

Ferrer, Ana M.; González, Francisco M.; Nesterova, Iuliia

Working Paper

Conspicuous consumption and visible inequality

Working Paper Series, No. 81

Provided in Cooperation with:

Canadian Labour Economics Forum (CLEF), University of Waterloo

Suggested Citation: Ferrer, Ana M.; González, Francisco M.; Nesterova, Iuliia (2024) : Conspicuous consumption and visible inequality, Working Paper Series, No. 81, University of Waterloo, Canadian Labour Economics Forum (CLEF), Waterloo

This Version is available at:

<https://hdl.handle.net/10419/307596>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Canadian Labour Economics Forum

WORKING PAPER SERIES

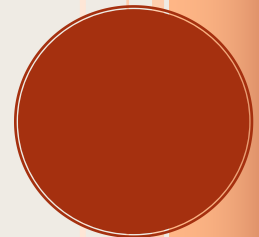
Conspicuous Consumption and Visible Inequality

Ana Ferrer (University of Waterloo)

Francisco M. Gonzalez (University of Waterloo)

Iuliia Nesterova (University of Waterloo)

CLEF WP #81



Conspicuous consumption and visible inequality*

Ana Ferrer, Francisco M. Gonzalez and Iuliia Nesterova

December 2024

Abstract

We document a significant effect of visible inequality on household spending in the United States from 2010 to 2018, but we do not find a comparable effect before the Great Recession. Our proposed definition of *visible inequality* refers to differences across households in their expenditures on highly noticeable consumption categories, including clothing, personal care, food away from home, and vehicles. Our empirical strategy exploits robust implications of a simple intertemporal model of conspicuous consumption, where a household's reference group consists of others in the same age group. Household spending is influenced by the distribution of permanent incomes within age groups, and visible inequality serves as a proxy for permanent income inequality. Our findings indicate that consumption externalities significantly distort household spending, which could have been up to 25 percent lower without these distortions. We hypothesize that low interest rates and social media contribute to these effects.

JEL classification: D1, E21.

Keywords: visible inequality, relative consumption, permanent income inequality.

*Ferrer: University of Waterloo and IZA (aferrer@uwaterloo.ca). Gonzalez: University of Waterloo (francisco.gonzalez@uwaterloo.ca). Nesterova: University of Waterloo (inesterova@uwaterloo.ca). We have benefited from comments by Francisco Alvarez-Cuadrado, Pierre Chausse, Curtis Eaton, Herb Emery, David Green, Rohini Pande, and Tom Parker. We gratefully acknowledge financial support from the Social Sciences and Humanities Research Council of Canada.

1 Introduction

People seem to care about relative consumption (Luttmer, 2005), but this does not necessarily mean that relative concerns will result in distorted consumption, let alone insufficient savings (Corneo and Jeanne, 1998; Arrow and Dasgupta, 2009). The reason is that consumption choices are inherently intertemporal. Even if current consumption enhances one’s relative standing today, it can diminish that standing in the future. From an empirical perspective, the challenge lies in distinguishing between the impact of group behavior on individual choices and the aggregation of similar individual behaviors (Manski, 1993). Additionally, standard consumption theory suggests that identifying these separate forces is particularly challenging in this context. The complication arises from the fact that household expenditures are linked to the household’s *permanent* income, which is not directly observable (Friedman, 1957).

In this paper, we use data from the Consumer Expenditure Survey (CE) 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. Instead, they are most likely the symptom of actual distortions associated with consumption externalities.

Our empirical strategy relies on a simple 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 four key features. First, it focuses on intertemporal choice, so one 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 *visible* goods and to the former as *non-visible* goods. Third, an individual’s reference group consists of other individuals in the same location/age group/year. Fourth, 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 visible spending involves the comparison of her consumption of visible 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.

Our model illustrates how equilibrium consumption behavior is influenced by both the households’ own permanent income and the distribution of visible spending within their reference

group. Additionally, the distribution of visible spending is a function of the distribution of permanent income. Our empirical analysis exploits two straightforward implications of these properties. First, consumption, by itself, is not a reliable proxy of permanent income in the presence of distortions associated with consumption externalities. This is because individuals with the same consumption, but different reference groups, must have different permanent incomes if consumption externalities are distorting consumption behavior. Second, controlling for group behavior in consumption regressions requires controlling for the distribution of *permanent*, rather than *current*, incomes. This is because consumers are forward looking.

The above implications are not unique to our specific model. More generally, they stem from optimal intertemporal decision-making in the presence of relative concerns.¹ However, they are routinely disregarded in existing empirical work on relative consumption. Instead, standard consumption regressions assume tacitly that the distribution of *current* incomes within the reference group is a reliable proxy for the relevant, observable group behavior, which is unlikely to be the case unless consumers are fully myopic. Furthermore, they either fail to control for permanent incomes (e.g., Bertrand and Morse, 2016) or assume that current consumption is a valid proxy for permanent income regardless of the reference group controls (e.g., Charles et al., 2009).

In practical terms, our model implies that an individual’s consumption expenditure is a valid proxy for her permanent income, conditional on the distribution of expenditure on visible goods across individuals in her reference group. This insight is the centerpiece of our empirical strategy. To implement this insight empirically, one needs a measure of visible spending, a definition of reference groups and a proxy for the distribution of visible spending within each reference group. We address these three issues as follows. First, following Charles et al. (2009) and Heffetz (2011), we define visible spending as spending on highly noticeable consumption categories. Our measure of visible spending 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. Specifically, 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 focus on the first two moments of the distribution of log spending on visible goods as the controls for visible group behavior. The second moment provides a measure of dispersion across households in their expenditures on highly noticeable consumption categories. In this sense, we say it is a measure of *visible inequality*. Restricting attention to the first two moments

¹Corneo and Jeanne, (1998), Arrow and Dasgupta (2009), Xia (2010), Alvarez-Cuadrado and Long (2011) and Ray and Robson (2012) are some examples of dynamic models of relative consumption.

is motivated by the observation that both CE and PSID data indicate that the distribution of consumption expenditures across U.S. households, within cohorts, is much closer to log normal than the distribution of income (Battistin et al., 2009). Standard intertemporal choice models of consumption suggest that the reason is that permanent income is closer to log normal than current income.

It is well understood that total expenditure and its component expenditures are jointly determined and therefore endogenous. 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). Recognizing that income is measured with error as well, particularly in the tails, we drop households in the top and bottom 5 percent of the income distribution. 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.

Identifying a causal effect running from group behavior to individual behavior requires that we disentangle 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 may drive consumption expenditure in those states, or that state-level variation in redistributive policies may drive variation in permanent-income inequality. Moreover, individuals may systematically sort into states in such a way that their spending behaviors are spuriously correlated. These are non-trivial concerns for analyses of conspicuous consumption that exploit state-level variation in the CE data (e.g., Charles et al., 2009; Bertrand and Morse, 2016).

To address these concerns, first we consider a fairly broad specification of reference groups, as explained above. Specifically, our focus on broad geographical regions, rather than states, serves to mitigate potential selection problems, at the expense of introducing within-group heterogeneity. Moreover, our main econometric specifications include a battery of fixed effects. It is noteworthy that the effect of dispersion in log visible spending on household expenditures over the period 2010–2018 is robust to controlling for fixed effects for both state-year interactions and age-group-year interactions. It is also robust to controlling for individual permanent incomes. This strategy makes it unlikely that the estimated effect of visible inequality on household expenditures in our regressions is driven by correlated variation in the distribution of unobserved permanent incomes.

In principle, despite the inclusion of a wide array of fixed effects in our regressions, and the fact that we control for the household’s permanent income, spending might still be affected by some factor omitted from the model. A further endogeneity concern arises because some categories of consumption are poorly reported in the CE and this reporting seems to have degraded over time (Bee et al., 2015). Expenditure underreporting is especially evident at the very top percentiles of income and expenditures (Sabelhaus et al., 2015). Accordingly, the estimated effects of the mean and the variance of log visible spending within the reference group might be biased. We address these endogeneity concerns in several ways.

First, we replace the mean and the variance of the distribution of log visible spending within a reference group with the median and the ratio of the 95th to the 5th percentile, our measure of visible inequality. Both of these measures are less sensitive to underreporting at the tails of the distribution. This strategy serves to limit the potential impact of non-classical measurement error.

Second, we argue that the effect of visible group spending on household spending can be consistently estimated by instrumenting median log visible spending and visible inequality with contemporaneous measures of income inequality and past values of visible inequality. Furthermore, our measure of income inequality uses data from the Current Population Survey (CPS), which is well known to have better measures of income at the top of the distribution.

Finally, we run our main specification using data from the 2011-2019 waves of the PSID, whose measures of income and expenditures at the top of the distribution are reputedly less susceptible to measurement error. Our results suggest that the estimated effect of visible inequality on household spending is driven by neither correlated effects nor measurement error.

Our analysis contributes to the literature on consumption inequality by calling attention to the effect of *visible* inequality on spending behavior in the U.S. over the period 2010–2018.² Qualitatively, we find that increases in visible inequality tend to increase visible spending and decrease both current non-visible spending and current savings. Quantitatively, our estimates indicate considerable relative consumption effects. Average household spending from 2010 to 2018 might have been up to 25% smaller in a counterfactual scenario without distortions from consumption externalities.

In Section 2, we discuss our empirical methodology in detail. In Section 3, we report estimates of the effect of visible inequality on visible spending over the period 1984–2018. In Section 4, 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 5, we focus on the period after the Great Recession and confirm the robustness of our main findings to alternative specifications. In Section 6 we discuss our main

²See Attanasio and Pistaferri (2016) for a survey of the literature on consumption inequality.

findings and conclude.

2 Data and empirical methodology

Our main data come from the interview sample of the Consumer Expenditure Survey (CE) of the U.S. Bureau of Labor Statistics. We use CE data from the survey years 1984-2018. Our sample restrictions follow Aguiar and Bils (2015). We start by aggregating to the 20 expenditure categories they consider, which are listed in Table A1. We measure consumption as the sum over four quarters of expenditure surveys for a given household. Our measures of after-tax income, which we use to proxy for the household permanent income, also follow Aguiar and Bils (2015). For more details on the CE, sample restrictions, and measures of consumption and income, see Appendix A.

Table A2 in Appendix A summarizes the main variables used in the analysis by age group. Average family income after tax is almost \$28,000 and total spending about \$22,000. Approximately, 29 percent of total spending is in housing and almost 22 percent goes to visible consumption goods (clothing, personal care, vehicles and food away from home). Older households have higher family income and total expenditure, but spend less on 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 less 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.

To identify the effect of visible (consumption) inequality, which is endogenous, we construct measures of the before-tax income distribution in each region/age-group/year cell to serve as instruments. Following Bertrand and Morse (2016), our data come from the large samples of the March Current Population Survey (CPS). We use data from the 1983–2017 surveys, which provide information on income for the previous calendar year. For details on the sample and the measures of income, see Appendix A.

Below, we confirm our main results using data from the 2011-2019 waves of the Panel Study of Income Dynamics (PSID). Despite covering the broad range of consumption expenditures for a shorter time period than the CE, 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. See Appendix A for more details.

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 are interested in understanding whether relative concerns cause distortions in individual consumption behavior. We need to address three key issues: (i) the empirical spec-

ification of conspicuous consumption, (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 issues in turn.

2.1 Visible versus non-visible spending

According to Veblen (1899), conspicuous consumption is a way of displaying wealth, which in turn confers social status. While some spending may be intended to signal unobservable wealth, here we think of “conspicuous consumption” more generally as observable spending on certain goods in a world in which people care about their consumption of those goods relative to that of other people. Our theoretical framework below formalizes this idea. For the purpose of our empirical analysis, a preliminary question is which categories of spending may reflect relative consumption effects.

Rather than aggregating consumption expenditure across all categories, it is important to distinguish conceptually between visible and non-visible spending. Not all categories of spending are equally noticeable, and the extent to which spending is noticeable is not an exogenous characteristic of a given consumption good. For instance, the rise of social media has made it possible for consumers to decide what to make observable, and to whom, to an unprecedented extent. Furthermore, spending on some consumption goods, even if it is noticeable, may be driven by considerations other than relative concerns. For instance, tobacco and alcohol consumption may be heavily influenced by addiction, and by complex cultural and social aspects. As noted by Heffetz (2011), one may even argue that not consuming these goods may be a signal of high social status. Accordingly, we exclude spending on both tobacco and alcohol from our working definition of visible spending. We also treat housing expenditure as a separate category. This is by far the largest expenditure for many households, and it is unlikely that it is driven primarily by relative concerns. For instance, differential treatment in the housing and mortgage markets might cause young individuals to have very different expenditures on housing, even in the absence of relative concerns.

With the above considerations in mind, we define visible spending as the sum of annual spending on clothing, personal care, vehicles and food away from home. All of these categories of spending rank among the most visible in both Charles et al.’s (2009) visible goods survey and Heffetz’s (2011) survey.³ Our definition differs from the one proposed in Charles et al. (2009) only in that we include food away from home among the visible goods. In this sense we follow

³Charles et al. (2009) conducted a survey of about 320 graduate students in the University of Chicago’s Harris School of Public Policy and the University of Chicago’s Graduate School of Business. Heffetz (2011) constructs a visibility index based on a U.S. national telephone survey of 480 individuals from May 2004 to February 2005.

Bertrand and Morse (2016), who rank the categories used by Charles et al. (2009) according to Heffetz’s (2011) visibility index and highlight the fact that food away from home is significantly more visible than food at home. Indeed, spending on food away from home plays an important role in our results.

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.

2.2 Reference group

A natural hypothesis is that individuals only respond to visible spending 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 may have heterogeneous attitudes towards conspicuous consumption. In this sense, it is useful to take as given some broad characteristic of reference groups as salient. 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 visible spending 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, it is well known that the dispersion of expenditures in the CE is increasing over the life cycle.⁴ Below, 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 CE. We make this choice so we

⁴In 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 CE actually declines over the life cycle, which undermines the view that dispersion is driven by idiosyncratic permanent income shocks.

have enough observations in a given region/age-group/year cell to construct reliable measures of the reference group’s consumption distribution. Furthermore, 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.

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

2.3 Empirical methodology

From an empirical standpoint, we need to address two fundamental problems. First, the influence of group behavior on individual behavior can be mistaken for the aggregation of similar individual behavior (Manski, 1993). Second, in the case of consumption externalities, an additional problem is that household expenditures should be related to the household’s expectations of future income, which are inherently unobservable. In the presence of these problems, the identification of 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.

2.3.1 Theoretical framework

In this section we present a simple model in which conspicuous consumption can lead to distortions in current and future consumption behavior. The formal model provides the basis for our empirical strategy in the sense that it helps clarify the specific assumptions that we make in order to be able to control for (unobservable) permanent income and to identify the endogenous effect of group behavior on individual behavior. Proofs of all formal statements are found in Appendix B.

The model has the following key features. First, it focuses on intertemporal choice, so we can understand the mapping between permanent income and consumption behavior. Second, for some goods people 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 *visible* goods and to the former as *non-visible* goods. Third, an individual’s reference group consists of other individuals in the same location/age-group/year, and for simplicity, we consider only one location.

The model

There is a continuum of consumers, with mass one, living for two periods. Each consumer faces a standard intertemporal choice problem. 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 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

$$c_i + \frac{c'_i}{1+r} + v_i + \frac{v'_i}{1+r} = Y_i, \quad (1)$$

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 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 *visible* consumption. Consumer i 's preferences are defined over consumption in the two periods according to

$$U_i = \alpha u(c_i) + E(u(v_i/v_j^\delta)) + \frac{1}{1+\rho} [\alpha u(c'_i) + E(u(v'_i/v'_j{}^\delta))], \quad (2)$$

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 visible consumption by others. Here, instead, we assume that utility from visible 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 visible consumption only through the average utility associated with positional concerns.

Since $v_i/v_j^\delta = v_i^{1-\delta} (v_i/v_j)^\delta$, positional concerns are such that consumer j 's visible consumption v_j affects consumer i 's 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 .⁵ The parameter δ reflects the strength 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

$$u(z) = \frac{z^{1-\sigma}}{1-\sigma}, \text{ with } \sigma > 1, \quad (3)$$

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 visible and non-visible consumption.

A subgame perfect equilibrium is characterized by a distribution of visible 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 visible and non-visible consumption, the precise manner in which 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) $\frac{\partial u(v_i/v_j^\delta)}{\partial v_j} < 0$, for all $\sigma > 0$;
- (ii) $\frac{\partial^2 u(v_i/v_j^\delta)}{\partial v_j^2} < 0$ if and only if $\sigma > 1 + 1/\delta$;
- (iii) $\frac{\partial^2 u(v_i/v_j^\delta)}{\partial v_i \partial v_j} > 0$ if and only if $\sigma > 1$;
- (iv) $\frac{\partial^3 u(v_i/v_j^\delta)}{\partial v_i \partial v_j^2} > 0$ if and only if $\sigma > 1 + 1/\delta$.

⁵In this sense, our specification of positional concerns is as in Abel (2005).

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 visible 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-visible consumption, but also with respect to visible 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 , $E(u(v'_i/v_j'^\delta))$, and thus U_i , will decrease, for all i , with an increase in the distribution of visible consumption in the sense of first-order stochastic dominance, since $u(v_i/v_j^\delta)$ is a decreasing function of v_j . Thus, everything else equal, individuals dislike it when visible consumption in their reference group grows.

Second, for a fixed level of v_i , $E(u(v'_i/v_j'^\delta))$, and thus U_i , will decrease, for all i , with a mean-preserving spread in the distribution of visible consumption if and only if $\sigma > 1 + 1/\delta$, since $u(v_i/v_j^\delta)$ 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 visible 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.

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 visible and non-visible consumption in period 1 are given by*

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

for all i , where the function H and the equilibrium outcome $V^* = E((v_j^*)^{\delta(\sigma-1)})$ are characterized in the Appendix as a function of the exogenous distribution of permanent incomes, F_Y . The levels of visible and non-visible consumption in period 2 for all i are given by

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

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 visible and non-visible consumption. Each of the two shares is a function of the distribution of visible 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 $V^* = E((v_j^*)^{\delta(\sigma-1)})$, and that V^* is uniquely determined as a function of $E(Y_i^{\delta(\sigma-1)})$, 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 visible and non-visible 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 visible or non-visible 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 front-load visible consumption relative to non-visible 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 visible consumption relative to non-visible consumption. Finally, note that saving rates are independent of the distribution of permanent incomes when equilibrium levels of visible and non-visible 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 visible consumption, the utility cost of not raising one's own visible 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 visible 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 visible consumption to increases in the inequality of permanent incomes. Thus, individuals engage in more visible 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 this sense, we formalize a notion of equilibrium “expenditure cascades” in the spirit of the original notion proposed in Frank et al. (2014).

In the Appendix we show that the marginal rate of substitution of visible for non-visible consumption is a monotone function of $V^* = E((v_j^*)^{\delta(\sigma-1)})$. Accordingly, when permanent incomes become more unequal, then individual non-visible consumption falls exactly when individual visible 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.⁶

2.3.2 Econometric approach

Our econometric approach relies explicitly on the above theoretical framework, which is special in several important respects. There is no risk, preferences are time-consistent and isoelastic, utility is additively separable across consumption goods and over time, and households are heterogeneous only with respect to permanent incomes. It is well known that these assumptions have a number of counterfactual implications (Deaton, 1992), and we do not wish to impose them on the data. However, we believe other broad features of our model are general. In particular, building on Proposition 1, our empirical strategy rests on the assumption that if visible spending distorts consumption behavior, then *equilibrium expenditures are a function of both the household's permanent income and the distribution of visible spending within the reference group*.

To study the effect of visible group behavior on individual spending we start by considering the following specification of the visible consumption process:

$$\ln v_{hgrt} = X_{hgrt}\beta + \alpha_1 E[\ln v_{jgrt-1}] + \alpha_2 \text{Var}[\ln v_{jgrt-1}] + \phi \ln Y_{hgrt} + \varepsilon_{hgrt}, \quad (4)$$

for $t \geq 0$, where v_{hgrt} is visible spending of household h belonging to generation g and located

⁶This is a common feature of alternative intertemporal models of conspicuous consumption (e.g., Corneo and Jeanne, 1998; Arrow and Dasgupta, 2009; Xia, 2010).

in region r at time t , Y_{hgrt} is her permanent income, and $E[\ln v_{jgrt-1}]$ and $\text{Var}[\ln v_{jgrt-1}]$ are the first and second moments of the cross-sectional distribution of households, indexed by j , who belong to generation g , and who are located in region r at time $t-1$. 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 household's size and the number of earners. Throughout the paper we refer to age-groups and generations interchangeably.

Equation (4) is motivated by Proposition 1, which characterizes log spending as a function of the household permanent income and the distribution of log visible spending for a set of homogeneous households belonging to the same reference group. Note that our theoretical framework provides a simple example in which the separate Euler equations for visible and non-visible spending may feature no social effects, even though they do distort consumption levels. Accordingly, our reduced-form approach focuses on the effect of consumption externalities on consumption levels, in contrast to the standard Euler equation approach in the consumption literature, which focuses on consumption growth rates.⁷ Additionally, we have made three assumptions. First, we allow for systematic differences in consumption levels across households with different socio-demographic characteristics, as explained above.

Second, we have assumed that the effect of group behavior on household spending occurs via the first two moments of the distribution of log visible spending. This assumption is based on data from the CE and the PSID, which suggest that the distribution of consumption expenditures across U.S. households, within cohorts, is much closer to log normal than the distribution of income (Battistini et al., 2009). Standard intertemporal choice models of consumption suggest that the reason is that the distribution of permanent income is closer to log normal than the distribution of current income. In that spirit, our basic model implies the following.

Proposition 4 *Suppose that the cross-sectional distribution of permanent incomes is log normal. Then the cross-sectional distributions of period-1 visible, non-visible and total consumption are log normal as well, with*

$$\begin{aligned} E[\ln v_j^*] &= \ln H(V^*) + E[\ln Y_j], \\ E[\ln c_j^*] &= \frac{-1}{\sigma} \ln(V^*/\alpha) + \ln H(V^*) + E[\ln Y_j], \\ E[\ln(c_j^* + v_j^*)] &= \ln\left(1 + (V^*/\alpha)^{-1/\sigma}\right) + \ln H(V^*) + E[\ln Y_j], \end{aligned}$$

⁷Attanasio and Weber (2010) discuss the advantages and limitations of each approach. Maurer and Meier (2008) and De Giorgi et al. (2020) are two interesting analyses of conspicuous consumption that use the Euler equation approach.

and

$$\text{Var} [\ln v_j^*] = \text{Var} [\ln c_j^*] = \text{Var} [\ln (c_j^* + v_j^*)] = \text{Var} [\ln Y_j],$$

where the function H and the equilibrium outcome $V^* = \text{E} [(v_j^*)^{\delta(\sigma-1)}]$ are characterized in Proposition 1.

Proposition 4 follows immediately from Proposition 1, which implies that visible, non-visible and total spending are all linear in the agent's own permanent income. The counterfactual implication that all consumption categories exhibit a unit elasticity with respect to permanent income is common to standard models of intertemporal consumption. It is driven by the additive separability across goods and periods, by the homotheticity of the subutility functions, and the fact that they are independent of age (Deaton, 1992). These assumptions, however, allow us to characterize the interaction between intertemporal choice and relative concerns. In particular, relative concerns imply a dependence of consumption behavior on the distribution of permanent incomes, even though our assumptions imply that there are neither luxuries nor necessities. One could allow for these, for example, by making the elasticities of substitution a function of age, or by considering Stone-Geary subutilities.

The third assumption in equation (4) is made in order to attenuate the effect of measurement error in consumption. To that end, we use lagged values of the moments of the distribution of log visible spending. Otherwise, our estimates might capture contemporaneous correlation between measurement errors in the aggregate and individual variables. Intuitively, it is plausible that visible group behavior becomes visible with some delay.

There are a number of issues that arise in estimating (4) using data from the CE. First, we should control for the household's permanent income, which is unobservable. Second, we should consider the possibility that unobservable determinants of household spending may be correlated across households within the same reference group, which is well known to complicate the identification of causal social effects. Third, we should consider the possibility of measurement error in the CE. We discuss how we address each of these challenges in turn.

Previous empirical analyses of (non-conspicuous) consumption rely on textbook consumption theory to argue that consumption expenditure is a valid proxy of permanent income (e.g., Aguiar and Bils, 2015). However, not only total expenditure and consumption decisions are jointly determined and therefore endogenous, but also measurement error in the components of consumption and expenditures are likely to be correlated. To address these problems, the usual strategy in the consumption literature is to instrument individual consumption expenditure with measures of current income (such as indicators for the household's income category and the log of after-tax income). This strategy exploits the fact that total expenditure reflects

permanent income and thus it will be correlated with current income.

The literature on conspicuous consumption has followed the standard strategy in proxying permanent income with expenditure and instrumenting this with current income (Charles et al., 2009). However, our theoretical framework illustrates why this approach becomes problematic in the presence of consumption externalities. Formally, Proposition 1 implies that changes in a household’s permanent income are reflected into changes in the consumption of the household *only when conditioned on the distribution of conspicuous consumption within her reference group*. The reason is intuitive. If we fail to account for group behavior, two individuals with identical consumption levels but exposed to different visible 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. This is not a concern for our study, because we are not asking whether people belonging to different reference groups have a relatively higher or lower propensity to respond to group spending. Rather, we are asking whether individuals vary their spending in response to changes in the distribution of visible spending within their own reference group.

A second challenge with estimating (4) lies in distinguishing between the impact of group behavior on individual choices and the aggregation of similar individual behaviors (Manski, 1993). In the context of conspicuous consumption, this challenge is amplified because the distribution of permanent incomes is unobservable, and so is the covariance between permanent incomes and household characteristics that can affect consumption behavior. Accordingly, estimates of the effect of group behavior may simply reflect the fact that both permanent incomes and household characteristics vary across groups, or that within-group permanent incomes are changing together over time. This is particularly likely to be the case if one fails to control for the individual households’ permanent incomes. Similarly, the distribution of preferences (e.g., discount factors, relative importance of visible versus non-visible goods) across households is unobservable, and it likely varies both within and across age groups and regions, and possibly over time. Furthermore, unobservable determinants of spending 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} in (4) 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.

Our identification strategy relies on a battery of fixed effects in order to minimize the potential problems caused by the above sources of confounding effects. Specifically, our starting

specification includes fixed effects for generation-region interactions, generation-year interactions, and region-year interactions, as well as the interaction between generation-region groups and a linear time trend to account for time-, location- and generation-specific confounding factors that could influence the reference group's distribution of expenditure. Furthermore, anticipating our results below, our main specifications shall include fixed effects for state-year interactions.

A third issue arises because some categories of consumption are poorly reported in the CE and this reporting seems to have degraded over time. Poorly measured categories include non-visible categories such as furniture and alcohol as well as visible categories such as food away from home and clothing. Well-measured categories of spending also include non-visible categories such as food at home as well as visible categories such as new motor vehicles (Bee et al., 2015). Expenditure underreporting is especially evident at the very top percentiles of income and expenditures (Sabelhaus et al., 2015). One concern is that non-classical measurement error may have an impact on the dependent variable in our regressions. However, Meyer and Sullivan's (2023) recent analysis of the evolution of consumption inequality in the CE does not support an important bias associated with non-classical measurement error, provided one excludes the extreme tails of the distribution of income and expenditure, as we do. Accordingly, our main concern is that the estimated effects of the mean and the variance of log visible spending within the reference group might be biased. We address this issue by replacing these two moments with the median and the ratio of the 95th to the 5th percentile, which are insensitive to spending at the very top. Furthermore, among our robustness checks below we shall consider an IV approach to address remaining endogeneity concerns about our measures of visible group behavior.

Motivated by the above considerations, we begin by estimating the following regression

$$\begin{aligned}
\ln(\text{visible spending})_{hgrt} = & \text{p50 of } \ln(\text{visible spending})_{grt-1} \\
& + \text{p95/p05 of } \ln(\text{visible spending})_{grt-1} \\
& + \ln(\text{total spending})_{hgrt} \\
& + (X)_{hgrt} \\
& + \text{generation}_g + \text{region}_r + \text{year}_t \\
& + (\text{generation}_g \times \text{region}_r) + (\text{region}_r \times \text{year}_t) + (\text{generation}_g \times \text{year}_t) \\
& + (\text{generation}_g \times \text{region}_r \times \text{linear trend}_t) + u_{hgrt},
\end{aligned} \tag{5}$$

where we instrument for the log of individual total spending using indicators for the household's income category and the log of after-tax income, as explained above.

3 Impact of visible inequality on visible spending

In this section we report estimates of the effect of the distribution of consumption of the reference group on individual household consumption under different specifications of equation (5). All regressions in Table 1 include controls for 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 in all regressions are clustered at the generation-region level.

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. As discussed above, OLS regressions are problematic because total spending contains the dependent variable and because it is subject to measurement error. Thus, we instrument total spending with measures of current income, recognizing that income is measured with error as well, particularly in the tails. Specifically, our instruments include dummy variables indicating the household's income category and a continuous variable accounting for the log of real after-tax income. We have dropped households in the top and bottom 5 percent of the income distribution.

Our specification in column (1) includes controls for the median and the ratio of the 95th percentile to the 5th percentile ($p95/p05$) of the distribution of log visible consumption. It also includes all fixed effects considered in equation (5). 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.

Despite the battery of fixed effects, a potential concern about the results in column (1) is that there may still be some factor that mechanically affects visible spending and is correlated with either the median or the dispersion of permanent incomes. For example, differences in housing prices, both in the cross-section and over time, may be correlated with average permanent incomes and also have a differentiated effect across age-groups within regions. When housing prices are high, conditional on permanent income, households in one age group may spend more on housing, and less on other goods, than households in a different age group. Accordingly, the systematic sorting by age into states with different housing prices may induce a spurious correlation between visible inequality and household visible spending.

To address the possibility that housing prices may be driving our results through differential effects on age groups, in column (2) 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.⁸ Note that, holding the distribution of visible spending in the reference group constant, an increase in individual housing expenditures reduces visible spending. The estimated effects of the distribution in group spending in columns (1) and (2) are similar.

Although the results in column (2) suggest that it is unlikely that house prices are driving our results, it is possible that they are driven by correlated state-level variation in other factors, such as relative prices of other goods, or redistributive taxation. To account for potentially confounding sources of state-level variation over time, we replace region and region-year interactions with state and state-year interactions in column (3). The potentially confounding sources likely have some differentiated effect across age groups, but we expect that the fixed effects for state-year interactions will absorb much of the variation due to unobserved variables affecting spending at the state level over time for both age groups.

In column (4) of Table 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 spending 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 spending show no significant effect on household’s visible spending in this case. This suggests that the naive OLS estimates are biased, masking the effect of visible inequality on individual spending.

4 Impact of visible inequality over time

Next, we explore whether the effect of visible inequality on household spending has changed over time. After all, the estimates in Table 1 refer to a relatively long period of time over which many things, including technology, have changed dramatically. In particular, both the visibility of certain consumption goods and people’s preferences over them are likely to have changed. We consider three different periods: 1984-1999, 2000-2006, and 2010-2018. The first period, is characterized by an increasing use of personal computers at home (up to 50% of households) and internet connections (up to 40%). However, given the absence of social media platforms, and given the still relative low levels of uptake, it was unlikely that technology had an effect on consumption. The period 2000-2006 is characterized by further growth in household computers and internet connection (up to 70% and 60% of households, respectively), coupled with

⁸We use the Federal Housing Finance Agency (HFPA) House Price Index, a broad measure of changes in single-family home values at state level and merge it with the CE data. <https://www.fhfa.gov/DataTools/Downloads/Pages/House-Price-Index.aspx>.

budding social platforms.⁹ The period 2010-2018 is characterized by the ubiquitous presence of computers and mobile devices as well as gigantic social media platforms that are extensively used.¹⁰ We exclude the period 2007-2009 because the Great Recession had a disproportionate impact on the consumption of many households.

Table 2 reports estimates of the effect of group visible spending interacted with different time periods using the specification of column (3) in Table 1. The estimates show 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.

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 consider this possibility in Table 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 4. Our results indicate that both of those effects are also present over the period 2010-2018, but not before.

Column (1) in Table 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 additional problem of how to control for permanent income. Here, again, we rely on our theoretical framework, which implies that both visible and non-visible spending are also good proxies for permanent income, once one controls for the distribution of visible consumption in the reference group. However, note that our simplifying assumptions imply that there are neither necessities nor luxuries in our model, as discussed above. By contrast, our

⁹Camille Ryan "Computer and Internet Use in the United States: 2016", American Community Survey Reports, Issued August 2018 ACS-3 <https://www.census.gov/history/pdf/comp-internetuse2016.pdf>).

¹⁰The first platforms developed in the late 1990s and early 2000s (Six Degrees (1997), Friendster (2001), Myspace (2003), LinkedIn (2002), Facebook (2004), Instagram (2010)). The iphone was launched in 2007.

empirical results indicate that visible consumption is a luxury, with an expenditure elasticity of 1.4 (see Table 3), whereas non-visible consumption is a necessity, with an expenditure elasticity of 0.9 (see column (1) in Table 4). These results also indicate that non-visible spending is much closer to a linear function of total spending. In turn, this suggests that non-visible spending is the more reliable proxy for total spending, provided that the relative price of non-visible consumption goods does not change substantially over time in a way that is not captured by the fixed effects.

Accordingly, we estimate a similar regression for total spending, where permanent income is proxied instead by non-visible spending, and similarly instrumented with indicators for income group and the log of after tax income. Coefficients shown in column (2) of the table indicate that increases in visible inequality tend to increase household total spending after the Great Recession, but not before then.

The hypothesis that our estimates of the effect of visible inequality on household spending reflect distortions associated with consumption externalities magnified by the rise of social media is intuitively plausible. The fact that our regressions include fixed effects for both state-year interactions and generation-year interactions, among others, is reassuring. However, our analysis so far treats the p50 and the ratio p95/p05 of log visible spending within the household's reference group as exogenous variables, although it is still possible that our estimates are driven by some other omitted factor. For example, the CE survey question on food away from home changed from a question about usual monthly spending to one about usual weekly spending, starting with the second quarter of 2007. Meyer and Sullivan (2023) note that the change resulted in a significant increase in reported spending on food away from home. In principle, our finding of significant effects of visible inequality on expenditure after 2010 might be driven by this change in reporting. Using data from 1964 to 2018, Meyer and Sullivan account for this by adjusting down consumption in food away from home after 2007. However, given our focus on visible spending, and our finding that the effect of visible inequality becomes evident only after 2010, we think it is important to rely on the better measures of spending on food away from home after 2010, which we expect to suffer less from underreporting. Accordingly, the results we have reported so far use the actual measures of spending on food away from home in the CE, without any adjustment.¹¹ Furthermore, in the rest of the paper we consider the period 2010-2018 in isolation, where the change in the survey question about food away from home is not a concern.

¹¹However, we have verified that our results remain qualitatively unchanged when we follow Meyer and Sullivan's adjustment procedure.

5 Impact of visible inequality after the Great Recession

In this section we turn our attention to the period after the Great Recession and show that the effect of visible inequality on household spending is robust to alternative specifications and alternative data sets.

Column (1) in Table 5 replicates our main specification in column (3) of Table 1, except that the sample period now is 2010-2018. The point estimates indicate that increasing within-group median visible spending by a magnitude equal to the average absolute value of annual changes over the sample increases visible spending of the average household by 2%. More interestingly, increasing the within-group p95/p05 ratio by a magnitude equal to the average absolute value of annual changes in this ratio also results in a 2% increase in the visible spending of the average household.¹²

In principle, despite the battery of fixed effects that are included in our regressions, and the fact that we control for the household’s permanent income, spending might still be affected by some factor omitted from the model. As a further robustness check, column (2) in Table 5 treats the p50 and the ratio p95/p05 of log visible spending within the household’s reference group, in addition to the household’s log spending — the proxy for her permanent income — as endogenous variables. These variables are instrumented with $\text{date-}t - 1$ values of *income* inequality, calculated using CPS data, and with $\text{date-}t - 2$ and $\text{date-}t - 3$ values of *visible* inequality, in addition to the indicators of the household’s income category plus the log of real after-tax income. Our choice of instruments tries to maximize the predictive power of the endogenous variables while minimizing the problem of weak instruments. The usual specification tests of weak instruments fail to reject the model at conventional levels of significance. The direct effect of visible inequality on the households’ visible spending becomes somewhat smaller, whereas the effect of the p50 of log visible spending in the reference group (which includes indirect effects of past visible inequality) becomes somewhat larger. Both effects remain significant, with comparable magnitudes.

The main purpose of the above IV strategy is to address any contemporaneous correlated effects that may confound the identification of the effects of true consumption externalities. To see why, note that aggregating equation (4) across households (indexed by j) within each

¹²Annual changes in within-group median visible spending have an average absolute value, across groups and over time, of 10%. Annual changes in the within-group p95/p05 ratio have an average absolute value, across groups and over time, of 0.05 units, which amounts to 4% of the average p95/p05 ratio.

reference group (indexed by grt) separately, and iterating through time, implies that

$$\begin{aligned} E[\ln v_{jgrt}] &= E[X_{jgrt}\beta] + \alpha_1 E[\ln v_{jgrt-1}] + \alpha_2 \text{Var}[\ln v_{jgrt-1}] + \phi E[\ln Y_{jgrt}] + E[\varepsilon_{jgrt}] \\ &= \sum_{s=0}^t \alpha_1^s \left\{ E[X_{jgrt-s}\beta] + \alpha_2 \text{Var}[\ln v_{jgrt-s-1}] + \phi E[\ln Y_{jgrt-s}] + E[\varepsilon_{jgrt-s}] \right\}, \end{aligned}$$

for all $t \geq 1$, where the operators E and Var denote the expectation and the variance, respectively, taken over all households j within the same generation g , region r and time t . As expected, mean log visible spending at any date $t \geq 1$ is a function of current and past values of both average permanent incomes and average household characteristics of the group. However, it is also a function of past values of the variance of log visible spending. Thus, to the extent that visible inequality influences spending behavior, average visible spending of the group does not only reflect average permanent incomes, but also visible inequality. Interestingly, this observation implies that the impact of average visible spending of the group on individual behavior cannot be identified separately from the impact of visible inequality. It also implies that past values of the dispersion of log visible spending may be good instruments of current values of average log visible spending within a given reference group. For example, our assumption is that the instrument $\text{Var}[\ln v_{jgrt-2}]$ is correlated with $E[\ln v_{jgrt-1}]$ but not with $E[X_{jgrt}\beta]$ and $E[X_{jgrt-1}\beta]$. In this sense, past values of visible inequality help disentangle the effects of average group behavior from the effects of average group characteristics.

Note that our IV strategy also helps addressing potential concerns that the p50 of log visible spending might be affected by non-classical measurement error. If such concerns arise from underreporting at the top of the income and expenditure distribution, then it is unlikely to affect the p50 of log visible spending. Otherwise, note that our instruments (p95/p05 ratios) are less likely to be affected by measurement error if the decline over time in the reporting of the components of spending is constant across income for each component, as suggested by Meyer and Sullivan (2023).

Similarly, equation (4) implies that the cross-sectional variance of log visible spending within a given reference group at any date t can be written as a function of the cross-sectional variance of household characteristics, the cross-sectional variance of permanent incomes, and the covariance between household characteristics and permanent incomes within the group:

$$\text{Var}[\ln v_{jgrt}] = \text{Var}[X_{jgrt}\beta] + \phi^2 \text{Var}[\ln Y_{jgrt}] + 2\text{Cov}[X_{jgrt}\beta, \phi \ln Y_{jgrt}] + \text{Var}[\varepsilon_{jgrt}].$$

Additionally, note that common econometric specifications of the income process suggest that current-income inequality can be a good instrument of permanent-income inequality.¹³ For

¹³See Meghir and Pistaferri (2011) for a review of alternative characterizations of the income process and

example, consider the standard specification of income as the sum of a random walk process and a transitory i.i.d. component:

$$\begin{aligned}\ln y_{hgrt} &= \ln Y_{hgrt} + \psi_{hgrt}, \\ \ln Y_{hgrt} &= Z_{hgrt}\beta + \ln Y_{hgrt-1} + \eta_{hgrt},\end{aligned}$$

where y_{hgrt} and Y_{hgrt} are transitory and permanent incomes, respectively; Z_{hgrt} is a vector of household characteristics that influence the growth of income; ψ_{hgrt} is the transitory shock in log current income, and η_{hgrt} is the shock to log permanent income. In this example, we have:

$$\text{Var} [\ln y_{hgrt}] = \text{Var} [\ln Y_{hgrt}] + \Sigma_{gr},$$

where $\text{Var} [\psi_{hgrt}] = \Sigma_{gr}$, for all t .

The above arguments suggest that the p95/p05 ratio of reference-group income at date $t - 1$ may be a good instrument of the p95/p05 ratio of log visible spending at date $t - 1$. The assumption is that the p95/p05 ratio of income at date $t - 1$ is correlated with the p95/p05 ratio of log visible spending at date $t - 1$, but uncorrelated with $\text{Var} [X_{hgrt-1}]$.

For our measures of income inequality, we use income data from the Current Population Survey (CPS) because it is well known that the CPS has better measures of income at the top of the distribution (Sabelhaus et al., 2015). It should be noted, however, that income in the CPS has been shown to be under-reported at the very bottom percentiles of income, and the extent of underreporting has increased over time (Meyer, Mok and Sullivan, 2015). Relying on the p95/p05 ratio, rather than the variance, serves to mitigate this concern.

Our key identifying assumption is that the standard errors in our regression are not systematically related to the effect of visible inequality. Identification comes from the fact that if group visible spending varies across generation-region groups and over time, then this variation is not driven by contemporaneous effects from some omitted factor. Our identification strategy relies on a battery of fixed effects, and the facts that we control for permanent income, and we instrument both measures of visible group behavior, in addition to the households' total spending.

It is useful to note the strengths and the weaknesses of our approach. First, the exclusion restrictions would be violated if $E[X_{hgrt-s}]$ or $\text{Var}[X_{hgrt-s}]$ are correlated with ε_{hgrt} , for some $s \geq 0$. Our IV strategy makes it unlikely that this is the case for $s = 0$ and $s = 1$.

Second, note that local tax policies would have to affect individual spending via a channel other than permanent incomes for them to have a confounding effect.

Third, confounding effects associated with relative price changes would have to be such that their relationship to consumption behavior.

their differential effect on the two age groups within a region is correlated with our instruments.

Fourth, note that individuals are not exogenously assigned to reference groups in our regressions. While the endogenous formation of reference groups is a serious concern in empirical research on conspicuous consumption in general, we believe it is unlikely that systematic sorting may drive our estimates of the effect of visible inequality on visible spending. Here, the concern is that, even though age is an exogenous characteristic, geographical location is not. It is possible that certain high visible-spending individuals in a particular age group may choose to locate in states with high average visible spending, or high visible inequality, within their age group. With respect to this, our focus on individuals younger than 64 serves to minimize concerns associated with the systematic sorting of retirees across states according to their discount factors, which in turn may be associated with their health status. Similarly, our focus on individuals older than 25 serves to minimize concerns associated with the systematic sorting of students across states according to their family wealth. Furthermore, our specification of reference groups at the broad region level serves to mitigate geographical sorting problems more generally.

In column (3) of Table 5 we replicate the regression in column (1) using data from the 2011-2019 waves of the PSID, and find that the estimated effects of visible inequality are similar in magnitude. Although the PSID collects less detailed information on consumption expenditures, coverage since 2005 includes 95 percent of CE spending (Andresky et al., 2014). Additionally, PSID measures of income and expenditures at the top of the distribution are reputedly less susceptible to measurement error. In columns (4) and (5) we exploit the panel structure of the PSID as a further robustness check. Specifically, a potential concern is that visible inequality across reference groups is correlated with unobserved heterogeneity regarding household preferences over consumption. For instance, if some individuals display more ostentatious behavior or have more expensive tastes than others, and if households with these preferences cluster together, our estimates of the impact of inequality on consumption might be biased. With respect to this, controlling for household fixed effects removes time-invariant heterogeneity of household preferences.

The specification in column (4) of Table 5 differs from that in column (3) only in that it includes household fixed effects. The magnitude of the estimated effect of visible inequality remains non-negligible. Not surprisingly, however, it is noisily estimated, both because of the variation absorbed by the household fixed effect and because of the loss of observations (those households that appear only once in the panel). Alternatively, in column (5) we replace household fixed effects with family fixed effects. This allows us to account for time-invariant characteristics, such as tastes or cultural preferences, that run in families — households with the same family identifier — while being less restrictive and keeping more observations. The

estimated effect of median log visible spending becomes negligible. However, estimates of the effect of visible inequality using family fixed effects remain large in magnitude, although they are not significant.

Table 6 replicates the five specifications in Table 5, except that the dependent variable is log non-visible spending, instead of log visible spending. Column (1) in Table 6 replicates the specification in column (1) of Table 5, by treating both the p50 and the ratio p95/p05 of log visible spending within the household's reference group as exogenous. We find that increasing visible inequality reduces spending in non-visible consumption.

Column (2) treats both the p50 and the ratio p95/p05 of log visible spending within the household's reference group, in addition to the proxy for the household's permanent income, as endogenous variables. These variables are instrumented exactly as it was done in column (2) of Table 5. The usual specification tests of weak instruments fail to reject the model at conventional levels of significance. We continue to find no effect of average log visible spending of the group. The effect of visible inequality, however, remains significant at the 10 percent confidence level, with a comparable magnitude.

We turn to the PSID again as a final robustness check. Columns (3)-(5) of Table 6 replicate the regressions in the last three columns of Table 5, except that the dependent variable is log non-visible spending, instead of log visible spending. Column (3) of Table 6 replicates the regression in column (1) using data from the 2011-2019 waves of the PSID, column (4) includes household fixed effects, and column (5) replaces household fixed effects with family fixed effects. Using data from the PSID and from the CE produces similar results. In particular, the estimated effect of visible inequality on household non-visible spending remains large and significant when family fixed effects are included in the regression. The only discrepancy is that the estimates of the elasticity of non-visible spending with respect to total spending are statistically significantly below one in the CE sample and above one in the PSID sample, given the precision of the estimates. We return to this issue below, when we discuss a similar discrepancy in the estimates reported in Table 7.

Next, in Table 7 we report estimates of the effect of visible inequality on total household spending. Since the dependent variable is total household spending, a preliminary question is how to proxy for the household permanent income. To address this question, note that, unlike the expenditure elasticity of visible spending (Table 5), the expenditure elasticity of non-visible spending (Table 6) is fairly close to one. In that sense, non-visible spending is a good proxy for total spending, provided that the relative price of non-visible consumption goods does not change substantially over time in a way that is not captured by the fixed effects. Accordingly, all specifications in Table 7 proxy the household's permanent income with the log of non-visible spending.

As in Table 5 and Table 6, column (1) in Table 7 treats both the p50 and the ratio p95/p05 of log visible spending within the household’s reference group as exogenous variables. Column (2) treats both group variables, in addition to the proxy for the household’s permanent income, as endogenous variables, using the same instruments as in the previous two tables. As before, the usual specification tests of weak instruments fail to reject the model at conventional levels of significance. We find no effect of median log visible spending of the group. The effect of visible inequality on total spending, just like its effect on non-visible spending, remains significant at the 10 percent confidence level, with a comparable magnitude. As a final robustness check, we turn to the PSID once more (columns (4)-(6) in Table 7) and find that the estimated effect of visible inequality on total household spending remains large and significant even when family fixed effects are included in the regression.

Note that the estimates of the elasticity of total spending with respect to our proxy of permanent income (non-visible spending) are statistically significantly above one in the CE sample and below one in the PSID sample, given the precision of the estimates. However, all estimated elasticities are close to one, which suggest that total spending is close to linear in permanent income in both samples. By contrast, Straub (2019) argues that a log-linear relationship between consumption and permanent income is a good fit to data from the PSID, but estimates an elasticity around 0.7, which indicates significant concavity. Straub (2019) uses data from 1999 to 2013 and focuses on the consumption expenditures that are available every year since 1999, covering around 70 percent of the categories covered in the CE and excluding the categories that were added to the PSID in 2005. Instead, we use PSID data over the period 2010-2018 and our measure of total spending accounts for most of the expenditures considered in the CE. Specifically, we include information on all consumption categories that were added in 2005, including telecommunication, home repairs and maintenance, household furnishings and equipment, clothing and apparel, trips and vacations, and recreation and entertainment (Andreski et al., 2014). However, if we subtract these categories from our definition of total spending (the dependent variable) and non-visible spending (the proxy for permanent income), we find an expenditure elasticity of 0.74 in the PSID sample, which is comparable to the estimates reported in Straub (2019).¹⁴

The results reported in column (3) of Table 7 illustrate why empirical estimates of the effects of inequality on consumption behavior should be interpreted with reference to an explicit theoretical framework. Specifically, recall that our model of conspicuous consumption illustrates why equilibrium consumption behavior is a function of the distribution of permanent, as opposed to current, incomes across households in the same reference group. It should be clear that this is not just a feature of our model, but more generally a feature of models of relative consumption

¹⁴By contrast, the estimated elasticity in the CE sample is 1.08 when we subtract the same categories.

as long as they also recognize that consumption decisions are inherently intertemporal. By contrast, empirical analyses of conspicuous consumption routinely take as given that individual behavior responds to changes in the distribution of *current*, rather than permanent, incomes (e.g., Charles et al., 2009; Bertrand and Morse, 2016). For a comparison, in column (3) we continue to control for the household permanent income, but we use the CPS data to replace the median and the p95/p05 ratio of log visible spending with the median and the p95/p05 ratio of log *income*. For details on the sample and the measures of income, see Appendix A.

Although the estimates of the effects of group variables in column (3) are noisy, they suggest a strong *negative* relationship between total spending and income inequality. By contrast, columns (1) and (2) suggest a strong *positive* relationship between the same variables.¹⁵ From the lens of our theoretical framework, this sharp discrepancy suggests that the distribution of current income in the reference group is not a reliable proxy of group visible behavior. As illustrated by our basic model of conspicuous consumption, the reason is that optimal intertemporal decision-making implies that visible consumption inequality is the more reliable proxy of inequality in *permanent* incomes.

Consider the potential economic significance of our estimated effect of visible inequality on total household spending. The estimates in column (1) of Table 7 imply that increasing the within-group p95/p05 ratio by a magnitude equal to the average absolute value of annual changes in this ratio over the sample (see footnote 12) increases total spending for the average household by 0.7%. While this effect is small relative to average levels of spending, it is relatively large when compared with average growth rates. A 0.7% change in total spending for the average household within each of the eight groups we consider amounts, respectively, to 10%, 9%, 14%, 20%, 15%, 25%, 12%, and 12% of the average absolute value of actual growth rates in within-group average spending over the period 2010-2018.

The magnitudes of these predicted effects suggest that the impact of visible inequality on consumption smoothing might be non-trivial. To get a sense of economic magnitude, we consider a simple counterfactual exercise under the assumption that visible inequality remains constant at the 2010 level. The exercise suggests that annual changes in visible inequality increased the standard deviation of mean household expenditures for three of the reference groups, by 3.5%, 4.2% and 0.6%, respectively. For the other five groups annual changes in visible inequality decreased the standard deviation of mean household expenditures, by 7.6%, 2.9%, 4.8%, 2.2% and 3.6%, respectively. From the lens of our theoretical framework, these effects can be understood as arising from shocks to the distribution of permanent incomes.

Even more starkly, consider a counterfactual exercise in which households simply disregard

¹⁵Although not reported in Table 5 and Table 6, we find similar discrepancies in the effects of visible inequality versus income inequality on both visible and non-visible spending.

the distribution of visible spending within their group. Assuming that our estimated model is well specified, the estimates in column (1) in Table 7 imply that average household spending from 2010 to 2018 might have been up to 25% smaller in this counterfactual scenario without distortions from consumption externalities. This indicates a sizable distortionary effect of consumption externalities. However, it is worth noting that visible inequality over the period 2010-2018 fails to exhibit a noticeable trend for any of the reference groups we consider. In this sense, it is unlikely that visible inequality has influenced the increasing trend in household savings in the U.S. over this period.

6 Discussion and concluding remarks

While the Great Recession may have had some structural effects on consumption behavior, we conjecture 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 historically low interest rates in the United States in the aftermath of the Great Recession are well documented. The Federal funds effective rate fell below 1 percent in October 2008, and it remained below 1 percent until June 2017. The model in Section 3 illustrates the role of interest rates in conspicuous consumption. Households have an incentive to rise or lower their conspicuous consumption when visible inequality increases 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. 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 visible inequality.

Moreover, the rise of social media has likely had far reaching effects on household spending. In particular, it 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.¹⁶ 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).¹⁷

¹⁶<https://www.statista.com/statistics/253577/number-of-monthly-active-instagram-users/>

¹⁷<https://www.businesswire.com/news/home/20190513005203/en/>

It is unlikely that social media has had any significant effect on 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 because the Great Recession had a disproportionate effect on consumption expenditures during 2008 and 2009.

Others have noted that household spending on cars accounts for a large share of visible spending (Charles et al., 2009; Bertrand and Morse, 2016).¹⁸ Not surprisingly, spending on cars does play a role in our results. More interestingly, however, 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 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 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).¹⁹ 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 1 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

¹⁸Previous 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.

¹⁹A. Atanasova: “The Psychology of Foodstagramming”, Published Nov. 9, 2016: <https://www.socialmediatoday.com/social-networks/psychology-foodstagramming>.

log spending on each visible category (Panel B). Note that panels A and B of Figure 1 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 (Panel A). Note that the CE survey question on food away from home changed from a question about usual monthly spending to one about usual weekly spending, starting with the second quarter of 2007. Others have noted that the change in the survey question resulted in a significant increase in reported spending on food away from home (Meyer and Sullivan, 2023). Arguably, this drives the increase in 2007 and 2008. However, note that FAFH is the only category of visible consumption whose median spending did not fall significantly during the Great Recession and its aftermath (Panel B). Interestingly, median log spending on FAFH surpassed median log spending on vehicles in 2007 and remained relatively higher until 2018.²⁰

It is also interesting to note the trends throughout the distribution of FAFH spending, as illustrated in Figure 2. At the bottom 5 percent of the distribution, FAFH spending went from 4.79 log points in 2010 to 5.19 in 2018, an increase of 8 percent. During the same period mean log FAFH rose 1.6 percent, while the top 5th percentile of the distribution rose by less than 1 percent. In this sense FAFH spending contributed to moderate visible inequality over the period 2010-2018.²¹

Finally, we illustrate the potential role of FAFH spending in our results by considering the response of households' spending to the spending of others on all visible categories other than FAFH. Results of this exercise are shown in Table 7. Column (1) replicates the specification in column (1) of Panel B in Table 6, except that visible group behavior now refers to visible spending minus FAFH. Similarly, column (2) replicates the specification in column (4) of Panel B in Table 6, except that visible group spending now excludes FAFH. There is no longer an effect of group spending on the households' total spending, which suggest that spending on FAFH by others has indeed had an effect on household consumption via visible inequality after 2010.

Overall our analysis suggests that moderate changes in visible inequality can have non-trivial effects on household spending. From the lens of our theoretical framework, visible inequality

²⁰Attanasio 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.

²¹Spending in vehicles also contributed to moderate visible inequality over the period 2010-2018. It increased by 13 percent at the bottom 5th percentile and by 2 percent at the top 5th percentile of the distribution of vehicle spending. By contrast, spending in clothing declined by 2.7 percent at the top 5th percentile and by 9 percent at the bottom 5th percentile.

is a symptom of permanent income inequality. Accordingly, further research on conspicuous consumption and visible inequality might shed new light on the dependence of consumption behavior on the distribution of permanent incomes. Our analysis also suggests that social media may have magnified the effects of visible inequality significantly. With respect to this, further research linking consumption data and data on social media usage might shed new light on the effects of visible inequality, and on the nature of conspicuous consumption. Additionally, the above arguments suggest that policies designed to influence nutritional outcomes might benefit from a better understanding of the effects of consumption externalities on FAFH spending, and the substitutability between FAFH and food at home spending.

References

- [1] Abel, Andrew B. (1990): “Asset prices under habit formation and catching up with the Joneses,” *American Economic Review* 80, 38-42.
- [2] Aguiar, Mark and Mark Bilal (2015): “Has consumption inequality mirrored income inequality?” *American Economic Review* 105, 2725-2756.
- [3] Aguiar, Mark and Erik Hurst (2013): “Deconstructing life cycle expenditure,” *Journal of Political Economy* 121(3), 437-492.
- [4] Alvarez-Cuadrado, Francisco and Ngo Van Long (2011): “The relative income hypothesis,” *Journal of Economic Dynamics and Control* 35, 1489-1501.
- [5] Andreski, Patricia, Geng Li, Mehmet Zahid Samancioglu, and Robert Schoeni (2014): “Estimates of annual consumption expenditures and its major components in the PSID in comparison to the CE,” *American Economic Review: Papers & Proceedings* 104(5), 132-135.
- [6] Arrow, Kenneth J. and Partha S. Dasgupta (2009): “Conspicuous consumption, inconspicuous leisure,” *Economic Journal* 119, F497-F516.
- [7] Attanasio, Orazio P. and Luigi Pistaferri (2016): “Consumption inequality,” *Journal of Economic Perspectives* 30, 3-28.
- [8] Attanasio, Orazio P., Kieran Larkin, Morten O. Ravn and Mario Padula (2022): “(S)Cars and the Great Recession,” *Econometrica* 90, 2319-2356.
- [9] Attanasio, Orazio P. and Guglielmo Weber (2010): “Consumption and saving: Models of intertemporal allocation and their implications for public policy,” *Journal of Economic Literature* 48(3), 693-751.
- [10] Barauskaitė, D., J. Gineikiene, V. Auruskeviciene, B. Fennis, M. Yamaguchi, and N. Kondo (2018). "Eating healthy to impress: How conspicuous consumption, perceived self-control motivation, and descriptive normative influence determine functional food choices" *Appetite* 131(1) . 10.1016/j.appet.2018.08.015.
- [11] Bee, Adam, Bruce D. Meyer and James X. Sullivan (2015): “The validity of consumption data: Are the Consumer Expenditure interview and diary surveys informative?” in *Improving the Measurement of Consumer Expenditures*, Carroll, Christopher D., Thomas F. Crossley, and John Sabelhaus (eds), University of Chicago Press, chapter 7.
- [12] Bertrand, Marianne and Adair Morse (2016): “Trickle-down consumption,” *Review of Economics and Statistics* 98(5), 863-879.
- [13] Charles, Kerwin K., Erik Hurst and Nikolai Roussanov (2009): “Conspicuous consumption and race,” *Quarterly Journal of Economics*, 124(2), 425-467.
- [14] Corneo, Giacomo and Olivier Jeanne (1998): “Social organization, status, and savings behavior,” *Journal of Public Economics* 70, 37-51.

- [15] De Giorgi, Giacomo, Anders Frederiksen and Luigi Pistaferri (2020): "Consumption network effects," *Review of Economic Studies* 87, 130-163.
- [16] Deaton, Angus and Christina Paxson (1994): "Intertemporal choice and inequality," *Journal of Political Economy* 102(3), 437-467.
- [17] Duesenberry, James S. (1949): *Income, Saving and the Theory of Consumption Behavior*, Harvard University Press, Cambridge MA.
- [18] Elitzak, H. and A.M. Okrent (): "A Retrospective of Food-Away-From-Home Expenditures From 1987 to 2017" in *America's Eating Habits: Food Away from Home.*" EIB-196 USDA, Economic Research Service
- [19] TVO today, "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) by Joshna Maharaj <https://www.tvo.org/article/conspicuous-consumption-why-are-we-photographing-our-food-instead-of-enjoying-it>
- [20] Frank, Robert H., Adam S. Levine and Oege Dijk (2014): "Expenditure cascades," *Review of Behavioral Economics* 1, 55-73.
- [21] Friedman, Milton (1957): *A theory of the Consumption Function*, Princeton University Press, Princeton NJ.
- [22] Grinblatt, Mark, Matti Keloharju and Seppo Ikaheimo (2008): "Social influence and consumption: evidence from the automobile purchases of neighbors," *Review of Economics and Statistics* 90, 735-753.
- [23] Heffetz, Ori (2011): "A test of conspicuous consumption: visibility and income elasticities," *Review of Economics and Statistics* 93, 1101-1117.
- [24] Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn (2011): "The effects of lottery prizes on winners and their neighbors: evidence from the Dutch postcode lottery," *American Economic Review* 101, 2226-2247.
- [25] Luttmer, Erzo F.P. (2005): "Neighbors as negatives: relative earnings and well-being," *Quarterly Journal of Economics* 120, 963-1002.
- [26] Manski, Charles F. (1993): "Identification of endogenous social effects: the reflection problem," *Review of Economic Studies* 60(3), 531-542.
- [27] Maurer, Jurgen and André Meier (2008): "Smooth it like the 'Joneses'? Estimating peer-group effects in intertemporal consumption choice," *The Economic Journal* 118, 454-476.
- [28] Meghir, Costas and Luigi Pistaferri (2011): "Earnings, consumption and life cycle choices," *Handbook of Labor Economics*, Volume 4b, chapter 9, 773-854.
- [29] Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan (2015): "Household surveys in crisis," *Journal of Economic Perspectives* 29(4), 199-226.

- [30] Meyer, Bruce D. and James X. Sullivan (2023): “Consumption and income inequality in the United States since the 1960s,” *Journal of Political Economy* 131(2), 247-284.
- [31] Modigliani, Franco and Richard Brumberg (1954): “Utility analysis and the consumption function: an interpretation of cross-section data,” in Kurihara K.K. (ed.), *Post-Keynesian Economics*, Rutgers University Press, New Brunswick, NJ, 388–436.
- [32] Ray, Debraj and Arthur Robson (2012): “Status, intertemporal choice and risk taking,” *Econometrica* 80, 1505-1531.
- [33] Sabelhaus, John, David Johnson, Stephen Ash, David Swanson, Thesia Garner, John Greenlees and Steve Henderson (2015): “Is the Consumer Expenditure Survey representative by income?” in *Improving the Measurement of Consumer Expenditures*, Carroll, Christopher D., Thomas F. Crossley, and John Sabelhaus (eds), University of Chicago Press, chapter 8.
- [34] Straub, Ludwig (2019): “Consumption, savings, and the distribution of permanent income,” unpublished manuscript, Harvard University.
- [35] US Department of Agriculture, (2018) "America's Eating Habits: Food Away From Home", Economic Information Bulletin No.196 (EIB-196).
- [36] Veblen, Thorstein B. (1899): *The Theory of the Leisure Class: An Economic Study of Institutions*, Modern Library, New York.
- [37] Xia, Bianjun (2010): “Status, inequality and intertemporal choice,” *The B.E. Journal of Theoretical Economics* 10,

Table 1: Impact of visible inequality on visible spending, 1984-2018

	(1)	(2)	(3)	(4)
IV ln(total spending)	1.407*** (0.038)	2.052*** (0.154)	1.434*** (0.037)	
ln(housing spending)		-0.641*** (0.147)		
p50 of ln(visible spending)	0.021 (0.012)	0.028** (0.011)	0.013 (0.013)	0.016 (0.017)
p95/p05 of ln(visible spending)	0.136** (0.045)	0.084** (0.025)	0.126** (0.037)	-0.007 (0.064)
household controls	Yes	Yes	Yes	Yes
generation FE \times region FE	Yes	Yes	Yes	Yes
generation FE \times year FE	Yes	Yes	Yes	Yes
generation FE \times region FE \times time trend	Yes	Yes	Yes	Yes
generation, region and year FEs	Yes	Yes		
region FE \times year FE	Yes	Yes		
generation, state and year FEs			Yes	Yes
state FE \times year FE			Yes	Yes
observations	69,091	64,258	64,245	64,245
R ²	0.621	0.706	0.624	0.271

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 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). Sample years 1984-2018.

Columns (1) to (3) control for log total spending of the household (instrumented with indicators for the household's income group and the log of after-tax income). Column (2) 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 log of total spending. All IV specifications have first-stage F statistics that fail to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

*** p<0.01, ** p<0.05, * p<0.1

Table 2: Impact of visible inequality on visible spending over time

IV $\ln(\text{total spending})$	1.444*** (0.038)
p50 of $\ln(\text{visible spending}) \times (1984-1999)$	-0.03 (0.029)
p50 of $\ln(\text{visible spending}) \times (2000-2006)$	-0.187** (0.058)
p50 of $\ln(\text{visible spending}) \times (2010-2018)$	0.151*** (0.039)
p95/p05 of $\ln(\text{visible spending}) \times (1984-1999)$	0.002 (0.077)
p95/p05 of $\ln(\text{visible spending}) \times (2000-2006)$	0.06 (0.039)
p95/p05 of $\ln(\text{visible spending}) \times (2010-2018)$	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 socio-demographic controls as in column (3) in Table 1. Total household expenditure instrumented with indicators for the household's income group and the log of after-tax income.

Years 2007-2009 removed from the sample.

First-stage F statistic fails to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Impact of non-visible inequality on visible spending over time

IV $\ln(\text{total spending})$	1.444*** (0.038)
p50 of $\ln(\text{non-visible spending}) \times (1984-1999)$	0.124 (0.077)
p50 of $\ln(\text{non-visible spending}) \times (2000-2006)$	0.206** (0.081)
p50 of $\ln(\text{non-visible spending}) \times (2010-2018)$	0.338*** (0.065)
p95/p05 of $\ln(\text{non-visible spending}) \times (1984-1999)$	-0.213 (0.334)
p95/p05 of $\ln(\text{non-visible spending}) \times (2000-2006)$	0.040 (0.374)
p95/p05 of $\ln(\text{non-visible spending}) \times (2010-2018)$	-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 same fixed effects and socio-demographic controls as in column (3) in Table 1. Total household expenditure instrumented with indicators for the household's income group and the log of after-tax income.

Years 2007-2009 removed from the sample.

First-stage F statistic fails to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

** $p < 0.01$, * $p < 0.05$, $p < 0.1$.

Table 4: Impact of visible inequality on non-visible and total spending over time

Dependent variable	(1) ln(non-visible spending)	(2) ln(total spending)
IV ln(total spending)	0.901*** (0.014)	
IV ln(non-visible spending)		1.110*** (0.017)
p50 of ln(visible spending) \times (1984-1999)	-0.007 (0.019)	0.008 (0.022)
p50 of ln(visible spending) \times (2000-2006)	0.110*** (0.016)	-0.122*** (0.017)
p50 of ln(visible spending) \times (2010-2018)	-0.012 (0.023)	0.013 (0.026)
p95/p05 of ln(visible spending) \times (1984-1999)	-0.037 (0.029)	0.041 (0.033)
p95/p05 of ln(visible spending) \times (2000-2006)	-0.035 (0.029)	0.039 (0.032)
p95/p05 of ln(visible spending) \times (2010-2018)	-0.111*** (0.014)	0.124*** (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. Dependent variable in column (2) is the natural log of total spending. This regression replaces 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 same fixed effects and socio-demographic controls as in column (3) in Table 1. Years 2007-2009 removed from the sample.

Both specifications have first-stage F statistics that fail to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Impact of visible inequality on visible spending after the Great Recession

	(1)	(2)	(3)	(4)	(5)
	CE sample 2010-2018		PSID sample 2011-2019		
IV ln(total spending)	1.403*** (0.052)	1.403*** (0.052)	1.282*** (0.034)	1.532*** (0.180)	1.367*** (0.046)
p50 of ln(visible spending)	0.190*** (0.038)		-0.214 (0.166)	-0.236 (0.192)	-0.054 (0.164)
p95/p05 of ln(visible spending)	0.420*** (0.067)		0.404* (0.194)	0.277 (0.227)	0.348 (0.201)
IV p50 of ln(visible spending)		0.349** (0.145)			
IV p95/p05 of ln(visible spending)		0.202* (0.107)			
household fixed effects				Yes	
family fixed effects					Yes
observations	16, 150	16, 150	29, 730	27, 609	29, 351
R ²	0.588	0.588	0.614	0.203	0.520

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.

Column (1) replicates our main specification in column (3) of Table 1, except that the sample period now is 2010-2018.

Column (2) replicates column (1) instrumenting (date- t) ln(total spending), (date- $t - 1$) p50 of ln(visible spending) and (date- $t - 1$) p95/p05 of ln(visible spending) with date- t indicators of the household's income category plus the log of real after-tax income, with date- $t - 1$ values of the p95/p05 of ln(CPS income), and with date- $t - 2$ and date- $t - 3$ values of the p95/p05 of ln(visible spending).

Column (3) replicates column (1) using PSID data for the sample years 2011-2019.

Column (4) replicates column (3) including household fixed effects.

Column (5) replicates column (4) replacing household fixed effects with family fixed effects.

All specifications have first-stage F statistics that fail to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 6: Impact of visible inequality on non-visible spending after the Great Recession

	(1)	(2)	(3)	(4)	(5)
	CE sample 2010-2018		PSID sample 2011-2019		
IV ln(total spending)	0.944*** (0.010)	0.944*** (0.010)	1.136*** (0.016)	1.350*** (0.118)	1.124*** (0.022)
p50 of ln(visible spending)	-0.021 (0.024)		-0.091 (0.133)	-0.235* (0.114)	-0.138 (0.107)
p95/p05 of ln(visible spending)	-0.130*** (0.020)		-0.219* (0.104)	-0.163 (0.089)	-0.255** (0.086)
IV p50 of ln(visible spending)		-0.069 (0.084)			
IV p95/p05 of ln(visible spending)		-0.206* (0.102)			
household fixed effects				Yes	
family fixed effects					Yes
observations	16,150	16,150	29,730	27,609	29,351
R ²	0.804	0.804	0.808	0.462	0.739

Note: Dependent variable in all columns is the natural log of non-visible spending, defined as the sum of total spending, minus visible spending, minus shelter.

Column (1) replicates column (1) in Table 5, but with log non-visible spending as the dependent variable.

Column (2) replicates column (1) instrumenting (date- t) ln(total spending), (date- $t - 1$) p50 of ln(visible spending) and (date- $t - 1$) p95/p05 of ln(visible spending) with date- t indicators of the household's income category plus the log of real after-tax income, with date- $t - 1$ values of the p95/p05 of ln(CPS income), and with date- $t - 2$ and date- $t - 3$ values of the p95/p05 of ln(visible spending).

Column (3) replicates column (1) using PSID data for the sample years 2011-2019.

Column (4) replicates column (3) including household fixed effects.

Column (5) replicates column (4) replacing household fixed effects with family fixed effects.

All specifications have first-stage F statistics that fail to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 7: Impact of visible inequality on total spending after the Great Recession

	(1)	(2)	(3)	(4)	(5)	(6)
	CE sample 2010-2018			PSID sample 2011-2019		
IV ln(non-visible spending)	1.059*** (0.011)	1.059*** (0.011)	1.059*** (0.011)	0.876*** (0.013)	0.727*** (0.075)	0.886*** (0.018)
p50 of ln(visible spending)	0.022 (0.026)			0.078 (0.117)	0.174* (0.080)	0.122 (0.096)
p95/p05 of ln(visible spending)	0.138*** (0.021)			0.192* (0.090)	0.119 (0.066)	0.226** (0.074)
IV p50 of ln(visible spending)		0.074 (0.088)				
IV p95/p05 of ln(visible spending)		0.219* (0.107)				
p50 of ln(CPS income)			0.060 (0.057)			
p95/p05 of ln(CPS income)			-0.147 (0.425)			
household fixed effects					Yes	
family fixed effects						Yes
observations	16,150	16,150	16,150	29,730	27,609	29,351
R ²	0.750	0.750	0.750	0.771	0.449	0.674

Note: Dependent variable in all columns is the natural log of total spending. Sample years 2010-2018.

Column (1) replicates column (1) in Table 5, and in Table 6, but now with log total spending as the dependent variable, replacing the control for household total spending with a control for the household's non-visible spending, instrumented with indicators for the household's income group and the log of after-tax income.

Column (2) replicates column (1) instrumenting (date- t) ln(non-visible spending), (date- $t - 1$) p50 of ln(visible spending) and (date- $t - 1$) p95/p05 of ln(visible spending) with date- t indicators of the household's income category plus the log of real after-tax income, with date- $t - 1$ values of the p95/p05 of ln(CPS income), and with date- $t - 2$ and date- $t - 3$ values of the p95/p05 of ln(visible spending).

Column (3) replaces the measures of the distribution of ln(visible spending) used in column (1) with measures of the distribution of current income in the reference group, calculated from CPS data.

Column (4) replicates column (1) using PSID data for the sample years 2011-2019.

Column (5) replicates column (4) including household fixed effects.

Column (6) replicates column (5) replacing household fixed effects with family fixed effects.

All specifications have first-stage F statistics that fail to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 8: Impact of group FAFH expenditures on total spending

	(1)	(2)
	CE sample 2010-2018	PSID sample 2011-2019
IV ln(non-visible spending)	1.059*** (0.010)	0.876*** (0.012)
p50 of ln(visible spending minus FAFH)	-0.009 (0.021)	-0.091 (0.115)
p95/p05 of ln(visible spending minus FAFH)	-0.003 (0.030)	0.048 (0.061)
observations	16, 150	29, 730
R ²	0.750	0.771

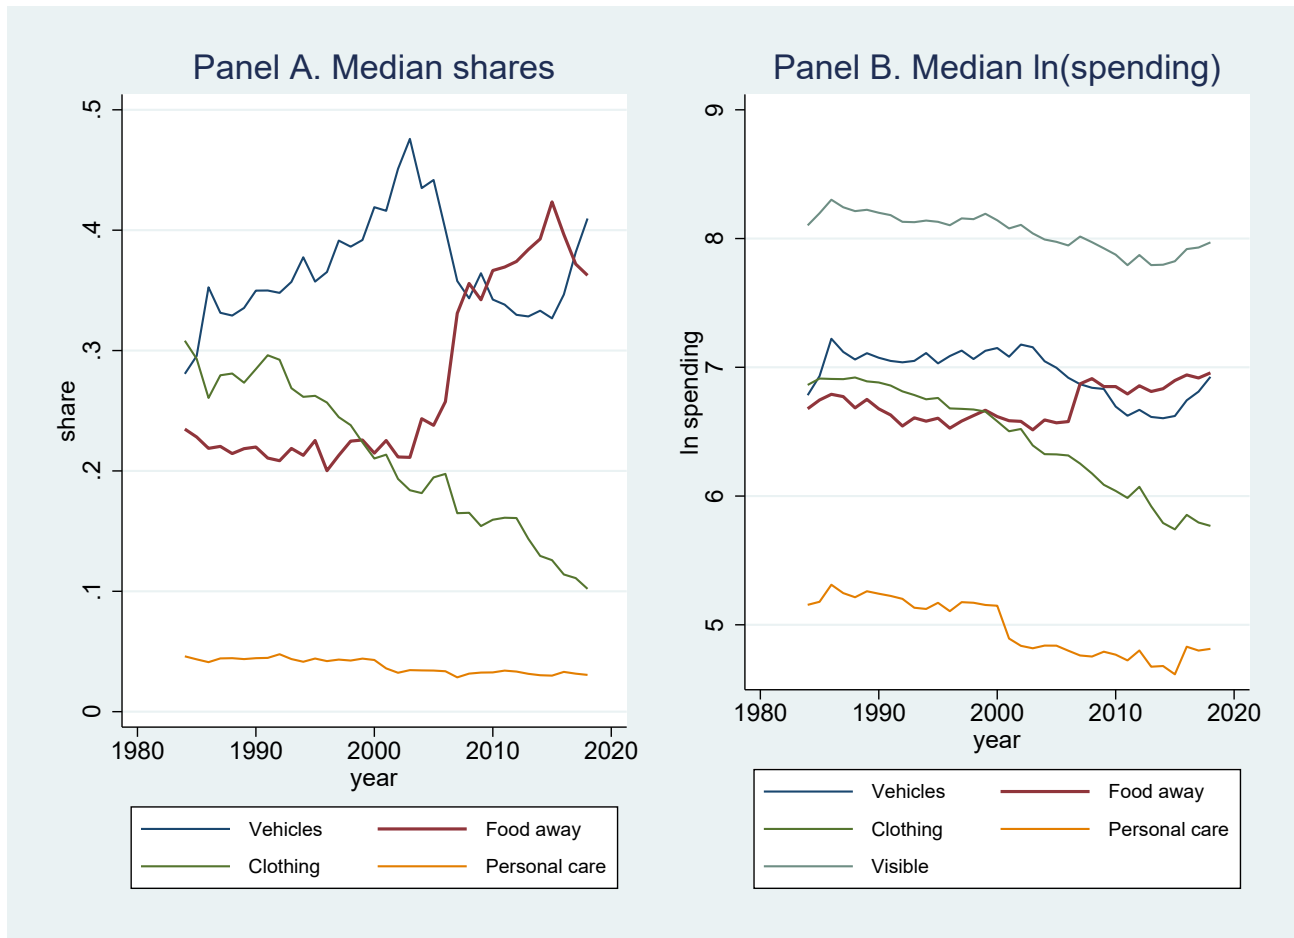
Note: Column (1) uses CE data for 2010-2018. Column (2) uses PSID data for 2011-2019. Dependent variable in both columns is total spending. Visible spending is the sum of personal care, clothing, vehicle and FAFH spending; non-visible spending is total spending minus visible spending minus shelter.

All columns include same fixed effects and socio-demographic controls as in column (3) in Table 1. Non-visible spending is instrumented with indicators for the household's income group and the log of after-tax income.

All specifications have first-stage F statistics that fail to reject the model at conventional levels of significance. Standard errors (in parentheses) are corrected for heteroskedasticity and clustered at the generation-region level.

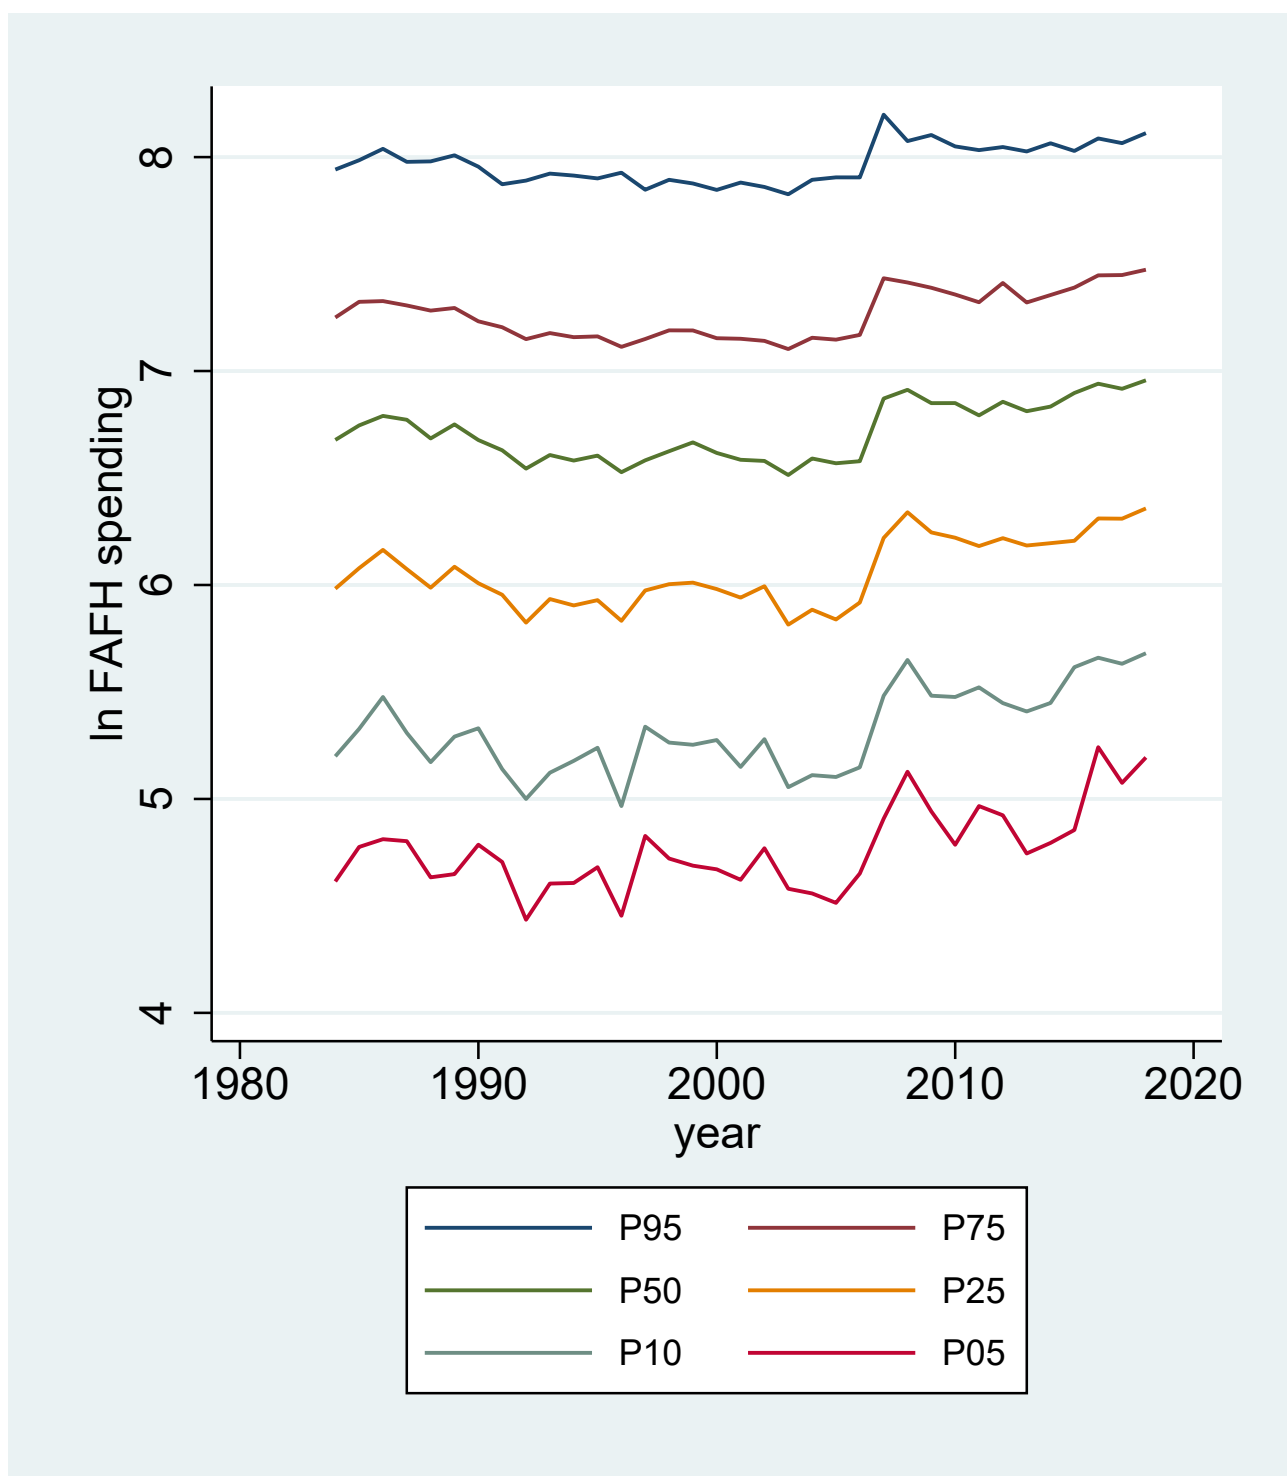
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1: Visible spending by category.



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

Figure 2: Distribution of $\ln(\text{FAFH spending})$.



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

Appendix A

In this Appendix we describe the data we use.

CE data

Our main data come from the interview sample of the Consumer Expenditure Survey (CE) of the U.S. Bureau of Labor Statistics. We use CE data from the survey years 1984-2018. Following an initial bounding interview, each household is interviewed up to four times over four consecutive calendar quarters before it is replaced by a new household unit. Each of the four interviews collects detailed information about expenditures over the previous three months. Households are asked about their annual income in the last interview. Approximately one-fourth of the households contacted each quarter are new to the survey in order to keep a balanced panel across interview quarters.

The sample restrictions we impose follow Aguiar and Bils (2015). We include only households who completed all four interviews and reported at least one source of income or benefits, which reduces our sample by approximately half of the observations. We aggregate expenditures for each household across the four interviews, and so each household appears only once in the sample. 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. 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. The fraction of households top coded on income fluctuates from about one to just over four percent across survey waves. Accordingly, we also eliminate households with before-tax income over the 95th and below the 5th percentile, 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 CE survey weights through the analysis to make our results representative of the national population.

We start by aggregating to the 20 expenditure categories considered in Aguiar and Bils (2015), which are listed in Table A1. We then further aggregate all categories other than housing to two broad categories, visible versus non-visible spending, as explained in the text. Their definitions of expenditures by good follow closely the definitions in the CE with a few exceptions. For renters, they define housing by rent paid. However, for home owners, they subtract spending on mortgage interest, property taxes, maintenance, repairs, home insurance, and other expenses, and add the self-reported rental equivalent of the home. For expenditures on vehicles, their measures differ from the CE measures in that they subtract the value of used vehicles that are sold by a household. Furthermore, the reported expenditures on food at home in surveys conducted after 1987 are noticeably larger due to a change in the wording for the question regarding spending on food at home. To correct for the effect of this change, Aguiar and Bils adjust food at home expenditures upward by 11 percent for the years before 1988.

Our measures of before-tax income are also taken from Aguiar and Bils (2015), without further adjustment. They subtract personal taxes, which include federal, state and local income

taxes, to arrive at a measure of after-tax income. We use this measure of after-tax income to proxy for the household permanent income.

The CE survey question on food away from home changed from a question about usual monthly spending to one about usual weekly spending, starting with the second quarter of 2007. Meyer and Sullivan (2023) note that the change resulted in a significant increase in reported spending on food away from home, and they adjust down spending on food away from home after 2007. As explained in the main text, we do not adjust the CE data to correct for the effect of this change in the questionnaire. The main part of our analysis focuses on the period 2010-2018, after the change in the questionnaire, which is precisely when underreporting on food away from home expenditures is less likely.

PSID data

We use data from the 2011-2019 waves of the Panel Study of Income Dynamics (PSID), which is a longitudinal survey of a representative sample of U.S. individuals and their families. Despite covering the broad range of consumption expenditures for a shorter time period than the CE, the PSID has reputedly better measures of income. Before 1999, only housing and food expenditures data were collected. Coverage of expenditures was expanded in 1999, to include about 70 percent of all CE expenditures, and again in 2005, to cover about 95 percent of CE spending. Although the PSID collects less detailed information than the CE, total household expenditure in the PSID and the CE are broadly consistent. However, there are significant differences for some subcategories of spending (Andresky et al., 2014). Unlike the CE, the PSID is a biennial survey since 1997.

We measure visible spending as the sum of food away from home, plus clothing and apparel, plus transportation expenditures. The PSID does not identify expenditures on personal care. The category “clothing and apparel” includes all household expenditure on clothing and apparel, including footwear, outerwear, and products such as watches or jewelry. The category “transportation” combines expenditures for vehicle loan, lease, and down payments, insurance, other vehicle expenditures, repairs and maintenance, gasoline, parking and car pool, bus fares and train fares, taxicabs and other transportation. Our measure of shelter combines expenditures on mortgage and loan payments, rent, property tax, insurance and utilities. For a detailed description of the mapping of the CE data into PSID spending categories, see Andresky et al. (2014, Online Appendix, Table 3).

CPS data

Our instruments include lagged measures of before-tax income distribution in each region/age-group/year. We use data from the large samples of the March Current Population Survey (CPS) for the survey years 1983–2017, which provide information on income for the previous calendar year for all households, including those without labor force participants. We impose a minimal number of restrictions on the sample. We keep only households who report non-negative incomes below. We eliminate households with real before-tax incomes above \$200,000. We also restrict the sample to households where the head is between 25 and 64 years of age.

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 compute the median and the ratio of the p95 to p05 percentiles of the household income distribution for each region/age-group/year cell, using the household weights provided in the CPS, and merge it with the corresponding CE households in that group.

Table A1: Average expenditure shares by consumption category

Housing	30.0
Visible spending	
Vehicle purchasing, leasing, insurance	10.1
Food away from home	4.7
Men's and women's clothing	1.9
Shoes and other apparel	1.2
Personal care	0.8
Children's' clothing (up to age 15)	0.8
Non-visible spending	
Food at home	12.7
Utilities	5.8
Health expenditures including insurance	5.4
Appliances, phone, computers	5.1
Entertainment equipment	3.8
Cash contributions	2.1
Entertainment fees, admissions, reading	1.7
Furniture and fixtures	1.2
Domestic services and childcare	1.2
Education	1.2
Tobacco, other smoking	1.1
All other transportation	1.1
Alcoholic beverages	1.0
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 A2: Sample summary statistics, overall and by age of the household head

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

Note: Sample: CE waves 1984-2018. Expenditure and income series deflated to 1983 dollars using CPI index. Summary statistics for the whole sample and for the two age groups. Number of observations: 70,952 (All), 34,493 (Old), 36,459 (Young). Last column indicates that the means between two age groups are different at 1% significance level)

Appendix B

Proof of Lemma 1

To verify Part (i), note that

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

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

To verify Part (ii), note that

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

where the elasticity of marginal utility is equal to

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

To verify Part (iii), note that

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

To verify Part (iv), note that

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

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

$$v'_i = c'_i (V'/\alpha)^{1/\sigma}, \text{ where } V' \equiv E \left((v'_i)^{\delta(\sigma-1)} \right), \quad (6)$$

which, using the fact that $c'_i = a'_i + y'_i - v'_i$, implies that

$$v'_i = \left[\frac{(V'/\alpha)^{1/\sigma}}{1 + (V'/\alpha)^{1/\sigma}} \right] (a'_i + y'_i), \quad (7)$$

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

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

equivalently,

$$V' \left[1 + (V'/\alpha)^{-1/\sigma} \right]^{\delta(\sigma-1)} = E \left((a'_i + y'_i)^{\delta(\sigma-1)} \right), \quad (8)$$

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 (8) 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 (8) with $V' > 0$. In turn, this implies that there is a unique solution to equation (7) 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:

$$\frac{c'_i}{c_i} = \left(\frac{1 + r}{1 + \rho} \right)^{1/\sigma} \equiv 1 + g_c. \quad (9)$$

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:

$$v_i = c_i (V/\alpha)^{1/\sigma}, \text{ where } V \equiv E \left(v_i^{\delta(\sigma-1)} \right). \quad (10)$$

Equations (9) and (10), together with (6) imply that

$$\frac{v'_i}{v_i} = (1 + g_c) \left(\frac{V'}{V} \right)^{1/\sigma}. \quad (11)$$

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

$$\frac{v'_i}{v_i} = (1 + g_c)^{\frac{\sigma}{\sigma-\delta(\sigma-1)}} \equiv 1 + g_v. \quad (12)$$

Corollary 1 follows immediately.

Next, note that the intertemporal budget constraint, given by (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 (10), we have

$$v_i = H(V) Y_i \quad (13)$$

and

$$c_i = (V/\alpha)^{-1/\sigma} H(V) Y_i, \quad (14)$$

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 (13) to the power of $\delta(\sigma - 1)$, and aggregate over all consumers to obtain

$$V [H(V)]^{-\delta(\sigma-1)} = E \left(Y_i^{\delta(\sigma-1)} \right). \quad (15)$$

Writing equation (15) 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 (15) with $V^* > 0$. In turn, this implies that there is a unique solution to equation (13) 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 (8) and equation (15) are satisfied simultaneously. To verify that this is the case, note that the consumer's budget constraint implies that

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

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

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

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 (8) 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 (11)). Hence equation (8) implies that the right sides of the above two equations must be equal, which together with equation (12) implies that equation (15) 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 (15). To that end, first note that in the proof of Proposition 1 we show that the left side of (15) 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**

Proof of Proposition 4

Proposition 1 implies that

$$\ln v_i^* = \ln H(V^*) + \ln Y_i,$$

$$\ln c_i^* = \frac{-1}{\sigma} \ln(V^*/\alpha) + \ln H(V^*) + \ln Y_i,$$

$$\ln(c_i^* + v_i^*) = \ln\left(1 + (V^*/\alpha)^{-1/\sigma}\right) + \ln H(V^*) + \ln Y_i.$$

The proposition follows immediately from noting that a linear combination of Normal random variables is Normally distributed. Calculating the mean and the variance of log spending in each of the three equations, as a function of V^* and the distribution of permanent income, is straightforward. **QED**