# Uncovering MPC Heterogeneity: Insights from US COVID-19 Stimulus Payments

Emily Barnes

2023

**Abstract**

This paper uncovers heterogeneity in the Marginal Propensity to Consume (MPC) from the three rounds of US COVID-19 stimulus payments. Using a novel clustering method to group households by consumption changes, we estimate the full unconditional MPC distribution via a two-stage least squares approach. This is in contrast to the prevalent practice of linking MPCs solely to observable household attributes. Controls for additional income fluctuations, COVID restrictions, and increased unemployment benefits and child tax credits refine our MPC estimates. Estimated MPCs are smaller than in previous literature, 0.13 to 0.27 for total expenditures with an average MPC of 0.07, indicating limited stimulus payment impact due to heightened uncertainty. MPC heterogeneity persist across expenditure categories, with greater MPCs in durables than nondurables. We identify correlations between the MPC and various observable household characteristics, encompassing income levels, educational attainment, liquid assets, and home ownership status. These insights bear implications for both policy formulation and economic modelling, underlining contextual influences on consumer behaviour and heterogeneity.

# 1 Introduction

Understanding households' consumption responses in a variety of economic environments is of utmost importance for both economic modelling and policymaking. The marginal propensity to consume (MPC) describes the change in household consumption expenditures resulting from a change in income, and is therefore central to understanding shock propagation and the effects of monetary and fiscal policy. Until relatively recently, the primary focus had been on estimating an homogeneous MPC for an entire population. Increased attention on heterogeneous agent models has shifted focus towards understanding heterogeneity in the MPC across consumers. In particular, Heterogeneous Agent New Keynesian (HANK) models document the importance of this heterogeneity for monetary policy transmission (Kaplan et al., 2018). For example, Auclert (2019) emphasised that consumption responses to monetary policy shocks are dependent upon the cross-sectional covariances between household MPCs and their exposure to aggregate shocks. It is therefore not just the magnitude of the MPC that is important, but also its distribution, and the subsequent policy implications.

The COVID-19 pandemic presented a rare opportunity to estimate the MPC via a natural experiment as many governments provided economic relief in the form of irregular stimulus payments. The United States (US) government was one such, providing three separate rounds of stimulus payments in 2020 and 2021 as a component of three Acts[1] which aimed to mitigate the economic impact of the pandemic and the associated economic slowdown. These stimulus payments represent unexpected transitory income shocks, allowing us to estimate the MPC from the resulting household consumption responses. The effect of these stimulus payments is directly determined by the induced consumption response and any ensuing effects on wider economic activity. Since a key aim of these payments was to boost consumer spending and stimulate the macroeconomy via a multiplier effect, estimating the distribution of the MPC is crucial to understanding

[1]These were the Coronavirus Aid, Relief, and Economic Security (CARES) Act 116th Congress of the United States (2020b), the Consolidated Appropriations Act 116th Congress of the United States (2020a), and the American Rescue Plan Act 117th Congress of the United States (2021).

the efficacy of such policies.

Estimation of the MPC is not straightforward due to the inherent complexity and variability of consumer behaviour; hence a variety of approaches have been implement in the existing literature, with varying results. Estimates of the distribution of the MPC are relatively scarce, with Lewis et al. (2019) providing one of very few. Most of the existing literature employs approaches which require a stance to be taken on the source of MPC heterogeneity in order to group households and estimate the MPC of these groups. As a result, these likely do not provide true estimates of the full distribution of the MPC since they will omit any unobservable (or simply overlooked) sources of heterogeneity[2]. We therefore follow Lewis et al. (2019) and estimate the MPC distribution by grouping households based on their consumption response, studying the observable drivers ex post.

A key contribution of this paper is its application of a Gaussian mixture instrumental variable regression (GMIVR) to the estimation of the MPC using a new dataset and set of controls. To our knowledge, Lewis et al. (2019) is the only existing paper to employ this methodology to the study of consumption responses. MPC estimation requires joint estimation of the coefficient on the rebate variable and each household's group membership. A Gaussian Mixture model (GMM) is a probabilistic approach used to assign group membership based on how well the group-specific parameters describe the households within the group. The fit of the model is improved by iterating through steps to optimise the parameters of each group and the group assignment of each household. An instrumental variable approach estimates the MPC for each household within these groups, taking receipt of the stimulus as the instrument. More specifically, we apply two-stage least squares to a standard regression of the change in consumption on the economic stimulus payment and a group of controls (Johnson et al., 2006; Parker et al., 2013).

We build on the method of Lewis et al. (2019) by incorporating additional controls into the GMIVR to control for additional factors impacting household consumption responses during the COVID-19 pandemic. Increased child tax credits and unemployment benefits during this period will have induced consumption responses of their own;

[2]Lewis et al. (2019) find that less than a quarter of MPC heterogeneity can be explained by observable characteristics.

controlling for this allows us to study the MPC out of the economic stimulus payments alone. Many households experienced significant income changes during the pandemic so the addition of a control for this change in income ensures that our MPC estimates are not picking up these separate income effects. Finally, governments around the world introduced COVID restrictions including closures of shops and restaurants, limits on social gathering, and stay-at-home orders. These restrictions presented a huge supply shock and altered consumers' expenditure opportunities. In order to mitigate the effects of this on our MPC estimates, we incorporate an index for the stringency of these restrictions as an additional control variable. We are not aware of any other studies controlling for government restrictions in this way.

Estimating the MPC distribution using our GMIVR requires household-level data for income, expenditures, and observable characteristics. We construct the majority of our variables from household responses to the Consumer Expenditure Survey, collected by the US Census Bureau for the Bureau of Labour Statistics[3]. We match each variable to the relevant household based on a household ID provided by the Consumer Expenditure Survey. Our measure of COVID restrictions is constructed as a monthly aggregate of the daily stringency index produced for each US state by the Oxford Coronavirus Government Response Tracker project (Hale et al., 2021). The stringency index is then merged with the rest of our data based on the household's US state of residence. The resulting dataset contains 4,222 households across all US states who were in the Consumer Expenditure Survey's stratified sample.

Through our estimation of the distributions of the MPC for total expenditures, durable expenditures, nondurable expenditures, and food expenditures, we find notable heterogeneity in the consumption response across both households and consumption types. The extent of this heterogeneity is greater in our study than it is in some others, particularly those studying income changes unrelated to stimulus payments (see Fagereng et al., 2021), reflecting the multifaceted nature of the COVID-19 shock and its diverse impact on consumers. The distribution of the MPC for total expenditures ranges from -0.13 to 0.27, with the highest frequency of estimates at this upper bound and an average MPC of 0.07. Despite being unusual, negative estimates for the MPC are not unheard

[3]www.bls.gov/cex/

of (see Misra and Surico, 2014). This observed decrease in consumption associated with the receipt of a stimulus payment for some households is likely due to pandemic-related uncertainty, virus fears, and households' formation of new routines and habits due to government restrictions. This is also a likely explanation for the small size of the rest of our MPC estimates relative to many other studies. Karger and Rajan (2020) also find that 12% of consumers decreased spending following receipt of a stimulus payment.

In line with much of the existing literature, our estimated MPC distributions for durable, nondurable, and food expenditures present smaller consumption responses for these smaller consumption categories. Similarly to the MPCs for total expenditures, these consumption categories produce high frequencies of zero-estimated MPCs and there are no negative MPC estimates for food.

Considering the sources of this MPC heterogeneity across households, we find a number of statistically significant relationships with observable household characteristics. Total household income is found to be one of the key determinants of the MPC, with lower incomes associated with larger MPCs. Our correlation coefficient for the MPC with liquid wealth is perhaps smaller than expected given other papers' findings that it is a significant source of MPC heterogeneity (see Parker et al., 2013), although this is likely due to the low response rate for liquid wealth in the consumer expenditure survey. At the same time, the lack of statistical significance for the relationship of the MPC with nonsalary income may also be due to its own low response rate. Other statistically significant drivers of MPC heterogeneity include salary income, household size, employment status, level of education, mortgage and housing values, and housing ownership status.

This paper joins a growing literature on the wider effects of the COVID-19 pandemic on the economy (see Crossley et al., 2021; Baker et al., 2023). The COVID-19 pandemic and the subsequent economic stimulus payments in the US present an unprecedented shock and a stimulus payment package much larger than those previously studied. The scale of the economic downturn, the associated uncertainty, and the speed of its onset provide a truly novel setting. Estimating the distribution of the MPC in this context therefore provides an insight into household responses in a very different environment to that of existing literature. Our study also contributes to the broader literature estimating households' responses to income shocks. While its unprecedented nature complicates the extrapolation of our results to stimulus payment policies in general, it is nonetheless vital to understanding the impacts of the specific shock caused by COVID-19 and building a clearer picture of consumer behaviour on the whole.

We identify a number of potential avenues for future research. Studying varying aggregations of consumption expenditures and drivers of MPC heterogeneity in greater detail would provide further insight into consumer spending behaviour. Similarly, applying our methodology to estimate the distribution of the MPC from other income shocks would allow for a study of the shock-related factors influencing consumption responses. Our results are useful for informing the growing literature on heterogeneous agent models, providing insight into MPC heterogeneity and its sources.

This paper proceeds as follows. Section 2 provides an overview of the relevant literature. Section 3 outlines the COVID-19 economic stimulus payments which provide the transitory income shock for our study while Section 4 outlines the datasets used. Section 5 describes the MPC specification used in the paper and Section 6 discusses our method for estimating the distribution of the MPC. Results are presented in Section 7, and possible next steps for research are discussed in Section 8. Finally, Section 9 concludes.

## 2 Related Literature

The existing literature studying consumption responses to income changes is vast. Until relatively recently, the literature predominantly focused on estimating a homogeneous MPC representing the average consumption response of households in the economy. This is in line with a standard life-cycle permanent income model in which consumption responses to an income shock are similar across households since consumption is proportional to permanent income. This is, however, in stark contrast to the burgeoning literature on heterogeneous agent models. Early examples of this are the models of incomplete markets and idiosyncratic risk of Bewley (1983), Imrohoroğlu (1989), Huggett (1993), and Aiyagari (1994), and more recently the Heterogeneous Agent New Keynesian (HANK) models such as that of Kaplan et al. (2018).

Focus has more recently turned to estimating the heterogeneous consumption responses across households, with many papers finding notable heterogeneity. Crawley and Kuchler (2023) for example, group households with similar MPCs; the first group has an MPC of 0.8, the second an MPC of 0.6, and the third an MPC of 0.3. Further, the distinct wealth characteristics of these groups are comparable to those proposed by Kaplan et al. (2014) named "poor hand-to-mouth", "wealth hand-to-mouth", and "wealthy" respectively. There is however, significant variation in MPC estimates between studies. Lewis et al. (2019) find that the majority of MPCs are relatively low at approximately 0.27, but at the same time 14% of households have an MPC of one or above. Using the same data on 2008 US economic stimulus payments, Misra and Surico (2014) estimate that a large proportion of households have an MPC of zero or lower. The stark variation in the findings of these two papers is arguably due to their differing methods for estimating the MPC.

Havranek and Sokolova (2020) conducts a meta-analysis of studies on consumption responses, documenting the variation in estimates. They postulate a number of sources of variation: (i) the size of the income change, (ii) the time horizon studied, (iii) the consumption and income measures, (iv) the characteristics of the data used, and (v) the estimation techniques used. Existing literature estimating the effects of the size of the change in income is limited, but Fuster et al. (2021) and Christelis et al. (2019) obtain qualitatively similar findings: a negative size effect on the intensive margin and a positive size effect on the extensive margin. Generally, MPC estimates covering a longer time horizon are larger, but this consumption response is notably larger in the period of the income change than it is in subsequent periods (see Lewis et al., 2019; Parker et al., 2013). Furthermore, the estimated value of the MPC tends to rise as increasingly aggregated consumption measures are studied. For example, Souleles (1999) finds MPCs of 0.0062 for food, 0.093 for strictly nondurable goods, and 0.640 for total consumption. The latter two sources of variation in MPC estimates proposed by Havranek and Sokolova (2020) can be considered components of the methodological design of the studies.

Estimation of the MPC uses one of three methods of obtaining consumption response data. The first, which we will follow, is to measure the consumption response to a change in income in the form of a natural experiment. Perhaps the most common of these are the US 2001 tax rebates (see Shapiro and Slemrod, 2003; Johnson et al., 2006) and the

US 2008 economic stimulus payments (see Parker et al., 2013; Lewis et al., 2019). These studies report increases in spending, albeit of different magnitudes, following the receipt of a payment. While most use consumer survey data, many exploit high-frequency transaction data to avoid the high measurement error associated with such surveys. Agarwal et al. (2007) use credit card accounts for the 2001 tax rebates, Broda and Parker (2014) use high-frequency scanner data for the 2008 stimulus payments, and Fagereng et al. (2021) work with third-party household balance sheet information for lottery winners. This type of high-frequency transaction data is typically very difficult to obtain and time consuming to work with. Given that this type of data is also relatively new, it does not facilitate meaningful comparison between more recent consumption responses and those from before this data type was available.

A second data-type relies on individuals' self-reported consumption response to an income change, either hypothetical or real. Fuster et al. (2021) study hypothetical income changes, finding heterogeneity across consumers, as well as for income changes of differing sizes, signs, and lead times. They find that the mean MPC falls as the income gain increases, but the MPC is larger for hypothetical losses. News of future income changes, whether a gain or a loss, elicits a smaller consumption response than a contemporaneous change. A large-scale US survey revealed that the average consumer reported spending (or planning to spend) approximately 40% of their COVID-19 stimulus payment (Coibion et al., 2020). There is discrepancy in the literature on the difference between intended spending and actual spending from income changes; Graziani et al. (2016) find that consumers spend more than they intended to, while Parker and Souleles (2019) estimate similar MPCs for each. The potential unreliability of the households' answers is a clear limitation for the accuracy of this approach.

The third data collection approach identifies the consumption response to income shocks in which the persistence varies by imposing covariance restrictions on consumption and income from panel data. The most widely cited example of this approach is that of Blundell et al. (2008) using Panel Study of Income Dynamics (PSID) data. They find a consumption response to transitory shocks of almost zero. More recently, Crawley and Kuchler (2023) implement a similar method using Danish registry data but find a notably larger consumption response to transitory income shocks. Crawley and Kuchler (2023) cite methodological sensitivities to assumptions about the consumption path in Blundell et al. (2008) as the reason for the variations in these estimates.

While Havranek and Sokolova (2020) concludes that there are no systemic effects of the choice of estimation technique on MPC estimation results, different techniques remain more (or less) applicable under different circumstances. For example, the use of quantile regressions (as in Broda and Parker, 2014) does not allow for estimation of the full MPC distribution. Instead, MPCs are estimated at different quantiles of the conditional distribution of the consumption response.

Estimates of the distribution of MPCs are particularly scarce (Lewis et al., 2019), but the techniques employed in the existing literature predominantly lie in two categories. The first simulates a distribution of MPCs using an estimated fully structural model. Violante and Kaplan (2014)'s two-asset model is an example of this in which many households can be classified as "wealthy hand-to-mouth", holding large amounts of illiquid wealth but little liquid wealth and therefore exhibiting a high MPC. Parameterising their model to the 2001 US economic stimulus payments allows Violante and Kaplan (2014) to show that these households are the predominant drivers of a consumption response in this model. The estimated coefficient on the rebate aligns with the evidence. The model of Carroll et al. (2017) includes a household-specific income process with a permanent and a transitory component, and preference heterogeneity to match the distribution of wealth in the US. The authors then use this to jointly estimate the wealth distribution and the MPC distribution, with many households holding little wealth and exhibiting a high MPC. An important contribution of Carroll et al. (2017) is the finding that the effects of an income shock will depend upon the distribution of the shock across households since this influences the aggregate MPC.

The second estimation technique assumes some observable characteristics and group households based on these, then estimate the MPC of each of these groups. Examples of the application of this technique include Johnson et al. (2006); Blundell et al. (2008); Parker et al. (2013); Kaplan et al. (2014); Fagereng et al. (2021); Crawley and Kuchler (2023). The specific set of characteristics used to group households and the subsequent number of groups varies, although the key observable characteristics typically used in these papers are liquid and illiquid wealth. Kaplan et al. (2014) place particular emphasis on this, putting households into three categories: the "not hand-to-mouth" hold a positive amount of liquid assets after consuming; the "poor hand-to-mouth" does not hold any liquid or illiquid assets after consuming; and the "wealthy hand-to-mouth" holds a positive amount of illiquid assets but no liquid assets after consuming. Further characteristics include income (e.g. Johnson et al. 2006), age (e.g. Parker et al. 2013), education (e.g. Crawley and Kuchler 2023), family size or composition (e.g. Blundell et al. 2008), and the size of the income change (e.g. Fagereng et al. 2021). The key weakness of both of these estimation techniques is the need to assume the sources of MPC heterogeneity ex ante, making it unlikely that all factors determining the MPC are accounted for and the full distribution is therefore not estimated.

Lewis et al. (2019) proposes the use of a Gaussian Mixture Regression to estimate the distribution of the MPC without having to make assumptions about the determinants of group membership. Households are instead grouped by their consumption responses and the MPCs estimated within these groups, allowing for estimation of the full unconditional distribution which can be driven by both observable and latent factors. In this context, GMMs are preferable to the common approach to modelling heterogeneity as unit-specific, time-invariant fixed effects. Using a fixed effects approach with a short panel dataset, such as ours, leads to poorly estimated fixed effects and an "incidental parameter" bias (Nickell, 1981). Standard fixed effects also have the undesirable assumption that unobserved heterogeneity is time-invariant (Bonhomme and Manresa, 2015). Späth (1979) develops clusterwise regressions to allow for separate regression functions and membership of a given number of clusters. Extending the work of Späth (1979), DeSarbo and Cron (1988) formulate a conditional mixture maximum likelihood estimation approach for performing clusterwise linear regression by implementing the Expectation-Maximisation algorithm of Dempster et al. (1977). The GMR approach used by Lewis et al. (2019), which we will follow, is based on this method proposed by DeSarbo and Cron (1988). Section 6 goes into greater detail on the method. The use of such models is more commonly used in fields outside of Economics, particularly in Robotics[4]. None of the studies in the meta-analysis of Havranek and Sokolova (2020) implement a switching

[4]See Calinon et al. (2007); Cederborg et al. (2010); Drews et al. (2013).

regression for estimating the MPC[5]

A number of existing papers have estimated the MPC during the COVID-19 pandemic. Baker et al. (2023) use high-frequency transaction data to estimate an MPC of 0.14 within the first week of receiving the stimulus payment and 0.25-0.30 over three months. Karger and Rajan (2020) use an separate transactional-level bank account dataset and estimate an average MPC of 0.46 in the two weeks following payment receipt, but consumption returned to normal levels after two weeks. However, they also find that 12% of consumers decreased their spending following receipt of a stimulus payment. The MPC distribution estimated by Coibion et al. (2020) using survey data for the value of consumers' self-reported spending of their stimulus payments spans the full range of values between zero and one, with many MPCs at each of the two extremes. MPC estimates for a hypothetical UK stimulus payment due to Crossley et al. (2021) are relatively modest at 0.11 on average.

# 3 Covid-19 Economic Stimulus Payments

The Coronavirus Aid, Relief, and Economic Security (CARES) Act (116th Congress of the United States, 2020b) was passed in the US in late March 2020. It provided a variety of initiatives to mitigate the economic fallout of the COVID-19 pandemic US Department of the Treasury (2022). One such initiative was direct economic assistance to US households in the form of economic impact payments[6]. This direct economic relief was extended in late December 2020 with the Consolidated Appropriations Act (116th Congress of the United States, 2020a), and in early March 2021 with the American Rescue

[5]Havranek and Sokolova (2020) do identify 81 excess sensitivity estimates from six papers using switching regressions, out of the 3127 total consumption response estimates studied from 144 papers. Note that the data collection for the meta-analysis was terminated before the publication of Lewis et al. (2019).

[6]Additional help for individuals in the CARES Act included increased and longer-term unemployment benefits, and suspended payments and interest accrual on student loans. These were smaller in scale than the stimulus payments.

Plan Act (117th Congress of the United States, 2021)[7]. These stimulus payments were mostly delivered by the Internal Revenue Service (IRS) in the form of direct deposits in recipients' bank accounts or sent out as pre-paid debit cards, although some households received them as cheques in the initial wave of payments. In aggregate the three COVID-related stimulus payments were of an unprecedented scale, totalling \$814 billion over 476 million payments (IRS, 2022).[8]

Each of the three rounds of payments was structured differently, with amounts also varying across households based on income, tax filing status, and number of dependants. As outlined by Crandall-Hollick (2021), the first round of payments included up to \$1,200 per eligible adult and \$500 per qualifying child under the age of 17. The second round provided up to \$600 per adult and \$600 per child. Payments were more extensive in the third and final round, with up to \$1,400 per adult and \$1,400 per dependent (including adult dependants this time).

The first and second payments were based upon an individual's Tax Year 2019 return, or their Tax Year 2018 return had this not been filed. For the third round of payments, eligibility was determined by an individual's Tax Year 2020 return, or their Tax Year 2019 return had this not been filed. An additional "plus-up" payment was available in the latter case once their Tax Year 2020 return had been filed if this indicated that they were eligible for a larger payment than they had originally received (IRS, 2022).

The first two rounds of payments were phased out by 5% of a household's adjusted gross income (AGI) for those exceeding a threshold of \$112,500 if an individual filed as a head of household, \$150,000 if married and filing a joint tax return (or a qualifying widow or widower), or \$75,000 for any other filing status (IRS, 2022). Similarly, the third round of payments was phased out rateably between given income levels: \$112,500 - \$120,000 for heads of households; \$150,000 - \$160,000 for those who are married and filing jointly; and \$75,000 - \$80,000 for those with any other filing status (Crandall-Hollick, 2021). The median household income in the US was \$69,560 in 2019 (Shrider et al., 2021) so these

[7]Further initiatives to provide additional support to individuals included increased child tax credit and unemployment benefits, and forgiven student loan debt.

[8]Approximately 130 million payments were made under the Economic Stimulus Act of 2008 which was the largest of its kind at the time (IRS, 2008).

thresholds were set relatively high and most households were eligible for the stimulus payments.

The COVID-related economic stimulus payments present a quasi-natural experiment, enabling estimation of the marginal propensity to consume. The IRS was instructed to make stimulus payments "as rapidly as possible" as part of the CARES Act, meaning that the stimulus payments were largely unexpected by consumers. The use of the COVID-related stimulus payments for this purpose is somewhat weakened as practical aspects of the stimulus payment roll out resulted in some non-randomness in the order which households received payments (Clark et al., 2023). For example, those who had not filed tax returns in 2018 or 2019 faced delayed stimulus payments. Households using a direct deposit received faster delivery of automatic stimulus payments since the IRS already had their information. Younger households, who tend to have lower incomes, most commonly use direct deposits. However, Clark et al. (2023) find that by the sixth week, these differences had largely disappeared.

Also, over half of recipients had received the first stimulus payment within the first week of payments, and approximately 95% received it within six weeks (Clark et al., 2023). Further, within nine weeks of the passing of the CARES Act, 160 million payments had been made to all households believed to be eligible at the time. This is in stark contrast to the timeline of the 2008 stimulus payments, in which it took 11 weeks for the first payments to be issued and 21 weeks for almost all payments to be issued. This rapid response and the relatively small window with which payments were received mitigates the effects of the aforementioned non-randomness of payments.

# 4 Data

## 4.1 The Consumer Expenditure Survey

The Consumer Expenditure survey (CEX) contains household-level data on consumer expenditures, income and household characteristics from a rotating stratified random sample of US households. This dataset is used by a number of existing studies (see Lewis et al., 2019; Parker et al., 2013) to calculate the MPC from the 2001 and 2008 US stimulus payments. Each household is surveyed four times over four consecutive quarters before it is dropped in place of a new household. Since new households enter the survey each month, our data is at a monthly frequency.

In the second quarter iteration of the 2020 CEX, questions concerning the COVID-19 economic stimulus payments were added to the survey. Households were asked whether they had received an economic stimulus payment from the government, and then to report the month in which this was received and the amount. The households' self-reported use of any stimulus payment and whether they received it via cheque or direct deposit were also included. These questions were included in subsequent survey rounds until the fourth quarter of the 2021 CEX.

Following Johnson et al. (2006), the amount of stimulus payment received in the preceding three-month period is summed for each household to give the total amount of stimulus payment received in each reference period by each household. This gives our economic stimulus payment variable, ESP. Only households which received a stimulus payment in at least one period are included in our sample. We use the 2019, 2020, 2021, and 2022 waves of the CEX interview survey to construct our dataset which runs from July 2019 to June 2022, allowing the inclusion of all four survey responses from all households which had at least on interview during the period covering the inclusion of stimulus payment questions.

Household characteristics, including the composition of the household, income and wealth measures, are also used in our dataset. Consumer expenditures of varying levels of aggregation are incorporated in line with previous literature using the CEX (such as Browning and Lusardi (1996); Parker et al. (2013); Lewis et al. (2019). These are food, nondurable expenditures, durable expenditures, and total expenditures. The durables category includes education, housing (excluding utilities), transportation (excluding public transportation and fuel), and entertainment. Nondurable expenditure encompasses food, household utilities, public transportation, fuel, personal care, tobacco, miscellaneous goods, apparel goods and services, health care, and reading materials. The food expenditure category includes food away from home, at home, and alcoholic beverages.

## 4.2 The Stringency Index

The Oxford Coronavirus Government Response Tracker (OxCGRT) project produced a stringency index which we use as a de jure measure of restrictions faced by consumers. The stringency index comprises nine metrics: school closures; workplace closures; cancellation of public events; restrictions on public gatherings; closures of public transport; stay-at-home requirements; public information campaigns; restrictions on internal movements; and international travel controls (Hale et al., 2021). A daily mean of these nine metrics is calculated, producing a stringency index taking values between 0 and 100. Stricter government responses warrant a higher value of the index. The index is constructed to mitigate the effects of variation in policy by vaccination status by weighting the value of the stringency index for restrictions placed on those who are vaccinated and on those who are unvaccinated based on the vaccination share of the population.

We aggregate the data to a monthly frequency, matching that of the consumer expenditure survey, and use the full available sample which runs from January 2020 to May 2021. The stringency index is then merged with the CEX data using each household's state and the month of their successful interview. Summary statistics for the variables used in the MPC estimation are shown in Table 1. Further details of the he resulting dataset of 4,222 households and the constructed variables are presented in Appendix A.

| Variable | Obs | Mean | Min | Max |
|---|---|---|---|---|
| Income | 4,220 | 85,694.58 | 0 | 679,083 |
| Age | 4,222 | 53.69 | 19 | 88 |
| Month of survey | 4,222 | 6.40 | 1 | 12 |
| Year of survey | 4,222 | 2020 | 2019 | 2021 |
| Stringency index | 4,222 | 48.52 | 0 | 92.56 |
| Stimulus payment | 4,222 | 969.93 | 0 | 17,000 |
| No. children | 4,222 | 2.15 | 0 | 14 |
| Change in total expenditure | 4,222 | -2,133.92 | - 119,301.70 | 118,371.70 |
| Change in durable expenditure | 4,222 | -1,215.58 | -119,301.70 | 97,281 |
| Change in nondurable expenditure | 4,222 | -918.34 | -29,614.95 | 115,170 |
| Change in food expenditure | 4,222 | -323.16 | -12,483 | 6,733.33 |
| Unemploment indicator | 4,222 | 0.31 | 0 | 1 |

**Table 1:** Summary statistics for the variables used in the estimation of the MPC.

# 5 Specification of the MPC

## 5.1 Homogeneous MPC

We take a homogeneous specification of the type used in much of the previous literature (for example, Browning and Lusardi 1996; Souleles 1999; Johnson et al. 2006; Parker et al. 2013) as a benchmark:

$$\Delta C_j = \beta' \omega_j + \lambda ESP_j + \alpha + \epsilon_j, \qquad j = 1, ..., N, \tag{1}$$

where $j = (i, t)$ represents household $i$ in period $t$ and $\Delta C_j$ is the consumption change from the previous period to the current period. The variable $ESP$ denotes the total amount of stimulus payments received by a household in the current period. The coefficient $\lambda$ describes the effects of the stimulus payment on consumption expenditures. In previous studies $\omega$ comprises a set of controls constructed from time dummies used to absorb seasonal variation in consumption and aggregate shocks, age and age squared, and change in the number of adults and the number of children within the household intended to absorb preference-driven consumption changes. We build on this specification by including the stringency index (Hale et al., 2021) as an additional control variable to account for variation in COVID restrictions between both time and US states. Since the three Acts which facilitated the stimulus payments also included other initiatives to mitigate the detrimental effects of the COVID-19 pandemic, we also incorporate controls for the number of children in a household and a dummy variable for unemployment. These are intended to absorb any consumption expenditure changes due to increased child tax credits and unemployment benefits. Despite the measures taken by the government, many households faced falling incomes during the COVID-19 pandemic. We therefore incorporate a control for changes in income between periods to absorb the consumption effects of changes in income outside of the stimulus payments.

Identification of $\lambda$ is enabled by the monthly sampling frequency of the CEX data and variation in the month in which households received their stimulus payments. We compare the change in consumption expenditures for households who received the stimulus payment with that of those who did not.

## 5.2 Heterogeneous MPC

Following Lewis et al. (2019), we generalise the homogeneous specification in Equation 1 to allow for heterogeneous consumption responses to stimulus payment receipt. We assume that households' heterogeneous consumption responses can be assigned into $G$ groups. Our heterogeneous specification is as follows:

$$\Delta C_j = \beta' \omega_j + \sum_{g=1}^{G} \mathbf{1}[j \in g](\lambda_g ESP_j + \alpha_g) + \epsilon_j, \tag{2}$$

$$j = 1, ..., N,$$
$$\forall g = 1, ..., G,$$
$$E[\epsilon_j | \omega_j, ESP_j, j \in g] = 0,$$

where the indicator function $\mathbf{1}[j \in g]$ is equal to 1 if household $i$ belongs to group $g = 1, ..., G$ in period $t$. The vector of coefficients $\{\alpha_g, \lambda_g\}$ includes heterogeneous intercepts $\alpha_g$, allowing the interpretation of $\lambda_g$ as the group-specific MPC. Heterogeneity in the MPC is therefore described by $\lambda = (\lambda_1 ... \lambda_G)'$, and combining this with the indicator function provides an approximation of the distribution of MPCs.

# 6 MPC estimation method

The MPC specification given in Equation 2 requires joint estimation of $\lambda$ and $\mathbf{1}[j \in g]$ through the assignment of households into groups. This is not a novel problem, and much of the pre-existing literature groups households by their observable characteristics. As highlighted by Lewis et al. (2019), this approach requires assumptions to be made about the determinants of the MPC a priori; hence it does not facilitate the estimation of the full MPC distribution. Instead, following Lewis et al. (2019) by grouping households based upon their latent consumption responses removes the need for such assumptions and therefore allows estimation of the full MPC distribution. We use a clustering method to fulfil this goal, assigning households into a pre-specified number of groups based solely upon their change in consumption in the period of payment receipt. This allows us to estimate the unknown parameters of the Gaussian distributions. The coefficient on the economic stimulus payment variable, interpreted as the MPC, can then be estimated for each household using an instrumental variable approach.

## 6.1 Gaussian Mixture Models

We build a Gaussian Mixture Model (GMM) to cluster the households into $G$ groups by identifying underlying patterns. As the name suggests, GMMs model the joint probability density of the data as a weighted sum of multiple normal distributions. Each cluster has its own Gaussian distribution and therefore its own mean, covariance, and weight. The weights represent the proportion of data points assigned to a given group, intended to account for the possibility of unequal sample sizes in each group and will sum to 1. Each household will be assigned to a group based on the probability that it belongs to a given Gaussian distribution.

Given this probabilistic group membership, the likelihood that household $j$ is part of group $g$ is expressed as:

$$\Pr(z_j = g|\Omega) = \frac{\pi_g N(z_j|\mu_g, \Sigma_g)}{\sum_{g=1}^{G} \pi_g N(z_j|\mu_g, \Sigma_g)} \tag{3}$$

where $\Omega = \{\mu_1, ..., \mu_G, \Sigma_1, ..., \Sigma_G, \pi_1, ..., \pi_G\}$ and $N(.)$ represents the normal (or Gaussian) distribution.

The likelihood of observing a household $j$ given that it is a member of group $g$ is:

$$\Pr(x_j|z_j = g, \mu_g, \Sigma_g) = N(x_j|\mu_g, \Sigma_g)$$

This can then be extended to give the likelihood of observing household $j$ given the mixture of Gaussian distributions:

$$L(\Omega) = \Pr(x|\Omega) = \prod_{j=1}^{J} \sum_{g=1}^{G} \pi_g N(x_j|\mu_g, \Sigma_g) \tag{4}$$

**Expectation-Maximisation algorithm** In order to find the mixture of normal distributions that most probably represents the data, we require the parameters $\Omega$ of each distribution in order to allocate our data into groups but estimating these parameters requires us to know which group each household is in. In other words, we seek to maximise the likelihood of the model parameters given by:

$$\Omega^* = \arg\max_{\Omega} L(\Omega)$$

Using the Expectation-Maximisation algorithm proposed by Dempster et al. (1977) (and applied to GMMs by Ghahramani and Jordan (1993)) provides a solution to this. The algorithm begins by initialising the set of parameters $\Omega$. The mean and covariance are initialised randomly, and weights are initially set to be equal for all clusters. The Expectation step proceeds by evaluating the likelihood that each data point comes from each distribution by calculating $\Pr(z_j = g|\Omega)$ in Equation 3 using the initial parameter estimates. These are known as the responsibilities $\gamma_{jg}$ of the model. The Maximisation step then uses these responsibilities to update the model parameters $\Omega$ to maximise the expected likelihood calculated in the previous step. Solving for the model parameters gives Equations 5 - 7 for this step.

$$\mu_g = \frac{\sum_{j=1}^{J} \gamma_{jg} x_j}{\sum_{j=1}^{J} \gamma_{jg}} \tag{5}$$

$$\Sigma_g = \frac{\sum_{j=1}^{J} \gamma_{jg}(x_j - \mu_g)(x_j - \mu_k)}{\sum_{j=1}^{J} \gamma_{jg}} \tag{6}$$

$$\pi_g = \frac{\sum_{j=1}^{J} \gamma_{jg}}{J} \tag{7}$$

We iterate between the Expectation step and the Maximisation step until the model parameters converge to their maximum likelihood estimates and the responsibilities remain unchanged from one iteration to the next.

**Model Implementation** Sci-kit learn is a Python library used for machine learning, such as implementing clustering algorithms via the sklearn.clustering module (Pedregosa et al., 2011). The sklearn.mixture package within this module fits and estimates a GMM. It assigns datapoints to clusters based on the mahalanobis distance to the centre of a cluster; updating these centres via the Expectation-Maximisation algorithm outlined above.

The choice of the number of groups to use in a clustering algorithm is particularly important, as the subsequent model and its results are highly dependent upon this choice. GMM is no exception. Too few groups will result in underfitting, while too many groups will overfit the data. In line with existing literature, we use the Bayesian Information Criterion (BIC). The BIC incorporates a penalty term for increasing parameters into its evaluation of the model's likelihood, reducing the risk of overfitting (Fraley, 1998). The optimal number of groups in the model is the one which minimises the BIC. We evaluate the BIC using the Sci-kit learn Python library for the full dataset.

A step-by-step breakdown of the algorithm used to implement this method in Python is provided in Appendix B.

## 6.2 Gaussian Mixture Instrumental Variable Regression

Having fitted the GMM to the data, we now run a regression analysis to predict the unknown stimulus payment coefficient $\lambda$. We implement two-stage least squares approach, as opposed to a linear regression, to address potential endogeneity issues arising from the value of the stimulus payment received by a household. An indicator for receipt of the payment in a given period is used as our instrument. The resulting Gaussian Mixture Instrumental Variable Regression (GMIVR) uses the GMM to estimate the unknown $\lambda$ by calculating its conditional probability distribution given the known features of the data. The first stage estimates the effect of our instrument on the potentially endogenous value of the stimulus, and the second stage estimates the effect of the instrumented variable on the change in consumption. The value of $\lambda$ is calculated as a linear combination of the means of the weights from this conditional distribution, given known features of the data in each group of the GMM (Ghahramani and Jordan, 1993).

The first stage estimates the following equation:

$$ESP_j = \alpha_g + \beta' \omega_j + \eta_j \tag{8}$$

which produces the instrumented value of $ESP$ variable, $ESP_{IV}$.

The second stage estimates the target equation, obtained by rewriting Equation 2 as:

$$\Delta C_j = \sum_{g=1}^{G} \mathbf{1}[j \in g] \psi_g^{G'} x_j + \epsilon_j, \tag{9}$$

where $x_j = \begin{bmatrix} 1 & ESP_{IV} & \omega_j' \end{bmatrix}'$. Note that we restrict the parameter values $\psi_g^{G'}$ which correspond to $\omega_j$ to be constant within a given group $g$. These are the parameters to be estimated.

The estimated group-level MPC and the household-specific weights are then used to compute the MPC for each household.

# 7 Results

We use the data described in Section 4 to estimate the distribution of the MPC via the GMIVR algorithm. Our findings reveal a notable amount of heterogeneity in the MPC and considerable variation in both the magnitude and distribution of the MPC for each of the consumption categories considered. Our estimated MPCs are smaller than those in much of the existing literature, with a material share around or below zero. We begin by estimating the distribution of the MPC for total expenditures then repeat this exercise for durable, nondurable, and food expenditures. Having estimated the MPC distribution, we study the observable determinants.

## 7.1 The Distribution of the Marginal Propensity to Consume

The results of the BIC described in Section 6.1 indicate that using six groups and assuming that the model's covariance structure is diagonal is the most appropriate parameterisation for the GMM. Appendix C contains these BIC scores. We implement the GMM, assigning each household to one of the six groups based on their consumption response. Figure 1 shows this assignment, with the colours representing each of the distinct groups of households, and plots the fitted Gaussian associated with each group. The GMM algorithm assigns households to the group whose distribution their change in consumption most likely belongs to. This is done by iterating through steps to update the groups' parameters (and hence the Gaussian distribution associated with it) and each household's group membership to find the optimal assignment.

The MPC distribution shown in Figure 2 is estimated by applying our IV approach to the GMM, as outlined in Section 6. Our results are shown in Appendix C.2 to be robust to alternative MPC specifications. The posterior weighted MPC is computed for each household using the group-specific MPCs and the household-specific weights. We observe notable heterogeneity in the MPC, with values ranging from -0.13 to 0.27, and the highest frequency of MPCs at this upper bound. In comparison to Lewis et al. (2019), the range of our MPC estimates is of a similar size but the distribution of MPCs over this range is perhaps more equal. This greater heterogeneity in our estimated MPCs is likely due to the diverse impact of the pandemic on households. The vertical red line in Figure 2 represents the average of the heterogeneous MPCs, calculated as the mean of the individual-level MPCs, which is estimated as 0.07. The vertical green line in Figure 2 represents the homogeneous MPC in Equation 1, calculated as the mean of the group-

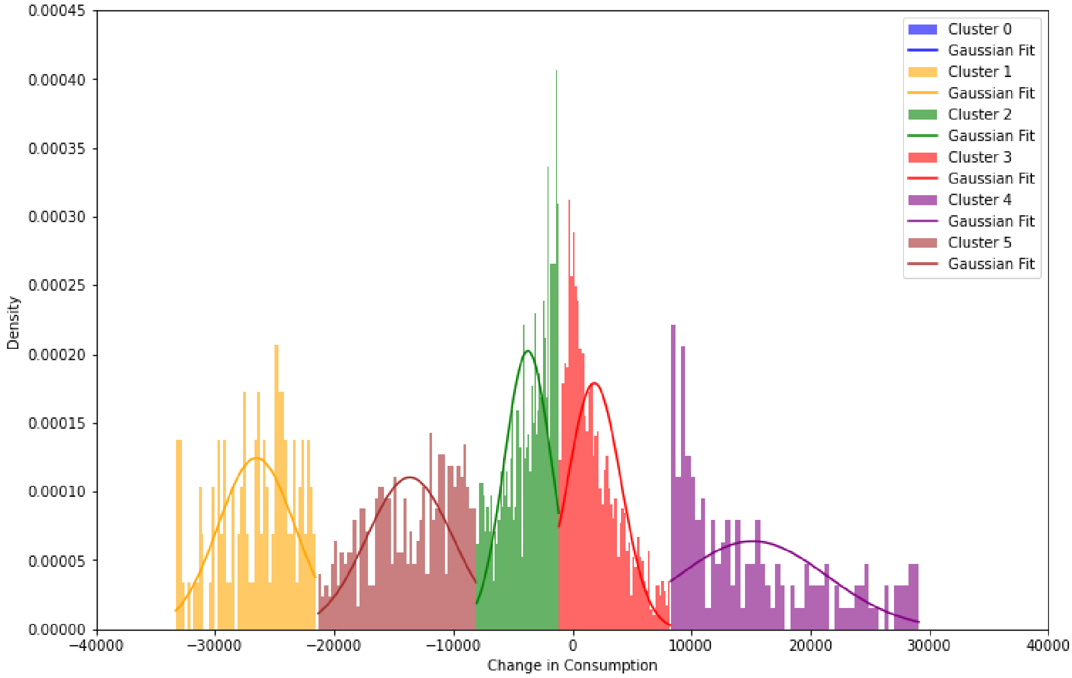


**Figure 1:** The GMM clustering and associated fitted Gaussian distribution of households' total consumption expenditure response to the receipt of the stimulus payment.

level MPCs, which is estimated as 0.02. This discrepancy between the values of the average heterogeneous MPC and the homogeneous MPC is due to the variation in group membership and their associated distributions, highlighting the importance of allowing for heterogeneity in the group-specific MPCs.

Figure 2 shows that a substantial number of households have an MPC of approximately 0, suggesting that they did not use the stimulus payment for any consumption expenditures. Furthermore, a non-negligible number of households have a negative MPC and therefore reduced their consumption upon receipt of the stimulus payment. These two observations are initially surprising, but upon further consideration there are a number of potential explanations. The first is the high uncertainty during the pandemic. Households faced regular changes in COVID restrictions and macroeconomic conditions over this period, which likely increased the precautionary savings motive for many households since they were now more uncertain about the future. As a result, households are more likely to use their stimulus payment to save (or payoff debts) rather than for consumption expenditures. The households for whom the MPC is negative present a particularly extreme case of this increased precautionary savings motive, since their con-

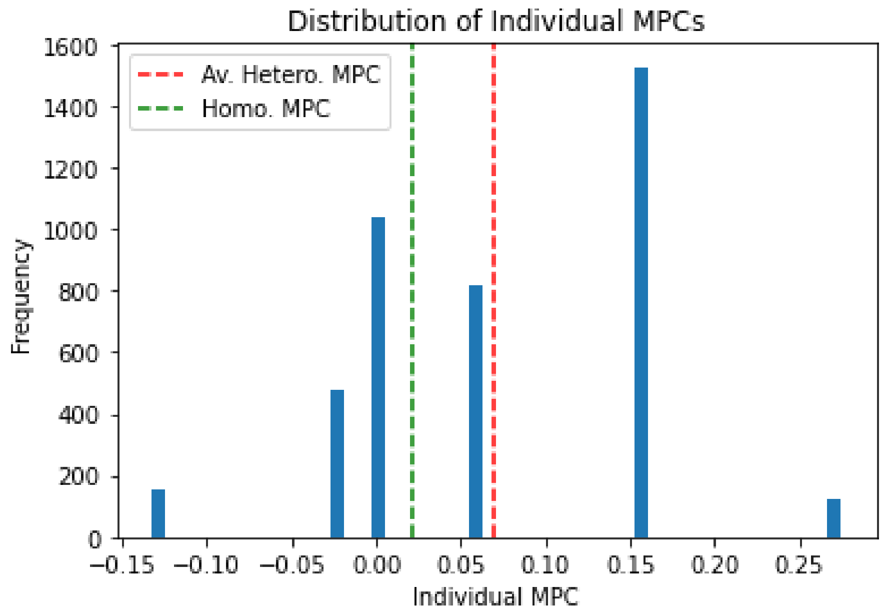


**Figure 2:** The estimated distribution of the MPC for total expenditures. The homogeneous MPC and average heterogeneous MPC are also plotted.

sumption expenditures are lowered upon receipt of the stimulus payment. In a similar vein, Carroll et al. (2021) suggests that unemployed individuals who expect their job search to be significantly hampered by the pandemic will exhibit a low MPC since they know that they might have to make their stimulus payment last longer.

The second and third potential factors influencing these observations are virus fears and households' formation of new routines and habits. Both of these factors would result in reduced consumption expenditure through people staying at home more. Individuals who were particularly fearful of contracting COVID-19 were more likely to stay at home, even when restrictions did not require them to do so, in order to reduce this risk. Similarly, many people's daily routines were significantly disrupted during COVID-19, leading them to establish new routines and habits that complied with government restrictions and were appropriate to the outlook of the pandemic. These new routines typically involved staying at home, travelling less, and visiting fewer shops or restaurants (if at all). Consequently, consumption expenditures on activities outside of the home, such as dining in a restaurant or paying for fuel to drive a car, were reduced. Both the negative MPCs and the general small value of our MPCs reflect this diminished preference for consumption.

Our estimated MPCs are closer to zero than the estimates in much of the existing literature. The mean MPC in the meta-analysis conducted by Havranek and Sokolova (2020) is 0.21 - comparable to our upper bound of 0.27. Applying the same approach as the present paper to the 2008 US stimulus payments, Lewis et al. (2019) find that all of their MPC estimates a much larger than zero, with a lower bound of approximately 0.27. Notably, our largest estimated MPC of 0.23 is equal to the average MPC of Lewis et al. (2019). Similarly, Crawley and Kuchler (2023)'s lowest estimated group-level MPC of 0.3 exceeds our highest estimate. Fagereng et al. (2021) and Olafsson and Pagel (2018) also find relatively large MPCs for Norway and Iceland respectively. Negative or zero estimates of the MPC are not unheard of however; Misra and Surico (2014) estimate that a large proportion of MPCs are negative or zero using the same data as Lewis et al. (2019). Despite being much smaller in magnitude, our finding that the average heterogeneous MPC is greater than the homogeneous MPC is in line with that of Lewis et al. (2019), although the discrepancy between our two estimates is larger.

In comparison to other papers estimating the MPC from the US COVID stimulus payments, our estimates are again at the smaller but to a slightly lesser degree. Baker et al. (2023) estimate an MPC of 0.14 within the first week of receiving the payment and 0.25-0.30 over three months. Karger and Rajan (2020) on the other hand estimate an average MPC of 0.46 in the two weeks following payment receipt, but consumption returned to normal levels after two weeks. They do however identify that 12% of consumers decreased their spending following receipt of a stimulus payment. Not only do the findings of these two studies differ from our estimates, the two of them also differ in terms of both the magnitude of MPCs and their trend in subsequent months. Unlike our MPC specification in Equation 2, neither of these papers include controls for unemployment or number of children. This is likely to have biased their MPC estimates upwards since the coefficient on the stimulus payment variable will also be absorbing effects on the consumption response due to the increased child tax credits and unemployment benefits that formed part of the three COVID economic relief Acts. Appendix C.2 presents our estimated distribution without these controls which contains slightly higher MPC estimates, highlighting the likely upward bias of these two papers. The MPC distribution estimated by Coibion et al. (2020) spans the full range of values between zero and one, with many MPCs at each of the two extremes. Almost 40% of their estimated MPCs are zero, which is comparable to our finding of approximately 30% of MPCs equalling zero. The presence of higher estimates in their distribution, and lack of negative values, is a clear contrast with our results. However, Coibion et al. (2020) use survey data for the value of consumers' self-reported spending of their stimulus payments, subsequently preventing MPCs below zero or above one. They therefore cannot account for the negative MPCs which we observe, perhaps biasing their estimates upwards. MPC estimates for a hypothetical UK stimulus payment due to Crossley et al. (2021) are relatively modest at 0.11 on average, only slightly higher than our average estimate. Their results could also be upwardly biased by their survey question in a similar way to Coibion et al. (2020).

### 7.2 The Distribution of the Marginal Propensity to Consume for Different Consumption Categories

In order to estimate the distributions of the MPC for different consumption goods, we must first check the BIC scores for each specification separately since they will each be represented by different probability distributions. The optimal number of groups according to the BIC is five with a diagonal covariance structure for durables, six with diagonal covariances for nondurables, and siz with full covariances for food (see Appendix D). Examples of durable expenditures include education and housing (excluding utilities). Nondurable expenditure examples include food and personal care items. The food expenditure category includes food away from home, at home, and alcoholic beverages. As with our MPC estimates for total expenditures, our estimates for alternative consumption categories are smaller than much of the existing literature.

Figure 1 presents the assignment of households into groups based on their consumption change (durables in Panel A, nondurables in Panel B, and food in Panel C). The change in consumption for durables has a notably larger range than that of nondurables, suggesting that the majority of the consumption response for total expenditures can be attributed to durable goods. Similarly, the distribution of the consumption change for durables is comparable to that of total expenditures, further supporting this conclusion.

The distribution of the MPC for each of the three alternative consumption specifications are presented in Figure 4. A high frequency of MPCs at or very close to zero

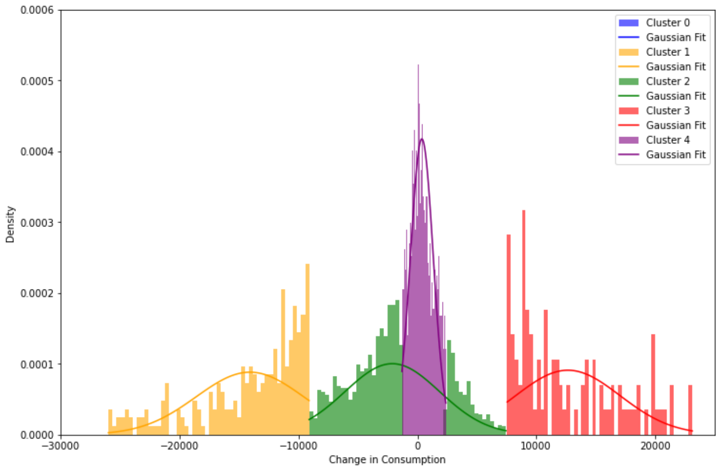


**(a)** Durable Expenditures

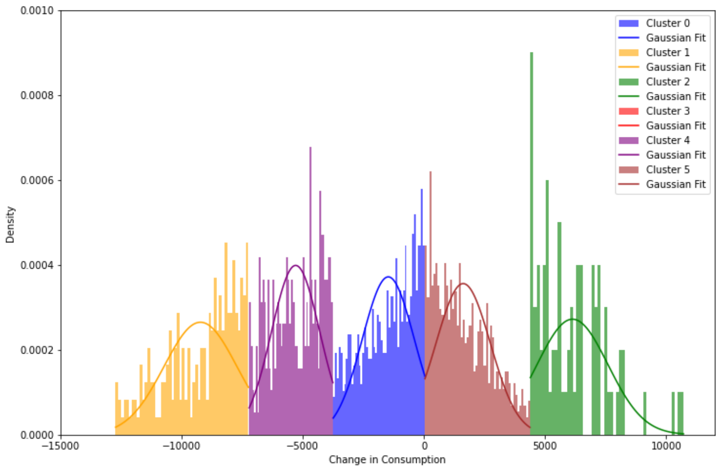


**(b)** Nondurable Expenditures

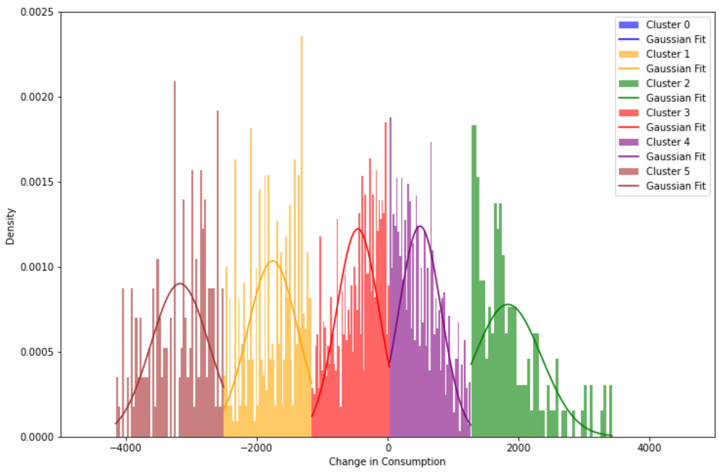


**(c)** Food Expenditures

**Figure 3:** The GMM clustering and associated fitted Gaussian distributions of households' consumption responses to the receipt of the stimulus payment.

are present for each consumption category, but particularly for nondurables. We observe less heterogeneity in the MPCs for durable, nondurable and food expenditures than is present for total expenditures, with both upper and lower bounds of each being closer to zero. Panels A & B show a smaller frequency of negative MPCs for durables and nondurables than total expenditures, and these values are far less negative. Interestingly, we do not find any negative MPC estimates for food expenditures in Panel C, suggesting that households spent at least as much on food when they received the stimulus payment as they had done previously. However, the average MPC for food expenditures is notably smaller than that of the other consumption specifications, implying a smaller adjustment in food expenditures. This is a logical finding given that food is a component of nondurables and is therefore a much smaller consumption category with less scope for spending. The average MPC of approximately 0.07 for durables is equal to the average MPC for total expenditures, and it is smaller for nondurables at approximately 0.03. This finding suggests that the change in consumer spending on durables upon receipt of the stimulus payments was the greatest.

Findings on relative spending on durables and nondurables are mixed in the literature. Similarly to our results, Parker et al. (2013) for example find that households on average increased food expenditures by about 2% of the 2008 stimulus payment, nondurables by 8% of the payment, and durables by 12% of the payment. These estimates are higher than ours which is to be expected since they estimate an increase in total expendituress of 52% of the stimulus payment which is much larger than our equivalent estimates. On the other hand, Baker et al. (2023) estimate that households spent more on nondurables during the pandemic, although they note that this finding is at odds with previous papers. This discrepancy in results may be driven by different relative spending habits of constrained and unconstrained households on durable and nondurable goods (Souleles, 1999).

### 7.3 Drivers of MPC Heterogeneity

We now turn to the drivers of our estimated MPC distribution. Since we estimate the MPCs without taking an ex ante stance on the characteristics which determine the dis-

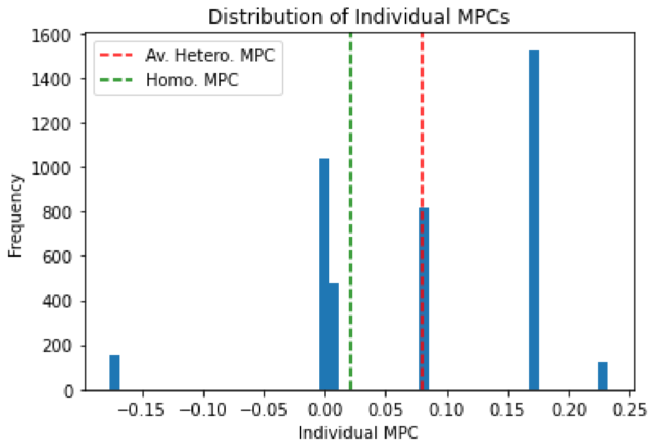


**(a)** Durable Expenditures

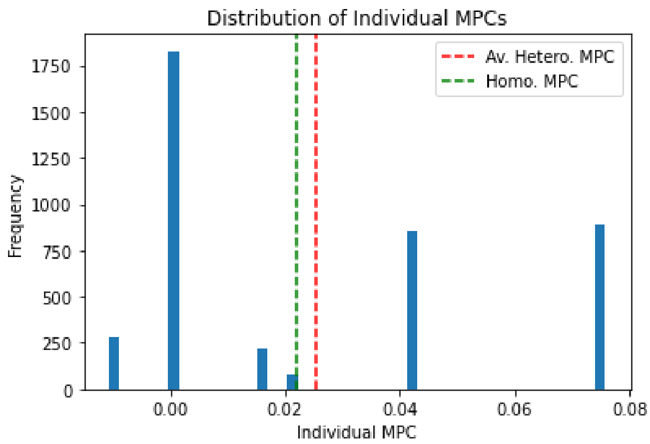


**(b)** Nondurable Expenditures

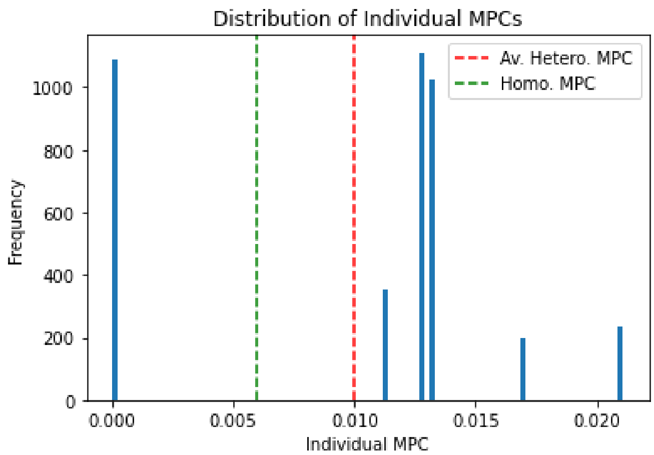


**(c)** Food Expenditures

**Figure 4:** The estimated distribution of the MPC for durable, nondurable, and food expenditures. The homogeneous MPC and average heterogeneous MPCs are also plotted.

tribution, we can study these relationships ex post. Table 2 presents our findings that a number of observable household characteristics have statistically significant correlations with our estimated MPCs.

| | Correlation Coefficient |
|---|---|
| Liquid wealth | -0.0307* |
| Total household income | -0.1170*** |
| Salary income | -0.0614** |
| Size | -0.0721*** |
| Unemployment | 0.0466*** |
| Education | -0.0724*** |
| Mortgage value | -0.0440** |
| Housing value | -0.2490* |
| Renters | 0.1065*** |

**Table 2:** The correlations between MPC estimates for total expenditures and observable household characteristics. Note *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels respectively.

Perhaps the most surprising result is the particularly small coefficient on liquid wealth of -0.03, implying a weaker relationship than is reported in much of the existing literature[9]. This could be due to the observation that consumers made fewer large purchases during the pandemic (Coibion et al., 2020) which they would have used savings for, and added their stimulus payments to these savings instead. The low response rate for liquid wealth and the associated potential non-response bias also likely contributes to this weaker estimated relationship. The sign of this coefficient however, is consistent with other papers including Karger and Rajan (2020) using COVID data. Households with little to no liquid wealth, defined as "hand-to-mouth" by Kaplan et al. (2014), consume all of their income and therefore have a higher MPC.

Household income and salary income both have negative correlations with the

[9]Baker et al. (2023) find that liquidity has the strongest relationship with the MPC for example.

MPC at -0.12 and -0.6 respectively, a finding which is consistent with existing literature (see Parker et al., 2011; Baker et al., 2023). Households with higher income will generally have higher liquid wealth (see Appendix E), lowering their MPC. Further, the stimulus payments will be smaller relative to income for those with higher incomes so will present a smaller income shock and therefore induce a smaller consumption response. This effect is further illustrated by the positive coefficient of 0.05 on unemployment since unemployed consumers typically have lower incomes (see Appendix E) so will exhibit higher MPCs. The lower MPC associated with higher education levels likely also works through this income mechanism since consumers with higher levels of education typically have higher incomes (see Appendix E), consistent with Coibion et al. (2020). We find no statistically significant relationship between nonsalary income and the MPC, a finding which is at odds with Lewis et al. (2019). As with liquid wealth, our data has a low response rate for nonsalary income so this relationship is also likely subject to non-response bias.

We find significant relationships between the MPC and a number of housing variables. The negative coefficient of -0.04 on the value of the household's mortgage (if they have one) implies a lower MPC for those with a higher mortgage, a finding also presented by Coibion et al. (2020). This is also likely related to the income mechanism since income and mortgage value are positively correlated (Appendix E). Mortgages are the only form of loan or credit found to have a statistically significant relationship with the MPC, likely because of the low response rate for these variables. The value of housing is strongly correlated with the value of a household's mortgage (Appendix E) which is a likely driver of the negative correlation between housing value and the MPC. Renting of housing is associated with a higher MPC, although this is unlikely to be related to the income and liquidity mechanisms discussed above since neither of these variables have a statistically significant relationship with renting. Those that own their housing are more likely to have a mortgage and therefore have incentives to use their stimulus payments to overpay their mortgage, which is not classed as an expenditure. Renters who put their stimulus payments towards housing costs, on the other hand, would be using it for an expenditure. This is consistent with Crossley et al. (2021).

Household size and the MPC are weakly negatively correlated, with a coefficient of -0.07. Larger households with more consumers received larger total stimulus payments, but this suggests that household expenditures do not scale up proportionally to the size of the household (Deaton and Paxson, 1998; Logan, 2011).

# 8 Next Steps for Future Research

## 8.1 MPC Estimation Approach

This study presents a number of avenues for future research. First, a number of improvements could be made to our approach. A key weakness of clustering methods, such as GMMs, is the need to pre-specify the number of clusters in the model which can affect results. Selecting the optimal model parameterisation according to the BIC scores is a useful, but imperfect solution. Variational Bayesian Gaussian Mixture Modelling is a Bayesian approach to clustering which extends GMM by incorporating Bayesian techniques for model selection and parameter estimation (Pedregosa et al., 2011). The number of clusters is not fixed a priori; instead the optimal number of clusters are automatically determined from the data using selection techniques such as the Dirichlet Process prior. Due to its high computational intensity, this method was not selected for use here, although it could provide an interesting comparison in future work.

Extending the consumption categories to include a wider range of consumption aggregations would shed greater light on consumer consumption behaviour. For example, automobile sales and the hospitality sector were amongst the industries most impacted by COVID-19 (Vidovic, 2022). Extracting consumption response data specific to these sectors and estimating their MPC distributions would provide further information on the drivers of heterogeneity in the MPC and consumer spending patterns. Similarly, analysis of the group membership assigned by the GMM would aid our understanding of the characteristics of the households within each group and their MPC.

A third improvement to the design of our study would be to separate our data into three datasets, one for each round of stimulus payments. This would enable an investigation into the effects of whether the payments are truly unexpected. The present paper combines the data for the whole pandemic into one dataset, potentially masking any differences in the consumption response to the three rounds of payments. It could be argued that the second and third stimulus payments, although paid swiftly after they were announced, were not entirely unexpected. Gaining an understanding of the extent of this would be highly valuable for evaluating their use in the estimation of the MPC. Estimating the MPC distributions of the three rounds of stimulus payments individually would also allow for a study of the evolution of the pandemic and the associated economic downturn.

## 8.2 Empirical Extensions

We also identify a number of empirical extensions to our study. The first would be to replicate Lewis et al. (2019)'s MPC estimation using the 2008 stimulus payments and then extend the methodology to the 2001 stimulus payments. This would allow for a comparison of consumer behaviour and the efficacy of stimulus payments in relation to three different shocks.

Our study uses a de jure measure of the restrictions on consumer spending during the pandemic in the form of the stringency index. Incorporating google mobility trends data as a de facto measure would complement our findings. A number of papers have used google mobility trends data to study the impacts of COVID-19, for example Bahaj et al. (2022) find links between business creation during the pandemic and the decline in retail footfall.

A third empirical extension would be to investigate any potential delayed impact of the stimulus payments. While estimates of the longer term spending effects of stimulus payments do exist, the COVID stimulus payments are unique in that they were disbursed at a time when governments had imposed many restrictions. It would be interesting to explore whether the imposing and then lifting of these restrictions lead to the rebate having a delayed impact, and whether this could have contributed to the persistent inflation experienced post-pandemic.

## 8.3 Economic Modelling Applications

Finally, from a modelling perspective, our results can be used to inform the development of Heterogeneous Agent models. Our findings on the sources of MPC heterogeneity, such as housing ownership, could be used to incorporate richer heterogeneity into existing models. For example, modelling households by grouping them by their housing ownership status, in a similar way to the approach to wealth liquidity of Kaplan et al. (2018), could provide interesting extensions to the literature studying housing markets in New Keynesian models (see Iacoviello, 2005). The aforementioned analysis of group membership assignments in the GMM would also be useful to inform such a model.

# 9 Conclusion

This paper studies the distribution of the MPC from the economic stimulus payments provided by the US government in response to the COVID-19 pandemic. Our key contribution hinges on the use of a clustering method to group households based on the change in their consumption. This method is preferable to those employed by much of the existing literature since it allows the full unconditional distribution of the MPC to be estimated, rather than estimating how the MPC varies with a set of observable household characteristics. In our analysis the MPC distribution is estimated via an instrumental variable approach, taking receipt of the rebate as our instrument, to address potential endogeneity concerns. Controls for non-stimulus changes in income, number of children, unemployment, and government restrictions are used to mitigate their effects on our MPC estimates from the stimulus payments.

Using household-level survey data, we find significant heterogeneity in the MPC with a distribution spanning values from -0.13 to 0.27 for total expenditures. Our estimates are of smaller MPCs relative to much of the existing literature, with an average MPC for total expenditures of 0.07. We find a non-negligible share of MPCs at or below zero, suggesting a limited impact of the stimulus payments. Negative MPC values have been observed by other studies, particularly those focused on the economic impacts of COVID-19, and likely reflect the heightened uncertainty, virus fears, and routine formation during the pandemic. Our estimated MPCs for both durable and nondurable expenditures exhibit similar heterogeneity and are also comparatively small, with durable expenditures showing the larger MPCs. The magnitude and heterogeneity of the MPC for food expenditures are considerably smaller than those of the other consumption categories considered, implying a limited effect of the stimulus payments on food expenditures.

We then study the drivers of MPC heterogeneity ex ante. We find statistically significant relationships between the MPC and a number of observable household characteristics, including household income, education levels, liquid wealth, and housing ownership status. The coefficient on liquid wealth is smaller than anticipated when compared to existing literature, likely due to the low response rate for this variable in the CEX survey. Our findings have implications for both policymaking and economic modelling, particularly for the targeting of stimulus payments and for heterogeneous agent models respectively. That our MPC estimates are smaller than many others highlights the difference in consumer behaviour depending on the shock and the wider economic environment. Greater emphasis should therefore be placed on this when assessing economic policy.

The opportunities for future research related to this paper are numerous. Improvements and adaptations to the estimation approach, such as using a Variational Bayesian Gaussian Mixture Model in place of a Gaussian Mixture Model, would be welcome additions. Investigating the effects of varying levels of aggregation for the three rounds of stimulus payments and the consumption categories studied here would shed greater light on consumer spending patterns. Finally, our study provides useful results for disciplining heterogeneous agent models.

# References


110th Congress of the United States (2008). Economic Stimulus Act of 2008.

116th Congress of the United States (2020a). Consolidated Appropriations Act, 2021.

116th Congress of the United States (2020b). Coronavirus Aid, Relief, and Economic Security Act of 2020.

117th Congress of the United States (2021). American Rescue Plan Act of 2021.

Agarwal, S., Liu, C., and Souleles, N. (2007). The Reaction of Consumer Spending and Debt to Tax Rebates—Evidence from Consumer Credit Data. *Journal of Political Economy*, 115(6):986–1019.

Aiyagari, S. R. (1994). Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684.

Auclert, A. (2019). Monetary Policy and the Redistribution Channel. *American Economic Review*, 109(6):2333–2367.

Bahaj, S., Piton, S., and Savagar, A. (2022). Business creation during Covid-19.

Baker, S. R., Farrokhnia, R. A., Meyer, S., Pagel, M., and Yannelis, C. (2023). Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments. *Review of Finance*.

Bewley, T. (1983). A Difficulty with the Optimum Quantity of Money. *Econometrica*, 51(5):1485.

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

Bonhomme, S. and Manresa, E. (2015). Grouped Patterns of Heterogeneity in Panel Data. *Econometrica*, 83(3):1147–1184.

Broda, C. and Parker, J. A. (2014). The Economic Stimulus Payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics*, 68:S20–S36.

Browning, M. and Lusardi, A. (1996). Household Saving: Micro Theories and Micro Facts. *Journal of Economic Literature*, 34(4):1797–1855.

Calinon, S., Guenter, F., and Billard, A. (2007). On Learning, Representing, and Generalizing a Task in a Humanoid Robot. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, 37(2):286–298.

Carroll, C., Slacalek, J., Tokuoka, K., and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.

Carroll, C. D., Crawley, E., Slacalek, J., and White, M. N. (2021). Modelling the Consumption Response to the CARES Act. *International Journal of Central Banking*, 17(1):107–141.

Cederborg, T., Ming Li, Baranes, A., and Oudeyer, P.-Y. (2010). Incremental local online Gaussian Mixture Regression for imitation learning of multiple tasks. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 267–274. IEEE.

Christelis, D., Georgarakos, D., Jappelli, T., Pistaferri, L., and van Rooij, M. (2019). Asymmetric Consumption Effects of Transitory Income Shocks. *The Economic Journal*, 129(622):2322–2341.

Clark, L. R., Cole, A. J., Eng, A., Meiselman, B. S., Pharris-Ciurej, N., Pierce, K., and Voorheis, J. (2023). The Demographics of the Recipients of the First Economic Impact Payment.

Coibion, O., Gorodnichenko, Y., and Weber, M. (2020). How Did U.S. Consumers Use Their Stimulus Payments? Technical report, National Bureau of Economic Research, Cambridge, MA.

Crandall-Hollick, M. L. (2021). COVID-19 and Direct Payments: Comparison of First and Second Round of "Stimulus Checks" to the Third Round in the American Rescue Plan Act of 2021 (ARPA; P.L. 117-2).

Crawley, E. and Kuchler, A. (2023). Consumption Heterogeneity: Micro Drivers and Macro Implications. *American Economic Journal: Macroeconomics*, 15(1):314–341.

Crossley, T. F., Fisher, P., Levell, P., and Low, H. (2021). MPCs in an economic crisis: Spending, saving and private transfers. *Journal of Public Economics Plus*, 2:100005.

Deaton, A. and Paxson, C. (1998). Economies of Scale, Household Size, and the Demand for Food. *Journal of Political Economy*, 106(5):897–930.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum Likelihood from Incomplete Data Via the EM Algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22.

DeSarbo, W. S. and Cron, W. L. (1988). A maximum likelihood methodology for clusterwise linear regression. *Journal of Classification*, 5(2):249–282.

Drews, P., Núñez, P., Rocha, R. P., Campos, M., and Dias, J. (2013). Novelty detection and segmentation based on Gaussian mixture models: A case study in 3D robotic laser mapping. *Robotics and Autonomous Systems*, 61(12):1696–1709.

Fagereng, A., Holm, M. B., and Natvik, G. J. (2021). MPC Heterogeneity and Household Balance Sheets. *American Economic Journal: Macroeconomics*, 13(4):1–54.

Fraley, C. (1998). How Many Clusters? Which Clustering Method? Answers Via Model-Based Cluster Analysis. *The Computer Journal*, 41(8):578–588.

Fuster, A., Kaplan, G., and Zafar, B. (2021). What Would You Do with $500? Spending Responses to Gains, Losses, News, and Loans. *The Review of Economic Studies*, 88(4):1760–1795.

Ghahramani, Z. and Jordan, M. I. (1993). Supervised learning from incomplete data via an EM approach. In Cowan, J., Tesauro, G., and Alspector, J., editors, *Advances in Neural Information Processing Systems*, volume 6, pages 120–127. Morgan Kaufmann Publishers Inc., Cambridge, MA.

Graziani, G., van der Klaauw, W., and Zafar, B. (2016). Workers' Spending Response to the 2011 Payroll Tax Cuts. *American Economic Journal: Economic Policy*, 8(4):124–159.

Hale, T., Angrist, N., Goldszmidt, R., Kira, B., Petherick, A., Phillips, T., Webster, S., Cameron-Blake, E., Hallas, L., Majumdar, S., and Tatlow, H. (2021). A global panel database of pandemic policies (Oxford COVID-19 Government Response Tracker). *Nature Human Behaviour*, 5(4):529–538.

Havranek, T. and Sokolova, A. (2020). Do consumers really follow a rule of thumb? Three thousand estimates from 144 studies say "probably not". *Review of Economic Dynamics*, 35:97–122.

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

Iacoviello, M. (2005). House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle. *American Economic Review*, 95(3):739–764.

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

IRS (2008). Facts about the 2008 Stimulus Payments. Technical report, Internal Revenue Service Media Relations Office, Washington DC.

IRS (2022). SOI Tax Stats - Coronavirus Aid, Relief, and Economic Security Act (CARES Act) Statistics.

Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, 96(5):1589–1610.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary Policy According to HANK. *American Economic Review*, 108(3):697–743.

Kaplan, G., Violante, G., and Weidner, J. (2014). The Wealthy Hand-to-Mouth.

Karger, E. and Rajan, A. (2020). Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Payments.

Lewis, D. J., Melcangi, D., and Pilossoph, L. (2019). Latent Heterogeneity in the Marginal Propensity to Consume. *SSRN Electronic Journal*.

Logan, T. D. (2011). Economies of Scale in the Household: Puzzles and Patterns from the American Past. *Economic Inquiry*, 49(4):1008–1028.

Misra, K. and Surico, P. (2014). Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs. *American Economic Journal: Macroeconomics*, 6(4):84–106.

Nickell, S. (1981). Biases in Dynamic Models with Fixed Effects. *Econometrica*, 49(6):1417.

Olafsson, A. and Pagel, M. (2018). The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software. *The Review of Financial Studies*, 31(11):4398–4446.

Parker, J., Souleles, N., Johnson, D., and McClelland, R. (2011). Consumer Spending and the Economic Stimulus Payments of 2008. Technical report, National Bureau of Economic Research, Cambridge, MA.

Parker, J. A. and Souleles, N. S. (2019). Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates. *American Economic Review: Insights*, 1(3):273–290.

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–2553.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.

Shapiro, M. D. and Slemrod, J. (2003). Consumer Response to Tax Rebates. *American Economic Review*, 93(1):381–396.

Shrider, E. A., Kollar, M., Chen, F., and Semega, J. (2021). Income and Poverty in the United States: 2020.

Souleles, N. S. (1999). The Response of Household Consumption to Income Tax Refunds. *The American Economic Review*, 89(4):947–958.

Späth, H. (1979). Algorithm 39 Clusterwise linear regression. *Computing*, 22(4):367–373.

US Department of the Treasury (2022). Economic Impact Payments. Technical report, United States Government, Washington DC.

Vidovic, L. (2022). Industries Most and Least Impacted by COVID-19 from a Probability of Default Perspective - January 2022 Update. Technical report, S&P Global Market Intelligence.

Violante, G. L. and Kaplan, G. (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica*, 82(4):1199–1239.

# A Description of dataset and variable construction

| Variable | Source & Formula | Description |
|---|---|---|
| ID | CEX FMLI: NEWID | Unique ID for consumer unit and the number of their interview |
| ESP | CEX CNT: CONTCODE = 800 | Amount of rebate received since the reference month |
| EXPENDITURE | CEX FMLI: TOTEXPCQ | Total expenditures this quarter |
| EXPENDITUREP | CEX FMLI: TOTEXPPQ | Total expenditures last quarter |
| NONDURABLES | CEX FMLI: TOTEXPCQ - DURABLES | Total expenditures on non-durable goods this quarter |
| DURABLES | CEX FMLI: HOUSCQ + TRANSCQ + ENTERTCQ + EDUCACQ | Total expenditures on durable goods this quarter |
| NONDURABLESP | CEX FMLI: TOTEXPPQ - DURABLESP | Total expenditures on non-durable goods last quarter |
| DURABLESP | CEX FMLI: HOUSPQ + TRANSPQ + ENTERTPQ + EDUCAPQ | Total expenditures on durable goods last quarter |
| FOODC | CEX FMLI: FOODCQ | Total expenditures on food this quarter |
| FOODP | CEX FMLI: FOODPQ | Total expenditures on food last quarter |
| SIZE | CEX FMLI: FAM_SIZE | Number of members of CU |
| AGE | CEX MEMI: (AGE + AGE_REF)/2 | (Average) age of the head of the household (and spouse) |
| CHILDREN | CEX MEMI: Calculated from AGE using indicator | Number of members of the household aged under 18 years old |
| MORTGAGE | CEX MOR: ORGMRTX | Value of mortgage currently held |
| OWNHOUSE | CEX RNT: OWNED | Does the CU own the housing (mortgage or outright)? |

| Variable | Source & Formula | Description |
|---|---|---|
| VALUE | CEX OPB: OWN_PURX | Total price of property including land and construction costs |
| EDUCATION | CEX FMLI: EDUC_REF | Highest level of schooling completed |
| EDUCATION | CEX FMLI: EDUCA2 | Highest level of schooling completed by spouse |
| SALARY | CEX MEMI: SALARYX | Pre-tax income from salary/ wages over the last 12 months |
| INCOME | CEX FMLI: FINCBTAX | Total CU income before tax over the last 12 months |
| NONSALARY | CEX MEMI/ FMLI: FINCBTAX - SALARYX | Income from non-salary sources over the last 12 months |
| LIQUID | CEX FMLI: LIQUIDX | Value of liquid assets currently |
| LIQUIDP | CEX FMLI: LIQUIDYRX | Value of liquid assets one year ago |
| ILLIQUID | CEX FMLI: IRAX + STOCKX + OTHASTX | Value of illiquid assets currently |
| ILLIQUIDP | CEX FMLI: IRAYRX + STOCKYRX + OTHSTYRX | Value of illiquid assets one year ago |
| CREDIT | CEX FMLI: CREDITX | Amount owed on credit cards currently |
| CREDITP | CEX FMLI: CREDTYRX | Amount owed on credit cards one year ago |
| STUDENT | CEX FMLI: STUDNTX | Amount owed in student loans currently |
| STUDENTP | CEX FMLI: STDNTYRX | Amount owed in student loans on year ago |
| OTHERLOAN | CEX FMLI: OTHLONX | Amount owed in other loans currently |
| OTHERLOANP | CEX FMLI: OTHLNYRX | Amount owed in other loans one year ago |

| Variable | Source & Formula | Description |
|---|---|---|
| DEBT | CEX FMLI: CREDITX + STUDNTX + OTHLONX | Amount of debt currently |
| DEBTP | CEX FMLI: CREDTYRX + STD-NTX + OTHLNYRX | Amount of debt one year ago |
| STATE | CEX FMLI: STATE | US State which the CU is in |
| STRINGENCY | OxCGRTUS: Aggregated monthly stringency by state | COVID-19 Stringency Index by state |

**Table 3:** Description of Variables

# B Implementation of GMIVR in Python

1. Calculate the BIC scores to select the optimal parameterisation of the GMM.
2. Fit a Gaussian Mixture Model with a specified number of groups to the variable for change in consumption.
3. Predict group labels for each data point using the trained GMM.
4. Calculate group-level MPCs using Ordinary Least Squares (OLS) for each GMM group.
5. Calculate individual-level MPCs using Two-Stage Least Squares (2SLS) for each data point. Loop through each household in the dataset:
   (a) Extract the ESP value for each
   (b) Determine the GMM group label for the current household's consumption change using the pre-fitted Gaussian Mixture Model
   (c) Estimate MPCs using instrumental variables

# C The Distribution of the MPC: Additional Results for Total Expenditures

## C.1 BIC scores

| Number of components | Type of covariance | BIC score |
|---|---|---|
| 6 | diag | 56985.209456 |
| 6 | full | 57715.745407 |
| 7 | full | 57806.998276 |
| 8 | full | 58723.391186 |
| 4 | full | 60498.671976 |

**Table 4:** BIC scores for the top five parameterisations of the GMM for total expenditures.

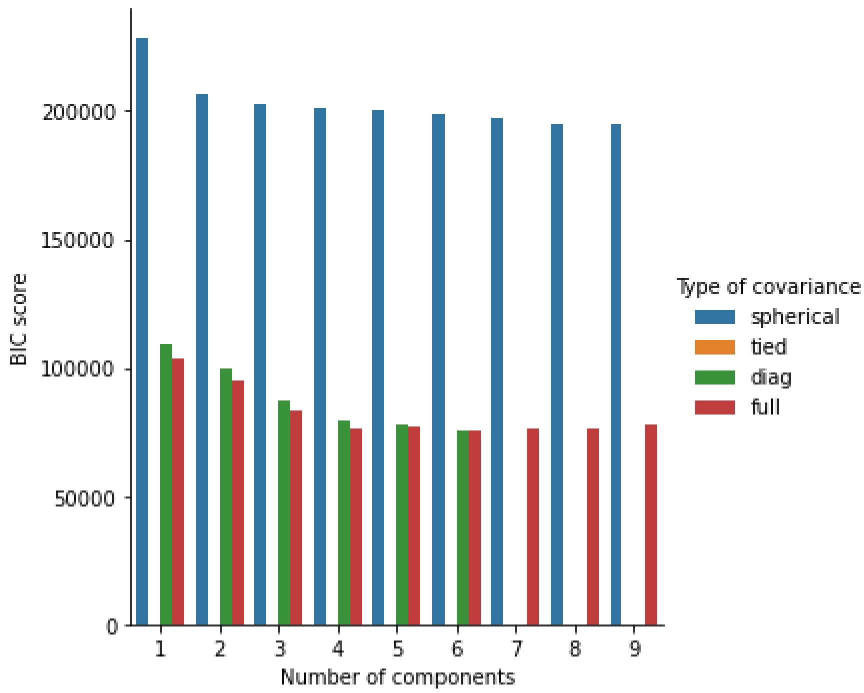


**Figure 5:** BIC scores for the top five parameterisations of the GMM for total expenditures.

## C.2 Alternative specifications

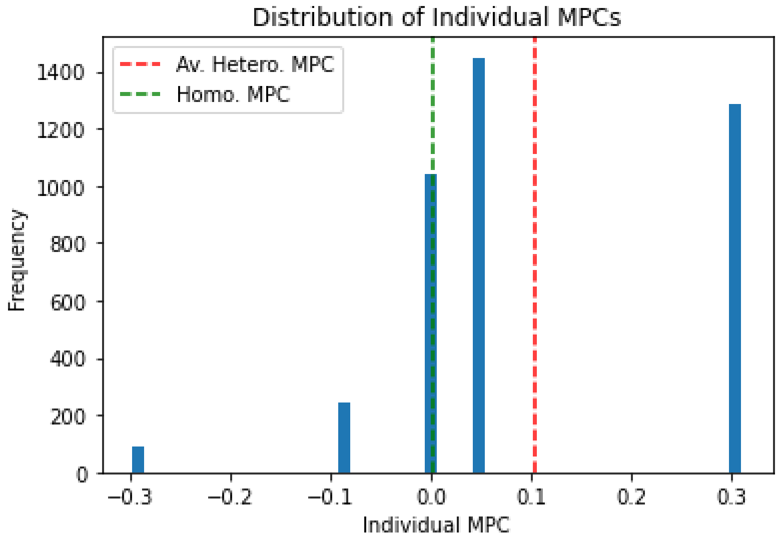


**Figure 6:** MPC distribution for total expenditures without controls for unemployment or number of children.

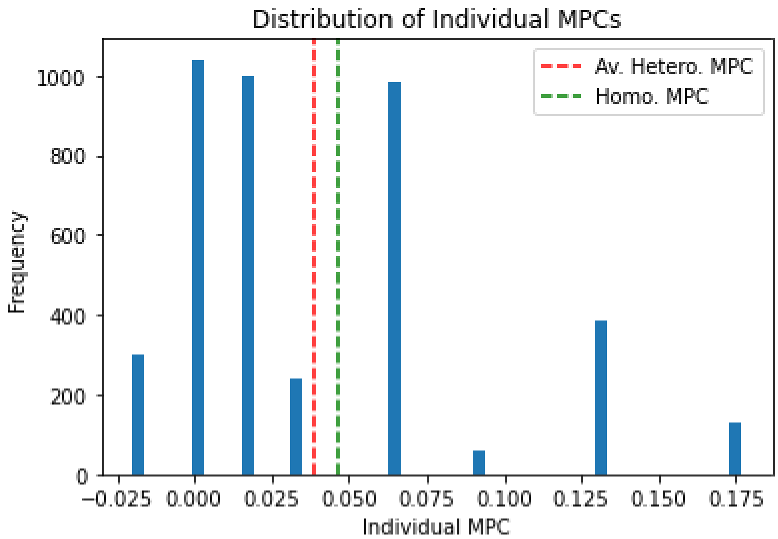


**Figure 7:** MPC distribution for total expenditures without controls for the stringency of government restrictions.

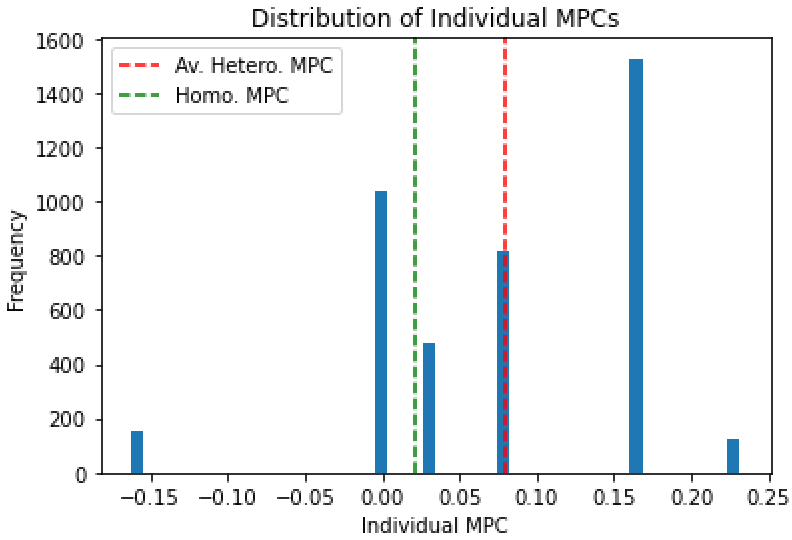


**Figure 8:** MPC distribution for total expenditures without controls for income change.

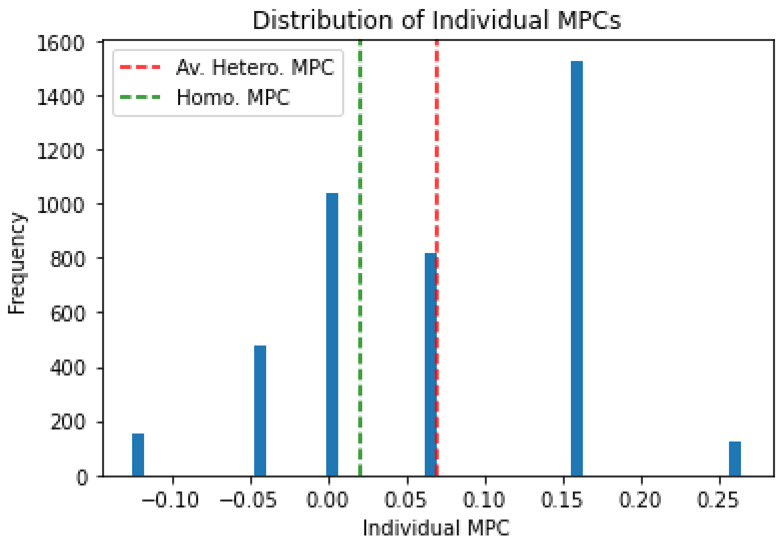


**Figure 9:** MPC distribution for total expenditures without controls for unemployment, number of children, stringency of restriction, or income change. Note that this results in the specification of Lewis et al. (2019).

# D The Distribution of the MPC: Additional Results for Durable, Nondurable, and Food Expenditures

## D.1 BIC scores for durable expenditures

| Number of components | Type of covariance | BIC score |
|---|---|---|
| 5 | diag | 57517.796593 |
| 5 | full | 57520.268234 |
| 6 | full | 57529.333857 |
| 8 | full | 58022.341832 |
| 4 | full | 59452.381583 |

**Table 5:** BIC scores for the top five parameterisations of the GMM for durable expenditures.

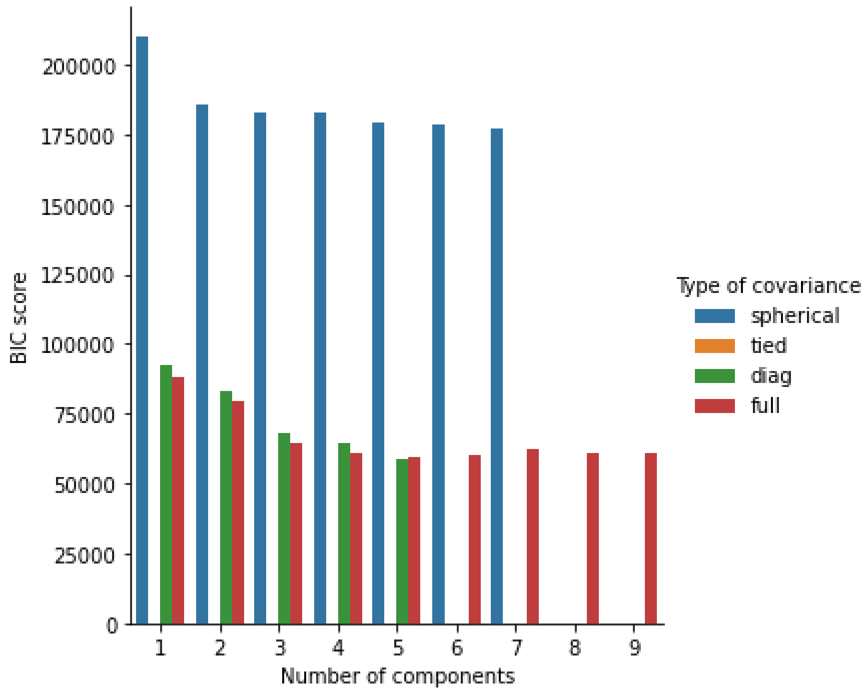


**Figure 10:** BIC scores for the top five parameterisations of the GMM for durable expenditures.

## D.2 BIC scores for nondurable expenditures

| Number of components | Type of covariance | BIC score |
|---|---|---|
| 6 | diag | 55212.501892 |
| 7 | full | 57427.739380 |
| 4 | full | 58293.102215 |
| 5 | diag | 58371.512180 |
| 9 | full | 58767.574306 |

**Table 6:** BIC scores for the top five parameterisations of the GMM for nondurable expenditures.

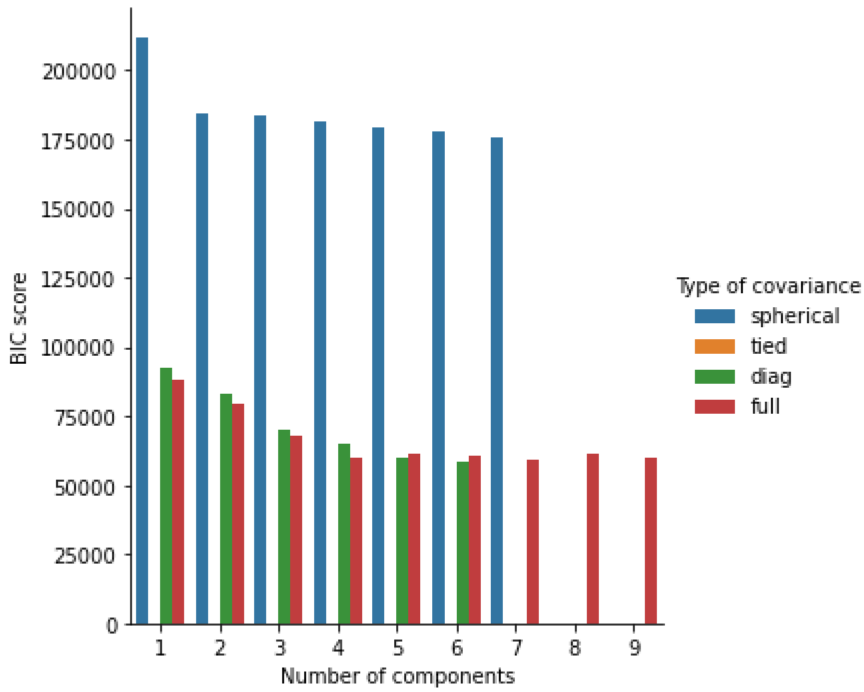


**Figure 11:** BIC scores for the top five parameterisations of the GMM for nondurable expenditures.

## D.3 BIC scores for food expenditures

| Number of components | Type of covariance | BIC score |
|---|---|---|
| 6 | full | 53390.340792 |
| 5 | full | 54016.217306 |
| 5 | diag | 54891.573714 |
| 6 | diag | 56129.442053 |
| 4 | full | 56806.138279 |

**Table 7:** BIC scores for the top five parameterisations of the GMM for food expenditures.

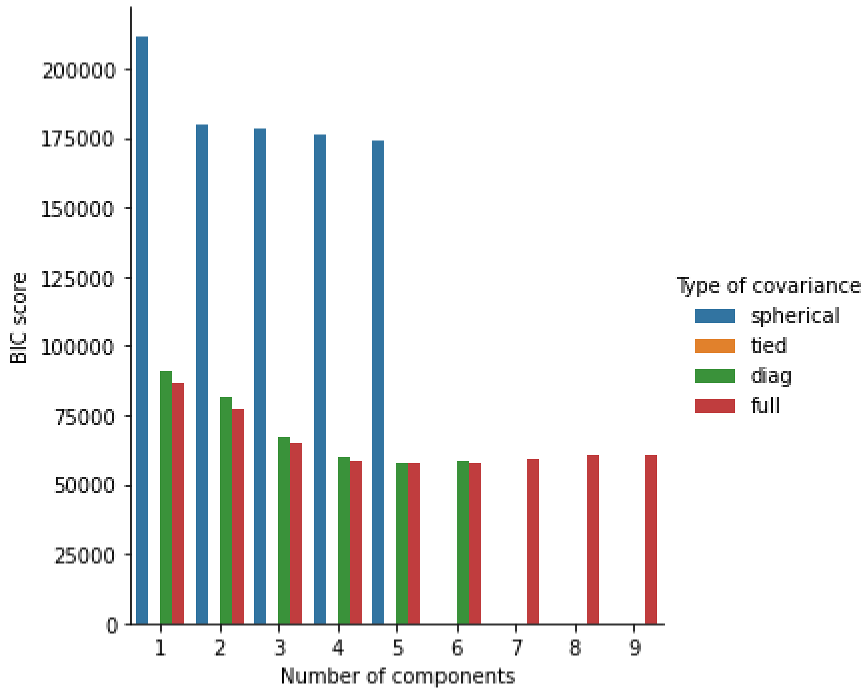


**Figure 12:** BIC scores for the top five parameterisations of the GMM for food expenditures.

# E Drivers of the MPC distribution: Additional correlations

| | Correlation coefficient with Income |
|---|---|
| Liquid wealth | 0.0388* |
| Unemployment | -0.3399*** |
| Mortgage value | 0.2921*** |

**Table 8:** The correlations between total household income and liquid wealth, unemployment, and mortgage value. Note *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels respectively.