be that households are systematically exposed to relatively higher local prices, or more redistributive policies, in some locations or time periods. This would imply that correlations among ε_{hgrt} are the reason why households tend to consume more when the distribution of visible consumption within their reference group has higher values of the mean or the variance. The presence of these correlated effects complicates the identification of endogenous social effects. In general, economic models of social interactions do not provide a logical basis for choosing instruments (Brock and Durlauf, 2001). Instead, we rely on fixed effects to achieve identification. In particular, we include year, region and generation fixed effects to account for time, location and generation-specific factors that could influence the reference group's distribution of expenditure. We also include the interaction of region and year fixed effects to address regional variation over time (such as time variation in local prices and redistributive politics at the level of the region). Finally, we include the interaction of generation and region fixed effects, generation and year fixed effects, as well as the interaction between generation-region groups and a linear time trend to account for systematic variation in preferences across generations, both across regions and over time.

Accordingly, we begin by estimating the following regression:

$$\begin{aligned}
\ln(\text{visible spending})_{hgrt} = {} & \text{p50 of} \ln(\text{group visible spending})_{grt-1} \\
& + \text{p95/p05 of} \ln(\text{group visible spending})_{grt-1} \\
& + \ln(\text{total spending})_{hgrt} \\
& + (\text{X})_{hgrt} \\
& + \text{generation}_g + \text{region}_r + \left(\text{generation}_g \times \text{region}_r\right) \\
& + \text{year}_t + (\text{region}_r \times \text{year}_t) + \left(\text{generation}_g \times \text{year}_t\right) \\
& + \left(\text{generation}_g \times \text{region}_r \times \text{trend}_t\right) + u_{hgrt}
\end{aligned} \tag{3.2}$$

and, like others, address concerns about measurement error in an individual's total spending by instrumenting this variable with dummy variables indicating the household's income category and a continuous variable accounting for the log real after-tax income (Aguiar and Bils, 2015).

In Table 3.1 we report estimates of the effect of the distribution of consumption of the reference group on individual household consumption under different specifications of equation (3.2). All regressions in Table 3.1 include controls for sociodemographic variables (gender, race, age and age squared and indicators for education, the number of individuals and the number of earners in the household).

Figure 3.1: Correlation between state-year consumption in 15 categories and total consumption.

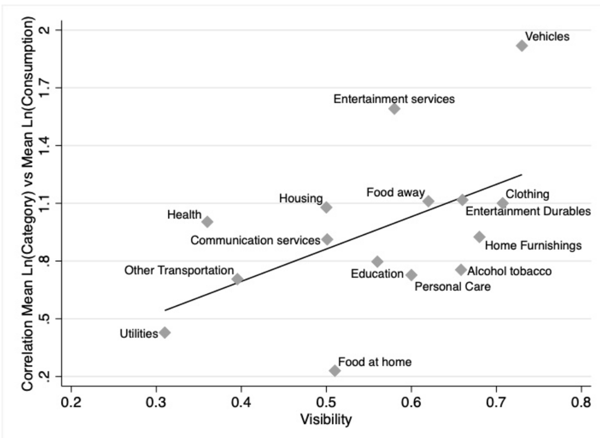

(a) Mean

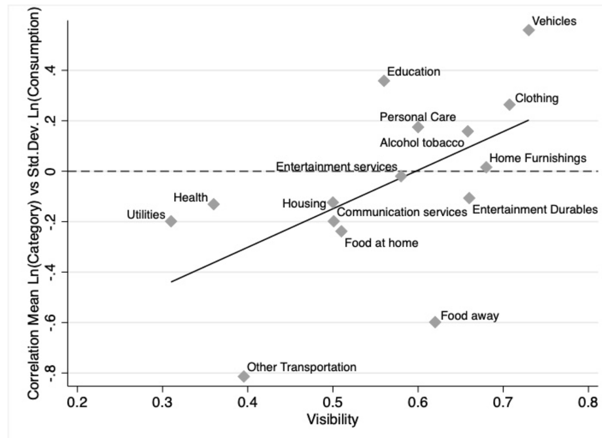

(b) Standard Deviation

Notes: Correlation between mean log spending in 15 consumption categories and the mean (left) and standard deviation (right) of log total consumption spending against visibility index. 15 categories are constructed to match the visibility index in Heffetz (2011). To that end, we aggregate adult clothing, children clothing, and shoes and apparel into a clothing category. We also sum expenditure on alcoholic beverages, tobacco, and other smoking to form one category. Appliances, phone, computers are depicted under communication category on the graph. Lastly, we do not plot domestic services and cash contributions due to the absence of direct link to the visibility index reported in Heffetz (2011). Data source: CEX waves, 1984–2018; authors calculations.

Figure 3.2: Age estimates for visible vs. non-visible spending

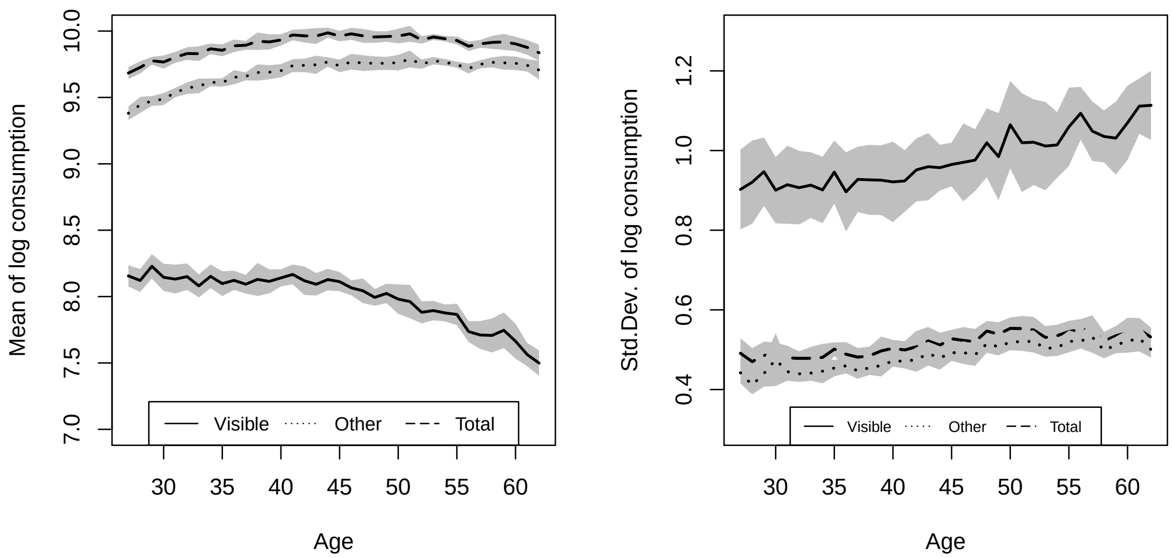

Notes: Source: authors' calculations.

Table 3.1: Impact of visible inequality on visible spending: identification

	(1)	(2)	(3)	(4)	(5)	(6)
p50 of ln(visible spending)	0.021	0.013	0.47	0.016		0.072
	(0.012)	(0.013)	(0.025)	(0.017)		(0.040)
p95/p05 of ln(visible spending)	0.136**	0.126**	0.124***	−0.007		
	(0.045)	(0.037)	(0.035)	(0.064)		
ln(housing spending)			−1.637***			
			(0.032)			
p50 of ln(CPS income)					−0.120	
					(0.123)	
p95/p05 of ln(CPS income)					−0.358	
					(0.248)	
IV p95/p05 of ln(CPS income)						0.426*
						(0.201)
household controls	Yes	Yes	Yes	Yes	Yes	Yes
generation FE × region FE	Yes	Yes	Yes	Yes	Yes	Yes
generation FE × year FE	Yes	Yes	Yes	Yes	Yes	Yes
generation FE × region FE × time trend	Yes	Yes	Yes	Yes	Yes	Yes
generation, region and year FEs	Yes					
region FE × year FE	Yes					
generation, state and year FEs		Yes	Yes	Yes	Yes	Yes
state FE × year FE		Yes	Yes	Yes	Yes	Yes
observations	69,091	64,245	64,245	64,245	64,245	64,245
R^2	0.621	0.624	0.579	0.271	0.624	0.624

Dependent variable, in all columns, is the natural log of visible spending, defined as the sum of personal care, clothing, vehicle spending and food away from home. All regressions include socio-demographic variables (gender, race, age and age squared and indicators for education, the number of individuals and the number of earners in the household). Standard errors clustered at the generation-region level. Sample years 1984-2018.
Columns (1) to (4) control for household total expenditure (instrumented with indicators for the household's income group and the log of after-tax income). Column (3) includes a control for shelter expenditure separately, instrumented with the mean level of housing prices in the individual's state of residence.
Column (4) shows estimates from an OLS model, where controls for current income (indicators for the household's income group and the log of after-tax income) replace instrumented total expenditure.
Column (5) uses CPS income data to construct measures of the distribution of current income in the reference group, which replace measures of the distribution of ln(visible spending).
Column (6) replicates column (2) instrumenting the p95/p05 of ln(visible spending) with the p95/p05 of ln(CPS income).
Robust standard errors in parentheses.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

From the perspective of our model, total spending is a valid proxy of a household's permanent income, conditional on the distribution of visible consumption within the reference group of the household. Accordingly, in columns (1)-(3) we control for permanent income, using total spending as a proxy, and instrumenting it with measures of current income (indicators for the household's income group and the log of after-tax income). More importantly, we also include controls for the median and the ratio of the 95^{th} percentile to the 5^{th} percentile (p95/p05) of the distribution of log visible consumption. Once we control for age of the households and for age-group fixed effects, only the p95/p05 of log visible spending has an effect on household's visible spending.

Our specification in column (1) includes all fixed effects considered in equation (8). Although this should take care of much of the spurious correlation that could affect our estimates of the distribution of group visible spending, there may still be some factor correlated with the distribution of visible spending within reference groups that mechanically affects household spending in visible goods. Specifically, state variation in relative prices or redistributive taxation might be correlated with the dispersion of permanent incomes. To account for potentially confounding sources of state level variation over time, we replace region and region-year interactions with state and state-years interaction fixed effects in column (2) and subsequent regressions in the analysis. The effect of visible inequality on the household's visible spending remains strong and significant after including the full battery of fixed effects, in addition to controlling for the average visible spending of the household's reference group. While state-year fixed effects will absorb much of the variation due to unobserved variables affecting spending at the state level over time, it remains the possibility that some of these variables have a differentiated effect across generation-groups within states. For example, when housing prices are high, conditional on permanent income, older households may spend more on comparable amounts of housing, and less on other goods, than younger households. States with high housing prices and an older distribution of the population may drive the relationship between inequality in group visible consumption and household visible consumption. To see if this is the case, in column (3), we control directly for the household's log of housing expenditures. Since housing spending is endogenous with respect to the household's spending decisions on other consumption categories, we instrument individual housing expenditures with the mean value of house prices in the household's state of residence.[3] Note that, holding the distribution of visible spending in the reference group constant, an increase in individual housing expenditures reduces visible spending. However, comparing the estimates for the distribution in group spending in columns (2) and (3), we find strikingly similar results, making it unlikely that housing prices are driving our results through the differential effect on generation-groups.

We have considered the possibility that the distribution of the reference group's permanent income is not exogenously assigned in our regressions. There could be systematic sorting of individuals across states in a way that gives rise to additional correlated effects that are unrelated to conspicuous consumption. For instance, individuals with high permanent incomes may choose to spend more on visible goods, and they may also choose to locate in certain states but not in others. While there is no clear strategy to deal with this potential source of endogeneity, we think it unlikely that it is a major driver of our results. It is not obvious to think of drivers of such sorting after conditioning on the permanent income of the household and the extensive array of fixed effects already included.

In column (4) of Table 3.1 we report estimates from a naive OLS regression where the log of household visible consumption is regressed against the distribution of the reference group visible consumption as before, but without controlling for permanent income, simply using measures of current income instead. The coefficients for the distribution of the reference group visible consumption show no significant effect on household's visible consumption in this case, suggesting that the naive OLS estimates are biased, masking the effect of consumption inequality of the reference group on individual consumption. Further, we also show that, once we control for permanent income, the distribution of visible spending by

[3]We use the Federal Housing Finance Agency (HFA) House Price Index, a broad measure of changes in single-family home values at state level and merge it with the CEX data. https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx.

the reference group has no effect on individual's consumption if the group effect is measured using the distribution of *income*, rather than consumption, as it is done in previous work on conspicuous consumption (e.g., Charles et al., 2009; Bertrand and Morse, 2016). These estimates are reported in column (5), which uses data from the CPS to construct group measures for the distribution of income. The coefficients of median and P95/P05 group income became not significant when the household's permanent income is added to the regression.

One remaining concern regarding the estimates of column (2) arises because of the evidence of non-classical measurement error in the CEX data set associated with underreporting at the top (Aguiar and Bils, 2015). This is unlikely to bias estimates of the effect of median visible spending of the reference group. However, we might expect a bias in the estimates of the effect of visible inequality. To provide robust estimates of the dispersion of the distribution of visible spending of the reference group, we instrument this measure with the P95/P05 of the distribution of log income of the reference group using income data from the CPS. It is well known that the CPS has better measures of income at the top of the distribution and are less likely to be affected by measurement error, and so they are preferred to construct the instrumental variable. The IV estimates in column (6) continues to show a positive (and rather large) relationship between group consumption inequality and conspicuous consumption, suggesting that estimates in column (2) may be biased downwards, and so they are a conservative estimate of the effect of visible inequality on the household's visible spending.

3.3 Impact of visible inequality over time

In this section, we argue that visible inequality has had a significant effect on household spending behavior over the period 2010–2018, but not before the Great Recession. Table 3.2 reports estimates of the effect of the group visible spending interacted with different time periods using the specification of column (2) in Table 3.1. We exclude the period

2007-2009 because the Great Recession had a disproportionate impact on the consumption of many households. The estimates reveal that when controlling for the full set of fixed effects, increasing visible inequality has a positive and significant effect on the households' visible spending, but only after the Great Recession.

Table 3.2: Impact of visible inequality on visible spending over time

p50 of ln(visible spending) × (1984-1999)	-0.03
	(0.029)
p50 of ln(visible spending) × (2000-2006)	-0.187^{**}
	(0.058)
p50 of ln(visible spending) × (2010-2018)	0.151^{***}
	(0.039)
p95/p05 of ln(visible spending) × (1984-1999)	0.002
	(0.077)
p95/p05 of ln(visible spending) × (2000-2006)	0.06
	(0.039)
p95/p05 of ln(visible spending) × (2010-2018)	0.270^{**}
	(0.097)
observations	57,184
R^2	0.630

Note: Dependent variable is the natural log of visible spending, defined as the sum of personal care, clothing, vehicle spending and food away from home. Regression includes same fixed effects and sociodemographic controls as in column(2) in Table 1. Total household expenditure instrumented with indicators for the household's income group and the log of after-tax income. Standard errors clustered at the generation-region level.
Years 2007-2009 removed from the sample.
Robust standard errors in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A natural question to consider is whether our estimates actually reflect the effect of the reference group on individual's visible spending. In principle, other unobserved variables could drive the sorting of households across states, resulting in the observed effect. It is not obvious what such driver could be that is not picked up by our extensive list of fixed effects and the inclusion of household permanent income in the regression. However, one could consider that within-generation variation in permanent-income inequality could be correlated with the prices of consumption goods more generally, both visible and/or non-visible. If this was the case, one may expect individual visible spending to respond to non-visible inequality as well. We explore this possibility in Table 3.3, which shows the impact of the reference group *non-visible* spending on an individual's visible spending over time. Contrary to what a price-driven sorting mechanism would imply, we find no effect of non-visible inequality on the households' visible spending, not even after the Great Depression.

Through the lens of our theoretical framework, a response of visible spending to rising visible inequality must come at the expense of either current consumption of non-visible consumption goods, or future consumption. We proceed to consider the opportunity cost of increases in visible inequality by looking at its effects on the household's non-visible consumption and total expenditure in Table 3.4. Our results indicate that both of those effects are also present over the period 2010-2018, but not before.

Column (1) in Table 3.4 shows that between 2010 and 2018, households reduced non-visible spending in response to increases in visible inequality. When turning to household's total spending, we face the problem of how to control for permanent income. Here, again, we rely on our theoretical framework, which implies that non-visible spending is also a good proxy for permanent income, once one controls for the distribution of visible consumption in the reference group. This allows us to estimate a similar regression for total spending, where permanent income is proxied instead by non-visible consumption, and similarly instrumented with indicators for income group and the log of after tax income. Coefficients are shown in column (2) of the table, revealing that increasing visible inequality increased household total spending after the Great Recession.

Table 3.3: Impact of non-visible inequality on visible spending over time

p50 of ln(non-visible spending) × (1984-1999)	0.124
	(0.077)
p50 of ln(non-visible spending) × (2000-2006)	0.206**
	(0.081)
p50 of ln(non-visible spending) × (2010-2018)	0.338***
	(0.065)
p95/p05 of ln(non-visible spending) × (1984-1999)	−0.213
	(0.334)
p95/p05 of ln(non-visible spending) × (2000-2006)	0.040
	(0.374)
p95/p05 of ln(non-visible spending) × (2010-2018)	−0.007
	(0.246)
observations	57,184
R^2	0.630

Note: Dependent variable is the natural log of visible spending, defined as the sum of personal care, clothing, vehicle spending and food away from home.
Regression includes the same fixed effects and sociodemographic controls as in column(2) in Table 1. Total household expenditure instrumented with indicators for the household's income group and the log of after-tax income. Standard errors clustered at the generation-region level.
Years 2007-2009 removed from the sample.
Robust standard errors in parentheses.
** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3.4: Opportunity cost of visible spending over time

Dependent variable	(1) ln(non-visible spending)	(2) ln(total spending)
p50 of ln(visible spending) × (1984-1999)	−0.007	0.008
	(0.019)	(0.022)
p50 of ln(visible spending) × (2000-2006)	0.110***	−0.122***
	(0.016)	(0.017)
p50 of ln(visible spending) × (2010-2018)	−0.012	0.013
	(0.023)	(0.026)
p95/p05 of ln(visible spending) × (1984-1999)	−0.037	0.041
	(0.029)	(0.033)
p95/p05 of ln(visible spending) × (2000-2006)	−0.035	0.039
	(0.029)	(0.032)
p95/p05 of ln(visible spending) × (2010-2018)	−0.111***	0.124***
	(0.014)	(0.016)
observations	57,184	57,184
R^2	0.786	0.715

Note: Dependent variable in column (1) is the natural log of non-visible spending, defined as total spending, minus visible spending, minus shelter. Household permanent income in column (1) is proxied with total household spending, instrumented with indicators for the household's income group and the log of after-tax income.
The dependent variable in column (2) is the natural log of total spending. This regression replaces a control for household total expenditure with a control for non-visible spending, instrumented with indicators for the household's income group and the log of after-tax income.
Both regressions include the same fixed effects and sociodemographic controls as in column (2) in Table 1. Standard errors clustered at the generation-region level. Years 2007-2009 removed from the sample.
Robust standard errors in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.4 Conspicuous spending after the Great Recession

In this section we turn our attention to the period after the Great Recession, in which we observe an effect of visible inequality on visible spending, and confirm the robustness of this result to alternative specifications.

Column (1) in Table 3.5 reports our main specification, as in column (2) in Table 3.1. Our estimates indicate that a 10 percent increase in median group visible spending increased household visible spending by 1.9 percent. More interestingly, an increase in the dispersion of group visible spending equivalent to 0.1 increase in the ratio P95/P05, increases individual visible consumption by 4 percent. This is in sharp contrast with results from a specification that uses the distribution of current income to capture group effects, as is done in previous work on conspicuous consumption. These estimates, reported in column (2), are large, but imprecise, and show a *negative* effect of group inequality on household visible consumption. In column (3) we instrument our measure of visible inequality with the p95/p05 of log income, using the CPS. The effect of visible inequality becomes even larger - a similar increase in group visible spending inequality will result in 8 percent in household visible consumption - and precisely estimated. Finally, in column (4) we replicate the regression in column (1) using PSID data, where concerns about non-classical measurement error have not been reported. We find an effect of visible inequality that is similar in magnitude to the effect reported in column (1).

Next, we consider the effect of visible inequality on non-visible and total spending. In Table 3.6 we report the corresponding estimates. Increasing visible inequality reduces spending in non-visible consumption and increases total spending (see column (1)). Note that measuring group effects using CPS income has no significant effect on either type of spending once all fixed effects are included in the regression (column (2)). Instrumenting visible (consumption) inequality using group CPS income inequality produces estimates that have the expected sign but are not precisely estimated (column 3). We turn to the PSID to check once more that the results are robust to considerations of under-reporting

Table 3.5: Impact of visible inequality on visible spending, 2010–2018

	(1)	(2)	(3)	(4)
p50 of ln(visible spending)	0.190***		0.306***	−0.214
	(0.038)		(0.072)	(0.166)
p95/p05 of ln(visible spending)	0.420***			0.404*
	(0.067)			(0.194)
p50 of ln(CPS income)		0.330		
		(0.299)		
p95/p05 of ln(CPS income)		−0.932		
		(0.496)		
IV p95/p05 of ln(CPS income)			0.805***	
			(0.182)	
observations	16,150	16,150	16,150	29,730
R^2	0.588	0.588	0.588	0.614

Note: Dependent variable in all columns is the natural log of visible spending, defined as the sum of personal care, clothing, vehicle spending and food away from home.
All columns include the same fixed effects and sociodemographic controls as in column (2) in Table 1. Total household expenditure instrumented with indicators for the household's income group and the log of after-tax income. Standard errors clustered at the generation-region level. Sample years 2010-2018
Column (2) uses CPS income data to construct measures of the distribution of current income in the reference group, which replace measures of the distribution of ln(visible spending).
Column (3) instruments p95/p05 of ln(visible spending) with the p95/p05 of Ln(CPS income).
Column (4) replicates column (1) using PSID data.
Sample years 2010-2018.
Robust standard errors in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

at the top. The results are in line with those reported in column (1) and are significant at 10 percent confidence level.

To get a sense of the potential economic significance of our estimates, consider a simple back of the envelope calculation using the results in Panel B of Table 3.6. The average p95/p05 ratio of Ln(visible spending), across reference groups and over time, is 1.51 in our sample. Increasing this ratio by one-tenth of a unit (an increase in inequality of 6.6%) will result in an increase in predicted total spending between $314 and $431, depending on whether one uses the estimates from column (1) or column (4) in Table 6. This is approximately between one third and 40% of what the median household spends on average on food away from home.

3.5 Discussion and conclusion

While the Great Recession may have had some structural effects on conspicuous consumption, our conjecture is that two main sources underlie the effect of visible inequality on household expenditures that we identify: historically low interest rates and the rise of social media. The unprecedented drop of interest rates in the United States during the Great Recession is widely known. The Federal funds effective rate felt below 1 percent in October 2008, and it remained below 1 percent until June 2017.

The model in Chapter 2 illustrates the role of interest rates in conspicuous consumption. Households have an incentive to rise or lower their conspicuous consumption when visible inequality increase depending on the interaction between interest rates and the households' time preferences. Intuitively, increasing one's conspicuous consumption improves one's relative position currently, but it will also weaken one's relative position in the future. The historically low interest rates after the Great Recession created the ideal circumstances under which households have an incentive to respond by increasing their conspicuous consumption at the expense of future consumption when facing an increase in

Table 3.6: Opportunity cost of visible spending, 2010–2018

	(1)	(2)	(3)	(4)
	Panel A. ln(non-visible spending)			
p50 of ln(visible spending)	−0.021		0.001	−0.091
	(0.024)		(0.056)	(0.133)
p95/p05 of ln(visible spending)	−0.130***			−0.219*
	(0.020)			(0.104)
p50 of ln(CPS income)		−0.057		
		(0.053)		
p95/p05 of ln(CPS income)		0.137		
		(0.402)		
IV p95/p05 of ln(CPS income)			−0.56	
			(0.211)	
observations	16,150	16,150	16,150	29,730
R^2	0.804	0.804	0.804	0.808
	Panel B. ln(total spending)			
p50 of ln(visible spending)	0.022		−0.001	0.078
	(0.026)		(0.059)	(0.117)
p95/p05 of ln(visible spending)	0.138***			0.192*
	(0.021)			(0.090)
p50 of ln(CPS income)		0.060		
		(0.057)		
p95/p05 of ln(CPS income)		−0.147		
		(0.425)		
IV p95/p05 of ln(CPS income)			0.062	
			(0.223)	
observations	16,150	16,150	16,150	29,730
R^2	0.750	0.750	0.750	0.771

Note: Columns include same fixed effects and sociodemographic controls as in column(2) in Table 1. Standard errors clustered at the generation-region level. Sample years 2010-2018.
In Panel A, the dependent variable is the natural log of non-visible spending, defined as the sum of total spending, minus visible spending, minus shelter. Total household expenditure is instrumented with indicators for the household's income group and the log of after-tax income.
In Panel B, the dependent variable is the natural log of total spending. This regression replaces a control for household total expenditure with a control for household non-visible spending, instrumented with indicators for the household's income group and the log of after-tax income.
Column (2) uses CPS income data to construct measures of the distribution of current income in the reference group, which replace measures of the distribution of ln(visible spending).
Column (3) replicates column (1) instrumenting the p95/p05 of ln(visible spending) with the p95/p05 of ln(CPS income).
Column (4) replicates column (1) using PSID data.
Robust standard errors in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

visible inequality.

Moreover, the encroaching of social media platforms in daily lives since 2007 has likely had far-reaching effects on household spending. In particular, the rise in social media has changed the visibility of expenditures on some consumption items, notably food away from home (we return to this point below). As of December 2021, there were 3.4 billion social media active users worldwide, a number that has been increasing steadily since 2007, with Instagram's 2 billion users as the fourth most popular social media site in the world.[4] According to Schwab's 2019 Modern Wealth Survey, more than a third of Americans admit their spending habits have been influenced by images and experiences shared by their friends on social media and confess they spend more than they can afford for fear of missing out (or FOMO).[5]

It is unlikely that social media has had any significant effect on conspicuous consumption before 2007. For example, Facebook membership did not become available to everyone with a valid email address until September 2006. Video sharing in YouTube was not widespread until after it was acquired by Google in November 2006. Twitter was not launched until July 2006 and Instagram not until October 2010. We emphasize the effect of inequality on household consumption after 2010, rather than 2007, because the 2008–2009 financial crisis had a disproportionate effect on consumption expenditures.

Others have noted that household spending on cars accounts for a large share of visible spending (Charles et al., 2009; Bertrand and Morse, 2016).[6] Not surprisingly, spending on cars does play a role in our results. More interestingly, however, we conjecture that spending on food away from home seems to be a likely driver of our results.

Using Heffetz's visibility index (Heffetz, 2011), food away from home (FAFH) already ranked in the top tiers of visibility within 31 categories of consumption, based on results

[4] https://www.statista.com/statistics/253577/number-of-monthly-active-instagram-users/

[5] https://www.businesswire.com/news/home/20190513005203/en/

[6] Previous work has reported strong social effects with respect to car purchases. Grinblatt, Keloharju and Ikaheimo (2008) argue that the timing of car purchases exhibits significant social effects. Kuhn et al. (2011) argue that neighbors of lottery winners are more likely to buy a car.

from a survey conducted in 2004. Since then, consumption of FAFH has steadily risen to become a substantial part of American diet. According to the United States Department of Agriculture (2018), its share of the food budget in 2010 was 50 percent (up from 41 percent in 1984), surpassing the share for food at home (FAH) for the first time. Examples of media coverage supporting the view that social media has fueled expenditures in food are easy to find. The hashtags *"food", "foodporn", "instafood"* and *"yummy"* are amongst the 100 most popular *direct* Instagram hashtags and are attached to over four hundred million photos on the platform, making food the most photographed subject on the network (Social Media today, November 9, 2016).[7] According to TVO, an agency of the Ontario Ministry of Education and a not-for-profit, social impact charity:

> "We have come to believe food is edible status, and photographs allow us to share that status with others. And if food is edible status, then photographing food is about conspicuous consumption. It's not enough to simply eat something delicious: ultimately, what we really want is the high that comes from other people knowing we've eaten something delicious."
> [TVO, December 2016]

Figure 3.3 plots the median of the distribution of individual shares of spending on each visible category as a fraction of visible consumption (Panel A) and the median of the distribution of log spending on each visible category (Panel B). Note that panels A and B of Figure 3.3 are not directly comparable because the individuals with the median share on a given category are not the same as those with the median log spending on that category. It is nevertheless instructive to consider both of them. The CE data shows that the median share of FAFH spending increased substantially after 2007, rising from slightly over 20 percent to around 40 for the median spender (see Panel A). Panel B shows that FAFH is the only category of visible consumption whose median spending did

[7]A. Atanasova: "The Psychology of Foodstagramming", Published Nov. 9, 2016: https://www.socialmediatoday.com/social-networks/psychology-foodstagramming.

not fall significantly during the Great Recession and its aftermath. Interestingly, median log spending on FAFH surpassed median log spending on vehicles in 2007 and remained relatively higher until 2018.[8]

It is also interesting to note that similar trends are observed throughout the distribution of FAFH spending. At the bottom decile, FAFH consumption went from 5.15 log points in 2006 to 5.68 in 2018, an increase of 10 percent. During the same period, log FAFH rose 2.6 percent at the top 5^{th} percentile of the distribution of FAFH expenses (Figure 4). Hence, from the point of view of inequality, spending in FAFH actually contributed to moderate the rising inequality in visible consumption spending, driven by increasing inequality in vehicle spending.[9]

We consider the potential role of FAFH in our results by using an alternative definition of visible group consumption that *e*xcludes spending in FAFH. If our results are driven by FAFH spending, we expect to see significantly weaker effects associated with this alternative definition of visible spending. Results of this exercise are shown in Table 3.7 for visible, non-visible and total spending over the period 2010–2018. Using CE data (Panel A) and PSID data (Panel B) produces similar results. Removing FAFH spending from the definition of visible spending significantly reduces the positive effect of visible inequality on household visible spending (column 1). Furthermore, there is no longer an effect of visible inequality on either non-visible or total spending (columns 2 and 3, respectively). These findings suggest that spending on FAFH has indeed played a role on conspicuous consumption after 2010.

[8] Attanasio et al. (2022) highlight the fact that American households that purchased a car during the Great Recession spent less than those that purchased a car prior to the crisis.

[9] Spending in vehicles declined by 7 percent at the bottom decile and by 1.5 percent and the top 5^{th} percentile of the distribution of vehicle spending

Figure 3.3: Visible spending by category.

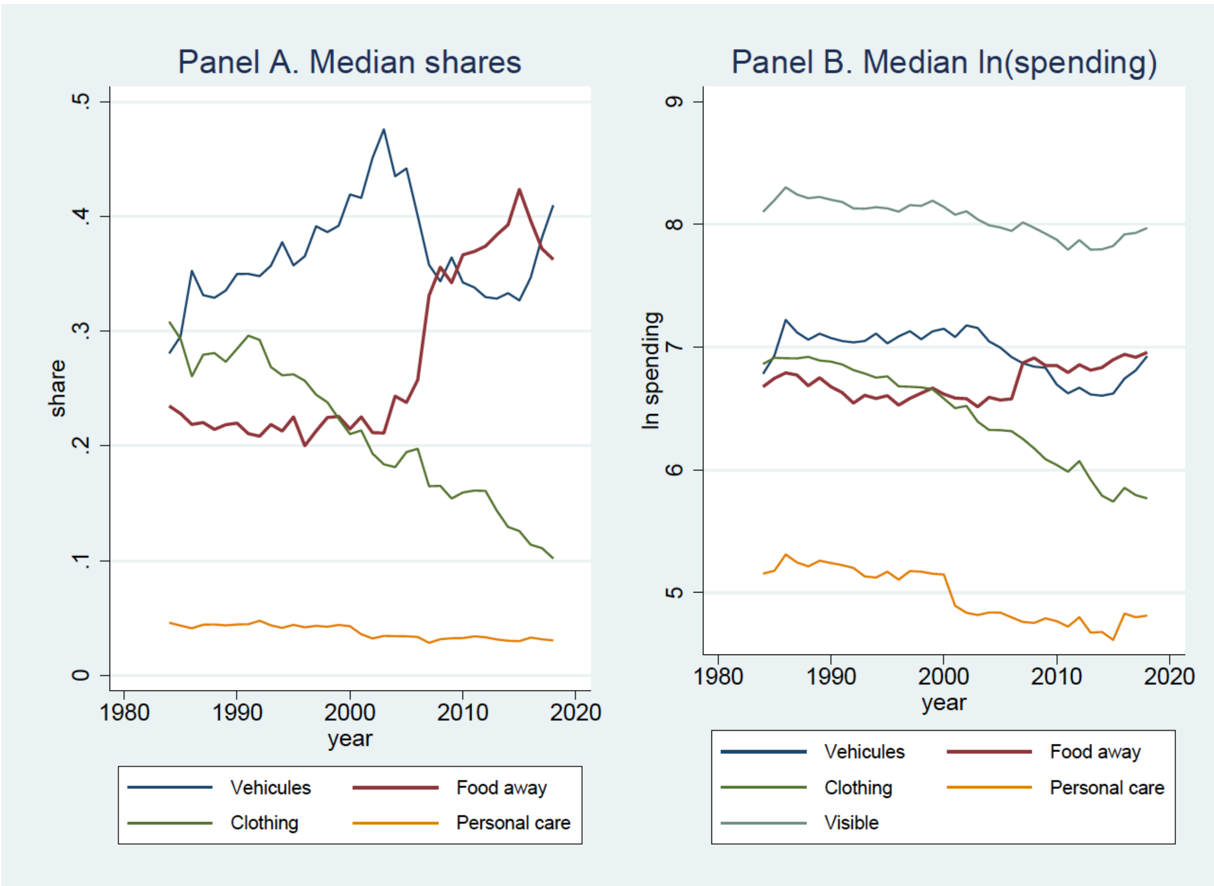

Data source: CE waves, 1984–2018; authors calculations.

Figure 3.4: Distribution of ln(FAFH spending).

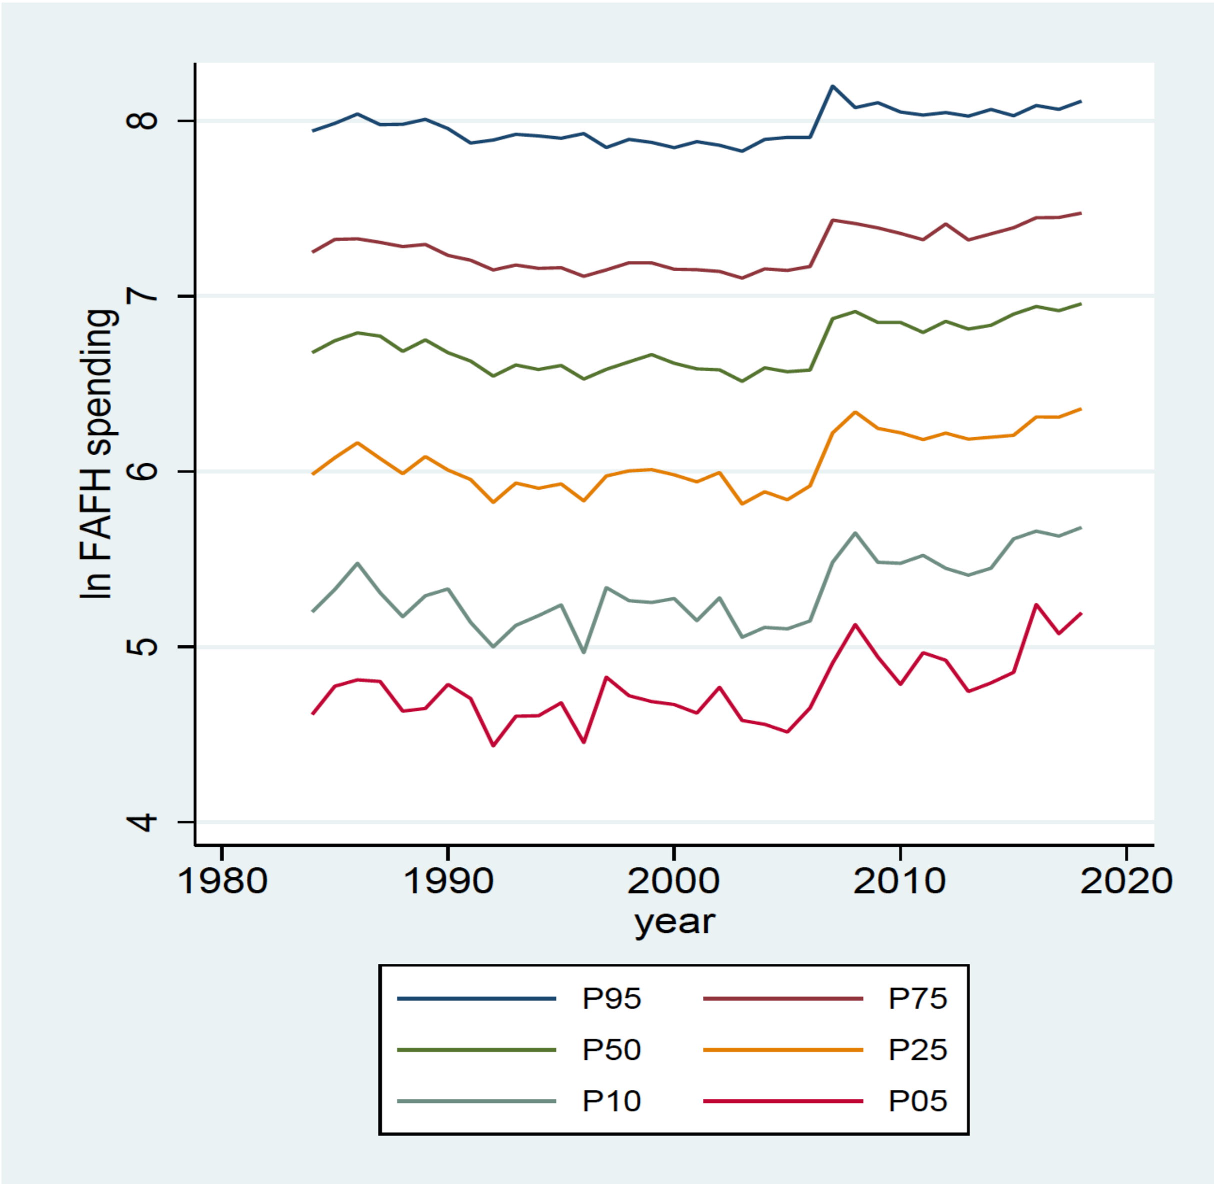

Notes: Variables are logs. Data source: CE waves, 1984–2018; authors calculations.

Table 3.7: The role of FAFH expenditures, 2010–2018

Dependent variable	(1) ln(visible spending minus FAFH spending)	(2) ln(non-visible spending plus FAFH spending)	(3) ln(total spending)
	Panel A. CE data, 2010–2018		
p50 of ln(visible spending minus FAFH spending)	0.057	0.024	−0.009
	(0.049)	(0.019)	(0.021)
p95/p05 of ln(visible spending minus FAFH spending)	0.149**	0.023	−0.003
	(0.060)	(0.025)	(0.030)
observations	16,048	16,048	16,048
R^2	0.475	0.822	0.750
	Panel B. PSID data, 2010–2018		
p50 of ln(visible spending minus FAFH spending)	−0.310*	0.102	−0.091
	(0.135)	(0.132)	(0.115)
p95/p05 of ln(visible spending minus FAFH spending)	0.298**	−0.055	0.048
	(0.109)	(0.069)	(0.061)
observations	29,730	29,730	29,730
R^2	0.614	0.808	0.771

Note: Panel A uses CEX data. Panel B uses PSID data. Visible spending is the sum of personal care, clothing, vehicle and FAFH spending; non-visible spending is total spending minus shelter minus visible spending.
All columns include the same fixed effects and socio-demographic controls as in column (2) in Table 1. Total household expenditure in columns (1) and (2) is instrumented with indicators for the household's income group and the log of after-tax income. The regression in column (3) replaces a control for total household expenditure with a control for non-visible spending, instrumented with indicators for the household's income group and the log of after-tax income.
Sample years 2010-2018. Standard errors clustered at the generation-region level. Robust standard errors in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Our analysis suggests that linking consumption data and data on social media usage is a potentially fruitful avenue to uncover the effects of social spillovers on individual consumption. In particular, in the spirit of Angrist (2014), we expect that increasingly available social media data will facilitate implementing research designs that manipulate peer characteristics in a manner unrelated to individual characteristics. We conjecture that policies designed to influence nutritional outcomes will likely benefit from a better understanding of the effects of consumption externalities on FAFH spending.

References

[1] Andrew B Abel. Asset prices under habit formation and catching up with the joneses, 1990.

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

[3] Francisco Alvarez-Cuadrado and Ngo Van Long. The relative income hypothesis. *Journal of Economic Dynamics and Control*, 35(9):1489–1501, 2011.

[4] Kenneth J Arrow and Partha S Dasgupta. Conspicuous consumption, inconspicuous leisure. *The Economic Journal*, 119(541):F497–F516, 2009.

[5] Orazio Attanasio, Kieran Larkin, Morten O Ravn, and Mario Padula. (s) cars and the great recession. *Econometrica*, 90(5):2319–2356, 2022.

[6] Orazio P Attanasio and Luigi Pistaferri. Consumption inequality. *Journal of Economic Perspectives*, 30(2):3–28, 2016.

[7] Dovile Barauskaite, Justina Gineikiene, Bob M Fennis, Vilte Auruskeviciene, Miwa Yamaguchi, and Naoki Kondo. Eating healthy to impress: How conspicuous consumption, perceived self-control motivation, and descriptive normative influence determine functional food choices. *Appetite*, 131:59–67, 2018.

[8] Erich Battistin, Richard Blundell, and Arthur Lewbel. Why is consumption more log normal than income? gibrat's law revisited. *Journal of Political Economy*, 117(6):1140–1154, 2009.

[9] Adam Bee, Bruce D Meyer, and James X Sullivan. The validity of consumption data: are the consumer expenditure interview and diary surveys informative? Technical report, National Bureau of Economic Research, 2012.

[10] Marianne Bertrand and Adair Morse. Trickle-down consumption. *Review of Economics and Statistics*, 98(5):863–879, 2016.

[11] Timo Boppart. Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196, 2014.

[12] Kerwin Kofi Charles, Erik Hurst, and Nikolai Roussanov. Conspicuous consumption and race. *The Quarterly Journal of Economics*, 124(2):425–467, 2009.

[13] Giacomo Corneo and Olivier Jeanne. Social organization, status, and savings behavior. *Journal of Public Economics*, 70(1):37–51, 1998.

[14] Javier Cravino and Sebastian Sotelo. Trade-induced structural change and the skill premium. *American Economic Journal: Macroeconomics*, 11(3):289–326, 2019.

[15] David M Cutler and Adriana Lleras-Muney. Understanding differences in health behaviors by education. *Journal of health economics*, 29(1):1–28, 2010.

[16] Angus Deaton and Christina Paxson. Intertemporal choice and inequality. *Journal of political economy*, 102(3):437–467, 1994.

[17] J.S. Duesenberry. *Income, Saving, and the Theory of Consumer Behavior*. A Galaxy book, GB 180. Harvard University Press, 1949.

[18] Karen E Dynan, Jonathan Skinner, and Stephen P Zeldes. Do the rich save more? *Journal of political economy*, 112(2):397–444, 2004.

[19] Howard Elitzak and Abigail M Okrent. A retrospective of food-away-from-home expenditures from 1987 to 2017. *America's Eating Habits: Food Away from Home*, EIB-196 USDA, 2018.

[20] Robert Frank, Adam Levine, and Oege Dijk. Expenditure cascades. 2014.

[21] Milton Friedman. *Theory of the consumption function.* Princeton university press, 1957.

[22] Mark Grinblatt, Matti Keloharju, and Seppo Ikäheimo. Social influence and consumption: Evidence from the automobile purchases of neighbors. *The review of Economics and Statistics*, 90(4):735–753, 2008.

[23] Fumio Hayashi. Is the japanese extended family altruistically linked? a test based on engel curves. *Journal of Political Economy*, 103(3):661–674, 1995.

[24] Ori Heffetz. A test of conspicuous consumption: Visibility and income elasticities. *Review of Economics and Statistics*, 93(4):1101–1117, 2011.

[25] Dirk Krueger and Fabrizio Perri. Understanding consumption smoothing: Evidence from the us consumer expenditure data. *Journal of the European Economic Association*, 3(2-3):340–349, 2005.

[26] Peter Kuhn, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn. The effects of lottery prizes on winners and their neighbors: Evidence from the dutch postcode lottery. *American Economic Review*, 101(5):2226–2247, 2011.

[27] Erzo FP Luttmer. Neighbors as negatives: Relative earnings and well-being. *The Quarterly journal of economics*, 120(3):963–1002, 2005.

[28] Joshna Maharaj. Conspicuous consumption: Why are we photographing our food instead of enjoying it? time to ditch the instagram filters, put away that phone, and dig in. December 5, 2016.

[29] Charles F Manski. Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542, 1993.

[30] Franco Modigliani and Richard Brumberg. Utility analysis and the consumption function: An interpretation of cross-section data. *Franco Modigliani*, 1(1):388–436, 1954.

[31] Geoffrey D Paulin. Expenditure patterns of older americans, 1984-97. *Monthly Lab. Rev.*, 123:3, 2000.

[32] Geoffrey D Paulin. Improving data quality in consumer expenditure survey with taxism. *Monthly Lab. Rev.*, 138:1, 2015.

[33] Debraj Ray and Arthur Robson. Status, intertemporal choice, and risk-taking. *Econometrica*, 80(4):1505–1531, 2012.

[34] John Sabelhaus, David Johnson, Stephen Ash, David Swanson, Thesia I Garner, John Greenlees, Steve Henderson, Johnson Sabelhaus, Swanson Ash, and Greenlees Garner. Is the consumer expenditure survey representative by income? 2015.

[35] Michelle J Saksena, Abigail M Okrent, Tobenna D Anekwe, Clare Cho, Christopher Dicken, Anne Effland, Howard Elitzak, Joanne Guthrie, Karen S Hamrick, Jeffrey Hyman, et al. America's eating habits: food away from home. Technical report, United States Department of Agriculture, Economic Research Service, 2018.

[36] Thorstein Veblen. 1994. the theory of the leisure class: an economic study of institutions, 1899.

[37] Bianjun Xia. Status, inequality and intertemporal choice. *The BE Journal of Theoretical Economics*, 10(1):0000102202193517041606, 2010.

APPENDICES

Appendix A

Appendix for Chapter 1: Descriptive Statistics

Table A.1: Trends in ratio of top income quintile to bottom income quintile over time

	1984-1986	1991-1993	2005-2007	2008-2010	log change 1984-1986/2005-2007	log change 2005-2007/2008-2010
Labour Earnings	7.83	8.48	7.89	8.59	0.01	0.09
Before Tax Income	5.52	5.8	6.38	6.41	0.14	0.00
After Tax Income	4.99	5.12	5.85	5.89	0.16	0.01
Total Expenditure	2.88	2.77	2.92	2.78	0.01	-0.05

Notes: Table presents trends in income and consumption inequality using a sample of Consumer Expenditure Survey 1984-2010. Inequality is measured as a ratio of high-income to low-income respondents. High income group includes households with before-tax income larger than 80^{th} percentile and smaller than 95^{th} percentile. Columns 1-4 report average ratio of top income quintile to bottom income quintile over three-year periods. Column 5 represent a log difference between the beginning of the sample and the last period preceding the recession. Column 6 summarizes log changes during the recession. All variables are converted into 1983 constant dollars before averaging.

Appendix B

Appendix for Chapter 2: Proofs

Proof of Lemma 1

To verify Part (i), note that

$$\frac{\partial u\left(v_i/v_j^{\delta}\right)}{\partial v_j} = -\delta v_i v_j^{-\delta-1} u'\left(v_i/v_j^{\delta}\right) < 0,$$

for all $\sigma > 0$, where $u'\left(v_i/v_j^{\delta}\right) > 0$, is the marginal utility of the composite $v_i^{1-\delta}\left(v_i/v_j\right)^{\delta}$.

To verify Part (ii), note that

$$\frac{\partial^2 u\left(v_i/v_j^{\delta}\right)}{\partial v_j^2} = \delta v_i v_j^{-\delta-2} u'\left(v_i/v_j^{\delta}\right)\left(1+\delta\left(1+\frac{v_i/v_j^{\delta} u''\left(v_i/v_j^{\delta}\right)}{u'\left(v_i/v_j^{\delta}\right)}\right)\right),$$

where the elasticity of marginal utility is equal to

$$\frac{v_i/v_j^{\delta} u''\left(v_i/v_j^{\delta}\right)}{u'\left(v_i/v_j^{\delta}\right)} = -\sigma.$$

To verify Part (iii), note that

$$\frac{\partial^2 u\left(v_i/v_j^{\delta}\right)}{\partial v_i \partial v_j} = -\delta v_j^{-\delta-1} u'\left(v_i/v_j^{\delta}\right)(1-\sigma).$$

To verify Part (iv), note that

$$\frac{\partial^3 u\left(v_i/v_j^{\delta}\right)}{\partial v_i \partial v_j^2} = (1-\sigma)\delta v_j^{-\delta-2} u'\left(v_i/v_j^{\delta}\right)\left(1+\delta(1-\sigma)\right). \textbf{ QED}$$

Proof of Proposition 1

We proceed by constructing the unique subgame perfect equilibrium. We consider period 2 first and prove the following result.

Lemma 2 *For any distribution of second-period wealth, with $a_i' + y_i' > 0$, for all i, there is a unique equilibrium distribution of second-period conspicuous consumption, with $v_i' \in (0, a_i' + y_i')$, for all i.*

Consumer i's first-order condition for an interior optimal choice is

$$v_i' = c_i' \left(V'/\alpha\right)^{1/\sigma}, \text{ where } V' \equiv E\left(\left(v_i'\right)^{\delta(\sigma-1)}\right), \tag{B.1}$$

which, using the fact that $c_i' = a_i' + y_i' - v_i'$, implies that

$$v_i' = \left[\frac{\left(V'/\alpha\right)^{1/\sigma}}{1+\left(V'/\alpha\right)^{1/\sigma}}\right]\left(a_i' + y_i'\right), \tag{B.2}$$

for all i. Raising both sides to the power of $\delta(\sigma-1)$, and aggregating over all consumers, we have

$$V' = \left[\frac{\left(V'/\alpha\right)^{1/\sigma}}{1+\left(V'/\alpha\right)^{1/\sigma}}\right]^{\delta(\sigma-1)} E\left(\left(a_i' + y_i'\right)^{\delta(\sigma-1)}\right),$$

equivalently,

$$V' \left[1 + (V'/\alpha)^{-1/\sigma}\right]^{\delta(\sigma-1)} = E\left((a_i' + y_i')^{\delta(\sigma-1)}\right), \tag{B.3}$$

which determines V' as a function of the distribution of wealth at the beginning of the second period.

Suppose that $a_i' + y_i' > 0$, for all i. The left side of equation (B.3) is equal to zero when $V' = 0$, and it approaches infinity as V' grows without bound, for $\sigma > 0$. One can also verify that it is an increasing function of V' if and only if

$$1 - \delta\left(\frac{\sigma-1}{\sigma}\right)\left(\frac{(V'/\alpha)^{-1/\sigma}}{1 + (V'/\alpha)^{-1/\sigma}}\right) > 0.$$

A necessary and sufficient condition is

$$1 + (V'/\alpha)^{1/\sigma} > \delta\left(\frac{\sigma-1}{\sigma}\right),$$

which is the case for all $V' > 0$, since $\delta \leq 1$, and $1 - (1/\sigma) < 1$, for all $\sigma > 1$. Hence, there is a unique solution to equation (B.3) with $V' > 0$. In turn, this implies that there is a unique solution to equation (B.2) with $v_i' \in (0, a_i' + y_i')$ for all i, as required. **QED**

It is easy to verify that an interior optimal choice of second-period assets must satisfy the standard Euler equation for non-conspicuous consumption:

$$\frac{c_i'}{c_i} = \left(\frac{1+r}{1+\rho}\right)^{1/\sigma} \equiv 1 + g_c. \tag{B.4}$$

One can also verify that an interior optimal choice of conspicuous consumption must equate the marginal returns to conspicuous and non-conspicuous consumption in period 1:

$$v_i = c_i\,(V/\alpha)^{1/\sigma}, \text{ where } V \equiv E\left(v_i^{\delta(\sigma-1)}\right). \tag{B.5}$$

Equations (B.4) and (B.5), together with (B.1) imply that

$$\frac{v_i'}{v_i} = (1 + g_c) \left(\frac{V'}{V} \right)^{1/\sigma}. \tag{B.6}$$

Raising both sides to the power of $\delta (\sigma - 1)$, and aggregating over all consumers, we have

$$\frac{V'}{V} = (1 + g_c)^{\frac{\sigma \delta(\sigma-1)}{\sigma - \delta(\sigma-1)}},$$

and therefore

$$\frac{v_i'}{v_i} = (1 + g_c)^{\frac{\sigma}{\sigma - \delta(\sigma-1)}} \equiv 1 + g_v. \tag{B.7}$$

Corollary 1 follows immediately.

Next, note that the intertemporal budget constraint, given by (2.1), implies that

$$\left(1 + \frac{1 + g_c}{1 + r}\right) c_i + \left(1 + \frac{1 + g_v}{1 + r}\right) v_i = Y_i.$$

where $Y_i \equiv a_i + y_i + \frac{y_i'}{1+r}$. Since $v_i = c_i (V/\alpha)^{1/\sigma}$, from equation (B.5), we have

$$v_i = H(V) Y_i \tag{B.8}$$

and

$$c_i = (V/\alpha)^{-1/\sigma} H(V) Y_i, \tag{B.9}$$

for all i, where

$$H(V) \equiv \left(\frac{1}{1 + \frac{1+g_v}{1+r} + \left(1 + \frac{1+g_c}{1+r}\right) (V/\alpha)^{-1/\sigma}} \right).$$

In order to solve for V, raise both sides of equation (B.8) to the power of $\delta (\sigma - 1)$, and aggregate over all consumers to obtain

$$V\left[H\left(V\right)\right]^{-\delta(\sigma-1)} = E\left(Y_i^{\delta(\sigma-1)}\right). \tag{B.10}$$

Writing equation (B.10) as

$$V^{1-\frac{\delta(\sigma-1)}{\sigma}}\left(\left(1+\frac{1+g_v}{1+r}\right)V^{1/\sigma}+\left(1+\frac{1+g_c}{1+r}\right)\alpha^{1/\sigma}\right)^{\delta(\sigma-1)} = E\left(Y_i^{\delta(\sigma-1)}\right),$$

it is easy to see that the left side of the equation is an increasing function of V, for all $\sigma > 1$. One can also verify that it is equal to zero when $V = 0$, and it approaches infinity as V grows without bound. Hence, there is a unique solution V^* to equation (B.10) with $V^* > 0$. In turn, this implies that there is a unique solution to equation (B.8) with $v_i \in (0, Y_i)$ for all i.

It remains to show that our candidate equilibrium is consistent with aggregation. In particular, we need to show that equation (B.3) and equation (B.10) are satisfied simultaneously. To verify that this is the case, note that the consumer's budget constraint implies that

$$\begin{aligned} a_i' + y_i' &= (1+r)\left(a_i + y_i - c_i - v_i\right) + y_i' \\ &= (1+r)\left(a_i + y_i + \frac{y_i'}{1+r} - (c_i + v_i)\right), \end{aligned}$$

which, noting that

$$c_i + v_i = a_i + y_i - \frac{a_i'}{1+r} = \left(1 + (V/\alpha)^{-1/\sigma}\right)H(V)Y_i,$$

can be written as

$$a_i' + y_i' = \left(\left(1 + g_v + (1+g_c)\left(V/\alpha\right)^{-1/\sigma}\right)H(V)\right)^{\delta(\sigma-1)} Y_i.$$

Raising both sides to the power of $\delta(\sigma-1)$, and aggregating over all consumers, we have

$$E\left((a_i'+y_i')^{\delta(\sigma-1)}\right)=\left(\left(1+g_v+(1+g_c)(V/\alpha)^{-1/\sigma}\right)H(V)\right)^{\delta(\sigma-1)}E\left(Y_i^{\delta(\sigma-1)}\right).$$

Next note that the left side of equation (B.3) satisfies

$$V'\left[1+(V'/\alpha)^{-1/\sigma}\right]^{\delta(\sigma-1)}=\left(\frac{1+g_v}{1+g_c}\right)^{\sigma}V\left(\frac{1+g_v+(1+g_c)(V/\alpha)^{-1/\sigma}}{1+g_v}\right)^{\delta(\sigma-1)},$$

where we have used the fact that $V'/V=\left(\frac{1+g_v}{1+g_c}\right)^{\sigma}$ (see equation (B.6)). Hence equation (B.3) implies that the right sides of the above two equations must be equal, which together with equation (B.7) implies that equation (B.10) must hold, as required. **QED**

Proof of Proposition 2 and Proposition 3

The proof of Proposition 1 implies that (i) c_i/Y_i falls with V, (ii) v_i/Y_i and $\frac{v_i}{c_i+v_i}$ rise with V, and (iii) $\frac{c_i+v_i}{Y_i}$ rises with V if and only if $g_c>g_v$.

Hence, in order to characterize the effects of changes in the distribution of permanent income on consumption and saving behavior, it remains to characterize the effect of such changes on the equilibrium value of V, which is given by equation (B.10). To that end, first note that in the proof of Proposition 1 we show that the left side of (B.10) is an increasing function of V.

Then, Proposition 2 follows from the standard definition of first-order stochastic dominance and the fact that $Y_i^{\delta(\sigma-1)}$ is an increasing function of Y_i for all $\sigma>1$. Similarly, Proposition 3 follows from the standard definition of a mean-preserving spread and the facts that $Y_i^{\delta(\sigma-1)}$ is an increasing function of Y_i for all $\sigma>1$, and it is also concave if $\delta(\sigma-1)<1$ and convex if $\delta(\sigma-1)>1$. **QED**

Appendix C

Appendix for Chapter 3: CEX Expenditure Categories

Table C.1: Description of Expenditure Categories

Expenditure category	Share of total expenditure (%)
Housing	30.05
	Visible Expenditure
Children's clothing	0.77
Shoes and other apparel	1.22
Adult clothing	1.88
Personal care	0.82
Food Away from Home	4.68
Vehicle purchase	10.14
	Non-Visible Expenditure
Food at home	12.67
Tobacco, other smoking	1.11
Alcoholic beverages	0.99
Appliances, phone, computers	5.13
Entertainment fees, admissions	1.72
Entertainment equipment	3.83
Utilities	5.82
Other transportation	1.06
Health expenditure	5.40
Education	1.22
Furniture and fixtures	1.24
Domestic services	1.24
Cash contributions	2.13
Observations	70,952

Note: Shares are computed as annual average expenditure in a category over annual average total expenditure. Then shares are averaged over the sample period 1984-2018. All series are deflated using CPI to 1983 dollars. Averages are computed using CE survey weights.

Table C.2: Overall sample summary statistics and by age of the household head

	ALL			OLD			Young			
Variable	**Obs**	**Mean**	**SD**	**Obs**	**Mean**	**SD**	**Obs**	**Mean**	**SD**	**Diff**
Visible	70,952	4,837	5,118	34,493	4,769	5,227	36,459	4,901	5,012	***
Non-Visible	70,952	10,937	5,894	34,493	11,483	6,445	36,459	10,426	5,276	***
Total	70,952	22,198	11,460	34,493	23,067	12,291	36,459	21,385	10,558	***
Housing	70,952	6,424	3,770	34,493	6,815	4,009	36,459	6,057	3,492	***
After-tax income	70,952	27,955	15,982	34,493	29,195	16,852	36,459	26,794	15,029	***
Visible share	70,952	0.12	0.09	34,493	0.11	0.09	36,459	0.13	0.09	
Ln(Visible)	70,952	8.03	1.00	34,493	7.98	1.05	36,459	8.09	0.94	
p50 ln(visible)	70,952	8.03	0.16	34,493	8.00	0.18	36,459	8.06	0.14	
p95/05 ln(visible)	70,952	1.50	0.07	34,493	1.55	0.07	36,459	1.46	0.05	
Education										
Never Attended	70,952	0.00	0.04	34,493	0.00	0.05	36,459	0.00	0.03	***
HS or less	70,952	0.38	0.49	34,493	0.41	0.49	36,459	0.35	0.48	***
Post-Secondary	70,952	0.49	0.50	34,493	0.46	0.50	36,459	0.53	0.50	***
MA/PhD	70,952	0.12	0.33	34,493	0.12	0.33	36,459	0.12	0.32	**
Number or persons in the household										
Less than 3	70,952	0.48	0.50	34,493	0.59	0.49	36,459	0.38	0.48	***
3 to 4	70,952	0.38	0.49	34,493	0.32	0.47	36,459	0.44	0.50	***
More than 4	70,952	0.14	0.35	34,493	0.09	0.29	36,459	0.18	0.39	***
Number of dual earner households	70,952	0.66	0.47	34,493	0.66	0.47	36,459	0.67	0.47	**
Gender										
Male	70,952	0.57	0.49	34,493	0.57	0.50	36,459	0.58	0.49	
Female	70,952	0.43	0.49	34,493	0.43	0.50	36,459	0.42	0.49	
Region										
NorthEast	70,952	0.20	0.40	34,493	0.21	0.41	36,459	0.19	0.39	***
South	70,952	0.23	0.42	34,493	0.23	0.42	36,459	0.23	0.42	
MidWest	70,952	0.33	0.47	34,493	0.33	0.47	36,459	0.33	0.47	
West	70,952	0.23	0.42	34,493	0.22	0.42	36,459	0.24	0.43	***
Race										
White	70,952	0.84	0.37	34,493	0.85	0.36	36,459	0.83	0.38	***
Black	70,952	0.11	0.32	34,493	0.11	0.31	36,459	0.12	0.32	**
Other	70,952	0.05	0.22	34,493	0.04	0.21	36,459	0.05	0.23	***
Age	70,952	44	11	34,493	53.81	5.73	36,459	35	5	***

Note: Sample: CEX waves 1980-2018. Expenditure and income series deflated to 1983 dollars using CPI index. Summary statistics for the whole sample and for the two age groups. Last column indicates that the means between two age groups are different at 1% significance level)