Robust Consumption Inference: extracting recession expectations from a restricted consumption basket

Kieran Larkin*

12 March 2026

Abstract

This paper proposes a method for recovering households' expectations from consumption data when large disruptions cause the consumption basket to change dramatically for non-economic reasons. I propose a methodology for computing a new measure: robust consumption that is informative about prevailing economic circumstances. I study consumption during the Covid recession, but the methodology is applicable to any large structural disruption, such as pandemics or trade wars. Combining data with a sophisticated life-cycle model, I estimate households' expectations during 2020-22. Household beliefs are shown to have an important effect on the strength of the consumption response to government transfer policies.

Keywords: Consumption, Expectations, Uncertainty, Shock Persistence, Covid, Heterogeneous Agents

JEL Codes: D14, D15, D31, E21, E32

*Institute for International Economic Studies (IIES). Contact: kieran.larkin@iies.su.se. Institute for International Economic Studies. Stockholm University, Stockholm, SE-106 91, Sweden. I thank Andrew Atkeson, Orazio Attanasio, Richard Foltyn, Alexandre Kohlhas, Per Krusell, Kurt Mitman, Morten Ravn, Vincent Sterk, and Kjetil Storesletten for helpful suggestions and input. I thank conferences and seminars presentations at Oslo Macro Conference, Barcelona Summer Forum, Society of Economic Dynamics, Midwest Macro, UCL, Värmdö Kurtosis Conference, University of Bonn, University of Glasgow, Stockholm School of Economics, and Uppsala.

1 Introduction

Consumption patterns are highly informative about the economic environment facing households (Hall (1987), Blundell et al. (2008)). Large economic disruptions typically affect consumption via two distinct channels. First, they alter households' expectations about future income and employment (Friedman (1957), Kimball (1990)). This is the standard channel through which consumption theory permits inference about beliefs. Second, they change which goods and services households choose to purchase, through supply restrictions, regulatory changes, or shifts in the feasibility of specific activities. When both channels operate simultaneously, aggregate consumption conflates expectations-driven responses with compositional shifts, and standard inference breaks down.

The Covid recession generated an unprecedented 18 percent consumption decline in April 2020—far exceeding any prior U.S. downturn. However, unlike in most previous recessions, aggregate consumption provides a misleading guide to the economic forces facing households, during this period. The pandemic reorganized commerce and production through lockdowns, remote work, and voluntary social distancing. These disruptions altered which goods households could purchase and wanted to consume, independently of underlying economic conditions. This makes it challenging to estimate household expectations—a critical input for understanding the recession and recovery and evaluating policy responses. Notably, in the recovery, consumption remained depressed, leading policymakers into concluding stimulus was still needed. I will show that the economically relevant measure of consumption had recovered, suggesting the transfers risked generating excess demand.

This paper develops a method for extracting household expectations from consumption data when a large disruption causes the consumption basket to change dramatically. I define a new measure: *robust consumption*. This is a data driven consumption index, applicable during periods of large disruption to the economy, that can be interpreted as the consumption response as if only economic drivers were present. The purpose is to enable the measurement of household beliefs in unstable settings, using the tools commonly used to measure beliefs in normal times. For the Covid recession, I construct a robust consumption and robust durable consumption measure from individual goods and services unaffected by pandemic-specific factors. Combining this measure with a quantitative life-cycle model, I identify how household expectations evolved, in particular: the persistence households assigned to the aggregate state and the level of idiosyncratic risk they perceived they faced. Computing these two beliefs enables me to conduct policy counterfactuals. While the Covid recession is a relevant application, the methodology applies more generally . During any large structural disruption – pandemics, trade wars, wars, natural disasters – forces that alter the composition of the consumption basket operate alongside the income and employment shocks. I apply the methodology to the early 1980s

recession, to demonstrate its broader applicability and provide a welfare comparison for the Covid recession.[1]

The first step is to partition aggregate consumption into two baskets using high-frequency data on narrow product categories. The partition identifies goods and services that experienced extreme consumption changes in April 2020, when pandemic restrictions peaked, while constraining both baskets to exhibit the same business cycle properties– volatility and cyclicality– prior to the Covid recession. This means the robust measure can be interpreted as responding to economic shocks during 2020-22 as the standard consumption basket would in a normal recession. The robust measure reveals patterns that aggregate data obscures. Robust non-durables and services consumption fell 1.1 percent at the trough, still substantial by historical standards, but rebounded strongly before the end of 2020. Robust durables consumption also declined sharply (14 percent).[2] I map the baskets to household-level data using the Consumer Expenditure Survey, documenting substantial heterogeneity by education. High-education households reduced robust consumption more persistently throughout the recovery, while low-education households cut durables more sharply at the recession's onset. These patterns are consistent with differential expectations about the recession's duration and income risk.

Having constructed a measure of robust consumption, I embed the series in a sophisticated heterogeneous-agent life-cycle model featuring three consumption goods (robust, excluded, and durables), stochastic adjustment costs for durables, borrowing constraints, and rich income dynamics.[3] The recession is simulated with the observed aggregate shocks to productivity, employment transitions, fiscal transfers, and prices. I allow changes in the relative price of the robust consumption basket to affect consumption patterns, capturing economic supply and demand factors during the period. I use model consistent, household decision-making to infer the beliefs implied by the path of robust consumption expenditures. This can be understood as the heterogeneous agent extension of inverting the household Euler equation, after accounting for the changing consumption basket. As motivation, I first present a simple example to show how robust consumption can be used to infer these beliefs. I am able to estimate quarterly paths of two key objects: the persistence households assign to the aggregate state, ϱ_t^e, and the level of idiosyncratic income uncertainty they perceive, $\sigma_{z,t}^e$, both by education level, e. Robust non-durable consumption primarily identifies perceived persistence, while robust durables consumption identifies uncertainty.[4]

[1] The 1979 energy crisis caused by the Iranian revolution is often cited as a contributing factor to the 1980s recession and can be considered a large structural disruption.

[2] This somewhat contradicts the conventional wisdom that Covid was purely a services-driven recession (Beraja and Wolf (2021)).

[3] Excluded consumption could still influence the inferred expectations if lower spending relaxes credit constraints facing households with low assets. Further, the substitutability between robust and excluded consumption goods will influence the mapping from revealed preferences to expectations.

[4] Durables sensitivity provides crucial identifying variation: durables respond more strongly than non-durables to changes in uncertainty, while non-durables reveal information about perceived persistence (Eberly (1994), Bertola et al. (2005), Harmenberg and Öberg (2021)).

The estimates point clearly to dynamic expectations: households gradually learned about the recession's severity rather than immediately anticipating its full trajectory. Initially, households perceived shocks as highly transitory, explaining the modest consumption response despite severe labor market disruptions. Uncertainty followed a different path, rising more gradually and peaking in early 2021—coinciding with the second fiscal stimulus. Beliefs differ systematically by education: high-education households consistently perceive greater shock persistence, while low-education households face substantially larger increases in economic uncertainty. One interpretation is that higher-education households anticipated longer-run changes to the way their jobs would be affected, while lower-education households were more affected by disruptions associated with short-run restrictions. The estimated household beliefs align closely with independent survey evidence on consumer expectations, validating the model's rationalization of the data.

Understanding household expectations proves critical for evaluating pandemic-era policies. First, regarding inflation expectations: the model shows that rising inflation in 2021 drove strong durables spending by reducing real interest rates. However, the observed durables response reveals households expected this inflation episode to be fairly transitory; otherwise, the durables boom would have been three times larger.[5] Second, on fiscal stimulus timing: while transfers substantially boosted consumption, the second stimulus, which took place in early 2021, arrived after robust consumption had already returned to trend. Total consumption remained depressed, misleading policymakers into thinking stimulus was still needed. I show robust consumption had recovered, suggesting the transfers risked generating excess demand. This helps explain the subsequent inflation surge. Examining consumption dynamics through the lens of robust consumption reveals the economy was closer to potential than the aggregate data suggested. Third, policy uncertainty: given plausible uncertainty about household beliefs in late 2020, policymakers faced a large range in predicted consumption responses to the second stimulus, particularly for durables. Better measures of household expectations could improve policy design. The method I propose could be implemented by policy makers in close to real time, subject to lags in data publication.[6]

I find the economic component of the Covid recession imposed welfare costs of 1.2-1.9 percent of lifetime consumption. Low-education households suffered greater average losses, though some high-education households experienced severe losses following job separations. Fiscal transfers offset 20-30 percent of these costs, indicating that the economic shocks of the recession were still large relative to the sizable government support. The welfare costs were around half of the 1980s recession.

[5]This inference aligns with Federal Reserve communication at the time (Powell (2021)).
[6]The key measure: monthly NIPA consumption expenditures is currently published with a six week lag, but this could be expedited.

This paper makes three contributions. First, methodologically, I develop a general framework for extracting household expectations from consumption when preferences or constraints shift dramatically. While there are standard methods for inferring economic circumstances from consumption choices, these only work in stable environments. I extend the methods to unstable environments, which are likely to occur at the time of large economic disruptions. Second, I provide the first structural estimates of how household expectations evolved during the Covid recession and recovery and document substantial heterogeneity across education groups. Third, I demonstrate that expectations are crucial for evaluating fiscal and monetary interventions in this period. Optimal countercyclical policy requires correctly measuring economic conditions *and* the evolution of household expectations.

Related literature: This paper contributes to the large literature that uses consumption data to understand the economic environment facing households and the decisions they make. At its core, such an approach goes back to the concept of revealed preferences (Samuelson, 1948). Aggregate consumption is one of the most closely studied business cycle variables (Blanchard (1993), Meyer and Sullivan (2013), Berger and Vavra (2015), Attanasio et al. (2022)), and consumption underpins welfare calculations in nearly all macroeconomic models (Lucas (2003), Baqaee et al. (2024)). Consumption patterns in the micro data have been studied extensively (Zeldes (1989), Attanasio and Weber (1995), Attanasio and Pistaferri (2016)).[7]

Consumption choices have been used to identify unobserved or difficult-to-measure economic processes. A leading example of this is the study of income processes and quantifying the risk facing households (Heathcote et al. (2014), Blundell and Preston (1998), Blundell et al. (2008), Low et al. (2010), Guvenen and Smith (2014), Commault (2022)). In this setting, consumption is useful for identification because the permanent income hypothesis predicts consumption should respond very differently depending on the transitory or permanent nature of an income shock (Friedman, 1957). Guvenen et al. (2024) study the implications of high-order income moments for consumption patterns and the responsiveness to stimulus checks. Fagereng et al. (2024) studies the consumption response to unemployment. Aggregate consumption has also been used extensively to study business cycles (Gorbachev, 2011), including the recent Covid recession (Stock and Watson, 2025).

This paper takes a structural approach.[8] In this paper, I study the joint non-durable and durable goods consumption patterns during a recession. Closely related to this is Berger

[7]The literature has argued that the canonical life-cycle models of household behavior can replicate the patterns observed in the data (Hubbard et al. (1995), Gourinchas and Parker (2002), Attanasio and Weber (2010)). An implication of this is that a quantitative structural consumption model can serve as a laboratory for interpreting patterns in the data and studying policy interventions.

[8]The incomplete markets model of the household problem goes back to Imrohoruglu (1989), Aiyagari (1994), Huggett (1993). Recent contributions have emphasized the heterogeneous consumption response in this class of models. In particular, the role of savings frictions and the importance of matching the empirical Marginal Propensity to Consume (MPC) (Kaplan and Violante (2014), Kaplan et al. (2016), Auclert et al. (2024)).

and Vavra (2015), who emphasize that durables adjust more sluggishly in response to recessions. Attanasio et al. (2022) study car spending in the Great Recession, arguing it provides a powerful way of identifying aggregate shocks. McKay and Wieland (2021) studies the implications for monetary policy. Further, I relate to a growing literature that uses survey data to measure household expectations. For an overview, see: D'Acunto et al. (2024). Burke and Ozdagli (2023) documents the link to consumption responses. D'Acunto et al. (2021) emphasizes the role of expectations in the design of policy. Pappa et al. (2023) investigates the interaction with incomplete markets. Guiso and Jappelli (2025) study the household expectations and consumption uncertainty. The approach here differs in that I infer expectations from spending data through the lens of a structural model.

In studying the Covid recession, this paper is closely related to the literature on the macroeconomics of a pandemic-driven recession, perhaps most notably Atkeson (2020) and Eichenbaum et al. (2021).[9] The methodology of these papers is to incorporate an SIR model from the epidemiological literature into a macroeconomic model to understand how pandemic dynamics might proceed and solve for the optimal policy. The most closely related paper is Eichenbaum et al. (2024), who also use consumption data to uncover household expectations, but focus on the expectations of the pandemic, particularly mortality. The approach here is quite different. I do not build an SIR model into the heterogeneous agents framework, but instead focus on the economic determinants of households' behavior during the recession. The approach is closer to a standard macroeconomic analysis of the recession, which has so far been overlooked in the literature. Similar in theoretical motivation is Lenza and Primiceri (2020), who addresses the approach to estimating VARs in the presence of extreme observations such as the Covid decline.

The research is also related to a set of papers that study real-time spending during the Covid recession. Many of these studies utilized proprietary transaction-level datasets. Chetty et al. (2023) build a large database studying economic activity during the Covid recession. They document high-frequency consumer spending patterns during the recession, emphasizing a larger decline in spending by high-income households. Cox et al. (2020) document similar patterns in a different dataset.[10]

Finally, the approach is conceptually related to the consumption-asset pricing literature where properties of the time-varying returns structure are estimated from the Euler equation (see: Hansen and Singleton (1983), Kleibergen and Zhan (2020), Di Tella et al. (2026)). Here, I also start from the Euler equation, with rich heterogeneity, but ask what household beliefs rationalize consumption patterns, given returns.

[9]Other examples include: Kaplan et al. (2020), Alvarez et al. (2021), Berger et al. (2022), Boppart et al. (2022), Faria-e Castro (2021), Farboodi et al. (2021) and Glover et al. (2023).

[10]In other countries: Andersen et al. (2020) study the response in Scandinavia, focusing on the role of restrictions, which varied by country, while Carvalho et al. (2021) study Spanish transaction data. The contribution of this paper is to combine empirical analysis with a quantitative structural model to interpret consumption patterns and infer household expectations.

The rest of the paper is structured as follows: Section 2 discusses the empirical methodology, Section 3 sets out the model, Section 4 uncovers household expectations, Section 5 evaluates policy, Section 6 concludes.

2 Consumption during the Covid recession and recovery

2.1 A simple robust consumption example

Before estimating the robust consumption basket, I present a simple example demonstrating how robust consumption can be used to infer households' expectations about economic shocks. Consider a two-period, two-good model where robust consumption, c_t^r, captures the component of consumption driven solely by economic shocks, net of non-economic factors.[11] A recession occurs as a negative income shock, z, in period t, with persistence ϱ affecting period $t+1$ income. Additionally, period $t+1$ income is uncertain. With probability λ income is σ lower than expected and $1-\lambda$ it is σ higher than expected. Let R denote the return on a bond and y the steady state endowment. The household's consumption Euler equation is:

$$\theta^{\frac{1-\rho}{\epsilon-1}}\left(c_t^r\right)^{-\rho} = \left(\frac{\Lambda}{\phi}\right)^{1-\rho}\Big[\lambda\left(y(1+R)-(\varrho+R)z-\sigma-Rc_t^r\right)^{-\rho}+ \\ (1-\lambda)\left(y(1+R)-(\varrho+R)z+\sigma-Rc_t^r\right)^{-\rho}\Big] \tag{1}$$

This equation reveals the identification strategy: given the current income shock z and observed robust consumption c_t^r, one can recover either the expected persistence ϱ or the uncertainty σ. Including durable goods, consumption provides sufficient variation to identify both parameters simultaneously. The term $(\Lambda/\phi)^{1-\rho}$ adjusts preference weights for changes in demand for excluded goods driven by non-economic factors; this adjustment is computed from the intratemporal optimality condition.[12] While equation (1) illustrates the core intuition, the actual estimation incorporates additional realistic features absent from this simple example, including borrowing constraints, a more general income process, and dynamic expectations over multiple periods, as detailed in the full life-cycle model in Section 3.

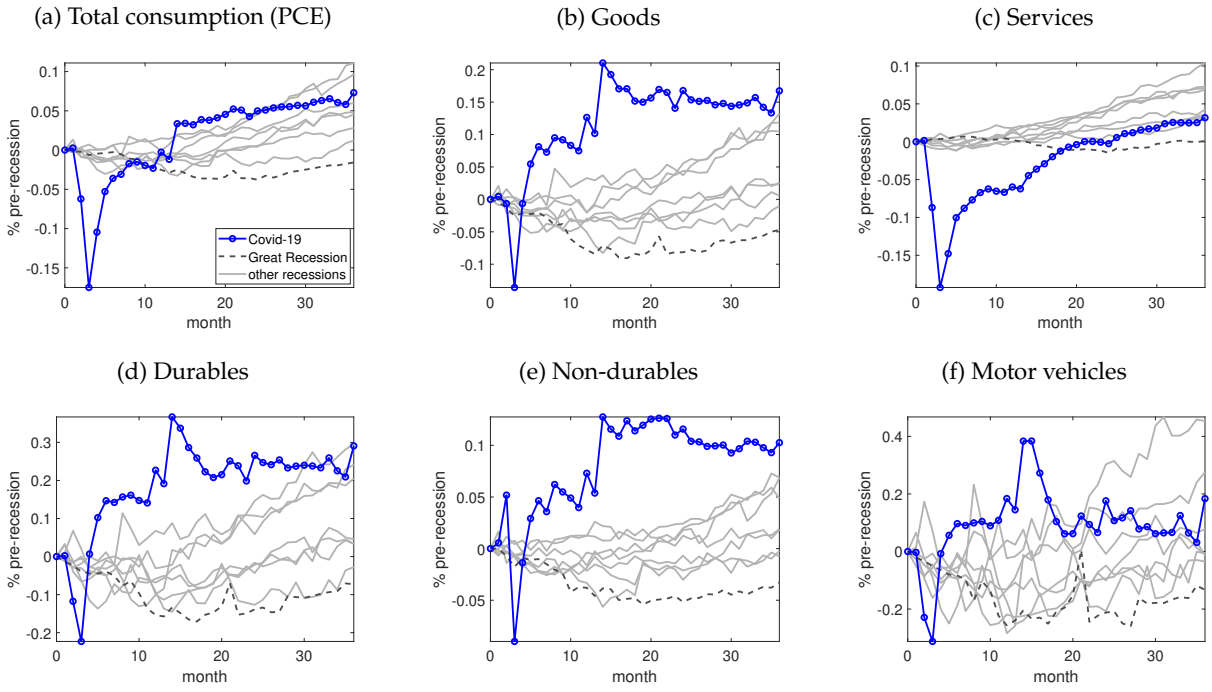
2.2 Components of consumption

Consumption fell dramatically at the onset of the Covid recession, by a magnitude not experienced in any recent downturn. Figure I plots consumption at a monthly frequency relative to its pre-recession level for the major consumption categories during the three years following the Covid downturn. The data is from the Bureau of Economic Analysis

[11] For the full derivation see Appendix Section A. The Euler equation only applies to an unconstrained household.
[12] I show for the case where the weight on the excluded good goes to zero in period t. $\chi \to 1$ in notation of Appendix.

(BEA) National Income and Product Accounts (NIPA).[13] Grey lines show other recessions, with the Great Recession highlighted by the dashed black line. Panel a plots the full consumption basket: Personal Consumption Expenditures (PCE). Aggregate consumption declined by more than 15 percent in April 2020 relative to its pre-recession level, dwarfing the typical recession decline and far exceeding the Great Recession's peak decline of 6 percent. A second notable feature is the rapid rebound: by the end of the three-year window, total consumption stands 5 percent above its initial level and exceeds the average recovery of previous recessions. In contrast, after the Great Recession, consumption remained below its pre-recession level for three years.

Figure I: Monthly consumption decline



Notes: Data from the BEA's NIPA accounts (1959-2022). Path of consumption during recession. Deviations relative to period prior to NBER recession date. Blue line is Covid recession, black dashed line is Great Recession, grey lines are other recessions.

Disaggregating consumption into major components reveals rich variation in household behavior during the recovery. Panels b and c compare goods versus services. While both fell initially, their recovery paths diverged sharply. Goods spending returned above its pre-recession level within months (Panel b), whereas service consumption remained depressed for far longer (Panel c). The initial services decline reflects restrictions on restaurants and live entertainment, which effectively closed this sector during the lockdown. More striking is the persistent weakness: even after 3 years, total service spending remained below the average recession recovery path. Panels d and e decompose goods into durables and non-durables. Durable goods underwent a substantial initial decline, larger than but comparable to previous recessions, contradicting the conventional characterization of Covid as a purely services-driven downturn (Panel d). This decline proved short-lived: after

[13]Monthly consumption data is available from 1959; it is seasonally adjusted.

12 months, durables spending exceeded pre-recession levels by 30 percent. Non-durable consumption displays evidence of precautionary stockpiling, with strong initial growth followed by a sharp decline (Panel e). Finally, Panel f examines motor vehicles, the largest durables type. The decline here resembles other significant recessions in magnitude. Notably, motor vehicle spending fell most sharply in the month before other spending components declined, suggesting forward-looking behavior by households anticipating the downturn.[14]

2.3 Recovering a robust consumption measure

2.3.1 Empirical approach

As shown in Section 2, consumption declined dramatically during the Covid recession, yet much of this decline cannot be attributed to economic shocks. Government restrictions and pandemic-induced preference shifts substantially altered both the basket of goods available to households and their desired consumption patterns. This section proposes a method for recovering a *robust consumption* measure: a consumption series whose dynamics reflect household behavior in a typical recession facing *economic shocks of the magnitude* observed during Covid, but absent the pandemic-specific disruptions. The approach partitions consumption into two baskets. The first contains goods and services affected by non-economic factors–supply restrictions, health concerns, and changed preferences driven by the pandemic itself (excluded consumption). The second contains goods and services unaffected by these pandemic-specific forces. I call this robust consumption. This basket isolates the consumption response to economic shocks and reveals households' forward-looking behavior and expectations about income, employment, and the recession's persistence. Section 4 uses this series to infer household expectations through a structural life-cycle model. The empirical implementation exploits monthly aggregate consumption data disaggregated into narrow sub-types. I partition the consumption basket by identifying goods and services that experienced extreme consumption changes in April 2020–when pandemic restrictions and health concerns simultaneously affected all households. Crucially, I impose the constraint that both baskets exhibit similar business cycle sensitivity before 2020. This constraint ensures the method does not simply separate high versus low business-cycle-sensitivity goods, but rather isolates the pandemic-specific component of the consumption decline from the economic-shock component. This method is a purely data-driven methodology and does not rely on categorizing goods by the type of spending.

Partitioning consumption: I implement the partition using detailed monthly consumption data disaggregated by goods and services sub-categories, also referred to as a type

[14]The Covid Recession also featured significant dislocation in the labor market (See: Cajner et al. (2020), Ganong et al. (2020), Hall and Kudlyak (2022), and Chetty et al. (2023)). Change in employee compensation, transfer receipts, and job flows are shown in the Appendix, Section E.3.1.

or product.[15] For a given consumption aggregate (e.g., non-durables and services, $\mathbf{C_t}$, or durables, $\mathbf{I_t^d}$), the underlying data contain spending on J narrow sub-types. The real value of spending on type j is denoted c_t^j. Conceptually, a partition of the sub-types into two baskets: robust consumption, $\mathbf{C_t^R}$, and excluded consumption, $\mathbf{C_t^O}$, can be written:

$$\mathbf{C_t} = \left(\mathbf{C_t^R} = \sum_{j \in C^R} c_t^j\right) + \left(\mathbf{C_t^O} = \sum_{j \in C^O} c_t^j\right) \equiv \sum_{j=1}^{J} c_t^j \tag{2}$$

where $C^R \equiv j \in \{1, \ldots, J^R\}$ and $C^O \equiv j \in \{1, \ldots, J^O\}$ denote the sets of product types in the robust and excluded basket, respectively. These are of size J^R and J^O, with $C^R \cup C^O = C$, $C^R \cap C^O = \emptyset$, and $J = J^R + J^O$. The sets C^R and C^O are unknown objects to be found.

Initial partition by size of shock: There are 181 sub-types for nondurable goods and services and 43 sub-types for durables, available monthly from 1959 (753 observations). I deflate each sub-type by its product-specific price deflator. An initial partition, $\hat{C}^R$ and $\hat{C}^O$, assigns product types based on whether they experienced extraordinarily large consumption changes in April 2020 relative to their historical volatility. For each sub-type j, I calculate the average log growth rate $\bar{g}^j = \frac{1}{T}\sum_{t=1}^{T} g_t^j$ and standard deviation $\sigma(g^j)$ using pre-2020 data, where $g_t^j = \log(c_t^j) - \log(c_{t-1}^j)$. I then assign sub-types to the excluded basket if their normalized consumption change in April 2020 exceeds γ times their historical standard deviation:[16]

$$j \in \begin{cases} \hat{C}^O & if\ |g_{t^*}^j| > \gamma\sigma(g_j) \\ \hat{C}^R & otherwise \end{cases} \tag{3}$$

where t^* = April 2020 for most types. The threshold parameter $\gamma = 3.8$ is calibrated through Monte Carlo analysis to ensure the robust basket exhibits similar business cycle properties to the full consumption aggregate in pre-Covid recessions (see Appendix B for details). This criterion ensures products with abnormally large April 2020 declines, by historical standards, are allocated to the excluded basket, isolating pandemic-specific disruptions. By focusing on the change in a single month, I allow other medium-run economic pandemic changes to still play a role, for example, through prices. I apply this procedure separately to non-durables and services ($\mathbf{C_t}$) and to durables ($\mathbf{D_t}$), yielding distinct robust and excluded baskets for each aggregate. For brevity, I sometimes refer to non-durables and services consumption as "consumption" and durable consumption as "durables".

Consistent business cycle properties: minimization step: The initial partition based on equation (3) does not guarantee that the robust basket, $\hat{\mathbf{C}}^\mathbf{R}$, and excluded basket, $\hat{\mathbf{C}}^\mathbf{O}$,

[15] These are reported in NIPA Table 2.4.4u. Note: not all sub-categories are available for the full series length. I use the longest length possible.

[16] To explain the methodology, I focus on products with extreme declines due to pandemic restrictions. But the method also excludes products with extreme increases (e.g., from stockpiling behavior). For motor vehicles, I use March 2020 as t^* because spending fell one month earlier, exhibiting anticipatory behavior.

have similar business cycle properties. If the baskets differ systematically in their income elasticities or volatilities, inference about household expectations from robust consumption would be confounded.[17] I therefore refine the initial partition through an optimization algorithm that reassigns product types to minimize differences in business cycle properties between the two baskets while maximizing their differential response during April 2020.

Starting from the initial partition, the algorithm sequentially reassigns products to minimize the following objective function:

$$\{\omega_j\}_{j=1}^{J} = \underset{\omega_j \in \{\varpi, 1-\varpi\}}{\operatorname{argmin}} \underbrace{\Gamma\left\{\left|\sigma_t(\mathbf{C_t^R})/\sigma_t(\mathbf{C_t^O}) - 1\right| + \left|\rho_t(\mathbf{C_t^R}, Y_t) - \rho_t(\mathbf{C_t^O}, Y_t)\right|\right\}}_{\text{business cycle properties}} \tag{4}$$

$$\underbrace{-\left\{\frac{1}{J^O}\sum_{m \in C^O}|\tilde{g}^m_{t^*}| - \frac{1}{J^R}\sum_{m \in C^R}|\tilde{g}^m_{t^*}|\right\}}_{\text{covid response difference}}$$

$$\underbrace{+\left\{\sigma_{\tilde{g}}\left(|\tilde{g}^m_{t^*}| \,\middle|\, m \in C^R\right) + \sigma_{\tilde{g}}\left(|\tilde{g}^m_{t^*}| \,\middle|\, m \in C^O\right)\right\}}_{\text{within group heterogeneity}}$$

$$\text{subject to:} \qquad \mathbf{C_t^R} = \textstyle\sum_{j=1}^{J} \omega_j c_t^j \qquad \text{and} \qquad \mathbf{C_t^O} = \textstyle\sum_{j=1}^{J} (1-\omega_j) c_t^j$$

where $\tilde{g}^j = |g^j_{t^*}|/\sigma(g_j)$ denotes the normalized change in April 2020. The algorithm assigns each product type j a weight $\omega_j \in \{\varpi, 1-\varpi\}$ in the robust basket, where $\varpi \approx 1$ effectively allocates the product to either robust ($\omega_j = \varpi$) or excluded ($\omega_j = 1-\varpi$) consumption.[18]

The first line of equation (4) minimizes differences in pre-2020 business cycle properties: the ratio of standard deviations, $\sigma_t(\cdot)$, and the difference in correlations with aggregate income $\rho_t(\cdot, Y_t)$. The parameter Γ governs the relative weight on matching business cycle properties versus maximizing the April 2020 differential response. I set $\Gamma = 7.75$ based on Monte Carlo analysis (see Appendix B).[19] The second line maximizes the difference in average normalized consumption declines between the two baskets in April 2020, ensuring the excluded basket captures products most affected by pandemic-specific factors. The third line minimizes within-basket variance in normalized consumption changes, producing baskets with internally homogeneous April 2020 responses.The algorithm iterates until no reassignment reduces the objective function. Appendix C provides further implementation details.

[17] For example, if robust consumption systematically contains less income-elastic goods, observed consumption dynamics would reflect compositional differences rather than expectational responses to the recession.

[18] Setting $\varpi < 1$ allows products to contribute partially to both baskets, improving the algorithm's ability to match business-cycle properties. Following Monte Carlo analysis (see Appendix B), I set $\varpi = 0.99$. I also experimented with choosing $\varpi_j \in [0.5, 1] \forall j$

[19] See Appendix Figure A.6 for robustness to alternative hyperparameters. For aggregate income we use compensation of employees.

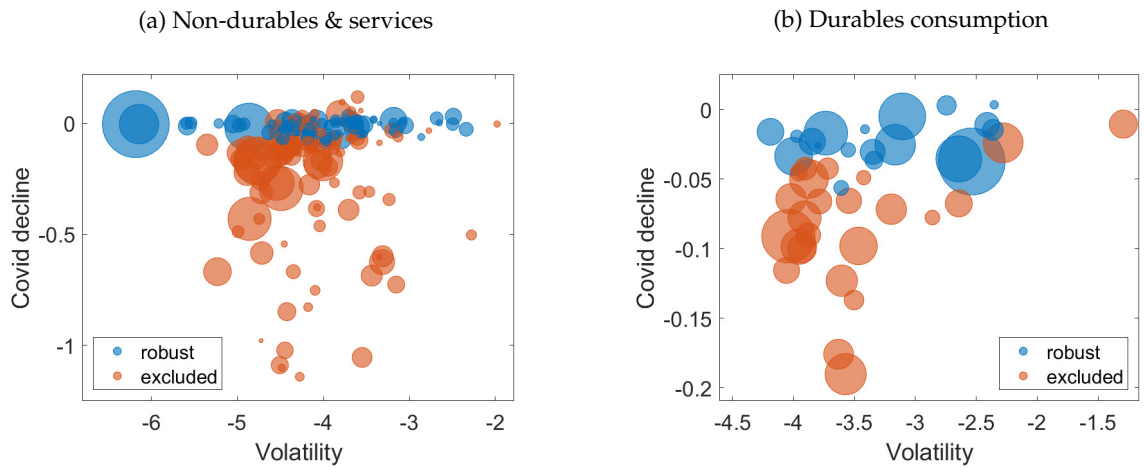
Robust Consumption definition: A *robust consumption* partition is a set of weights $\{\omega_j\}_{j=1}^J$ for J consumption sub-types, which allocates each sub-type j to the robust $\omega_j = \varpi > 0.5$ or excluded basket. The basket maximizes the normalized average difference in consumption growth $|\tilde{g}^j_{t^*}|$ at date t^*, while i.) minimizing within group variance in $|\tilde{g}^j_{t^*}|$ and ii.) constraining the aggregate baskets $\mathbf{C^R_t}$ and $\mathbf{C^O_t}$ to have the same standard deviation $\sigma_t(\mathbf{C^x_t})$ and co-movement with aggregate income $\rho_t(\mathbf{C^x_t}, Y_t)$ for $t < t^*$. The weights $\{\omega_j\}_{j=1}^J$ satisfy the objective function (4) and aggregation requires that all consumption goods are allocated satisfying (2).

2.4 A robust consumption measure

2.4.1 Empirical consumption baskets

Figure II displays the resulting partition. Each bubble represents a consumption sub-category, with its size indicating average monthly expenditure between 2010 and 2020. Blue bubbles denote products allocated to the robust basket; red bubbles denote products in the excluded basket. The x-axis plots the log standard deviation of pre-2020 growth rates, $\log \sigma_t(g^j)$. The y-axis plots the April 2020 consumption change, $g^j_{t^*}$.

Figure II: Properties of consumption partition



Notes: Partition of consumption goods into robust and excluded baskets. X-axis is (log for visual clarity) average standard deviation prior to Covid recession, $\log(\sigma)$. Y-axis decline in 3rd month of recession. Area is consumption share of sub-type. I use A few products with extreme declines $g^j_{t^*} < -15$ are omitted from the figure.

The figure reveals several key features of the partition. First, the magnitude of the Covid recession's impact is striking: numerous non-durables and services experienced declines exceeding 50 log points (Panel a). Second, the volatility distributions of the two baskets substantially overlap, suggesting that they exhibit similar business cycle sensitivity before Covid. This overlap validates the optimization procedure: the partition successfully isolates pandemic-specific disruptions without simply separating high- and low-cyclicality products. Third, the asymmetry in the partition differs across consumption types. For durables (Panel b), all excluded products experienced April 2020 declines. For non-durables and services (Panel a), the excluded basket also contains products with

large *increases*—likely reflecting stockpiling behavior rather than responses to economic fundamentals.[20] Fourth, two large, low-volatility expenditures appear in the robust basket: rental of tenant-occupied property and imputed rental of owner-occupied housing (Panel a). While their allocation is economically sensible—housing services were largely unaffected by pandemic restrictions—one might worry these stable expenditures dominate the robust basket's dynamics. However, they account for only 38 percent of robust consumption spending and contribute negligibly to short-run fluctuations, ensuring the robust measure captures meaningful variation in household behavior.

Table I: Consumption partition summary statistics

	Non-durables & services		Durables consumption	
Moment	$\mathbf{C^R}$	$\mathbf{C^O}$	$\mathbf{ID^R}$	$\mathbf{ID^O}$
Share of consumption basket	0.39	0.61	0.34	0.66
Standard dev. relative to income	0.53	0.57	2.91	2.91
Correlation with all consumption	0.84	0.94	0.69	0.73
Correlation with income	0.81	0.72	0.61	0.64
Correlation with $\mathbf{C^R}$	1.00	0.67	0.49	0.57

Notes: Constructed robust and excluded goods. Data from NIPA.

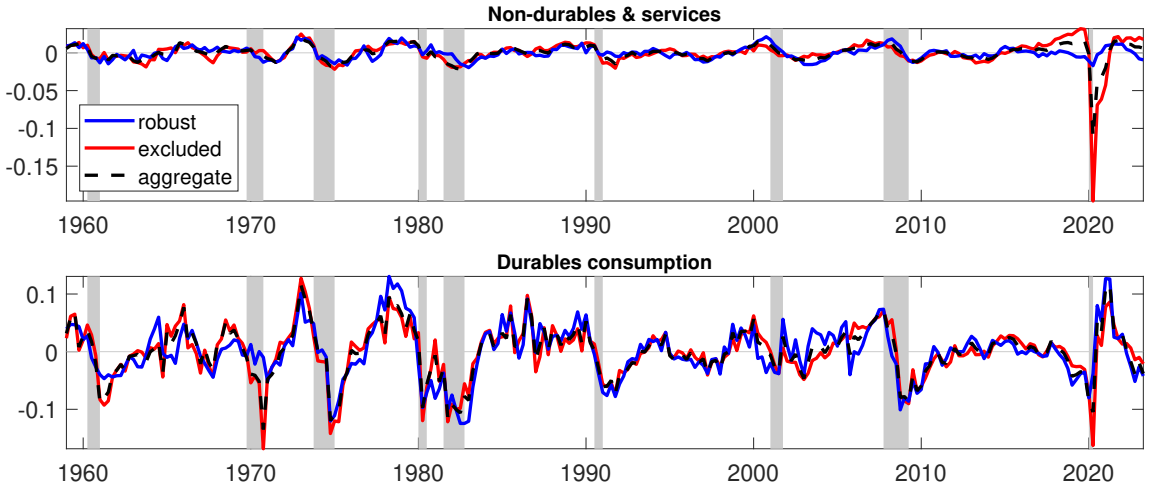
Table I reports summary statistics for the partition. The robust partition accounts for 39 percent of non-durables and services spending, indicating that a substantial share of the consumption basket remains informative about household expectations despite excluding pandemic-affected types. The table also compares pre-2020 business cycle properties of the two baskets. The robust basket exhibits slightly lower volatility (0.53 versus 0.57 for the excluded basket) and a stronger correlation with income (0.81 versus 0.72 for the excluded basket), but these differences are economically small, confirming that the optimization successfully balanced matching business cycle properties with isolating pandemic-specific effects.

Figure III provides the most compelling validation by comparing HP-filtered consumption paths for the two baskets. Prior to the Covid recession, the series move closely together, demonstrating similar cyclical behavior across previous recessions. During Covid, however, the paths diverge sharply: the excluded basket exhibits the extreme volatility characteristic of pandemic disruptions, while the robust basket follows a more typical recession trajectory.[21] This stark divergence during Covid, combined with close co-movement beforehand, confirms that the robust consumption measure successfully isolates the component of consumption driven by economic shocks while excluding pandemic-specific disturbances.

[20] Excluding these high-growth products prevents upward bias in the robust consumption measure from pandemic-specific behaviors uninformative about household expectations.

[21] The slight positive deviation in the robust series immediately before Covid reflects the HP filter's local smoothing window. This artifact disappears when estimation uses only pre-2020 data.

Figure III: HP filtered consumption series



Notes: HP filtered monthly data on robust (blue) and excluded (red) consumption. NBER recession dates are shaded. Top panel consumption, bottom panel durables consumption. Data from NIPA.

2.4.2 Response during the recession

I now examine what the partition reveals about households' responses to economic versus pandemic-specific shocks during the Covid recession. Figure IV plots robust and excluded consumption for non-durables and services (top row) and durables (bottom row). For brevity non-durables and services is referred to as "consumption."

Non-durables and services: Panel a shows that the robust measure undergoes a modest decline at the recession's onset, falling only 1.1 percent below its pre-recession level in March 2020. In contrast, the excluded composite declines by nearly 30 percent (Panel b). Of course, this divergence reflects the partition methodology. Yet the peak 1.1 percent decline is still historically large. Only one prior recession exhibits a significantly larger drop in the robust basket at the onset of the recession: 1960, with a 3.2 percent peak decline. The robust measure decline in other recessions is similar to the decline of the full basket.[22] One might be concerned that the construction of the robust partition mechanically generates a mild recession by removing sharp declining types. This is not the case: quantitatively, it is important to recognize that the initial partition cut-off used, calibrated using Monte Carlo simulations, implies extremely large declines by historical standards. Only 0.6 percent of observations during prior NBER recessions exceed this value (See Appendix Section D.1 for details).

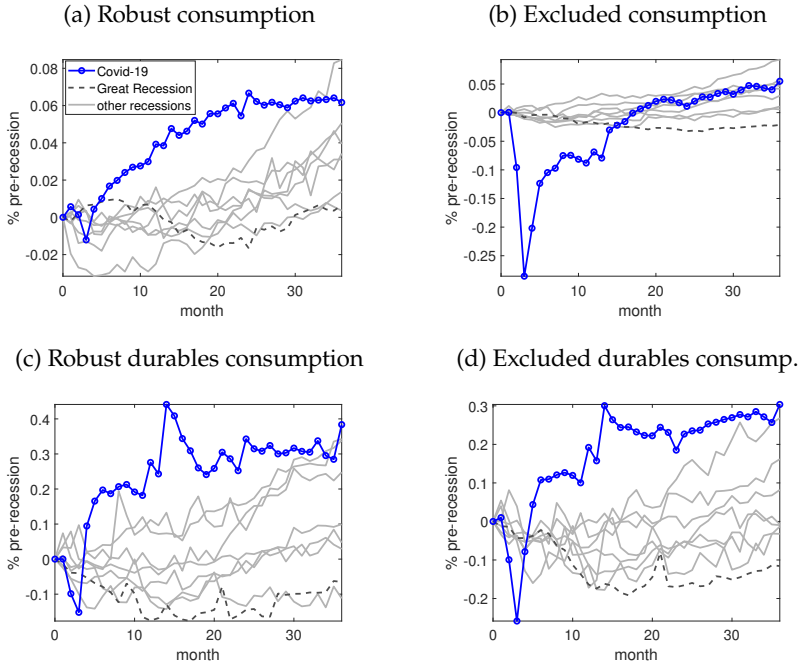
Crucially, the excluded basket's behavior in previous recessions closely tracks the robust measure, confirming the partition does *not* simply separate high- versus low-income-sensitivity goods (Appendix, Figure A.5). Rather, it isolates products uniquely disrupted by pandemic restrictions, leaving a robust measure informative about household expectations of economic conditions.

[22] During the Great Recession the maximum decline of the robust measure for non-durables and services is -1.6 percent for the full basket it is -2.4 percent. For durables the declines are essentially the same (robust: -17.7, full: -17.1).

The dynamics during the recovery reinforce this interpretation. The robust measure recovers quickly, exceeding its pre-recession level after eight months—a pace matching the fastest historical recoveries. This contrasts sharply with the Great Recession, when the robust basket remained below its initial level even 25 months into the recovery, reflecting that recession's unusually persistent consumption effects. Meanwhile, the excluded basket remains more depressed throughout the three-year window, reflecting ongoing pandemic restrictions and voluntary social distancing. This persistent weakness, while substantial, provides little information about underlying economic conditions or household expectations about income and employment prospects.

Durables consumption: Panels c and d show a similar but less extreme pattern for durables. The robust measure declines 14.2 percent three months into the recession—historically large but smaller than the excluded basket's 25.7 percent decline. This more moderate divergence indicates durables were less affected by pandemic-specific disruptions than non-durables and services, consistent with conventional wisdom about the Covid experience. Both durables measures recover strongly, exceeding their pre-recession levels by over 30 percent around 15 months into the recovery. The robust measure recovers somewhat faster, but the magnitudes are similar and far exceed any previous recession's durables recovery. This exceptional strength suggests households responded to economic conditions—perhaps including fiscal transfers and low interest rates. In Appendix Section D.4, I present a series of robustness checks and alternative specifications. The results are very consistent.

Figure IV: Robust consumption path in recessions



Notes: Recession deviations of robust and excluded consumption. Data from NIPA

2.5 Analysis of robust consumption by education in the micro data

The aggregate analysis masks potentially important heterogeneity in household responses to the Covid recession. I now examine consumption patterns by education group, which prior research showed experienced sharply different labor market outcomes during Covid (Chetty et al. (2023), Cox et al. (2020)). Understanding whether the robust consumption decline reflects concerns concentrated among specific groups provides insight into the distribution of households' recession expectations.

I use the Consumer Expenditure Survey (CEX) to construct household-level robust and excluded consumption measures. The CEX provides comprehensive consumption data for approximately 6,000 households per wave, though some components suffer from measurement error and coverage issues (Attanasio et al., 2014). My sample spans 2013-2020, yielding 195,000 household-quarter observations. To map the product-level weights ω_j from Section 2.4.1 to CEX categories, I use the BLS crosswalk to identify close analogs for most goods. When CEX aggregates multiple NIPA categories, I construct a weighted average of the NIPA weights, using each type's average 2010-2020 expenditure share as the weight.[23]

Table II reports summary statistics by education level. High-education households have substantially higher incomes: after-tax income from all sources averages 80 percent higher than for low-education households. Average quarterly robust consumption spending is \$5,300 for low-education households and \$7,900 for high-education households. Despite this difference in levels, the robust consumption *share* is remarkably similar across groups (approximately 38 percent for both), fairly closely matching the aggregate NIPA share from Table I. This consistency suggests the partition's economic logic—isolating pandemic-disrupted products—applies similarly across the education distribution.[24] Durables consumption exhibits similar patterns across education groups and partition types and the robust share is similar to that in the aggregate data. As expected quarterly spending levels are considerably lower than for non-durables and services.

Table II: Consumption partition summary statistics

Moment	Low education	High education
Share of population	0.58	0.42
Average income	\$48,463	\$88,140
Robust consumption expenditure	\$5,301	\$7,930
Robust consumption share	0.37	0.38
Robust durables expenditure	\$920	\$1,340
Robust durables share	0.49	0.44

Notes: Data from Consumer Expenditure Survey (CEX). Consumption measure composite from interview and diary survey. Income is average annual income from all source after tax. All spending on consumption goods is reported at a quarterly frequency. All values in 2019 prices. For interview only statistics see Table A.3.

[23] I divide households into low and high education groups based on college completion, consistent with the CPS analysis.
[24] Figure A.13 in the Appendix plots Engel curves for the baskets.

Figure V reveals striking heterogeneity in consumption responses by education group during the Covid recession.[25] Panel a shows sharply divergent robust consumption paths: high-education households reduced robust consumption after the recession began. Low-education households also reduce robust consumption on impact, but after five months, their consumption has rebounded. This contrast is surprising given that low-education households experienced more severe and immediate labor market disruptions (Figure A.17). Several mechanisms could explain this pattern. Low-education households may have spent fiscal transfers disproportionately on robust consumption categories, exhibited lower elasticity of substitution between robust and excluded goods, or held less pessimistic expectations about the recession's persistence relative to high-education households, who faced less immediate employment risk but greater uncertainty about duration.[26] Finally, the dashed line shows the patterns excluding housing; this is informative as the type of housing spending differs significantly by education level. The results consistent and the difference in the consumption dynamics between the two education groups is even larger. Panel b shows that both groups reduced spending on excluded consumption, but high-education households cut spending more sharply. This asymmetry likely reflects compositional differences: high-education households allocate larger budget shares to services heavily affected by pandemic restrictions—such as dining out, recreation, and travel—making their excluded consumption more sensitive to lockdown measures.

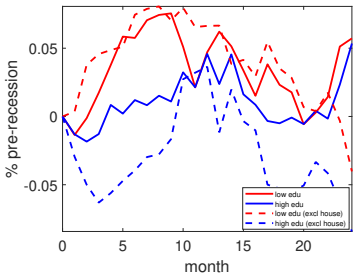
The durables partition (Panel c) reveals a third pattern: low-education households cut durables spending more aggressively at the start of the recession. Correspondingly, low education households also reduce the probability of durables adjustment by more than high education households (See Appendix Figure A.9). This is the opposite of the pattern for robust consumption. These education-specific patterns across consumption categories provide identifying variation for inferring household expectations. The stronger response of robust consumption for low education households, despite worse employment outcomes, suggests differing expectations of the economic duration of the recession. In addition, fiscal policy may have been more effective at offsetting the negative consumption decline for low education households. Conversely, the large durables decline for low education households suggests heightened uncertainty, particularly at the recession's onset. Section 4 formalizes this intuition by estimating a structural model that maps these heterogeneous consumption responses to underlying beliefs about the recession's duration and income uncertainty.

[25] I seasonally adjust the aggregate CEX data and apply 3-month moving averages to smooth high-frequency variation. Appendix Figure A.10 shows that the CEX aggregates closely track the NIPA patterns from Figure IV.

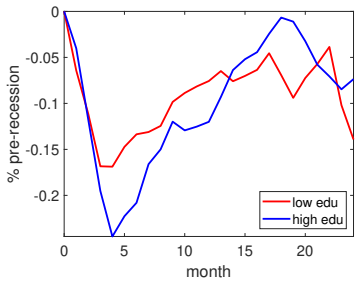
[26] Increased work from home by high education households is unlikely to be the explanation, as the major work from home-related expenditures – telecommunication costs, computers etc – are categorized in the robust basket.

Figure V: Consumption by education level

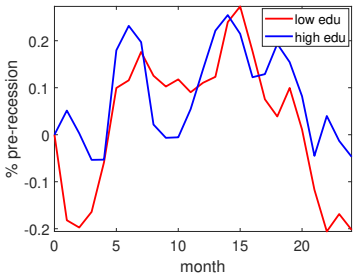
(a) Robust consumption



(b) Excluded consumption



(c) Robust durables consumption



Notes: Data from Consumer Expenditure Survey (CEX). Consumption measure composite from interview and diary survey. High education household (blue), low education households (red).

3 Life-cycle model of consumption

This section sets out the life-cycle model that I will use to study household choices during the Covid recession and to infer household expectations. The model is a partial equilibrium model with aggregate shocks. For computational expediency, I do not model the production side of the economy. Aggregate shocks and prices follow exogenous processes consistent with household expectations.

3.1 Household problem

3.1.1 Household preferences

Households are finitely lived. They enter at age $a = 1$ and each period a household survives until the next period with probability π_a. Households have a maximum lifespan of $\mathcal{A}$ periods after which they die with probability 1. Upon death, either stochastically or in the last period of life, households receive utility in the form of a warm glow bequest, $\mathcal{B}(\cdot)$. During the first $\mathcal{A}^w$ periods of life the household is working, after this they retire. Each period, households make a consumption-savings decision. They value consumption over three goods: *robust* non-durables consumption, $c^R_{a,s}$, *excluded* non-durables consumption, $c^O_{a,s}$ and a service flow from their durables stock, $d_{a,s}$. For this analysis, all durables consumption is considered to be *robust*.[27] The household felicity utility function $u_a(c^R, c^O, d'; \chi)$ is allowed to depend on age to capture life-cycle features. Utility over consumption goods is also affected by a time-varying preference shifter, χ_t, which primarily captures the utility derived from the *excluded* non-durables consumption good, $c^O_{a,s}$. Warm glow bequests are defined over end-of-period assets, $b_{a+1,s+1}$, and durable stocks, $d_{a+1,s+1}$. Given this setting, households make choices to maximize the expected

[27] Extending the analysis to partition the durables choice would be interesting and potentially informative. However, it is computationally infeasible as it would require keeping track of an additional household state variable.

value of lifetime utility at time t:

$$V_{a,t} = \mathbf{E_t} \sum_{s=0}^{\mathcal{A}-a} \beta^s \Bigg(\pi_{a+s} u_a(c^R_{a+s,t+s}, c^O_{a+s,t+s}, d_{a+s,t+s+1}; \boldsymbol{\chi}_{t+s}) + (1-\pi_{a+s}) \mathcal{B}(b_{a+s,t+s+1}, d_{a+s,t+s+1}) \Bigg)$$

3.1.2 Household durables choice

Households durables stock follows a standard law of motion $d_{a,s+1} = (1-\delta)d_{a-1,s} + i^d_{a,s}$ where $i^d_{a,s}$ if the period durables consumption, or investment in the durables stock, and δ captures per period durables depreciation. The durables stock is subject to adjustment costs, which are stochastic. At the beginning of the period, households draw an i.i.d. cost of adjustment, $\psi_s \in \mathcal{P}$. This captures variation in marginal value of time required as an input to the durables decision. Households wishing to adjust their durables stock holdings must pay a cost $\psi_s d_{a-1,s}$ whereas households that allow their durables stock to depreciate, beyond maintenance, ϑ, incur no adjustment costs. The specification of the durables adjustment cost function is:

$$\Psi(d_{a-1,s}, d_{a,s+1}, \psi_s) = \begin{cases} 0 & if\ d_{a,s+1} = (1-\delta(1-\vartheta))d_{a-1,s} \\ \psi_s d_{a-1,s} & \text{otherwise} \end{cases} \tag{5}$$

Such adjustment non-convex costs deliver lumpy durables investment and (S,s) policy functions in the current period durables stock. The stochastic adjustment cost generates greater persistence of durable consumption in response to aggregate shocks. It can be optimal for households to delay their response to a shock if they currently face a high cost.

3.2 Asset choice

Households are able to save for precautionary and life-cycle motives in a risk free bond, $b_{a,s+1}$, which pays the real return r. Borrowing is allowed for the purchase of durables. Constrained households may take up a loan against their durables purchase, with downpayment $(1-\zeta)d_{a,s+1}$, such that they face the collateral constraint: $b_{a,s+1} \geq -\zeta d_{a,s+1}$. Borrowing can also be undertaken at the risk free rate. At the beginning of the lifecycle, households receive an endowment of assets $a_0 \sim \log \mathcal{N}(\mu_{a_0}, \sigma^2_{a_0})$ and durables $d_0 \sim \log \mathcal{N}(\mu_{d_0}, \sigma^2_{d_0})$. Households receive a warm glow bequest $\mathcal{B}(b_{a,s+1}, d_{a,s+1})$ for assets held at the end of life or following death.

3.3 Income process

Household income follows a rich specification capturing persistent, permanent, and transitory components as well as aggregate and idiosyncratic shocks. Denote the income of a working age household with education level e as $y^e_{a,s}(Z_{e,s}, z^e_{a,s}, \iota_{a,s})$. Where $Z_{e,s}$ is

aggregate productivity, $z_{a,s}$ is the persistent component of income, and $\iota_{a,s}$ is an indicator function denoting employment status. When $\iota_{a,s} = 1$ a household is employed; when $\iota_{a,s} = 0$ a household is unemployed. Income has the following functional form:

$$y^e_{a,s}(Z_{e,s}, z_{a,s}, \iota_{a,s}) = \begin{cases} Z_{e,s} \exp(\alpha^e_a + z^e_{a,s}) & if\ \iota_{a,s} = 1 \\ \mu Z_{e,s} \exp(\alpha^e_a + z^e_{a,s}) & if\ \iota_{a,s} = 0 \end{cases} \tag{6}$$

Household income follows a deterministic path $\{\alpha^e_a\}^{\mathcal{A}^r}_{a=1}$ which differs by education level. Unemployed households receive the replacement rate μ of their income while in employment. This captures unemployment insurance.

Persistent income shocks: The persistent component of income follows an AR(1) process: $z^e_{a,s} = \rho_{z^e} z^e_{a-1,s-1} + \varepsilon^e_{a,s}$ with $\varepsilon^e_{a,s} \sim \mathcal{N}(0, \sigma^2_{z^e,s})$. I allow the persistence ρ_{z^e} and variance of $\sigma^2_{z^e,s}$ of the process to vary by education level. In the main analysis, I also allow for the variance of shocks $\sigma^2_{z^e,s}$, to vary deterministically over time, capturing *increased uncertainty* facing households during a recession.

Transitory income shocks: Households experience periods of unemployment where their income falls to μ percent of their in-work income. The probability of transition into and out of unemployment follows a Markov chain $\mathcal{P}(\iota_{a,s}, \iota_{a+1,s+1}; a, e, Z_s)$, where the notation indicates dependence on age, education, and the aggregate state. The transitory shocks can be summarized by transition probabilities into λ^0_{a,e,Z_s} and out of λ^1_{a,e,Z_s} the unemployment state, such that the Markov matrix can be written as:

$$\mathcal{P}(\iota_{a,s}, \iota_{a+1,s+1}; a, e, Z_s) = \begin{bmatrix} 1 - \lambda^1_{a,e,Z_s} & \lambda^1_{a,e,Z_s} \\ \lambda^0_{a,e,Z_s} & 1 - \lambda^0_{a,e,Z_s} \end{bmatrix} \tag{7}$$

Aggregate shocks: Households also experience aggregate income shocks $Z_{e,s}$ and can be interpreted as TFP shocks.[28] Aggregate shocks follow a Markov matrix $\mathcal{P}(Z_s, Z_{s+1})$.

Retirement: After $\mathcal{A}^w$ periods, households enter retirement. In the retirement state, households receive a pension which is an education-specific share, κ^e of their final salary.[29] Pension income is then constant for the rest of the household's life.

3.4 Budget constraint

I can now write out the household budget constraint. Households choose a consumption basket that includes robust non-durables, $c^R_{a,s}$, excluded non-durables, $c^O_{a,s}$, and durables consumption, $i^d_{a,s}$. The latter two goods have the relative prices p_s and p^d_s, respectively. They also make a savings decision, $b_{a,s+1}$. The resources each period are income $y^e_{a,s}(Z_{e,s}, z_{a,s})$, if of working age, or pension; the bonds saved from last period $b_{a-1,s}$ after interest; and a

[28] Aggregate productivity is shown here depending on education, as in the Covid application, we allow for education-specific TFP shocks. In the stochastic simulation to calibrate the initial conditions, there is a single aggregate TFP process, the aggregate Markov processes are not education-specific.

[29] For technical reasons, productivity in the first period of retirement is also stochastic following the same persistent productivity process for $z^e_{a+1,s+1}$. Retirees receive a share κ^e of this income level.

government transfer, $\mathcal{T}_{a,e,s}$. Should the household adjust its durables stock they must also pay the adjustment cost $\Psi(d_{a-1,s}, d_{a,s+1}, \psi_{a,s})$.

$$c^R_{a,s} + p_s c^O_{a,s} + b_{a,s+1} + p^d_s i^d_{a,s} = \tag{8}$$
$$y^e_{a,s}(Z_{e,s}, z_{a,s}, \iota_{a,s}) + (1+r)b_{a-1,s} - \Psi(d_{a-1,s}, d_{a,s+1}, \psi_{a,s}) + \mathcal{T}_{a,e,s}$$

3.5 Recursive formulation

The household choice problem can be presented in a recursive formulation. For this purpose, I suppress the dependency on education level, e, and age, a, and time s subscripts, except where it provides clarification, and use prime notation for next period's values. There are value functions for households that adjust their durables stock, $V^{adj}(\cdot)$, and for those that do not, $V^{nadj}(\cdot)$. The idiosyncratic state variables for the household are bond holdings, b, beginning-of-period durables stock, d, the persistent income state, z, employment status, ι, and current cost of adjustment, ψ. The aggregate state variables are productivity, Z, a vector of relative prices, $\mathbf{p} = \{p, p^d\}$, and a vector of preference shifters, $\boldsymbol{\chi}$, defined below.

The value function for the working age household that adjusts is, $V^{adj}(\cdot)$, is:

$$V^{adj}_a(b, d, z, \iota, \psi; Z, \mathbf{p}, \boldsymbol{\chi}) = \max_{c^R, c^O, d', b'} u(c^R, c^O, d'; \boldsymbol{\chi}) + (1-\pi_a)\mathcal{B}(b', d') +$$
$$\beta\pi_a \mathbf{E} V_{a+1}(b', d', z', \iota', \psi'; Z', \mathbf{p}', \boldsymbol{\chi}')$$

subject to

$$c^R + pc^O + b' + p^d d' = y(Z, z, \iota) + (1+r)b + p^d(1-\delta)d - \psi d + \mathcal{T}$$
$$b' \geq -\zeta d'$$

The value function for the working age household that does not is, $V^{nadj}(\cdot)$, is:

$$V^{nadj}_a(b, d, z, \iota; Z, \mathbf{p}, \boldsymbol{\chi}) = \max_{c^R, c^O, b'} u(c^R, c^O, (1-\delta(1-\vartheta))d; \boldsymbol{\chi}) + (1-\pi_a)\mathcal{B}(b', (1-\delta(1-\vartheta))d) +$$
$$\beta\pi_a \mathbf{E} V_{a+1}(b', (1-\delta)d, z', \iota', \psi'; Z', \mathbf{p}', \boldsymbol{\chi}')$$

subject to

$$c^R + pc^O + b' = y(Z, z, \iota) + (1+r)b - \delta p^d \vartheta d + \mathcal{T}$$
$$b' \geq -\zeta(1-\delta(1-\vartheta))d$$

Where the next period value, $V(\cdot)$ is the envelope over the decision to adjust or not.[30]

$$V_a(b,d,z,\iota,\psi;Z,\mathbf{p},\boldsymbol{\chi}) = \max\left\{V_a^{adj}(b,d,z,\iota,\psi;Z,\mathbf{p},\boldsymbol{\chi}), V_a^{nadj}(b,d,z,\iota;Z,\mathbf{p},\boldsymbol{\chi})\right\}$$

Retired households' value functions are similar, with the exception that the value of the household's pension income, $\kappa^e y^e(z_{a=\mathcal{A}^{w+1}})$ replaces working age income. Retired households' only exposure to stochastic variables is via aggregate prices, $\mathbf{p}$, and the preference shock, $\boldsymbol{\chi}$.

3.6 Calibration and numerical implementation

3.6.1 Utility function form

Consumption preferences: Households in the model derive consumption over three goods: robust consumption, $c^R_{a,t}$, excluded consumption, $c^O_{a,t}$, and durable goods, $d_{a,t+1}$. I adopt a nested CES structure to allow for rich and flexible interaction in the demand for goods. Define the felicity utility function:

$$u_a(c^R_{a,t}, c^O_{a,t}, d_{a,t+1};\boldsymbol{\chi}_t) = u\left(C(c^R_{a,t}, c^O_{a,t};\chi_t), d_{a,t+1};\boldsymbol{\chi}_t\right) \tag{9}$$

With preferences over consumption and durables consumption defined by

$$u(C(c^R_{a,t}, c^O_{a,t};\chi_t), d_{a,t+1};\boldsymbol{\chi}_t) =$$
$$\mathcal{X}_t \frac{\left[(\theta+\hat{\chi}_t)^{\frac{1}{\hat{\epsilon}}}(C(c^R_{a,t}, c^O_{a,t};\chi_t))^{\frac{\hat{\epsilon}-1}{\hat{\epsilon}}} + (1-\theta-\hat{\chi}_t)^{\frac{1}{\hat{\epsilon}}}(d_{a,t+1})^{\frac{\hat{\epsilon}-1}{\hat{\epsilon}}}\right]^{\frac{(1-\rho)\hat{\epsilon}}{\hat{\epsilon}-1}} - 1}{1-\rho} \tag{10}$$

And the sub-utility function over consumption given by:

$$C(c^R_{a,t}, c^O_{a,t};\chi_t) = \left[(\theta_c+\chi_t)^{\frac{1}{\epsilon}}(c^R_{a,t})^{\frac{\epsilon-1}{\epsilon}} + (1-\theta_c-\chi_t)^{\frac{1}{\epsilon}}(c^O_{a,t})^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}}$$

The parameters that govern the utility function are the intertemporal elasticity of substitution, $\rho^{-1} > 0$, the share of consumption, $\theta \in (0,1)$, the elasticity of substitution between consumption and durables, $\hat{\epsilon}$, the share of consumption allocated to robust consumption $\theta_c \in (0,1)$, the elasticity of substitution between robust and excluded consumption ϵ.

The vector of preference shifters $\boldsymbol{\chi}_t = \{\chi_t, \hat{\chi}_t, \mathcal{X}_t\}$ captures three margins by which consumption responds to the non-economic forces during the pandemic. Firstly and most importantly, the preference shifter of the sub-utility function, χ_t, changes the robust consumption share in non-durables and services consumption. A fall in the value of χ_t captures the decline in value of certain goods during the Covid recession. It raises demand for the robust good and lowers demand for the excluded consumption good. In addition, I

[30]The cost of adjustment is irrelevant for the non-adjustment value function and does not appear as a state.

allow for the preference shifter to reallocate demand between non-durables and durable consumption, $\hat{\chi}_t$, and to change total consumption demand with a utility shifter, $\mathcal{X}_t$. The preference shifters only play a role during the Covid recession. Discussion of its functional form is deferred until Section 4.1.1.

Bequest function: Households also value warm glow bequests, following De Nardi (2004). I use the non-homothetic warm glow bequest function:

$$\mathcal{B}_e(b_{a,t+1}, d_{a,t+1}) = \frac{\mathcal{B}_0^e\big((1 + (b_{a,t+1} + (1-\delta)d_{a,t+1}))/\mathcal{B}_1^e\big)^{1-\rho_{\mathcal{B}}}) - 1}{1 - \rho_{\mathcal{B}}} \tag{11}$$

The parameters that govern the utility function are the value of bequests $\mathcal{B}_0^e$, the importance of luxury bequests $\mathcal{B}_1^e$, and the curvature of the bequest function $\rho_{\mathcal{B}}$.[31]

3.6.2 Externally calibrated parameters

The model is rich with a large number of parameters. I choose some parameters from external sources and internally calibrate the remaining parameters, using the Simulated Method of Moments (SMM). The time period of the model is a quarter. I set the inverse of elasticity of intertemporal substitution, ρ, to 1.5. This is a common value in the literature and consistent with evidence from Attanasio and Weber (1995). I assume consumption and durables are Cobb-Douglas. This is convenient as it simplifies the estimation of preference shocks later and is widely used in the literature (Berger and Vavra (2015), Eichenbaum et al. (2022)).[32]. The steady state preference shock, χ, is set to 0. The durables depreciation rate, δ is set to 0.047, in line with BEA estimates and maintenance, ϑ, to 0.25. Mirroring the empirical analysis in the micro data, I use two education groups. The low type is defined as no- or some-college and the high type as completed college.[33] I estimate the persistent income process separately for high and low education heads of household in the PSID. The values are similar with high education households facing slightly higher variance shocks. I then apply a five-state discrete approximation using Rouwenhorst methods.[34] The unemployment replacement rate is set to μ following Shimer (2000). I set the pensions replacement rate, κ, for high education households, to 0.6 based on evidence in Bernheim et al. (2001). I set the rate higher for the low education household, $\kappa = 0.65$.

The interest rate is set to a standard value of 4 percent on an annual basis. I require a downpayment on durables of 20 percent, setting the collateral constraint parameter $\zeta = 0.8$. The curvature of the bequest function is set equal to the curvature of the utility function, $\rho_{\mathcal{B}} = \rho$. For the aggregate shock, I estimate processes in the data. For productivity, I assume a three-state process, capturing booms, recessions, and deep recessions. I use the

[31] Bequests over durables are valued after depreciation in keeping with the timing assumption that households first derive utility from consumption of the durables stock today.
[32] It also ensures the Euler equation can be inverted when solving household policy functions with the Endogenous Grid Method, speeding up computation (See Appendix Section E.2.1)
[33] From the CPS this implies 65 percent of households are of low education type.
[34] The use of Rouwenhort discretization methods is particularly important for delivering an accurate approximation when subjecting the economy to an uncertainty shock.

Table III: Externally calibrated parameters

Parameter		Value
Elasticity of intertemporal substitution^{-1}	ρ	1.5
Elasticity of substitution (durables)	$\hat{\epsilon}$	1.0
Steady state preference shock	χ	0.0
Durables depreciation rate (annualized)	δ	0.17
Durables maintenance	ϑ	0.25
Share of education type	e	$\{0.65, 0.35\}$
Persistence of income	ρ_{z^e}	$\{0.978, 0.979\}$
Variance of income shocks	$\sigma^2_{z^e}$	$\{6.8e^{-4}, 9.3e^{-4}\}$
Unemployment replacement rate	μ	0.4
Pensions replacement rate	κ	$\{0.65, 0.6\}$
Interest rate (annualized)	r	0.04
Collateral constraint	ζ	0.8
Productivity shocks	Z	$\{0.96, 0.98, 1.01\}$
Persistence of consumption price	ρ^p	0.83
Variance of consumption price shocks	σ^2_p	0.005

TFP estimates of Fernald (2021), defining a deep recession as a 3σ deviation and recession and booms as $-\sigma$ and $+\sigma$ events, respectively.[35] I then fit a Markov chain based on the nearest neighbour fit of each data point to this grid at each period in time. The resulting process mainly fluctuates between recessions and booms, spending around 5 percent of the time in deep recessions.[36] For the aggregate price of the excluded consumption good, I estimate an AR(1) process from the implied data series. This gives a persistence, ρ^p, of 0.87 and a variance of 0.005. I assume a constant durable price $p^d = 1$. All externally calibrated parameter values are given in Table V.

Table IV: Targeted moments

Moment		Parameter	Data	Model
Robust consumption ratio	C^R/C^O	θ_c	0.65	0.66
Durable consumption share	$I^d/(C^R + C^O + I^d)$	θ	0.05	0.05
Correlation between consumption goods	$\rho(C^R, C^O)$	ϵ	0.67	0.47
Std dev of durable consumption	$\sigma(I^d)/\sigma(Y)$	Ψ	2.91	2.68
Percent of households adjusting	$\Pr(i^d > 0)$	Ψ	0.21	0.11
Kurtosis of durable consump.	$(I^d)^4$	$\mathcal{P}$	3.8	4.3
Median wealth:inc of working age	$\sum_{a\in W} B_a^{med} : Y_a$	β^e		
Median wealth:inc of retired	$\sum_{a\in R} B_a^{med} : Y_a$	$\mathcal{B}_0^e$	Figure A.15	
Mean wealth:inc of retired	$\sum_{a\in R} \overline{B}_a : Y_a$	$\mathcal{B}_1^e$		

I also set several external parameters that vary over the life-cycle. In particular, I take pre-pandemic survival probabilities, π_a, from the National Center for Health Statistics (NCHS) and estimate the exogenous life-cycle income profile, α_a^e, of low and high education types in the PSID. I compute age×education×aggregate productivity level specific labor market transitions for the job finding, $\lambda^0_{a,e,Z}$, and job separation rate, $\lambda^1_{a,e,Z}$, in the CPS.[37]

[35] I use the measure of TFP unadjusted for utilization, which has more reasonable dynamics in the Great Recession.
[36] Both the Great Recession and Covid recession would have been categorized as deep recessions.
[37] These transition rates are estimated on pre- Covid recession data. The aggregate productivity level is defined as a boom or recession.

These are shown in the Appendix Figure A.14. Separation rates and finding rates fall over the life-cycle, with the exception of the separation rate of the high education type. Separation rates are much higher for low education types, whereas the difference between the job finding rates is less pronounced.[38]

3.6.3 Model estimation: distribution subject to aggregate shocks

I estimate several parameters within the model using the SMM. To do this, I simulate a non-stochastic distribution of households, subjected to a sequence of aggregate shocks. The moments related to this calibration exercise are shown in Table IV. The robust share of consumption and durables consumption share from Section 2.4.1 identify the preference parameters θ and θ_c. To calibrate the elasticity of substitution between the robust consumption good and excluded consumption good, ϵ, I target their business cycle correlation.[39] The model features a shock to the relative price of the excluded good, which (by driving a wedge between these two goods) provides the identification for this parameter.

I also need to target the adjustment cost parameters. I choose a two-state specification for the adjustment cost realizations. This leaves three parameters. The fixed cost in each state, ψ, and the probability of drawing the low state, $\mathcal{P}$. For the costs, I target both the probability of adjusting measured in the CEX and the standard deviation of the robust durables consumption good in NIPA.[40] In the data, 27 percent of households adjust their robust durables consumption per quarter. To target the probability of drawing the low adjustment cost, I target the kurtosis of durables consumption in NIPA. By capturing the size of larger adjustments, this determines the value of waiting to draw the low cost.

Table V: Internally calibrated parameters

Parameter		Value
Consumption share	θ	0.94
Robust consumption share	θ_c	0.40
Robust consump. elasticity	ϵ	1.03
Adjustment costs	Ψ	$\{0.02, 0.07\}$
Low adj. cost share	$\mathcal{P}$	0.39
Discount factor	β^e	$\{0.98, 0.99\}$
Bequest motive level	$\mathcal{B}_0^e$	$\{1.57, 2.00\}$
Bequest motive luxury	$\mathcal{B}_1^e$	$\{1.98, 1.23\}$

Finally, I calibrate the household discount factors and bequest parameters. I target the life-cycle median wealth to mean income ratio, measured in the 2019 Survey of Consumer Finance. I use the full life-cycle path by education as seen in the Appendix Figure A.15. In practice, the working-age portion informs the discount factor, β^e, while the retirement-age

[38] The effect of recessions is very non-linear, with the change between booms and recessions substantially smaller than the change between recessions and deep recessions.
[39] This elasticity is important to calibrate as the price of excluded goods varies during the recession.
[40] Intuitively, the cost of adjustment affects the frequency of adjustment and the volatility of the associated durables consumption series.

portion informs the bequest motive, $(\mathcal{B}_0^e, \mathcal{B}_1^e)$. The model generates a very good fit for most moment targets. The resulting parameters are shown in Table V.

4 Extracting expectations

4.1 Recession shocks

The main quantitative exercise combines the robust consumption measures from Section 2.3–robust non-durables and services ($\mathbf{C_t^R}$), excluded consumption ($\mathbf{C_t^O}$), and robust durables ($\mathbf{I_t^R}$)–with the life-cycle model from Section 3 to extract information about household expectations during the Covid recession. While changes in excluded consumption do not respond in an interpretable way to economic shocks, they could still influence the inferred expectations if lower spending relaxes credit constraints facing households with low assets.[41] The estimation exploits revealed preferences: observed consumption choices, interpreted through the model's optimality conditions, to identify three key features of household beliefs:

1. the level of economic uncertainty households perceived;
2. how persistent they expected the recession's economic effects to be;
3. how these expectations evolved as the recession unfolded.

I take as given the observed paths of productivity, labor market transition rates, prices, the preference shifter for excluded consumption, and government transfers. Figure VI displays these series, estimated from the data. Productivity shocks, labor market transitions, and transfers are education-specific, denoted by subscript e. The model then identifies the paths of education-specific income uncertainty, $\sigma^2_{z^e,e,t}$, and shock persistence, $\varrho_{e,t}$, required to match four empirical targets: aggregate robust consumption, aggregate robust durables consumption, and the consumption differences between education groups.[42] An implicit assumption is that to extract beliefs correctly, it is necessary that the empirical exercise correctly identifies the true robust and excluded basket partition.

The quantitative exercise begins in 2020:Q1, following the NBER recession dating. I compute impulse responses for 30 quarters following the recession's onset.[43] Appendix E.3.2 provides precise mathematical definitions of all shocks and their timing under each information structure.

[41] The model only features one category of durables but this is of less importance for credit constraints due to lower spending. I target the behavior of the robust durables empirical series.
[42] I subtract trend growth from consumption series using the average growth rate from the two years preceding the recession to maintain consistency with the model's stationarity assumption.
[43] For shocks that have not returned to trend by the final observation, I assume they revert with persistence 0.85. Households also experience Covid mortality shocks, taken from the data.

Productivity shocks: Productivity shocks, $\{Z_t^e\}_{t=1}^T$, are used to match the change in wages by education type. These are computed in the CPS. I constrain the resulting aggregate series to follow negative income deviations in NIPA.[44]

Labor market shocks: The job finding, $\{\lambda_t^{0,e}\}_{t=1}^T$, and separation rate, $\{\lambda_t^{1,e}\}_{t=1}^T$, shocks are computed from the CPS. Shocks differ by education level and are computed as deviations from the rate in 2019.

Price shocks: Price shocks, $\{p_t\}_{t=1}^T$, are fed in as deviations in the relative price of the excluded consumption good (see Figure A.8). Note that this means price changes due to economic supply and demand channels are allowed to impact the consumption choices.[45]

Preference shocks: Preference shocks, $\{\chi_t\}_{t=1}^T$, change the consumption weights in the utility function. Given this specification, χ_t is given by the intratemporal consumption choice, which can be rearranged as:

$$\chi_t = \frac{1 - \theta_c\left(1 + p_t^{\epsilon}\frac{\mathbf{C_t^O}}{\mathbf{C_t^R}}\right)}{1 + p_t^{\epsilon}\frac{\mathbf{C_t^O}}{\mathbf{C_t^R}}} \tag{12}$$

I solve for the value of χ_t each period that delivers the deviations in robust consumption relative to excluded consumption, given the deviation in the relative price.

Real interest rate: In Section 4.4, I extend the analysis to include a fall in the real interest rate $\{r_t\}_{t=1}^T$ associated with rising inflation. This is computed as the difference between the Fed Funds rate and the inflation rate. As the increase in inflation was not anticipated by households at the beginning of the Covid recession, I restrict the implementation of this shock to the case of dynamic expectations.[46]

4.1.1 Calibrating the remaining preference shock parameters

The robust consumption preference shock, χ_t, can be directly computed from the data with series for $\mathbf{C_t^O}$ and $\mathbf{C_t^R}$ (see Equation 21). In addition, the preference specification incorporates the effect of the preference shock on two other margins, the consumption-durables choice, $\hat{\chi}$, and a demand/utility shifter, $\mathcal{X}$.[47] I flexibly parameterize these two other shifters as a function of the robust consumption preference shock, χ_t, with the following functional forms: $\hat{\chi}(\chi_t) = -\xi_1\chi_t$ and $\mathcal{X}(\chi_t) = \exp(-\xi_2|\chi_t|)$. To calibrate the parameters $\{\xi_1, \xi_2\}$, I use cross-sectional variation in the CEX at the household level. I regress robust- consumption and durables consumption growth for household i in 2020, on the change in the ratio of

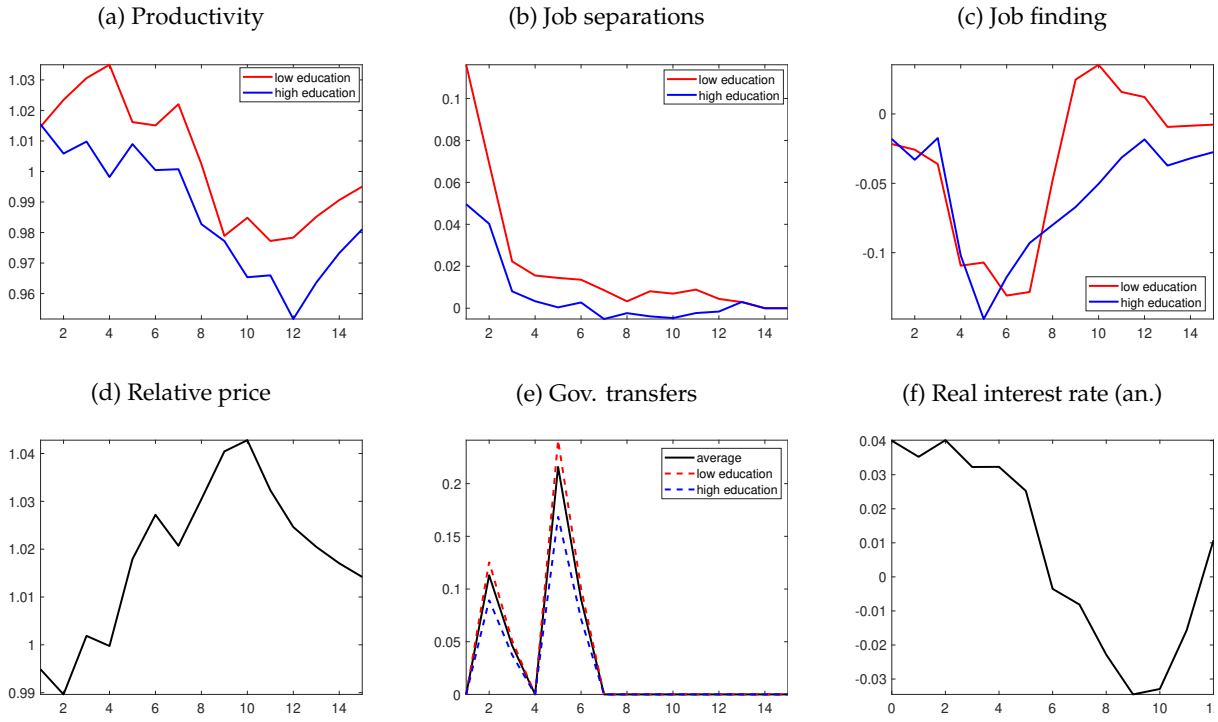
[44] I adjust productivity to offset the aggregate effect of a change in $\sigma^2_{z^e,e,t}$ and ensure any rise in uncertainty is mean preserving.

[45] I experimented with allowing the durables price p^d to vary following the data series. However, under the static information setting, due to full information, it generated counterfactual paths of durables expenditures. It seem unlikely that households anticipated the increase in durables prices that occurred in 2021 in 2020.

[46] Inflation also reduces real wages and the real value of retirees' pensions. These are included through the productivity shocks and occur in both static and dynamic cases.

[47] The model ratio $C_t^O : C_t^R$ is independent of $\hat{\chi}$ and $\mathcal{X}$.

Figure VI: Covid recession shocks



Notes: Recession shocks fed into the model. Productivity from income in CEX. Job separations and finding rate from CPS. Relative price from NIPA. Government transfers constructed. Real interest rate is fed funds rate minus inflation.

excluded consumption to robust consumption.[48] I also include a set of household controls, including income, and state-level fixed effects, $\mathbf{X_i}$. The specification for robust consumption is as follows.

$$\Delta \log C_i^R = \beta_0 + \beta_1 \Delta \frac{C_i^O}{C_i^R} + \gamma \mathbf{X_i} + \epsilon_i \tag{13}$$

The estimation reveals an increase in the C^O/C^R ratio reduces robust consumption demand and has no statistically significant effect on durable consumption. To map the semi-elasticities to the model, I target the annual response of an impulse response function targeting a semi-elasticity of -10 percent for robust consumption and no on-impact effect for durables.[49]

4.2 Source of identification

To understand how the model can identify expectations during a recession, it is informative to show how different shocks generate distinct consumption responses. Section 2.1 showed this intuition analytically using a simplified two-period model. I now demonstrate identification in the full quantitative life-cycle model through impulse response functions. Figure VII plots consumption responses to four shocks: a preference shifter for excluded

[48] To address concerns about measurement error, I winsorize both the change in the consumption ratio and the change in the dependent variable at the 5% and 95% level.

[49] For the IRFs, I use a decline in the preference shock by 5 percent, matching the change calculated from equation 21. I target the data estimates that exclude housing, as otherwise, this may capture additional channels, such as housing wealth effects.

Table VI: Calibration targets for preference parameter

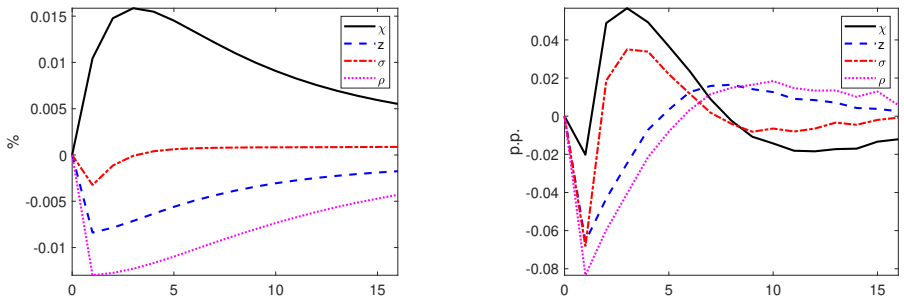
	Robust consumption			Durable consumption		
	(1)	(2)	(3)	(4)	(5)	(6)
Coefficient on $(C^O : C^R)$	-0.09	-0.10		0.16	0.01	
	(0.01)	(0.00)		(0.16)	(0.06)	
Excluding housing		✓			✓	
R^2	0.15	0.14		0.02	0.02	
Observations	3,327	3,376		2,330	2,352	
Model			-0.10			0.01

Notes: Cross-section regression run in Consumer Expenditure Survey in 2020 on change in robust or durables consumption on change in ratio of excluded to robust consumption. Model is an annually aggregated model response to preference shock of same magnitude as data.

consumption, a productivity decline, an increase in income uncertainty, and an increase in shock persistence.[50] Panel a shows the robust consumption responses; Panel b shows the durables:robust-consumption ratio. The shocks generate distinct patterns in both magnitude and dynamics, indicating how revealed preferences can be used to uncover multi-dimensional household expectations. Further, combining the two consumption margins strengthens identification.

Figure VII: Impulse response functions to shocks

(a) Robust consumption (b) Durable:robust consump. ratio



Notes: Impulse response function to one time shock in model. Preference shock (black); productivity (blue dash); uncertainty (red dash dot); persistence (magenta dot). Persistence is an increase in persistence of productivity shock.

Preference shock: A negative preference shock reduces the desirability of excluded consumption relative to robust consumption, shifting demand towards the robust basket. Robust consumption increases on impact, while durables consumption remains unchanged (Table VI).[51] This shock thus generates a decline in the durables-to-robust ratio.

Productivity versus uncertainty: Both shocks reduce robust consumption, but uncertainty generates a smaller, less persistent decline (Panel a). The durables ratio exhibits similar initial responses but sharply different dynamics: the uncertainty effect dissipates quickly and overshoots, while the productivity effect persists (Panel b). Intuitively, heightened

[50] Shock sizes are: preference = −5%, productivity = −5%, uncertainty = $\sigma_{z^e,e}$, persistence = $\Delta\varrho_e = 0.2$. All shocks have baseline persistence 0.57, matching the estimated preference shock persistence.

[51] Over the medium run, durables consumption rises because the preference shifter reduces current consumption's value relative to future consumption, and the durables stock delivers greater future consumption services.

uncertainty induces precautionary saving and delays durable purchases temporarily, but households resume spending once uncertainty resolves. In contrast, a productivity decline reduces permanent income, generating persistent consumption reductions across both margins.

Persistence shock: A positive persistence shock represents a productivity decline with longer expected duration. As anticipated, this amplifies the robust consumption decline relative to a transitory productivity shock. Crucially, Panel b shows that greater persistence also increases the durables-to-robust ratio response—households cut durables more aggressively when expecting a prolonged downturn. This differential sensitivity of durables to expected duration provides key identifying variation. These distinct impulse responses demonstrate how the combination of robust consumption and durables consumption helps identify household expectations.[52] Uncertainty primarily affects the short-run dynamics and degree of overshooting, while persistence governs the magnitude and duration of consumption responses. Education-specific consumption patterns provide additional identifying variation.

4.3 Static expectations

I consider two information structures: *static expectations* and *dynamic expectations*. Dynamic expectations will be the benchmark, but the static case provides a useful comparator. In particular, one can compare whether consumption patterns reflect immediate recognition of the pandemic's economic implications or gradual learning about the recession's severity and duration. Under static expectations, households observe the complete future path of all shocks — productivity, labor market transitions, preferences, and transfers in 2020:Q1. The only unknowns are the level and persistence of uncertainty households perceived at the recession's onset. There are estimated to fit the data. I estimate the parameters (Θ) by minimizing the distance between model predictions and data for four consumption series: aggregate robust consumption, $\hat{\mathbf{C}}_{\mathbf{t}}^{\mathbf{R}}$, the aggregate robust durables stock $\hat{\mathbf{D}}_{\mathbf{t}}$, and the education differences in these quantities $\Delta_e \hat{\mathbf{C}}_{\mathbf{t}}^{\mathbf{R}} = \hat{\mathbf{C}}_{\mathbf{t}}^{\mathbf{R,e=L}} - \hat{\mathbf{C}}_{\mathbf{t}}^{\mathbf{R,e=H}}$ and $\Delta_e \hat{\mathbf{D}}_{\mathbf{t}} = \hat{\mathbf{D}}_{\mathbf{t}}^{\mathbf{e=L}} - \hat{\mathbf{D}}_{\mathbf{t}}^{\mathbf{e=H}}$, from the micro data.[53]

The static model provides a fairly good fit of robust consumption, slightly overestimating the initial decline, but a poor fit for durables (Figure VIII, red line). Full details of the estimation and fit of the static expectations case is presented in the Appendix Section E.4

4.4 A dynamic model of expectations

I now consider a dynamic model of household expectations to see how well this information structure explains consumption patterns during the recession and what expectations it implies. For the dynamic model, I estimate a *sequence* of household expectations.

[52] In the quantitative exercise, the impact magnitudes of most shocks (except uncertainty) are calibrated directly from observed data, leaving uncertainty and persistence as the primary parameters identified from consumption dynamics.
[53] Hat notation $\hat{\mathbf{X}}$ denotes deviations from steady state (model) or trend (data).

4.4.1 Estimation exercise

Households begin in the high-productivity steady state. The model is solved forward sequentially, with households learning about the recession quarter by quarter. Each period t, the current shock realizations—productivity, labor market transitions, prices, preferences, and transfers—become observable. Rather than knowing future shock paths, households form expectations assuming all aggregate shocks follow a common mean-reverting process with persistence ϱ_o^e, where subscript o denotes the period when the shock first hits and superscript e denotes education group.[54] For example, when the productivity shock Z_1^e hits in 2020:Q1, households expect future productivity to evolve as $Z_t^e = (1 - \varrho_1^e)\bar{Z} + \varrho_1^e Z_{t-1}^e$. Each period there is a new realization of additional uncertainty, $\hat{\sigma}_{z,o}^e$, which households also expect to persist with the coefficient ϱ_o^e.

To avoid under-identification, I parametrize education heterogeneity through constant gaps rather than estimating separate parameters each period:

$$\varrho_o^{e=H} = \varrho_o^{e=L} + \alpha_\varrho \tag{14}$$

$$(\hat{\sigma}_{z,o}^{e=H})^2 = \alpha_\sigma \times (\hat{\sigma}_{z,o}^{e=L})^2 \tag{15}$$

This specification assumes high-education households consistently perceive (i) a persistence gap of α_ϱ relative to low-education households and (ii) uncertainty scaled by factor α_σ. These gaps capture systematic differences in how education groups experience and perceive economic shocks. I employ a nested estimation strategy with inner and outer loops, where the outer loop solves for the education differences.

Inner loop: First, conditional on the education gaps, I estimates period-by-period belief parameters $\left\{\Theta_o = \{\hat{\sigma}_{z,o}^e, \varrho_o^e\}_{e=1}^2\right\}_{o=1}^T$, with high-education parameters determined by gaps α_ϱ and α_σ. For each period o, I minimize squared deviations between model and data for current-period aggregate consumption, given the model-implied future path of exogenous variables:

$$F_o^{inner}(\Theta_o) = M_o^{inner}(\Theta_o) M_o^{inner}(\Theta_o)'$$

$$M_o^{inner}(\Theta_o) = [\hat{C}_{o\ data}^R - \hat{C}_o^R(\Theta_o)_{model}, \hat{D}_{o\ data}^R - \hat{D}_o(\Theta_o)_{model}]$$

This identifies households' beliefs in the period o given observed shocks.[55] The model is then solved forward one period using these policy functions to generate the next period's initial distribution over states.[56]

[54] The shocks following the common persistence parameter are: productivity (Z_t^e), labor market transition rates ($\lambda_{0,t}^e, \lambda_{1,t}^e$), real interest rate ($r_t$), real pension value ($\kappa_t^e$), and income uncertainty ($\sigma_{z,t}^e$). Two exceptions are the relative price (p_t) and preference shifter (χ_t), which follow the processes estimated directly from data and apply to both education groups.

[55] To ensure a non-explosive income process, the persistence of shocks is capped at 0.99 during estimation.

[56] In each iteration, I solve value functions backward from period T to the current period o under the assumed persistence ϱ_o^e for all $t \geq o$.

Outer loop: In the outer loop, I estimate the education gaps $\tilde{\Theta} = \{\alpha_\varrho, \alpha_\sigma\}$ by minimizing the distance between model and data across all periods for both aggregate moments and education-specific differences:

$$\mathbf{F}^{outer}(\tilde{\Theta}) = W \times |\mathbf{M}^{outer}(\tilde{\Theta})|$$

$$\mathbf{M}^{outer}(\tilde{\Theta}) = \Big[\hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}\,data} - \hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}(\tilde{\Theta})_{model}, \hat{\mathbf{D}}^{\mathbf{R}}_{\mathbf{t}\,data} - \hat{\mathbf{D}}_{\mathbf{t}}(\tilde{\Theta})_{model}, \Delta_e\hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}\,data} - \Delta_e\hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}(\tilde{\Theta})_{model}, \Delta_e\hat{\mathbf{D}}^{\mathbf{R}}_{\mathbf{t}\,data} - \Delta_e\hat{\mathbf{D}}_{\mathbf{t}}(\tilde{\Theta})_{model}\Big]$$

W is the observation weights where observations are weighted by the inverse of the square root of the number of periods ($t^{-1/2}$).[57] Each evaluation of the outer loop objective requires solving the inner loop for the full sequence $\{\Theta_o\}_{o=1}^{T}$ given candidate gaps. In practice, the inner loop is solved with a nelder-mead algorithm, while the outer loop is solved on a grid of parameters.[58]

Dynamic shock timing: I assume the arrival of the government transfers is anticipated one period in advance. This implies the first set of transfers is expected at the beginning of the recession (2020.Q1), while the second set of transfers is not anticipated until 2020.Q4. Relative to the static model, I also feed in a real rate shock. In mid-2021, the rise in inflation resulted in a decline in the real interest rate (See Figure VI Panel e).[59]

4.4.2 Dynamic model results:

Figure VIII compares the dynamic expectations model (black solid line) to the data (black dashed) and the static model from Section 4.3 (red dotted line). Dashed magenta lines show the consumption paths in the dynamic model households expected at time t given their period-t beliefs about shock persistence. The dynamic model substantially outperforms the static specification, particularly for durables, and reveals how household expectations evolved throughout the recession.

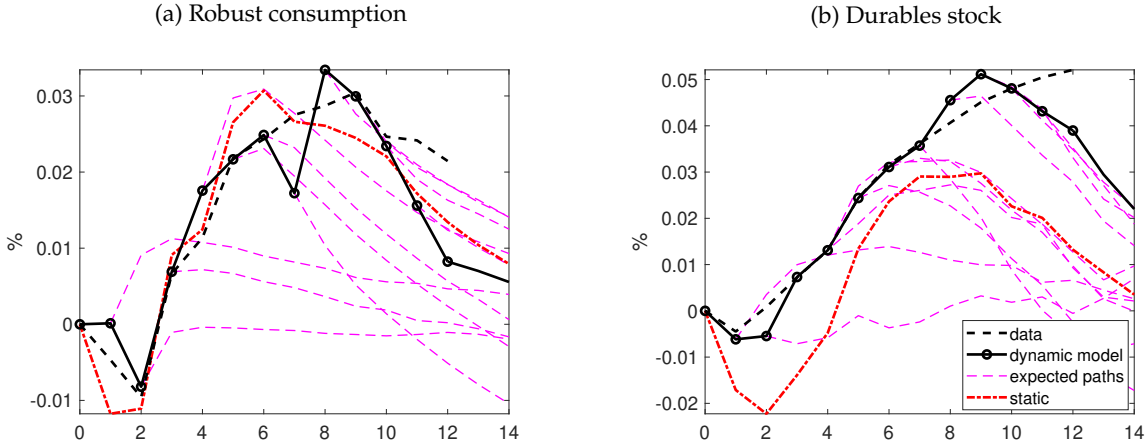
Robust consumption: Panel a shows that the dynamic model captures the robust consumption trajectory in the data very accurately. It captures the magnitude of the decline in the second quarter (2020.Q2) and subsequent boom and recovery. The model with dynamic expectations provides a better fit than the static model. The improvement is modest, the static model already fits reasonably well, but the static specification over-predicts the initial decline and exaggerates the recovery-phase boom. The fit for excluded consumption also improves (Appendix Figure A.24).

[57] The aggregate series from the inner loop are included as there does not necessarily exist an exact solution to the non-linear model. The education gaps are also weighted at 0.5.

[58] Model predictions depend on both $\tilde{\Theta}$ and the implied sequence $\{\Theta_o\}_{o=1}^{T}$. A discussion of the identification from the education-specific differences is provided in the Appendix, Section E.6.

[59] The increase in inflation is also assumed to lower the real value of retired households' pensions. This is already included in the static estimation exercise. Full details on the construction of the shocks are in the Appendix Section E.3.2.

Figure VIII: Consumption dynamics in a dynamic model of expectations



Notes: Response to Covid recession with dynamic expectations. Data is dashed black line. Static expectations (red); Dynamic expectations (solid black); Expected path in each period (magenta).

Durables consumption: The dynamic model's superior fit to the data is more evident for durables. Panel b plots the durables stock.[60] The dynamic model matches the initial decline in durables closely and, while some overshooting remains in 2020:Q2, the fit dramatically improves relative to the static model's steep predicted drop. More strikingly, the dynamic model successfully replicates the large durables boom during 2021-2022. This is a feature the static model entirely misses. Two factors drive this success. First, gradual learning means households initially cut durables cautiously, avoiding the excessive retrenchment predicted under immediate full information. Second, the mid-2021 decline in real interest rates—driven by rising inflation—encourages households to shift from financial savings toward durable goods during the recovery. The static model does not feature this economic driver and cannot produce the observed late-recovery boom.[61] A full quantitative analysis of the model fit is provided in section E.7.

Evolution of expectations: The magenta dashed lines in Figure VIII reveal how household beliefs evolved. Early in the recession (2020:Q2-Q3), households held pessimistic expectations: given their perceived shock persistence, they anticipated consumption paths well below the actual trajectories. These pessimistic beliefs reflected both the severity of observed shocks and expectations of prolonged weakness. As the recession progressed, positive surprises—including substantial fiscal transfers and faster-than-expected labor market recovery—led households to revise beliefs upward. By late 2021, expected consumption paths converge to realized values, indicating household expectations aligned with the actual recovery trajectory. This pattern of initial pessimism followed by positive surprises and belief updating cannot be captured by the static model, where households hold correct beliefs from the outset.

[60] I impute the durables stock in the data using observed durables consumption and the model's depreciation rate. Durables stock provides a clearer comparison than the volatile consumption flow. See Appendix Figure A.24 for durables consumption.

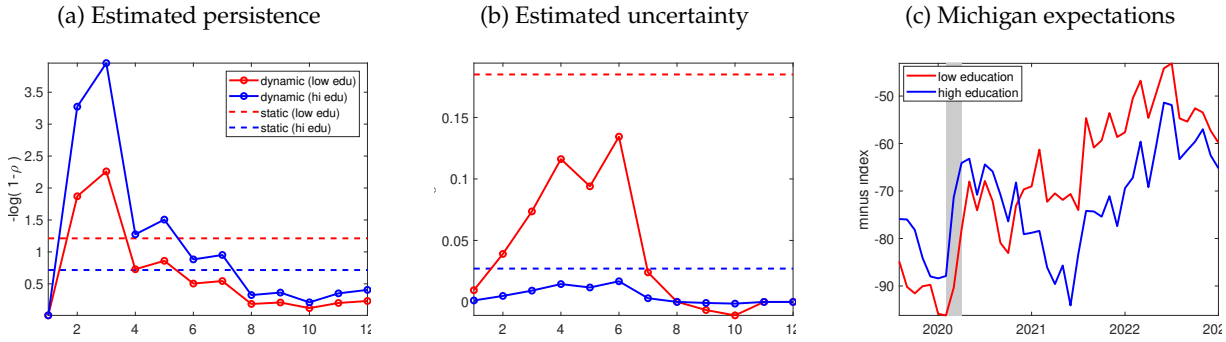
[61] Alternatively, a static model that reveals this interest rate decline in 2020:Q1, generates implausibly large durables holdings early in the recession.

4.4.3 Path of household expectations

Figure IX displays the estimated paths of household beliefs that rationalize observed consumption dynamics. Expected shock persistence and perceived income uncertainty exhibit distinct dynamics throughout the recession.

Shock persistence: Panel a shows households initially perceive shocks as highly transitory in 2020:Q1, explaining the modest consumption response despite severe labor market disruptions. Perceived persistence rises sharply as the recession unfolds, peaking at 0.895 for low-education households (high education: 0.98) in the third quarter of the recession. This is driven by the significant aggregate fall in robust consumption, with only a minor fall in TFP, and labor market transitions recovering. Persistence expectations then decline gradually, returning to transitory levels by 2022 as the recovery solidifies. The model identifies important heterogeneity by education: high-education households perceive greater persistence than low-education households throughout the recession. This differential stems from the consumption data, where robust consumption of the high-education group recovers more slowly, indicating these households anticipated a more prolonged downturn despite facing less severe employment shocks.

Figure IX: Implied household expectations



Notes: Panel a and b, implied household expectations in dynamic expectations model. Panel a is persistence ρ shown in log form for readability, Panel b is uncertainty σ. High education (blue), low education (red). Panel c is Consumer Expectations Index from Michigan Survey of Consumers.

Income uncertainty: Panel b plots the evolution of perceived income uncertainty, which follows a different trajectory. Uncertainty rises more gradually than persistence, peaking in 2021. Low-education households face substantially larger uncertainty increases, consistent with their more precarious labor market positions.[62] Notably, the uncertainty peak coincides with the receipt of fiscal transfers. The model reconciles the modest observed consumption response to these transfers by inferring elevated uncertainty: households saved substantial portions of transfers rather than spending. This interpretation aligns with survey evidence from Coibion et al. (2020), who document low intended marginal

[62]Note: I hold fixed the relative difference in the increases in the variance of uncertainty. See equation 15.

propensities to consume from CARES Act payments, with many households reporting plans to save transfers or pay down debt.[63]

Validation with survey data. Panel c compares model-implied expectations to the Michigan Survey of Consumers' Index of Consumer Expectations, reversed so increases indicate deteriorating expectations.[64] The survey index tracks the model's estimated paths remarkably closely, validating both the time-series dynamics and education heterogeneity. From 2021 onward, low-education survey respondents report substantially worse expectations than high-education respondents, precisely mirroring the model's prediction of greater uncertainty for low-education households. This independent validation strengthens confidence that the model correctly identifies the evolution of household beliefs. Section E.5 presents additional supporting results for spending expectations. Section E.6 in the Appendix explores the implications of this education heterogeneity in greater depth.

4.4.4 Inflation expectations: additional predictions of the dynamic model

The model generates predictions for variables beyond the consumption moments used in estimation, providing out-of-sample validation tests. These untargeted moments help assess whether the dynamic or static expectations specification better captures household behavior.

Durables adjustment probability: Panel a of Figure X compares the model-predicted probability that households adjust their durables stock to the adjustment rate observed in the CEX.[65] The dynamic expectations model closely tracks both the magnitude and timing of adjustment probability fluctuations during the recession. While the model exhibits a slight lag, it captures the initial decline, the subsequent recovery, and the elevated adjustment rates in 2021-2022. In contrast, the static model overpredicts the initial decline and generates an implausibly smooth time path—lacking the volatility observed in the data because households hold fixed beliefs after 2020:Q1. The dynamic model's success in matching adjustment timing provides independent evidence that households gradually updated beliefs rather than immediately recognizing the recession's full trajectory.

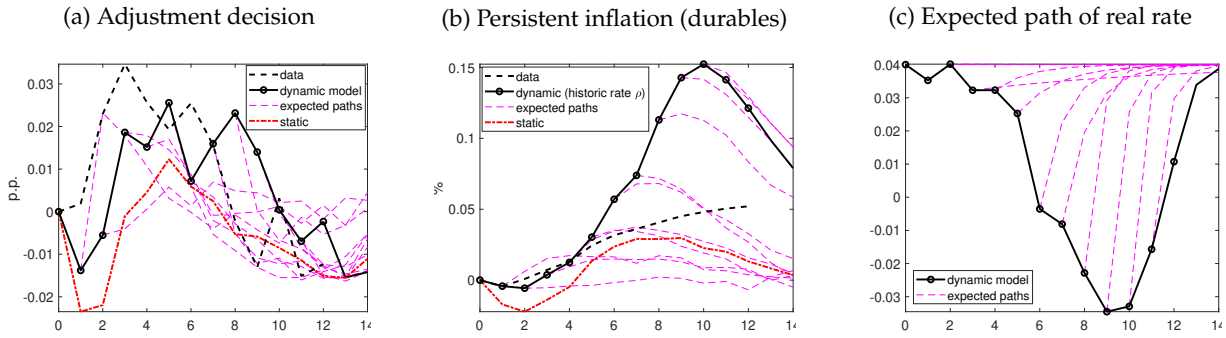
Inflation expectations: The model also provides insight into household inflation expectations. The declining real interest rate in 2021—driven by rising inflation while nominal rates remained near zero—improves the model's fit by stimulating durables consumption during the recovery (Figure VIII, Panel b). Lower real rates make durables a more attractive savings vehicle for transferring resources to the future relative to bonds. However, this channel's strength depends critically on whether households expect the real rate decline

[63] Baker et al. (2020) further shows spending responses varied substantially with household liquidity, consistent with the model's heterogeneous uncertainty estimates.

[64] The survey asks about a household's financial conditions and business conditions in the next year, and the likelihood of widespread unemployment or an economic depression during the next five years. The Index of Consumer Expectations combines these measures.

[65] I calculate the adjustment probability in the CEX as the fraction of households reporting positive durables purchases each quarter.

Figure X: Consumption dynamics in a dynamic model of expectations



Notes: Model of dynamic household expectations. See Figure VIII. Panel a is the probability of adjustment. Panel b is the durables stock in a model solved with data persistence of the real interest rate. Panel c is the path of real interest rates and expected path in each period.

to persist. Panel b of Figure X shows a counterfactual where households expect the real rate to mean-revert with its historical persistence of 0.75.[66] Under this assumption, the model predicts a durables boom three times larger than observed—households would aggressively substitute toward durables, anticipating prolonged low real rates. The data, therefore, reveal that households initially expected inflation (and thus the real rate decline) to be temporary.[67] Panel c shows the implied evolution of household beliefs about real rate persistence (magenta dashed lines). Although the Fed's tightening was delayed, households' expected speed of real rate normalization—measured from the 2021 trough—closely matches the eventual policy path. This consistency suggests household expectations tracked Fed guidance: initially anticipating temporary inflation, then gradually updating as inflation proved more persistent and the Fed signaled tighter policy.

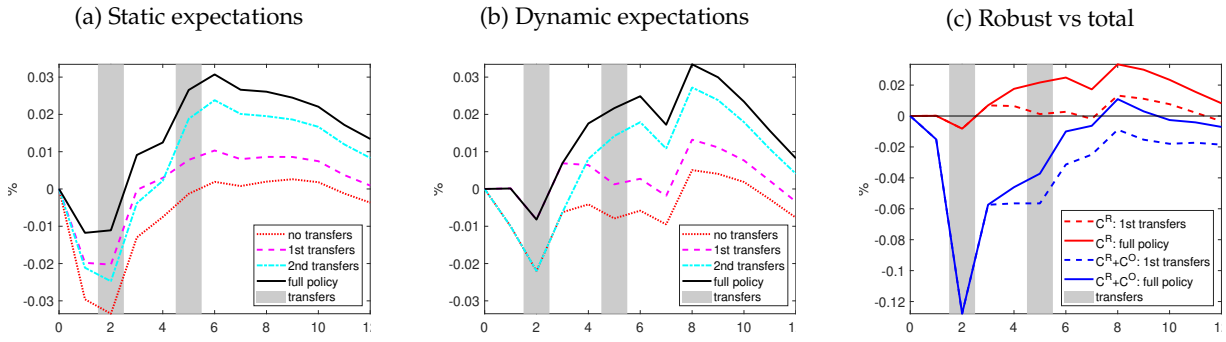
5 Robust consumption and policy analysis

5.0.1 Impact of transfers policy

The Covid-era fiscal transfer program was unprecedented in scale (Figure A.17). This section quantifies the contribution of transfers to consumption using the estimated model. I compare baseline consumption paths to counterfactuals where one or both stimulus rounds are not enacted, holding fixed households' expectations of other economic shocks.[68] The analysis reveals that the timing of information about transfers critically affects their estimated impact, with important implications for evaluating counter-cyclical fiscal policy. In the Appendix, Section E.8, I decompose the effect of the other shocks during the recession.

[66] I estimate this persistence using quarterly real interest rate data from 1960-2019.
[67] This inference aligns with Federal Reserve communication during 2021. Fed Chair Powell and other officials repeatedly characterized inflation as "transitory," suggesting monetary policy would not respond aggressively in the near term (Powell, 2021).
[68] Households correctly perceive whether they will receive transfers.

Figure XI: Robust consumption dynamics with and without transfers



Notes: Panels a and b: Response of robust consumption under alternative fiscal transfer schemes. Panel c: Comparison of both stimulus versus first stimulus. Robust consumption (red), Total consumption is robust plus excluded consumption (blue).

Consumption paths with and without transfers: Figure XI displays consumption under baseline and no-transfer scenarios. The information structure matters substantially. Under static expectations (Panel a), households anticipate both transfer rounds in 2020:Q1, causing non-constrained households to increase consumption immediately—before receiving any payments. Transfers thus appear highly effective at offsetting the initial recession impact. Under dynamic expectations (Panel b), transfer timing becomes crucial: households learn about the first round in 2020:Q1 but do not anticipate the second round until 2020:Q4. Consequently, the counterfactual no-transfer decline in early 2020 is more modest than under static expectations. This distinction affects assessments of recession severity: incorrectly assuming static expectations overstates how much transfers mitigated the downturn because it attributes the entire early-2020 consumption response to anticipated transfers rather than recognizing that the second round was unanticipated.

Quantifying transfer effectiveness: Table VII reports consumption responses in the three-quarter window centered on transfer receipt.[69] The magnitudes are substantial. The first stimulus round increases robust consumption by 3.6-3.7 percent regardless of the expectations assumption. Durables consumption rises more dramatically: 23.6 percent under static expectations and 35.0 percent under dynamic expectations. The second round, being larger, generates bigger responses: robust consumption increases 5.2-5.8 percent, while durables consumption rises 33.4-44.0 percent. Dynamic expectations consistently produce larger consumption responses because households face tighter borrowing constraints when they cannot smooth consumption against unanticipated future transfers.

Non-linearities and diminishing returns: The combination of the two stimulus rounds is non-linear (columns (3) and (6)). When households have already received the first transfer, the response of robust consumption in the window around second transfer is larger, albeit less than the additive sum. For durables, the pattern is even more non-linear with the

[69] Under dynamic expectations, households learn about transfers one quarter in advance, so consumption responds before actual receipt.

Table VII: Contribution of government transfers to consumption

	Stimulus Response					
	Static			Dynamic		
	(1)	(2)	(3)	(4)	(5)	(6)
Stimulus round	1^{st}	2^{nd}	2^{nd}	1^{st}	2^{nd}	2^{nd}
Counterfactual:	none	none	1^{st}	none	none	1^{st}
Robust Consumption						
Baseline	3.6	5.2	7.7	3.7	5.8	8.2
No uncertainty	3.7	5.4	8.0	3.8	6.1	8.5
Durables consumption						
Baseline	23.6	33.4	28.6	35.0	44.0	14.7
No uncertainty	28.9	37.2	28.2	38.2	58.2	25.1

Notes: Effect of transfers with date t on consumption response in period $t-1, t, t+1$. Calculated as baseline model minus model without transfers. (1) and (4) are response to the 1^{st} transfers, (2) and (5) response to the 2^{nd} transfers where the counterfactual is no transfers. (3) and (6) are the response to the 2^{nd} transfers, where the counterfactual is only 1^{st} transfers. No uncertainty is a model with no increase in household uncertainty in counterfactual or baseline.

marginal spending response declining. The first stimulus induces households to accelerate durable purchases, bringing their stocks closer to desired levels.[70] The second stimulus, therefore, requires smaller additional expenditures to maintain target durables stocks. Under dynamic expectations, the second-round durables response falls by 29 percentage points when the first round is present, illustrating how the timing of durable adjustment magnifies nonlinear effects.

Finally, elevated uncertainty substantially reduces transfer effectiveness. The bottom row of Table VII shows consumption responses under a counterfactual with no uncertainty increase. Average consumption responses would be 4 percent larger without heightened uncertainty and 70 percent higher for durables consumption, as precautionary saving motives would be weaker. This finding underscores that fiscal policy effectiveness depends critically on household expectations—not only about future transfers but also about income risk and recession duration.

5.0.2 Robust consumption as a measure of the output gap

The robust consumption measure also reveals a potential policy mistake during the recovery. Standard aggregate consumption data may have misled policymakers about the appropriate timing and magnitude of the second stimulus round. Figure XI, Panel c compares robust consumption to total consumption (the sum of robust and excluded consumption) under two scenarios: both stimulus rounds versus only the first round.

The first stimulus clearly provides appropriate support: both consumption measures remain below trend after the transfer, matching observed data. The second stimulus presents a different picture depending on which consumption measure guides policy. Total consumption remains below its pre-recession level in early 2021, suggesting continued economic weakness and justifying additional fiscal support. Robust consumption, in

[70] This is similar to the logic of McKay and Wieland (2021).

contrast, has already returned to its steady-state level before the second stimulus arrives. The additional transfers push robust consumption further above trend, potentially generating excess demand.[71]

This distinction has important implications for inflation. In a standard New Keynesian framework, the output gap—the deviation of actual output from potential—determines inflationary pressure. If robust consumption better measures the output gap by filtering out pandemic-specific consumption distortions, then the economy was closer to potential in early 2021 than total consumption suggested. The second stimulus would then generate larger inflationary pressure than policymakers anticipated when using total consumption as their guide. A caveat is that the analysis is partial equilibrium, so it is not possible to definitively conclude that the second stimulus increased inflation.[72] This analysis illustrates the value of robust consumption measures for counter-cyclical policy design during structural shocks that distort consumption baskets. When consumption composition shifts dramatically, aggregate consumption can provide misleading signals about economic slack. Partitioning out non-economic distortions provides a clearer view of underlying demand pressures and the required fiscal stimulus.

5.0.3 Impact of expectations on effectiveness of transfers

The analysis above shows that household expectations substantially affect consumption responses to transfers (Table VII). This creates a challenge for policymakers: when designing the second stimulus in late 2020, they faced uncertainty not only about economic conditions but also about household beliefs. How sensitive was the anticipated transfer response to potential misperceptions of household expectations? I address this question by computing the range of consumption responses a policymaker might have predicted in 2020:Q4, conditional on different beliefs about household uncertainty and persistence expectations. Relative to the previous section, these are expectations at time t, assuming no future shocks (Figure XII Panel a). The exercise proceeds as follows. I fix the economic shock path at 2020:Q4 values and compute the expected consumption response to the second transfer (arriving in 2021:Q1) under alternative assumptions about household beliefs.

The baseline case (black line in Panels b and c) uses the beliefs estimated from the dynamic model in 2020:Q4—the policymaker's best guess given observed consumption. I then recompute the transfer response under every belief combination estimated across all periods of the dynamic model, plus the static expectations beliefs, and the no-uncertainty case. The shaded region shows the envelope of these responses, representing the range of transfer effectiveness a policymaker might have anticipated given uncertainty about household beliefs. Panel b shows the robust consumption response varies within a narrow range: the impact effect spans 1.9 to 2.2 percent in 2021.Q1 across all belief specifications. This relative

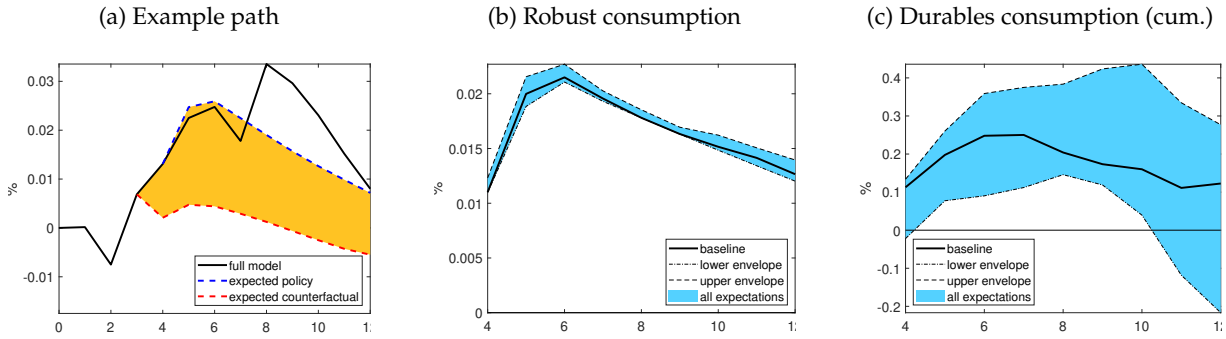
[71] Giannone and Primiceri (2024) also attribute the inflation episode to demand side forces. Bernanke and Blanchard (2025) stress the importance of supply-side restrictions in addition to fiscal policy.

[72] I do not explicitly model the inflation response here, but the mechanism operates through standard New Keynesian channels where positive output gaps generate demand-pull inflation.

stability, although non-trivial, reflects that non-durable consumption is less sensitive to uncertainty and persistence expectations. Panel c reveals much wider uncertainty for durables consumption. Two mechanisms drive this dispersion. First, durables purchases are highly sensitive to perceived uncertainty—heightened uncertainty induces households to delay lumpy purchases. Second, state dependence amplifies forecast uncertainty: the timing of past adjustments affects current adjustment probabilities, creating path dependence. Initially, the baseline belief specification (black line) produces a durables response near the upper bound—households perceive moderate uncertainty and adjust durables relatively freely. Over subsequent quarters, alternative belief paths generate larger cumulative responses as they imply lower perceived uncertainty or less persistent shocks.

The policy implication is striking: given plausible uncertainty about household beliefs in late 2020, policymakers faced substantial forecast uncertainty about the durables response to the second stimulus—roughly a three-fold range in cumulative responses. If policymakers overestimated household uncertainty, they would have underpredicted the durables boom. Conversely, underestimating uncertainty would lead to overpredicting the response. This sensitivity helps explain why the strength of durables spending during 2021 may have surprised policymakers: correctly forecasting transfer effectiveness required not just understanding economic conditions but also accurately measuring household beliefs about recession persistence and income risk.

Figure XII: Household expectations and transfer effectiveness



Notes: Panel (a) illustrates the exercise: the marginal response is the difference between the dashed blue line (with second transfer) and dashed red line (without second transfer), computed from period 4 forward.

5.1 Welfare cost of recession

Finally, I quantify the welfare cost of the Covid recession's economic shocks using consumption-equivalent variation: the percentage reduction in steady-state consumption that leaves households indifferent to experiencing the recession. Because the model includes preference shocks (χ_t) that alter the desirability of excluded consumption, I construct the counterfactual as an economy with only the preference shock versus the full recession (all economic shocks plus the preference shock). This isolates the welfare cost of economic

disruptions—productivity declines, labor market deterioration, and uncertainty—from pandemic-induced preference changes that do not represent true welfare losses.[73]

Welfare costs depend on household expectations about the recession's severity and persistence. I compute three estimates corresponding to different information structures, all evaluated in 2020:Q2 when the recession's impact peaks. First, the *static expectations* case uses beliefs from Section 4.3, where households immediately recognize the full shock path. Second, the *dynamic expectations* case uses 2020:Q2 beliefs from Section 4.4, where households gradually learn. Third, the *perfect foresight* case assumes households correctly anticipate the entire sequence of shocks, including uncertainty shocks as estimated in the dynamic model. This is a static (or perfect foresight) interpretation of the dynamic model's path of expectations.

Table VIII: Welfare cost of recession

	Consumption equivalent loss		
	Static	Dynamic	Perfect Foresight
Baseline	-1.9	-1.2	-1.4
low education	-2.4	-1.5	-2.2
high education	-1.2	-0.9	-0.5
No transfers	-2.5	-1.4	-1.7
Alternative (1980s)	-2.4	-1.4	-3.1

Notes: welfare cost of recession calculated in model given assumption on household expectations. Static: static expectations, Dynamic: dynamic expectations, Perfect Foresight: Implied dynamic path of expectations, with information revealed at time zero.

Aggregate welfare costs: Table VIII, row 1, reports consumption-equivalent welfare losses. The recession imposed substantial costs, ranging from 1.2 to 1.9 percent of steady-state consumption depending on the information structure. Static expectations generate the largest welfare loss (1.9 percent) due to the large uncertainty increase perceived by low-education households, combined with immediate recognition of shock severity. The dynamic specification yields the lowest estimate (1.2 percent), as economic expectations and realized wages continue to deteriorate after the time welfare is calculated, leading to larger losses in the static and perfect foresight cases. The perfect foresight assumption produces an intermediate loss (1.4 percent), mainly because the implied increase in uncertainty is smaller than in the static case. Row 4 shows that fiscal transfers reduced welfare costs but had a modest impact relative to recession losses. Across specifications, transfers improved welfare by 20-30 percent of the economic cost, lowering the welfare loss from 1.7 to 1.4 percent in the perfect foresight case. This reflects that the economic shocks of the Covid recession were still large.

[73]Implicitly, I do not include welfare losses due to the consumption choices being constrained, as this is assumed to reflect preferences.

Heterogeneity in welfare costs: It is also interesting to consider heterogeneity in welfare losses. Figure A.25 in the Appendix displays the distribution of consumption-equivalent losses by education and employment status. Working-age households suffer substantially larger losses than retirees, whose income sources are less cyclically sensitive. Among working-age households, low-education households experience greater losses than high-education households throughout most of the distribution, reflecting larger job separation rates and tighter borrowing constraints. Importantly, there is a tail of households that experience large losses, exceeding 3 percent of consumption. In particular, a small share of high-education households suffer severe losses exceeding the largest losses of low-education households. These are high-education households who experienced job loss—a low-probability but significant event for this group. They suffer disproportionately because they have large illiquid durable stocks and receive relatively smaller transfers. Further in the dynamic case they perceive the recession as more persistent than low-education households.[74]

5.1.1 Application to the 1980s recession

Comparing the Covid recession to other episodes with large consumption basket distortions provides context for the welfare estimates and demonstrates the framework's broader applicability. I apply the methodology to the early 1980s recessions (1980-82), where inflation cost shocks can be interpreted as similarly affecting the composition of the household consumption basket.[75] Here, the robust consumption framework can differentiate economic responses from compositional distortions.

I construct robust and excluded consumption baskets for the 1980s using the same methodology as for Covid. This implies a different set of CES parameters. The estimated expectations reveal households initially perceived the productivity and labor market shocks as moderately persistent but revised these beliefs upward as the downturn extended.[76] Table VIII, row 5, reports welfare costs for the 1980s recessions: which at 1.4 to 3.1 percent of steady-state consumption, depending on the information structure, are higher than the Covid recession estimates. Two factors explain this difference. First, the 1980s downturn lasted longer: consumption remained depressed for three years versus roughly one year for Covid. Persistent income losses generate larger welfare costs than short-lived disruptions even if the initial shock magnitudes are similar. Second, the 1980s featured no large-scale fiscal transfers to cushion consumption losses, whereas Covid-era transfers offset roughly 20-30 percent of welfare costs (row 4). Most importantly, this comparison provides a useful benchmark against which to assess the welfare cost of the Covid recession.

[74] In the static model, high-education households face a greater increase in uncertainty than low-education households, generating large welfare losses via a different channel.

[75] In addition, as nominal interest rates peaked above 15 percent in response to Federal Reserve anti-inflation policy, households facing higher credit costs substantially reduced purchases of interest-sensitive durables—particularly housing and automobiles. I partition goods based on the consumption change relative to historical volatility in April 1980.

[76] Figure A.26 presents the model fit for consumption paths and period-by-period expectations.

6 Conclusion

This paper has proposed a robust measure of consumption for inferring household expectations during the Covid recession and recovery. It proposes a way of computing this by partitioning aggregate consumption and studying a component that is informative about underlying economic conditions. It combines the new consumption measure with a sophisticated life cycle model to extract information on households' expectations during this period. Robust consumption declined by 1.1 percent at the beginning of the recession, with the structural model implying households experienced a significant increase in uncertainty.

The analysis allows us to quantify the role of government stimulus in boosting consumption. In the absence of the transfers, consumer spending would have been 2 to 3 percentage points lower, and the transfers reduced the welfare cost of the recession by 10-15 percent. As robust consumption had largely recovered by the time the second stimulus took place, this also supports the view that the transfers contributed to inflationary pressures. A mitigating consideration highlighted by the research is that, due to the importance of household expectations for consumption responses, policymakers' forecasts for the impact of stimulus were subject to significant forecast error. Further, as the inflation response is not endogenous in the model, one can not definitively conclude that the transfers caused the inflationary spike. Finally, the research remains silent on the sources of information households use to form their beliefs. Clearly, this is of importance, but beyond the scope of the current paper and is left for future research.

References

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

Alvarez, F., Argente, D., and Lippi, F. (2021). A simple planning problem for covid-19 lock-down, testing, and tracing. *American Economic Review: Insights*, 3(3):367–82.

Andersen, A. L., Hansen, E. T., Johannesen, N., and Sheridan, A. (2020). Pandemic, Shutdown and Consumer Spending: Lessons from Scandinavian Policy Responses to COVID-19. Papers 2005.04630, arXiv.org.

Atkeson, A. (2020). What will be the economic impact of covid-19 in the us? rough estimates of disease scenarios. Staff Report 595, Federal Reserve Bank of Minneapolis.

Attanasio, O., Hurst, E., and Pistaferri, L. (2014). The Evolution of Income, Consumption, and Leisure Inequality in the United States, 1980–2010. In *Improving the Measurement of Consumer Expenditures*, NBER Chapters, pages 100–140. National Bureau of Economic Research, Inc.

Attanasio, O., Larkin, K., Ravn, M. O., and Padula, M. (2022). (S)Cars and the Great Recession. *Econometrica*, 90(5):2319–2356.

Attanasio, O. P. and Pistaferri, L. (2016). Consumption inequality. *Journal of Economic Perspectives*, 30(2):3–28.

Attanasio, O. P. and Weber, G. (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, 103(6):1121–1157.

Attanasio, O. P. and Weber, G. (2010). Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature*, 48(3):693–751.

Auclert, A., Rognlie, M., and Straub, L. (2024). The intertemporal keynesian cross. *Journal of Political Economy*, 132(12):4068–4121.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2020). Income, liquidity, and the

consumption response to the 2020 economic stimulus payments. Working Paper 27097, National Bureau of Economic Research.

Baqaee, D. R., Burstein, A. T., and Koike-Mori, Y. (2024). Measuring welfare by matching households across time. *The Quarterly Journal of Economics*, 139(1):533–573.

Beraja, M. and Wolf, C. K. (2021). Demand composition and the strength of recoveries. Working Paper 29304, National Bureau of Economic Research.

Berger, D., Herkenhoff, K., Huang, C., and Mongey, S. (2022). Testing and reopening in an seir model. *Review of Economic Dynamics*, 43:1–21.

Berger, D. and Vavra, J. (2015). Consumption dynamics during recessions. *Econometrica*, 83(1):101–154.

Bernanke, B. and Blanchard, O. (2025). What caused the us pandemic-era inflation? *American Economic Journal: Macroeconomics*, 17(3):1–35.

Bernheim, B. D., Skinner, J., and Weinberg, S. (2001). What Accounts for the Variation in Retirement Wealth among U.S. Households? *American Economic Review*, 91(4):832–857.

Bertola, G., Guiso, L., and Pistaferri, L. (2005). Uncertainty and Consumer Durables Adjustment. *Review of Economic Studies*, 72(4):973–1007.

Blanchard, O. (1993). Consumption and the recession of 1990-1991. *The American Economic Review*, 83(2):270–274.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, 98(5):1887–1921.

Blundell, R. and Preston, I. (1998). Consumption Inequality and Income Uncertainty. *The Quarterly Journal of Economics*, 113(2):603–640.

Boppart, T., Harmenberg, K., Krusell, P., and Olsson, J. (2022). Integrated epi-econ assessment of vaccination. *Journal of Economic Dynamics and Control*, 140:104308. Covid-19 Economics.

Burke, M. A. and Ozdagli, A. (2023). Household inflation expectations and consumer spending: Evidence from panel data. *The Review of Economics and Statistics*, 105(4):948–961.

Cajner, T., Crane, L. D., Decker, R. A., Grigsby, J., Hamins-Puertolas, A., Hurst, E., Kurz, C., and Yildirmaz, A. (2020). The us labor market during the beginning of the pandemic recession. *Brookings Papers on Economic Activity*, 51(2 (Summer):3–33.

Carvalho, V. M., Garcia, J. R., Hansen, S., Ortiz, , Rodrigo, T., Rodríguez Mora, J. V., and Ruiz, P. (2021). Tracking the covid-19 crisis with high-resolution transaction data. *Royal Society Open Science*, 8(8):210218.

Chetty, R., Friedman, J. N., Stepner, M., and Team, O. I. (2023). The economic impacts of covid-19: Evidence from a new public database built using private sector data*. *The Quarterly Journal of Economics*, 139(2):829–889.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). How did u.s. consumers use their stimulus payments? Working Paper 27693, National Bureau of Economic Research.

Commault, J. (2022). Does consumption respond to transitory shocks? reconciling natural experiments and semistructural methods. *American Economic Journal: Macroeconomics*, 14(2):96–122.

Cox, N., Ganong, P., Noel, P., Vavra, J., Wong, A., Farrell, D., Greig, F., and Deadman, E. (2020). Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data. *Brookings Papers on Economic Activity*, pages 35–69.

De Nardi, M. (2004). Wealth Inequality and Intergenerational Links. *Review of Economic Studies*, 71(3):743–768.

Di Tella, S., Hébert, B., Kurlat, P., and Wang, Q. (2026). The zero-beta interest rate. *Journal of Political Economy*, forthcoming.

D'Acunto, F., Charalambakis, E., Georgarakos, D., Kenny, G., Meyer, J., and Weber, M. (2024). Household inflation expectations: An overview of recent insights for monetary policy. Working Paper 32488, National Bureau of Economic Research.

D'Acunto, F., Hoang, D., and Weber, M. (2021). Managing households' expectations with unconventional policies. *The Review of Financial Studies*, 35(4):1597–1642.

Eberly, J. C. (1994). Adjustment of Consumers' Durables Stocks: Evidence from Automobile Purchases. *Journal of Political Economy*, 102(3):403–436.

Eichenbaum, M., Godinho de Matos, M., Lima, F., Rebelo, S., and Trabandt, M. (2024). Expectations, infections, and economic activity. *Journal of Political Economy*, 132(8):2571–2611.

Eichenbaum, M., Rebelo, S., and Wong, A. (2022). State-dependent effects of monetary policy: The refinancing channel. *American Economic Review*, 112(3):721–61.

Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2021). The macroeconomics of epidemics. *The Review of Financial Studies*, 34(11):5149–5187.

Fagereng, A., Onshuus, H., and Torstensen, K. N. (2024). The consumption expenditure response to unemployment: Evidence from norwegian households. *Journal of Monetary Economics*, 146:103578.

Farboodi, M., Jarosch, G., and Shimer, R. (2021). Internal and external effects of social distancing in a pandemic. *Journal of Economic Theory*, 196:105293.

Faria-e Castro, M. (2021). Fiscal policy during a pandemic. *Journal of Economic Dynamics and Control*, 125:104088.

Friedman, M. (1957). *The Permanent Income Hypothesis*. Princeton University Press.

Ganong, P., Noel, P., and Vavra, J. (2020). Us unemployment insurance replacement rates during the pandemic. *Journal of Public Economics*, 191:104273.

Giannone, D. and Primiceri, G. (2024). The drivers of post-pandemic inflation. NBER Working Papers 32859, National Bureau of Economic Research, Inc.

Glover, A., Heathcote, J., Krueger, D., and Ríos-Rull, J.-V. (2023). Health versus wealth: On the distributional effects of controlling a pandemic. *Journal of Monetary Economics*, 140:34–59.

Gorbachev, O. (2011). Did household consumption become more volatile? *American Economic Review*, 101(5):2248–70.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption Over the Life Cycle. *Econometrica*, 70(1):47–89.

Guiso, L. and Jappelli, T. (2025). Anatomy of consumption risk. Technical report, Centre for Studies in Economics and Finance.

Guvenen, F., Ozkan, S., and Madera, R. (2024). Consumption dynamics and welfare under non-gaussian earnings risk. *Journal of Economic Dynamics and Control*, 169:104945. Heterogeneity and Macroeconomics of Labor Markets.

Guvenen, F. and Smith, A. A. (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica*, 82(6):2085–2129.

Hall, R. E. (1987). Consumption. NBER Working Papers 2265, National Bureau of Economic Research, Inc.

Hall, R. E. and Kudlyak, M. (2022). The unemployed with jobs and without jobs. *Labour Economics*, 79:102244.

Hansen, L. P. and Singleton, K. J. (1983). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, 91(2):249–265.

Harmenberg, K. and Öberg, E. (2021). Consumption dynamics under time-varying unemployment risk. *Journal of Monetary Economics*, 118:350–365.

Heathcote, J., Storesletten, K., and Violante, G. L. (2014). Consumption and Labor Supply with Partial Insurance: An Analytical Framework. *American Economic Review*, 104(7):2075–2126.

Hubbard, R. G., Skinner, J., and Zeldes, S. P. (1995). Precautionary saving and social insurance. *Journal of Political Economy*, 103(2):360–399.

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

Imrohoruglu, A. (1989). Cost of Business Cycles with Indivisibilities and Liquidity Constraints. *Journal of Political Economy*, 97(6):1364–1383.

Kaplan, G., Moll, B., and Violante, G. L. (2016). Monetary Policy According to HANK. NBER Working Papers 21897, National Bureau of Economic Research, Inc.

Kaplan, G., Moll, B., and Violante, G. L. (2020). The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the u.s. Working Paper 27794, National Bureau of Economic Research.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kimball, M. S. (1990). Precautionary saving in the small and in the large. *Econometrica*, 58(1):53–73.

Kleibergen, F. and Zhan, Z. (2020). Robust inference for consumption-based asset pricing. *The Journal of Finance*, 75(1):507–550.

Lenza, M. and Primiceri, G. E. (2020). How to Estimate a VAR after March 2020. NBER Working Papers 27771, National Bureau of Economic Research, Inc.

Low, H., Meghir, C., and Pistaferri, L. (2010). Wage Risk and Employment Risk over the Life Cycle. *American Economic Review*, 100(4):1432–1467.

Lucas, Robert E., J. (2003). Macroeconomic priorities. *American Economic Review*, 93(1):1–14.

McKay, A. and Wieland, J. F. (2021). Lumpy durable consumption demand and the limited ammunition of monetary policy. *Econometrica*, 89(6):2717–2749.

Meyer, B. D. and Sullivan, J. X. (2013). Consumption and income inequality and the great recession. *American Economic Review*, 103(3):178–83.

Pappa, E., Ravn, M. O., and Sterk, V. (2023). Expectations and incomplete markets. In **Bachmann, R., Topa, G., and van der Klaauw, W.**, editors, *Handbook of Economic Expectations*, pages 569–611. Academic Press.

Powell, J. (2021). Monetary policy in the time of covid. "Macroeconomic Policy in an Uneven Economy," economic policy symposium sponsored, Jackson Hole, Wyoming.

Samuelson, P. A. (1948). Consumption theory in terms of revealed preference. *Economica*, 15(60):243–253.

Stock, J. H. and Watson, M. W. (2025). Recovering from covid. NBER Working Papers 33857, National Bureau of Economic Research, Inc.

Zeldes, S. P. (1989). Consumption and Liquidity Constraints: An Empirical Investigation. *Journal of Political Economy*, 97(2):305–346.

Online Appendix: Not for publication

A A simple example of identification with robust consumption

A.1 Setting

The household consumes two goods: robust, c_t^r, and excluded consumption, c_t^o. With felicity utility function $u(c_t^r, c_t^o, \chi_t)$ which could depend on time through χ_t. There are two periods. In steady state, the household receives income, y, each period. There is a one period bond for saving, with return $R = 1/\beta$.

Consider a recession in the first period where income falls by z. The household forms expectations about income in period 2, where the income shock is expected to persist with auto regressive coefficient ϱ. Households face income uncertainty. With probability λ, they additionally get a negative income shock σ. With probability $1-\lambda$, they receive a positive income shock of the same size.

The household problem can be written as:

$$\begin{aligned}\max_{c_t^r, c_t^o, b_{t+1}, c_{t+1}^r, c_{t+1}^o} \; & u(c_t^r, c_t^o, \chi_t) + \beta\left[\lambda u(\hat{c}_{t+1}^r, \hat{c}_{t+1}^o, \chi_{t+1}) + (1-\lambda)u(\tilde{c}_{t+1}^r \tilde{c}_{t+1}^o, \chi_{t+1})\right] \\ c_t^r + pc_t^o + b_{t+1} &= y - z \\ \hat{c}_{t+1}^r + p\hat{c}_{t+1}^o &= y - \varrho z - \sigma + Rb_{t+1} \\ \tilde{c}_{t+1}^r + p\tilde{c}_{t+1}^o &= y - \varrho z + \sigma + Rb_{t+1}\end{aligned}$$

where $\hat{c}(\tilde{c})$ is consumption in the negative (positive) income shock state, respectively. p is a constant relative price of the excluded consumption good. Assume that the utility function is CES CRRA:

$$u(c_t^r, c_t^o, \chi_t) = \frac{\left(\{\theta^{\frac{1}{\epsilon}}(c_t^r)^{\frac{\epsilon-1}{\epsilon}} + (1-\chi_t)^{\frac{1}{\epsilon}}(c_t^o)^{\frac{\epsilon-1}{\epsilon}}\}^{\frac{\epsilon}{\epsilon-1}}\right)^{1-\rho} - 1}{1-\rho}$$

In steady state without a recession ($z = 0, \sigma = 0, \chi_t = \theta$), it is trivial to show robust consumption is given by: $c^r = y/[1 + p^{1-\epsilon}(1-\theta)/\theta]$ and $c^o = (y - c^r)/p$. Define the following functions of utility parameters: $\phi_t = 1 + p^{1-\epsilon}(1-\chi_t)/\theta$, $\hat{\phi}_t = 1 + p^{-\epsilon}(1-\chi_t)/\theta$ and $\Lambda_t = (\theta^{1/\epsilon} + p^{1-\epsilon}(1-\chi_t)/\theta^{\frac{\epsilon-1}{\epsilon}})^{\frac{\epsilon}{\epsilon-1}}$. The utility function can be reformulated in terms of aggregate consumption $C_t = c_t^r + c_t^o$ with utility given by: $U(C_t, \chi_t) = \frac{(\Lambda_t/\hat{\phi}_t C_t)^{1-\rho} - 1}{1-\rho}$.

A.2 Euler equation

A.2.1 Standard recession

If there is a recession is the first period, $z > 0$, but the preference shifter remains at its steady state value $\chi_t = \theta$, consumption is given by the aggregate Euler equation.

$$\begin{aligned}C_t^{-\rho} =& \beta R\left[\lambda \hat{C}_{t+1}^{-\rho} + (1-\lambda)\tilde{C}_{t+1}^{-\rho}\right] \\ =& \lambda(y - \varrho z - \sigma + Rb_{t+1})^{-\rho} + (1-\lambda)(y - \varrho z + \sigma + Rb_{t+1})^{-\rho}\end{aligned}$$

As can be seen, aggregate consumption, C_t, is informative about households' expectations of the next period's income.

This is the standard case, where consumption growth is informative about income shocks. However, as $\Delta C_t = \Delta c_t^r + \Delta c_t^o$, if c_t^o declines for another reason, the change in C_t and aggregate Euler equation no longer identifies ϱ.

In the example, there are two sources of variation in households' expectations, whereas the equation only uniquely identifies one of these. An additional restriction is required to separately identify ϱ for σ. In the quantitative exercise, I use durables to do this.

A.2.2 Recession with change in consumption basket

Now, in addition to the fall in income due to the recession, I assume there is also a change in the consumption basket due to a preference shock in period t, which reverts to $\chi_{t+1} = \theta$ in $t+1$. I write the Euler equation in terms of the robust consumption good:

$$\begin{aligned}\frac{\Lambda_t^{1-\rho}}{\phi_t}(c_t^r)^{-\rho} =& \frac{\Lambda^{1-\rho}}{\phi}\beta R\left[\lambda(\hat{c}_{t+1}^r)^{-\rho} + (1-\lambda)(\tilde{c}_{t+1}^r)^{-\rho}\right] \\ =& \left(\frac{\Lambda}{\phi}\right)^{1-\rho}\left[\lambda(y - \varrho z - \sigma + Rb_{t+1})^{-\rho} + (1-\lambda)(y - \varrho z + \sigma + Rb_{t+1})^{-\rho}\right] \\ =& \left(\frac{\Lambda}{\phi}\right)^{1-\rho}\left[\lambda(y(1+R) - (\varrho + R)z - \sigma - R\phi_t c_t^r)^{-\rho} + (1-\lambda)(y(1+R) - (\varrho + R)z + \sigma - R\phi_t c_t^r)^{-\rho}\right]\end{aligned}$$

It can be seen that information on $c_{r,t}$ could be informative for households' expectations. If $\chi_t \to 1$ one can recover the same Euler relationship using only c^r, in particular ϱ is identified from $\log(c_{r,t}) - \log(c_r)$.

$$\theta^{\frac{1-\rho}{\epsilon-1}}(c_t^r)^{-\rho} = \left(\frac{\Lambda}{\phi}\right)^{1-\rho}\left[\lambda(y(1+R) - (\varrho + R)z - \sigma - Rc_t^r)^{-\rho} + (1-\lambda)(y(1+R) - (\varrho + R)z + \sigma - Rc_t^r)^{-\rho}\right]$$

In general, this is not the case, as you also need to observe Λ_t and ϕ_t. Fortunately, this term can be constructed from the intra-temporal optimality condition:

$$c_t^r = p^\epsilon \theta/(1-\chi_t)c_r^o$$

Having, found χ_t I can construct Δ_t and ϕ_t, and use c_t^r to find ϱ.

B Monte Carlo simulation of robust consumption partitioning algorithm

B.1 Parameterization

I use Monte Carlo analysis to simulate multiple draws of the Covid shock using simulated data. This is useful for understanding the properties of the partitioning algorithm and for calibrating the algorithm's hyperparameters. The model is calibrated to replicate features of the NIPA data. Let x_t be the aggregate business cycle state which follows the law of motion:

$$x_{t+1} = \rho^x x_t + \epsilon^x_{t+1} \tag{16}$$

with $\epsilon^x_t \sim \mathcal{N}(0, (\sigma^x)^2)$. In period t^*, a negative shock is realized $\epsilon^x_{t^*} < 0$

Let p_t be the Covid state. $p_t = 0$ in all periods $t < t^*$. When the Covid shock hits, $p_t = \epsilon_{t^*} < 0$ and then follows:

$$p_{t+1} = \rho^p p_t + \epsilon^p_{t+1} \tag{17}$$

The log of an individual consumption good $c^j_t = \log(C^j_t)$ follows:

$$c^j_{t+1} = z^j_{t+1} + \gamma^x_j x_{t+1} + \gamma^p_j \mathbb{1}[j \in C^O] p_{t+1} \tag{18}$$

$$z^j_{t+1} = \rho^{z,j} z^j_t + \epsilon^{z,j}_{t+1} \tag{19}$$

where z^j_t captures idiosyncratic factors that drive consumption of good j.

The parameters γ^x_j and γ^p_j capture the consumption good-specific loading of the aggregate shock and Covid shock. These are assumed to follow:

$$\gamma^x_j \sim \mathcal{N}(\mu^{\gamma^x}, (\sigma^{\gamma^x})^2)$$
$$\gamma^p_j - \underline{\gamma^p} \sim \log \mathcal{N}(\mu^{\gamma^p}, (\sigma^{\gamma^p})^2)$$

Demand for a good is affected by Covid with $Pr(j \in C^O) = \alpha$, else the Covid shock has no impact on consumption demand.[77] The parameters of the idiosyncratic process follow $1 - \rho^{z,j} \sim \log \mathcal{N}(\mu^{\rho^z}, (\sigma^{\gamma^z})^2)$ and $\epsilon^{z,j}_t \sim \mathcal{N}(0, (\sigma^z)^2)$. Finally, an auxiliary series y_t captures a noisy signal of the underlying aggregate state x_t. This replicates the use of income in the data.

$$y_{t+1} = \hat{y}_{t+1} + x_{t+1}$$
$$\hat{y}_{t+1} = \rho^y \hat{y}_t + \epsilon^y_{t+1}$$

with $\epsilon^y_t \sim \mathcal{N}(0, (\sigma^y)^2)$.

All model parameters are calibrated to match the monthly NIPA consumption data. Given uncertainty, I also allow for parameter noise and simulate the economy for N^d draws of the shocks and N^P draws of the parameters.

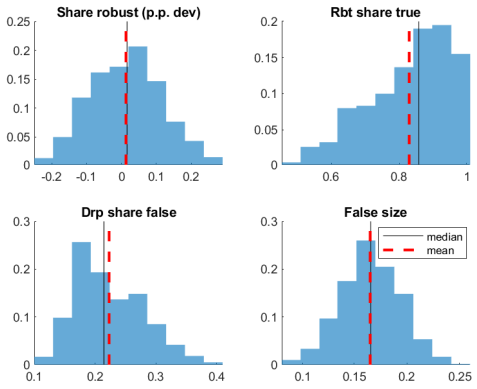
B.2 Results

The Monte Carlo model is used to calibrate hyperparameters in the algorithm, for use on the real data realizations. The true identification of the robust and excluded consumption goods implies identifying correctly the set $j \in C^O$ based on the shocks in period t^*. The cutoff γ is calibrated as the fall of the total consumption basket (data equivalent: PCE) in period t^*, relative to the series average standard deviation. This is chosen to minimize the sum of the share of robust goods incorrectly allocated to the excluded share (Type I error) and the share of excluded goods allocated to the robust share (Type II error). Monte Carlo simulation with $N^d = 15$ and $N^P = 15$, suggests a parameter of $\gamma = 3.8$

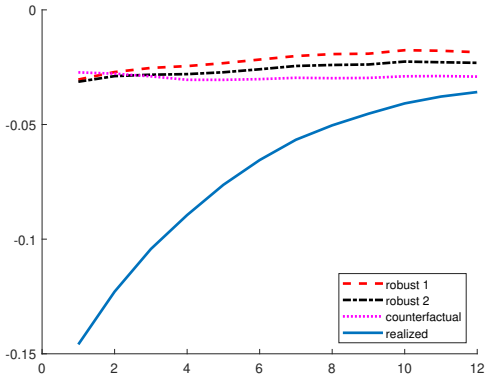
[77] For each j, α is allowed to depend on γ^x_j to induce correlation between the responsiveness to aggregate shocks and the likelihood to being affected by the Covid shock.

Figure A.1: Simulated robust consumption algorithm

(a) Allocation



(b) Example path in simulated data



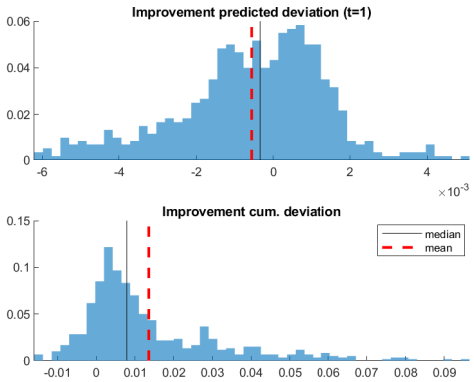
Notes: Results from a Monte Carlo simulation of Covid and recession shocks. In Panel b, *robust 1* is robust consumption after the initial split; *robust 2* is robust consumption after the sorting algorithm; *counterfactual* is the true path of economic shocks; *realized* is the simulated aggregate data series.

In a second stage, I use the Monte Carlo simulation to choose the parameter Γ which governs the weight of the aggregate statistics in the minimization step. I choose this to target the largest improvement in the R^2 of the aggregate state regressed on the composite robust consumption good and the difference in the log change in period t^* of the robust consumption good and aggregate state.[78] Simulation suggests a parameter choice of $\Gamma = 7.75$.

The results of the Monte Carlo simulation are shown in Figure A.1. Panel a shows that the algorithm does a good job of identifying the robust consumption goods in the simulated data. The median share of robust goods correctly identified exceeds 80 percent, while less than 20 percent of excluded goods are incorrectly classified as robust. Further, the size of the pandemic shock for the excluded goods incorrectly allocated to the robust basket is small. It is around 16 percent of the average pandemic shock. Panel b shows the estimated robust consumption path, before and after the sorting algorithm, against its true counterfactual and realized aggregate consumption in the Monte Carlo simulation.

Figure A.2 shows the distribution of improvements in the second stage of the algorithm. On average, the minimization step lowers the predicted cumulative deviations significantly, while the effect on impact is close to zero.

Figure A.2: Prediction improvement



Notes: Results from a Monte Carlo simulation of Covid and recession shocks. Panel a at time 1, Panel b cumulative across horizon.

[78] This essentially puts additional weight on the first period deviation.

C A sorting algorithm to find the robust consumption basket

To find the robust consumption basket, I solve for the allocation that minimizes the objective function in Equation 20. The weights ω_j are from $\omega_j \in \{\varpi, 1-\varpi\}$ (see Section B), so the minimization is over sets of goods allocated to robust or excluded.

$$\begin{aligned}\{\omega_j\}_{j=1}^{J} = \underset{\omega_j \in \{\varpi, 1-\varpi\}}{\operatorname{argmin}} \; & \Gamma\Big\{\Big|\sigma_t(\mathbf{C_t^R})/\sigma_t(\mathbf{C_t^O}) - 1\Big| + \Big|\rho_t(\mathbf{C_t^R}, Y_t) - \rho_t(\mathbf{C_t^O}, Y_t)\Big|\Big\} \\ & - \Big\{\frac{1}{J^O}\sum_{m \in C^O} |\tilde{g}^m_{t^*}| - \frac{1}{J^R}\sum_{m \in C^R} |\tilde{g}^m_{t^*}|\Big\} \\ & + \Big\{\sigma_{\tilde{g}}\big(|\tilde{g}^m_{t^*}| \,\big|\, m \in C^R\big) + \sigma_{\tilde{g}}\big(|\tilde{g}^m_{t^*}| \,\big|\, m \in C^O\big)\Big\}\end{aligned} \tag{20}$$

The idea of the algorithm is to find the allocation that satisfies the property that there exists no alternative allocation with one good reassigned to the excluded set that would lower the value of the objective function.

Starting from an initial allocation ω_j^0, with superscript 0 denoting the first iteration.

1. Given the allocation ω_j^k evaluate the objective function, which gives value f^{*k}.

2. Sort the robust and excluded shocks by normalized absolute size of the Covid decline: $\tilde{g}_s^j = |g_{t^*}^j|/\sigma_t(g_j)$ where $s \in R, O$ indicates (R)obust or excluded (O). All robust shocks are negative. This gives sets:

$$\mathbf{G}^R = \{\tilde{g}_R^1 < \tilde{g}_R^2 < \ldots < \tilde{g}_R^{N_R}\}$$
$$\mathbf{G}^O = \{\tilde{g}_O^1 < \tilde{g}_O^2 < \ldots < \tilde{g}_O^{N_O}\}$$

3. Starting with the largest negative shock of the robust set $\tilde{g}_R^{N_R}$, evaluate the objective function when reassigning one robust shock to the excluded set $f_{R,n}$. If $f_{R,n} < f^{*k}$ `stop` else continue to $f_{R,n-1}$.

4. Starting with the smallest negative shock of the excluded set $\tilde{g}_O^1$ evaluate the objective function when reassigning one excluded shock to the robust set $f_{O,m}$. If $f_{O,m} < f^{*k}$ `stop` else continue to $f_{O,m+1}$.

5. If there exists a $f_{R,n} < f^{*k}$ and $f_{O,m} < f^{*k}$ check if the union satisfies $f_{n \cup m} < \min\{f_{R,n}, f_{O,m}\}$. If `yes` reallocate $\tilde{g}^n$ and $\tilde{g}^m$ to complementary set. If `no` reallocate $\tilde{g}^n$ or $\tilde{g}^m$ which has lower objective function. Given new allocation ω_j^{k+1} `return to step 1`.

6. If $f_{R,n} < f^{*k}$ or $f_{O,m} < f^{*k}$ reallocate $\tilde{g}^n$ or $\tilde{g}^m$ whichever satisfies $f_s < f^{*k}$ Given new allocation ω_j^{k+1} `return to step 1`.

7. If $f_{R,n} > f^{*k}$ and $f_{O,m} > f^{*k}$, for all $n \in N_R$ and $m \in N_R$, there exists no one-step improvements. `End minimization routine`.

D Additional robust consumption material

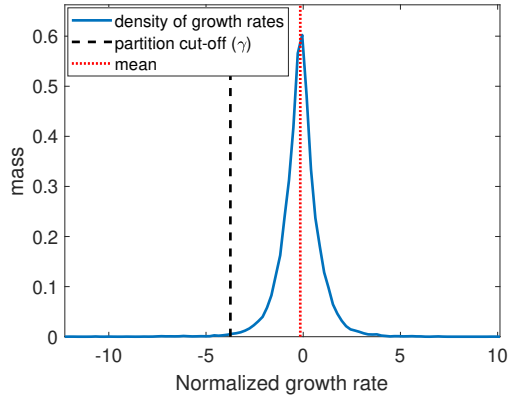
D.1 Growth rate of types during standard recession

Figure A.3 plots the normalized growth rates of items during NBER recessions pre-2020. Growth rates are normalized by the mean and standard deviation of the series in the full time series pre-2020. Panel a shows non-durables and services and Panel b shows durables consumption. The mean for non-durables and services is -0.15 and for durables it is -0.18. I exclude items which are subtracted from consumption, but this has a small impact on the results.

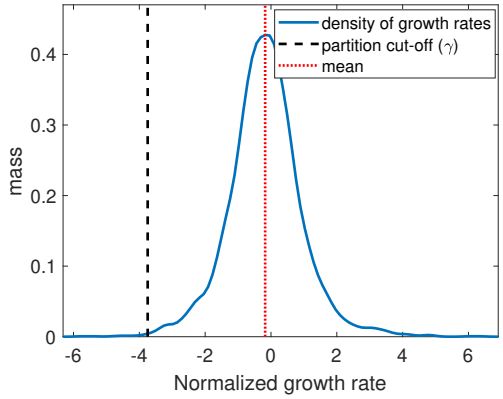
In each picture also plotted is a vertical line consistent with a $\gamma = 3.8$ standard deviation decline. This is the cut-off used to categorize items as belonging to the excluded category in the initial partition. As can be seen, during a normal recession declines exceeding this value are *exceedingly* rare. For consumption during recessions is accounts for 0.57 percent of the observations. For durables it is 0.24 percent.

Figure A.3: Distribution of growth rates during NBER recessions

(a) Non-durables & services



(b) Durables consumption



Notes: Data from NIPA. Normalized growth rates of items during NBER recession periods (pre 2020). Panel a non-durables and services. Panel b durables. Dashed line indicates initial partition cut-off. Red line is mean.

D.2 Sub-sector components of robust consumption

Tables A.1 and A.2 report the five largest share items allocated to the robust and excluded consumption and durables consumption baskets. For robust consumption, important categories include housing rents and utilities such as electricity and cellular telephone services. Significant excluded items include physician services and purchased meals. For durables, robust items include trucks and computing equipment. Excluded items include furniture and new cars.

Table A.1: Largest sub-components: Non-durables and services

	Robust		**Excluded**	
	item	share	item	share
1.	Imputed rent of owner housing	0.29	Receipts non-profit hospitals (NPISH)	0.07
2.	Output of non-profit hospitals (NPISH)	0.14	Non-profit hospitals	0.07
3.	Rent of tenant housing	0.09	Other purchased meal	0.06
4.	Electricity	0.04	Physician services	0.06
5.	Cellular telephones services	0.03	Gasoline and other motor fuel	0.05

Notes: Data from NIPA. Largest sub-components of robust and excluded consumption basket.

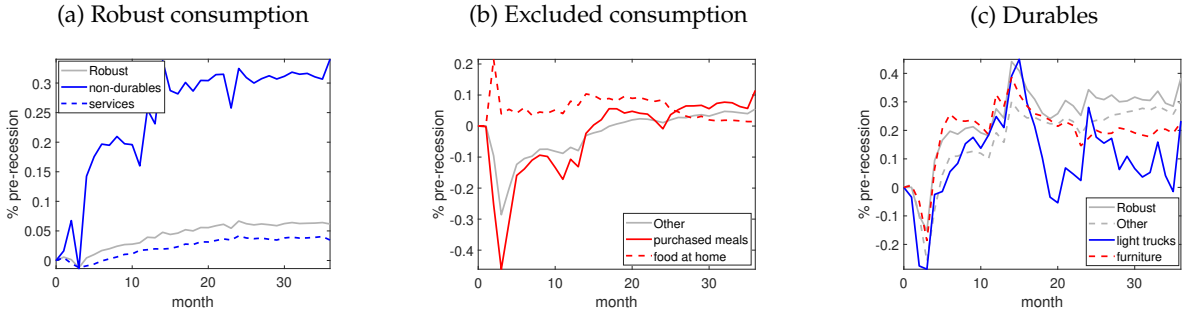
Figure A.4 shows the path of sub-components or individual goods relative to the aggregate baskets. Panel a shows robust consumption decomposed into non-durables and services. Goods rebound more quickly than services. Panel b shows the relative path of food at home and food out, both components of the excluded baskets. Food out falls precipitously while food at home jumps up, showing clear patterns of substitution. Panel c shows durables. Despite being allocated to the robust basket, light trucks falls significantly.

Table A.2: Largest sub-components: Durable consumption

	Robust		Excluded	
	item	share	item	share
1.	New light trucks	0.26	Furniture	0.14
2.	Computer software and accessories	0.12	Jewelry	0.08
3.	Used light trucks	0.12	New domestic autos	0.07
4.	Sporting equipment and guns	0.11	Televisions	0.07
5.	Personal computers and tablets	0.09	Telephone equip	0.07

Notes: Data from NIPA. Largest sub-components of robust and excluded durables consumption basket.

Figure A.4: Sub-sector examples

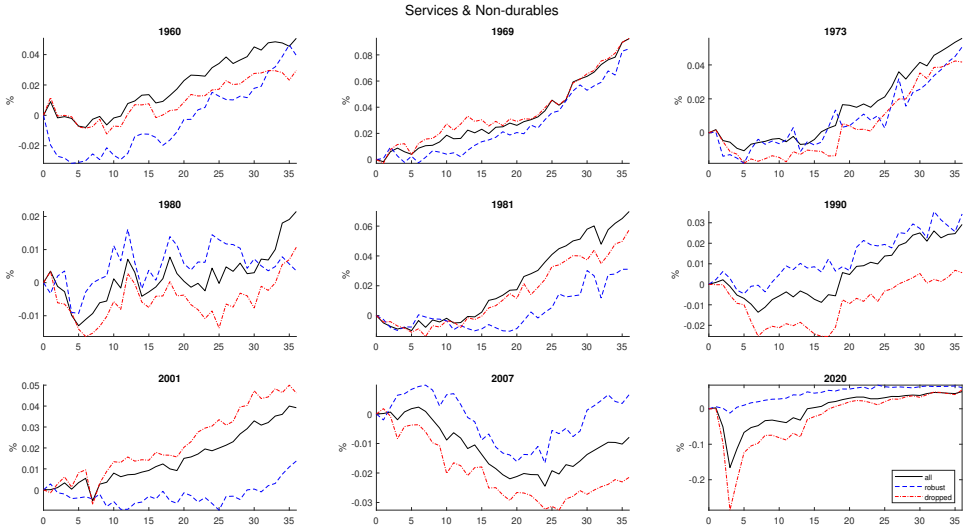


Notes: Data from NIPA. Panel a decomposes robust consumption in goods and services. Panel b shows the path of different food categories, both allocated to excluded. Panel c compares durables items.

D.3 Robust consumption in other recessions

Figure A.5 plots robust, excluded and total consumption in all recessions in the monthly data series. Across the recessions, there is no consistent pattern in the magnitude of the decline between robust and excluded consumption and in most cases, the dynamics are extremely similar. The exception if the Covid recession where, by design, the excluded basket falls by substantially more than aggregate consumption while the robust basket falls by less.

Figure A.5: Comparing Recessions



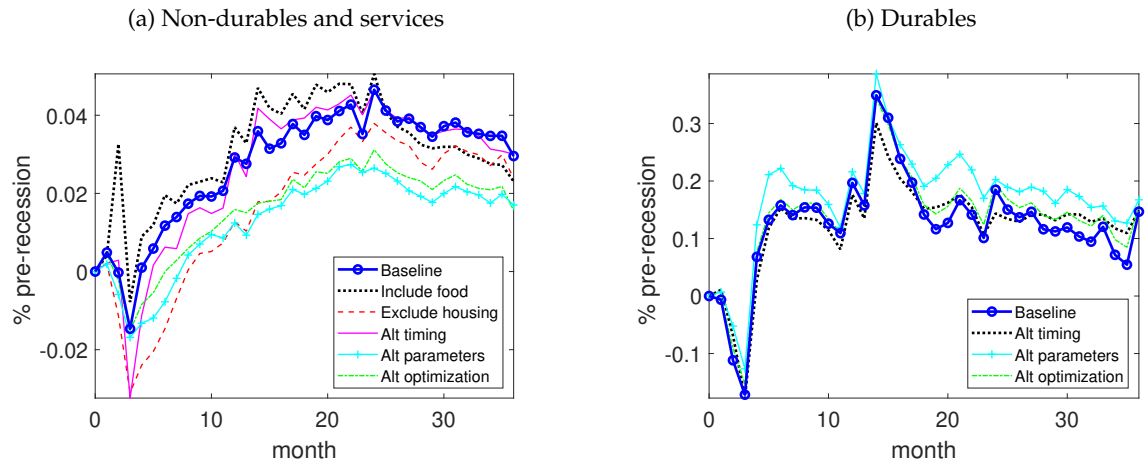
Notes: Data from NIPA. Robust versus excluded consumption in all recessions. All (black), robust (blue dashed), and excluded (red) dot-dashed.

D.4 Robustness and alternative partitions of robust consumption

In this section, I present the robustness of the empirical measure of robust consumption to alternative assumptions and partition methodologies. Overall, the baseline measure is fairly robust to alternative assumptions. In particular, I evaluate possible confounding factors and bias in the measure of robust consumption. One concern is that information was revealed about the economic effect of the Covid recession in the same period as the largest shock. I argue the threat is not severe. Firstly, the monthly frequency at which I assess the peak consumption changes means that it is likely that important information about the severity of the recession was revealed outside of this window. Secondly, the change in consumption is measured relative to the average standard deviation of a given consumption type, ensuring volatile sub-components remain in the robust basket.

Alternative assumptions: Figure A.6 presents a series of alternative specifications. The headline result is that for most series, the decline and dynamics look similar, and where they are not, the baseline seems more plausible. Panel a shows the effect of including food, excluding housing, using the quarterly change in consumption, using alternative hyperparameter, or minimizing a set of continuous convex weights (rather than allocating to a basket). Including food leads to a precautionary jump, the alternative timing leads to a larger transitory fall, while excluding the service flow of housing leads to a more persistent decline. The alternative parameterization of the hyperparameters: less restrictive on the decline during the recession and a lower penalty for deviating from pre-Covid business cycle ($\gamma = 2.6,\ \Gamma = 3.9$) and alternative optimization step leads to very similar dynamics, with slightly weaker recoveries. Panel b shows durables. The series all strongly co-move. I also tried equal weighting ($\Gamma = 1$) and the qualitative results were very similar, but the aggregate series capture pre-recession business cycle dynamics less well. Alternative partition definitions–including partitions based on the raw size of the decline rather than the normalized change, or partitions applied to residuals from regressions of consumption growth on income–yield similar basket compositions.

Figure A.6: Alternative robust specification



Notes: Data from NIPA. Panel a is path of robust consumption during Covid recession under alternative data: Baseline (blue), include food consumption (black), exclude housing (red), alternative timing, change over one quarter (magenta), alternative hyperparameters ($\gamma = 2.6,\ \Gamma = 3.9$) (cyan), alternative solution: convex minimization of weights ω_j (cyan). Panel b is durables consumption.

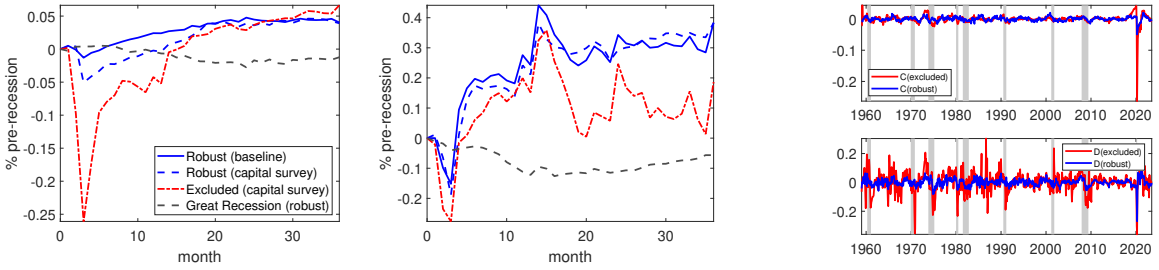
Sector disruption: Figure A.7 takes a different approach, grouping consumption goods by sectors based on responses to the BEA's capital expenditure survey. This survey recorded the share of the sector's respondents that reported "There was no effect on normal level of operations" and planned expenditures in 2020 due to Covid. The idea is to investigate a more direct measure of Covid disruption. After mapping the sectors to consumption types, I partition non-durables based on the response of the 39th percentile of consumption types ordered by disruption severity (equivalent to the share allocated to robust consumption). For durables, I allocate types exceeding the median type respondent.[79] This is a less data-driven approach and suggests a larger decline in robust consumption during the recession, but scale is not dramatically different. The peak decline is 5 percent and the recovery dynamics are similar. For durables, the series follows very closely with the preferred robust durables measure. Importantly, for both, the time series

[79] Durables production sectors reported larger disruptions, so a lower bound is necessary to ensure robust durables is not empty.

behavior at a business cycle frequency prior to the recession shows very different dynamics. This suggests the bundles partitioned by the severity to the sectoral shock are not a good measure of how consumption would have responded to economic shocks during the recession, excluding the effect of Covid.

Figure A.7: Alternative robust specification

(a) Non-durables and services (b) Durables (c) Non-durables and services



Notes: Data from NIPA. Partition based on sectors hardest hit by pandemic based on reported impact on capital expenditures in 2020. Panel a is non-durable and services. Panel b. Durables. Robust (blue dashed), excluded (red), Baseline robust (blue), Great Recession comparator (grey). Panel c is HP filtered series.

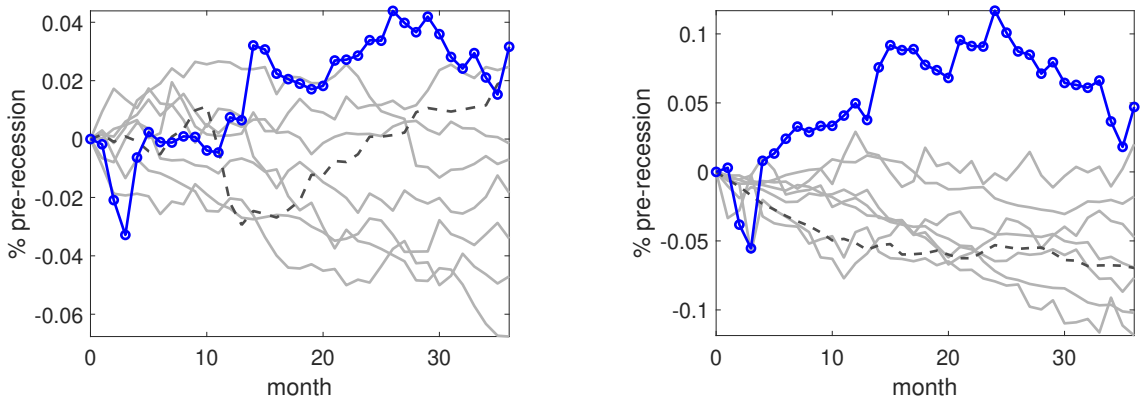
D.5 Robust consumption prices

In addition to partitioning consumption quantities, I also apply the same weights to the price of the baskets. I use the implied price deflators for each type, weighted by the consumption in each period. The prices are then computed relative to the price of the robust consumption good. The results are shown in Figure A.8. Panel a shows the behavior of the price of excluded consumption goods. At the onset of the recession, there is a sharp decline in the relative price of these goods, with a peak-to-trough fall of around 5 percent. This likely reflects the discounting of existing stocks of goods and services. Panel b shows the relative price movements for the robust durables composite. The initial price decline is smaller, after which prices begin to rise sharply. After 15-months the price of this basket of goods is 10 percent higher than it was at the beginning of the recession. The price rise likely reflects supply shortages.

One notable feature of both panels is that the patterns in the previous recessions are very heterogeneous, in contrast to those for the quantities. This again indicates that while the goods allocated to the baskets are specific to the dynamics of the Covid recession, based on the relative prices, there is nothing particularly specific to them in all recessions. Thus, the robust consumption basket should provide a good summary of unrestricted behavior in response to economic shocks.

Figure A.8: Relative price changes

(a) Price: excluded consumption (b) Price: Robust durables



Notes: Notes: Recession deviations of relative prices of excluded consumption and robust durables consumption, based on price indexes. Data from NIPA

D.6 Robust consumption in the micro data

D.6.1 Technical details of data construction

1. I start by taking the allocated sub-categories and weights ω_j from the baskets estimated in the NIPA data. These define the *robust* and *excluded* partitions.

2. All items are deflated by the NIPA product specific price index. I use both the Interview and Diary surveys

3. For each product type in the NIPA data – 181 services and non-durables, 43 durables – I identify the equivalent lowest set of categories in the CEX that maps to the NIPA sub-categories.

4. The crosswalk comes from the BLS UCC to NIPA product categories mapping. For most sub-categories it is possible to find an exact mapping.
   See: `https://www.bls.gov/cex/pce-ce-distributions.htm`

5. Where it is not possible to match the baskets exactly, as the CEX has a high aggregation (for example, food at home). I assign a proportion based on the average consumption share in NIPA.

6. Implied homeowner rents are subject to missing observations. For occupier-owned housing consumption, I compute an implied rent for all households based on housing characteristics. Where imputed rent is missing, I use this value. Otherwise, I take an average of BLS imputed home-owner rent and my imputed rent measure. This is closer to the share of expenditure reported in NIPA. I also report measures excluding homeownership.

7. Baskets are aggregated at the household level. Using the allocation and weights ω_j at a monthly frequency, based on the month of expenditure. A durables adjustment is an expenditure of more than $200 in 2019 dollars.

8. Aggregate series are computed by household head education level separately for the interview and diary series. The aggregate series are seasonally adjusted and smoothed with a moving average at a quarterly frequency.

9. I combine the aggregate interview and diary series. For non-durables and services consumption, I use a weight of 1 to 4 for the interview to diary series, reflecting that the diary survey only captures one week of spending for a reference month. For durables, I use equal weights given that durables are lumpy purchases.

10. I compute changes relative to aggregate consumption in December 2019.

D.6.2 Further results

Table A.3: Consumption partition summary statistics (Interview survey)

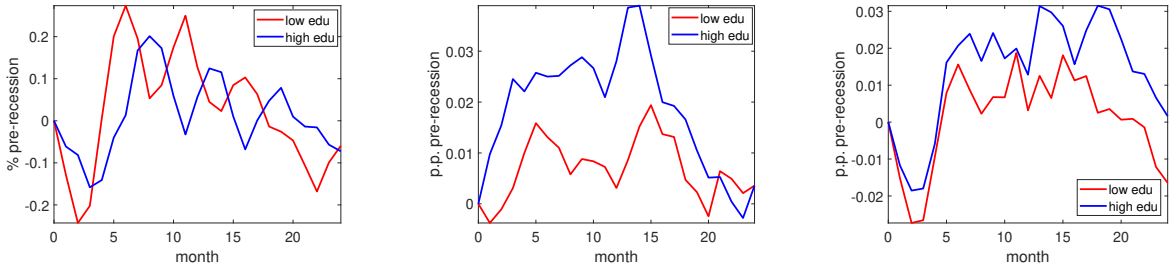
Moment	Low education	High education
Share of population	0.58	0.42
Average income	$48,463	$88,140
Robust consumption expenditure	$4,112	$6,452
Robust consumption share	0.48	0.49
Robust durables expenditure	$764	$1,070
Robust durables share	0.38	0.40

Notes: Data from Consumer Expenditure Survey (CEX). Consumption basket base only on interview survey. Income is average annual income from all source after tax. All spending on consumption goods is reported at a quarterly frequency. All values in 2019 prices.

Figure A.9 presents additional results from the construction of robust baskets in the CEX microdata. Panel a presents the excluded durables consumption basket. Overall, the pattern is more similar to the robust basket than for non-durables. Low-education households also cut excluded durables consumption more sharply. Panel b and c show the probability of adjustment. For robust durables, there is a clear gap, with

Figure A.9: Durables consumption by education level, additional results

(a) Excluded durables (b) Robust durables % adj. (c) Excluded durables % adj.



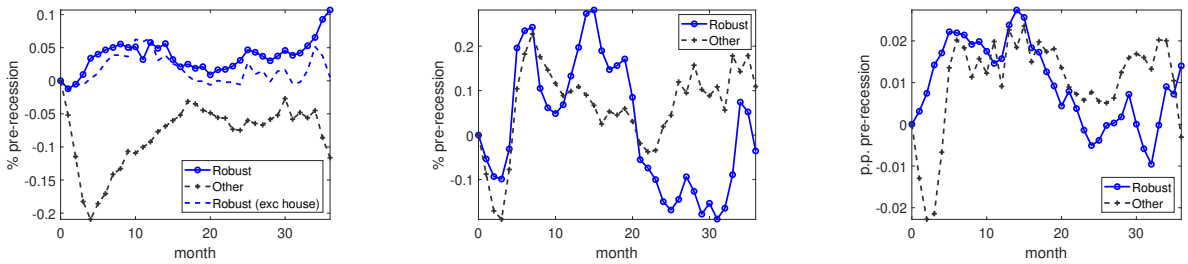
Notes: Data from Consumer Expenditure Survey (CEX). Consumption measure composite from interview and diary survey. High education household (blue), low education households (red).

high education households adjusting significantly more than low education households, particularly in 2021. For the excluded durables type, the response is more similar.

Figure A.10 shows the implied aggregates for consumption, durables and % adjusting durables. Comparing the patterns to Figure IV in the main text, it can be seen that reassuringly the aggregates constructed in the micro data line up closely with the dynamics and magnitude of the decline in the NIPA data. Interestingly, there is a large increase in the share of households adjusting their durables stock as the recession progresses. This serves as an additional time series moment to compare the model predictions to.

Figure A.10: Consumer Expenditure Survey aggregates

(a) Consumption (b) Durables (c) % Adjust



Notes: Data from Consumer Expenditure Survey (CEX). Aggregates based on household level data. Robust (blue), Other (black dashed), Robust excluding housing (blue dashed).

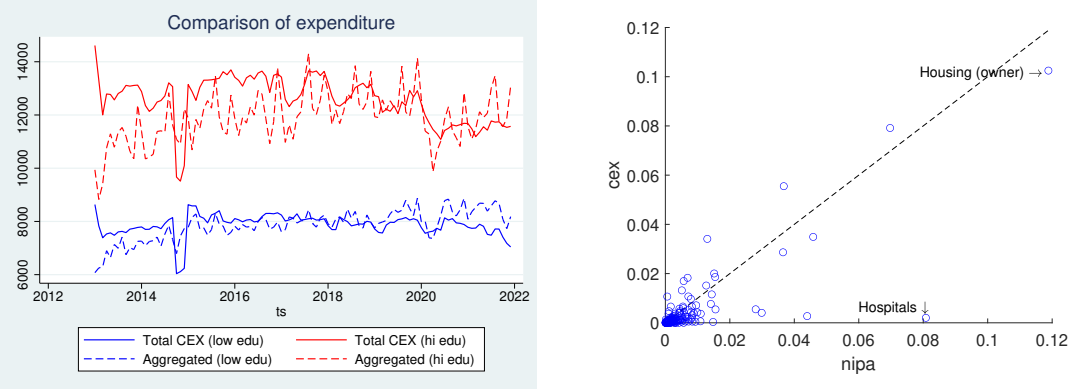
Figure A.11 provides additional checks of the sub-components of CEX micro data. Panel a shows the aggregated sub-components, selected to be consistent with the NIPA sub-component definition, line up closely with the total expenditure measure for both education groups. Panel b compares basket shares of the sub-components in NIPA and the CEX, and it can be seen that overall, the basket shares are fairly consistent. One exception is health expenditure, which is known to be under-reported in the CEX.

Figure A.12 split the CEX sample by states where Covid restrictions were above or below the median. A concern is that differences in education really reflect exposure to covid-related restrictions. Instead, the consumption series actually follow each other fairly closely with some evidence of an early easing of restriction (and higher consumption) in lower restriction states. The differences by education are substantially larger.

Figure A.13 plots Engel curves for the robust and excluded baskets. While there is some evidence of non-homothetic preferences, the difference between the robust and excluded series is not that significant emphasizing that robust consumption is a good proxy for consumption behavior. The switch from non-durables to durables is more evident.

Figure A.11: CEX basket consistency

(b) Basket shares (interview survey)

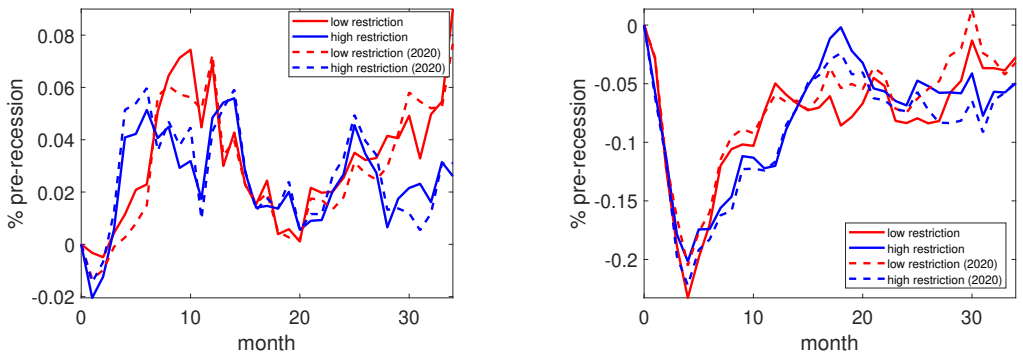


Notes: Panel a data from Consumer Expenditure Survey (CEX). Comparison of total expenditure and aggregate baskets. High education (red), Low education (blue). Panel b data from NIPA and CEX. Share of sub-components in aggregate basket and average share in household baskets.

Figure A.12: Split by Covid restrictions

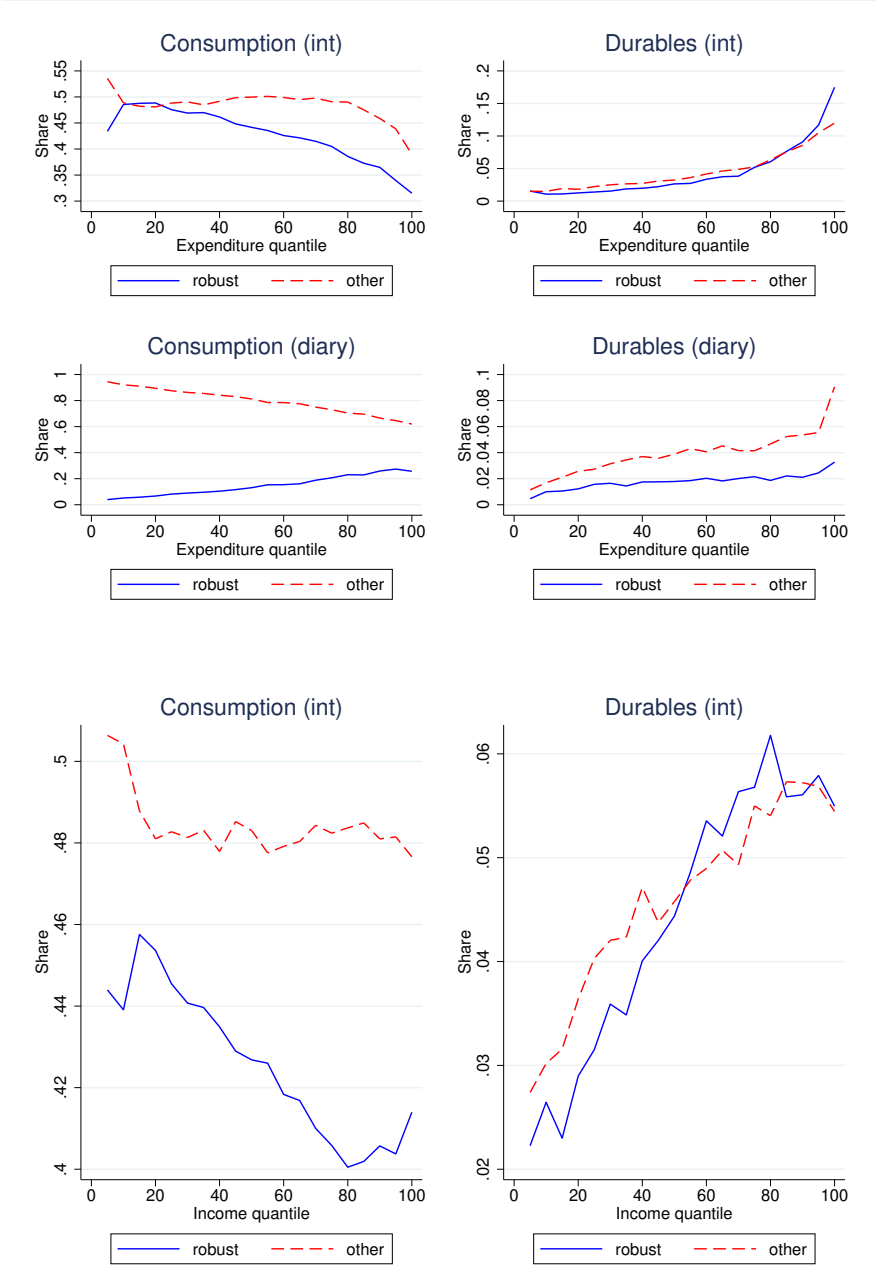
(a) Robust consumption

(b) Other consumption



Notes: data from Consumer Expenditure Survey (CEX). Split by severity of restriction in 2020 or 2020-22. Low restrictions (red), high restrictions (blue).

Figure A.13: Engel curves



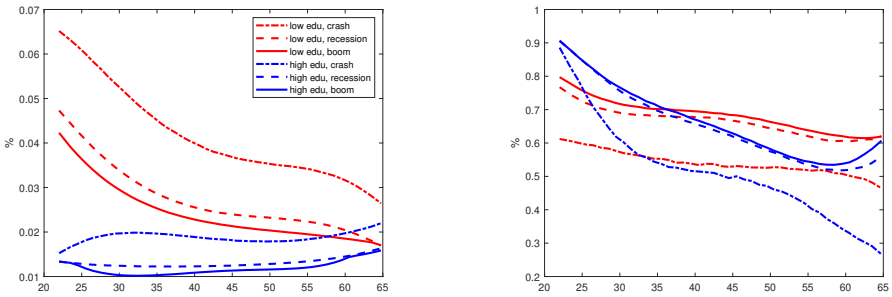
Notes: data from Consumer Expenditure Survey (CEX). Panel a: Basket shares by total expenditure diary and interview survey by total expenditure. Panel b: Basket shares in interview by income. Robust (blue), Excluded (red).

E Additional model material

E.1 Household life-cycle

Figure A.14: Job separation and job finding over lifecycle

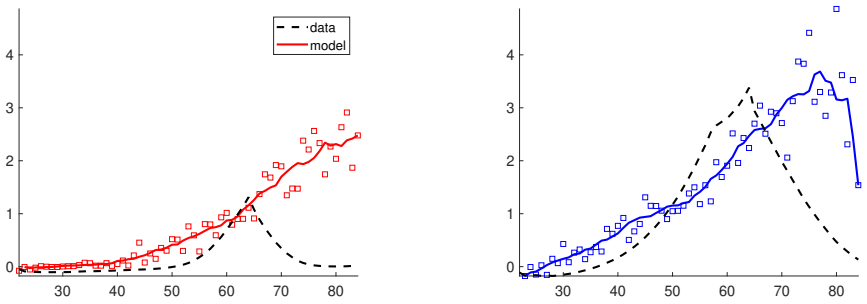
(a) Job separation (b) Job finding



Notes: Panel a and b. Data from CPS job separation and job finding rates along life cycle by age and state of the aggregate economy.

Figure A.15: Life-cycle targets

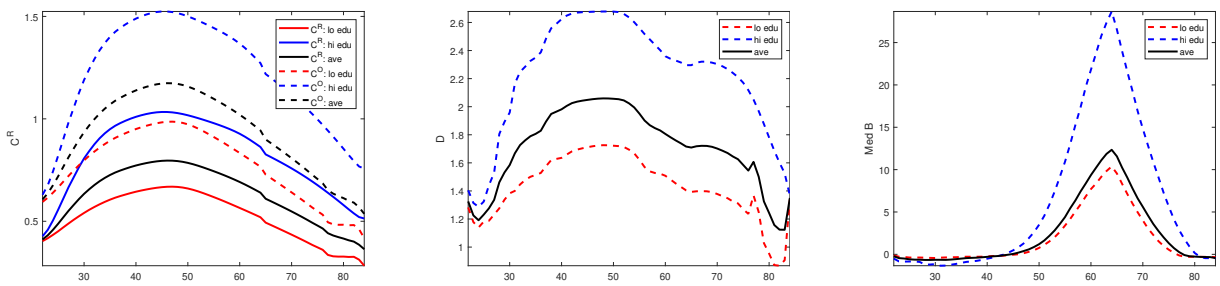
(a) Wealth:inc, low edu (b) Wealth:inc, hi edu



Notes: Median household net wealth (excluding pensions) to average income from 2019 Survey of Consumer Finance (SCF), by education. Panel a low education, Panel b high education. Data is estimated at each age, line is smoothed local average.

Figure A.16: Life-cycle profiles

(a) Consumption (b) Durables (c) Median assets



Notes: Life cycle plots for consumption, durables and asset from model. High education (blue), low education (red).

Figure A.16 displays lifecycle profiles for consumption goods, the durables stock and assets. Given the specification of preferences and fixed prices, robust and excluded consumption follow similar paths over the lifecycle. The path of robust consumption is lower, reflecting the smaller expenditure share. The model also

predicts that high education household both consume more and have steeper consumption profiles at the beginning of life. A similar pattern is seen for the durables stock in Panel b. Low education households tend to economize more on durables later in life, partly reflecting the weaker bequest motive. Finally, Panel c plots median asset holdings. As is common in these models, these are mainly held to smooth consumption during retirement. High education households accrue significantly more savings than low education households, in keeping with the targeted data moment. At the median, both household types are fairly constrained during the first 20 years of life.

E.2 Solution method

E.2.1 Household problem and policy functions

The model is solved with the Generalized EGM (GEGM) method (Fella, 2014; Ishakov, 2017) solved in the space of the consumption aggregator C.

A simplified version of the household problem is:

$$V(a,d,y) = \max_{c^r,c^b,d'} u(C(c^r,c^b),d') + \beta \mathbb{E}V(a',d',y')$$

subject to:

$$\begin{aligned} c^r + p^c c^b + p^d d' &= (1+r)a + (1-\delta)p^d d + y \\ a' &\geq -\Theta d' \end{aligned}$$

The solution satisfies the following first order conditions:

$$\begin{aligned} -u_c(C(\cdot),d')C_{c^r}(\cdot) + \beta \mathbb{E}V_a(a',d',y') &= 0 \\ -u_c(C(\cdot),d')C_{c^r}(\cdot)p^c + u_c(C(\cdot),d')C_{c^b}(\cdot) &= 0 \\ -u_c(C(\cdot),d')C_{c^r}(\cdot)p^d + u_d(C(\cdot),d') + \beta \mathbb{E}V_d(a',d',y') &= 0 \end{aligned}$$

And the envelope conditions:

$$\begin{aligned} V_a(a,d,y) &= u_c(C(\cdot),d')C_{c^r}(\cdot)(1+r) \\ V_d(a,d,y) &= u_c(C(\cdot),d')C_{c^r}(\cdot)(1-\delta)p^d \end{aligned}$$

Assuming Cobb-Douglas between the consumption aggregator and durables and a CES function between robust and excluded consumption, the sub-utility function is:

$$C(c^r,c^b;\chi) = \left[(\theta_c+\chi)^{\frac{1}{\epsilon}}(c^r)^{\frac{\epsilon-1}{\epsilon}} + (1-\theta_c-\chi)^{\frac{1}{\epsilon}}(c^b)^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}}$$

The consumption-savings first-order condition is:

$$\mathcal{X}(\theta+\hat{\chi})C(\cdot)^{(\theta+\hat{\chi})(1-\rho)-1}(d')^{(1-\theta-\hat{\chi})(1-\rho)} \times (\theta_c+\chi)^{\frac{1}{\epsilon}}(c^r)^{\frac{\epsilon-1}{\epsilon}-1}\left[(\theta_c+\chi)^{\frac{1}{\epsilon}}(c^r)^{\frac{\epsilon-1}{\epsilon}} + (1-\theta_c-\chi)^{\frac{1}{\epsilon}}(c^b)^{\frac{\epsilon-1}{\epsilon}}\right]^{\frac{\epsilon}{\epsilon-1}-1} = \tilde{V}_a$$

While inverting the Euler equation gives solves for the consumption aggregator

$$\mathcal{X}(\theta+\hat{\chi})C(\cdot)^{(\theta+\hat{\chi})(1-\rho)-1}(d')^{(1-\theta-\hat{\chi})(1-\rho)}(\theta_c+\chi)^{\frac{1}{\epsilon}}\left[(\theta_c+\chi)^{\frac{1}{\epsilon}} + \left(\frac{1}{p^c}\right)^{\epsilon-1}\frac{1-\theta_c-\chi}{(\theta_c+\chi)^{\frac{\epsilon-1}{\epsilon}}}\right]^{\frac{\epsilon}{\epsilon-1}-1} = \tilde{V}_a(a',d',y)$$

$$C^*(\cdot) = \left[\frac{\tilde{V}_a(a',d',y)}{\mathcal{X}(\theta+\hat{\chi})(d')^{(1-\theta-\hat{\chi})(1-\rho)}(\theta_c+\chi)^{\frac{1}{\epsilon}}\left[(\theta_c+\chi)^{\frac{1}{\epsilon}} + \left(\frac{1}{p^c}\right)^{\epsilon-1}\frac{1-\theta_c-\chi}{(\theta_c+\chi)^{\frac{\epsilon-1}{\epsilon}}}\right]^{\frac{\epsilon}{\epsilon-1}-1}}\right]^{\frac{1}{(\theta+\hat{\chi})(1-\rho)-1}}$$

Total consumption expenditures are:

$$C^x(C^*) = C^*(\cdot)\left[1 + \left(\frac{1}{p^c}\right)^{\epsilon-1}\frac{1-\theta_c-\chi}{\theta_c+\chi}\right]\left[\theta_c^{\frac{1}{\epsilon}} + \left(\frac{1}{p^c}\right)^{\epsilon-1}\frac{1-\theta_c-\chi}{(\theta_c+\chi)^{\frac{\epsilon-1}{\epsilon}}}\right]^{\frac{\epsilon}{1-\epsilon}}$$

Conditional cash on hand can be solved as:

$$X(d, d', a, y) = C^x(C^*) + a'$$

Robust consumption is then:

$$c^{r*}(C^*; \chi) = C^* \left[(\theta_c + \chi)^{\frac{1}{\epsilon}} + \left(\frac{1}{p^c}\right)^{\epsilon-1} \frac{1 - \theta_c - \chi}{(\theta_c + \chi)^{\frac{\epsilon-1}{\epsilon}}} \right]^{\frac{\epsilon}{1-\epsilon}}$$

Excluded consumption is

$$c^b(c^r) = \left(\frac{1}{p^c}\right)^{\epsilon} \frac{1 - \theta_c - \chi}{\theta_c + \chi} c^r$$

E.2.2 Algorithm details

- The model is solved sequentially back from the terminal age J.
- The GEGM step is applied each period, the only difference is that we solved for the consumption aggregator $(C^*; \chi)$ and have to adjust the total expenditures in the budget constraint $C^x(C^*)$ to find cash on hand.
- As in GEGM, each period we check the upper envelope condition to rule out solutions to the f.o.c. which are not solutions to the maximum problem.
- We then take the max over $d' \in D$ choices to find the unconditional policy functions.
- The adjusted envelop condition for assets gives next periods $\tilde{V}_a(a', d', y)$

E.2.3 Distribution

- The idiosyncratic state space for the problem is assets, a, durables, d, education, e, persistent income, z, unemployment (when working), ι, age, j.
- The aggregate state space is productivity Z and relative price p.
- This gives distributions on the working age $\lambda(a, d, e, z, \iota, j; Z, p)$ and retired $\lambda^R(a, d, e, z, j; p)$.
- We solve the distributions forward given household policy functions, for a sequence of Z and p.
- The dimensions are $N_a = 800$, $N_d = 11$, $N_e = 2$, $N_y = 5$, $N_\iota = 2, N_{J^w} = 43$, $N_{J^r} = 20$ and $N_Z = 3, N_p = 3$.
- While we solve for the household policies on a quarterly basis, we only evaluate the distribution once at each age. Households age stochastically, reaching $j + 1$ with probability=0.25. We use the policy function in the first quarter of each household age.
- **Initialized distribution:** The initialized distribution is found by solving for J periods with $Z = Z^{boom}$ and $p = 1$.
- **Stochastic simulation:** To calibrate the model, we simulate it for 1,000 years, allowing Z and p to be drawn from their Markov processes consistent with household expectations.

E.3 Recession shocks

E.3.1 Labor market outcome

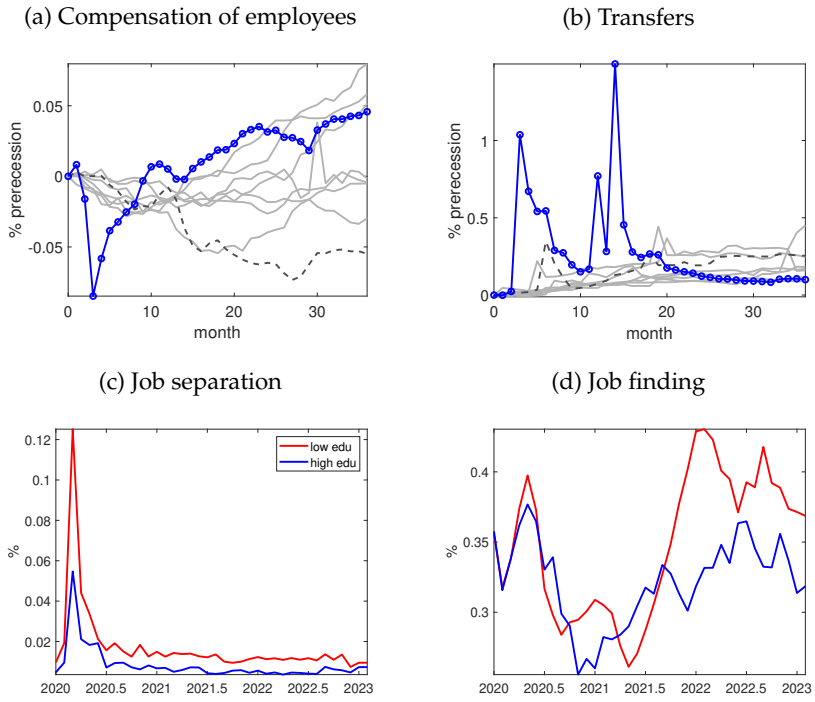
The top two panels of Figure A.17 show income data from NIPA. Panel a shows a large real fall in the compensation of employees during the Covid recession.[80] From peak-to-trough this declined by around 7.5 percent. Panel b shows government transfers to households. Two major spikes are evident in April 2020 and

[80] Deflated by the Personal Consumption Expenditures (PCE) price index.

April 2021. It is striking how large these transfers were relative to previous recessions. Transfers include extended Unemployment Insurance benefits and stimulus checks.[81]. As a result, household incomes were actually fairly well protected, while the checks provided a major source of fiscal stimulus.

The lower half of the panel shows labor market data from the Current Population Survey (CPS). In Panel c, I plot the job separation rate by low and high education group.[82] Separation rates increased substantially for both education groups, but peaked higher for the low education types (12 percent). These increases were short-lived and separation rates had returned to pre-recession levels by the end of 2020. Job finding rates fell initially before recovering as many workers were able to return to employment, following the initial separation shocks (Panel d). As the recession progressed, the job finding rate fell again, at a similar magnitude for high and low-education workers. From late-2021 onwards, there is evidence of a strong recovery in the job finding rate of low education workers.

Figure A.17: Monthly labor market



Notes: Data in Panels a and b from NIPA accounts. Transfers is Personal current transfer receipts. Data in Panels c and d from CPS.

E.3.2 Shocks construction

Productivity shocks: Productivity shocks, $\{Z_t^e\}_{t=1}^T$, are used to match the change in wages by education type. These are computed in the CPS, using weekly earnings deflated by the CPI and seasonally adjusted with monthly dummies. To account for selection effects which lead to an increase in measured wages in this period, I compute occupation×industry×education×gender wages and hold fixed the cell shares prior to the recession. I detrend with the growth rate in the previous 4 years. Additionally, I constrain the resulting implied aggregate series (after multiplying by model implied unemployment rates) to be less than the negative income deviations of compensation of employees in NIPA, scaling the education level deviation to match this aggregate. Finally, in the model, to ensure any rise in uncertainty is mean preserving, I adjust productivity to offset the aggregate effect of a change in $\sigma^2_{z^e,e,t}$. Retired households' pensions are assumed to be fixed nominally and deflated at the equivalent CPI price level.

Labor market shocks: The job finding, $\{\lambda_t^{0,e}\}_{t=1}^T$, and separation rate, $\{\lambda_t^{1,e}\}_{t=1}^T$, shocks are calculated from the CPS (See A.17 Panel c and d). Shocks differ by education level and are computed as deviations from the rate in 2019. Labor market shocks are assumed to have a uniform (percentage point) impact across the age-varying life-cycle transition profiles.

[81] Increased UI benefits reflect both increased uptake and generosity

[82] Low education is defined as no or incomplete college. High education is completed college. Job separations are computed as households that report being unemployed in the month following a month they reported being employed. Job finding is employment following unemployment.

Price shocks: Price shocks, $\{p_t^{\epsilon}\}_{t=1}^{T}$, are fed in as deviations in the relative price of the excluded consumption good relative to the robust good (see Figure A.8).

Preference shocks: Preference shocks, $\{\chi_t\}_{t=1}^{T}$, change the consumption weights in the utility function. Given this specification, χ_t, is given by the intra-temporal consumption choice, which can be rearranged as:

$$\chi_t = \frac{1 - \theta_c\left(1 + p_t^{\epsilon}\frac{\mathbf{C_t^O}}{\mathbf{C_t^R}}\right)}{1 + p_t^{\epsilon}\frac{\mathbf{C_t^O}}{\mathbf{C_t^R}}} \tag{21}$$

I solve for the value of χ_t each period that delivers the deviations in robust consumption relative to excluded consumption, given the deviation in the relative price.

Government transfers: Government transfers, $\{\mathcal{T}_t^{e}\}_{t=1}^{T}$, capture the three waves of fiscal stimulus provided during the Covid recession. Based on the average family size, this equates to a per household transfer of \$3,400 in March 2020, \$2,400 in December 2020, and \$5,600 in March 2021. The transfers had different payments for adults and children. I use the average family size by education level to scale the transfers to an age-education specific transfers. Transfer receipt is subject to a maximum income threshold of \$75,000 per adult household.

Capturing serial correlation in receipts at the aggregate level (see Figure A.17 Panel b), households receive 70 percent of the transfer value in the quarter of impact and the rest in the following quarter. In both static and dynamic cases, I assume households anticipate the transfers (which arrive in Q2) at the start of the recession. With static expectations, the second round of transfers, which occurs in 2021.Q1 are also anticipated at the beginning of the recession. With dynamic expectations, households learn about the second round of transfers with one quarter's advance notice. Households anticipate the transfers will be funded with lump-sum taxes levied in the last 4 quarters of the impulse response horizon. The cost is cumulative of interest at the real rate. Part of the cost will be met by households yet to enter the economy.

Real interest rate: I feed in deviations in the difference between the Fed Funds rate and inflation. I assume households base their expectation of inflation in month t based on the inflation rate over the past 12 months. Implicitly, I am assuming the household did not correctly predict the increase in inflation experienced, which given the large unexpected increase, seems appropriate. On a monthly basis the real rate is computed as $r_t = FF_t - (cpi_t - cpi_{t-12})$, where cpi_t is the log of the price level. Monthly real rates are averaged to get the quarterly value, and I compute deviations from the real rate in 2019.Q4.

As the increase in inflation was not anticipated by households at the beginning of the Covid recession, I restrict the implementation of this shock to the case of dynamic expectations. Inflation also reduces real wages and the real value of retirees' pensions. These are included through the productivity shocks and occur in both static and dynamic cases.

Morbidity: I compute Covid death shocks from the CDC Covid deaths tables. To calculate the implied increase in morbidity, I compute Covid deaths relative to total deaths minus Covid deaths per quarter by 7 broad age categories. This percentage deviation is multiplied by the steady state life-cycle death probability. Households perfectly perceive the death shocks under both information structures. Between 2020.Q1 and 2022.Q2, this averages a 12.6 percent increase in the mortality rate.

1980s shocks: The 1980s shocks follow the same methodology as the shocks for the Covid recession. In the 1980s, we assume no government transfers of morbidity shocks. The sample window captures two NBER-dated recessions (January-July 1980 and July 1981-November 1982), reflecting the Federal Reserve's stop-start approach to disinflation.

E.4 Static expectations: one-time increase in uncertainty

I now estimate household expectations under a static information setting. The unobserved objects are the education-specific increase in income uncertainty, $\sigma^2_{z^e,e,0}$, and the persistence of that uncertainty increase, $\varrho_{e,0}$, where subscript 0 denotes the initial quarter of the recession (2020:Q1). Under static expectations, households observe the complete future path of all shocks—productivity, labor market transitions, preferences, and transfers—in 2020:Q1. The only unknowns are the level and persistence of uncertainty households perceived at the recession's onset. The persistence parameter $\varrho_{e,0}$ governs how quickly the uncertainty increase dissipates. I model uncertainty as applying to idiosyncratic labor income risk, though it could capture broader economic uncertainty. This specification implies that uncertainty affects primarily working-age households who face labor income risk. I estimate $\Theta = \{\sigma^2_{z^e,e,0}, \varrho_{e,0}\}_{e\in\{\text{low, high}\}}$ by minimizing the distance between model-predicted and observed consumption paths.

E.4.1 Estimation exercise

The model is initialized at steady state, with households distributed over assets, income states, and employment status according to the ergodic distribution in the high-productivity state. The steady state corresponds to the ergodic distribution when households use stochastic model policy functions, but only the high-productivity state and steady-state relative price ($p = 1$) are realized. In 2020:Q1, households learn the complete paths of productivity, labor market transitions, the relative price of excluded consumption, and government transfers. In 2020:Q2, two shocks hit: the preference shock for excluded consumption and an unexpected, persistent increase in income uncertainty. I use 2020:Q2 for the uncertainty shock because this quarter exhibits the largest declines in both robust and excluded consumption, effectively the trough of the recession, giving the static specification the best chance of matching these targets.

The uncertainty shock increases the variance of idiosyncratic income shocks through the following mean-reverting process:

$$\sigma_{z,e,t} = \bar{\sigma}_{z,e} + \hat{\sigma}_{z,e,t}$$

$$\hat{\sigma}_{z,e,t} = \begin{cases} 0 & if\ t = 1 \\ \varepsilon^{\sigma}_{e} & if\ t = 2 \\ \varrho_e \hat{\sigma}_{z,e,t-1} & if\ t > 2 \end{cases}$$

where shocks to persistent income are drawn from $\epsilon^{e}_{a,t} \sim \mathcal{N}(\mu_{z,e,t}, \sigma^2_{z,e,t})$. The standard deviation jumps to $\bar{\sigma}_{z,e} + \varepsilon^{\sigma}_{e}$ in 2020:Q2, then decays toward its steady-state level $\bar{\sigma}_{z,e}$ with persistence ϱ_e. Households learn about the increased variance in 2020:Q2 before income is realized that quarter, so there are no unexpected income realizations in 2020:Q2 itself. The mean $\mu_{z,e,t}$ adjusts to keep the distribution mean-preserving, taking a negative value when uncertainty increases. Both the magnitude and persistence of the uncertainty increase are education-specific, yielding four parameters to estimate: $\Theta = \{\varepsilon^{\sigma}_{e}, \varrho_e\}_{e\in\{\text{low, high}\}}$.

I estimate Θ by minimizing the distance between model predictions and data for four consumption series: aggregate robust consumption, $\hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}$, the aggregate robust durables stock $\hat{\mathbf{D}}_{\mathbf{t}}$, and the education differences in these quantities $\Delta_e \hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}} = \hat{\mathbf{C}}^{\mathbf{R,e=L}}_{\mathbf{t}} - \hat{\mathbf{C}}^{\mathbf{R,e=H}}_{\mathbf{t}}$ and $\Delta_e \hat{\mathbf{D}}_{\mathbf{t}} = \hat{\mathbf{D}}^{\mathbf{e=L}}_{\mathbf{t}} - \hat{\mathbf{D}}^{\mathbf{e=H}}_{\mathbf{t}}$, from the micro data. Hat notation $\hat{\mathbf{X}}$ denotes deviations from steady state (model) or trend (data). I construct the durables stock in the data by forward simulation using observed durables consumption and the model's depreciation rate, which produces a smoother series than durables consumption itself. Formally, I minimize:

$$F(\Theta) = \mathbf{M}(\Theta) W \mathbf{M}(\Theta)'$$

where W is a weighting matrix that assigns linearly decreasing weights to observations farther in the future, and $\mathbf{M}(\Theta)$ stacks the four deviation vectors:

$$\mathbf{M}(\Theta) = \Bigg[\hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}{}_{data} - \hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}(\Theta)_{model}, \hat{\mathbf{D}}^{\mathbf{R}}_{\mathbf{t}}{}_{data} - \hat{\mathbf{D}}_{\mathbf{t}}(\Theta)_{model}$$

$$\Delta_e \hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}{}_{data} - \Delta_e \hat{\mathbf{C}}^{\mathbf{R}}_{\mathbf{t}}(\Theta)_{model}, \Delta_e \hat{\mathbf{D}}^{\mathbf{R}}_{\mathbf{t}}{}_{data} - \Delta_e \hat{\mathbf{D}}_{\mathbf{t}}(\Theta)_{model} \Bigg]$$

The first two moments come from aggregate NIPA consumption; the latter two come from education-specific CEX consumption.[83]

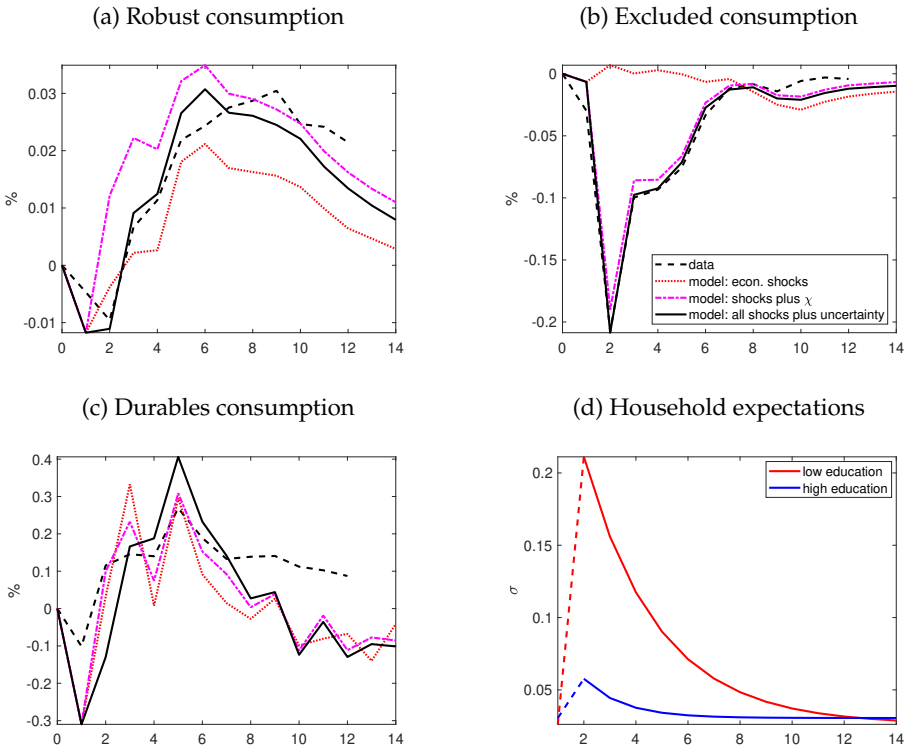
[83] The model features a single durable good (hence no robust superscript on D), while the data moments use the robust durables measure. A discussion of the identification from the education-specific differences is provided in Section E.6.

E.4.2 Static model results

Figure A.18 presents the static expectations model results, alongside two alternative specifications that isolate the preference and uncertainty shocks' contribution. The red line shows predictions using only the observed economic shocks (productivity, labor market transitions, and transfers) without the preference or uncertainty shocks. This specification matches robust consumption reasonably well (Panel a) but counterfactually predicts a boom in excluded consumption: the declining relative price of excluded goods induces substitution that overwhelms the income effect, generating rising consumption when the data show a collapse.

The magenta line adds the estimated preference shock, which shifts demand away from excluded consumption and towards robust consumption. This addition successfully captures the excluded consumption decline (Panel b) but overpredicts robust consumption, which rises more than observed. The black line shows the full static model, adding the estimated uncertainty shock. Heightened uncertainty induces precautionary saving and dampens the robust consumption response, particularly in the early part of the recession, substantially improving the fit. Despite these improvements, the static model exhibits two key failures. First, it overpredicts the initial robust consumption decline in 2020:Q1 (Panel a). Under static expectations, households immediately recognize the recession's full severity and cut consumption sharply. The data's more modest initial decline suggests households did not initially anticipate the size and persistence of the Covid economic shocks. Second, the model slightly overpredicts the robust consumption rebound during recovery. For durables (Panel c), the model tracks the overall trajectory after 2020:Q2 but again overpredicts the 2020:Q1 decline and underpredicts consumption during the later recovery. The sharp initial durables drop reflects households lowering the probability of adjusting their durables stock upon the initial realization of the economic shocks, again suggesting incomplete initial information. Appendix Figure A.23 shows that the model successfully replicates the large increase in household saving rates following the 2020-21 transfer payments.

Figure A.18: Consumption dynamics with a one-time increase in uncertainty



Notes: Response to Covid recession with static expectations. Data is dashed black line. Only economic shocks (red), Preference shock added (magenta), Additional uncertainty shock (solid black). Panel d shows estimated household expectations.

Panel d displays the estimated uncertainty paths. The implied uncertainty increases are substantial: more than doubling for high-education households and increasing by even more for low-education households. The estimates also indicate greater persistence for low-education households, consistent with their more precarious labor market position. These findings highlight the important role of economic uncertainty in the consumption dynamics. However, the systematic overprediction of the initial consumption responses, for both robust consumption and durables, indicates households did not possess complete information in 2020:Q1. A dynamic model of expectations may better capture the data.

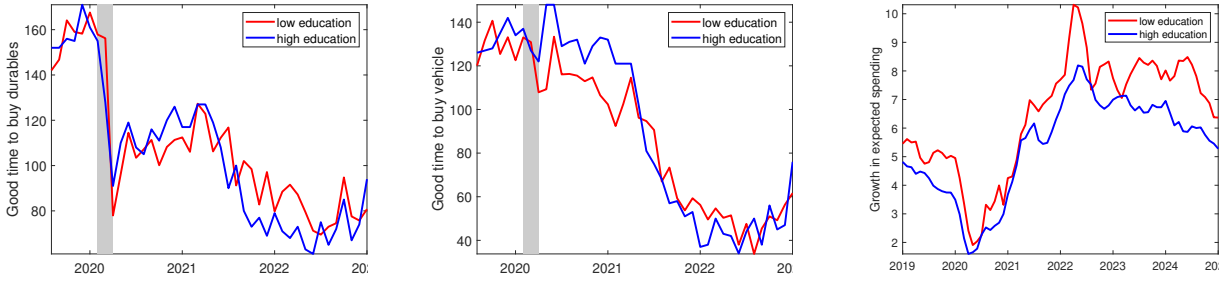
E.5 Additional results in household expectations surveys

Figure A.19 presents additional results from household expectation surveys. All measures strongly support the findings and beliefs of the main quantitative exercise. Panels a and b also use data from the Michigan survey. They show a measure of whether households think it is is a good time but durables. Panel a is for durables generally, Panel b is specifically vehicles. There is a decline in sentiment at the start of the recession, notably for durables. Further, just as in the CEX data the response is more negative for low education households, consistent with a greater increase in uncertainty.

Panel c plots a measure from the Federal Reserve Bank of New York's Survey of Consumer Expectations. The results report the percentage spending growth households expect during the next 12-months. Consistent with the evidence for robust consumption in the CEX, low education households have stronger expected spending growth in 2021 to 2023.

Figure A.19: Consumption and expectation in surveys

(a) Durables (Michigan) (b) Vehicles (Michigan) (c) Expected spending (FRBNY)



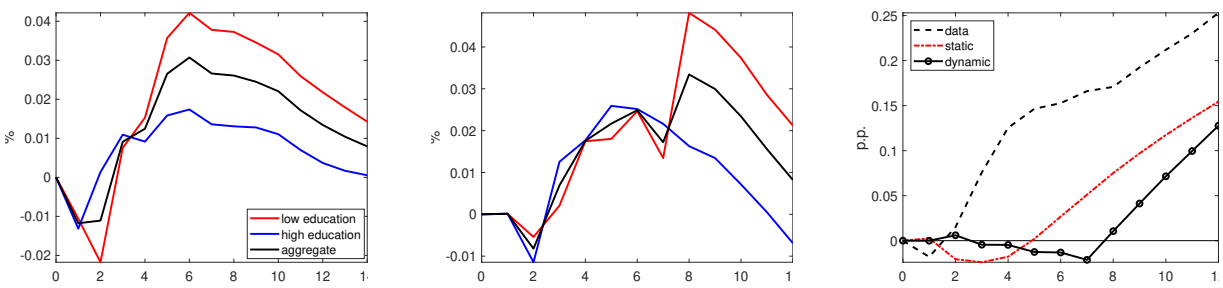
Notes: Panel a and b: are responses from the Michigan survey on whether it is a good time to but durable (Panel a) or vehicles (Panel b). Panel c is expected spending growth in the next 12 months in the FRBNY survey of consumer sentiments.

E.6 Analysis by education level

In the model, the identification of household expectations was disciplined by variation in the response by education level. An important feature of the CEX micro data is the more persistent negative response of robust consumption by the high education group and a larger initial decline in durables consumption for low education households. Figure A.20 plots the response by education group for robust consumption. Panel a shows the static model, while Panel b reports the results for the dynamic model. Panel a shows that, given the estimated expectations, the model is qualitatively able to generate these dynamics. Robust consumption by high education households is below the low education households from quarter 5 onward. Panel c shows the consumption difference estimated in the model compared with the data. The static model does a better job of matching the initially larger decline in robust consumption by the high education households. With dynamic expectations, this effect is delayed.

Figure A.20: Consumption and durables dynamics by education level

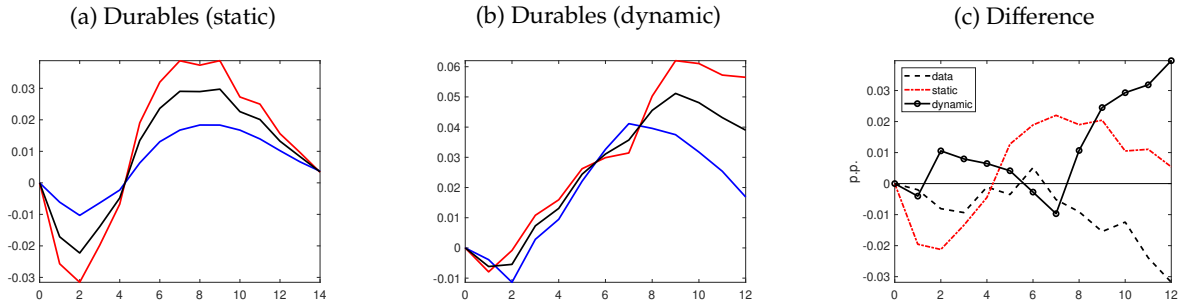
(a) Consumption (static) (b) Consumption (dynamic) (c) Difference



Notes: Panel a and c: consumption response in dynamic model by education group. High education (blue), low education (red). Panels b and d: targeted difference in consumption response by education. Data (dashed black), static model (red), dynamic model (solid black).

Figure A.21 shows the path of durables consumption. Durables of the low education households falls by more and is persistently below the high education group's expenditures. For durables consumption, the fit is better with dynamic expectations. The key mechanism is that high uncertainty facing low education households causes them to postpone durables adjustment, thus their durables consumption is below that of the high education group.

Figure A.21: Consumption and durables dynamics by education level



Notes: Panel a and c: consumption response in dynamic model by education group. High education (blue), low education (red). Panels b and d: targeted difference in consumption response by education. Data (dashed black), static model (red), dynamic model (solid black).

E.7 Model fit

Table A.4: Model fit

	Aggregate			Δ Education		
	$\mathbf{C^R}$	$\mathbf{C^O}$	**D**	$\mathbf{C^R}$	**D**	**Total**
Economic shocks only	0.047	0.351	0.107	0.266	0.091	0.683
+ taste	0.050	0.066	0.085	0.151	0.101	0.327
+ uncertainty	0.023	0.044	0.099	0.487	0.104	0.462
Dynamic	0.022	0.027	0.016	0.572	0.108	0.405

Notes: Sum of squared deviations between model series and data. Periods are weighed by inverse of time since the recession began t^{-1}. Gaps by education are weighted by 0.5 in the total objective.

I also directly assess model fit by comparing the implied aggregates of the model to the targeted data series. The results are presented in Table A.5. The first three columns show the fit to the aggregate consumption variables, the next two to the differences by education group, while the last column presents the total. The results are weighted consistent with the estimation strategy in the static expectations exercise. The first row shows the model fit when only the economic shocks are included. This has the largest deviation, driven by the poor fit of the excluded consumption good $\mathbf{C^O}$. Adding the taste shock improves the fit of the excluded consumption good (and the difference between robust consumption of low and high education groups), but worsens the fit of the robust consumption good. Adding the shock to uncertainty improves the fit of both consumption goods; it slightly worsens the fit of the durables stock and consumption difference by education, however. The bottom row shows the result for the dynamic fit. The overall fit of the dynamic model is much improved. In particular, the path of the durables stock is much better approximated. The only shortcoming is the difference between the robust consumption by education groups, which increases in the dynamic estimation. The model generates a positive gap in robust consumption between low and high education households, later relative to when it occurs in the data. The results without weighting are shown in Table A.5, the fit is even better as the model captures differential behavior by educational groups with a delay and this fit is better reflected in the unweighted model fit measures.

E.8 Decomposition of economic shocks

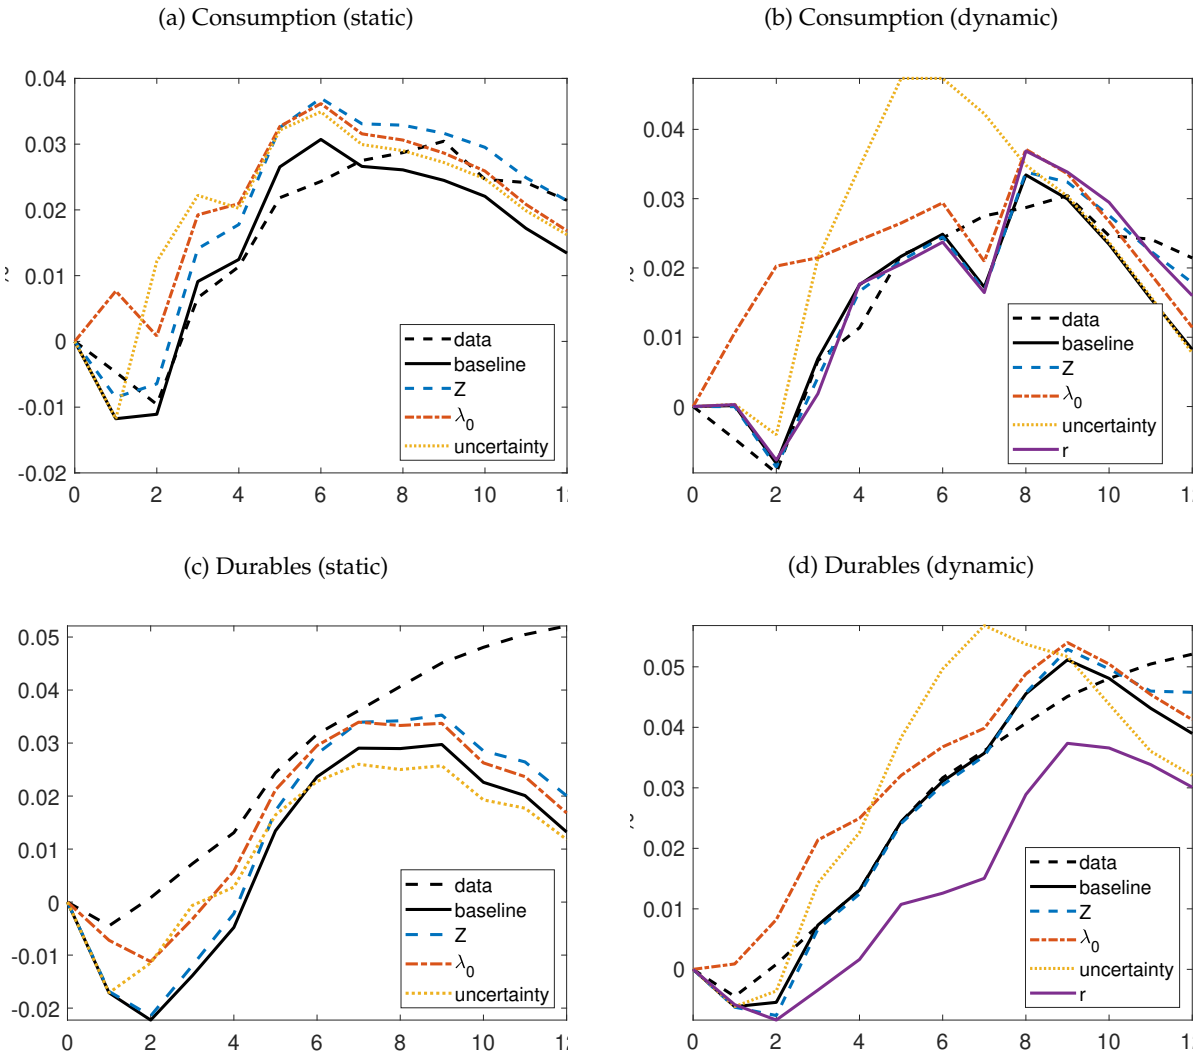
Next, I decompose the economic causes of the Covid recession. To implement this, I re-run the model holding expectations fixed, shutting down each shock, one at a time. I then compare robust consumption and the durables stock path to the responses in the baseline model. The time series for these experiments

Table A.5: Model fit (unweighted)

	Aggregate			Δ Education		
	C^R	C^O	D	C^R	D	Total
Economic shocks only	0.009	0.051	0.022	0.052	0.018	0.118
+ taste	0.007	0.010	0.017	0.027	0.020	0.058
+ uncertainty	0.004	0.007	0.019	0.097	0.021	0.089
Dynamic	0.004	0.006	0.003	0.121	0.025	0.086

Notes: Sum of squared deviations between model series and data. Gaps by education are weighted by 0.5 in the total objective.

Figure A.22: Decomposition of recession by economic shocks



Notes: Path of consumption excluding each shock. Data (black dash), baseline (black), Productivity [Z] (blue), job separation [λ_0] (red), uncertainty (yellow), interest rate [r] (purple).

is shown in Figure A.22. This is the counterfactual path of the economy if a particular shock has not been realized. In Table A.6, I calculate summary statistics for the importance of each shock. The top panel shows the sum of absolute deviations over the 12-quarter horizon of the model. This gives a measure of the average contribution of each shock. The bottom panel shows the sum of negative deviations of the baseline model.

This shows the contribution of the shocks that lowered consumption during the recession.[84] This provides information on the economic factors most important for the downturn.[85]

Table A.6: Decomposition of recession by economic shocks

	Shock						
	Z	λ_0	λ_1	price	rate	transfer	σ
Sum of deviations							
$\mathbf{C^R}$ (Static)	-7.5	-8.6	-2.0	2.1	0.0	26.6	-6.9
D (Static)	-4.6	-8.0	-2.1	3.5	0.0	11.4	-1.7
$\mathbf{C^R}$ (dynamic)	-1.9	-8.9	-1.4	3.1	-2.3	25.4	-11.6
D (dynamic)	-0.5	-7.8	-1.5	1.0	13.3	15.7	-7.3
Negative contribution							
$\mathbf{C^R}$ (Static)	-7.5	-8.6	-2.0	-1.0	0.0	0.0	-6.9
D (Static)	-4.6	-8.0	-2.2	-0.1	0.0	-1.2	-3.5
$\mathbf{C^R}$ (dynamic)	-2.4	-8.9	-1.4	-0.2	-2.9	0.0	-11.6
D (dynamic)	-1.1	-7.8	-1.5	-0.1	-0.1	-0.3	-8.6

Notes: Contribution of each shock to recession. Measure is deviation relative to baseline. Sum of deviations is total contribution, negative contribution is absolute value of shocks that lower consumption below the baseline. The numbers are multiplied by minus 1 such that a negative number indicates the addition of the shock lowers consumption. Productivity [Z], job separation [λ_0], job finding [λ_1], relative price, interest rate, government transfers, uncertainty [σ].

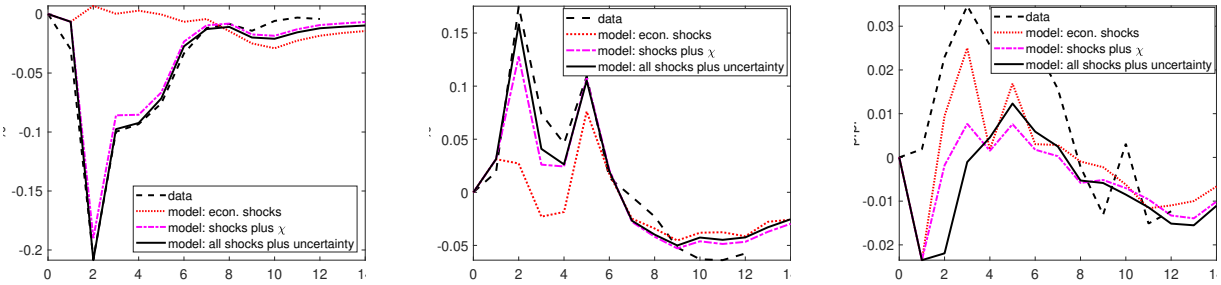
From Figure A.22, one can see that the most important contributor to the consumption decline was the increase in the job separation rate, λ_0. In the counterfactual economy, consumption is substantially higher if this shock is excluded. In total, it contributes a 9 percent cumulative decline in robust consumption and an 8 percent decline in the durables stock during this period (Table A.6). In comparison, the decline in the job finding rate had a far more modest effect. Another significant factor was falling productivity, although the importance of this differs across the specifications. In the static model, it plays a greater role. This is because a large share of the negative consumption effect in the static model is associated with a fall in real wages that occurred at the end of the window as inflation increased. In the dynamic model, this force is not anticipated at the beginning of the recession. As a consequence, the model implied dynamic expectations find a larger role for uncertainty, which is particularly important for lowering robust consumption. The transfers positively contribute to the consumption and durables. This is discussed further in Section 5.0.1.

[84] Mechanically, if a shock weakly lowers consumption in all periods, the sum of deviations equals the negative contribution.
[85] Negative contribution is calculated as: $\sum_t (baseline_t - counterfactual_t) \times 1[baseline_t < counterfactual_t]$, such that a negative number indicates a shock lowered consumption during the recession.

F Additional Figures

Figure A.23: Consumption dynamics in static model of expectations

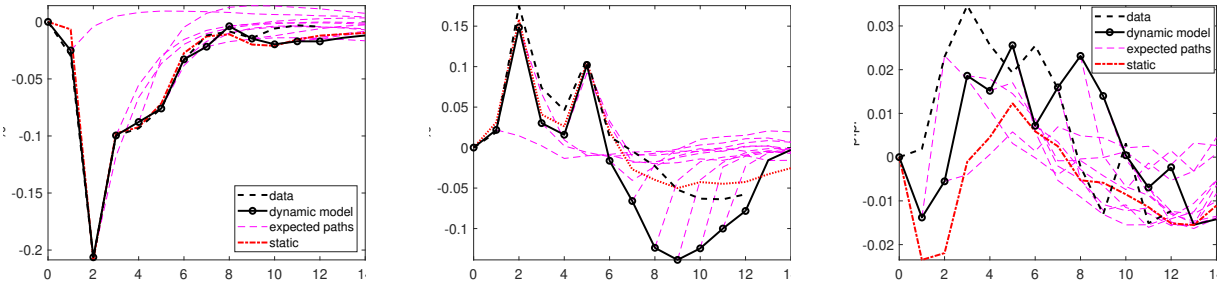
(a) Excluded consumption (b) Savings rate (c) Adjustment



Notes: Response to Covid recession with static expectations. Data is dashed black line.

Figure A.24: Consumption dynamics in a dynamic model of expectations

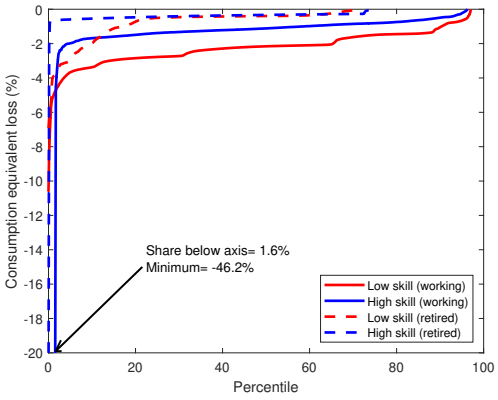
(a) Excluded consumption (b) Savings rate (c) Adjustment



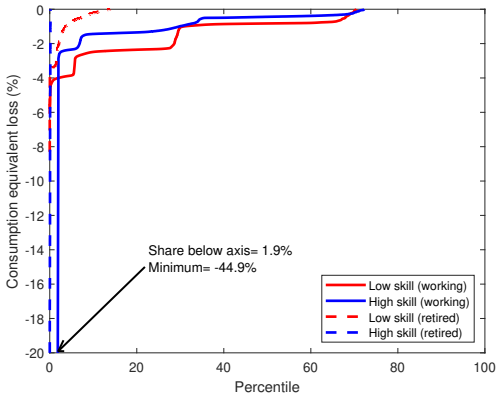
Notes: Response to Covid recession with dynamic expectations. Data is dashed black line. Static expectations (red); Dynamic expectations (solid black); Expected path in each period (magenta).

Figure A.25: Welfare across the distribution

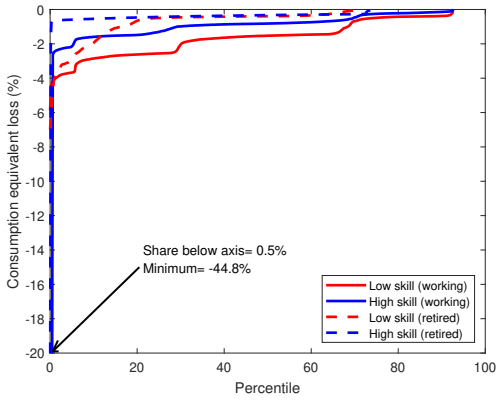
(a) Static



(b) Dynamic



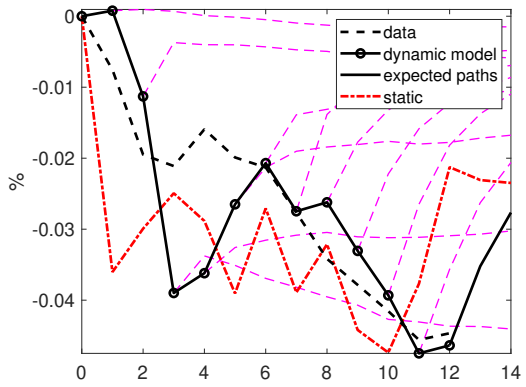
(c) Perfect foresight



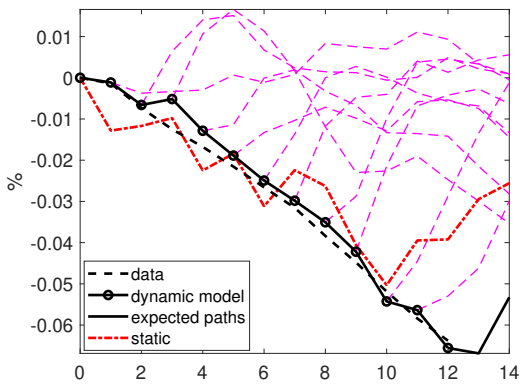
Notes: welfare loss in consumption equivalence units ordered by size of loss. High education (blue), low education (red), working (solid), retired (dashed). Panel a: static expectations, panel b: dynamic expectations, panel c: perfect foresight. Only shown for those with negative welfare.

Figure A.26: Consumption in the 1980s recession

(a) Robust consumption



(b) Durables stock



Notes: Response to 1980s recession with dynamic expectations. Data is dashed black line. Static expectations (red); Dynamic expectations (solid black); Expected path in each period (magenta).