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

Kieran Larkin*

Click here for latest version

16 February 2022

Abstract

This paper investigates what information can be recovered on households expectations from consumption data, during a period where the consumption basket changes acutely. We study consumption behavior during the Covid-19 recession, when household consumption declined significantly: in part due to Government restrictions and changing behavior. We propose a methodology for computing a measure of "robust" consumption that identifies a subset of goods and services that can still provide us with information on prevailing economic circumstances. Prior to the Covid-19 recession this measure has broadly similar business cycle dynamics to those in the restricted basket. We combine this measure of consumption with a sophisticated life-cycle model, featuring three consumption goods and a rich income process, and use it to estimate households expectations during the Covid-19 recession when subject to a range of shocks. In particular, we use the framework to understand the extent to which households experienced increased uncertainty and to estimate changing household perceptions of the persistence of the economic shock during the course of the recession.

Keywords: Heterogeneous Agents, Consumption, Covid-19, Uncertainty, Persistence

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

*Institute for International Economic Studies (IIES). Contact: kieranlarkin1@gmail.com. Institute for International Economic Studies. Stockholm University, Stockholm, SE-106 91, Sweden.

1 Introduction

The Covid-19 recession had a dramatic effect on the economy. In the second quarter of 2020, GDP declined on an annualized basis by 9.3 percent. This is by far the largest decline on record. Evidence of the extreme nature of the crisis is widespread. For example, the S&P500 fell by 30 percent between February and March 2020. The pandemic radically reorganized commerce and production. Restrictions of the ability of people to meet, attend shops and broad requirements to work from home reshaped households' lives in a way there were difficult to imagine, ex ante. Even in the absence of mandated restrictions there was a large scale voluntary re-ordering of the work and leisure activities households chose to perform. Most significantly was the huge loss of life caused by the pandemic.

The Covid-19 recession was also highly unusual from the perspective of the economics profession. In many past recessions the precise causes of the slow down are often vigorously debated, but in the Covid-19 recession the cause was almost self evident. Further, variables which have previously thought to be informative about about features of economic activity were obfuscate by the the wide ranging impact of lockdowns or changing behavior. Although output declined precipitously at the start of the crisis the economy subsequently rebounded strongly along certain dimensions. While the health cost of the Covid-19 recession has persisted, the economic effects seems less persistent.

In this paper we seek to understand what can be learnt from household consumption patterns during the Covid-19 recessionConsumption has typically been thought to be highly informative about the economic environment facing households. This line of thinking extends back at least to Friedman's Permanant Income Hypothesis. More broadly, a rich literature has used consumption to identify otherwise difficult to measure features of the economic environment, such as the contribution of permanant and transitory shocks in a households income process. Consumption is informative because it is a forward looking household choice that should update to reflect changes in their expectations.

The Covid-19 recession presents a particular challenge in this regard. As documented in Section 2, total personal consumption expenditures fell dramatically with the onset of the Covid-19 recession and the patterns of change were somewhat unusual. For example, as discussed by x the Covid-19 was particularly focused on a reduction in services rather than the usually highly sensitive durables consumption. This is the opposite of a typical recession. Clearly, a large part of the explanation for these movement was not primarily the *economic* considerations of households but instead considerations specific to the pandemic recession. Indeed, it may seem that consumption choices are not information about households economic circumstances during this period. This paper seeks to present a full analysis of the household consumption choices during the Covid-19 recession. It proposes a methodology for extracting what we describe as a "robust" measure of consumption from the available data, making use of high frequency (publicly available) aggregate and micro economic data. We argue that this measure of consumption can be thought to act as a unbiased proxy for the broader measure of household consumption that has previously been used to undertake inference on important features of the economic environment affecting households. We construct this measure by identifying sub-components of the household consumption basket that underwent extreme changes during the window in which the majority of health related information about the pandemic was revealed to households. We in effect partition total consumption into two baskets: a "robust" measure and "other" consumption goods. Further, our method constrains these two baskets to have comparable time series properties, thus ensuring our "robust" measure of consumption in not simply composed of goods and services that would not be informative of more typical recessions. Relative to the literature we address this problem in an entirely data-driven and agnostic procedure. This not only means that we do not impose any *a prori* assumptions about the way the pandemic would impact the consumption basket, but that the methodology is generalizable to other economic events which the household basket might undergo a substantial change, for example an exchange rate shock.

Based on this partition we are able to assess households response to the Covid-19 recession. We see our "robust" measure of non-durables consumption and services still falls dramatically at the start of the Covid-19 recession, falling by 1.1 percent at the trough, but rebounds fairly strongly. Possibly, indicating improved household expecations for the future. We also compute a "robust" measure of durables consumption. The dynamics of the "robust" durables measure are more similar to the "other" durables consumption series. Importantly, at a monthly frequency we see this robust measure did fall dramatically at the onset of the Covid-19 recession. Similarly, to the consumption measure the "robust" durables measure undergoes significant growth about one year after the recession begins. We also construct the same consumption basket in data derived from the US Consumer Expenditure Survey

While the "robust" measures of consumption are of interest in of themselves, they are not sufficient to attempt to infer households expectations during the Covid-19 recession. In part because the goods may have important substitutable or complementary properties. In particular, the reduced expenditure that households save on their restricted basket may inflate our "robust" measure and when interpreting the data only mask a larger decline in expectations. Further, the Covid-19 recession saw extremely large level of fiscal stimulus. Such feature are important to disentangle from households expectations. To interpret these consumption dynamics we build a sophisticated heterogenous agent life-cycle model featuring three consumption goods: "robust" consumption, other consumption and durables consumption and a consumption-savings decision. The model features a rich income process that allows us to capture permanent, persistent and transitory shocks. We allow the shocks to differ by age and education. We also incorporate aggregate shocks to income and prices. This enables us to credibly calibrate the model to the data and estimate preference parameters of interest by the Simulated Method of Moments.

Having constructed this model we use it to analyse the consumption dynamics of the Covid-19 recession. We do this by estimating the period by period shocks experienced by households. We treat the difference between the model generated consumption series and estimated "robust" consumption series as informative for underlying household expectations. We use this to estimate changing measures of household *uncertainty* and household expectations of the *persistence* of the economic component of the Covid-19 recession over its duration. We find that households beliefs deteriorated over the course of the recession. While initially the expected persistence of shocks was low, this lengthed dramatically as more information became available to households.

1.1 Related literature

This paper is of course closely related to the large literature which since the start of the Covid-19 seeking to understand the macroeconomics in pandemic driven recession. This literature is already large and growing, but perhaps some of the most notable examples include Atkeson (2020),Eichenbaum et al. (2020), Boppart et al. (2020), Farboodi et al. (2021), and Glover et al. (2020). The method of these papers is to apply the SIR (or extended version) from the epidemiological literature to in most cases a fairly reduced macroeconomic model to understand how an pandemic dynamics might proceed and to consider optimal policy. Kaplan et al. (2020) is an exception in that the combine the SIR model with a rich heterogenous agents framework. The approach here is quite different, in that we do note build and SIR model into our rich heterogenous agents framework, but instead focus on the economic determinants of households behaviour during the Covid-19 recession and what can be learnt about these features of the data. Our work is related to those papers that studied real time spending in proprietary datasets during the Covid-19 recessions. Examples from this literature include Bachas et al. (2020), Andersen et al. (2020), Chetty et al. (2020) and Carvalho et al. (2020). Many of the insight related to the nature of the consumption decline in these papers are similar. The novelty of our approach is to combine empirical analysis with a model driven interpretation of the consumption patterns to retirieve information about the economic shock during 2020. Perhaps closest of the Covid-19 literature to ours is Beraja and Wolf (2021) who identify and consider the implications of a services rather than durables led recession, but the empirical analysis and theoretical contribution is quite different.

The approach in this paper follows that of Heathcote et al. (2014), Blundell et al. (2008) and Guvenen and Smith (2014) which treat consumption as a useful measure of understanding the labor income risk facing households. use spending on cars in particular to understand the shocks associated with the Great Recession. The model studied is a rich heteregenous agents model. Version of this model have been analysed by Kaplan and Violante (2014), Berger and Vavra (2015) and Berger et al. (2018). Our approach differs in that the focus of our analysis is understanding consumption in the Covid-19 recession.

2 Empirical section

2.1 Consumption fall during the Covid-19 recession

This section investigate the effect of the Covid-19 recession on consumption and compares it to the experience of previous recession. We make use of BEA NIPA data on consumption by type,which is available at a monthly frequency. Figure 7 plots consumption dynamics by major type during the Covid-19 recession. We also show all other recessions since 1959 and highlight the Great Recession. Panel (a.) shows the full consumption basket: Personal Consumption Expenditures. As can be seen consumption underwent a dramatic decline during the recession falling by more than 15 percent relative to its pre-recession level. This dwarfs the typical decline in consumption seen in recessions. However, the series does rebound. By the end of the window the total consumption basket is fairly close to a trend recovery. In comparison by this point in the Great Recession consumption was still declining.

Next we analyse the consumption patterns in the major sub groups. In panel (b) we see goods displayed an even more significant decline in March 2020. After this the pattern of goods consumption rebounded strongly. For non-durables, in panel (e.) we see evidence of precautionary storage behavior, with strong growth in this consumption type at the very beginning of the recession, followed by a dramatic decline. The one type of consumption that stands out in particular is services. This also fell dramatical as the recession began, but has completely failed to recover. This major component of consumption continues to be substantially below almost any historical recession decline more than 18 months after the recession began. While this is perhaps not surprising as the restrictions placed on many services, such as cafes, restaurants and live entertainment, essentially closed this sector during the lockdown period. Even as the pandemic has continued many of these industries have faced restricted access or preferences for their consumption has fallen. However, the magnitude is stark and indicates any broad measure of consumption is likely to impacted by household restricted basket.

Figure 1: Monthly consumption decline

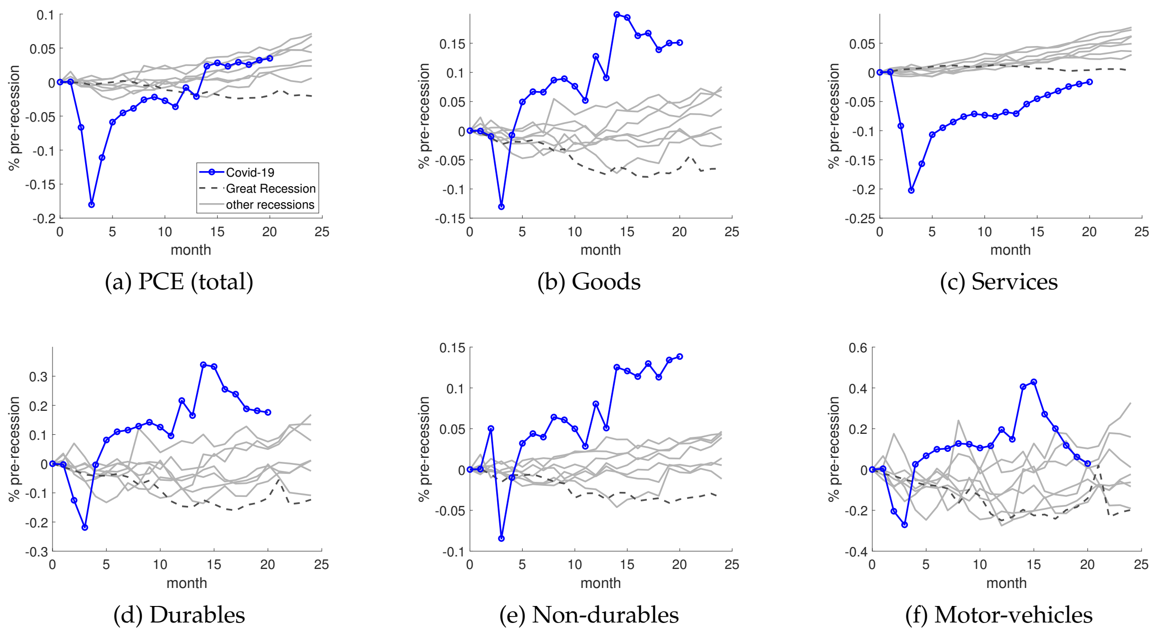

(a) PCE (total) (b) Goods (c) Services

(d) Durables (e) Non-durables (f) Motor-vehicles

2.2 Heterogenous labor market effects

The changes in the labor market during the Covid-19 recession have been extensively documented, so we only briefly describe them here. The top two panels of Figure 2.2 show income data from the NIPA. Panel (a) shows the large fall in compensation of employees during the Covid-19 recession. From peak to trough this declined by around 7.5 percent. Panel (b) shows government transfers to households. It is striking how large these transfer were relative to previous recession experiences, this constituted a major fiscal stimulus. These transfer include extended Unemployment Insurance benefit associated with higher unemployment but also the significant fiscal stimulus cheques. The lower half of the panel shows labor market data from the Current Population Survey. In Panel (c) we plot the job separation rate, by low and high education group.[1] We see the separation rate increased substantially for both education groups, but peaked substantially higher for the low education types. Job finding in Panel (d) fell initially before spiking as many workers were able to return to employment. As the recession progressed the job finding rate fell in a similar way for high and low education workers.

Figure 2: Monthly labor market

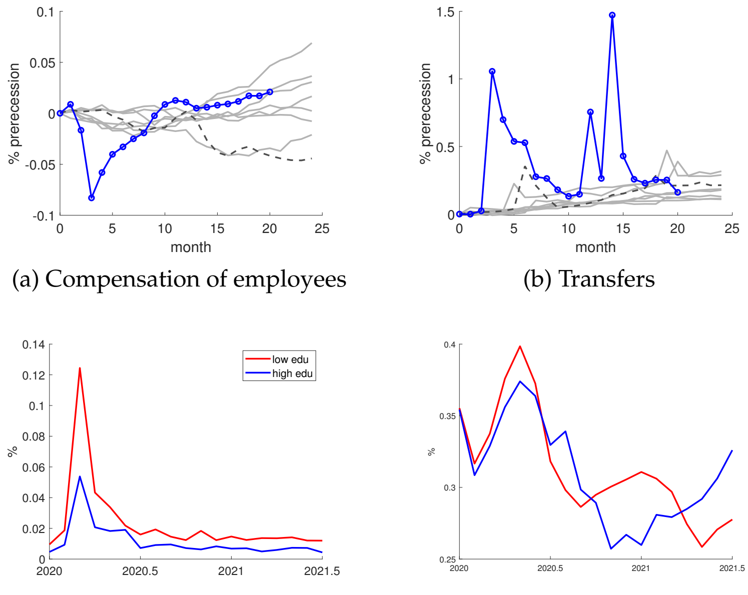

(a) Compensation of employees (b) Transfers

(c) Job separation (d) Job finding

3 Recovering a robust consumption measure

3.1 Empirical approach

The size the consumption decline seen during the Covid-19 recession was extremely large. It is also clear that the size of the decline cannot be accounted for by purely economic factors and the forward looking behavior of households. Due to a combination of changing tastes and government restrictions the basket of goods available to households changed substantially.

[1] 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 computed analogously.

Therefore, one might little information on the economic consequence of the recession could still be extracted from household decisions.

Our approach is to instead form two consumption baskets. One basket of goods likely to be affected by non-economic factors such as changes in behaviours and restrictions and one basket of good that should be relatively unaffected by these considerations and therefore contains information on the forward looking behavior of households. We categorize these baskets as "other" and "robust" consumption respectively. To do this we take advantage the high-frequency data provided by the Bureau of Econmic Affairs for consumption categories and the fact that information about the pandemic side of the recession was revealed in a discrete manner. In practice our approach is to partion the consumption basket based on goods and services that underwent extremely large changes in consumption in March 2020. Further, we constrain the groupings such that both sets of goods have similar business cycle properties prior to the Covid-19 recession. This ensures the method does not simply recover high- and low- business cycle sensitivity goods.

3.1.1 Partitioning consumption by size of shock

For a given consumption aggregate, $\mathbf{C_t}$ (Non-durables and services and durables consumption) we take advantage of the rich information available on spending by type of product, c_t^j in the underlying data NIPA table (Table 2.4.4u). For durables this is 43 sub-types of products. For nondurable goods and services this is 181 sub-types. This data is available from 1959, or 753 monthly observations. All components are deflated by their sub-product specific price deflator. Our approach is then to categorize each sub-product into a basket of "robust", $\mathbf{C_t^R}$, or "other", $\mathbf{C_t^O}$, consumption goods:

$$\mathbf{C_t} = \mathbf{C_t^R} \cup \mathbf{C_t^O} \equiv \sum_{j=1}^{J} c_t^j$$

For each sub-type we calculate the time series properties of each sub-type prior to 2020.[2] In particular we calculate the average log growth rate $\bar{g}^j = \sum_{t=1}^{T} g_t^j = \log(c_t^j) - \log(c_{t-1}^j)$ and the standard deviation of the growth rates $\sigma(g^j)$. We make an initial partition of the the aggregate consumption category, based on whether the normalized absolute change in the period associated with the largest *aggregate* consumption fall, t^* is more than γ the average standard deviation of the growth rate of that consumption sub-type. We set $t^* =$ March 2020 except for Motor vehicles which show anticipatory effect so we set t^* one month earlier. Here, γ is a parameter which governs the sensitivity of the partition. We set $\gamma = 3$, implying goods and services that underwent extremely large falls, by historical standards are allocated to the other basket. Whilst we are mainly interested in goods that experienced a very large fall due to the pandemic restrictions, we also want to exclude those goods that underwent extremely large increase. This could have been due to stockpiling behavior, for example. Thus our

[2]Note: not all sub-types are available for the full series length. We use the longest length possible.

defintion of the consumption partition can be written as:

$$c_j \in \begin{cases} \mathbf{C^O} & if\ |g^j_{t^*}| > \gamma\sigma(g_j) \\ \mathbf{C^R} & otherwise \end{cases} \tag{1}$$

We do this for both non-durables and services and durables consumption.[3]

3.1.2 Consistent business cycle properties: minimization step

In our method for partitioning consumption sub-types into "robust" and "other" groups we use the normalized growth rate of a consumption good rather than the change seen at the peak of the Covid-19 recession. This should ensure that the partition does not simply split the sample into high- and low- business cycle sensitive goods. This would confound our analysis as it would mean the subset of the consumption goods that we attempt to extract information from would be less elastic to income changes, for example. This might imply incorrect inference about the information the "robust" provided about the current and future state of the economy.

Ideally, we would like that the two consumption baskets have similar business cycle properties. To ensure this is the case after the inital partition, we apply a second minimization step. The idea here is based on the initial split of the sample we apply a minimisation routine to reassign weights $\omega_j \in (0,1)$ to each sub-product, such that the allocation minimizes the difference between the business cycle properties of the two goods while maximising the difference in the (absolute) decline at the peak of the covid recession. In particular, we choose weights than minimize the following objective function:

$$\{\omega_j\}_{j=1}^J = \text{argmax} \left|\sigma_t(\mathbf{C^R_t}) - \sigma_t(\mathbf{C^O_t}, \mathbf{Y})\right| + \left|\rho_t(\mathbf{C^R_t}, Y_t) - \rho_t(\mathbf{C^O_t})\right| - \tag{2}$$

$$\left|\sum_{j=1}^J \omega_j |g^j_{t^*}| - \sum_{j=1}^J (1-\omega_j)|g^j_{t^*}|\right| + \left|\sigma_j(\omega_j|g^j_{t^*}|) - \sigma_j((1-\omega_j)|g^j_{t^*}|)\right| \tag{3}$$

$$s.t. \tag{4}$$

$$\mathbf{C^R_t} = \sum_{j=1}^J \omega_j c^j_t \tag{5}$$

$$\mathbf{C^O_t} = \sum_{j=1}^J (1-\omega_j) c^j_t \tag{6}$$

$$\omega_j \in (0,1) \tag{7}$$

The first line of equation (2) defines a set of time series properties of the aggregate "robust" consumption series and "other" series. We choose weights to minimize the difference in the standard deviation of the two series and the correlation with a measure of income Y_t. Both of these moments relate to the properties prior to 2020. The second line refers to the growth rate

[3]We experimented with alternative partition definitions, including the raw size of the fall, rather than the normalized change, and the method applied to the residuals of consumption growth on income. While the results differ somewhat, the overall basket definitions are similar.

of the consumption sub-type in one month of the Covid-19 recession. Here the weights are chosen to maximize the difference between the average decline in consumption of the two groups, while minimising the within variance of the growth rates.[4]

3.2 A "robust" consumption measure

Table 1 provides some summary statistics of the final allocation between "robust" and "other" consumption goods for nondurables and services. The results for durable consumption can be found in the appendix. The robust consumption share accounts for 43 percent of the nondurables and services consumption basket. Therefore, a substantial component of consumption could still be useful for infering household expectations. We also report some summary statistics regarding the time series properties of the two baskets. While the "robust" basket does appear to be somewhat less volatile than the "other" set of goods and therefore somewhat less sensitive to business cycle movements, overall this problem does not seem too severe. In particular the quarterly correlation with income is very similar for both series. Therefore, we conclude the methodology is able to identify an informative component of consumption, which should be unaffected by the non-economic disturbances associated with the pandemic.

Table 1: Consumption partition summary statistics

Moment	$\mathbf{C^R}$	$\mathbf{C^O}$
Share of consumption basket	0.43	0.67
Standard dev. relative to income	0.47	0.70
Correlation with all consumption	0.75	0.95
Correlation with all income	0.72	0.70
Correlation with $\mathbf{C^R}$	1.0	0.57

We now turn to what these series can tell us about the behavior of consumption during the Covid-19 recession. Figure 3 plots measure of "robust" and "other" consumption for non-durables and services (top row) and durables consumption (bottom row). Considering non-durables and services first, in Panel (a) we see the "robust" measure undergoes a much small decline at the start of the Covid-19 recession. Its peak decline still occurs in March 2020, but is only 1.1 percent below the level of this aggregate at the beginning of the recession. By contrast the other composite consumption good undergoes a decline of almost 30 percent, see Panel (b). The substantial difference between these two is of course largely mechanical given the method of partitioning the series. However, the 1.1 percent decline in the "robust" consumption measure is still large by historical standard. Indeed, in only one other recession did the basket decline by this much.[5] We also see the decline in the "other" measure of consumption in previous recession is more in keeping with movements in the "robust" measure. Again, this is largely by design in that the methodology sought

[4]This method has similarities to the Random Forest approach in that we are defining two groups based on mimimising the within-group variance. However, here we are also penalizing the function to ensure the two groups have similar attributes in some other dimension.

[5]In the 1960 recession the peak decline was 2.5 percent

to identify consumption goods particularly hit in the pandemic. However, it highlights that what has not been identified is goods and services *generally* sensitive to movements in income or expecations for the future. Thus the "robust" basket should be useful in extracting information about such expectations.

Figure 3: Robust consumption

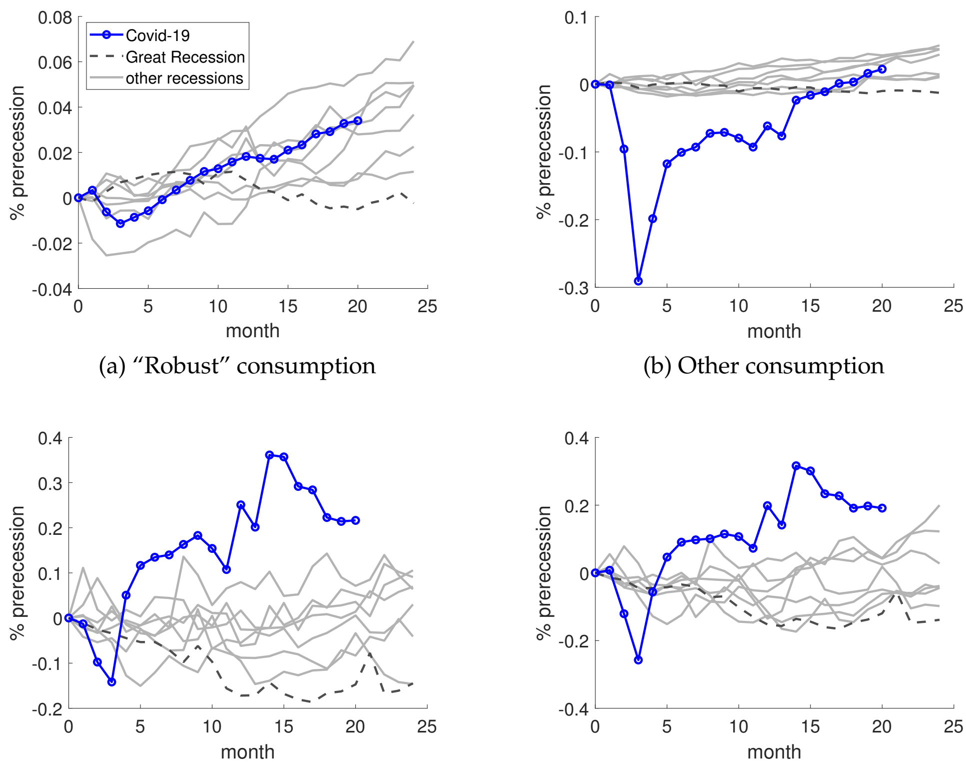

(a) "Robust" consumption (b) Other consumption

(c) "Robust" durables consumption (d) Other durables consumption

Considering, the behavior of the series during the course of the recession we notice the "robust" measure makes a reasonably quick recovery. After eight months it has exceeded the level at the beginning of the recession. Its growth rate during the recession is broadly in keeping with the fastest recoveries. This contrasts with the behavior of the "robust" basket during the Great Recession. As has been documented, this recession had particularly large and persistent effects on consumption. After 25 month this month had barely exceed its level at the beginning of the downturn. The "other" consumption measure remains depressed for most of the time horizon, reflecting the ongoing effect of restrictions and voluntary social distancing. The continuing weakness of this component is large but ultimately is relatively uniformative about the ongoing economic conditions.

The lower two panels show the same partition exercise for durables consumption. Here a similar but less extreme divergence between the two series is observable. In Panel (c) we see the "robust" measure of durables consumption undergoes a large decline of 14.2 percent three months after the recession begins. As with the "robust" consumption measure this is still a historically large decline. The other durables consumption composite falls by 25.7 percent, this is a larger but more similar decline to the "robust" measure. Thus we infer that while durables were affected by the Covid-19 specific restrictions, the effect was smaller than for non-durables and services. These patterns fit with the conventional wisdom of the Covid-19 experience. After the initial decline both durables consumption measures make a strong recovery. The recovery of the "robust" measure is somewhat faster but the magnitudes are similar, with durables consumption exceeding 30 percent above the level at the start of the recession around 15 months after it begins. This growth is far stronger than the recovery or either of the durables consumption measures in previous recessions.

In addition to partitioning consumption quantities, we also apply the same weights to the price of the baskets. Here we use the implied price deflators for each category 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 4. Panel (a) shows the behavior of the price of "other" 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 discounting of existing stocks of goods and services that retailers were no longer able to sell given the restrictions. After this initial decline lower prices persist before recovering after around one year. 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 5 percent higher than it was at the beginning of the recession. In this case, the price rise likely reflect supply shortages in the face of growth in demand. One final notable feature of both panels is that the patterns in the previous recessions are very heterogenous, 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-19 recession, based on the relative prices there is nothing particularly specfic to them in all recession. Thus, the "robust" consumption basket should provide a good summary behavior of unrestricted behavior in responseto economic shocks.

Figure 4: Relative price changes

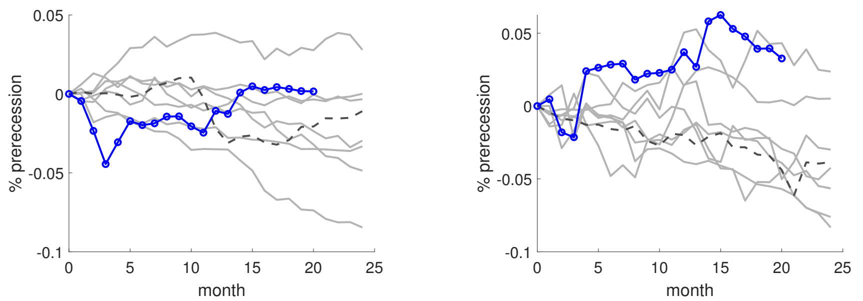

(a) Price: Other consumption (b) Price: "Robust" durables

3.3 Analysis of "robust" consumption in the micro data

To be completed

3.4 Accounting for bias in the "robust" consumption measure

4 Life-cycle model of consumption

This section sets out the life-cycle model that we will use to study household choices during the Covid-19 Recession and to infer household expectations. The model is a partial equilibrium model with aggregate shocks. For computational expediency we do not model the production side of the economy. Aggregate shocks and prices follow processes consistent with household expectations.

4.1 Household problem

4.1.1 Household preferences

Households are finitely lived. They enter at age $a = 1$ and each period a household survive until the next period with probability π_{a+s-t}. 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, $c^R_{a,s}$, *"other"* non-durables consumption, $c^O_{a,s}$ and a service flow from their durables stock, $d_{a,s}$. For the purpose of this analysis all durables consumption is considered to be *"robust"*.[6] The houshold felicity utility function u_a is allowed to depend on age to capture life-cycle features. Utility over consumption goods is also affected by a time-varying taste shifter, χ_s, which captures the utility derived from the *"other"* non-durables consumption good, $c^O_{a,s}$. A fall in the value of χ_s captures the decline in value of certain goods during the Covid-19 recession.

Warm glow bequests are defined over end of period assets, $b_{a+1,s+1}$, (discussed later) and durable stocks, $d_{a+1,s+1}$. Given this setting households make choices to maximize the expected value of lifetime utility:

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

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

4.1.2 Household durables choice

Households durables stock follows a standard law of motion:

$$d_{a+1,s+1} = (1-\delta)d_{a,s} + i^d_{a,s}$$

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. Households wishing to adjust their durables stock holdings must pay a cost $\Psi(d_t, d_{t+1})$ whereas households that allow their durables stock to depreciate incur no adjustment costs. The specification of the durables adjustment cost function is:

$$\Psi(d_t, d_{t+1}) = \begin{cases} 0 & if\ d_{a+1,s+1} = (1-\delta)d_{a,s} \\ \Psi d_{a,s} & otherwise \end{cases} \tag{8}$$

Such adjustment non-convex costs deliver lumpy durables investment and (S,s) style policy function in the current period durables stock.

4.2 Asset choice

Households are able to save for precautionary and life-cycle purposes 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 make take up a loan against their durables purchase, with downpayment $(1-\xi)d_{a,s+1}$, such that they face the collateral constraint:

$$b_{a,s+1} \geq -\xi d_{a,s+1} \tag{9}$$

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 $d_0 \sim \log \mathcal{N}(\mu_{d_0}, \sigma^2_{d_0})$. Household 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.

4.3 Income process

Household income follows a rich specfication 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 $y^e_{a,s}(\mathcal{Z}_s, z^e_{a,s}, \iota_{a,s})$. Where $\mathcal{Z}_s$ 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 $\iota_{a,s} = 0$ the household is unemployed. Income has the following functional form:

$$y^e_{a,s}(\mathcal{Z}_s, z_{a,s}, \iota_{a,s}) = \begin{cases} \mathcal{Z}_s \exp(\alpha^e_a + z^e_{a,s}) & if\ \iota_{a,s} = 1 \\ \mu \mathcal{Z}_s \exp(\alpha^e_a + z^e_{a,s}) & if\ \iota_{a,s} = 0 \end{cases} \tag{10}$$

Household income follows a deterministic path $\{\alpha_a^e\}_{a=1}^{\mathcal{A}^r}$ 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:

$$z_{a+1,s+1}^e = \rho_{z^e} z_{a+1,s+1}^e + \varepsilon_{a,s}^e \tag{11}$$

With $\epsilon_{a,s}^e \sim \mathcal{N}(0, \sigma_{z^e,s}^2)$. We allow the persistence ρ_{z^e} and variance of $\sigma_{z^e,s}^2$ of the process to vary by education level. In our main analysis we also allow for the variance of shocks $\sigma_{z^e,s}^2$, to vary deterministically over time capturing *increased uncertainty* facing households during a recession.

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

$$\mathcal{P}(\iota_{a,s}, \iota_{a+1,s+1}; a, e, \mathcal{Z}_s) = \begin{bmatrix} 1 - \lambda_{a,e,\mathcal{Z}_s}^1 & \lambda_{a,e,\mathcal{Z}_s}^1 \\ \lambda_{a,e,\mathcal{Z}_s}^0 & 1 - \lambda_{a,e,\mathcal{Z}_s}^0 \end{bmatrix} \tag{12}$$

Aggregate shocks: Households also experience aggregate income shocks $\mathcal{Z}_s$ and can be interpreted as TFP shocks. Aggregate shocks follow a Markov matrix $\mathcal{P}(\mathcal{Z}_s, \mathcal{Z}_{s+1})$.

Retirement: After $\mathcal{A}^w$ periods, households enter retirement. In the retirment state household receive a pension which is an education specific share, κ^e of their final salary.[7] Pension income is then constant for the rest of the household life.

4.4 Budget constraint

We can now write out the household budget constraint. Households choose a consumption basket that includes robust non-durables, $c_{a,s}^R$, other non-durables, $c_{a,s}^O$, and durables consumption, $i_{a,s}^d$. The later two goods have the relative prices p_s and p_s^d, respectively. They also make a savings decision, $b_{a,s+1}$. The resources each period are income $y_{a,s}^e(\mathcal{Z}_s, z_{a,s}$, if of working age, or pension; the bonds saved from last period $b_{a-1,s}$ after interest; and a government transfer, $\mathcal{T}_s$. Should the household adjust their durables stock the must also pay the adjustment cost $\Psi(d_{a-1,s}, d_{a,s+1})$.

$$c_{a,s}^R + p_s c_{a,s}^O + b_{a,s+1} + p_s^d i_{a,s}^d = y_{a,s}^e(\mathcal{Z}_s, z_{a,s}, \iota_{a,s}) + (1+r)b_{a-1,s} + \Psi(d_{a-1,s}, d_{a,s+1}) + \mathcal{T}_s \tag{13}$$

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

4.5 Recursive formulation

The household choice problem can be presented in a recursive formulation. For this purpose we supress 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 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 and employment status, ι. The aggregate state variables are productivity, $\mathcal{Z}$, a vector of relative prices, $\mathbf{p} = \{p, p^d\}$, and a taste shifter for other non-durable consumption goods, χ.

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

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

subject to

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

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

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

subject to

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

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

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

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

4.5.1 Discussion

4.6 Calibration and numerical implementation

4.6.1 Utility function form

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

$$u_a(c^R_{a,t}, c^O_{a,t}, d_{a,t+1}; \chi_s) = u\Big(\mathcal{C}(c^R_{a,t}, c^O_{a,t}; \chi_s), d_{a,t+1}\Big) \tag{14}$$

With the preferences of consumption and durables consumption defined by

$$u(\mathcal{C}(c^R_{a,t}, c^O_{a,t}; \chi_s), d_{a,t+1}) = \frac{\left[\theta^{\frac{1}{\hat{\epsilon}}}(\mathcal{C}(c^R_{a,t}, c^O_{a,t}; \chi_s))^{\frac{\hat{\epsilon}-1}{\hat{\epsilon}}} + (1-\theta)^{\frac{1}{\hat{\epsilon}}}(d_{a,t+1})^{\frac{\hat{\epsilon}-1}{\hat{\epsilon}}}\right]^{\frac{(1-\rho)\hat{\epsilon}}{\hat{\epsilon}-1}} - 1}{1-\rho} \tag{15}$$

And the sub-utility function over consumption given by:

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

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, $\hat{\epsilon} \in (0,1)$, the share of consumption allocated to "robust" consumption $\theta_c \in (0,1)$, the elasticity of substitution between robust and other consumption $\epsilon \in (0,1)$.

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

$$\mathcal{B}(b_{a,t+1}, d_{a,t+1}) = \frac{\mathcal{B}_0(\mathcal{B}_1 + (b_{a,t+1} + (1-\delta)d_{a,t+1}))^{1-\rho_{\mathcal{B}}}) - 1}{1-\rho_{\mathcal{B}}} \tag{16}$$

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

4.6.2 Externally calibrated parameters

The model we propose is rich with a large number parameters. We calibrate the model with a mixture of choices from external sources and internally calibrated parameters with the method of simulated moments. A time period in the model is a quarter. Starting with the inverse of elasticity of intertemporal substitution, ρ we set this to 1.5. This is a common value in the literature and consistent with evidence from Attanasio (2000). We assume consumption and durables are a Cobb-Douglas. This is convenient as it somewhat simplifies the estimation of taste shocks later. The steady state taste shock, χ, is set to 1. The durables depreciation rate, δ is set to 0.047, in line with BEA estimates. In keeping with our prior data analysis we use two education groups, defining the low type as no-college or some college and the high type as complete college. From the CPS this implies 65 percent of households are of low education type. We estimate the persistent income process separately for high and low education individuals in the PSID. The value are similar with higher education households facing slightly higher variance shocks. We then apply a five-state discrete approximation using standard Tauchen methods. The unemployment replacement rate is set to μ following Shimer (2000). We set the pensions replacement rate, κ, for high educated households,to 0.6 based on evidence in Bernheim (2004). We set the rate a little higher for the low educated household, $\kappa = 0.65$.

Table 2: Externally calibrated parameters

Parameter		Value
Elasticity of intertemporal substitution^{-1}	ρ	1.5
Elasticity of substitution (durables)	$\hat{\epsilon}$	1
Steady state taste shock	χ	1
Durables depreciation rate	δ	0.047
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
Curvature of bequest motive	$\rho_{\mathcal{B}}$	1.5
Luxury value of bequest	$\mathcal{B}_1$	0
Productivity shocks	$\mathcal{Z}$	$\{0.96, 0.98, 1.01\}$
Productivity shocks probabilities	$\mathcal{P}(\mathcal{Z}, \mathcal{Z}')$	see appendix
Persistence of consumption price	ρ^p	0.87
Variance of consumption price shocks	σ^2_p	0.005

The interest rate is set to a standard value of 4 percent on an annual basis. We 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, $\rho_{\mathcal{B}} = \rho$ while we assume no luxury bequests.

For the aggregate shock we estimate processes in the data. For productivity we assume a three state process, capturing booms, recessions and deep recessions. We use the un-utilisation adjusted estimates of Fernald (2021), defining a deep recession as a 3σ deviation and recession and booms as $-\sigma$ and $+\sigma$ events respectively. We then fit a Markov chain based on the neares 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 of which both the Great Recession and Covid-19 recession would have been categorized as. For the aggregate price of the "other" consumption good we 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. All externally calibrated parameter values are given in Table 5. We also target a

Table 3: Life-cycle parameters

Parameter		Source
Survival probabilities	π_a	NCHS
Lifecycle income profile	α_a^e	PSID
Job separation rate	$\lambda^0_{a,e,\mathcal{Z}}$	CPS
Job finding rate	$\lambda^1_{a,e,\mathcal{Z}}$	CPS

number of parameters which vary over the life-cycle. In particular we take pre-pandemic survival probabilities from the National Center for Health Statistics (NCHS) and estimate the exogenous life-cycle profile of low and high education types from the PSID, see Table 3. We also compute age-education-productivity level specific labor market transitions from the CPS.[8] These are shown in the top two panels of Figure 5. As can be seen separation rates and finding rate fall over the life-cycle, with the exception of the high education separation rate which is fairly flat. Separation rates are much high for low education types, whereas the difference between the job finding rates of high and low education households is less pronounced. Finally, the effect of recessions in very non-linear with change between booms and recession substantially smaller than the change between recessions and deep recessions.

Figure 5: Life-cycle targets

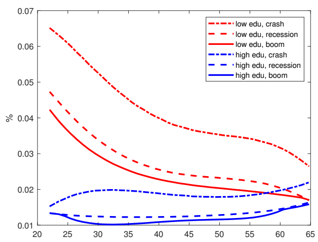

(a) Job separation

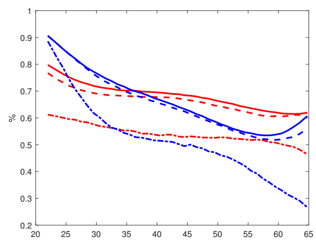

(b) Job finding

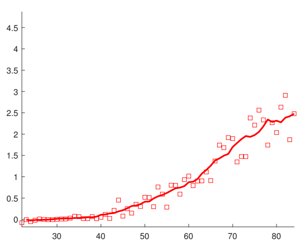

(c) Wealth:inc, low edu

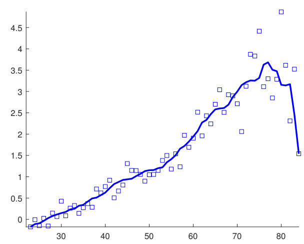

(d) Wealth:inc, hi edu

4.6.3 Model estimation

We also calibrate a number of parameters within the model using the Simulated Method of Moments approach. To do this we simulate a non-stochastic cross-sectional distribution of

[8]These transition rates are estimated on pre- Covid-19 recession data.

households, subjected to a series of aggregate shocks. We use this simulation to calibrate the "robust" share of consumption and durables share of consumption and identify the parameters θ and θ_c. We also choose the adjustment cost on durables, Ψ, to target the relative standard deviation of durables consumption and the correlation between the "robust" and other consumption good to identify the elasticity of substituion, ϵ. These moments related to the results of our calibration exercise are shown Table 4 and the resulting parameters are shown in Table 5.

Table 4: Targeted moments

Moment		Parameter	Data	Model
Robust consumption ratio	C^R/C^O	θ_c	0.68	0.68
Durables consumption share	$I^d/(C^R + C^O + I^d)$	θ	0.05	0.06
Std dev of durables consumption	$\sigma(I^d)/\sigma(Y)$	Ψ	2.89	3.0
Correlation between consumption goods	$\rho(C^R, C^O)$	ϵ	0.57	0.55

Table 5: Internally calibrated parameters

Parameter		Value
Share of consumption	θ	0.94
Share of robust consumption	θ_c	0.404
Elasticity of substitution (consumption)	$\hat{\epsilon}$	1.05
Adjustment cost of durables	γ	0.08
Discount factor	β	$\{0.985, 0.988\}$
Value of bequest	$\mathcal{B}_0$	$\{0.1, 0.1\}$

4.7 Household life-cycle

Figure 6 displays lifecycle profiles for the consumtion goods, durables stock and assets. Given the nature of the preferences and fixed prices "robust" and other consumption follow similar paths during the lifecycle. We see that the path of "robust" consumption is lower reflecting the smaller expenditure share. We also see the model predicts high education household to consume significantly more and to have more steeply increasing profiles at the beginning of life. A similar pattern is seen for the durables stock in Panel (b) although we notice households economize more on this in later life. Finally, Panel (c) plot median asset holdings. As is common in these model, these are mainly held to smooth consumption during retirement. We see high education households accrue significantly more savings than low education type, in keeping with the targeted data moment.

5 Extracting expectations

The main exercise undertaken in this research is to combine the statistical estimate of "robust" consumption during the Covid-19 recession from section 3 (as well as measures for "other" consumption and "robust" durables consumption) with our sophisticated life-cycle model in

Figure 6: Life-cycle profiles

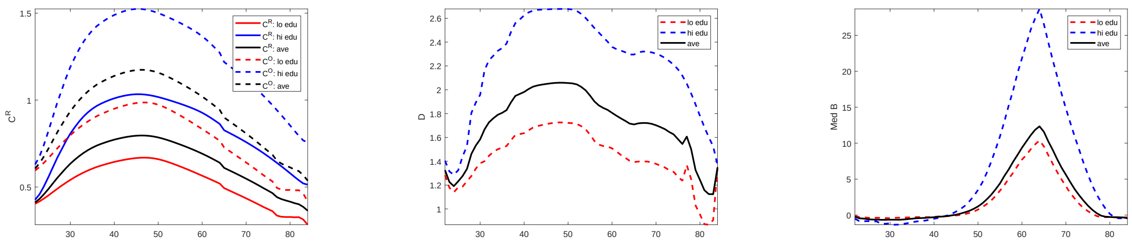

(a) Consumption (b) Durables (c) Median assets

order to understand household expectations during this period. In particular, we are interested in (i) the uncertainty faced by households; (ii) how persistent household's expectations for the economic effects of the recession were; and (iii) how perhaps most importantly how these expectations evolved.

To do this we feed in shocks to productivity, $\mathcal{Z}$, the labor market transition rates, λ^0 and λ^1, prices, p, the taste for "other" consumption goods, χ, and government transfers, $\mathcal{T}$. These series are estimated from the data and explained below. We then infer what shock to uncertainty, σ_0^2, and recession persistence ϱ is required to best match the "robust" consumption and durables consumption series.[9] The recession begins in January 2020 following the NBER definition. We compute impulse response functions for 30 periods after the recession begins.[10] The exact timing of the shocks depends on the details of the specific exercise and will be discussed in greater detail where appropriate. We now give a brief description of each shock. A precise definition is provided in the appendix.

Productivity shocks: Productivity shocks are used to match the aggregate (detrended) decline in compensation of employees shown in Figure 2 from the NIPA tables. The model counterpart excludes the effect associated with unemployment. Therefore, we construct the series as the residual after feeding in labor market shocks. If the measure of uncertainty is positive we also adjust productivity downwards, to ensure the uncertainty rise is mean preserving.

Labor market shocks: The job finding and job separation rate shocks are calculated from the CPS. Shock differ by education level and are computed as deviations to the level in previous year. We assume labor market shocks have a uniform (percentage point) impact across the lifecycle transition profiles.

Price shocks: Price shocks are fed in as deviations in the relative price of "other" consumption goods from Figure 4.

Taste shocks: We compute taste shocks as the implied decline in χ_s required to match the deviations in "robust" consumption, "other" consumption and the relative price. Given our

[9] To match the stationarity of the model we subtract a trend growth rate from these series consistent with the growth rate during the two years prior to the recession.

[10] In exercise 5.1 if the shock has not returned to trend by the end of the data series we assume the variable returns to trend with a persistence of 0.85.

preference specification we can rearrange the intratemporal consumption choice as:

$$\chi_t = \left(\frac{1-\theta_c}{\theta_c} \frac{\mathbf{C_t^R}}{\mathbf{C_t^O}} p_t^{-\epsilon} \right)^{1/(1-\epsilon)} \tag{17}$$

Government transfers: Government transfers capture the three waves of fiscal stimulus provided during the Covid-19 recession. Based on an 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.

Figure 7: Monthly consumption decline

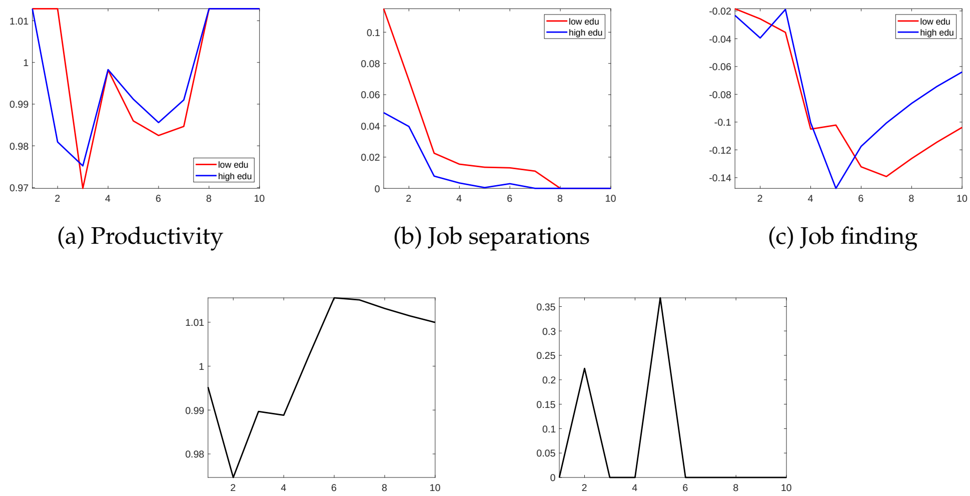

(a) Productivity (b) Job separations (c) Job finding

(d) Relative price (e) Gov. transfers

5.1 A one-time increase in uncertainty

In our first exercise we consider to what extent the remaining fall in consumption can be explained by a persistent increase in uncertainty.[11] We start the model at a steady state distribution over assets, income and unemployment.[12] In period 1 full information is revealed and households know the full path of the future shocks. In addition there is an increase in uncertainty. More precisely we assume the following process for uncertainty:

$$\sigma_{z,e,t} = \bar{\sigma}_{z,e} + \hat{\sigma}_{z,t} \tag{18}$$

$$\hat{\sigma}_{z,t} = \begin{cases} \varepsilon_1^{\sigma} & if\ t = 1 \\ \varrho \hat{\sigma}_{z,t-1} & \text{otherwise} \end{cases} \tag{19}$$

[11]We model this as an increase in labor income uncertainty, but it could be thought to capture a broader definition. However, as is the focus of this research we are seeking to capture *economic* determinants of consumption (and household expectations of) rather than duration of the medical consequences of the Covid-19 infection.

[12]The steady state distribution is such that households use the policy functions of the stochastic model, in the ergodic distribution only the high productivity state and steady state relative price, $p = 1$ is realized.

Where the standard deviation of the persistent income process increases to $\bar{\sigma}_{z,e} + \hat{\sigma}_{z,t}$ at $t = 1$ before decaying back to its steady state level with persistence ϱ.[13] This leaves two parameters to estimate $\Theta = \{\hat{\sigma}_{z,t}, \varrho\}$. These are then chosen to minimize the distance between the model and date estimates of $\mathbf{C_t^R}$ and $\mathbf{I_t^{d,R}}$. We minimize:

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

where W is a weighting matrix that assigns linearly decreasing weights to observations at greater time horizons and $\mathbf{M}(\Theta)$ is the stacked vector of deviations:

$$\mathbf{M}(\Theta) = \left[\mathbf{C_t^R}_{data} - \mathbf{C_t^R}(\Theta)_{model}, \mathbf{I_t^{d,R}}_{data} - \mathbf{I_t^{d,R}}(\Theta)_{model}. \right]$$

The results of the model with a one time shock to uncertainty are shown in Figure 8. The red line shows the model with just the shocks the black with an additional increase in uncertainty that raised the standard deviation of income by 0.05 and persists with a decay of 0.8. We can see the additional uncertainty helps the model better fit the robust consumption series, which otherwise would not decline sufficiently given the shocks. On the other hand the increase in uncertainty tends to exaggerate the durables decline relative to that seen in the data. Later in the series the model tends to recover too quickly, partly due to the large transfers. This suggests that household might be experiencing deteriorating expectations of the economic environment.

Figure 8: Consumption dynamics with a one-time increase in uncertainty

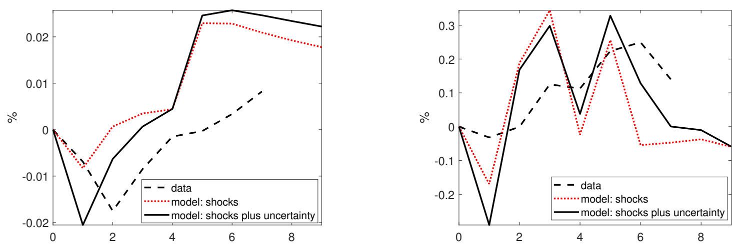

(a) "Robust" consumption (b) Durables consumption

5.2 A dynamic model of expectations

The limited success of the model featuring a one-time increase in uncertainty suggests that there could be a role for new information and the dynamic updating of expectations. To see if this is better able to explain household consumption patterns during the Covid-19 recession we undertake the following exercise. We again start from the steady state distribution, but we now estimate a series of household expectations. Each period the shock estimated in the data is realized, but rather than knowing the full path of future shocks, households also hold an expectation over the persistence with the shock it will return to trend from its current level.

[13]The timing is such that in the first period households learn future shock realisations will increase. I.e. there is no suprise income effects in period 1.

For example, in period 1, the aggregate productivity shock $\mathcal{Z}_1$ is realized and households expect the future values to follow:

$$\mathcal{Z}_t = (1 - \varrho_o)\bar{\mathcal{Z}} + \mathcal{Z}_{t-1}$$

Where the o subscript denotes the period in which the initial shock is realized. Each period we assume ϱ_o is shared by all variables, with the expectation of the price p which is assumed to follow its stochatic process from the data. As in the previous exercise we also allow for an increase in uncertainty each period $\hat{\sigma}_{z,t,o}$. Therefore, starting at an initial steady state we estimate a sequence of parameters $\{\Theta_o = \{\hat{\sigma}_{z,t,o}, \varrho_o\}\}_{o=1}^{T}$, mimimizing just the current period "robust" consumption and non-durables consumption deviations $\mathbf{M}_o(\Theta_o) = [\mathbf{C^R_o}_{data} - \mathbf{C^R_o}(\Theta_o)_{model}, \mathbf{I^{d,R}_o}_{data} - \mathbf{I^{d,R}_o}(\Theta_o)_{model}]$.

The results of the dynamic model are show in Figure **??**. The model is able capture quite closely the "robust consumption dynamics at the start of the Covid-19 recession. Here the model implies the expected persistence increases as the recession progresses. The model also does a slightly better job of delaying the rebound of durable consumption the second and third quarter of the recession. However, towards the end of the series the model still tends to predict a slightly too optimistic view of household expectations than the "robust" consumption series implies.

Figure 9: Consumption dynamics in a dynamic model of expectations

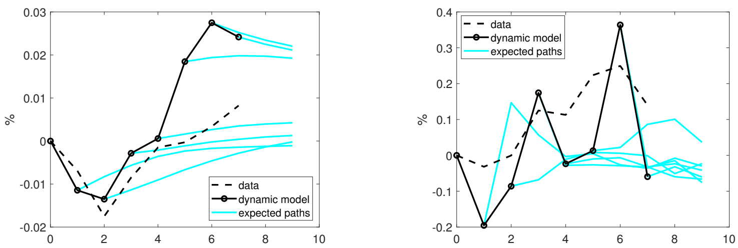

(a) "Robust" consumption (b) Durables consumption

5.3 Analysis by education level

To be completed

6 Conclusion

This paper has proposed a "robust" measure of consumption for understanding household expectations during the Covid-19 recession. It computed this measure by partitioning aggregate consumption based on the magnitude of the decline of a large subset of consumption types at the peak of the pandemic. It then integrated this measure with a sophisticated lifecycle model to study the evolution of consumption dynamics during the Covid-19 recession and the extract information of households expectations during this period.

References

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. NBER Working Papers 26867, National Bureau of Economic Research, Inc.

Bachas, N., Ganong, P., Noel, P. J., Vavra, J. S., Wong, A., Farrell, D., and Greig, F. E. (2020). Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data. Working Paper 27617, National Bureau of Economic Research.

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

Berger, D., Guerrieri, V., Lorenzoni, G., and Vavra, J. (2018). House Prices and Consumer Spending. Technical Report 3.

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

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

Boppart, T., Harmenberg, K., Hassler, J., Krusell, P., and Olsson, J. (2020). Integrated Epi-Econ Assessment. NBER Working Papers 28282, National Bureau of Economic Research, Inc.

Carvalho, V. M., Hansen, S., Ortiz, l., Ramón García, J., Rodrigo, T., Rodriguez Mora, S., and Ruiz, J. (2020). Tracking the COVID-19 Crisis with High-Resolution Transaction Data. CEPR Discussion Papers 14642, C.E.P.R. Discussion Papers.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., and Team, T. O. I. (2020). The economic impacts of covid-19: Evidence from a new public database built using private sector data. Working Paper 27431, National Bureau of Economic Research.

Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2020). The Macroeconomics of Epidemics. NBER Working Papers 26882, National Bureau of Economic Research, Inc.

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

Glover, A., Heathcote, J., Krueger, D., and Ríos-Rull, J.-V. (2020). Health versus wealth: On the distributional effects of controlling a pandemic. Working Papers 2020-038, Human Capital and Economic Opportunity Working Group.

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

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.

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.

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