Unequal Business Cycles

Marcos Sonnervig*

July 10, 2023

[click here for most recent version]

Abstract

Standard incomplete-markets (SIM) models predict that the consumption of low-skilled households is more cyclical than the consumption of high-skilled because they hold fewer liquid assets and experience larger income changes. I use the Consumer Expenditure Survey to show that the opposite is true in the data: lower-skilled households experience smaller consumption changes over the business cycles. I also show that this difference in consumption cyclicality is explained by the fact that high-skilled households consume relatively more luxuries. Motivated by these facts, I extend the SIM model to allow for non-homothetic preferences over goods, which I discipline using cross-sectional data on how consumption shares of luxuries and necessities vary with income (Engel curves). The model reproduces the observation that consumption is more cyclical for high-skilled households because they spend a larger share of their income on luxuries, which are easier to substitute over time. The model also predicts that the welfare costs of recessions are larger for low-skilled households, despite their lower consumption declines.

*I am especially grateful to Virgiliu Midrigan, Thomas Sargent, and Ricardo Lagos for their advice and support throughout this project. I would also like to thank Jarda Borovička, Simon Gilchrist, Katarína Borovičková, Jess Benhabib, Danial Lashkari, Camilo Morales-Jimenez, Mark Gertler, Venky Venkateswaran, Diego Perez, Aditya Aladangady, Corina Boar, Bálint Szőke, Lars Ljungqvist, Seb Graves, Will Gamber, Guido Menzio, Rafa Guntin, Cristina Fuentes-Albero, Cynthia Doniger, Isabel Cairó, Tomaz Cajner, Jake Orchard, Nils Gornemann, and participants at the NYU Macro lunch and the Federal Reserve Board for insightful comments.

Email: marcos.sonnervig@nyu.edu

1 Introduction

Business cycles affect different households in very different ways. The income of workers without a college degree (hereafter low-skilled) is substantially more sensitive to business cycles than the income of college-educated workers (high-skilled)[1]. Low-skilled workers also tend to have lower savings relative to income, which limits their ability to self-insure. In spite of these facts, as I document in this paper, the consumption expenditures of low-skilled households are less sensitive to business cycles than the consumption of their high-skilled counterparts.

The standard incomplete markets (SIM) models used in the literature are at odds with these features of the data[2]. These models predict that the consumption of low-skilled households is more sensitive to the business cycles, owing to their more cyclical income and lack of savings to smooth income fluctuations. My paper reconciles the model's predictions with the data by recognizing that richer, high-skilled households spend a larger fraction of their expenditures on luxury goods, which are easier to substitute intertemporally. Therefore, they are more sensitive to changes in interest rates, and this amplifies their consumption changes. In particular, I extend the SIM model to allow for non-homothetic preferences over goods and show that the model is capable of reproducing the consumption cyclicality observed in the data. I then use the model to re-evaluate the welfare costs of recessions and their distribution across households.

The first contribution of this paper is empirical. Using data from the Consumer Expenditure Survey (CEX), I document that consumption expenditure is less cyclical for low-skilled households (relative to high-skilled), and I show that this gap is explained by differences in their consumption baskets; in particular, by the fact that high-skilled households spend a larger share on luxuries.

To do so, I first estimate the degree of "essentiality" for different consumer goods in the CEX. This is done by estimating how their budget share varies with the household's total consumption (Engel curves). More essential goods are those whose budget shares are larger for poorer, low-consumption households, such as "food at home" and "electricity". Analogously, less essential goods (or more luxurious ones) are those whose budget shares are larger for richer, high-consumption households. Equipped with this classification, I

[1] Section 2.2 estimates income cyclicality for high- and low-skilled households using CPS-ASEC data. Guvenen, Karahan, Ozkan, and Song [2017], Guvenen, Ozkan, and Song [2014] document similar disparities along the income distribution. Focusing on the employment margin, Cairó and Cajner [2018] document sizeable differences across skill groups.

[2] For some examples, Imrohoroğlu [1989]; Krebs [2003]; Krusell, Mukoyama, Şahin, and Smith Jr [2009]; Mukoyama and Şahin [2006]; and Krueger, Mitman, and Perri [2016a] study the welfare costs of business cycles and recessions for different households. Bhandari, Evans, Golosov, and Sargent [2021] and Le Grand and Ragot [2022] study the implications of household heterogeneity for the design of fiscal and monetary policies. The models used in these papers are built on Bewley [1983]; Imrohoroğlu [1989]; Huggett [1993]; and Aiyagari [1994].

construct consumption baskets that differ in their degree of "essentiality." Then, using monthly household-level consumption changes, I estimate the sensitivity of consumption expenditures to the business cycles and, importantly, how these differ between high- and low-skilled households.

I find that high-skilled households display higher consumption sensitivity than low-skilled households when comparing their total expenditures (with and without durables). However, when comparing their expenditures on more essential baskets - that is, when excluding luxuries - the difference between groups disappears. This finding suggests that taking the composition of consumption bundles into account is crucial in explaining differences in consumption dynamics between high- and low-skilled households.

The second contribution of this paper is to examine the implication of heterogeneous consumption baskets for consumption dynamics over the business cycles. In particular, I extend the SIM model to allow for non-homothetic preferences over goods. I assume that preferences over different goods have different income elasticities, which I estimate using cross-sectional data on how the consumption shares of luxuries and necessities vary with household income. In the model, a consumption good characterized by having higher income elasticity also has higher intertemporal elasticity. This reflects the theoretical prediction of a large class of consumer preferences and the intuitive idea that "luxuries are easier to postpone"[3]. As a consequence, households that spend a larger share of their budget on luxuries also have a higher elasticity of intertemporal substitution (EIS) for their total consumption expenditures.

This endogenous difference in EIS across households allows the model to reproduce the fact that consumption is less cyclical for low-skilled households. Intuitively, because low-income households spend a larger fraction of their expenditure on necessities, they are relatively more averse to consumption changes. This creates gains from trade and opportunities for risk-sharing between poorer and richer households.

In my model, I also take into account the fact that low-skilled households have lower liquid assets relative to their income. In particular, I target the average liquid asset-to-income ratios for high- and low-skilled households, as reported in the Survey of Consumer Finances (SCF). As a result, the model predicts that the marginal propensity to consume (MPC) out of transitory income changes is higher for low-skilled households[4]. In equilibrium, therefore, two opposing forces operate: on one hand, high-skilled households have higher liquidity and lower MPCs, making their consumption less responsive to income

[3] Browning and Crossley [2000] proves this statement for a class of preferences. Deaton [1974], Atkeson and Ogaki [1996] and Crossley and Low [2011] also discuss the tight link between income and intertemporal elasticities.

[4] This is consistent with Parker, Souleles, Johnson, and McClelland [2013] and Misra and Surico [2014] who find that MPCs estimates are large and heterogeneous across households. Relatedly, Patterson [2019] and Ganong, Jones, Noel, Greig, Farrell, and Wheat [2020] estimate higher values for low-income and low-liquidity households.

changes. On the other hand, because these households consume more luxuries, they also have a higher EIS, making their consumption more responsive to income (and interest rate) changes. Hence, the resulting difference in consumption dynamics across household groups depends on how strong each force is relative to the other.

The calibrated model predicts that differences in EIS across households dominate and, as a consequence, high-skilled households display more cyclical and volatile consumption than low-skilled households. As in the data, the model predicts that this gap is explained by differences in the composition of consumption baskets across household groups. Notably, when comparing the consumption of specific goods (luxuries or necessities) separately, the model predicts the opposite: Low-skilled households' consumption of both goods are actually more cyclical. This is explained by the fact that luxuries account for a larger share of high-skilled households' consumption, and this category is significantly more cyclical than necessities.

In this environment, general equilibrium interest rate responses play an important role in explaining consumption dynamics. This is revealed by decomposing consumption responses into two parts: the *direct response* driven by income changes and the *indirect response* driven by (equilibrium) interest rate changes. I find that the direct effect of income changes is stronger for low-skilled workers - which reflects their higher MPCs - and the indirect effect of interest rates is stronger for high-skilled workers - which reflects their higher EIS.

Finally, as the paper's third contribution, I use the calibrated model as a laboratory to investigate the welfare costs of recessions and how they vary across households. The model predicts that during recessions, low-skilled households' consumption falls less than that of high-skilled households. Despite this, the welfare costs for the less skilled are significantly larger than those for the more skilled. This is explained by the fact that, with non-homothetic preferences, the welfare cost of a consumption decline crucially depends on the composition of consumption, i.e., the shares of necessities and luxuries. Since low-skilled households consume relatively more necessities, a consumption decline is more costly for them.

Literature. This paper is related to several strands of a large literature that studies the welfare costs of business cycles, consumption dynamics, and inequality.

First, my paper contributes to a large body of research that studies the welfare costs of aggregate fluctuations and how they differ across households. Starting with Lucas [1987], who examines this topic in a representative agent economy, a substantial literature has revisited this subject, using models that take into account household heterogeneity. Important contributions include Imrohoroğlu [1989], Krebs [2003], Krusell et al. [2009] and Mukoyama and Şahin [2006][5]. Relatedly, Krueger et al. [2016a] studies the costs of

[5]Mukoyama and Şahin [2006] explicitly accounts for differences in the exposure to business cycles

experiencing a particular recession[6]. Overall, these papers find substantial heterogeneity in the cost of aggregate fluctuations, with low-wealth households suffering the most. My paper contributes to this literature in two ways. First, I use direct household consumption data and document differences in consumption dynamics across skill groups[7]. Second, I show that the SIM model used in this literature can be reconciled with the data once it accounts for the observed non-homotheticity in consumption preferences.

This paper also relates to a large literature on asset pricing that studies differences in consumption dynamics across households, with the objective of understanding the *equity premium*[8] and as a potential resolution of the *equity premium puzzle* (Mehra and Prescott [1985]). As discussed in Lucas [2003], the welfare cost of business cycles is closely related to the *equity premium* because both concepts reflect how much people are willing to pay to avoid aggregate risk.

In an influential study, Vissing-Jørgensen [2002] documents that the consumption of asset holders is more volatile and more sensitive to changes in interest rates than the consumption of non–asset holders. Parker and Vissing-Jorgensen [2009], Parker and Vissing-Jorgensen [2010], and Malloy, Moskowitz, and Vissing-Jørgensen [2009] find that the consumption of high-consumption households is significantly more sensitive to aggregate fluctuations than that of the average household.

Consumption expenditure dynamics also differ across goods. Using CEX expenditures data on 57 categories, Bils and Klenow [1998] show that expenditures on luxuries are substantially more cyclical than expenditures on necessities. Similarly, using sales data on imported luxury automobiles (BMW, Mercedes-Benz, Jaguar, etc.), luxury retailers (Tiffany, Gucci, etc.), and other luxury items, Ait-Sahalia, Parker, and Yogo [2004] find that the consumption of luxuries covaries significantly more with stock returns than does average consumption. Orchard [2022] shows that this different cyclicality across goods is also reflected in their relative prices: In recessions, necessities are relatively more expensive than luxuries.

Guntin, Ottonello, and Perez [2020] study household consumption adjustments during the Euro crisis and sudden stops in emerging markets, and show that rich and poor households significantly adjust consumption relative to their income. They argue that this is consistent with the view that these crises involve a large contraction of households' permanent income.

Regarding the disparities in income dynamics across households, this paper relates to

across skill groups.

[6] As opposed to the ex-ante welfare analysis in Lucas [1987].

[7] Krueger, Mitman, and Perri [2016b] also use consumption data from the Panel Study of Income Dynamics (PSID) and document consumption changes for households in different wealth quintiles during the Great Recession. In my paper, I use monthly-frequency consumption data, as opposed to biennial, and I focus on differences across skill groups.

[8] The difference between the expected return on equities and safe assets.

Guvenen et al. [2014], Guvenen et al. [2017], Kramer [2022], and Heathcote, Perri, and Violante [2010b], who study differences in income exposure to business cycles across workers and households. I contribute to this literature by documenting differences in the cyclicality of disposable income - that is, total income plus government taxes and transfers.

Finally, this paper also relates to an extensive and diverse literature that analyzes consumption dynamics with non-homothetic preferences. The vast majority of studies focus on long-run dynamics and the implications of non-homoteticity for structural transformation, development, and growth. For a recent example, Comin, Lashkari, and Mestieri [2021] exploit differences in the shapes of Engel curves across sectors to determine the contribution of supply and demand channels to structural changes. Atkeson, Ogaki, et al. [1990] and Steger [2000] show that non-homothetic preferences can explain differences in savings rates across rich and poor countries.

Less attention has been devoted to the consequences of non-homotheticity for consumption dynamics at business-cycle frequencies. One exception is Ravn, Schmitt-Grohé, and Uribe [2008], who study the consequences of non-homothetic preferences for the behavior of markups. This paper contributes to this literature by showing that non-homotheticity is important in explaining *heterogeneity* in consumption dynamics at business cycle frequencies.

Layout. Section 2 establishes the main empirical results on consumption and income cyclicality. Section 3 introduces the model. Section 4 describes the parametrization of the model, including the estimation of non-homothetic preferences. Section 5 presents the main quantitative results and quantifies welfare costs of recession and section 6 concludes.

2 Data

In this section, I document three facts. First, that consumption is more cyclical and more volatile for high-skilled households relative to low-skilled. Second, I show that this difference disappears once we exclude expenditure on more luxurious consumption categories. Third, I document that disposable income is more cyclical for low-skilled households.

2.1 Consumption expenditures

The consumption measures used in this paper come from detailed household-level data from the Consumer Expenditure Survey (CEX). The CEX contains detailed information on household consumption in the US, and it is used by the Bureau of Labor Statistics for constructing weights of the Consumer Price Index (CPI).

The CEX is a monthly rotating panel in which households are selected to be representative of the US population and is available since 1980. Every month, around 2000 households are interviewed. An individual household is interviewed at most four times, once every three months - not necessarily matching calendar quarters. In each interview, households report consumption expenditures for the previous three months. As a result, for each household, we observe monthly expenditure for at most twelve consecutive months.

I use information on the educational attainment of the household head as a proxy for skill level. In particular, I define two broad groups of households based on this information: college-graduates (high-skilled) and no-college (low-skilled) households.

I follow Coibion, Gorodnichenko, Kueng, and Silvia [2017] in the mapping between consumption categories in the CEX and their counterparts in the National Income Product Accounts (NIPA). More specifically, I construct measures of *services, durables* and *non-durables*, excluding automobile and housing. I deflate each of these categories by their appropriate CPI inflation measures. Following the recommendation by the BLS, I sum expenditures that occur in the same month but are reported by households in different interviews.

I consider two measures of consumption, including and excluding durables. Because expenditures on durables do not reflect the actual consumption of their service flows, I take the consumption measure without durables as my baseline. This is also consistent with the definition used in most of the literature (Attanasio and Browning [1993], Aguiar and Hurst [2005], Coibion et al. [2017], Vissing-Jørgensen [2002]).

As mentioned, in the CEX, households report consumption expenditures for at most twelve months. I restrict attention to households that reported consumption expenditures for all these months. I then compute 6-month consumption changes for each household, as in Vissing-Jørgensen [2002]:

$$\Delta c_{i,t} = \ln \frac{c_{i,t} + c_{i,t-1} + c_{i,t-2} + c_{i,t-3} + c_{i,t-4} + c_{i,t-5}}{c_{i,t-6} + c_{i,t-7} + c_{i,t-8} + c_{i,t-9} + c_{i,t-10} + c_{i,t-11}}$$

Using this measure of consumption changes at the household level, I estimate the consumption cyclicality (the co-movement with business cycles), the consumption volatility and how these differ across households with different skill levels. I follow an approach similar to Doniger [2021].

2.1.1 Consumption cyclicality

In this section, I present evidence that consumption is more cyclical for high-skilled workers. For different measures of consumption and different indicators of business cycles (Z_t), I estimate the following linear projection:

$$\Delta c_{i,t} = \gamma \Delta Z_t + \beta \Delta Z_t \times 1_i^{college} + \phi X_i + \phi_c X_i \times 1_i^{college} + \sum_{\tau=1}^{11} \gamma^{\tau} 1_{month(t)=\tau} + \varepsilon_{i,t} \qquad (1)$$

Where $\Delta c_{i,t}$ represents household i consumption change between period $t-1$ and t, X_i is a vector of household characteristics - it includes age, sex, family size, number of earners, and region dummies. I also include an interaction of these variables with a college indicator and control for group-specific trend in consumption growth. Finally, to account for seasonality, I also include monthly dummies in the regression.

ΔZ_t represents changes in the business cycle indicator. I consider three indicators: aggregate consumption, unemployment rate, and the NBER recession indicator. For the aggregate consumption, I consider the aggregate real consumption expenditures (computed with CEX data): $\Delta Z_t = \log(C_t^{agg}/C_{t-6}^{agg})$[9]. For the unemployment indicator, I consider changes in the detrended civilian unemployment rate. More specifically, I first detrend the unemployment rate using HP filter[10] and then compute 6-month changes: $\Delta Z_t = U_t - U_{t-6}$, where U_t is the cyclical component of the unemployment rate. The recession indicator (Recession_t) is equal to one if the economy was in a recession in any of the six months prior to t.

Table 1: Consumption cyclicality: nondurables + services

	$\Delta c_{i,t}$		
	(1)	(2)	(3)
$\Delta \log C_t^{agg}$	0.92***		
	(0.044)		
$\Delta \log C_t^{agg} \times 1_{college}$	0.21**		
	(0.067)		
ΔU_t		-2.14***	
		(0.29)	
$\Delta U_t \times 1_{college}$		-1.61***	
		(0.42)	
Recession_t			-0.0013***
			(0.004)
$\text{Recession}_t \times 1_{college}$			-0.002*
			(0.001)
Observations	57,255	57,255	57,255
Adjusted R^2	0.007	0.007	0.006

Note: Newey-West standard errors in parentheses. ***p < 0.01, **p<0.05, *p<0.1.
Source: Consumer Expenditure Survey; author's calculations.

[9] Using total consumption expenditure measure: services + nondurables + durables.

[10] I use a smoothing parameter equal to 1600×3^4, as suggested by Ravn and Uhlig [2002]. In appendix A I report results for alternative smoothing parameters.

Table 1 reports results using the baseline consumption measure (services and non-durables). The first column reports estimates for γ and β, when using the aggregate consumption as cyclical indicator. The coefficients 0.92 and 0.21 indicate that a 1% increase in aggregate consumption is associated with an average 0.92% increase in consumption for low-skilled households and a 1.13% increase for high-skilled. The difference between the two groups is statistically significant. The second column reports estimates when using the civilian unemployment rate as cyclical indicator. The coefficients -2.14 and -1.61 indicate that a 1 percentage point increase in the unemployment rate is associated with an average 2.14% decrease in consumption for low-skilled households and a 3.75% fall for consumption of high-skilled. The difference between the two groups is also statistically significant. Finally, the third column reports estimates when using the recession indicator. The coefficients -0.013 and -0.002 indicate that during recessions, 6-month consumption changes are, on average, 1.3% lower for low-skilled workers and 1.5% lower for high-skilled.

Table 2 reports results using the total consumption measure: nondurables + services + durables. Qualitatively results are similar, with total consumption expenditures being more cyclical for high-skilled households.

Table 2: Consumption Cyclicality: Total Consumption

	$\Delta c_{i,t}$		
	(1)	(2)	(3)
$\Delta \log C_t^{agg}$	0.94***		
	(0.041)		
$\Delta \log C_t^{agg} \times 1_{college}$	0.18**		
	(0.072)		
ΔU_t		-2.39***	
		(0.40)	
$\Delta U_t \times 1_{college}$		-2.1***	
		(0.56)	
Recession_t			-0.0017***
			(0.005)
$\text{Recession}_t \times 1_{college}$			-0.002
			(0.001)
Observations	57,255	57,255	57,255
Adjusted R^2	0.005	0.004	0.005

Note: Newey-West standard errors in parentheses. ***p < 0.01, **p<0.05, *p<0.1.
Source: Consumer Expenditure Survey; author's calculations.

Overall the estimation results in this section provide evidence that consumption expenditures (with and without durables) of high-skilled workers are more cyclical than those of low-skilled. The next section, studies the role of luxuries in explaining this difference in consumption cyclicality.

2.1.2 The role of luxuries and necessities

In this section, I evaluate the role of differences in consumption baskets in explaining the consumption cyclicality gap between high- and low-skilled households. In particular, I show that differences in the shares of luxuries and necessities consumed by each group explain the unequal consumption sensitivity.

I proceed as follows. First, I classify consumption categories in terms of how essential they are. Following Bils and Klenow [1998], for each consumption category j (e.g. food at home, gasoline, books, etc.) I compute a *"luxury index" (ℓ_j)* defined by the following estimate:

$$S_i^j = \alpha_j + \ell_j \log c_i + \varepsilon_i$$

where c_i is the per-capita total consumption expenditure of household i. S_i^j is the budget share of category j for household i. The *"luxury index" (ℓ_j)*, therefore, measures how the share of a specific consumption category changes with household per-capita consumption. The higher is ℓ_j, the more "luxurious" (or less essential) that category is. Analogously, if the estimated ℓ_j for a given category is low (for many, it is actually negative) that category can is interpreted as a *necessity*.

This procedure allows us to rank consumption categories, based on how luxurious or essential they are. Table 3 reports the outcome of this classification when restricting the consumption to non-durables and services. Not surprisingly, categories such as *food at home* and *electricity* are found to be the more essential (or least luxurious), and categories such as *food away from home* and *entertainment services* are found to have the higher luxury index - within non-durable and services.

Table 3: Luxuries and necessities (non-durables and services)

Category	share low-skilled	share high-skilled	Luxury index (x100)
Food at Home	35%	25%	-19.2
Electricity	11%	9%	-5.6
Gasoline	12%	10%	-4.4
Telephone services	7%	6%	-2.7
Water and sewerage maintenance	3%	3%	-1.2
Tobacco products	2%	2%	-1.0
Intercity transportation	1%	1%	-0.5
Clothing	4%	4%	-0.4
Clothing Services	1%	1%	-0.3
Elderly expenses	0%	0%	0.0
Personal care	0%	0%	0.0
Books	0%	0%	0.3
Finance charges	1%	1%	1.3
Personal care	1%	2%	1.3
Alcohol away from home	1%	2%	1.4
Day care and preschool	2%	3%	2.5
Alcohol	0%	2%	2.9
Airline and taxi fares	2%	4%	4.7
Lodging away from home	1%	3%	5.1
Entertainment services	6%	9%	6.7
Food away from home	11%	15%	8.7

Source: CEX and author's own calculations. Sample restricted to prime-age heads.
Note: Expenditure shares are defined within non-durables and services.

I then define two baskets: necessities and luxuries. Necessities contain the most essential categories. Specifically, I include all categories with a negative luxury-index: *food at home, electricity, gasoline, ..., clothing services*. Luxuries contain all the other categories. Now, to assess the role of the difference in consumption baskets across high and low-skilled, I re-estimate Equation 1, using expenditures on necessities only. In other words, I estimate the same regressions, but now excluding from the consumption expenditure measures ($c_{i,t}$), the amount spent on luxury categories.

The results are reported in Table 4. To facilitate comparison, the first two columns repeat the estimates of Table 1, and the last two columns report the new estimates using the new consumption measure (Ex-luxuries). Because the differential cyclicality β was already insignificant at 5% significance level, when using the recession indicator, I focus on the other two indicators of business cycles: aggregate consumption and unemployment rate.

The estimates of the differential sensitivity β, for both aggregate consumption ΔC^{agg}, and unemployment rate ΔU are now insignificant, and the point estimates have the opposite sign. While the sensitivity to aggregate consumption was 0.21 higher for high-skilled households when luxuries are excluded, the differential sensitivity is negative and statistically indistinguishable from zero. Likewise, the higher sensitivity of high-skilled to changes in the unemployment rate (-1.61) reverses and is no longer significant.

Table 4: Consumption Cyclicality: The Role of Luxuries

	$\Delta c_{i,t}$			
	All services and non-durables		Ex-luxuries	
	(1)	(2)	(3)	(4)
$\Delta \log C_t^{agg}$	0.92***		0.84***	
	(0.044)		(0.037)	
$\Delta \log C_t^{agg} \times 1_{college}$	0.21**		-0.01	
	(0.067)		(0.067)	
ΔU_t		-2.14***		-2.01***
		(0.29)		(0.24)
$\Delta U_t \times 1_{college}$		-1.61***		0.061
		(0.42)		(0.45)
Observations	57,255	57,255	57,255	57,255
Adjusted R^2	0.005	0.004	0.005	0.009

Note: Newey-West standard errors in parentheses. ***p $<$ 0.01, **p$<$0.05, *p$<$0.1.
Source: Consumer Expenditure Survey; author's calculations.

One potential concern with the results just described is the fact that the cutoff used in the definition of necessities and luxuries was arbitrary. Focusing on the unemployment rate as a cyclical indicator ΔU_t, Figure 1 plots the estimated differential cyclicality β for alternative cutoffs. Starting from the original basket, which contains all categories, and subtracting luxury goods sequentially until there is only one good remaining, the least luxurious category, *food at home*.

Each dot in Figure 1 represents one estimate for β, and the black lines represent 90% confidence intervals. The first dot on the left represents the estimate for the consumption basket that includes all categories and, therefore, its value is exactly the one reported in Table 1: -1.61. As we move to the right, luxury categories are excluded sequentially, starting from the most luxurious *food away from home*. The upward-sloping trend indicates that once we start excluding luxuries, the difference in cyclicality between high and low-skilled households dies out. In other words, the more we restrict expenditures to more essential goods, the lower is the difference in consumption cyclicality between high and low-skilled households.

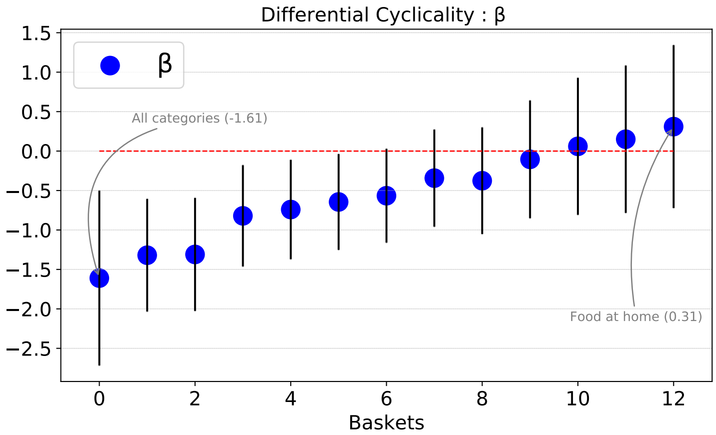

Figure 1: Difference in consumption cyclicality, between high-skilled and low-skilled, for different baskets. Starting from the all categories, luxury goods are removed sequentially, starting from the most luxurious to the least (*food at home*). Blue dots represent point estimates of β in equation 1, using the unemployment rate as cyclical indicator $\Delta Z_t = \Delta U_t$. Black lines represent 90% confidence intervals.

2.1.3 Consumption volatility

In the previous section, I showed that consumption is more cyclical for high-skilled workers relative to low-skilled. In this section, I focus on another important moment of their consumption process, the volatility, and document that consumption expenditure is also more volatile for high-skilled workers. These estimates serve as additional moments and will be used in the evaluation and validation of the model.

More specifically, I now estimate consumption expenditure volatility for both groups of households, also considering different consumption categories. Because in the CEX households are not observed over a long period of time, we cannot estimate consumption volatility at the individual level. Following the literature (e.g. Vissing-Jørgensen and Attanasio [2003], Chen, Favilukis, and Ludvigson [2013]) I address this limitation of the data by grouping individuals based on some observable characteristics, in my case, their education level. I then construct, for each group, a time-series of semi-annual consumption changes computed as the intra-cohort means. That is, the (average) semi-annual consumption change of college graduates in a given month is computed by averaging the consumption changes of the N_t^{high} high-skilled individuals observed in that given month:

$$\Delta \ln C_t^{high} = \frac{1}{N_t^{high}} \sum_{i=1}^{N_t^{high}} \Delta \ln c_{i,t}$$

I then compute the standard deviation of the time series $\{\Delta \ln C_t^{high}\}$, after controlling for seasonality. More specifically, I compute the standard deviation of the residuals of the following projection:

$$\Delta \ln C_t^{high} = \alpha + \sum_{\tau=1}^{11} \gamma^\tau 1_{month(t)=\tau} + \varepsilon_t$$

The results are summarized in Table 5. The resulting estimates show that for all categories, the volatility of consumption is higher for high-skilled households. This reaffirms the evidence that consumption is more stable for low-skilled households. Because the number of individuals within households may potentially vary between groups, I also restrict the sample to single-individual households and find similar results Table 16 in Appendix B report these estimates.

Table 5: Consumption volatility across skill groups and consumption categories

	Consumption	Services + Non-durables	Non-durables	Services
low-skilled	4.2	3.8	4.4	5.3
high-skilled	4.9	4.7	5.2	6.8
Ratio	**0.9**	**0.8**	**0.8**	**0.8**

Source: Consumer Expenditure Survey; author's calculations. Numbers in the first two lines represent standard deviation of consumption expenditures, for a given household group. The third line shows the ratio between low-skilled and high-skilled.

2.2 Income cyclicality

In this section, I analyze the income dynamics of households with different education levels, using annual data from the Annual Social and Economic Supplement (ASEC) of the Current Population Survey (CPS). I document that income is more cyclical for low-skilled households, even after all government transfers and taxes are taken into account.

The Current Population Survey is the official source of U.S. government statistics on employment, and is designed to be representative of the civilian non-institutional U.S. population. The ASEC supplement applies to the sample interviewed in March, and expands the set of questions to include detailed on income. The basic unit of observation is a household[11]. Households are interviewed at most two times, in two consecutive years. I focus on households that are interviewed twice, 12 months apart, always in March. I use the information on the household head to determine the sample. As in the CEX, I focus on prime-age heads (age between 25 and 65), and I split households based on education: college graduates (high-skilled) and no-college (low-skilled). I exclude the

[11] Strictly speaking, the unit of observation in the survey is a "housing unit" which is defined as all persons, related or unrelated, living together in a dwelling unit.

Covid-19 recession and focus on samples between 1979 - 2020. My final sample contains an average of 13.000 households per year.

Following Heathcote, Perri, and Violante [2010a], I define disposable income as: (i) pre-tax income, which includes all income from wages, salaries, business, self-employment and capital; plus (ii) government transfers, such as unemployment insurance (UI), welfare, social security, and pensions; minus (iii) taxes. Taxes are imputed using the TAXIM program[12]. I explain the income measures in more detail in Appendix B.

I drop households with less than \$500 annual income. Since I'm including all government transfers, these represents less than 2% of the sample.

As done with the consumption data, I compute income changes at the household level $\Delta y_{i,t}$, and I estimate the following projection in Equation 2:

$$\Delta y_{i,t} = \log\left(y_{i,t}\right) - \log\left(y_{i,t-1}\right)$$

$$\Delta y_{i,t} = \gamma \Delta Z_t + \beta \Delta Z_t \times 1_i^{college} + \phi X_i + \phi_c X_i \times 1_i^{college} + \varepsilon_{i,t} \qquad (2)$$

where $\Delta y_{i,t}$ represents household i income change between years $t-1$ and t, X_i is a vector of household characteristics - it includes age, sex, family size, and region dummies. I also include an interaction of these variables with a college indicator, and also control for group-specific trend in income growth. ΔZ_t represents changes in the business cycle indicator. I consider three indicators: aggregate income, unemployment rate, and the NBER recession indicator. For aggregate income, I consider aggregate real income, computed in the ASEC sample: $\Delta Z_t = \log(Y_t^{agg}/Y_{t-1}^{agg})$. For the unemployment indicator, I consider changes in de-trended civilian unemployment rate. More specifically, I first de-trend the unemployment rate using HP filter[13] and then compute 12-month changes: $\Delta Z_t = U_t - U_{t-1}$, where U_t is the cyclical component of the unemployment rate. The recession indicator (Recession_t) is equal to one if the economy was in a recession in any of the twelve months prior to t.

[12]TAXSIM is the NBER's program for calculating U.S. Federal and State income taxes from individual data.

[13]I use smoothing parameter equal to 6.25, as suggested for annual data by Ravn and Uhlig [2002].

Table 6: Disposable income cyclicality

	$\Delta y_{i,t}$		
	(1)	(2)	(3)
$\Delta \log Y_t^{agg}$	1.15***		
	(0.06)		
$\Delta \log Y_t^{agg} \times 1_{college}$	-0.20**		
	(0.07)		
ΔU_t		-2.12***	
		(0.17)	
$\Delta U_t \times 1_{college}$		0.47*	
		(0.28)	
Recession$_t$			-0.026***
			(0.001)
Recession$_t \times 1_{college}$			0.005*
			(0.002)
Observations	703,290	703,290	703,290
Adjusted R^2	0.005	0.004	0.004

Note: Newey-West standard errors in parentheses. ***p $<$ 0.01, **p<0.05, *p<0.1.
Source: Annual Social and Economic Supplement, CPS-ASEC and author's calculations.

Table 6 reports the main results for income cyclicality. The first column reports estimates for γ and β, when using the aggregate income as the cyclical indicator. The coefficients 1.15 and -0.20 indicate that 1 percent increase in aggregate income is associated with an average 1.15% increase in income for low-skilled households and a 0.95% increase for high-skilled. The difference between the two groups is statistically significant. The second column reports estimates when using the civilian unemployment rate as cyclical indicator. The coefficients -2.12 and 0.47 indicate that 1 percentage point increase in the unemployment rate is associated with an average 2.12% fall in income for low-skilled households and a 1.65% fall for high-skilled. Finally, the third column reports estimates when using the recession indicator. The coefficients -0.026 and 0.005 indicate that during recessions, 1-year income changes are on average 2.6% lower for low-skilled workers and 2.1% lower for high-skilled.

Overall, results in Table 6 show that disposable income is more cyclical for low-skilled workers. This is consistent with Guvenen et al. [2014], Guvenen et al. [2017] who show that income is more cyclical for low-income workers. It is also consistent with the fact that employment is substantially more cyclical for low-skilled workers, as studied in Cairó and Cajner [2018].

3 Model

This section describes the environment. The model serves two purposes. First, it is used to explore the consequences of non-homothetic preferences to consumption dynamics. Second, it is used as a laboratory to evaluate the welfare costs of recessions and their distribution across households.

The model is a standard Bewley–Huggett–Aiyagari-type dynamic general equilibrium model with incomplete markets (Bewley [1983] ,Huggett [1993], Aiyagari [1994]). I approximate the equilibrium with aggregate uncertainty using impulse responses to unexpected (or "MIT") shocks, as in Boppart, Krusell, and Mitman [2018].

3.1 Setup

Time is discrete. The economy is populated by a continuum of households, with measure 1. There are two (permanent) types of households $i \in \{low, high\}$. A fraction ω are high-skilled ($high$) and a fraction $1-\omega$ are low-skilled (low). These different types of households differ in (i) their income process, (ii) their exposure to aggregate shocks, and (iii) their return on savings.

Households have non-homothetic preferences over two types of goods: necessities (c_N) and luxuries (c_L), and maximize:

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u\left(c_{N,t}, c_{L,t}\right) \tag{3}$$

where $\mathbb{E}_t$ is a mathematical expectation operator conditioned on time t, $\beta \in (0,1)$ is the time discount factor and u is the utility function and it satisfies Inada conditions. Households are endowed with one unit of labor which is supplied inelastically. Following the standard practice in the literature, I assume that the log of idiosyncratic productivity z follows an AR(1) process:

$$\log z' = \mu_t^i + \rho \log z + \sigma \varepsilon' \tag{4}$$

Innovations are normally distributed: $\varepsilon \sim N(0,1)$. High-skilled workers have a higher average productivity $\mu_t^{high} > \mu_t^{low}$.

There is only one asset in this economy: one-period government bonds. Households cannot borrow. I assume that the rate of return on bonds differs across types: while high-skilled workers are paid $(1+r_t)$ units of c_N in period t for every unit invested in period $t-1$, the return for low-skilled households is $(1+r_t)(1-\phi)$, with $\phi \in (0,1)$[14].

[14] This can be interpreted as unmodeled differences in investment sophistication and/or the fact that high-skilled households have access to superior financial products. This difference is important because it allows the model to match the lower savings rates of low-skilled households.

Production:

Firms are competitive and operate a linear technology in the production of both goods:

$$Y_t^N = Z_t L_t^N \tag{5}$$

$$Y_t^L = Z_t L_t^L \tag{6}$$

Fiscal policy:

The government has an outstanding amount of debt D_t, sets a linear labor tax rate τ_t and lump-sum transfers T and is subject to the following budget constraint:

$$(1 + r_t)\, D_t + T = D_{t+1} + \tau_t w_t L_t + (1 + r_t)\phi \int a'(a, z, low) n_{t-1}(a, z, low) \tag{7}$$

where $L_t = \int zn(z, i)$ is the total amount of efficient units at time t. The government adjusts its government debt in response to changes in TFP in the following way:

$$\log(D_t/\bar{D}) = \rho_d \log D_{t-1} + \sigma_d \log(Z_t/Z) \tag{8}$$

Given D_t, the tax-rate τ_t adjusts to satisfy the budget constraint.

Recursive formulation:

The recursive value of a household of type i with assets a, productivity z, at time t is the following:

$$V_t(a, z, i) = \max_{c_N, c_L, a' \geq 0} \left\{ u(c_N, c_L) + \beta E_{z'} V_t\,(a', z', i) \right\} \tag{9}$$

Subject to:

$$c_N + p_t c_L + a' = (1 + r_t)(1 - \phi^i)a + z w_t (1 - \tau_t) + T \tag{10}$$

where w_t is the wage per efficiency unit, p_t is the relative price of good c_L, and ϕ^i is the spread on savings return ($\phi^{high} = 0$).

Aggregate shocks:

The only aggregate shock in the baseline economy TFP. Following the literature, I assume it follows an AR(1) process:

$$\log(Z_t/Z) = \rho_Z \log Z_{t-1} + \sigma_Z \varepsilon_t \tag{11}$$

Differential exposure to aggregate shocks:

To capture the different income cyclicality between high-skilled and low-skilled households, I assume that low-skilled workers are more exposed to aggregate shocks:

$$\mu_t^{low} = \mu^{low} + f_Z \log(Z_t/Z) \tag{12}$$
$$\mu^{high} = \mu^{high} \tag{13}$$

3.2 Equilibrium

Before defining the equilibrium, note that the fact that firms operate in a competitive environment and that both goods will be consumed in equilibrium[15] implies that the wage paid in each sector j is the same: $w_t = Z_t$. As a consequence, the relative price between the two goods will be equal to $p_t = 1$[16].

A **competitive equilibrium** consists of: (i) interest rates r_t, (ii) consumption, saving decisions of households $c_t^N(a,z,i), c_t^L(a,z,i), a_t'(a,z,i)$, (iii) labor and output choices of firms L_t^j, Y_t^j, and (iv) measures of households over their idiosyncratic states $n_t(a,z,i)$, such that:

1. Given prices, households and firms solve their optimization problems.
2. The measure $n_t(a,z)$ evolves according to an equilibrium mapping dictated by the households' optimal choices and the stochastic process for individual productivity.
3. The budget constraint of the government is satisfied every period.
4. Markets clear:
$$D_t = \int a_t'(a,z,i) dn_t(a,z,i).$$
$$L_t^N + L_t^L = L_t$$
$$Y_t^L = \int c_t^L(a,z,i) dn_t(a,z,i).$$
The market for c_t^N clears by Walras' Law.

4 Parametrization

4.1 Estimation of non-homothetic preferences

I assume that preferences over necessities (c_N) and luxuries (c_L) are separable and iso-elastic:

[15] As a consequence of the Inada conditions.

[16] Firm's profits are $\pi^j = p^j Y^j - w^j L^j$. If wages were different, all households would work for the sector that pays higher wages. Now, with equal wages, zero profits require $p = p^L/p^N = 1$.

$$u(c_N, c_L) = B\frac{c_N^{1-\frac{1}{\varepsilon}}}{1-\frac{1}{\varepsilon}} + \frac{c_L^{1-\frac{1}{\eta}}}{1-\frac{1}{\eta}} \tag{14}$$

Let E be the total expenditures and p the relative price across goods.

$$E = c_N + pc_L$$

Under these assumptions, the intra-temporal optimality condition takes a convenient log-linear form:

$$c_L^{-1/\eta} = \frac{B}{p}c_N^{-1/\varepsilon} \implies \log c_L = \log B/p + \frac{\eta}{\varepsilon}\log c_N \tag{15}$$

Note that if $\varepsilon = \eta$ then c_L and c_N are proportional and there is no meaningful distinction between goods in this economy. If $\varepsilon < \eta$ then the budget share of c_L increases with the amount of total consumption - that is, c_L is a luxury.

The log-linear relationship between consumption of luxuires and necessities implied by Equation 15 allows us to estimate ratio of elasticities $\theta := \frac{\eta}{\varepsilon}$, using cross-sectional data on consumption expenditures. In particular, I use the same CEX data described in section 2.1, restricting the sample to 2018. One year of data is enough to precisely estimate and θ and the estimate do not vary significantly when using different sample periods. More specifically, I estimate the following linear projection:

$$\log c_i^L = \alpha + \theta \log c_i^N + \beta X_i + \pi_i^{ra} + \varepsilon_i \tag{16}$$

where c_i^L represents expenditures by household i on "luxuries" - defined as in Section 2.1.2; and c_i^N the expenditures on "necessities".X_i represents household controls and include: age, sex, household size, region, and number of earners. π_i^{ra} represents a set of interactions between age and region. Since I'm not controlling for prices, an important identifying assumption is that households in the same region, same year and with the same age face the same prices and choose consumption baskets that are equal, up to the heterogeneity allowed in household characteristics X_i.

Following, Aguiar and Bils [2015] and Comin et al. [2021], I instrument c_i^N with household annual income. This is important to reduce variation in consumption of necessities that are driven by household-specific tastes and it also mitigates measurement errors.

Table 7: Non-homotheticity - estimation

	$\log c_i^L$	
	(1)	(2)
$logc_i^N$	3.731***	4.574***
	(0.0570)	(0.0829)
Controls	N	Y
Observations	15,771	15,159

Note: Standard errors in parentheses. ***p<0.01, **p<0.05, *p<0.1.
Source: Consumer Expenditure Survey (CEX)

Table 7 reports the estimated results. The first column reports the estimate of θ when no controls are included and the second column reports the estimate when controls are included. The fact that θ is significantly higher than one implies that expenditures on luxuries increase more than proportionally with increases in expenditures on necessities - as expected, given that categories included in the luxury bundle, by construction, are those whose budget shares are more elastic to total consumption. Figure 2 plots the (binned) residuals after all controls have been partialled-out from the instrumented consumption expenditures measures. As implied by the model, the residual variation in log luxuries is well approximated by a log-linear function of residual log necessities.

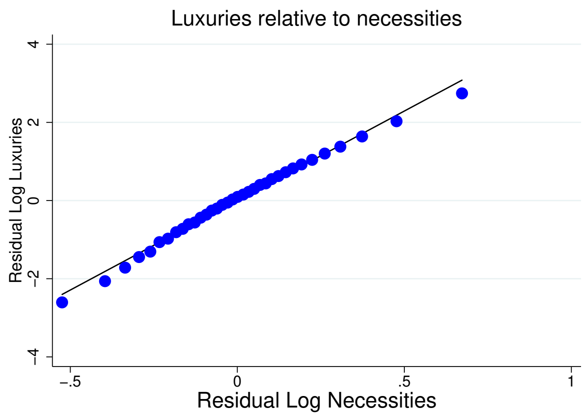

Figure 2: Partial correlation of the log consumption expenditures on luxuries and log consumption expenditures on necessities. Data from CEX. Notes: These plots depict the (binned) residuals corresponding to the average value of 35 equal-sized bins of the data. The black line depicts the linear regression between the residualized variables.

4.2 Other parameters

Having described the estimation of main preference parameter, I now describe on the calibration procedure for the other parameters of the baseline economy.

Assigned parameters A period is one quarter. A number of parameters need to be assigned: population share of high-skilled, ω, idiosyncratic productivity shock parameters, ρ and σ, the tax rate τ

The population share of high-skilled is set to 0.45 to match the average share of college-educated prime-age workers in 2020, according to data from the Current Population Survey (CPS). The parameters for the idiosyncratic productivity process, ρ and σ are set to 0.977 and 0.15. Aggregated annually, this process corresponds to one with a persistence of 0.91 and a standard deviation of 0.20, which I take as representative values from this literature (Floden and Lindé [2001], Krusell, Mukoyama, Rogerson, and Şahin [2017], Auclert, Rognlie, and Straub [2018]). For simplicity, I also assume that households cannot borrow: $\underline{a} = 0$. In line with the literature, I set the labor tax rate $\tau = 0.30$.

Regarding the process for (aggregate) TFP shocks, I choose persistence of 0.9 and standard deviation of 0.01, also standard values in the literature. Table 8 summarizes assigned parameters in the baseline economy.

Table 8: Assigned parameters of the benchmark economy.

Parameter	Description	Value
ω_{LS}	Pop. Share (Low Skilled)	0.55
ρ	Idiosyncratic Shock (persistence)	0.977
σ	Idiosyncratic Shock (std. dev)	0.15
τ	Labor Tax Rate	0.30
T	Gov. Transfers / GDP	0.29
$\underline{\text{a}}$	Ad-hoc Borrowing constraints	0
ρ_z	TFP process: persistence	0.9
σ_z	TFP process: std. dev TFP	0.01

I choose the remaining parameters so that the non-stochastic steady state equilibrium matches specific targets.

Liquid asset to income ratios:

An important moment in my analysis is the amount of liquid wealth relative to income. In this class of models, agents with lower liquid-wealth to income ratios are more likely to hit the borrowing constraint in the future. As a result, their ability to self-insure against income shocks is more limited.

Important determinants of the savings behavior are the discount factor and the return on savings. I first fix the return on savings $r = 4\%$ and then set both the discount factor

β, which is common across types, and the spread on savings return ϕ of low-skilled to match the average liquid wealth to income ratio for each group of households.

Using data from the 2018 sample of the Consumer Expenditure Survey, I compute average asset to income ratios for high-skilled and low-skilled households[17]. In the baseline calibration, I consider liquid assets, as the sum of transaction accounts (checking, savings, money market accounts), and I target the mean of the distribution for each type of household (high-skilled or low-skilled).

Table 9 summarizes the asset-to-income distribution in the SCF, for each group of households and for different measures of wealth. Regardless of the choice of assets and whether we compare the mean or median across groups, low-skilled workers have lower asset to income ratios. In the baseline calibration I target the mean of liquid assets for each group.

As shown in the first two lines of Table 9, the resulting value for the discount factor is $\beta = 0.98$ and for the spread on the savings return $\phi = 0.031$. Several studies document the existence of substantial and persistent differences in rates of return across households. Using Norwegian administrative tax records, Fagereng, Holm, Moll, and Natvik [2019] find that the rate of return for someone at the 10th percentile of the wealth distribution is 18 percentage points lower than the return of someone at the 90th percentile. Bach, Calvet, and Sodini [2020] and Smith, Zidar, and Zwick [2022] using respectively Swedish and U.S. data also find substantial differences in return along the wealth distribution.

Table 9: Asset to (annual) income ratios

	Asset Type					
	Liquid		Liquid + Stock		Net Worth	
	Mean	Median	Mean	Median	Mean	Median
High Skilled	34%	14%	70%	17%	571%	204%
Low Skilled	16%	5%	21%	5%	287%	97%

Note: Survey of Consumer Finances (SCF), 2019. Liquid assets include all checking, savings, money market accounts. High-skilled and low- skilled are households (head) with and without a college degree.

Given the total amount of government bonds, the interest rate, and the tax rate, the lump transfers T is chosen so that the government budget balances in the non-stochastic steady-state equilibrium.

Non-homotheticity:

In the previous section, I described in detail the estimation of the ratio between income elasticities of luxuries and necessities. The utility function in Equation 14 has two other

[17] I use the education level of the household head, and, following the definition used in the CEX and CPS-ASEC, high-skilled are the college graduates and low-skilled those without a college degree.

parameters: the elasticity of necessities ε and the scale parameter B. These two are chosen to match the average budget share of luxuries and the cumulative interest rate response to a 1% TFP shock. The first moment is directly computed using expenditure shares from CEX data. I calibrate it to match the average budget share of luxuries in 2018: 0.35.

For the interest rate response to a TFP shock, I rely on estimates from Winberry [2021]. This paper estimates the impulse response function of the real interest rate to a TFP shock using a simple vector autoregression VAR. The real interest rate is measured as the return on 90-day Treasury bills adjusted for realized CPI inflation. The measure of TFP, is taken from Fernald [2014]. TFP shocks are identified by assuming that shocks to the interest rate equation do not affect TFP on impact.

A 1% negative TFP shock leads to a positive response of interest rate, depicted in Figure 8, and the cumulative response over 20 quarters is equal to 3.8. I set ε targeting this estimate.

Table 10: Calibrated Parameters - baseline economy

Parameter	Value	Description	Data	Model
Preferences:				
β	0.98	Liquid Wealth to Income Ratio HS	0.34	0.34
ϕ	0.031	Liquid Wealth to Income Ratio LS	0.16	0.16
θ	4.57	Non-homotheticity	4.57	4.57
B	3.1	Average luxury share	0.35	0.35
ε	0.25	IRF (cummulative) Interest Rate	3.8	3.4
Income Process:				
$\mu^{high} - \mu^{low}$	-	Average Earnings Ratio: $\frac{Y^{HS}}{Y^{LS}}$	2.1	2.1
f_z	1.3	Differential Exposure	1.3	1.3
Government Debit Rule:				
ρ_d	0.9	IRF Interest Rate Shape	-	-
σ_d	0.01		-	-

Regarding the Government debt rule described by Equation 8, I set ρ_d, σ_d to match the shape of the interest rate impulse response function (IRF). In this economy, the interest rate is very sensitive to changes in the supply of assets, so a small change in the amount of government bonds has substantial effects on the shape of the IRF.

The last parameters are the ones governing the average income gap between high-skilled and low-skilled (the college premium) and the unequal exposure of low-skilled to aggregate TFP shocks f_Z. I normalize $\mu_{low} = 0$ and set μ^{high} (from in Equation 4) to match the average pre-tax income ratio between high-skilled and low-skilled. The differential exposure f_Z is set to match the estimated difference in post-tax income cyclicality, described in Section 2.

4.3 Wealth-varying EIS

Under the assumption of separable preferences, there is a tight link between the elasticities and the intertemporal elasticities across goods. In particular, luxury goods have a higher EIS. This reflects the intuitive idea that *luxuries are easier to postpone*. Browning and Crossley [2000] proves this statement for a more general class of preferences. Deaton [1974], Atkeson and Ogaki [1996] and Crossley and Low [2011] also discuss this connection.

Moreover, it can be easily shown (Section A) that in this model, the EIS for total consumption c can be written as an average of the two elasticities ε and η weighted by the consumption shares in necessities s_n and luxuries $1 - s_n$:

$$EIS\,(c) = \varepsilon s_n\,(c) + \eta\,(1 - s_n\,(c)) \tag{17}$$

Furthermore, if $\eta > \epsilon$, then c_L is a luxury and its budget share increases with c. Equation 17 reveals that a higher budget share on luxuries (lower s_n) leads to a higher EIS. Hence, in this economy, wealthier households have higher EIS. Importantly, the key parameter governing how the EIS varies with wealth across households is the ratio between elasticities: θ, which was estimated. To see this, Equation 17 can be re-written as:

$$EIS(c) = \varepsilon\big[s_n(c) + \theta(1 - s_n(c))\big] \tag{18}$$

Now note that the ratio between the total consumption EIS for households with consumption levels c_1 and c_2 is only a function of consumption shares and the estimated parameter θ:

$$\frac{EIS(c_1)}{EIS(c_2)} = \frac{s_n(c_1) + \theta\,[1 - s_n(c_1)]}{s_n(c_2) + \theta\,[1 - s_n(c_2)]}$$

Given that consumption shares $s_n(c)$ are directly observable in the data, we can test the prediction of the model and the ones implied by the data. Figure 3 plots the implied EIS from equation 18, using consumption shares in the data and the ones predicted by the model. Given the fact that the log-linear relationship between necessities and luxuries approximates the data very well, as shown in Figure 2, the calibrated model also matches consumption shares very precisely and, as a consequence, the model matches the implied EIS, as seen in Figure 3.

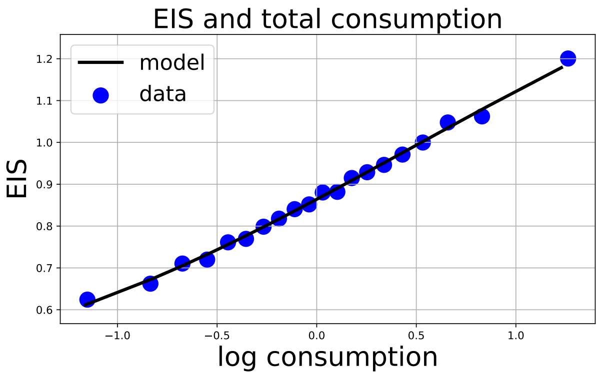

Figure 3: Elasticity of Inetertemporal Substitution as a function of consumption deviation from mean. EIS is computed with Equation 18, using shares s_n from the data and the ones implied by the model.

4.4 Homothetic (standard) model

I compare the dynamics of the baseline model described in the previous section with the a standard homothetic model. This section describes the parametrization of this standard model.

The homothetic preferences benchmark is the special case of my model when $\theta = 1$. I re-calibrate the parameters β, ϕ and ε to match the same set of moments of table 10. The new parameter values and targets are reported in Table 11.

Table 11: Calibrated Parameters of the homothetic benchmark

Parameter	Value	Description	Data	Model
β	0.98	Liquid Wealth to Income Ratio HS	0.34	0.34
ϕ	0.022	Liquid Wealth to Income Ratio LS	0.16	0.16
ϵ	0.91	IRF (cummulative) Interest Rate	3.8	3.4

The EIS ε is chosen to match the same cumulative interest rate response 3.4. Note that the spread on savings return ϕ is lower - and this is driven by the fact that low-skilled workers have a lower risk-aversion (higher EIS) and the self-insurance motive is dampened. Hence, to match the same level of asset-to-income ratio, the spread has to be smaller.

Relative to the baseline economy, there is little difference in the discount factor, but it is slightly lower here. This reflects the fact that high-skilled have slightly higher risk-aversion here (lower EIS).

5 Unequal Business Cycles

5.1 Unequal Business Cycles moments

In this section, I report business cycles moments of the model and compare them with the estimated using CEX data, in Section 2.1.1. The mechanisms driving this results are analyzed in the following section where I consider the dynamic responses to a single TFP shock.

Given the process for TFP, I simulate 10.000 quarters of data and estimate some moments, that are compared with the estimated using CEX data, in section 2.1.1. In particular, I estimate the projection

$$\Delta c_{i,t} = \alpha + \beta \Delta C_t^{agg} + \gamma \Delta C_t^{agg} \times 1_{college} + \eta_t$$

where $c_{i,t}$ is the total consumption expenditure of household i at time t, C_t^{agg} is the aggregate consumption and $1_{college}$ is an indicator for whether the household is high-skilled.

I compare the estimates for γ in the model and in the data (Equation 1). Point estimates are reported in the first column of Table 12. The model is successful in generating a consumption process more cyclical for high-skilled workers $\gamma > 0$. It explains around 76% of the differential cyclicality observed in the data. It is worth stressing that homothetic model predict the opposite of what is estimated in the data. I analyze the reasons for this in more detail in the next section.

The second column reports the estimates when instead of total consumption expenditures, the projection is estimated using only expenditures on necessities (excluding luxuries). As in the data, in this case, the consumption of high-skilled workers is *less* cyclical[18]. Because in the homothetic model goods are homogeneous, the model has no prediction for the cyclicality of necessities.

The third columns of Table 12 reports the ratio between the standard deviation of total consumption changes of high-skilled and low-skilled. Similarly, the model is able to reproduce the fact that consumption is more volatile for high-skilled workers and the homothetic benchmark fails to do so.

Finally, the last column reports the ratio between standard deviations of necessities and total consumption (including luxuries). As in the data, the model reproduces the fact that necessities are less volatile than total consumption. This is a testable prediction of the model and directly reflects the estimated ratio of income elasticities θ can be used to evaluate the assumption. In fact, an alternative way of estimating θ is by targeting this moment. The trade-off, however, is that estimates using noisier consumption time-series

[18] In the data, we cannot reject the hypothesis that $\gamma = 0$.

data are much less precise.

Table 12: Untargeted moment

	γ	$\gamma\vert_{necessities}$	$\frac{\sigma\left(c^{\text{high}}\right)}{\sigma\left(c^{\text{low}}\right)}$	$\frac{\sigma\left(c^{\text{necess}}\right)}{\sigma\left(c^{\text{total}}\right)}$
Data	0.21	-0.01	1.19	0.7
Baseline Model	0.16	-0.03	1.11	0.5
Homethetic Model	-0.08		0.85	

5.2 Recession dynamics

In this section, I examine the main experiment and application of the model: the recession dynamics. I first focus on studying the forces behind household consumption dynamics. I then compare the predictions of my baseline model with the homothetic benchmark. Finally, I assess the welfare cost of recessions and the differences between skill groups (and models).

Figure 4 reports the dynamic responses to a negative 1% TFP shock of output, wages, interest rate, total consumption expenditures (for high-skilled in blue and low-skilled in red), assets and the consumption gap between high-skilled and low-skilled.

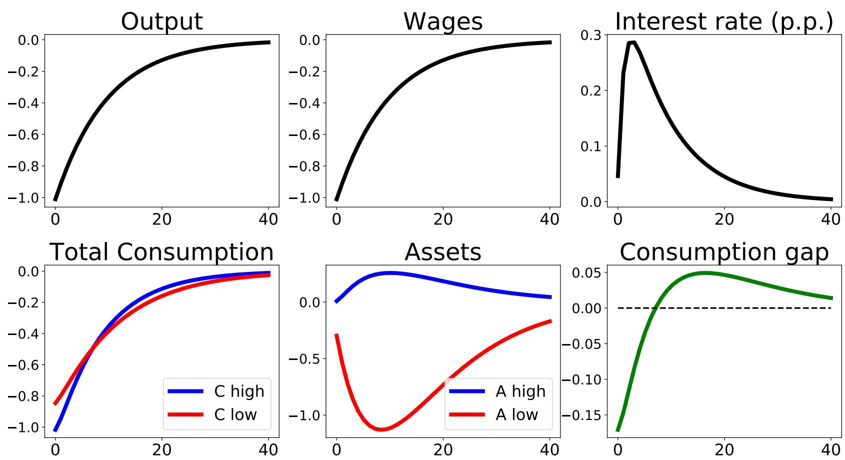

Figure 4: Impulse response functions in the baseline model - negative 1% TFP shock

The direct consequence of the negative TFP shock at time $t = 0$ is the decline in household income. Households would like to smooth consumption across time, but because the supply of goods is lower, the equilibrium price of consumption goods at $t = 0$ increase relative to the future (when the supply of goods is higher). In other words, the real interest rate increases. This change in relative prices affects the consumption-savings behavior of households differently: those with higher EIS respond more. As a consequence, in equilibrium, high-skilled households, who have a higher EIS, end up buying assets from the low-skilled. Also, as represented by the green line, the consumption falls relatively more for the high-skilled households: the *Consumption gap* falls.

Figure 5 depicts the consumption responses, breaking down the composition between luxuries and necessities. It reveals three things. First, that consumption expenditure on luxuries is substantially more responsive than expenditure on necessities. Second, within each consumption category, the response of low-skilled workers is higher, reflecting the fact that these individuals are more likely to be constrained. Third, it shows that the differential response of total consumption across households is driven by the fact that high-skilled households consume relatively more luxuries.

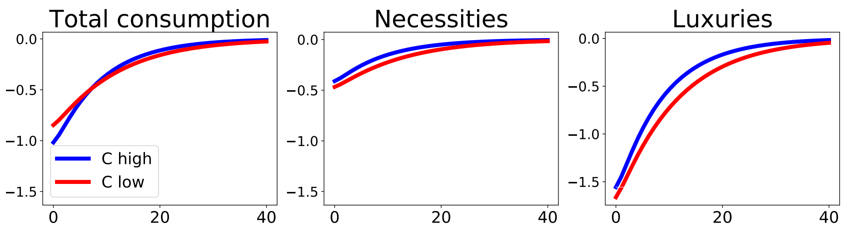

Figure 5: Consumption responses of luxuries and necessities in general equilibrium - IRF to negative 1% TFP shock

5.3 Partial equilibrium - the role of interest rates

To further understand the forces driving the consumption responses, consider what would be the consumption responses if interest rates were constant. Figure 6 plots this partial equilibrium experiment. Now the consumption of low-skilled workers is more responsive to the shock. Furthermore, focusing only in the response of necessities, the difference is even more pronounced.

This reflects the fact that low-skilled workers have lower liquid assets and, as a result, tend to be more constrained. This experiment also relates to the consumption responses estimated in the empirical literature on MPCs (Parker et al. [2013], Patterson [2019], Ganong et al. [2020]). Because these papers focus on non-durable consumption expenditure, which contains fewer luxury categories, the closer analog in my model is the response of necessities. Also consistent with this literature, the response of total consumption is higher than the consumption of necessities (Parker et al. [2013]).

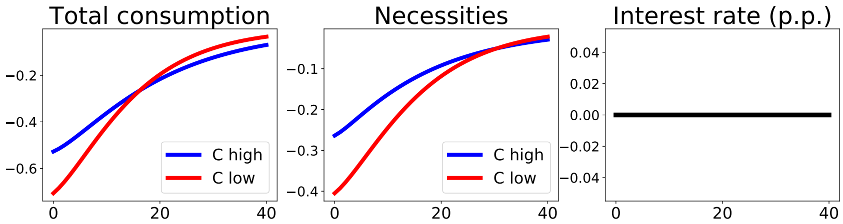

Figure 6: Partial equilibrium: constant interest rate - IRF to negative 1% TFP shock

Figure 6 plots the consumption responses for the two types in three different experiments: (i) General equilibrium (GE), (ii) constant interest rate, discussed in the previous paragraph and (iii) constant income (only changing interest rate). For the third experiment, I compute the consumption responses to the observed interest rate changes, holding labor income constant. From this experiment, it is clear the differential impact of these price changes on the consumption-savings behavior of households: the high-skilled respond substantially more.

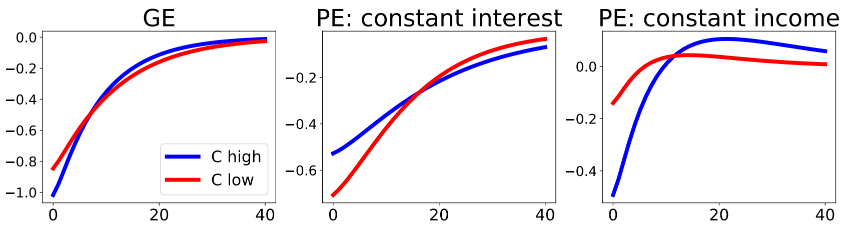

Figure 7: Consumption responses: General equilibrium, constant interest rate, constant income - IRF to negative 1% TFP shock

Given the importance of price changes to the behavior of households in this economy, it is important that the prediction of the model for interest rate dynamics matches the data. Figure 8 compares the dynamics of the model and data to a negative 1% TFP shock. The estimated IRF is taken from Winberry [2021][19].

[19]This paper estimates the impulse response function of the real interest rate to a TFP shock using a simple vector autoregression VAR. The real interest rate is measured as the return on 90-day Treasury bills adjusted for realized CPI inflation. The measure of TFP, is taken from Fernald [2014]. TFP shocks are identified by assuming that shocks to the interest rate equation do not affect TFP on impact.

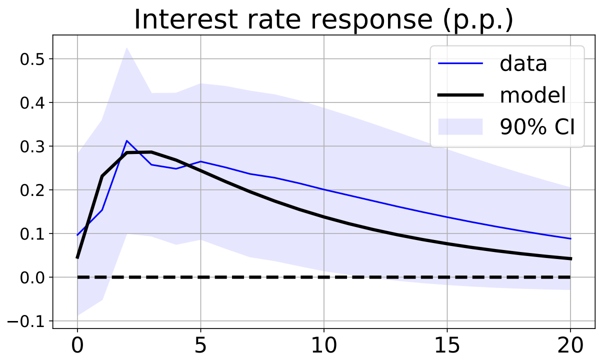

Figure 8: Winberry [2021]'s estimate of impulse response of the real interest rate to a negative 1% TFP shock and the predicted response in my baseline model. Empirical (90% CI) refers to the empirical impulse response and 90 percent error bands.

5.4 Standard homothetic model

The left plot in Figure 9 depicts the dynamic consumption responses of high- and low-skilled workers to a negative 1% TFP shock in the homothetic benchmark model. The right plot represents the difference between high- and low-skilled' consumption responses. Figure 9 shows that in the standard homothetic model, consumption of low-skilled households is more sensitive to recessions than consumption of high-skilled households. In particular, the average consumption of low-skilled households falls, on impact, around 0.25 p.p. more than avergate high-skilled households' consumption.

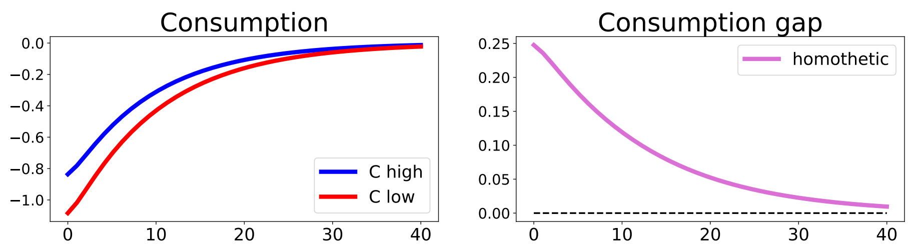

Figure 9: Consumption responses to a negative 1% TFP shock in the baseline model and in the homothetic benchmark.

In order to highlight the differences between the baseline non-homothetic model and the standard homothetic benchmark, in Figure 10 plots the difference between high- and low-skilled households' consumption responses to a negative 1% TFP shock in the baseline model and in the homothetic benchmark. As discussed, the model with homothetic preferences is unable to generate a higher sensitivity for the consumption of high-skilled (negative consumption gap). This is because, in this model, the only reason why consumption dynamics differ across groups is: (i) their different income cyclicality - which is higher for the low-skilled and (ii) the different amount of liquid assets - which is lower for the low-skilled. Hence, these two forces operate in the same direction and, as a consequence, the model predicts this counterfactual result.

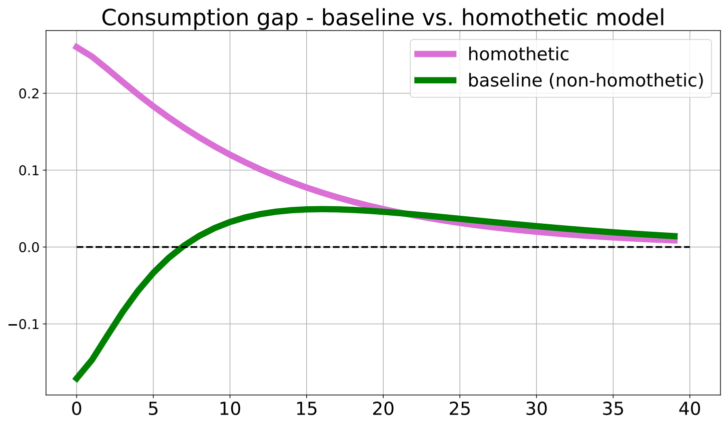

Figure 10: Difference between high- and low-skilled households' consumption responses to a negative 1% TFP shock in the baseline model and in the homothetic benchmark.

5.5 Welfare cost of recessions

Finally, I use the model to compare the welfare cost from recessions across skill groups. In particular, I compute the welfare costs to the same persistent negative 1% TFP shock studied in the last section. Welfare is measured in terms of % of lifetime total consumption expenditures. Table 13 reports the results. In the baseline model, the average welfare loss of low-skilled households is close to 0.30% of their lifetime consumption, whereas for high-skilled the cost is 0.08%. Comparing the average loss between groups, the model predict that low-skilled suffer 3.6 times more. Compared with the prediction of the homothetic model, this difference is slightly higher, and within each group of households, the baseline model predicts larger welfare in the non-homothetic economy.

Table 13: Welfare costs of recession

	low-skilled	high-skilled	ratio
Baseline	-0.29%	-0.08%	3.6
Homothetic	-0.21%	-0.07%	3.1

6 Conclusion

In this paper I study the differences in consumption cyclicality between high and low-skilled households and revisit the analysis of the welfare costs from aggregate fluctuations. The starting point of this paper is the observation that consumption is more cyclical for high-skilled households. This is important because the models used in the literature are inconsistent with this fact, as they predict that low-wealth households have more cyclical consumption (especially if they also have more cyclical income, as in the case of low-skilled workers).

The first contribution of this paper is to document, using data from CEX, that consumption expenditure is less cyclical for low-skilled households (relative to high-skilled), and that this gap is explained by the fact that they consume different products. In particular, high-skilled households spend a larger share on luxuries, which exhibit higher cyclicality. I show this using household-level monthly data on detailed expenditures. I construct consumption baskets with different degrees of "essentiality" and estimate their sensitivity to the business cycles. I find that when comparing total expenditures, high-skilled households display more cyclical consumption. However, excluding the more luxurious goods and comparing the expenditure on more essential items, I find no significant difference in their consumption sensitivities.

Motivated by these facts, I then extend the standard incomplete-markets model (Bewley-Hugget-Aiyagari) to allow for non-homothetic preferences over goods, and I use the model to (i) explore the consequences of heterogeneous consumption baskets for consumption dynamics; and (ii) the evaluation of welfare costs of recessions. I parameterize the degree of non-homotheticity with estimates of Engel curves for luxuries and necessities. I also match liquid-assets to income ratios for both low and high-skilled workers. These moments discipline the key parameters driving the positive and normative results.

The second contribution of this paper is to demonstrate, quantitatively, that non-homothetic preferences are important in explaining differences in consumption dynamics across households. In particular, contrary to the standard model used in the literature, the extended version with non-homothetic preferences is able to reproduce the observation that consumption is more cyclical for high-skilled households. Moreover, my model has testable predictions for the consumption volatility of luxuries and necessities, and these are also consistent with the estimates. Finally, in my model, the main drivers of consumption dynamics over the business cycles differ across high and low-skilled: low-skilled households are more sensitive to *income* changes, and high-skilled households are more sensitive to *interest rate* changes.

Finally, as the third contribution, I evaluate welfare costs of aggregate fluctuations and how they differ across households. Relative to the literature, I study this topic taking into account the fact that low-skilled households display *lower* consumption cyclicality. I find that the welfare costs of recessions are higher for low-skilled households despite the fact that these households exhibit lower *total* consumption changes. Therefore, the model highlights the importance of taking into account differences in consumption baskets on welfare analysis.

References

Mark Aguiar and Mark Bils. Has consumption inequality mirrored income inequality? *American Economic Review*, 105(9):2725–56, 2015.

Mark Aguiar and Erik Hurst. Consumption versus expenditure. *Journal of political Economy*, 113(5):919–948, 2005.

Yacine Ait-Sahalia, Jonathan A Parker, and Motohiro Yogo. Luxury goods and the equity premium. *The Journal of Finance*, 59(6):2959–3004, 2004.

S Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994.

Andrew Atkeson and Masao Ogaki. Wealth-varying intertemporal elasticities of substitution: Evidence from panel and aggregate data. *Journal of monetary Economics*, 38 (3):507–534, 1996.

Andrew Atkeson, Masao Ogaki, et al. Engel's law and saving. Technical report, University of Rochester-Center for Economic Research (RCER), 1990.

Orazio Attanasio and Martin Browning. Consumption over the life cycle and over the business cycle, 1993.

Adrien Auclert, Matthew Rognlie, and Ludwig Straub. The intertemporal keynesian cross. Technical report, January 2017 working paper, 2018.

Laurent Bach, Laurent E Calvet, and Paolo Sodini. Rich pickings? risk, return, and skill in household wealth. *American Economic Review*, 110(9):2703–47, 2020.

Truman Bewley. A difficulty with the optimum quantity of money. *Econometrica: Journal of the Econometric Society*, pages 1485–1504, 1983.

Anmol Bhandari, David Evans, Mikhail Golosov, and Thomas J Sargent. Inequality, business cycles, and monetary-fiscal policy. *Econometrica*, 89(6):2559–2599, 2021.

Mark Bils and Peter J Klenow. Using consumer theory to test competing business cycles models. *Journal of Political Economy*, 106(2):233–261, 1998.

Timo Boppart, Per Krusell, and Kurt Mitman. Exploiting mit shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, 89:68–92, 2018.

Martin Browning and Thomas F Crossley. Luxuries are easier to postpone: A proof. *Journal of political Economy*, 108(5):1022–1026, 2000.

Isabel Cairó and Tomaz Cajner. Human capital and unemployment dynamics: Why more educated workers enjoy greater employment stability. *The Economic Journal*, 128(609): 652–682, 2018.

Xiaohong Chen, Jack Favilukis, and Sydney C Ludvigson. An estimation of economic models with recursive preferences. *Quantitative Economics*, 4(1):39–83, 2013.

Olivier Coibion, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89, 2017.

Diego Comin, Danial Lashkari, and Martí Mestieri. Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374, 2021.

Thomas F Crossley and Hamish W Low. Is the elasticity of intertemporal substitution constant? *Journal of the European Economic Association*, 9(1):87–105, 2011.

Angus Deaton. A reconsideration of the empirical implications of additive preferences. *The Economic Journal*, 84(334):338–348, 1974.

Cynthia Doniger. The ways the cookie crumbles: Education and the margins of cyclical adjustment in the labor market. 2021.

Andreas Fagereng, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik. Saving behavior across the wealth distribution: The importance of capital gains. Technical report, National Bureau of Economic Research, 2019.

John Fernald. A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco, 2014.

Martin Floden and Jesper Lindé. Idiosyncratic risk in the united states and sweden: Is there a role for government insurance? *Review of Economic dynamics*, 4(2):406–437, 2001.

Peter Ganong, Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat. Wealth, race, and consumption smoothing of typical income shocks. Technical report, National Bureau of Economic Research, 2020.

Rafael Guntin, Pablo Ottonello, and Diego Perez. The micro anatomy of macro consumption adjustments. Technical report, National Bureau of Economic Research, 2020.

Fatih Guvenen, Serdar Ozkan, and Jae Song. The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660, 2014.

Fatih Guvenen, Fatih Karahan, Serdar Ozkan, and Jae Song. Heterogeneous scarring effects of full-year nonemployment. *American Economic Review*, 107(5):369–73, 2017.

Jonathan Heathcote, Fabrizio Perri, and Giovanni L Violante. Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics*, 13(1):15–51, 2010a.

Jonathan Heathcote, Fabrizio Perri, and Giovanni L Violante. Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics*, 13(1):15–51, 2010b.

Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control*, 17(5-6):953–969, 1993.

Ayşe Imrohoroğlu. Cost of business cycles with indivisibilities and liquidity constraints. *Journal of Political economy*, 97(6):1364–1383, 1989.

John Kramer. The cyclicality of earnings growth along the distribution-causes and consequences. 2022.

Tom Krebs. Growth and welfare effects of business cycles in economies with idiosyncratic human capital risk. *Review of Economic Dynamics*, 6(4):846–868, 2003.

Dirk Krueger, Kurt Mitman, and Fabrizio Perri. On the distribution of the welfare losses of large recessions. Technical report, National Bureau of Economic Research, 2016a.

Dirk Krueger, Kurt Mitman, and Fabrizio Perri. Macroeconomics and household heterogeneity. In *Handbook of macroeconomics*, volume 2, pages 843–921. Elsevier, 2016b.

Per Krusell, Toshihiko Mukoyama, Ayşegül Şahin, and Anthony A Smith Jr. Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, 12(3): 393–404, 2009.

Per Krusell, Toshihiko Mukoyama, Richard Rogerson, and Ayşegül Şahin. Gross worker flows over the business cycle. *American Economic Review*, 107(11):3447–76, 2017.

François Le Grand and Xavier Ragot. Managing inequality over business cycles: Optimal policies with heterogeneous agents and aggregate shocks. *International Economic Review*, 63(1):511–540, 2022.

Robbert E Lucas. Models of business cycles basil blackwell. *New York*, 1987.

Robert E Lucas. Macroeconomic priorities. *American economic review*, 93(1):1–14, 2003.

Christopher J Malloy, Tobias J Moskowitz, and Annette Vissing-Jørgensen. Long-run stockholder consumption risk and asset returns. *The Journal of Finance*, 64(6):2427–2479, 2009.

Rajnish Mehra and Edward C Prescott. The equity premium: A puzzle. *Journal of monetary Economics*, 15(2):145–161, 1985.

Kanishka Misra and Paolo Surico. Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4):84–106, 2014.

Toshihiko Mukoyama and Ayşegül Şahin. Costs of business cycles for unskilled workers. *Journal of Monetary Economics*, 53(8):2179–2193, 2006.

Jacob Orchard. Cyclical demand shifts and cost of living inequality. *Available at SSRN 4033572*, 2022.

Jonathan A Parker and Annette Vissing-Jorgensen. Who bears aggregate fluctuations and how? *American Economic Review*, 99(2):399–405, 2009.

Jonathan A Parker and Annette Vissing-Jorgensen. The increase in income cyclicality of high-income households and its relation to the rise in top income shares. Technical report, National Bureau of Economic Research, 2010.

Jonathan A Parker, Nicholas S Souleles, David S Johnson, and Robert McClelland. Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6):2530–53, 2013.

Christina Patterson. The matching multiplier and the amplification of recessions. *Unpublished Manuscript, Northwestern University*, 2019.

Morten O Ravn and Harald Uhlig. On adjusting the hodrick-prescott filter for the frequency of observations. *Review of economics and statistics*, 84(2):371–376, 2002.

Morten O Ravn, Stephanie Schmitt-Grohé, and Martin Uribe. Macroeconomics of subsistence points. *Macroeconomic Dynamics*, 12(S1):136–147, 2008.

Matthew Smith, Owen Zidar, and Eric Zwick. Top wealth in america: New estimates under heterogeneous returns. *The Quarterly Journal of Economics*, 2022.

Thomas Steger. Economic growth with subsistence consumption. In *Transitional Dynamics and Economic Growth in Developing Countries*, pages 21–60. Springer, 2000.

Annette Vissing-Jørgensen. Limited asset market participation and the elasticity of intertemporal substitution. *Journal of political Economy*, 110(4):825–853, 2002.

Annette Vissing-Jørgensen and Orazio P Attanasio. Stock-market participation, intertemporal substitution, and risk-aversion. *American Economic Review*, 93(2):383–391, 2003.

Thomas Winberry. Lumpy investment, business cycles, and stimulus policy. *American Economic Review*, 111(1):364–96, 2021.

A Complete Markets

In this section I present a very simple model that serves two purposes: (i) it highlights the relation between non-homothetic preferences and consumption volatility, and (ii) it allows the derivation of useful analytical results.

Time is discrete. Markets are complete. There are types of agents: high-skilled and low-skilled. Agents differ in their endowment process: $\{y^{HS}(s^t)\}$ and $\{y^{LS}(s^t)\}$ are the endowments of high and low-skilled as a function of history s^t.

There are two goods: luxuries and necessities. Agents have access to a linear technology that allows them to produce one unit of each one of these goods using one unit of their endowments:

$$\begin{aligned} y_N =& y \\ y_L =& y \end{aligned}$$

Preferences are separable in time, state and goods, and agents maximize:

$$E_0 \sum \beta^t u(c_N(s^t), c_L(s^t))$$

$$u(c_N, c_L) = B\frac{c_N^{1-\frac{1}{\varepsilon}}}{1-\frac{1}{\varepsilon}} + \frac{c_L^{1-\frac{1}{\eta}}}{1-\frac{1}{\eta}}$$

subject to the time-0 budget constraint:

$$\sum_{t=0}^{\infty} \sum_{st} q^0\left(s^t\right) c^i\left(s^t\right) \leq \sum_{t=0}^{\infty} \sum_{st} q^0\left(s^t\right) y^i\left(s^t\right)$$

where $c^i(s^t) = c^i_N(s^t) + c^i_L(s^t)$ is the total consumption of individual i in history s^t.

Definition: A competitive equilibrium is a feasible allocation, $c^i = \{c^i\left(s^t\right)\}_{t=0}^{\infty}$, and a price system, $\{q^0\left(s^t\right)\}_{t=0}^{\infty}$, such that, given the price system, the allocation solves each household's problem.

Next I discuss some properties of the preferences and also characterize some predictions of the model for the consumption volatility of different types of households.

Consumption volatility

Now define the total consumption volatility for each type of household as the standard deviation of their log consumption:

$$\sigma^i := \sigma(\log c^i(s^t))$$

The following Lemma establishes that there is no difference in consumption volatility across types, when comparing the consumption of specific categories. It also establishes that the volatility of luxuries is higher than the volatility of necessities ($\theta > 1$).

Lemma 2: *In equilibrium:*

$$(i) \quad \sigma(\log c_N^{HS}(s^t)) = \sigma(\log c_N^{HS}(s^t)) = \sigma_N$$
$$(ii) \quad \sigma_L = \theta \sigma_N$$

The proof is left to Appendix D. □

Let $\bar{c}^i$ be the mean of the stochastic process $c^i(s^t)$, and $\bar{c}_N^i$ the mean of $c_N^i(s^t)$. Also, define the average consumption share of necessity for each type $s_n^i := \frac{\bar{c}_N^i}{\bar{c}^i}$.

The following Proposition derives an analytical expression for the ratio between the total consumption volatility across household types. It also establishes that θ and the consumption shares (s_n^i) are *sufficient statistics* for it.

Proposition 2: *To a first order:*

$$\frac{\sigma^{HS}}{\sigma^{LS}} \approx \frac{s_n^{HS} + (1 - s_n^{HS})\theta}{s_n^{LS} + (1 - s_n^{LS})\theta} \tag{19}$$

The proof is left to Appendix D. □

In this simple model, the only reason why low-skilled and high-skilled households may have different (total) consumption volatility is because they consume different baskets. In particular, if low-skilled households have lower consumption level, then, because of non-homotheticity ($\theta > 1$), they will have a larger share of their budget spent on necessities: $s_n^{LS} > s_n^{HS}$. As a consequence, Proposition 1 implies $\sigma^{LS} < \sigma^{HS}$. This is formalized in the following corollary.

Let $Y^i = \sum_{t=0}^{\infty} \sum_{st} q^0 (s^t) y^i (s^t)$ be the present value of income of a household of type i.

Corollary 2: *If preferences are non-homothetic ($\theta > 1$) and high-skilled households have a higher income level ($Y^{HS} > Y^{LS}$), then the consumption of high-skilled households is more volatile:*

$$\sigma^{HS} > \sigma^{LS}$$

This simple model illustrates a natural reason for why consumption dynamics may be higher for high-skilled households: they consume more luxuries, which are less costly to substitute over time. As a consequence, if different types are able to trade claims on every state of the world (markets are complete), then perfect risk-sharing across types naturally implies more consumption volatility for high-skilled.

Moreover, as shown in Proposition 1, the unequal consumption volatility across types is a direct function of $\theta, s_n^{LS}, s_n^{HS}$, which can be estimated from cross-sectional data on household expenditures, as I showed in Section 4.

B Data Appendix

B.1 Income measures - ASEC-CPS

As pointed in Heathcote et al. [2010b], over the sample period considered here, the set of income-related questions included in the March CPS has undergone two significant revisions, the first of which started with the 1975 income year and the second with the 1987 income year. However, neither the total income nor how it is distributed among the various income classes has been much impacted by these changes. Except in the case of private transactions.

Here I describe more specifically the measures used to construct my main income measure. This is also following closely Heathcote et al. [2010b]:

Labor Income: income from wage and salary

Self Employment Income: income from non-farm self-employment + income from farm or non-incorporated self-employment+income from own business, self-employment

Earnings plus: labor income +2/3 self-employment income + private transfers

Net Asset Income: income from interest, dividends and net rentals+rents and trusts

Public Transfers = income from public assistance or welfare+ pensions + income from social security + income from supplemental security+ income from educational assistance+ income from worker's compensation+ income from unemployment compensation.

Disposable Income = earnings plus + net asset income + public transfers - taxes, where taxes are imputed with TAXIM.

B.2 Consumption

B.2.1 Alternative smoothing parameters

Table 14: Consumption cyclicality: nondurables + services

	$\Delta c_{i,t}$		
	(1)	(2)	(3)
$\Delta \log C_t^{agg}$	0.92***		
	(0.044)		
$\Delta \log C_t^{agg} \times 1_{college}$	0.21**		
	(0.067)		
ΔU_t		-3.15***	
		(0.41)	
$\Delta U_t \times 1_{college}$		-1.87***	
		(0.80)	
Recession$_t$			-0.0013***
			(0.004)
Recession$_t \times 1_{college}$			-0.002*
			(0.001)
Observations	57,255	57,255	57,255
Adjusted R^2	0.007	0.007	0.006

Note: Newey-West standard errors in parentheses. ***p < 0.01, **p<0.05, *p<0.1.
Source: Consumer Expenditure Survey; author's calculations.

Table 15: Consumption cyclicality: nondurables + services

	$\Delta c_{i,t}$		
	(1)	(2)	(3)
$\Delta \log C_t^{agg}$	0.92***		
	(0.044)		
$\Delta \log C_t^{agg} \times 1_{college}$	0.21**		
	(0.067)		
ΔU_t		-2.56***	
		(0.37)	
$\Delta U_t \times 1_{college}$		-1.23***	
		(0.34)	
Recession$_t$			-0.0013***
			(0.004)
Recession$_t \times 1_{college}$			-0.002*
			(0.001)
Observations	57,255	57,255	57,255
Adjusted R^2	0.007	0.007	0.006

Note: Newey-West standard errors in parentheses. ***p < 0.01, **p<0.05, *p<0.1.
Source: Consumer Expenditure Survey; author's calculations.

Table 16: Consumption volatility across skill groups and consumption categories

	Consumption (1)	Services + Non-durables (2)	Non-durables (3)	Services (4)
All Households				
low-skilled	4.2	3.8	4.4	5.3
high-skilled	4.9	4.7	5.2	6.8
Ratio	**0.9**	**0.8**	**0.8**	**0.8**
Single Households				
low-skilled	8.7	7.7	8.6	12.3
high-skilled	10.1	9.7	10	15.2
Ratio	**0.9**	**0.8**	**0.9**	**0.8**

Source: Consumer Expenditure Survey; author's calculations. Numbers in the first two lines represent standard deviation of consumption expenditures, for a given household group. The third line shows the ratio between low-skilled and high-skilled.

C Wealth-varying elasticity of intertemporal substitution (EIS)

The optimal intra-temporal consumption allocation requires:

$$c_L^i(s^t) = B^{-\eta}\left(c_N^i(s^t)\right)^{\frac{\eta}{\varepsilon}} \tag{20}$$

Define the ratio between elasticities: $\theta := \frac{\eta}{\varepsilon}$. If $\theta = 1$ then c_L and c_N are proportional and there is no meaninful distinction between goods in this economy. If $\theta > 1$ then the budget share of c_L increases with the amount of total consumption - that is, c_L is a luxury.

Under the assumption of separable preferences, there is a tight connection between the ratio of income elasticities and the ratio of intertemporal elasticities across goods. In particular, luxury goods have a higher EIS. This reflects the intuitive idea that *luxuries are easier to postpone*. Browning and Crossley [2000] proves this statement for a more general class of preferences. Deaton [1974], Atkeson and Ogaki [1996] and Crossley and Low [2011] also discuss this connection.

To see this, consider the rate of return from s^t to period s^{t+1}: $R_{t,t+1} = \frac{q^0(s^t)}{q^0(s^{t+1})}$. It is straigtfoward to show that the elasticity of intertemporal substitution (EIS) for luxuries and necessities is:

$$EIS^L := -\frac{\partial \ln\left(c_L^i(s^{t+1})/c_L^i(s^t)\right)}{\partial \ln R_{t,t+1}} = \eta$$

$$EIS^N := -\frac{\partial \ln \left(c_N^i(s^{t+1})/c_N^i(s^t)\right)}{\partial \ln R_{t,t+1}} = \varepsilon$$

Note that these elasticities are constant and equal across agents. However, because households may have different consumption shares, the EIS for total consumption may vary. To see this, define the EIS for total consumption:

$$EIS(c_t^i) := -\frac{\partial \ln \left(c_{t+1}^i(s^{t+1})/c_t^i(s^t)\right)}{\partial \ln R_{t,t+1}}$$

Define the budget share of necessities for a household of type i as $s_n^i(s^t) := \frac{c_N^i(s^t)}{c^i(s^t)}$. The following Lemma, shows that we can write the EIS for total consumption as a weighted average between the elasticities of necessities (ε) and luxuries (η):

Lemma 1: *The EIS for total consumption can be written as:*

$$EIS(c_t^i) = \varepsilon s_n^i(c_t^i) + \eta(1 - s_n^i(c_t^i))$$

The proof is left to Appendix D. □

Lemma 1 is equivalent to Corollary 2 in Browning and Crossley [2000]. Given the assumption that c_L is a luxury ($\eta > \varepsilon$), Lemma 1 also reveals that households with a higher consumption level (and hence a higher luxury budget share), have higher elasticity of intertemporal substitution.

An EIS that varies with the level of consumption is an inherent feature of non-homothetic preferences, as discussed in Crossley and Low [2011],Comin et al. [2021], and Ait-Sahalia et al. [2004][20]

Under the assumption of separable preferences, not only the EIS is increasing with the level of consumption, but the rate of increase depends only on θ and the consumption shares, which, as I show in the next section, can be precisely estimated from cross-secional data. This is formalized in the following Corollary.

Corollary 1: *The ratio between the total consumption EIS for households with consumption levels c_1 and c_2 is:*

$$\frac{EIS(c_1)}{EIS(c_2)} = \frac{s_n(c_1) + \theta\left[1 - s_n(c_1)\right]}{s_n(c_2) + \theta\left[1 - s_n(c_2)\right]}$$

Corollary 1 follows immediately from Lemma 1, after sustituting $\theta = \frac{\eta}{\varepsilon}$ and rearranging

[20] For general class of preferences, Crossley and Low [2011] show that the assumption of constant EIS imposes strong restrictions on the shapes of Engel curves, and those restrictions are clearly rejected from the data. This is also discussed in Comin et al. [2021].

terms.

D Proofs

Proof of Lemma 1:

For a given household of type i, the intertemporal optimality condition implies:

$$\beta \left(\frac{c_L^i(s^{t+1})}{c_L^i(s^t)} \right)^{-1/\eta} R_{t,t+1} = 1$$

$$\beta \left(\frac{c_N^i(s^{t+1})}{c_N^i(s^t)} \right)^{-1/\varepsilon} R_{t,t+1} = 1$$

For a given consumption level: $c^i(s^t) = c_N^i(s^t) + c_L^i(s^t)$:

$$EIS(c_t^i) = \frac{c_N^i(s^t)}{c^i(s^t)} \left(-\frac{\partial \ln \left(c_N^i(s^{t+1})/c_N^i(s^t) \right)}{\partial \ln R_{t,t+1}} \right) + \frac{c_L^i(s^t)}{c^i(s^t)} \left(-\frac{\partial \ln \left(c_L^i(s^{t+1})/c_L^i(s^t) \right)}{\partial \ln R_{t,t+1}} \right)$$

Substituting the following conditions, establishes the result:

$$-\frac{\partial \ln \left(c_L^i(s^{t+1})/c_L^i(s^t) \right)}{\partial \ln R_{t,t+1}} = \eta$$

$$-\frac{\partial \ln \left(c_N^i(s^{t+1})/c_N^i(s^t) \right)}{\partial \ln R_{t,t+1}} = \varepsilon$$

□

Proof of Lemma 2:

Let λ^i be the Lagrange multiplier associated to individual's budget constraint. Complete markets implies:

$$\frac{c_L^{HS}(s^t)}{c_L^{LS}(s^t)} = \left(\frac{\lambda^{HS}}{\lambda^{LS}} \right)^{-\eta} \tag{21}$$

$$\frac{c_N^{HS}(s^t)}{c_N^{LS}(s^t)} = \left(\frac{\lambda^{HS}}{\lambda^{LS}} \right)^{-\varepsilon} \tag{22}$$

that is, the ratios between luxuries and necessities between high-skilled and low-skilled are constant at all times and states.

Equation 22 implies:

$$\sigma(\log c_N^{HS}(s^t)) = \sigma(\log c_N^{HS}(s^t)) = \sigma_N$$

From Equation 20:

$$\log c_L^i(s^t) = \alpha + \theta \log c_N^i(s^t)$$

where α is a constant. The standard deviation of luxuries is then:

$$\sigma_L = \theta \sigma_N$$

$\square$

Let $\bar{c}^i$ be the mean of the stochastic process $c^i(s^t)$, and $\bar{c}_N^i$ the mean of $c_N^i(s^t)$. Also, define the average consumption share of necessity for each type $s_n^i := \frac{\bar{c}_N^i}{\bar{c}^i}$.

The following Proposition establishes that the difference in consumption volatility across household types, up to a first order, can be explained solely by θ and consumption shares.

Proposition 2: *To a first order:*

$$\frac{\sigma^{HS}}{\sigma^{LS}} \approx \frac{s_n^{HS} + (1 - s_n^{HS})\theta}{s_n^{LS} + (1 - s_n^{LS})\theta}$$

Proof of Proposition 2:

The log of total consumption of an individual of type i is:

$$\log c^i(s^t) = \log \left[c_N^i(s^t) + B^{-\eta} \left(c_N^i(s^t) \right)^{\theta} \right]$$

The first-order Taylor approximation (around the mean) of the equation above yields:

$$\log c^i(s^t) - \log \bar{c}^i \approx \left[s_n^i + (1 - s_n^i)\theta \right] \left(\log c_N^i(s^t) - \log \bar{c}_N^i \right)$$

Note first that $[s_n^i + (1 - s_n^i)\theta]$ is a constant within types. Also, and from Lemma 1, $\sigma((\log c_N^i(s^t) - \log \bar{c}_N^i)) = \sigma_N$. Then:

$$\sigma^i \approx \left[s_n^i + (1 - s_n^i)\theta \right] \sigma_N$$

$\square$