WORKING PAPER

HOW MUCH DOES ANTICIPATION MATTER FOR MPCS? EVIDENCE FROM SEASONAL BONUSES AND LOTTERY WINS

David Cesar Heymann

Charlie Nourse

Work In Progress

Click here for the latest version

How Much Does Anticipation Matter for MPCs? Evidence from Seasonal Bonuses and Lottery Wins

David Cesar Heymann and Charlie Nourse

November 2026

ABSTRACT

We study how household consumption responds to two very different income changes: one fully anticipated and recurring, the other entirely unanticipated and one-off. Using bank account records for 250,000 Spanish households, we compare workers who receive contractually fixed seasonal bonuses every June and December with workers paid evenly through the year. MPCs out of these bonuses average 21 percent over three months; they decline with liquidity but remain substantial across most of the distribution of liquidity and income. MPCs out of lottery winnings in the same panel are 29 percent over three months, and for households who receive both, the two MPCs are indistinguishable. Anticipation therefore matters far less than standard theory predicts: spending responds to cash flows more than to news about lifetime resources. Standard models predict seasonal bonus MPCs far below lottery MPCs, and near zero for liquid households. Models with a large share of impatient, constrained households can bring the two MPCs together, but only by making too many households hand-to-mouth and leaving MPCs near zero for everyone else. A model in which households dislike running down their savings links consumption more closely to current income and makes MPCs largely insensitive to anticipation. It matches both MPCs, the hand-to-mouth share, and several other moments. If this mechanism, or something like it, is important, policies that change current income flows will move spending more than those acting on wealth or credit, even for households with ample liquidity.

David Cesar Heymann
Faculty of Economics
University of St Gallen
david.cesarheymann@unisg.ch

Charlie Nourse
Department of Economics
London School of Economics
c.r.nourse@lse.ac.uk

1 Introduction

Household consumption is the largest component of GDP, and modelling it has therefore been a central preoccupation for macroeconomists. In all macroeconomic models, the response of consumption to income changes is a crucial object, determining the impact of fiscal and monetary policies, governing saving behaviour, shaping inflation dynamics, and model predictions about inequality. High quality evidence on this is therefore crucial to the project of accurately modelling the macroeconomy.

This paper asks how much anticipation matters for the consumption response to income. We assemble a large and rich dataset based on the bank account records of hundreds of thousands of Spanish households, and use it to compare, within the same panel, how households respond to two polar opposite types of income change. The first is seasonal bonuses ('pagas extraordinarias'), which are contractually fixed additional salary payments, and arrive every June and December. The bonuses are normally equal to the monthly salary in other months, and are paid on the same date each year, so the timing and magnitude are fully anticipated. The second is winning the lottery.

Theory makes strong, and very different, predictions about responses to these two income changes. A lottery win is a real windfall: it raises lifetime resources and releases precautionary savings. The arrival of a seasonal bonus carries no news - it is fully anticipated. A household that knows a bonus will arrive every June and December has no reason not to spread it over the following six months, even when facing borrowing constraints. Standard models therefore predict that the MPC out of a seasonal bonus should be an order of magnitude smaller than the MPC out of a lottery win of the same size.

We find that anticipation matters far less than theory suggests. Using difference-in-differences, we find that MPCs out of fully anticipated seasonal bonuses are in fact quite large: 21% on average, over three months. Our empirical strategy is viable because some Spanish firms pay seasonal bonuses whilst others do not, largely for historical reasons. Therefore, we observe workers receiving seasonal bonuses every six months, but also a control group of workers who are paid evenly from whom we can impute the counter-factual for our treatment group. Since bonus status is set at the firm level, the potential for selection is limited. The spending of the two groups is parallel in the months before each bonus and diverges precisely when bonuses are paid, every year.

We also estimate MPCs out of comparable-sized lottery wins, where identification comes from the randomized timing, and we find these MPCs are a little higher, 29% on average over three months: nowhere near the order of magnitude difference suggested by standard models. Restricting the sample to households who receive both seasonal bonuses and comparably-sized lottery wins, the MPCs out of the two are indistinguishable: both are 25% over three months. The two responses also have the same shape: spending jumps immediately after the money arrives and returns to baseline within three to four months. Both MPCs decline with liquidity, but remain substantial across almost all of the distribution of deposits and income. Our data allows us to break down spending into categories, and using this we show that our results are not driven by shifts in payment timing which do not correspond to consumption, nor are they driven by spending on durables. For example, spending on restaurants alone responds to the bonus in proportion to its share of the budget.

These facts suggest a simple interpretation: spending tracks income, not solely because of occasionally binding constraints or precautionary saving, but because households like to spend new cash when it arrives, whether or not they knew about it in advance, and whether or not they hold existing savings. The response is to cash flows, not information. This behaviour could arise from a number of plausible behavioural mechanisms: budgeting based on current inflows, rules of thumb, or mental accounts treating new income differently to accrued savings. It also matches how survey respondents rationalize their spending decisions (Colarieti, Mei, and Stantcheva (2024)). MPCs estimated from windfalls such as lottery wins, stimulus payments and tax rebates are a primary empirical input used to discipline macroeconomic models of consumption, and they are typically interpreted through models where borrowing constraints and precautionary saving are the key mechanisms. If simpler, behavioural mechanisms in fact explain most of the consumption response, that interpretation, and the policy conclusions drawn from it, need revisiting.

We show that our evidence is very difficult to rationalize in standard models. Modern macroeconomic models have at their core a 'buffer-stock' consumer who faces idiosyncratic income risk which they cannot fully insure against. Borrowing constraints, prudence, or some combination of the two, makes the consumption function of these households concave, which in turn can generate relatively high MPCs out of income shocks like lottery wins. We show that this does not extend to anticipated and recurring income. The Euler inequality restricts the shape of consumption responses to income in all

models of this class: whilst consumption may jump up upon income receipt due to constraints, it must then fall slowly. This restriction is inconsistent with the short and sharp shape of the MPCs we estimate. Applied to fully anticipated, recurring payments, it implies a tight bound on the size of the MPC unless households are very impatient, and it implies that a household which is not close to its borrowing constraint over the bonus cycle should not respond to the bonus at all.

We consider a range of modifications to the baseline buffer stock model which have been used to match MPCs: two asset models, heterogeneous discount factors, and present bias. Two asset models are unable to match our data: lottery MPCs are always much higher than seasonal bonus MPCs, and neither is high for liquid households. Matching total (illiquid) wealth requires relatively high patience, which in turn restricts the size of MPCs out of the fully anticipated seasonal bonuses (consistent with Aguiar, Bils, and Boar (2025)). Households can also plan their allocation to liquid assets around the known bonus schedule, reducing the MPC when it arrives. Models with heterogeneous discounting and models with present bias are able to reduce the gap between the anticipated and unanticipated MPCs under certain calibrations. But they do so by making an implausibly high fraction of households hold almost no liquid assets, and they predict MPCs very close to zero for all other households. This is because the core mechanism generating higher MPCs in these models is making a significant mass of households either very impatient, frequently liquidity-constrained, or both. Consequently, none of these variants can match our finding that households holding large buffers of liquid assets respond significantly to fully anticipated and recurring payments.

A different mechanism for bringing the two different MPCs together comes from models where spending is somehow directly anchored to income. Trying to keep spending within income month by month is a natural and intuitive way to budget, and it is exactly what much popular financial advice recommends. We therefore consider a type of 'spend on arrival' model, in which spending responds directly to current cash flows because of mental accounting, or dis-saving aversion (Graham and McDowall 2026; Gimeno-Ribes 2026). In this model, the response to a payment depends on when the cash arrives, not on when the household learned about it, because households dislike drawing on savings to fund consumption beyond their current income. MPCs are therefore much less dependent on anticipation - the property our data requires. We calibrate the degree of dis-saving aversion in this model to match the seasonal bonus MPC - the same strategy used with heterogeneous discount

factor and present bias models. With a single behavioural parameter chosen in this way, the model simultaneously matches the lottery MPC, the short and sharp shape of both MPCs, substantial MPCs across the liquidity distribution, and realistic levels of liquid wealth. It also predicts that households on lumpy pay schedules hold more deposits, as we observe in the data. Its main shortcoming is that the seasonal bonus MPC declines too little with liquidity.

The mechanism through which high MPCs are generated has important implications for policy. If it is impatience or constraints which generates high MPCs, then policies altering cash-flows for recipients who are not hand to mouth have little impact. Extending credit to constrained households is highly stimulative. On the other hand, if high MPCs are generated by spending tracking income - as our evidence suggests - then cash flow effects will impact spending for liquid households. For example, changes in mortgage payments do transmit to the consumption of relatively liquid households. Extending credit becomes far less effective than policies targeting cash flow directly. Debt relief is more stimulating if delivered via payment reductions than via reductions in principal. Responses to fiscal stimulus are also much more front-loaded in models where spending tracks income. There are also implications for saving behaviour: households who anchor to inflows save more when pay is lumpy.

The literature on how consumption responds to changes in income is very large; three excellent review papers are Jappelli and Pistaferri (2010), Sokolova (2023), and Crawley and Theloudis (2024). A wealth of evidence on MPCs out of windfalls and one-off payments suggests they are relatively large (Johnson, Parker, and Souleles (2006), Parker et al. (2013) and Parker (2017) have been very influential), but the magnitudes vary: Boehm, Fize, and Jaravel (2025) estimate an average MPC of 0.23, concentrated in the month after receiving the income, while Fagereng, Holm, and Natvik (2021) estimate an MPC which reaches almost 1 for low liquidity households after two years. In the meta-analysis of Sokolova (2023), the mean quarterly MPC is 0.35. Sokolova also finds that MPCs out of regular or anticipated payments are somewhat lower, but the evidence on this point is quite mixed. A consistent finding in the literature is higher MPCs for low-liquidity households, in line with the predictions of models with borrowing constraints. MPCs are also often found to be decreasing in payment size.

The literature studying anticipated income changes is smaller. As emphasized in Jappelli and Pistaferri

(2010), it is quite difficult to isolate instances where we can be sure that households really do fully anticipate changes in their income. Souleles (1999) is an influential paper: studying MPCs out of tax refunds he found clear evidence of excess sensitivity, as did Parker (1999) for predictable changes in social security taxes. Baugh et al. (2021) find evidence of asymmetric responses: households increase consumption when they receive refunds, but they do not reduce consumption when they have to make tax payments. A number of papers document failures to smooth predictable income changes that are hard to reconcile with standard models, including Ganong and Noel (2019) at the exhaustion of unemployment insurance and Gerard and Naritomi (2021) for severance payments. Kueng (2018) uses bank transactions data to study Alaska Permanent Fund dividends and finds an average MPC of 25% within a quarter (over-turning the results of Hsieh (2003), a previous paper on the same setting which used survey data). Closest to our paper in message is Graham and McDowall (2026), who use US bank transactions data to show that households spend a substantial share of predictable paychecks and bonuses quickly, even when liquid, and propose a mental accounting model to explain this.

Our core empirical contribution is to compare MPCs out of fully anticipated, recurring payments with MPCs out of fully unanticipated windfalls, in the same panel of households. The seasonal bonus institution gives us an experiment which is quite distinctive among anticipated income changes. Most of the existing literature studies settings like tax refunds, bonuses, or regular dividends where inattention and uncertainty about the timing or magnitude of the income clouds the interpretation of such payments as fully anticipated. Not so in our setting: seasonal bonuses are written into employment contracts, and many employees will have been paid in this manner for many years. The timing and magnitude are fully anticipated: the bonuses are paid on the same date each year, and typically exactly equal the base salary in non-bonus months. A second unique aspect of our setting is the recurrence of the bonuses. Much of the existing literature on anticipated income changes looks at one-off, infrequent, episodes, for example the expiry of benefits. In these settings inattention and limited salience, as well as potentially limited comprehension, are alternative explanations for the failure of consumption to respond in line with standard theory. None of these are plausible for contractual bonuses recurring every six months. Recurrence every six months also substantially shortens the period over which the bonuses can be smoothed - we will show that this recurrence significantly tightens the theoretical upper bound on how much consumption responds in a buffer-stock model. Our setting also offers advantages for identification: we have available a control group

consisting of evenly-paid workers in the same country and labour market, observed at the same time, which allows us to separate the response to the bonus from the seasonality of spending around the months when bonuses are paid. This is unique to the Spanish context: although seasonal bonuses exist in other countries, they are most often universally mandated. Finally, and importantly, our setting allows us to compare the seasonal bonus MPCs to a fully unanticipated income shock - winning the lottery. This is the most distinctive aspect of our paper, allowing us to compare anticipated and unanticipated MPCs in the same panel and within the same households.

There is a small existing literature studying how consumption responds to large, predictable, recurring payments. Hori and Shimizutani (2012) study bonus payments to public employees in Japan, exploiting the abolition of one of the bonuses, and find excess sensitivity concentrated in durable-type spending; Stephens Jr and Unayama (2011) exploit a change in the payment frequency of Japanese public pensions. Both exploit changes in the payment schedule from one year to another, looking at how the seasonal pattern of consumption changed. We believe that our approach (comparing two similar groups at each point in time, only one of which receives seasonal bonuses) offers a more robust test of the consumption response. The paper closest to ours in setting is Browning and Collado (2001), who study the same Spanish bonuses with a similar design using survey data, and find no excess sensitivity. Their data cover a period roughly three decades earlier and are no longer collected at sub-annual frequency, so we cannot compare estimates for the same period. We discuss possible reasons for the difference, including sample size and the classification of bonus households, in Appendix A.12.

We also contribute to the literature on modelling consumption in macroeconomics. We show explicitly that the Euler inequality implies restrictions on the shape of the intertemporal MPC path, and that this in turn implies a restriction on the size of MPCs out of fully anticipated and recurring payments. In doing this, we exploit the unique advantages of our empirical setting: modelling income changes as following a fully deterministic cycle would not be plausible in most other settings. Here we draw on the extensive theoretical literature on consumption going back to Friedman (1957), through Hall (1978), Deaton (1991), C. D. Carroll (1997), Kaplan and Violante (2014), Debortoli and Galí (2024), Kaplan, Moll, and Violante (2018), Auclert (2019), Auclert, Rognlie, and Straub (2024). The idea that the rate at which consumption falls after a payment is informative about impatience goes back

to Shapiro (2005); we turn it into a bound that holds for constrained households, aggregates across heterogeneous households under precautionary saving, and restricts both the shape of MPC paths and the size of MPCs out of recurring payments. Furthermore, we show that our combination of moments helps to separate models according to how they generate high MPCs, siding clearly with ones that do so by somehow anchoring spending to income.

Finally we build on the behavioural economics literature on mental accounting: R. Thaler (1985), Shefrin and R. H. Thaler (1988), Milkman and Beshears (2009), and a more recent literature applying this directly to consumption in macroeconomics: Baugh et al. (2021), Graham and McDowall (2026), Gimeno-Ribes (2026). Our contribution here is to show that the comparison between anticipated and unanticipated MPCs is a sharp test of this class of models: they predict that the two should be similar, standard models predict that they should be very different, and the data side clearly with the former. We also show that, in a single calibration, the models proposed by Graham and McDowall (2026) and Gimeno-Ribes (2026) match both MPCs, the short and sharp shape of the responses, and realistic levels of liquid wealth. Kueng (2018) interprets his findings on excess sensitivity to anticipated Permanent-Fund dividends as near-rational: deviations from smoothing are largest where their utility cost is smallest. Because seasonal bonuses scale with salary, the utility cost of failing to smooth them varies much less across households than for a flat per-person dividend. Using Kueng's formula, we find these costs to be small. However, we view his explanation as complementary to the behavioural model we study: given the small costs of the deviations from full smoothing we observe, it is plausible that even small psychological frictions might generate them.

In our theoretical section we also draw on the emerging literature on expenditure shocks: De Nardi, French, and Jones (2010), Fulford and Low (2024), Miranda-Pinto et al. (2025), Bergholt, Furlanetto, and Mori (2025). Two other strands of literature are closely related. Firstly, a small literature on the macroeconomic implications of seasonality: Barsky and Miron (1989), Beaulieu and Miron (1992), Wen (2002). Secondly, the literature on payday effects: Stephens Jr (2003), Shapiro (2005), Mastrobuoni and Weinberg (2009), Olafsson and Pagel (2018).

Table 1: Comparison of our main variables with official data sources.

	Source	2021	2022	2023	2024
Household consumption					
Our data	CaixaBank	20,374	21,941	23,122	23,794
Spending excl. rent and durables	EPF (INE)	22,397	24,440	24,981	25,623
All cash spending (alternative)	EPF (INE)	25,774	28,031	28,684	29,528
Consumption, national accounts (macro)	National accounts	36,354	39,873	42,003	44,354
After-tax salary					
Our data	CaixaBank	26,717	27,129	28,520	29,336
Net salary, main earner	ECV (INE)	23,552	24,950	25,579	26,370
Net salary, all earners (alternative)	ECV (INE)	32,835	35,230	36,589	38,031
Net salary, one-adult households (alternative)	ECV (INE)	20,192	21,646	22,443	23,274
Gross wages, national accounts (macro)	National accounts	25,088	26,935	28,642	30,540
Current-account deposits					
Our data	CaixaBank	21,369	22,382	23,090	22,471
Payment-account balances	EFF (Banco de España)	–	20,578	–	–
Liquid assets (alternative)	EFF (Banco de España)	–	31,419	–	–

Note: In each section, the first row shows the average annual value in our data, the next row shows the most comparable value from official data, and subsequent rows show other official statistics for context. Annual means in euros per household; deposits are balances. Official survey rows refer to comparable households: main person aged 25–65 whose main income is salary. ECV figures are aligned on the income year (survey year minus one). EFF figures exist only for survey waves, so deposits are compared in 2022 only. Sources: INE (EPF, ECV microdata), Banco de España (EFF microdata), Eurostat national accounts.

2 Data

We use a novel dataset constructed from anonymized transactions and account data, aggregated at the household level. The data come from CaixaBank, the leading bank in Spain by total assets and share in direct payroll deposits and originates from financial transactions registered in bank accounts.[1] Individual bank accounts are then aggregated into households based on their address. We draw on the entire set of CaixaBank customers and construct a panel from January 2019 to June 2025, filtering down to households who are active in every month over our analysis period.

We define our main consumption measure as the sum of card spending, bills, and cash withdrawals. We exclude spending via bank transactions because these are too difficult to categorize into consumption versus transfers into savings accounts or individuals. This means we do not capture categories such as rent, or large purchases of durable goods such as cars or home refurbishments.

Card spending includes both debit and credit card spending. In both cases, consumption is counted on the month of the expenditure. In case the customer chooses to fraction a payment, consumption is

[1] CaixaBank has 18.9 million clients in Spain, and 36% of Spanish paychecks are paid into CaixaBank accounts, according to the Q1 2026 results presentation.

still counted on the month of the purchase. Cash withdrawals are the sum of total withdrawals from ATM machines or bank offices in each month. We assume withdrawals are used for consumption in the same month. Bills are calculated from direct-debit payments. Essential bills include water, electricity, and gas. Non-essential bills are all other direct-debit payments (e.g. subscriptions, internet etc.)

Our sample includes income from paychecks, state pensions, unemployment benefits, and other public benefits. For the self-employed, income is the sum of incoming transfers minus outgoing transfers. In our main analysis, we focus on households receiving income from paychecks and no income from either unemployment benefits or pensions. Household age is determined by choosing the age of the household member with the highest income in a given month. Checking deposits are exclusively deposits in checking accounts (depositos a vista). We take the average value across the given month. Finally, we also observe lottery winnings. These are identified based on the matching the transaction tag of the payment in with the names of the main lottery providers in Spain.

Table 1 compares household average annual totals in our dataset with official data sources, in Appendix A.1 we compare the full distributions. Overall, our data appear to capture the income and spending of the main earner in each household very well. Our income variable lines up very closely with official data on the net (after tax) salary of the main earner, whilst our total spending variable is on average about 10 percent lower than the comparable official measure which is monetary spending excluding rent and big-ticket durables. This is probably mainly a consequence of imperfect household matching, meaning we miss some of the spending done by the other adults in the household. Our data appears to capture the average level, and the distribution, of checking deposits well. Checking deposits make up about two-thirds of total liquid assets in the official data.

3 Seasonal Bonus MPCs

In this section we provide the evidence for our main empirical facts about consumption and seasonal pay. Spanish law entitles all workers to receive two 'extraordinary payments' each year. The typical form this takes is that a worker's annual salary is divided into 14 components: they are paid 1 part in 14 most months, but in June and December they receive 2 parts in 14. This means that their monthly paycheck predictably doubles twice per year. The law also permits for these payments to be pro-rated

over the course of the year, so that pay is constant from month to month. Whether this happens is determined by collective bargaining agreements which are fixed sometimes at the sectoral level either nationally or regionally, and sometimes at the firm level. This means that workers do not choose how they are paid, and there is a lot of variation within sectors and regions in terms of the pay schedule.

Extraordinary payments, particularly at Christmas, appear to have deep roots, and there is evidence of firms paying their workers in this fashion back to the 1910s. The system was institutionalized in the 1940s when first Christmas, and then summer extra payments became mandatory. This was subsequently relaxed to allow for pro-rating.

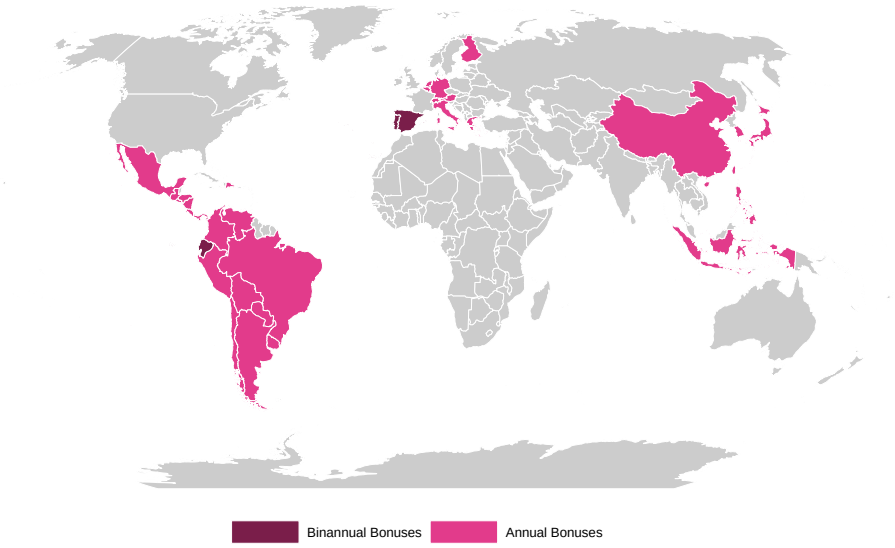


Figure 1: Countries where seasonal pay is customary for a large proportion of employees.

Regular, recurring seasonal bonuses are common around the world, in many cases mandated by law. A core empirical problem in studying them is that consumption preferences are also highly seasonal, and in aggregate data the seasonal preference shifts are not separately identified from the consumption response to the bonuses. The Spanish setting is uniquely useful for economists because pro-rating is allowed, so seasonal and flat pay coexist. From households with non-seasonal income, we can estimate the seasonal preference shifts, allowing us to identify the causal impact of seasonal bonuses from the difference between them and the households getting seasonal bonuses. Furthermore, bonus status is set at the firm level, and the collective bargaining agreements which govern it are the product of long and complex histories of negotiations between employers and different unions. This produces

a patchwork system, with a wide cross-section of workers receiving both types of pay.

3.1 Identification

We divide our time period into 12 event windows of six months each. Each event window runs from two months before to three months after a bonus payment. The first window covers April 2019 to September 2019 and the final window covers October 2024 to March 2025. Within each window, we restrict the sample to a balanced panel of households which are either categorized as bonus or non-bonus households for the full duration of the window. We then estimate the response of spending to bonus receipt in each event window using difference-in-differences: essentially comparing the growth of spending across the bonus and non-bonus groups around months of bonus payment. The core idea is that we can use the non-bonus group to impute the missing counter-factual for the bonus group: what would their spending have been had they not received a bonus?

Assuming the bonus group get their bonuses in period T, the basic difference-in-differences calculation is:

$$DID_{T+h} = E\left[C_{i,T+h} - C_{i,T-1} | i \in \text{Bonus Group}\right] - E\left[C_{i,T+h} - C_{i,T-1} | i \in \text{Control Group}\right]$$

Where $C_{i,t}$ is consumption of household i in period t.

Let $C_{i,t}(1)$ be the potential outcome of consumption for household i in period t if they receive a bonus in period T. Our parameter of interest is the average treatment effect on consumption of getting a seasonal bonus, for the treated group of bonus households:

$$ATT_{T+h} = E\left[C_{i,T+h}(1) - C_{i,T+h}(0) | i \in \text{Bonus Group}\right]$$

The difference-in-differences calculation DID_T is equal to the ATT_T under the standard assumption of parallel counter-factual trends and the ancillary assumptions that there are no anticipation effects ($C_{i,T-1}(1) = C_{i,T-1}(0)$) and no spillovers from treated to controls.

Parallel counter-factual trends requires that, in the absence of treatment, the average consumption of the bonus group would have evolved parallel to that of the control group:

$$E\left[C_{i,T}(0) - C_{i,T-1}(0) | i \in \text{Bonus Group}\right] = E\left[C_{i,T}(0) - C_{i,T-1}(0) | i \in \text{Control Group}\right]$$

Probably the most convincing evidence in favour of this identification assumption comes from simply plotting the data. Figure 2 shows average salary income over time for the two groups. The divergence each June and December is very clear. [2] Figure 3 shows average total spending for the two groups. The two series move more or less in parallel in the months before bonus receipt, before diverging each and every time bonuses are paid. Spending is seasonal in both groups, but it is notably more seasonal in the bonus group. December bonuses are generally paid halfway through the month, and in line with this the divergence in spending occurs mostly in the first month. In June, bonuses are typically paid towards the end of the month, so the largest divergence in spending is seen in July. Both income and total spending have similar levels across the groups in the months before bonuses are paid - this is not the result of any normalization, it rather reflects the fact that the groups are similar on average.

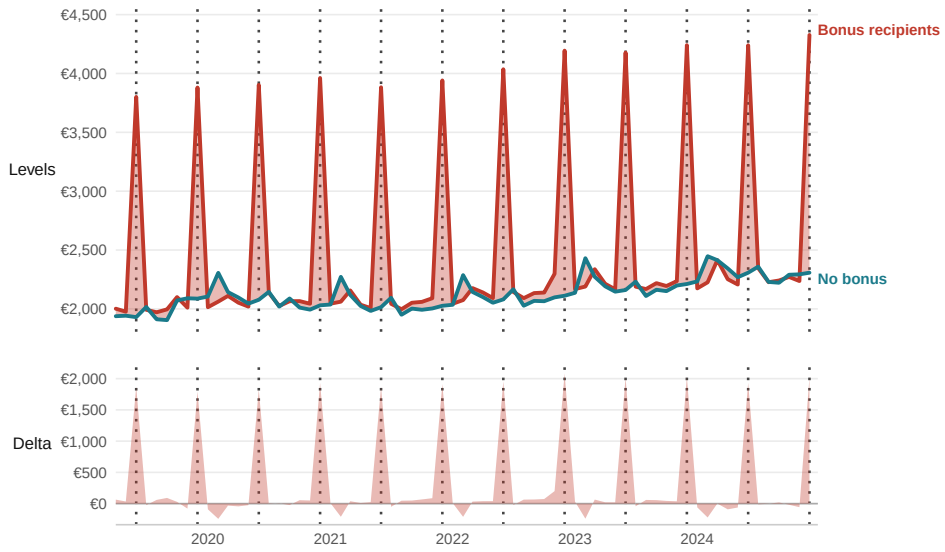


Figure 2: Monthly salary across the bonus and non-bonus groups (upper panel), and the delta between them (lower panel).

[2] There is a slight upwards jump in the salaries of the non-bonus group in July. This is likely because some workers in this group receive some smaller bonus payment in July.

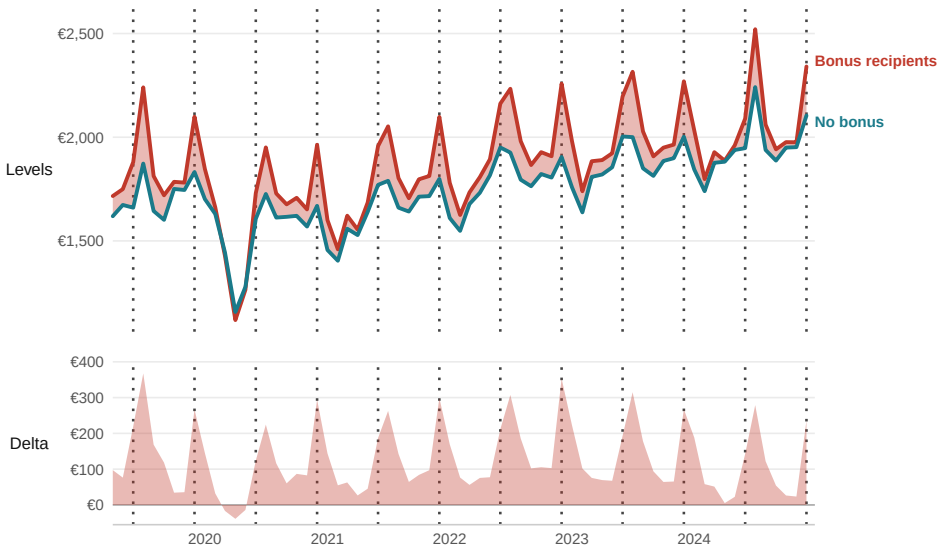


Figure 3: Monthly total spending across the bonus and non-bonus groups (upper panel), and the delta between them (lower panel).

One potential concern would be that bonus and non-bonus households are somehow selected or sorted into these groups in a way that correlates with the seasonality of their consumption patterns. We do not think this is very likely because bonus status is not chosen by the individual, it is determined at the level of the firm. Spanish law stipulates that all workers are entitled to receive seasonal bonuses twice per year, but it allows for collective bargaining agreements to decide whether these bonuses should be pro-rated across the year or not. Collective bargaining agreements are sometimes set nationally for a whole sector, sometimes with separate regional agreements within a sector, and sometimes at the firm level. These agreements mandate whether each firm pays the bonuses as lump-sums or pro-rated, meaning all of the variation in bonus scheme is at the firm (and sometimes the sector) level.

Checking for balance in our data, the bonus and non-bonus groups are relatively similar across most dimensions we measure (Table 2). The two exceptions are age - bonus households are on average about 4 years older - and checking deposits. The difference in checking deposits is reduced, but remains, if we control for age and income. We will discuss in the theoretical section that it is possible that this difference is a consequence of the bonus pay structure: in a model with mental accounting, households save more when their pay is unevenly distributed. We also highlight at this point that the

average level of checking deposits is quite high: equal to about 8 months salary for the non-bonus group and 10 months for the bonus group. This will be an important moment for evaluating models in the next section. As a robustness check, we repeat our main analysis under the stronger assumption that parallel trends holds only conditional on the lagged value of the covariates age, income, and deposits. The results are very similar.

The assumption of no anticipation effects requires that treatment has no effect on potential outcomes before receipt: $C_{i,T-1}(1) = C_{i,T-1}(0)$. This is not the assumption that households fail to anticipate their bonuses — which would be indefensible given the features we have emphasized about this institution — but the weaker assumption that their spending in May and November is not altered by that expectation.

We use our estimates from this section for two purposes, and we should be clear that the anticipation effects assumption is only critical for one of them. The first use is that the estimate serves as a moment for discriminating between consumption models. For that purpose we compute the identical statistic on model-simulated data, comparing simulated bonus-regime and flat-regime households around bonus months exactly as we compare them in the data. The no-anticipation assumption clearly fails in standard macroeconomic models, where forward-looking households choose consumption from their entire future income path and therefore do adjust before receipt. But this is not a nuisance to be corrected for: it is part of the substantive prediction we are evaluating. Those models predict that households smooth around the bonus on both sides — adjusting in advance of receipt, and spreading whatever remains over the months that follow, since concentrating the entire bonus into one or two months is not optimal even for households who have been liquidity constrained. The difference-in-differences statistic measures how far observed spending departs from that predicted path, and because the same estimator is applied to both model and data, any smoothing is treated symmetrically on both sides of the comparison. In this sense the statistic is best understood as a measure of the excess sensitivity of consumption to bonus receipt under a given pay regime, and the model–data comparison does not rest on the no-anticipation assumption at all.

The other use of our estimate is as a conventional MPC estimate, to be set alongside the existing literature. Here the assumption does matter, and we offer two pieces of evidence for it. First, prior work finds that households generally wait until anticipated transfers arrive before spending them

Table 2: Balance Across Groups

	Non-bonus HH (N = 125,889)		Bonus HH (N = 107,386)		
	Mean	SD	Mean	SD	Normalized Difference
Age	44.46	9.92	48.30	9.34	-0.282
Number of adults	1.35	0.69	1.43	0.75	-0.085
Monthly salary	2,176	1,795	2,491	1,495	-0.135
Bonus month salary	2,162	1,659	4,156	2,012	-0.765
Non-bonus month salary	2,164	1,627	2,159	1,064	0.002
Total spending	1,808	1,752	1,927	1,535	-0.051
Card spending	1,136	1,148	1,202	1,029	-0.043
Cash withdrawals	330	571	355	576	-0.031
Checking deposits	17,394	43,601	24,426	43,493	-0.114
Other CBK assets	33,118	76,420	33,763	65,812	-0.006
Checking deposits / income	7.93	28.05	10.66	25.93	-0.071
Borrowing from CBK	23,234	53,734	20,233	42,513	0.044

(Kueng (2018), Baugh et al. (2021), Fuster, Kaplan, and Zafar (2021)). Second, Figure 3 shows little divergence in spending between the groups in the months immediately before bonus receipt.This is reflected in our pre-trend estimates, which are close to zero, though their utility is limited in our setting because treatment repeats every six months and we cannot therefore add many pre-period coefficients without running into the lagged effect of the previous bonus. The two most plausible ways the assumption could fail also have opposite signs. Anticipatory smoothing would raise $C_{i,T-1}(1)$ above $C_{i,T-1}(0)$ and bias our estimates downwards, making 0.21 a lower bound on the total response. Deferral of planned purchases into the bonus month would do the reverse, and would represent re-timing rather than additional consumption. We present evidence below that the MPC appears to be mostly driven by non-durables and services where deferral is unlikely to be a major factor.

The final important identification assumption is that there are no spillovers from the treatment group to the controls: the spending of the non-bonus households isn't affected by the fact that other households are getting bonuses. There may be reasons for this not holding: social conventions may have evolved around bonus dates which induce all households to spend more money, firms might also respond through pricing or in other ways which affect the non-bonus households. Such spillover effects seem likely to increase spending in the non-bonus group and would therefore push us towards underestimating the consumption response to bonus receipt.

3.2 Estimation

We implement our difference-in-differences estimation using a two-way fixed effects regression.

$$C_{i,t} = \alpha_i + \lambda_t + \sum_{s \neq -1} \tau_s D_{i,t-s} + u_{i,t}$$

The regression includes household fixed effects α_i and time fixed effects λ_t. $D_{i,t}$ is a 'treatment' indicator: taking the value one for bonus households in the month when they received their bonus (which corresponds to s=0) and zero otherwise. We include leads and lags of $D_{i,t}$, omitting the period before bonus receipt (May or November) as the base period.[3] The coefficients τ_s estimate the *ATT* of bonus receipt on spending (in euros) h periods after bonus receipt.

Figure A.9 reproduces all of the individual event study regressions for each event window. For parsimony, here we report the results from a stacked event study where we summarize all of the events in a single regression, using household-by-event fixed effects.

$$C_{i,t} = \alpha_{i,e} + \lambda_t + \sum_{s \neq -1} \tau_s D_{i,t-s} + u^c_{i,t}$$

To calculate the MPC, we also need to estimate the average change in income induced by bonus receipt, so we also estimate the following regression:

$$Y_{i,t} = \alpha_{i,e} + \lambda_t + \sum_{s \neq -1} \pi_s D_{i,t-s} + u^y_{i,t}$$

where $Y_{i,t}$ is salary income. The coefficient π_0 estimates the average size of the seasonal bonuses received.

The MPC is calculated at the one-month and the three-month horizons as follows:

[3] We estimate only one 'pre-trend' coefficient relative to the base period in each window because of the repeated treatment every six months. The coefficient on $D_{i,t-3}$ in one event window is the same as the coefficient on $D_{i,t+4}$ in the next event window (except for small variations in sample selection). This regression therefore computes 2x2 difference-in-differences comparisons relative to the base period. Under our identification assumptions, τ_s captures the (euro) change in consumption induced by receiving a seasonal bonus (we discuss below how to convert this object into an MPC).

$$MPC_{1m} = \frac{\tau_0}{\pi_0}$$

$$MPC_{3m} = \frac{\tau_0 + \tau_1 + \tau_2}{\pi_0}$$

Standard errors for the MPCs are calculated using the delta method, taking account of the cross-equation correlation of the estimators for τ and π.

3.3 Main results

Our main event study results are shown in Figure 4, which plots the estimates of π_s in the left panel, τ_s in the middle panel, and the (non-cumulative) MPC on the right. Total spending rises by about 175 euros in the month of bonus receipt and the month after, with the effect falling to near 50 euros in the second month after before falling to zero. The MPC is 9% in the first month and 21% over the three months after receipt. The spending response is concentrated in the first two months after receipt, and is short and sharp. Figure A.9 in the Appendix A.4 shows each of the event studies separately, and it is clear from this that for December bonuses most of the effect occurs in the first month, whilst for June bonuses most of the effect occurs in the second month. This is largely a measurement artifact: December bonuses are typically paid earlier in the month, because of Christmas, whilst June bonuses mostly arrive towards the end of the month. This means that the measured one-month MPC of 9% is somewhat misleading: it averages a response from the last week of June with the whole of December - in fact most of the spending response is concentrated in the first month after bonus receipt. For this reason, we focus mainly on the three month estimate in the rest of the paper.

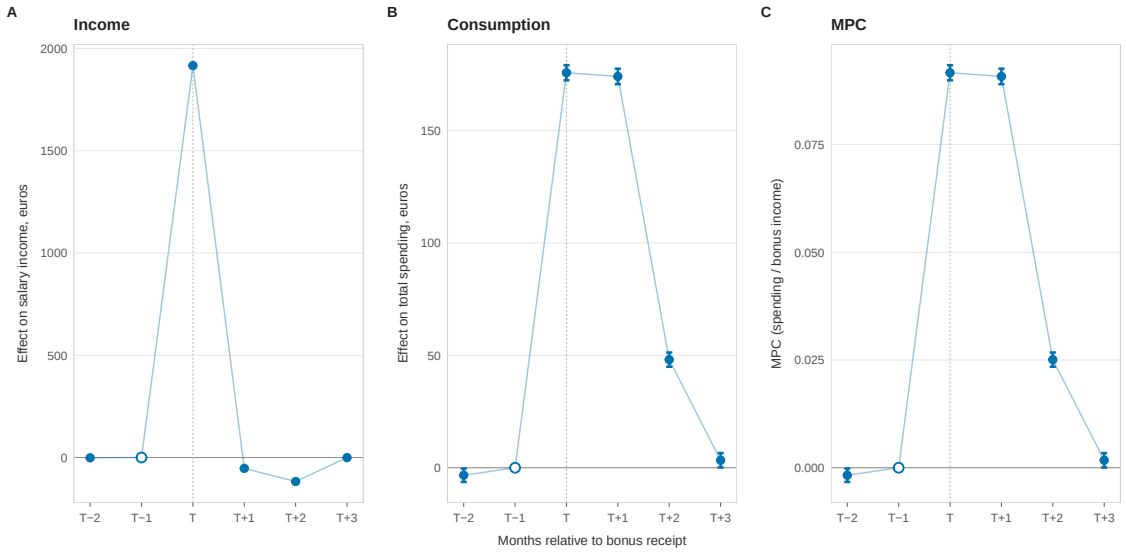


Figure 4: Event study plots based on a stacked regression pooling all events. 95% confidence intervals shown, based on standard errors clustered at household by event and date level.

As a robustness exercise, we repeat our main analysis under the stronger assumption that parallel trends holds only conditional on the (lagged) values of age, income and deposits. We implement this using the doubly-robust method of Sant'Anna and Zhao (2020). The results are in Appendix A.2: the MPC is very similar, falling slightly to 0.19.

Our results are quite different to Browning and Collado (2001), who found no evidence that bonus status affects spending (except for small durables). In Appendix A.12, we replicate their study and try to reconcile their results with ours. We believe the main reason for the difference is simply power: the dataset they had available was small, and they did not have high power to detect an MPC of the size we estimate.

3.4 Consumption or expenditures?

In this section, we address the concern that our data captures expenditures rather than consumption by exploiting data breaking down households' spending into categories. Our transactions data measure payments made from users' bank accounts, and there are several instances where payment may not coincide with the actual utility-generating consumption. Our consumption measure is built by summing card spending, cash withdrawals, and some direct debits. The main thing excluded is bank

transfers, because these are very difficult to classify - they can be payments for goods and services, but they can also be transfers between accounts or individuals, or paying off balances on pre-paid cards provided by some retailers. This means that two important categories of consumption are not in our data. The first is rent payments. These are typically made using bank transfers. The second is big-ticket durables. Debit and credit cards in Spain typically have limits on per-transaction and per-day spending, meaning very large purchases like cars, or large renovations, are made using bank transfers. These are not in our spending measure, and to the extent that this kind of spending may be particularly responsive to bonus receipt, we may in fact be under-estimating the true response. As we will discuss below, this issue is likely to be even more severe when we try to measure the consumption response to very large lottery winnings - other evidence suggests this is likely to dis-proportionately be on big ticket items.

Two concerns about measurement might remain. The first is that much of the spending response is concentrated in sectors where payment is separated temporally from receiving the goods or services, for example that bonus recipients disproportionately book (and pay for) their holidays right after getting the bonus, even though all households holiday at the same time. The second is that the observed spending response might disproportionately reflect durable goods, complicating our interpretation of it as a marginal propensity to **consume**. As we explained in the previous paragraph, the largest and lumpiest durables purchases are excluded from our dataset in any case, but this issue could still be relevant.

To address these concerns, we dis-aggregate our total spending measure into its components: cash withdrawals, different types of bill payments (direct debits), and a decomposition of card spending. Card spending is broken down into categories based on the merchant to whom the payment is sent. Each merchant in Spain has a code indicating its line of business, and these are used to categorize card payments. While this may lead to some measurement error as larger retailers sell many different types of product, it is likely to give a very good approximation of the type of good or service being purchased.

Table 3 shows results from decomposing our headline MPC into categories. This is done by running our main regression separately for each category of spending and summing the T, $T+1$ and $T+2$ coefficients to estimate the effect of bonus receipt on spending within each category. The table

Table 3: Spending response decomposed by category

Category	Effect	S.E.	Share (%)	Cumulative (%)
Cash withdrawals	65.5***	(5.6)	20.7	20.7
Food, drinks, tobacco	40.6***	(1.9)	12.9	33.6
Restaurants	24.1***	(1.6)	7.6	41.2
Other retail	22.9***	(2.3)	7.2	48.4
Fashion	20.5***	(2.0)	6.5	54.9
Insurance bills	17.2***	(5.7)	5.4	60.3
Other bills	12.0*	(6.3)	3.8	64.1
Transport	11.8***	(2.5)	3.7	67.9
Utility bills	10.0***	(1.2)	3.2	71.1
Culture and shows	7.3***	(2.0)	2.3	73.4
Gas stations	6.5***	(0.7)	2.1	75.4
Furniture	5.2***	(1.3)	1.6	77.1
Hotels	5.1***	(1.7)	1.6	78.7
Travel agencies	4.9***	(1.8)	1.6	80.2
Home appliances, IT	3.2**	(1.6)	1.0	81.2
Pharmacies	2.3***	(0.5)	0.7	81.9
Other essential goods	-0.6	(1.7)	-0.2	81.7
Residual	57.7		18.3	
TOTAL (Total spending)	316.2***	(15.1)	100.0	

Notes: Sum of coefficients on the bonus indicator at t, t+1 and t+2. Standard errors clustered by household × event. * p<0.10, ** p<0.05, *** p<0.01.

collects all of the results and presents a full decomposition.

The results suggest that the spending response is broad-based, with the largest categories being cash withdrawals, food, drinks and tobacco, and restaurants. Durable goods categories (furniture, home appliances and IT) account for a small amount of the total. Bills collectively account for 12.4 percent of the total MPC, and it is possible that this largely reflects shifting of payment timing rather than a consumption response. Many insurance bills are paid annually, for instance, and it is possible that bonus recipients are more likely to time these after bonus payments. Hotels and travel agencies could also be cases where payment and consumption timing differ, although it is not necessarily the case that the spending response here does not reflect genuine consumption. More pertinently, these potentially problematic categories make up at most just under 20 percent of the total. Not all vendors can be categorized by sector, leaving a residual of spending which is not tagged.

Should cash be included in our measure, given that we do not observe what it is spent on? Cash is still fairly important as a means of payment in Spain, particularly for smaller, point-of-sale payments. Cash is most prevalent in payment at restaurants, vending machines, and cultural activities - important categories of consumption which seem likely to be responsive to bonus receipt (European Central Bank (2024), Rodríguez Molina and Esandi Erro (2025)). Since the average amount of cash held by Spanish consumers is low relative to the use of cash in transactions (Rodríguez Molina and Esandi Erro (2025)), the average holding period is likely to be short, and assuming that the majority of cash withdrawn is spent within three months after withdrawal seems reasonable.

Overall, measurement issues mean that our estimated spending response could either under-estimate the true MPC (due to omitting big-ticket items) or over-estimate it (due to including some payments not fully coinciding with consumption). The evidence presented here does not suggest that we substantially over-estimate the MPC due to these measurement issues. Since our theoretical results on the MPCs out of seasonal bonuses are quite strong - they suggest that model MPCs are an order of magnitude lower than our empirical estimates - our core message is robust to even quite a large degree of mis-measurement, and the evidence here does not support this.

4 Lottery Win MPCs

In this section we estimate consumption responses to lottery winnings. These estimates provide a useful contrast with our results in the previous section - while seasonal bonuses are fully anticipated and recur every six months, lotteries are entirely unanticipated, one-off events. As we will discuss in the following section, standard models predict that consumption responds very differently to these two different types of income change.

Lottery winnings are identified in our data by matching the source of the transfer with the major lottery providers in Spain. The lottery in Spain is dominated by two providers, the state-run Sociedad Estatal Loterías y Apuestas del Estado (SELAE), and the charity Organización Nacional de Ciegos Españoles (ONCE), we identify payments from both of these as well as a number of smaller providers. We observe when winnings are paid into the account of the winner.

A high proportion of Spanish adults play the lottery: 31 million adults participate in SELAE lotteries (about 85% of adults) and 8.3 million in ONCE lotteries (Gómez Yáñez (2025)). In 2024, SELAE paid out 6.3bn euros and ONCE 1.6bn euros in lottery winnings. SELAE's Christmas lottery, "El Gordo" has the largest prize pot of any lottery in the world: it distributed 2.7bn euros in the 2025 draw, and is played by the majority of Spanish adults.

Table 4 shows summary statistics on lottery winnings in our data. Overall, we observe 30,556 lottery wins (after excluding repeat winners). We exclude the top 1 percent of prizes, which include several in the tens of millions of euros. The remaining prizes lie between 2 and 328,000 euros. One useful feature of the lottery in Spain is that a relatively large number of medium-sized prizes are paid out, rather than a small numbers of huge prizes. This means we observe a large number of winnings that are in the same size range as the seasonal bonus, allowing us to compare consumption responses across anticipated and unanticipated payments which are size-matched. While the mean prize is 14,474 euros, this is heavily inflated by a few very large prizes; the median prize is 3000 euros.

4.1 Identification and Estimation

Identification in this setting is in principle straightforward: among ticket-holders in any given month, whether you win the lottery is randomized. Comparing the spending of winners with a control group of non-winners can therefore identify the causal effect of winning the lottery. Lottery players

Table 4: Lottery statistics.

Group	N	Min (euros)	Mean (euros)	Median (euros)	Max (euros)
All lottery winners	30,556	2	14,474	3,000	328,000
Prize quintile 1	6,137	2	294	250	574
Prize quintile 2	6,136	577	1,664	2,000	2,376
Prize quintile 3	8,941	2,376	4,167	3,557	6,000
Prize quintile 4	3,332	6,001	9,508	8,645	18,066
Prize quintile 5	6,010	18,094	60,118	35,000	328,000
Prizes matched to seasonal bonuses	11,029	658	2,281	2,301	4,062

may have different spending patterns, in terms of time trends or seasonality, to non-players, and we therefore want to restrict the control group to other lottery-players only. Since we do not observe whether or not our households play the lottery unless they win, we restrict our sample to households which win the lottery at some point in our sample. Within this sample of ever-winners, we run an event-study, comparing winners with not-yet-winners (we exclude the small number of households which won multiple times).[4]

Lottery wins are staggered over time, so we use the imputation estimator of Borusyak, Jaravel, and Spiess (2024) rather than the two-way fixed-effects (TWFE) estimator we used in the previous section. TWFE identifies each month's time fixed effect partly from households that have already won, netting out their impact on these with the average effect at each horizon. If a treatment cohort's response path differs from that average, the difference is absorbed into the time effect, which is then wrong as a counterfactual for treated observations (Borusyak, Jaravel, and Spiess (2024); Sun and Abraham (2021)). For example, if December winners are still spending more in February than the average treated cohort does at that horizon, February's time effect is overstated for the other groups, and the other estimates using that February effect are understated (for instance households treated in February for whom the February effect is used to compute the counter-factual). TWFE is therefore valid only if responses follow the same path across cohorts, which is implausible here. Wins are seasonally clustered, and draws differ in prize size and in when within the month they pay out; El Gordo, for instance, pays large prizes late in December. The imputation estimator is robust to this: household and time effects are estimated only from observations before the household has won, so no treatment effect enters the counter-factual.

[4] In Appendix A.8, we repeat the analysis using a random sample of never-winning households as the control group instead. There is no meaningful impact on the results.

Specifically consider a household i winning the lottery in period T_i. Denote their spending in the period they win as C_{i,T_i}. The treatment effect of winning on their spending is $\tau_i = C_{i,T_i} - C_{i,T_i}(0)$, where $C_{i,T_i}(0)$ is their potential outcome if they had not won. The estimator imputes this counterfactual by estimating $\hat{C}_{i,T_i}(0) = \hat{\alpha}_i + \hat{\lambda}_{T_i}$. The unit fixed effect is estimated using the household's pre-win observations, and the time fixed effect is estimated using the observations of all ever-winning households who have not yet won in period T_i. The estimated treatment effect is therefore $\hat{\tau}_i = C_{i,T_i} - \hat{C}_{i,T_i}(0)$. The dynamic treatment effects in the periods after treatment are estimated by $\hat{\tau}_{i,h} = C_{i,T_i+h} - \hat{C}_{i,T_i+h}(0)$. The average treatment effect at each horizon is then calculated by averaging over all households and periods, in our case with equal weights to give the average additional euros of spending at each horizon: $\hat{\tau}(h) = \sum_{i \in I_h} \frac{1}{n_h} \hat{\tau}_{i,T_i+h}$, where I_h is the set of observations of winners h periods after winning, and n_h is the number of them.

Finally the period MPC at horizon h is calculated by summing all of the spending responses, and dividing by the total amount of winnings: $\hat{\text{MPC}}(h) = \sum_{i \in I_h} \hat{\tau}_{i,T_i+h} / \sum_{i \in I_h} P_{i,T_i}$, where P_{i,T_i} is the lottery winnings of household i in period T_i.[5] Using a binary win indicator, and then dividing total euros of additional spending by the total euros of winnings recovers correctly the true aggregate MPC: the total additional spending induced per euro of winnings. Standard errors are estimated following Borusyak, Jaravel, and Spiess (2024), with the denominator, lottery winnings, treated as non-stochastic because we observe its exact value. Our MPC estimates are strongly dependent on the size of the lottery win, so in what follows we always estimate MPCs within bins of prize-size. When we do this, all of the data we use for estimation is drawn from winners whose prizes fell in that bin. So the imputation of the counter-factuals $\hat{C}_{i,T_i}(0) = \hat{\alpha}_i + \hat{\lambda}_{T_i}$ is done using only not-yet-winners whose winning subsequently falls within that prize bin.

We estimate pre-trend coefficients by the method proposed in Borusyak, Jaravel, and Spiess (2024): using not-yet-treated observations, we estimate the equation $C_{i,t} = \alpha_i + \lambda_t + \sum_{h=1}^{H} \tau_{-h} D_{i,T_i+h}$ where D_{i,T_i+h} is an indicator for being treated in h periods, and the coefficients τ_{-h} are the estimates of the pre-trend effects relative to a base period of observations more than H periods before treatment. This essentially estimates whether units start to diverge from the level of spending predicted by the fixed

[5] We are careful to ensure that we use the same set of households for both the numerator and denominator at each horizon. A household winning near the end of our sample would not be able to contribute to the estimated response all the way to H, and we make sure that we drop their prize from the denominator whenever their response does not appear in the numerator.

effects model in the periods before receiving treatment. In some of our specifications, these reveal evidence of a small effect in the month before lottery winnings. This could be an anticipation effect: we observe when lottery winnings get paid into a winner's account, but they may learn that they won before this. Because of this, when estimating $\hat{C}_{it}(0) = \hat{\alpha}_i + \hat{\lambda}_t$, we use observations from two (rather than one) periods or more before winning.

In their influential study of Norwegian lottery winners, Fagereng, Holm, and Natvik (2021) use a somewhat different empirical strategy: restricting also to households which ever win in their sample, they run a separate regression at each horizon h of spending on a continuous lottery winnings variable. At each horizon they use only treated observations at that horizon, and control observations which are not yet treated, thereby also avoiding forbidden comparisons. We repeat our main analyses here using their estimator in Appendix A.8. For small and moderate prizes, the results are very similar. For large prizes, the BJS estimator yields significantly higher MPCs. This is entirely explicable: when using a continuous regressor, the coefficient gives a weighted average of dose-specific effects where the weights are proportional to the square of the dosage: Fagereng, Holm, and Natvik (2021) note this explicitly. In a sample with a few very large prizes, these results get weighted by their square, and completely dominate the whole effect. This lowers the estimated effects substantially, because in our sample we estimate very small MPCs for winners of huge prizes. This is in large part a measurement problem: since we do not use purchases made via bank transfer (it being very difficult to distinguish them from transfers for other purposes), we do not see a lot of the kind of spending someone who won a massive lottery prize might engage in: housing, cars, boats, and so on.

4.2 Results

In theory and in prior evidence, MPCs are strongly dependent on the size of the windfall, so we start by estimating the path of MPCs separately across the 5 quintiles of the lottery prize distribution. Rows 2-6 of Table 4 shows summary statistics for wins within each of these 5 quintiles. The group sizes are uneven because a large number of wins have the same value (at round numbers), and all of these fall into one quintile only.

Figure 5 shows the estimated MPC paths within each bin. The MPCs decrease sharply with prize size (notice that the scales change across the panels) from 0.52 in the first month in the bottom quintile to 0.019 in the top quintile. In the top quintile of prizes, the estimated MPCs are very small. This is

partly simply because the prizes are large - the 0.019 first month MPC still corresponds to an average of 1140 euros of spending. But it is also related to the aforementioned measurement issue - we do not observe big-ticket purchases made with bank transfers which are particularly likely after winning a large prize. A distinctive feature of the MPC paths across the bottom 4 prize quintiles is that they are not persistent: consumption jumps in the first month after winning but fall very rapidly.

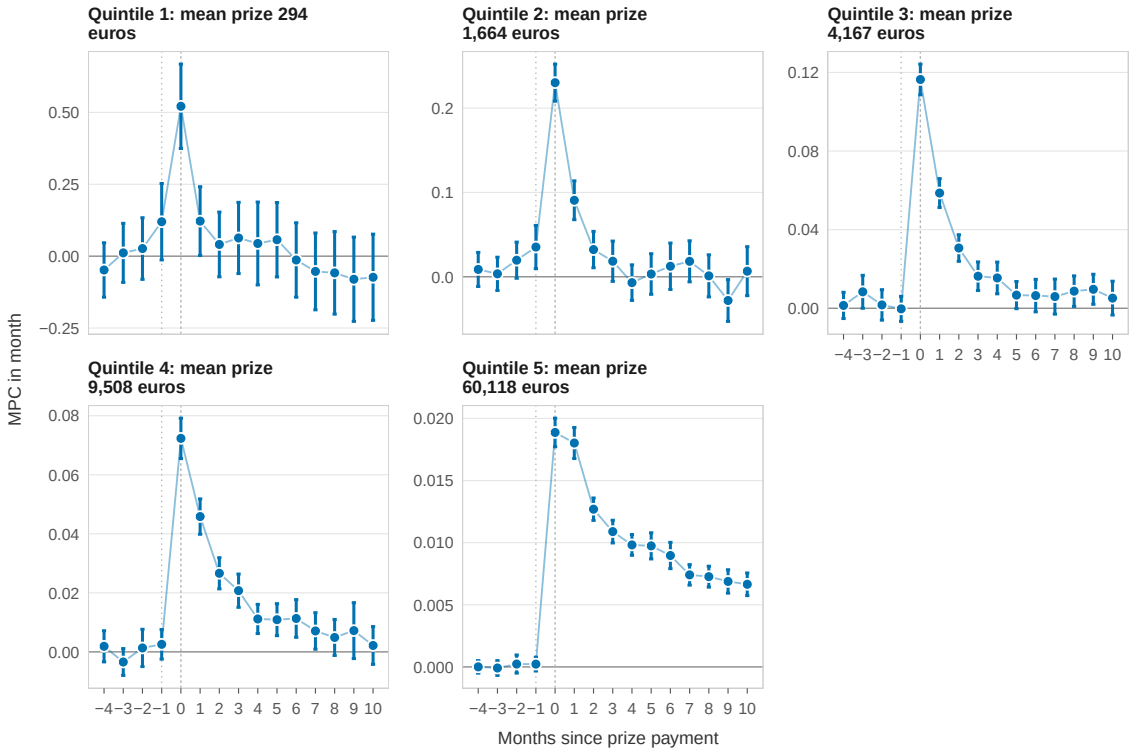


Figure 5: MPC event studies for lottery winnings, shown separately for each quintile of size of lottery win.

To compare to our seasonal bonus MPCs, we estimate lottery MPCs out of prizes which are comparable in size to the seasonal bonuses. Specifically, we retain only prizes which lie between the 5th and the 95th percentiles of the distribution of seasonal bonuses in our data - 657 euros to 4065 euros. The penultimate row of Table 4 show summary statistics for these prizes: we observe 11,131 wins falling in that range, and the mean 2296 euros is close to the mean seasonal bonus, which is just below 2000 euros.

Figure 6 shows the lottery MPCs for this comparable group. Spending jumps up rapidly: the

Table 5: Characteristics of lottery winners and seasonal bonus recipients compared.

Group	N	Age (years)	Monthly salary (euros)
All lottery winners	30,556	59.1	1,180
Prizes matched to seasonal bonuses	11,029	58.8	1,189
Seasonal bonus recipients	107,386	48.3	2,491

first-month MPC is 0.18. But spending also reverts rapidly: by four months after winning, the spending response is not significantly different from zero. This short-and-sharp pattern is consistent with other recent evidence on MPC (Boehm, Fize, and Jaravel (2025)). The cumulative MPC over this period (the first five months) is 0.32. The three month MPC is 0.29.

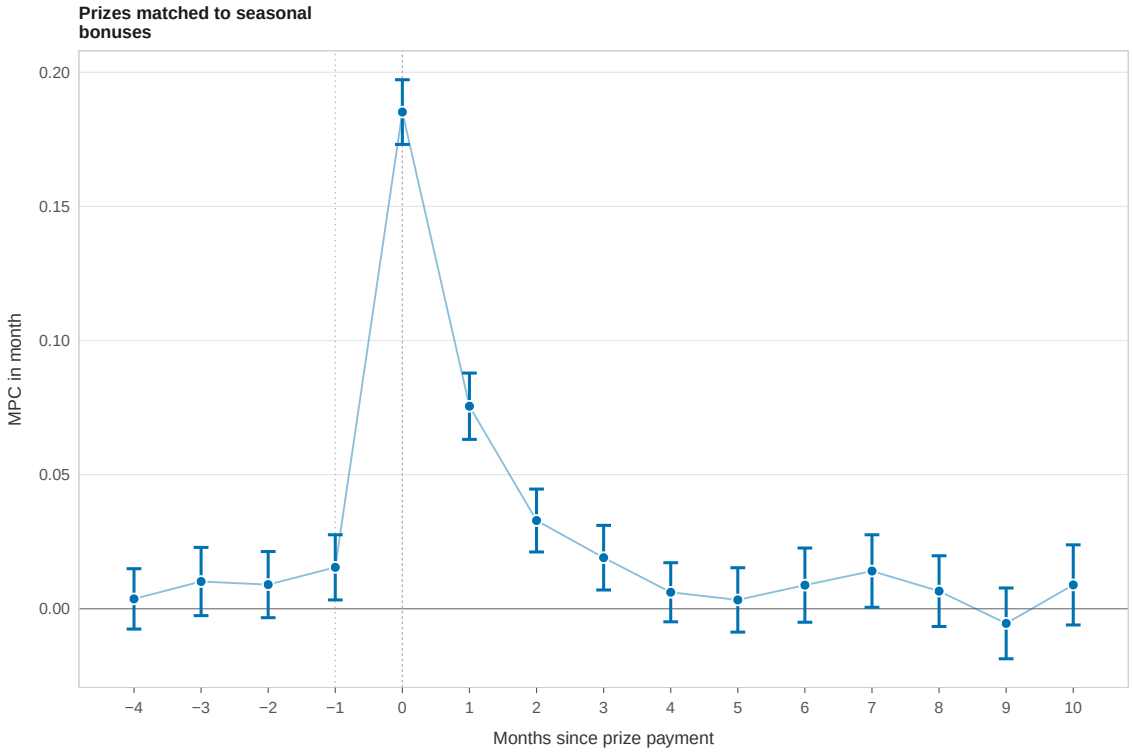


Figure 6: MPC event studies for lottery winnings of comparable size to seasonal bonuses.

5 Comparing MPCs: How Much Does Anticipation Matter?

Based on their own-population estimates, the three-month seasonal bonus MPC (0.21) is about a quarter smaller than the lottery win MPC (0.29), for income changes of comparable size. Given that

these are types of income change on opposite ends of the spectrum in their characteristics, this is a strikingly small difference. There are two reasons which could explain the remaining difference. The first, consistent with standard models of consumption, is that the nature of the income changes is very different: households should respond more to a one-off surprise windfall than they do to income changes which are fully anticipated and recurring. The second is that the population of seasonal bonus recipients is different to the population of lottery winners. The former are mostly full-time workers, the latter are a self-selected group who have chosen to play the lottery, they are older and lower income on average.

In this section, we present further evidence that the MPCs out of these two very different income lumps are similar, and that the 8 percentage point difference in the MPCs is at least as much due to differences in the populations as it is to the characteristics of the income changes. Our primary evidence for this comes from estimating both MPCs for a common set of households, therefore controlling for all unobservable household characteristics. We also show that differences in the distribution of observables across the two populations explains a large part of the difference.

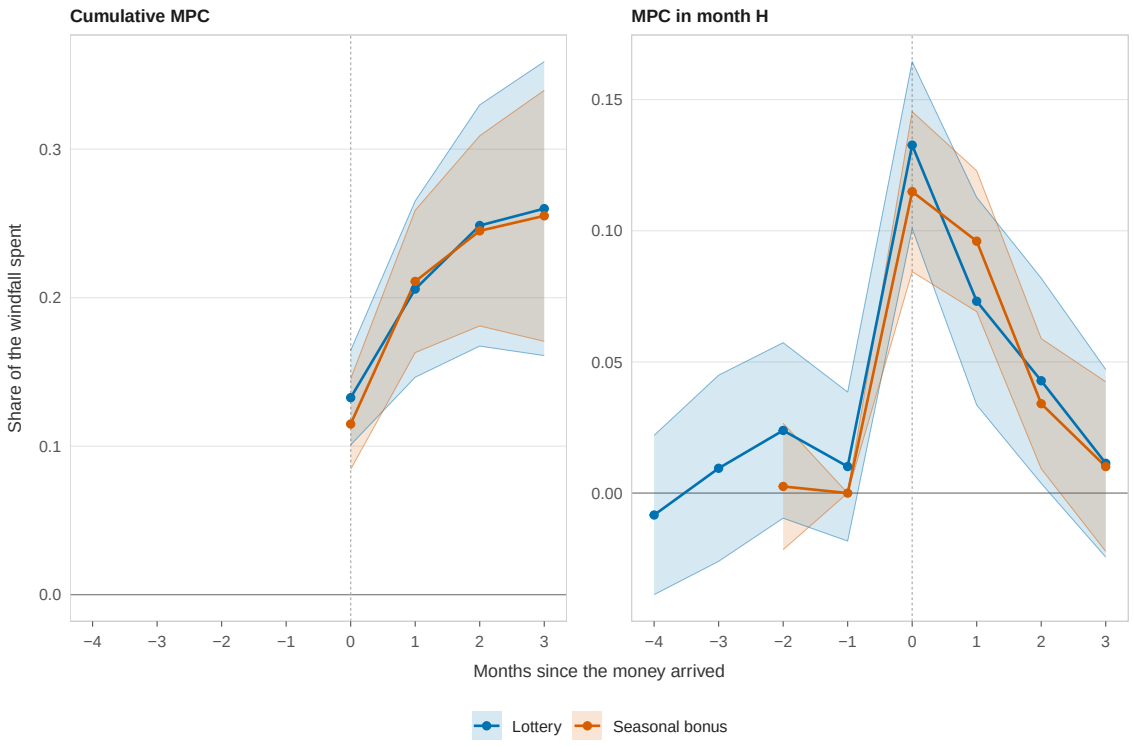


Figure 7: MPC event studies for lottery winnings and seasonal bonus, same households. N = 1078.

Table 6: Characteristics of households who receive both seasonal bonuses and lottery wins.

Variable	N	Mean	Median
Age (years)	1,078	50.8	51.8
Monthly income (euros)	1,078	2,594	2,303
Checking deposits (euros)	1,078	23,231	9,785
Checking deposits / monthly income	1,078	9.03	4.26
Monthly spending (euros)	1,078	2,200	1,940

Exploiting our large dataset, we can estimate both MPCs for the same sample of households. Identifying households which win a seasonal-bonus-sized lottery prize (657 to 4065 euros), and are also observed receiving seasonal bonuses before that lottery win (to avoid contamination), we form a sample of households for whom we can estimate both MPCs. There are 1078 such households. We estimate their seasonal bonus MPC as in the previous section, and their lottery MPC with the same method outlined in this section. Figure 7 shows the results from this exercise. The MPCs are strikingly similar, both in magnitude and shape. Both cumulate to 0.25 over the first three months after impact.

Table 6 provides summary statistics on this group of households: they are, on average, middle-aged, high income, and carry ample liquid assets. These are not, on the whole, hand to mouth households. The fact that their spending out of fully anticipated bonuses is much the same as out of winning the lottery is therefore very striking, and in sharp contrast to the prediction of standard models where spending should respond to news, not the arrival of cash.

Using the same households controls for all unobservable variation, but it leaves us with a small sample. Although the point estimates for these households are very similar, the confidence intervals are wide. Therefore, we now ask how much of the difference between the MPCs can be accounted for by differences in observable characteristics. In doing this, we also study how both MPCs vary with respect to these observables, allowing us to demonstrate that the MPCs are not driven by a small group of hand to mouth households.

To analyse heterogeneity, we stratify our sample with respect to the value of various covariates. To facilitate comparison across the two MPCs, we use common covariate buckets. That is, we use our whole panel - all households we observe, whether or not they get a seasonal bonus or lottery win - to

calculate the cut-offs which define the quantiles of each of our covariates. For example, for income, we estimate the income cut-offs which define the population quantiles of the income distribution for all households, and then we apply these common cut-offs to both our seasonal bonus and our lottery samples. Within each quantile, we then estimate both MPCs. In both cases, we draw both treatment and control households from within that quintile, and we categorize observations by the average value of the covariate in the six months before receiving the bonus or lottery win, with the exception of deposits, where we use the value in the month before.

Figure 8 shows how both MPCs vary across the distribution of income, and a measure of liquidity: the observed current account balance. Both MPCs are somewhat decreasing in income; the gradient is steeper for the seasonal bonus MPC. In the bottom (population) decile of income, there are few seasonal bonus recipients, so the confidence intervals are much larger. Both MPCs remain above 10% in all income deciles, and at or above 20% in all but the top two income deciles.

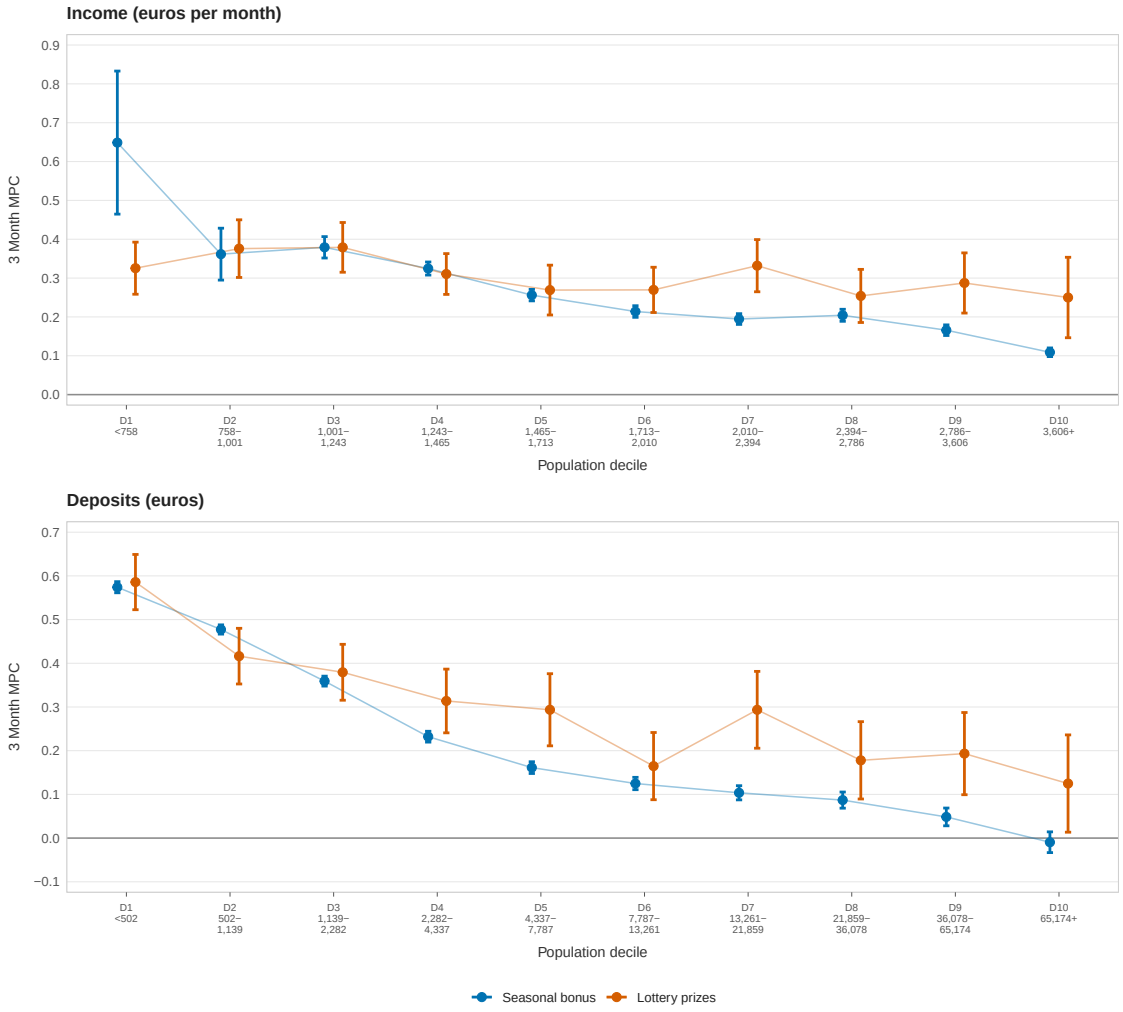


Figure 8: 3-Month MPCs: heterogeneity with respect to income and deposits.

Consistent with much of the previous empirical literature, we find that liquid assets is the covariate which predicts the most variation in MPCs.[6] Both MPCs decline as deposits increase. Lottery MPCs remain above 10% even in the top decile of deposits, while seasonal bonus MPCs drop below 10% only in the 8th decile, and to zero only in the top decile. The two MPCs are very similar at the bottom of the liquidity distribution; they come apart towards the top, but this is where the two populations are particularly different in other respects (see Figure 10 below).

The negative relationship between MPCs and liquid assets is usually interpreted as evidence for the

[6]Data on deposits are available from April 2020 only, so the results stratifying by deposits use only the bonus events and lottery wins after that date.

importance of liquidity constraints. Whilst our evidence is certainly consistent with this story, it also strongly suggests a role for factors beyond liquidity alone. Both MPCs remain positive, and fairly sizeable, across most of the liquidity distribution. From the middle of the distribution and above, the measured levels of deposits are very high: above 13,000 euros from the 7th decile, and this measure - based on current account alone - is a lower bound for the liquid assets of these households. It is not plausible that liquidity is a major driver of the spending responses of households with such high balances.

We also decompose variation in liquidity into persistent and temporary components, in the spirit of Gelman (2021). For each household, we calculate their 'long run' liquidity as the average level of deposits held over the whole of the panel (for lottery winners, the average over all observations before the win). Then we calculate the deviation from that long run level in the month before each bonus receipt or lottery win (for bonus recipients, the deviation from the usual level in the month before a bonus, to remove pay-cycle effects). We then stratify all bonus/lottery events by the long-term liquidity quintile of the household, and within each long-term liquidity quintile, by the terciles of deviation of short-term deposits from the long term level. Figure 9 shows the results from this exercise: it demonstrates that the vast majority of the variation in MPCs is driven by differences in long-term liquidity, rather than short-run deviations. Across most of the long-term liquidity quintiles, for both MPCs, households with temporarily below normal levels of liquidity do not even have higher MPCs than households with normal levels of liquidity. The standard mechanism in buffer-stock models, on the other hand, works entirely through temporary deviations in liquidity following a run of bad income realisations. This evidence is therefore suggestive of fixed household characteristics driving both MPCs and liquidity choices, as emphasised in for example Aguiar, Bils, and Boar (2025).

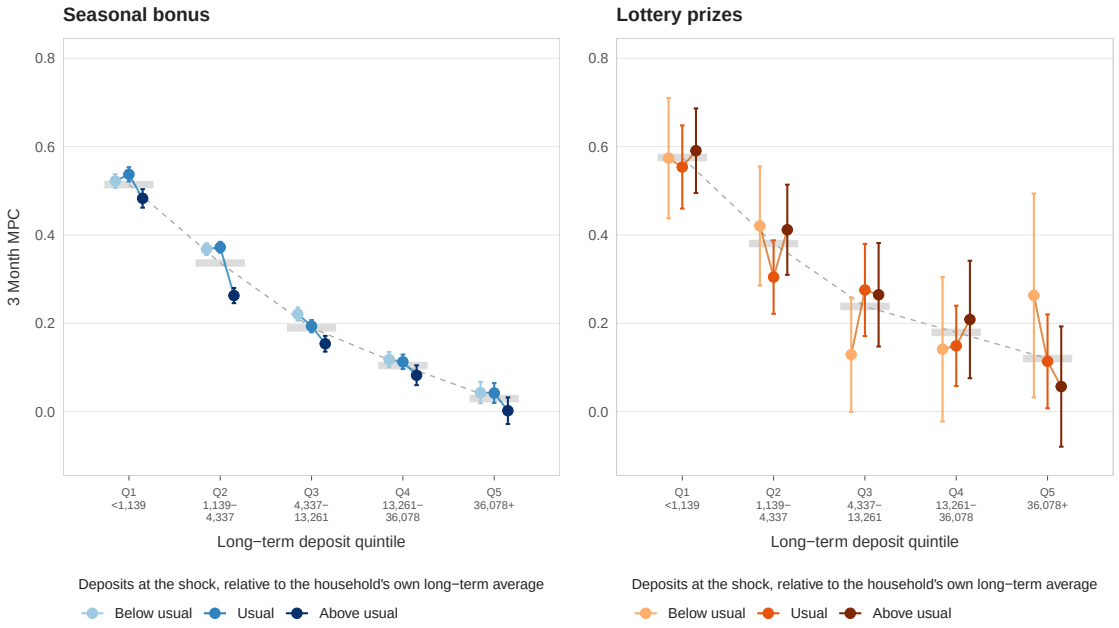


Figure 9: 3-Month MPCs and permanent versus transitory liquidity. Below usual: 1st tercile of deviation from long term liquidity; Usual: middle tercile; Above usual: 3rd tercile.

The seasonal bonus and lottery MPCs separate towards the top of the liquidity distribution: lottery MPCs stay positive whilst seasonal bonus MPCs eventually go to zero in the top 10%. But it is important to emphasize that even conditional on liquidity, lottery winners remain quite different to seasonal bonus recipients. Figure 10 shows that within liquidity quintiles, lottery winners are older, derive much less of their income from salaries (and more from pensions), and have higher levels of both spending and spending relative to (total) income. It seems likely, both a priori and based on this evidence, that playing the lottery is correlated with underlying traits which also correlate with MPCs.

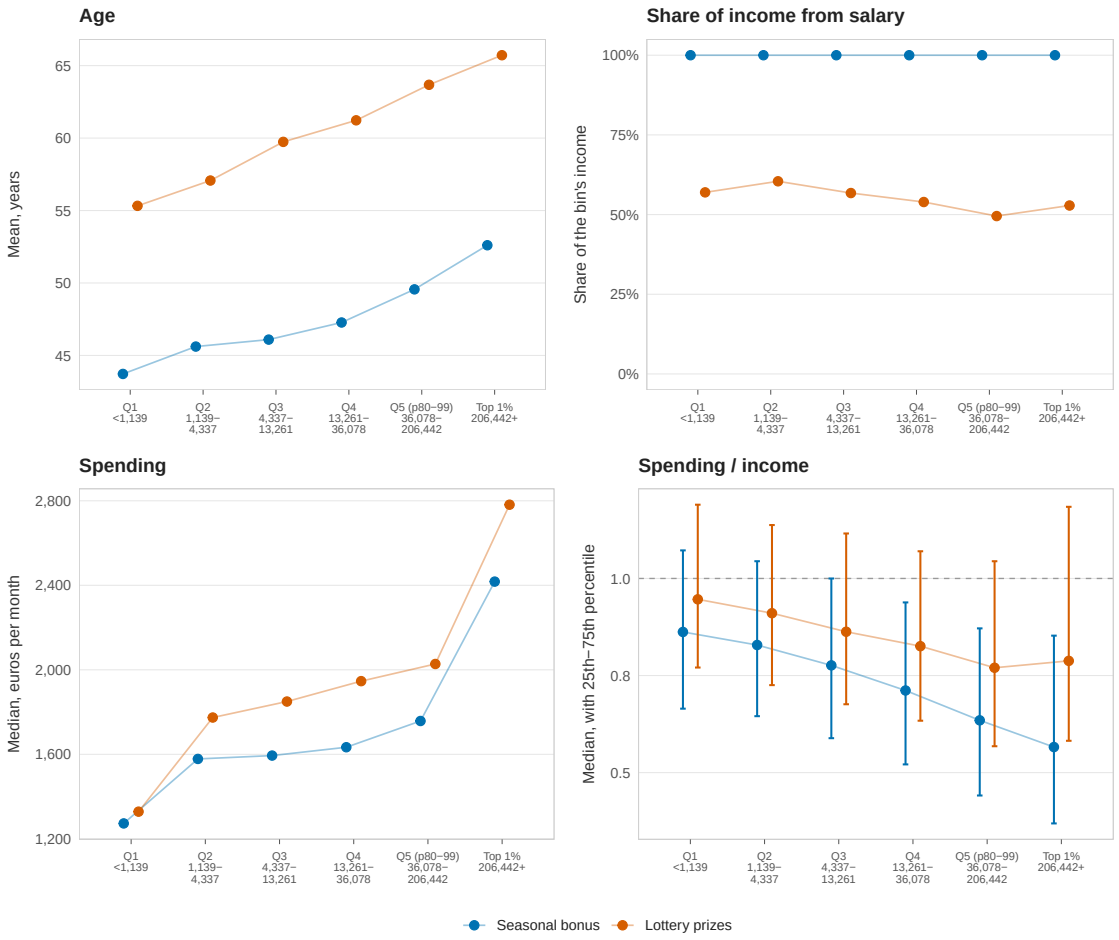


Figure 10: Comparing covariate values of seasonal bonus recipients and lottery winners across deposit quintiles.

Figure 11 shows how the MPCs vary across the distribution of age, and the ratio of (measured) spending to income - a proxy for the average propensity to consume (APC). The latter is calculated using the average of spending and income over the whole of the panel (over all pre-win observations for lottery wins), so it reflects a long-run characteristic of households. Higher age weakly predicts lower MPCs. No seasonal bonus recipients sit in the top two deciles of the overall age distribution, because these households are aged 73 and above, and only employed workers get seasonal bonuses. Both MPCs are also positively correlated with the spending to income ratio. It is possible that this partly reflects measurement: for some households we see a larger fraction of spending, and therefore a higher MPC. But this interpretation would not explain why the profile is much flatter for lottery

MPCs.

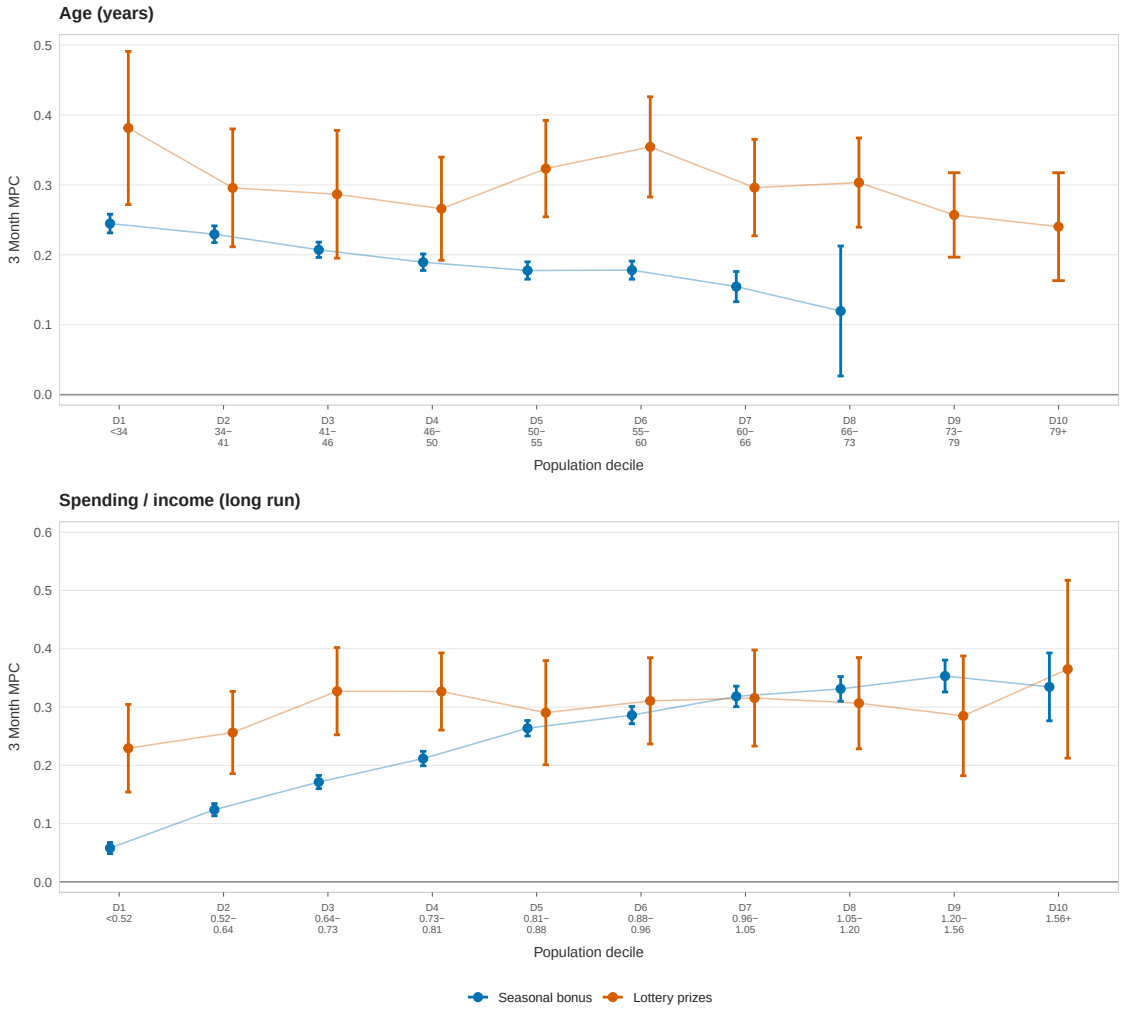


Figure 11: 3-Month MPCs: heterogeneity with respect to age and long term spending to income ratio.

How much of the difference in the aggregate, own-population MPCs is explained by covariates? We divide the seasonal bonus observations into cells according to the population quintiles of deposits, age, and spending relative to income. From this, we calculate weights based on the relative frequency of seasonal bonus and lottery observations in each cell. Using these weights, we re-estimate the seasonal bonus MPC, weighting it to reflect the same population distribution as the lottery winners (this is very similar to the procedure proposed by DiNardo, Fortin, and Lemieux (1996)). Because the age distributions do not share a common support, we restrict our attention to the subset of lottery

winners (with seasonal bonus-sized wins) between the ages of 25 and 65 and with no pension income. The three month lottery MPC for this group overall is 0.30, very close to the overall estimate.[7]

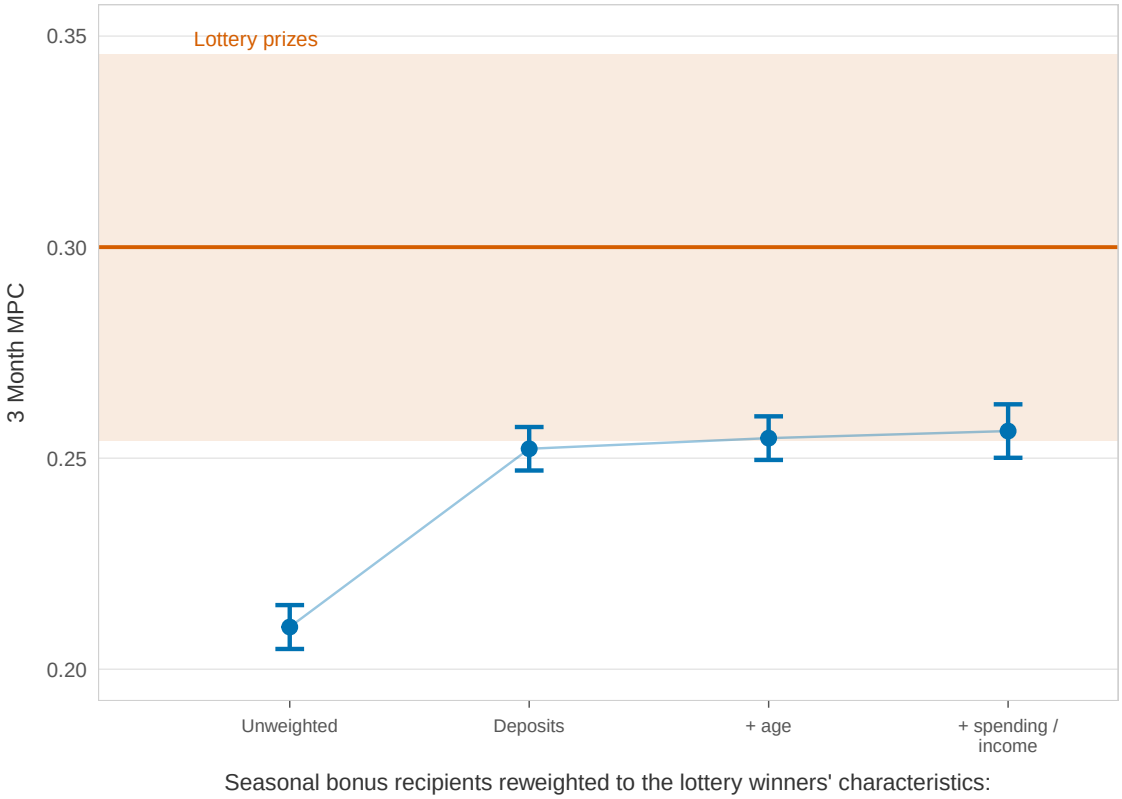


Figure 12: Re-weighting seasonal bonus MPCs to lottery covariate distribution.

Figure 12 presents the results from this re-weighting exercise. The different distributions of these three covariates explains about half of the total difference between the MPCs. Most of this comes from deposits: lottery winners hold less liquidity on average. There are likely other, unobservable, differences between the two groups of household. Lottery winners have chosen to enter a contract with negative expected value, and the propensity to do this is plausibly correlated with deviation from textbook optimizing behaviour.

Taking stock, we can summarize the key points in our evidence on MPCs as follows:

[7]Appendix A.9 shows that this procedure is robust to alternative methods for re-weighting including a propensity score method. Changes are estimated on bonus events from April 2020 with all covariates observed, and are applied to the full-sample estimate, because data on deposits are unavailable before April 2020.

1) The average three-month MPC out of seasonal bonuses is 0.21.
2) The average three-month MPC out of size-matched lottery winnings is 0.29.
3) For the same households, the two MPCs are indistinguishable.
4) Re-weighting by observables explains half of the total gap between the two MPCs.
5) Both MPCs have very low persistence: consumption reverts fully within 3-4 months.
6) Both MPCs are declining as deposits increase.
7) But both MPCs remain substantial across all of the income distribution, and most of the liquidity distribution.

These results are striking. Regular, anticipated, recurring seasonal bonuses could not be more different as a type of income change to winning the lottery. Yet our evidence suggests that the response of household spending is very similar. Cash flows matter, not information.

Moreover, this is not a result driven only by hand to mouth households. A model where some fraction of households are rule-of-thumb and simply spend all of their income each period would trivially predict identical MPCs for these two types of income. But our evidence shows both MPCs being positive across all of the income distribution and most of the liquidity distribution.

In the next section, we turn our attention to analyzing these two different MPCs in theoretical models of consumption. We show that our results are very helpful for adjudicating between models: a wide range of models are not able to replicate the behaviour we observe.

6 Buffer Stock Models

In this section, we assess the extent to which theoretical models can match the empirical facts documented above. It is well-known that in models with no uncertainty, full insurance, or quadratic utility and no borrowing constraints, consumption should not respond at all to fully anticipated changes in income, and should respond only very marginally to income shocks (Jappelli and Pistaferri (2017)). These are cases of the classic permanent income hypothesis result. In light of the accumulation of empirical evidence against these predictions, a large literature has emerged using buffer-stock models - where uninsurable idiosyncratic risk, borrowing constraints, and prudence are added - to explain why households would have higher MPCs out of income shocks. Much of the focus in this

literature has been on the MPC out of unanticipated income shocks, where borrowing constraints and precautionary saving both play a role.

We start with a standard buffer-stock model modified to match our seasonal-pay setting. We show analytically, and confirm by simulation, that this model cannot produce an MPC out of seasonal bonuses anywhere near the size we estimate without a discount factor so low that households hold almost no wealth. This conclusion survives a number of extensions: seasonal preferences, expense shocks, richer income processes, and adding a second (illiquid) asset.

6.1 A baseline buffer stock model

We start our theoretical analysis with a simple one-asset buffer stock consumption model. This is the model of Deaton (1991) and C. D. Carroll (1997), and the household block of modern quantitative heterogeneous-agent models (Kaplan, Moll, and Violante (2018); Auclert, Rognlie, and Straub (2024)). Its predictions for how consumption responds to *unanticipated* income shocks are well understood: borrowing constraints and a precautionary motive together generate MPCs that are substantial for households with little liquid wealth and that decline with the buffer, averaging perhaps 0.15–0.25 in quantitative work once the model is calibrated to realistic liquid balances. Our aim in this section is to show the model's predictions for seasonal bonus MPCs. Seasonal bonuses are neither unanticipated nor one-off, and it is not obvious what the same model implies for the response to a payment that is large, perfectly foreseen, and received on the same schedule every year. This subsection analyzes that question analytically, before we turn to simulations.

Households choose an infinite stream of consumption c_t to maximize expected utility. Their period utility function is constant relative risk aversion (CRRA), and importantly exhibits prudence ($u'''(c) > 0$) so that there is a precautionary motive for saving. They can save a_t in a single, risk-free asset which pays an exogenous and constant interest rate r, with corresponding gross interest rate $R = 1 + r$. There is a borrowing constraint at zero: households cannot borrow. Finally, households receive an income z_t which is subject to idiosyncratic risk: the logarithm of z_t evolves according to an AR(1) process. We will always calibrate the model such that $\beta R < 1$, so that buffer-stock behaviour emerges. Each period in the model is one month, and households have rational expectations - they understand the true distribution of the exogeneous state variables.

The household problem is:

$$\max_{\{c_t\}_0^\infty} E_0 \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$$

subject to

$$c_t + a_{t+1} = Ra_t + z_t$$

$$a_{t+1} \geq 0,\ a_0 = 0$$

$$log(z_{t+1}) = \rho log(z_t) + e_t,\ e_t \sim N(0, \sigma_e^2)$$

$$\beta R < 1$$

To this standard model, we add seasonal pay: we specify two types of household corresponding to the two pay schemes in our data. These households differ only in their budget constraints: for the group g = flat, the income process is exactly as specified above for the baseline model. On the other hand, for g = bonus, income doubles every sixth period (each period in the model represents a month). This doubling is entirely deterministic, and therefore fully anticipated by these households in advance.

More precisely (and dropping subscripts for notational simplicity):

$$c + a' = Ra + \kappa_{j,g} \cdot z$$

$$log(z') = \rho log(z) + e,\ e \sim N(0, \sigma_e^2)$$

For flat pay households: $\kappa_{j,flat} = 1$ always.

For bonus households:

$$\kappa_{j,bonus} = \begin{cases} 2w^* & \text{if } j \bmod 6 = 0 \\ w^* & \text{otherwise} \end{cases}$$

The multiplier w^* is normalized so that the present discounted value of income is identical across groups so that comparisons we make are comparing like-with-like.[8]

We calculate two different MPCs. The first is what we call the MPC-SB: the MPC out of a seasonal

[8] In our data, the bonus group have higher annual incomes on average. Carrying this feature over to the model would make bonus households hit their borrowing constraint less often, making MPCs in the buffer stock model even lower. Therefore this assumption is conservative with respect to our main conclusion.

bonus. This is calculated by solving our model for the bonus pay and the flat pay households, and then performing the same difference-in-differences calculation which we used on our data in the empirical section. Intuitively, this simply compares the growth rates between the bonus households and the non-bonus households around the months when bonuses are received. Let $\bar{C}_j^g$ denote mean consumption of group $g \in \{\text{b}, \text{f}\}$ (b for bonus group, f for flat) in month j of the seasonal cycle, in the model's stationary (cycle-periodic) distribution, and let month T be a bonus month. The model counterpart of our empirical MPC-SB is:

$$\text{MPC-SB}(T) = \frac{\left(\bar{C}_T^{\text{bon}} - \bar{C}_{T-1}^{\text{bon}}\right) - \left(\bar{C}_T^{\text{flat}} - \bar{C}_{T-1}^{\text{flat}}\right)}{w^*}$$

We also calculate what we call the MPC-IS. This is the standard MPC out of an unanticipated income shock, calculated as the numerical derivative of the consumption policy function. This is the model counterpart of our estimated MPCs out of lottery winnings, as these are entirely unanticipated windfalls. Due to the concavity of the consumption function, MPCs depend on the size of the windfall, so for comparability with our MPC-SB, we calculate the MPC-IS out of an income shock of size $z.w^*$. This means that the average size of the transfer is the same as the seasonal bonuses, and the windfall scales with income, as in the case of the seasonal bonuses.

$$\text{MPC-IS}(a, z, j) = \frac{c^*(a + z.w^*, z, j) - c^*(a, z, j)}{E(z).w^*}, \; E(z) = 1$$

We report the average MPC-IS, summing over the stationary distribution of the state variables:

$$\text{MPC-IS} = \int_a \int_z \int_j \text{MPC-IS}(a, z, j) dA \; dZ \; dJ$$

For the bonus group, the MPC-IS depends on the model period J: it is lower in the months when bonuses are received and higher in the months before. In our results tables we report the MPC-IS for the flat pay group, but this is never significantly different to that of the bonus group.

6.2 A restriction on the shape of buffer-stock MPC paths

Here we show that in this buffer stock model, the path that MPCs can follow over time is restricted. Most of the literature has focused on how consumption can jump up in the models upon receipt of a windfall: this is one of the core results in the model. Here we focus on the path of consumption **after** getting a windfall, which we show is restricted by the Euler inequality, with implications for the shape of consumption responses to lottery winnings and for the size of consumption responses to seasonal bonuses.

The idea that the rate at which consumption declines after a payment is informative about impatience goes back to Shapiro (2005), who estimates a daily discount factor based on the rate at which calorie consumption falls over the month after receipt of food stamps, showing that the observed rate of decline implies something like quasi-hyperbolic discounting. We use similar logic to derive a bound that holds for constrained households, via the Euler inequality, aggregates across households under precautionary saving, via Jensen's inequality, restricts the path of iMPCs, and, combined with periodic stationarity, bounds the size of MPCs out of recurring payments.

The first order conditions of the model imply the well-known Euler inequality:

$$u'(c_t) \;\geq\; \beta R\, E_t\left[u'(c_{t+1})\right].$$

For an unconstrained household, the marginal utility of consuming this month, $u'(c_t)$ should be equalized with the expected marginal utility of saving, increasing resources by a factor R, and consuming next month: $\beta R\, E_t\left[u'(c_{t+1})\right]$. However, a constrained household may not be able to achieve this if their borrowing constraint is binding, meaning their consumption this month is restricted and the marginal value of consumption this month consequently higher than the marginal value of saving. The reverse cannot occur: if the marginal value of saving is greater than the marginal value of consuming today, then there is nothing to stop the household from saving more, and they will do so.

The Euler inequality implies a restriction on how fast consumption can fall in this model. When combined with the periodic nature of bonus events, this leads to a restriction on the size of the MPC-SB.

Lemma 1: mean consumption cannot fall quickly. Define

$$R_\beta \;\equiv\; (\beta R)^{1/\sigma}.$$

Because $\beta R < 1$ and $\sigma > 0$, we have $R_\beta < 1$. Under the Euler inequality and CRRA utility, for every household and date,

$$E_t\big[c_{i,t+1}\big] \;\geq\; R_\beta\, c_{it}.$$

Derivation. Write the Euler inequality with $u'(c) = c^{-\sigma}$:

$$E_t\big[c_{t+1}^{-\sigma}\big] \;\leq\; \frac{c_t^{-\sigma}}{\beta R}.$$

Now apply Jensen's inequality to the function $g(x) = x^{-1/\sigma}$, which is convex and decreasing on $x > 0$. Convexity gives

$$E_t[c_{t+1}] \;=\; E_t\big[g\big(c_{t+1}^{-\sigma}\big)\big] \;\geq\; g\big(E_t\big[c_{t+1}^{-\sigma}\big]\big),$$

Since g is decreasing, applying it to both sides of the Euler inequality reverses the inequality:

$$g\big(E_t\big[c_{t+1}^{-\sigma}\big]\big) \;\geq\; g\Big(\frac{c_t^{-\sigma}}{\beta R}\Big)$$

Combining these leads to:

$$E_t[c_{t+1}] \;=\; E_t\big[g\big(c_{t+1}^{-\sigma}\big)\big] \;\geq\; g\big(E_t\big[c_{t+1}^{-\sigma}\big]\big) \;\geq\; g\Big(\frac{c_t^{-\sigma}}{\beta R}\Big) = (\beta R)^{1/\sigma}\, c_t = R_\beta\, c_t \qquad \square$$

Taking expectations across all states for a group of households, assuming rational expectations:

$$E\big[\,E_t[c_{i,t+1}]\,\big] \;=\; \bar{C}_{t+1},$$

where the outer expectation is over the cross-sectional distribution of household states at date t. Averaging next month's forecasts across households, using this month's distribution of states, gives next month's realised cross-sectional mean.

Combining this with Lemma 1 leads to a restriction on the average growth rate of all households.

$$\bar{C}_{t+1} \;\geq\; R_\beta \bar{C}_t, \qquad \text{and iterating,} \qquad \bar{C}_{t+s} \;\geq\; R_\beta^s \bar{C}_t \quad \text{for all } s \geq 0.$$

The growth rate of mean consumption is unbounded above. Borrowing constraints (via the Euler inequality) and precautionary saving (via the Jensen's inequality step above) can increase the growth rate above R_β. But the growth rate is bounded below: forward-looking buffer stock consumers will never choose a consumption path which falls at a rate greater than R_β each period. This restriction on the rate at which consumption can fall, in turn allows us to derive a bound on the shape of the impulse response function of consumption to an income shock. Notice the role of precautionary saving: it is one of the forces, via the Jensen's inequality step, which makes the inequality non-binding. The stronger the precautionary motive, the further the distance between this inequality and the actual rate at which consumption falls. The consequence of this is that the bound we derive below will be further from binding the more precautionary motive there is.

Suppose a treated group receives a one-off windfall of x per capita at date T (it can be anticipated or unanticipated), and write their mean consumption at horizon s in terms of potential outcomes: $\bar{C}_{T+s}(1)$ is mean consumption in the state of the world in which the windfall is received at T, and $\bar{C}_{T+s}(0)$ is mean consumption for the same households in the counterfactual state in which it is not. The impulse response of consumption to the windfall (Δ_s) and the associated path of MPCs (θ_s) are:

$$\Delta_s \;\equiv\; \bar{C}_{T+s}(1) - \bar{C}_{T+s}(0), \qquad \theta_s \;\equiv\; \frac{\Delta_s}{x},$$

Lemma 1 holds along both potential-outcome paths: each is an optimal consumption path for the same households in the respective state of the world, and the one-sided Euler inequality it is derived from requires only that a household can freely save at gross return R. Applying it along the treated path from T to $T + s$

$$\bar{C}_{T+s}(1) \;\geq\; R_\beta^s \bar{C}_T(1), \qquad R_\beta \equiv (\beta R)^{1/\sigma}.$$

Substituting $\bar{C}_T(1) = \bar{C}_T(0) + \Delta_0$ gives:

$$\Delta_s \geq R_\beta^s\left(\bar{C}_T(0) + \Delta_0\right) - \bar{C}_{T+s}(0).$$

The first term is a lower bound on the level of consumption in the treated group: it cannot fall faster than a geometric decay at rate R_β^s from its level in the period the windfall arrives, $\bar{C}_T(0) + \Delta_0$. The second term traces out the counter-factual no-bonus outcome.

To translate this result into MPC space, divide by x:

$$\frac{\Delta_s}{x} \equiv \theta_s \geq R_\beta^s \frac{\bar{C}_T(0) + \Delta_0}{x} - \frac{\bar{C}_{T+s}(0)}{x}.$$

This is Proposition 1.

Proposition 1. In any model to which Lemma 1 applies, the path of MPCs for $s \in 0, 1, 2, ...$ is bounded below by:

$$\theta_s \geq R_\beta^s\left(\theta_0 + k_0\right) - k_s, \qquad k_s \equiv \frac{\bar{C}_{T+s}(0)}{x}.$$

This result places a joint restriction on the size of the impact MPC and the horizon over which it dies down. Because consumption can decline at a rate of at most R_β per month, a large impact MPC requires a long horizon s to fully dissipate. Another way to view this proposition is that it places a restriction on the permissible shape of the intertemporal MPC matrix (Auclert, Rognlie, and Straub (2024)). An immediate implication is that MPCs (out of any kind of windfall) cannot be both large and short-lived. This is why the response of consumption to income shocks in HANK models is very persistent: geometric discounting does not permit any other shape.

6.2.1 Implications for lottery wins

We can take Proposition 1 to the data, and show that the buffer stock model is inconsistent with our estimates in the previous section. Specifically, we show that given our estimate of the impact MPC θ_0, our estimates of the subsequent MPC dynamic path $(\theta_1, \theta_2, ...)$ are far too low. To do this, we need estimates of the counter-factual no-treatment path of consumption, $\bar{C}_{T+s}(0)$. Fortunately,

our identification assumptions identify this path (indeed any strategy to identify causal effects must do this). The parallel trends assumption supplies a model for the no-treatment potential outcomes: $E[C_{it}(0)] = \alpha_i + \lambda_t$, and the no anticipation assumption allows us to estimate this factor model using not yet treated observations.

Under the imputation estimator of Borusyak, Jaravel, and Spiess (2024), which is what we used, the counterfactual is estimated explicitly. The fixed effects $\hat{\alpha}_i$ and $\hat{\lambda}_t$ are fitted on the untreated observations, each treated observation's no-treatment potential outcome is imputed as $\hat{C}_{it}(0) = \hat{\alpha}_i + \hat{\lambda}_t$, and the estimated effect is $\hat{\tau}_{it} = C_{it} - \hat{C}_{it}(0)$. All objects in Proposition 1 are therefore estimated already:

$$\hat{\theta}_s \;=\; \frac{\hat{\tau}_s}{x}, \qquad \hat{k}_s \;=\; \frac{1}{x}\sum_{it\in\Gamma(s)} w_{it}\hat{C}_{i,T_i+s}(0) \;=\; \frac{1}{x}\sum_{it\in\Gamma(s)} w_{it}\big(C_{i,T_i+s} - \hat{\tau}_{i,T_i+s}\big),$$

where $\Gamma(s)$ is the set of observations corresponding to treated units at horizon s after treatment, and w_{it} are weights, in our case equal to $1/N$ (assuming x is the average lottery size). In Figure 13 below, we use our estimate of the impact MPC, $\hat{\theta}_0$, and then trace out for several different values of β the minimum permitted value of $\theta_1, \theta_2, \ldots$ under Proposition 1, according to:

$$\theta_s \;\geq\; R_\beta^s\big(\hat{\theta}_0 + \hat{k}_0\big) \;-\; \hat{k}_s$$

We assume r corresponds to a 2% annual rate and that $\sigma = 1.5$.

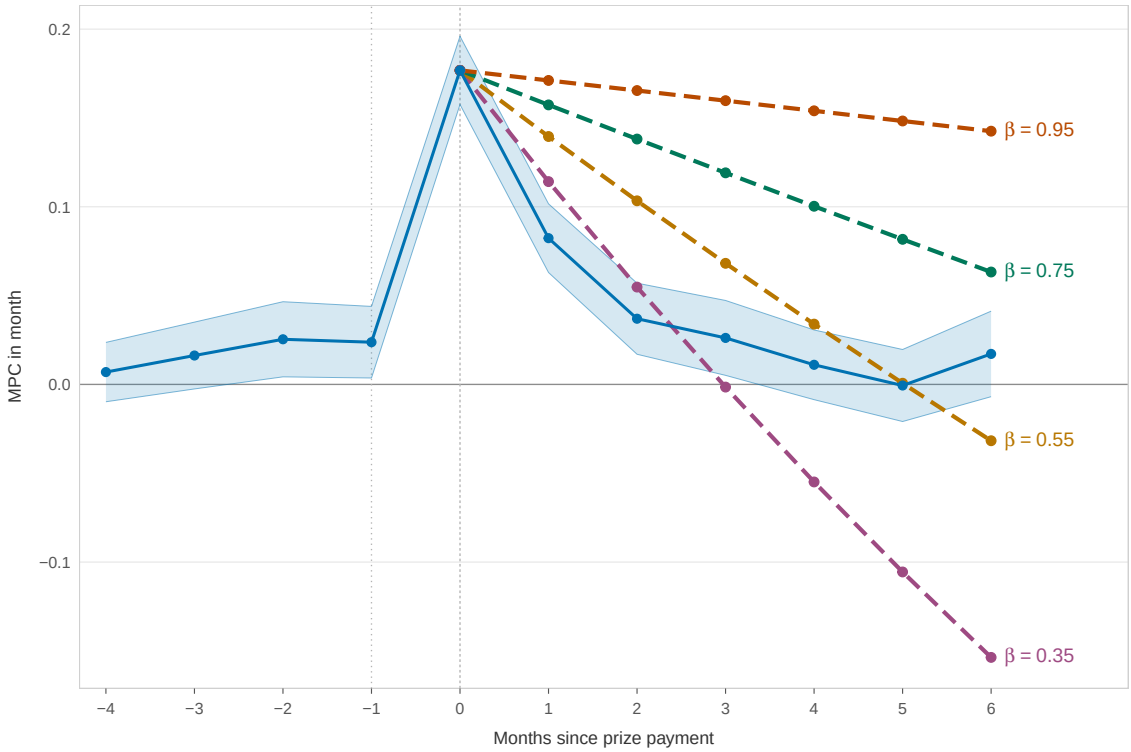


Figure 13: Path of MPC out of a lottery winnings, with the theoretical minimum path given the MPC at t=0 traced out for different discount factors.

Figure 13 shows our empirical estimates of the path of MPCs out of lotteries of comparable size to the seasonal bonuses, with the bound implied by Proposition 1 traced out for several values of β. This illustrates the result clearly: for an impact MPC of just 0.2, a typical annual discount factor like $\beta = 0.95$ implies a very gradual drift downwards (and a very large cumulative MPC). Even at $\beta = 0.35$, the empirical path is somewhat below the theoretical bound. Our empirical evidence on lottery win MPCs is therefore quite inconsistent with the intertemporal MPC path implied by the buffer stock model.

This is not an esoteric point: the persistence of consumption responses to an income shock is clearly just as important for the overall magnitude as the impact effect. Indeed, Auclert, Rognlie, and Straub (2024) show that it is the full dynamic path of intertemporal MPCs (iMPCs) which determines the impact of a fiscal stimulus. Heterogeneous agent models are often calibrated by choosing β to match the impact MPC, with the model tracing out the remaining iMPC profile. The fact that the buffer stock model predicts much more persistent consumption responses than we see empirically means

this calibration approach will significantly over-state the total consumption response.

6.2.2 Implications for seasonal pay

The restriction on impulse-response shape implied by Proposition 1 applies equally to the seasonal bonus MPC. In fact, in the context of seasonal pay, Proposition 1 implies a restriction on the size, not just the shape, of the MPC. Because the seasonal bonuses are recurring and fully anticipated, our model exhibits periodic stationarity: although consumption evolves within the six months of the bonus cycle, the cycle repeats and so the distribution of consumption in both groups is stationary from cycle to cycle. Periodic stationarity in turn places a sharp restriction on the persistence of the consumption response to receiving the seasonal bonus: it must have fully died down within six months, because average consumption must return to the same place: $\bar{C}^g_{t+6} = \bar{C}^g_t$. Combined with Proposition 1, which restricts how fast consumption can fall, periodic stationarity restricts the overall size of the MPC-SB. If mean consumption must return to its periodic stationary value within six months after the bonus arrives, but cannot fall rapidly, then it cannot jump up very much when the bonus arrives. We make this precise in Theorem 1.

Theorem 1. In the periodic stationary distribution, with an aseasonal control group, the MPC-SB is bounded.

$$\text{MPC-SB} \;\leq\; \frac{\bar{C}^b_{T-1}}{w^*}\left(R_\beta^{-5} - 1\right), \qquad R_\beta = (\beta R)^{1/\sigma}.$$

Derivation. Recall that we defined the MPC-SB by consumption in the bonus group relative to the flat pay group (with the latter serving, in the model, to impute the counterfactual of what the bonus group would have consumed if they were paid the same present value in even instalments):

$$\text{MPC-SB}(T) = \frac{\left(\bar{C}^{\text{b}}_T - \bar{C}^{\text{b}}_{T-1}\right) - \left(\bar{C}^{\text{f}}_T - \bar{C}^{\text{f}}_{T-1}\right)}{w^*}$$

By assumption, the flat group's consumption is constant over time, so $\bar{C}^{\text{f}}_T - \bar{C}^{\text{f}}_{T-1} = 0$.

By Lemma 1, $\bar{C}^b_{T+5} \geq R_\beta^5 \bar{C}^b_T$, and by periodic stationarity, $\bar{C}^b_{T+5} = \bar{C}^b_{T-1}$. Combining these, $\bar{C}^b_{T-1} \geq R_\beta^5 \bar{C}^b_T$, and inverting this, $\bar{C}^b_T \leq \bar{C}^b_{T-1} R_\beta^{-5}$. Therefore, $\bar{C}^{\text{b}}_T - \bar{C}^{\text{b}}_{T-1} \leq \bar{C}^b_{T-1} R_\beta^{-5} - \bar{C}^b_{T-1}$.

Plugging into the expression for the MPC-SB, the result follows. □

We can extend the result to provide a bound on the 3-month MPC-SB. Applying Lemma 1 from $T+h$ to $T+5$, a span of $5-h$ months, gives $\bar{C}^b_{T+h} \leq R_\beta^{-(5-h)} \bar{C}^b_{T-1}$ for $h = 0, 1, 2$. Summing the three over the common $T-1$ baseline gives the bound on the three-month cumulative MPC-SB:

$$\text{MPC-SB}(3m) \;\leq\; \frac{\bar{C}^b_{T-1}}{w^*}\Big[\left(R_\beta^{-5}-1\right) + \left(R_\beta^{-4}-1\right) + \left(R_\beta^{-3}-1\right)\Big]$$

The tightness of this bound depends crucially on the size of the bonus relative to consumption, $\frac{\bar{C}^b_{T-1}}{w^*}$. Consider a household which is constrained before they get the bonus. If the bonus is small relative to their consumption, they might consume it all, and remain constrained afterwards. If the bonus is large, a consumption-smoothing household will not spend it all in one period, they will spread it out, with the rate of spending down constrained by the Euler inequality.

How tight is this bound quantitatively? In the model, the bonus w^* is equal to $6/7$ of average monthly income, so if average consumption was equal to average monthly income, then $\frac{\bar{C}^b_{T-1}}{w^*} = 7/6$. At an annual discount rate of $\beta = 0.95$, annual interest rate of 2%, and $\sigma = 1.5$, $R_\beta = 0.998$ and the bound would be that the one-month MPC-SB cannot exceed 1%, and the three month cannot exceed 2.5%. In the data, $\frac{\bar{C}^b_{T-1}}{w^*}$ is roughly 0.9, so the implied bound is even tighter.

Although this result depends on $\bar{C}^b_{T-1}$, an endogenous model object, in practice this does not vary much across calibrations. In the stationary equilibrium, mean assets are the same at each point of the pay cycle year after year: the bonus group builds a buffer when the bonus arrives and draws it down before the next one, but there is no net accumulation from one cycle to the next. Averaging across households and summing the budget constraint from T to T+5 then implies that mean consumption equals mean income plus the interest earned on assets. From this we can derive an upper bound on $\bar{C}^b_{T-1}$ as a function of assets:[9]

$$\bar{C}^b_{T-1} \;\leq\; \Phi\left(1 + r\,\bar{a}^b_T\right), \qquad \Phi \;\equiv\; \frac{\sum_{s=0}^{5} R^{-s}}{\sum_{s=0}^{4} R^{-s} R_\beta^{s+1} \;+\; R^{-5}}$$

[9]The full derivation is in Appendix A.10.

Table 7: Upper bound on the MPC-SB at different discount factors

β (annual)	Φ	MPC-SB(1m), A/Y = 10	MPC-SB(3m), A/Y = 10	MPC-SB(1m), A/Y = 60	MPC-SB(3m), A/Y = 60
0.975	1.001	0.002	0.004	0.002	0.005
0.950	1.004	0.010	0.025	0.011	0.027
0.900	1.012	0.029	0.069	0.031	0.075
0.850	1.020	0.049	0.117	0.053	0.127
0.800	1.028	0.071	0.169	0.077	0.183
0.750	1.038	0.095	0.227	0.103	0.245
0.700	1.047	0.122	0.290	0.132	0.314
0.650	1.058	0.152	0.361	0.164	0.390
0.600	1.069	0.185	0.440	0.201	0.476
0.550	1.082	0.224	0.529	0.242	0.572
0.500	1.096	0.267	0.632	0.289	0.683

And therefore we can re-express the bound on the MPC-SB in terms of $\bar{a}_T^b$:

$$\text{MPC-SB}(3m) \;\leq\; \Phi \cdot \frac{1 + r\,\bar{a}_T^b}{w^*} \sum_{h=0}^{2} \left(R_\beta^{-(5-h)} - 1 \right)$$

The only endogenous object in this bound is $\bar{a}_T^b$, and importantly the bound is not sensitive to this number as it appears multiplied by the monthly interest rate. In all of our quantitative exercises, we use an annual real interest rate of 2%, implying a monthly rate of $r = (1.02)^{1/12} - 1 \approx 0.0017$. In Table 7, we show how this bound depends on the discount factor for two different values of total assets. We assume $\sigma = 1.5$. At $\beta = 0.95$ the three month bound only loosens by 0.002 when raising assets from 10 to 60 months of income.

The table illustrates that this bound is very tight for high discount factors, indeed as $\beta R \longrightarrow 1$ the bound approaches zero exactly. At $\beta = 0.95$, a fairly typical value, the 3 month MPC-SB cannot exceed 2.5%. Taking our empirical estimate that the 3 month MPC-SB is 0.21 and inverting Theorem 1 suggests that the maximum permissible discount factor is 0.76. Importantly, this bound does not depend at all on the specific income process except through the impact of the income process on total asset accumulation. One parameter which the bound is quite sensitive to is σ, the curvature parameter. At $\sigma = 1$ the bound is somewhat looser: raising the upper bound on the three month MPC-SB above 0.21 requires a discount factor just below 0.85. On the other hand, using a much larger value makes the bound very tight. If we were to calibrate based on Best et al. (2020), which estimates an elasticity of intertemporal substitution as low as 0.1 (implying $\sigma = 10$), then even at a discount factor of 0.5 the 3 month MPC-SB cannot be higher than 0.09.

Theorem 1 shows that the MPC-SB must be small at the kind of discount factors normally used to calibrate models. It shows that for any buffer-stock model to be able to hit a three month MPC-SB of 0.21 would require an annual discount factor somewhat below 0.8. This applies for any income process. Are discount factors below 0.8 necessarily problematic? A household with a discount factor of 0.8 discounts consumption in ten years time by 90%, and consequently simulating a buffer stock model with most standard income processes at a discount factor below 0.8 would predict very little wealth accumulation. The obvious way to counteract this would be with a very risky income process, creating a strong precautionary motive. But as we will show below, increasing income risk pushes the MPC-SB further below the theoretical bound for any given value of β, meaning this is not a viable way of matching realistic levels of wealth accumulation. Making households more risk averse, which would create a stronger precautionary motive, is also not viable: with CRRA preferences higher risk aversion means a higher value of σ. As discussed above, higher σ makes the theoretical bound much tighter. Therefore Theorem 1 helps to establish that it is a general property of buffer-stock models that they will not be able to simultaneously match substantial MPC-SB values and realistic levels of wealth accumulation.

In Appendix A.10, we show that allowing consumption to grow in both groups at the same rate over time (equal, moderate growth rates is what we observe in the data) relaxes the bound somewhat. Adding 2% annual consumption growth raises the maximum discount factor which allows for $MPC - SB(3m) = 0.21$ from 0.76 to 0.79. What if preferences also follow a seasonal cycle in the model, as they do in practice? As long as the preference cycle is the same across groups, the core idea of our result remains: seasonally-adjusted consumption cannot fall quickly given the Euler restriction, and so MPC-SBs should not be large. However, it becomes harder to put a simple and precise restriction on the MPC-SB. In Appendix A.10 we show one way to modify Theorem 1 for the case of seasonal preferences, implying a moderate loosening of the bound. Furthermore, in our simulations below we extensively test the robustness of our results to adding a seasonal preference cycle calibrated to our data, and we find very little impact on the MPC-SB.

Given that we must ultimately rely on simulations to establish the applicability of our results across a range of model variations, what is the utility of the results in this section? Firstly, they highlight the core mechanism for why buffer stock models predict persistent consumption responses and low

seasonal bonus MPCs: consumption smoothing **after** getting a windfall. Secondly, they show that these features are general, and will apply across all income processes. Finally, they signal clearly which elements of the model will need to be relaxed to match the data, a point we will return to in Section 7.

6.3 Buffer stock simulations

We now simulate the model over a range of values of β running from 0.1 up to just below 0.975 (the buffer-stock condition requires $\beta \leq 1/(1+r) = 0.9804$ - for higher values of β asset accumulation explodes). We set $\sigma = 1.5$ and $1 + r = 1.0017$, corresponding to a 2% annual rate. Following Auclert, Rognlie, and Straub (2024), we set the persistence of log income to be 0.91 at the annual level, implying a monthly value of $\rho = 0.99$. We then match the standard deviation of log income in our income process to an estimate from the data. Using a total income measure which sums salary and unemployment benefits, we calculate the monthly cross-sectional standard deviation of log income and average across months. We calculate the standard deviation both unconditionally (0.77), and after residualizing for household fixed effects (0.42), the latter removing all unobserved, permanent household heterogeneity in earnings. This allows us to consider income risk calibrations at the upper and lower end of what the data can reasonably support, which means we can see how the degree of income risk impacts our simulations.

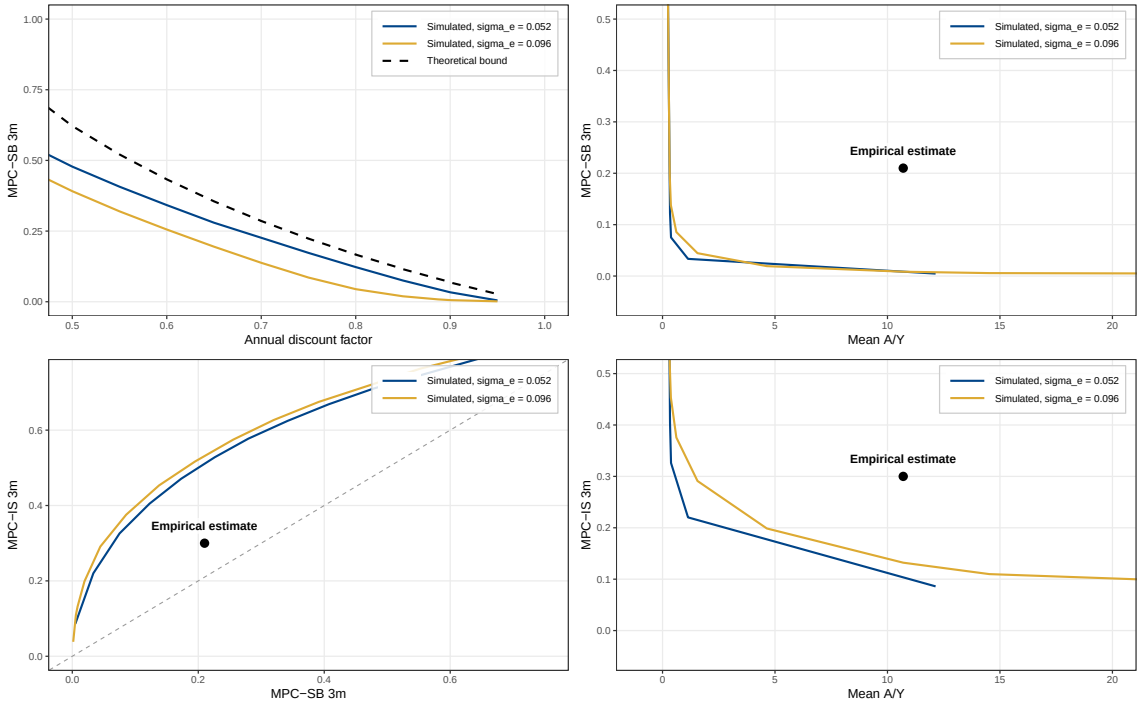


Figure 14: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with high and low income variance.

Figure 14 shows the results from our simulations. Several things are apparent. Looking at the top left panel, the simulated values of the MPC-SB lie comfortably below our theoretical bound. The distance below the theoretical bound is increasing with income risk: this is expected, because as we mentioned earlier, the distance between the minimum rate at which consumption is permitted to fall and the actual rate at which it falls is increasing in the strength of the precautionary motive, through the Jensen's inequality step. This is important: our theoretical bound already suggests that high MPC-SBs are impossible except at low discount factors, the simulations suggest that with income risk the trade-off is significantly more severe.

Turning to the top right panel, we see there is a sharp trade-off between the level of assets the model is able to hit and the MPC-SB. The MPC-SB is very small except at extremely low levels of average asset holdings. We also plot on the graph our empirical estimate of the MPC-SB, and a measure of wealth from our dataset: the ratio of mean current account deposits to mean income for the seasonal pay group. Table A.3 shows that this is very close to the average ratio of deposits to monthly income in official data, and somewhat lower than the average ratio of liquid wealth to monthly income.

Therefore the measure we plot is a lower bound on the amount of average wealth distribution the model should generate: the ratio of total net wealth to monthly income in the data is 125, more than ten times higher. Despite this, our empirical estimate lies very far outside of the frontier of possible values in the simulations. This panel therefore establishes clearly the empirical failure of the buffer stock model in this context: if we choose β to match A/Y in the model to the current account balance only (a so called 'cash only' calibration), the MPC-SB is an order of magnitude too small.[10]

In the bottom right panel, we show the same figure for the MPC-IS. Here the same trade-off exists: it is actually quite difficult to hit a reasonable level of asset accumulation and high MPCs simultaneously - this is one of the core messages of Kaplan and Violante (2022). However, the MPC-asset frontier lies much further out for the MPC-IS: the model can generate reasonably high MPC-IS values at moderate wealth levels. Moreover, the frontier of possible combinations is much more sensitive to income risk for the MPC-IS. This is why, using a somewhat richer income process, Kaplan and Violante (2022) are able to match empirical MPC(-IS) estimates using a cash only calibration. Achieving the same with the MPC-SB is impossible: not only is the trade-off between wealth accumulation and the MPC-SB much more severe, but it is also insensitive to income risk.

Finally, the bottom left panel of Figure 14 shows the relationship between the MPC-SB and the MPC-IS across different simulations. The line lies well above the 45 degree line: the MPC-SB is always much lower than the MPC-IS for a given discount factor. The empirical value, by contrast, is much closer to the 45 degree line - the model predicts a much greater discrepancy between the two than the data supports. Indeed, we argued in the previous section that a large part of the discrepancy between the two MPCs in the data is due to the different populations rather than the same households responding differently to different kinds of shock. For households getting both seasonal bonuses and lottery wins, the MPCs were identical.

The model MPC-IS is much larger than the MPC-SB for three reasons: news, planning, and precautionary savings. Firstly, an income shock carries news, a seasonal bonus does not. The MPC-IS represents the consumption response to a genuine windfall which increases lifetime resources, whilst

[10]This result is not dependent on using the mean as a summary of the wealth distribution, which is heavily skewed. In Appendix A.11 we show the same analysis comparing median wealth in the model to median deposits in the data. The take-away is exactly the same: the empirical point lies way outside the frontier of what any of the models can achieve. If you match median wealth, the MPC-SB is close to zero. If you match the MPC-SB, median wealth is close to zero.

the MPC-SB is a response to a fully anticipated income change with no implications for lifetime resources. Even in a certainty-equivalence model with no borrowing constraints, households should increase consumption somewhat after getting a windfall (consuming the perpetuity value each period), but not so for a fully anticipated fluctuation. Secondly, because of the bonus cycle is regular and anticipated, households can plan around it, saving most of it and running it down evenly until the next bonus date.

The final reason is that the precautionary saving motive increases MPCs out of income shocks for all households, whilst it does not play a direct role in the MPC-SB. Because of prudence ($u'''(c) > 0$), buffer stock households have a precautionary saving motive. In essence, their absolute risk aversion is higher at lower levels of consumption, meaning they suffer more from income risk when their consumption is low. This means they can reduce the damage of uncertain income by building up a precautionary savings buffer, and the intensity of this motive to save increases as their assets approach zero. For a low asset household, an income windfall pushes them to a point on their consumption function where the precautionary motive is reduced: they therefore spend down not just the annuity value of the windfall, but also some of precautionary buffer which they no longer need. This effect fades out gradually as assets increase, but it is a force which increases the MPC-IS across the whole range of asset holding. However, this effect does not exist with respect to seasonal bonuses. Receiving a seasonal bonus - because fully anticipated - does not change the households calculus regarding its precautionary buffer. In other words, the seasonal bonus induces no revision of the desired buffer, so no precautionary assets are liberated for spending by the income receipt. The only impact of precautionary saving on the MPC-SB is through its impact on the overall level of assets held: higher precautionary buffers lower the probability of any household hitting its borrowing constraint. This is why higher income risk and higher risk aversion both reduce the MPC-SB.

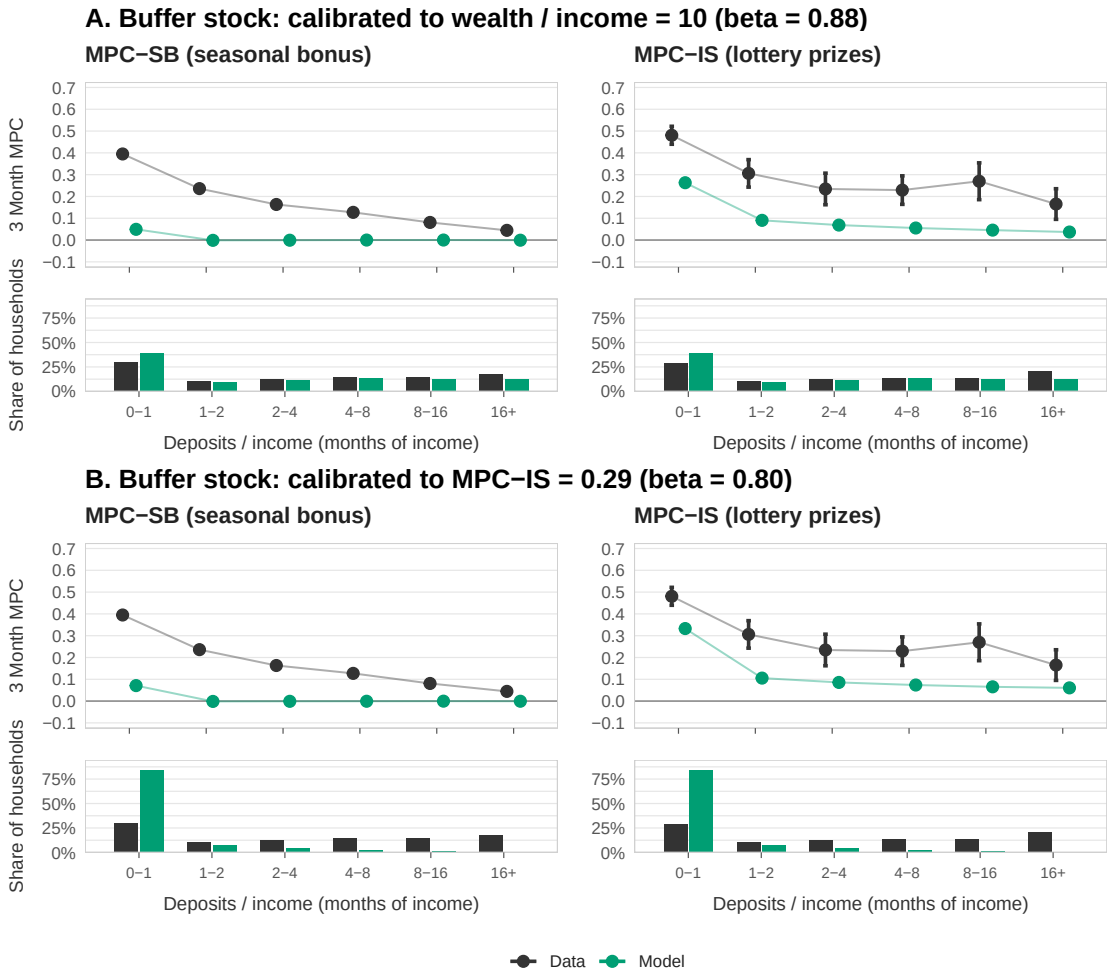


Figure 15: Distribution of MPCs and asset distribution in buffer stock model and the data.

Figure 15 demonstrates this point. It shows the MPC-SB (left) and the MPC-IS (right) for two calibrations of the buffer stock model. Each panel plots the MPCs across the distribution of assets/deposits relative to income in the model/data, and below the liquidity distribution of households in the model and data. The first row (A) sets $\beta = 0.88$ in order to match mean wealth/income of 10, the empirical (deposits) value. The second row (B) sets $\beta = 0.80$ in order to match an MPC-IS of 0.29, the empirical value. In both rows, the MPC-IS declines across the asset distribution, remaining positive up to the top liquidity bucket. The MPC-SB, on the other hand, is a step function: it is positive only in the bottom liquidity bucket. 100% of MPC-SB is generated from households with assets/income less than 1. This is in sharp contrast to our empirical results, where the MPC-SB

declines more smoothly across the liquid wealth distribution, remaining substantial over much of the distribution.

In Appendix A.11, we show how our results change if we adopt an income process with fat-tails and substantial negative skew based on Guvenen et al. (2021). The simulations produce much more wealth accumulation, and consequently much lower MPC-SBs, at each given β. They also raise the MPC-IS very significantly at any given level of assets. The MPC-SB remains very small except at very low levels of wealth, and the ratio of the MPC-SB to MPC-IS shifts much further away from its empirical value. Therefore, if anything adopting a more realistic income process pushes the model predictions even further from our findings.

6.3.1 Seasonal preferences

One potential limitation of the preceding analysis is that whilst it introduces seasonal pay, it does not reflect the fact that there is also a seasonal dimension to preferences. Intuitively, consumption should vary over the months of the year due to holidays, social customs, weather, and other factors. Indeed, in our data, spending is highly seasonal even for households whose pay is not seasonal, and their spending peaks at the same times of the year that the bonuses are paid. This is of course not a coincidence, as we have discussed the institution of seasonal bonuses evolved to provide extra liquidity at times of the year when consumption demands are particularly high. This seasonality in consumption preferences could interact with seasonality in pay and borrowing constraints in a way which produces higher MPCs out of seasonal bonuses.

This section therefore shows that adding seasonal preferences does not fundamentally alter the picture: the buffer-stock model still predicts very small MPC-SBs at reasonable wealth levels. We first introduce deterministic seasonal preference shifters whereby households derive higher marginal utility in certain months, driving them to vary their consumption over the months of the year like we observe in the data. We then allow these shifters to be stochastic, with marginal utility varying over a seasonal cycle subject to random shocks. Neither of these modifications substantially alter our conclusions in the previous section.

Firstly, we introduce a seasonal preference shifter: the variable γ_j evolves deterministically according to a 12 month cycle, making consumption more desirable in certain months than others. This variable

is calibrated so that the seasonal pattern of the non-bonus group in the model replicates the seasonal pattern of the non-bonus group in our data.[11]

The budget constraints are unchanged, and the period utility function becomes:

$$u(c, \gamma_j) = \gamma_j \cdot \frac{c^{1-\sigma}}{1-\sigma}$$

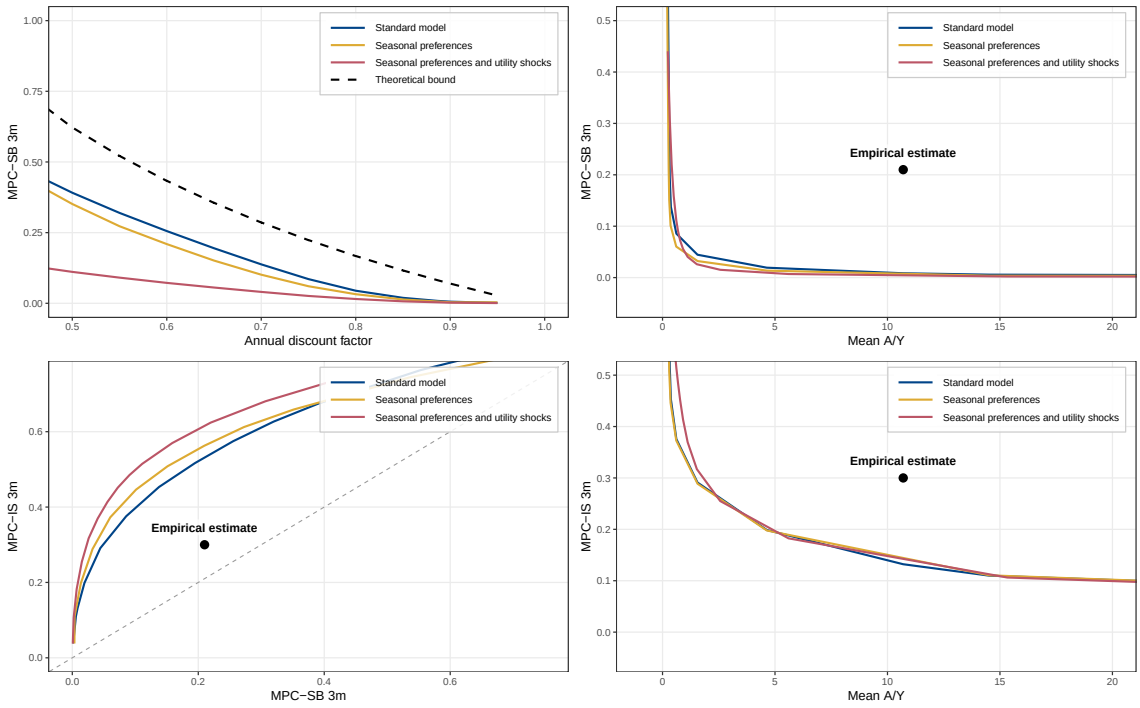


Figure 16: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with high and low income variance.

The results appear Figure 16. Overall, adding a deterministic seasonal preference cycle does not change our core conclusions: the model's MPC-SB remains very small at any reasonable level of assets. The reason is that households anticipate their seasonal preference cycle and they save appropriately - the average level of assets varies over the months of the year to facilitate spending more in months with higher γ_j. One consequence is that the MPC-SB becomes higher in December

[11] Consumption of the non-bonus group varies from 94% of the January level in February to 116% in July and 114% in December.

and lower in June.[12] The reason is that in May through to July there is a run of high γ_j months, and households build up higher savings in advance of this because being constrained is more costly at this time of year. Carrying higher buffers in the summer means that fewer households are close to their borrowing constraint when the June bonus arrives. In contrast, December has a high γ_j, but November and January do not, so households on average carry somewhat lower savings at this time of year, making it more likely that some of them are constrained when they received the December bonus. This prediction is counter-factual: in our empirical analysis, the MPC-SB is higher in June.

One alternative way of modelling seasonal preferences is to assume that there is some subsistence level of consumption which varies over the months of the year - a seasonal variation on Stone-Geary preferences. This version of the model would have the utility function:

$$u(c, \gamma_j) = \frac{(c - \gamma_j \, \bar{c})^{1-\sigma}}{1-\sigma}$$

In Appendix A.11, we present some results from this model. To summarize, it is very difficult for such a model to match the observed seasonality of the data without creating a non-zero chance that households with low assets and low income realizations fall below the subsistence level. To deal with this, we have to attenuate the magnitude of the seasonal cycle in γ_j somewhat. This model produces somewhat higher MPC-SBs, but the magnitude is still very far off what would be needed to match our data.

A further modelling device we consider is to allow the seasonal preference shifts to have a stochastic element, so that marginal utility follows a deterministic seasonal cycle over the months of the year, but is also subject to random shocks. There are two ways to interpret this. One interpretation is that seasonal preference variation has both a deterministic component based on things like annual holidays, but also a stochastic component, driven by for instance the weather and year to year variation in spending opportunities. A different, but related, interpretation, is that the stochastic aspect of the utility cycle represents 'expense shocks': unanticipated variation in consumption demands due to medical issues, or the need to make repairs to your home or car.

Our approach here is related to an emerging literature attempting to carefully quantify and model

[12]The MPC-SB we report is the average across the two.

expense shocks: for example Fulford and Low (2024) and Miranda-Pinto et al. (2025). For the sake of parsimony, however, we keep our model simple: marginal utility is subject to exogeneous, *iid* lognormal shocks. Their variance σ_r^2 is chosen to match the ratio of the variance of consumption to the variance of income in our data. [13]

The period utility function becomes:

$$u(c, \gamma_j, r) = r \cdot \gamma_j \cdot \frac{c^{1-\sigma}}{1-\sigma}$$

Where $\ln r \sim iid\ \mathcal{N}(0, \sigma_r^2)$.

The results are also displayed in Figure 16. Adding expense shocks significantly lowers the MPC-SB in this model for a given value of β. This is because the expense shocks add an additional precautionary motive to the model, inducing households to hold somewhat higher levels of assets. This means there are fewer households close to the borrowing constraint, and as we have emphasized - since the MPC-SB is purely a borrowing-constraint phenomenon - it is only such households which have a substantial MPC-SB.

6.4 Two asset models

In their review of MPCs in heterogeneous agent models, Kaplan and Violante (2022) note that it is hard for heterogeneous agent models to simultaneously match empirical MPC estimates and realistic level of wealth: either MPCs, or wealth accumulation, is too small. We have encountered a more extreme version of this same trade-off, because the buffer stock model predicts even lower MPCs out of seasonal bonuses for any level of wealth. Kaplan and Violante recommend two-asset models as the most promising way of resolving this trade-off. Two-asset models can generate the coexistence of high MPCs and much higher levels of total wealth, mostly held in illiquid forms. The mechanism is that wealthy households choose to keep most of their wealth in illiquid forms due to their higher returns, consequently holding relatively small liquid buffers, which in turn can generate higher MPCs.

In this section we therefore consider a two-asset version of the buffer stock model. We follow very closely the approach of Kaplan and Violante (2022). The only significant departures are: (i) we build

[13] In our data, the ratio Var(C)/Var(Y) is close to 1. Models without expense shocks all produce ratios significantly lower than this.

the model in discrete time, (ii) we use the same income process as in the previous section (the AR(1)), (iii) we add in the seasonal pay structure for one group of households. The only reason for these three departures is to keep the model as comparable as possible with the previous section.

Households hold two assets. Liquid wealth b_t earns a low return r_b, and illiquid wealth k_t earns a higher return $r_k > r_b$. Moving funds between the two accounts is costly. Each month a household chooses consumption c_t, next month's liquid wealth b_{t+1}, and a transfer d_t into the illiquid account ($d_t < 0$ is a withdrawal). Consumption can only be financed using liquid wealth. The problem is:

$$\max_{\{c_t, b_{t+1}, d_t\}} E_0 \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$$

subject to

$$c_t + b_{t+1} + d_t + \Phi(d_t, k_t) = R_b b_t + \kappa_{j,g} z_t$$

$$k_{t+1} = R_k k_t + d_t$$

$$b_{t+1} \geq 0, \quad k_{t+1} \geq 0, \quad d_t = 0 \text{ if } \xi_t = 0$$

where $R_b = 1 + r_b$ and $R_k = 1 + r_k$. The illiquid return accrues inside the illiquid account, so it can only be consumed by paying to withdraw it. The income process, the seasonal-pay multiplier $\kappa_{j,g}$ and its normalization w^* are exactly as in the baseline model.

Two frictions create an incentive to hold any balance in the lower-return liquid account. First, households can only adjust the illiquid amount when a rebalancing opportunity arrives ($\xi_t = 1$), which happens independently each month with probability p. Second, an adjustment carries a fixed cost:

$$\Phi(d, k) = \phi \cdot \mathbb{1}\{d \neq 0\}.$$

The MPC-SB and MPC-IS are defined exactly as in the baseline model. The windfall underlying the MPC-IS, $z \cdot w^*$, is paid into the liquid account.

The core mechanism in this model, first introduced by Kaplan and Violante (2014), is that the return gap $r_k - r_b$ leads households to hold most of their wealth in the illiquid account and only a small liquid buffer. This produces wealthy hand-to-mouth households, who hold substantial illiquid wealth

but little liquid wealth, alongside the poor hand-to-mouth of the one-asset model, who hold little of either. Following Kaplan and Violante (2014), a household is defined as hand-to-mouth if its liquid wealth is below half a month of its own income; it is poor hand-to-mouth if it also has no illiquid wealth, and wealthy hand-to-mouth otherwise. Because the wealthy hand-to-mouth have high MPCs out of liquid windfalls, two-asset models can deliver higher average MPCs out of income shocks, while matching total wealth. We ask whether this mechanism also raises MPCs out of fully anticipated and recurring payments, as our data requires.

Our baseline calibration follows Kaplan and Violante (2022): the liquid return is $r_b = -2\%$ a year (roughly a zero nominal return), the fixed cost of adjustment ϕ is set as 0.26 months of mean income[14], the expected number of re-balancing opportunities per month is 1 with a Poisson arrival rate - implying a probability of 0.63 ($p = 1 - e^{-1} = 0.63$). In the baseline, we target the same (US) wealth targets: mean net worth equal to 4.1 years of mean earnings (49.2 months), in the 2019 Survey of Consumer Finances excluding the wealthiest 5% of households, and a total hand-to-mouth share of 41%.[15]

We consider a number of alternative calibrations:

- (a) Rebalancing every month ($p = 1$). This removes the timing friction, so only the fixed cost deters re-balancing. β and r_k are recalibrated to the same targets.
- (b) Spanish wealth. We recalibrate to Spanish data. We use the 2022 EFF survey, restricting to the bottom 95% of the net wealth distribution, as Kaplan and Violante do. Mean net worth is 98.6 months of mean earnings and mean current-account balances are 10.9 months (the latter very close to our ratio of deposits to monthly income in our dataset). Net wealth relative to earnings is substantially higher for Spanish than US households (98.6 against 49.2 months), mostly because of higher house prices and lower earnings. We target total wealth of 98.6 months and then also liquid wealth of 10.9 months (rather than hand-to-mouth share). For this calibration, we set the liquid return to 2%, as in our one-asset models. β and r_k are re-calibrated to hit these targets. This calibration is the

[14] Kaplan and Violante use a fixed cost of 1500 dollars. Converting this into a ratio of monthly earnings gives 0.26.

[15] Kaplan and Violante define households as hand-to-mouth if they hold less liquid assets worth less than 0.5 times their monthly income. Later in this paper, we define hand-to-mouth status slightly differently: as households who hold liquid assets worth less than **1** times their monthly income.

most directly matched to the Spanish data.

- (c) Matching the MPC-IS. As we will see, the baseline calibrations fail to generate a sufficiently large MPC-IS. To see what it required to push the model to do so, we drop the hand-to-mouth target and instead choose the liquid return so that the three-month MPC-IS reaches the data's 0.29, with net worth still at 49.2 months and the illiquid return held at its baseline value. Doing so requires a liquid return of −25% a year. A negative real return of this size makes liquid wealth expensive to hold, so almost all households keep their liquid buffer close to zero.

Table 8 shows the results from our two-asset model simulations. These models can generate a lot of wealth accumulation alongside substantial MPCs from income shocks. But they cannot match our empirical MPC estimates. The baseline calibration produces an MPC-IS of 13%, less than half of our lottery estimate.[16] The MPC out of the seasonal bonus is 2%. Allowing monthly re-balancing makes almost no difference to the results. To match the Spanish wealth data, the third column required a higher discount factor and a far smaller gap between the returns of the liquid and illiquid assets. Liquid wealth consequently rises significantly, and both MPCs become very small: 4% for income shocks, zero for seasonal bonuses.[17] The distribution of liquid wealth in this calibration is a clear miss: only 8% of households hold less than a month's worth of liquid assets, much less than in the data. The final column shows what is required to match the income shock MPC: the liquid return must be -25%. Even with this extreme calibration, the MPC from the seasonal bonus is only 13%.

Why can't two asset models match our data? Two asset models are able to generate a larger share of hand to mouth households who can contribute to the aggregate MPC. But in any version of the buffer-stock model, even constrained households smooth windfalls at least to some extent, particularly when they are large. This means that the hand to mouth households have relatively moderate MPCs.

[16] At first glance, the MPC-IS of 13% in the baseline (KV) calibration might appear too low, but it is consistent with the results in Kaplan and Violante (2022), who find a quarterly MPC of 16% out of an income shock of $500. The difference between our 13% and their 16% is entirely due to the size of the income shock: our MPC-IS is calculated as the MPC out of a windfall equal to about a month's worth of income, matching the size of the seasonal bonus (and the size of the lottery wins we study empirically). In Table 7 of Kaplan and Violante (2022), the MPC out of a $5000 windfall is 12.3%. $5000 is about 0.9 months of mean earnings, so close to the size of shock we study.

[17] In a sense, these results are not surprising given that we showed in the previous section that a 'cash only' calibration of the buffer-stock could not generate high MPCs.

Table 8: Two Asset Model Simulation Results

Statistic	Data	KV baseline	Monthly rebalancing	Spanish wealth	MPC-IS target
Calibration					
β (annual)		0.894	0.897	0.948	0.892
r liquid		-2.0%	-2.0%	2.0%	-25.0%
r illiquid		8.2%	7.7%	3.4%	8.2%
Rebalancing probability		0.63	1.00	0.63	0.63
Fixed adjustment cost		0.26	0.26	0.26	0.26
Wealth					
Illiquid wealth / income		47.6	48.1	91.7	48.6
Liquid wealth / income	9.8	3.1	2.5	8.7	1.2
MPCs					
MPC-SB (3 months)	0.21	0.02	0.02	0.00	0.13
MPC-IS (3 months)	0.29	0.13	0.12	0.04	0.29
MPC-IS, 1m / 3m		0.38	0.37	0.36	0.41
Hand-to-mouth					
Share of MPC-SB from HtM	55%	97%	97%	100%	98%
Share of MPC-IS from HtM	46%	72%	70%	36%	85%
Households with < 1 month	30% / 29%	40%	40%	8%	65%

Note:
Share of MPC from HtM means the share of the total MPC accounted for by households with liquid assets < 1 months of income. Data for households with < 1 months is for the seasonal bonus sample / lottery sample in our bank data. In the Spanish Survey of Household Finances the share of households with any earnings who hold deposits less than 1 times their monthly earnings is 29.8 percent (falling to 27.1 percent for all liquid assets). Data for wealth / income is mean deposits divided by mean salary income in our bank data, for the seasonal pay group. In the Spanish Survey of Household Finances the ratio of mean deposits to mean earnings is 10.9.

This is especially so given that two-asset models, in trying to match total wealth, must have high discount factors. For the seasonal bonuses, arriving every six months, this smoothing makes the MPCs very small. Wealthy hand to mouth households can also plan their liquidity around the known bonus schedule, reducing any deviation from full smoothing. This means that anticipation does matter in these models.

Matching our data in a buffer-stock type model therefore requires more than a high share of hand to mouth households: it also requires these households to be impatient, so they spend down a large fraction of any windfall after receiving it, and to be poor, so they cannot plan their liquidity around anticipated payments. In the next section, we study models which do exactly this: making many households poor, constrained and impatient, as well as models which generate high MPCs through an entirely different mechanism which does not rely on households being constrained.

7 Three Models Which Match Some of Our Facts

Given the inability of one and two asset buffer stock models to match our empirical facts, we turn to models with ex ante heterogeneity or behavioural features. We ask whether these models can close the gap between anticipated and unanticipated MPCs, and whether they do this in a way that is consistent with the rest of our MPC facts. It would be impossible to comprehensively test the full range of possible alternative models, so we focus on three representative models which have featured in existing papers and are each able to match a number of our MPC facts simultaneously. Each model retains the core buffer-stock structure of a forward-looking consumer with income risk and a borrowing constraint, but layers some additional feature on top of this. Each one therefore introduces one new free parameter which is hard to discipline directly with external data, so we calibrate this in each case to hit our seasonal bonus MPC and see whether the model matches our other moments.

The first option we consider is a model with heterogeneous discount factors: one patient type which accumulates wealth, and one impatient type which lives hand to mouth, and has high MPCs. Second, we consider a model with present bias, or quasi-hyperbolic discounting. Finally, we consider a model with mental accounting of past and present income.

Heterogeneous discounting and present bias are both able to match the seasonal bonus MPC as well as liquid wealth by having a large group of high-MPC, low-liquidity households. The aggregate

MPC is driven entirely by these low-liquidity households, and MPCs are near zero for everyone else. Because all of the households contributing to the MPC are constrained, anticipation does not matter much and the gap between the MPC-SB and the MPC-IS is smaller than in buffer stock models. But this mechanism for narrowing the anticipation gap has a number of counter-factual implications: too many households are hand to mouth, and MPCs are too small (they are zero) for households that are not hand to mouth. Also, neither of these models produce the front-loaded MPC shape we see in the data.

The model with mental accounting generates MPCs through an entirely different mechanism: households have a degree of dis-saving aversion, so they prefer to spend as and when income arrives. In using this different mechanism, this model brings the seasonal bonus MPC and the lottery MPC closer together than any of the other models. It also does so without producing the same counter-factual implications as the previous two models: the hand to mouth share is close to the data, and (both) MPCs remain positive across the liquid wealth distribution. The one substantive failure of the model is that it cannot replicate the steep relationship between MPCs and liquid assets in the data.

Much of the literature on behavioural macroeconomics has focused on information frictions like sticky information or rational inattention, and belief distortions like diagnostic or extrapolative expectations (Mankiw and Reis (2002), Bordalo et al. (2020), C. D. Carroll, Crawley, Slacalek, et al. (2020)). We do not pursue this direction here because inattention or limited anticipation are simply not plausible with respect to the seasonal bonuses in Spain. As we have explained, this is a long-running, institutionalized system, and many Spanish workers will have been paid this way - with bonuses arriving in the same months - for their whole careers.

We now set out the details of the three different models. All retain most of the features of the (one-asset) buffer stock model we introduced above: a borrowing constraint at zero, risky AR(1) income, CRRA utility, and a flat-pay and a seasonal-pay group. We focus on 'cash only' calibrations, aiming to match assets in the model to liquid wealth only in the data.[18]

[18] One could in principle combine any of these models with the two asset structure to try to match liquid and total wealth. This would have different implications in each of the models, which we do not explore at length here.

7.1 Heterogeneous discount factor model

In Section 5, we presented evidence that MPCs were related more to persistent differences in the long-term tendency to hold liquidity than to short-term deviations. This pattern is more consistent with a model where permanent differences in household types drives MPCs than the temporary constraints mechanism of the standard buffer stock model. Motivated by similar evidence, a number of papers have advocated for models with ex ante heterogeneity, often in discount factors (for example Gelman (2021) and Aguiar, Bils, and Boar (2025)). Heterogeneous discounting has also been used in prior work as a means of better matching the wealth distribution and MPCs (Krusell and Smith (1998), C. Carroll et al. (2017)). The basic strategy is to use a highly impatient type, which holds very little wealth, to generate MPCs, whilst having a highly patient type generate sufficient wealth to hit aggregate wealth targets.

We consider a version of the basic buffer stock model where 50% of households have a high discount factor β_H and the other 50% have a low discount factor β_L. We impose the 50/50 split because allowing both the share and the new discount factor to vary would give the model an additional free parameter relative to the other two we consider.[19] We calibrate the two discount factors so that we simultaneously match the 3m MPC-SB and a ratio of mean assets to mean income of 10 (a cash-only calibration to mean deposits / mean income in the data).

7.2 Present bias model

The second model we consider is one with present bias (Laibson (1997), O'donoghue and Rabin (1999), Angeletos et al. (2001)). Present bias (quasi-hyperbolic discounting) applies an additional factor $\beta_0 \in (0, 1]$ between the current month and the next, so the discount factor from today to tomorrow is $\beta_0\beta$, while β is still used between any two future months:

$$u(c_j) + \beta_0\beta\, u(c_{j+1}) + \beta_0\beta^2\, u(c_{j+2}) + \dots$$

The idea is that households have time-inconsistent preferences: they discount the future overall relative to today more so than they discount between successive future periods. These preferences are

[19] An alternative approach is to calibrate the two discount rates and the share to hit some statistics on the overall wealth distribution as in Auclert, Rognlie, and Straub (2025).

time-inconsistent, so there are two versions: naive, where households do not internalize the fact that their future selves will also be present biased, and sophisticated, where they do. We focus on the former.

Let V_j be the ordinary, exponential-discounts value function,

$$V_j(a, z, e) = \max_{a' \geq 0} \left\{ u(c, e, \gamma_j) + \beta \, \mathbb{E}_j \left[V_{j+1}(a', z', e') \right] \right\}.$$

Then the present-biased household's value function $\tilde{V}_j$ is

$$\tilde{V}_j(a, z, e) = \max_{a' \geq 0} \left\{ u(c, e, \gamma_j) + \beta_0 \beta \, \mathbb{E}_j \left[V_{j+1}(a', z', e') \right] \right\}.$$

Note that it is V_{j+1}, the exponential-discounting value function, which appears on the right hand side: the household evaluates its continuation value assuming it will not be present-biased. This fact is what helps the present bias model to bypass Lemma 1: the discount factor governing the trade-off between the bonus month and the month after is given by $\beta_0 \beta < \beta$, but long-term asset accumulation is still regulated by V_{j+1}, the standard exponential-discounting value function. This allows households to have higher MPCs whilst still accumulating reasonable amounts of wealth.

We calibrate β_0 and β to hit both the (3m) MPC-SB and the same cash-only wealth target of 10 months of income.

7.3 Mental accounting model

The final model we consider is a model featuring dis-saving aversion which can be viewed as a kind of mental accounting. The idea of mental accounting goes back to R. Thaler (1985), and was developed in particular in Shefrin and R. H. Thaler (1988). The core idea is that households view resources from different sources as non-fungible, and their marginal utility from consuming is dependent on how they fund this consumption. In macroeconomics, the relevant non-fungibility is typically between past, current, and future resources, with households preferring to spend out of current resources than either past resources, which they bracket as 'savings' and dislike drawing on, or future resources, which they are also reluctant to spend even if able to.

To introduce mental accounting into our buffer stock model, we adopt the approach of Graham and

McDowall (2026) and Gimeno-Ribes (2026).[20] Households categorize their resources into two, separate 'mental accounts', based on when the resources arrived:

- a **current-income account**, credited each period with labour income y_t (including seasonal bonuses and lottery winnings, excluding asset income);
- a **savings account**, holding previously accumulated income a_t.

Because of the borrowing constraint, future income is unavailable. Households then follow a funding rule: consumption draws on current income first, and on savings only once current income is exhausted:

$$c_t = c_t^y + c_t^a, \qquad c_t^y = \min\{c_t,\, y_t\}, \qquad c_t^a = \max\{c_t - y_t,\, 0\}$$

The crucial behavioural assumption is then that a marginal unit of consumption funded from savings is worth a factor $(1 - \psi)$ less than a marginal unit funded from current income. Letting $v(c)$ be the period utility function:

$$v'(c) \;=\; \begin{cases} u'(c) & c \le y \\ (1-\psi)\, u'(c) & c > y \end{cases}$$

The implied utility function is a special case of reference-dependent preferences (Kőszegi and Rabin (2006)) where the reference point is current income. A simple interpretation of this is that it reflects budgeting behaviour, where households aim to keep their spending within their income month by month. They can violate this rule if they have savings, but they enjoy consumption less when they do so. A related interpretation is that this is an endowment effect, where once savings have been sitting in their accounts for a while, households are less willing to spend them down than when they first arrive as income.

$$v(c) = \begin{cases} u(c) & \text{if } c \le y \\ u(c) - \psi\,[u(c) - u(y)] & \text{if } c > y \end{cases}$$

[20] Gimeno-Ribes shows that the two formulations are isomorphic under certain parameter restrictions.

This utility function is standard whenever consumption is below current income, but as soon as consumption exceeds income, marginal utility is scaled down by a constant factor $1 - \psi$. One way to think about this is as a penalty on any amount of dis-saving. The interpretation is that spending is valued differently depending on which mental account it is drawn from. The first y euros come out of the current-income account and are valued as usual; anything beyond that must be drawn from the asset account, and is worth less at the margin. Setting $\psi = 0$ recovers standard preferences.

We calibrate ψ and β, as with the other models, to match both the (3m) MPC-SB and an asset to income ratio of 10.

7.4 Results

Table 9 shows our results from simulating the three different models, alongside our empirical estimates and the buffer stock calibration for reference. All three of the new models are able to match both the wealth target and the three month seasonal bonus MPC. The way they achieve this, however, is very different. Heterogeneous discounting and present bias hit the MPC-SB by making a large mass of households both hand to mouth (i.e. holding very low liquidity) and impatient (therefore having large MPCs). The table show the percentage of the MPC generated from hand to mouth households, defined as those holding assets worth less than 1 month of income. Both heterogeneous discounting and present bias generate 100% of the MPC-SB from hand to mouth households, and 95% of the MPC-IS. This compares with 55% and 46% respectively in the data.

Figure 17 shows visually why this is, comparing MPCs and the liquidity distribution in each model to the data. Heterogeneous discounting and present bias share two features: (i) both MPCs are a step function - large for hand to mouth households, and very low for all others, (ii) these models both have a large mass of hand to mouth households. These hand to mouth households have relatively similar MPCs out of seasonal bonuses as out of income shocks, because they are tightly constrained. The mental accounting model is quite different: it generates high MPCs across the whole of the liquidity distribution, and it does not create an excess mass of low liquidity households. It therefore generates only 50% of the MPC-IS from hand to mouth households, much closer to the value in the data. It generates only 21% of the MPC-SB from hand to mouth households - too low relative to the data - a point we will return to.

Table 9: Comparison of various models which can match the MPC-SB.

Statistic	Data	Buffer stock	Heterogeneous beta	Present bias	Mental accounting
MPCs					
MPC-SB (3 months)	0.21	0.01	0.20	0.21	0.21
MPC-IS (3 months)	0.29	0.14	0.38	0.34	0.30
MPC-IS, 1m / 3m		0.35	0.39	0.36	0.80
Hand-to-mouth					
Share of MPC-SB from HtM	55%	100%	100%	100%	21%
Share of MPC-IS from HtM	46%	75%	95%	95%	50%
Households with < 1 month	30% / 29%	40%	61%	58%	36%
Calibration					
Wealth / income	9.8	10.0	10.0	10.0	10.1
Parameters		$\beta = 0.88$	$\beta = \{0.50, 0.91\}$, 50/50 split	$\beta = 0.92, \beta_0 = 0.80$	$\beta = 0.89, \psi = 0.25$

Note:
Share of MPC from HtM means the share of the total MPC accounted for by households with assets < 1 months of income. Data for households with < 1 months is for the seasonal bonus sample / lottery sample in our bank data. In the Spanish Survey of Household Finances the share of households with any earnings who hold deposits less than 1 times their monthly earnings is 29.8 percent (falling to 27.1 percent for all liquid assets). Data for wealth / income is mean deposits divided by mean salary income in our bank data, for the seasonal pay group. In the Spanish Survey of Household Finances the ratio of mean deposits to mean earnings is 10.9.

Overall, the mental accounting model matches a range of moments better than the other models because it neutralizes anticipation without forcing all of both MPCs to come through hand to mouth households. It captures the profile of MPCs across the liquidity distribution better than the other models: predicting positive MPCs even for high liquidity households. For the MPC-IS, the fit is very close.[21] The mental accounting model also gets closest to matching the relative magnitude of the two MPCs. Heterogeneous discounting predicts that the MPC-IS is almost double the MPC-SB, not consistent with our evidence, while present bias overshoots by about 5 percentage points. The mental accounting model is the closest model in predicting the hand to mouth share of households. It is also the only model to predict the front-loading of MPCs we saw in the data: 80% of the three month MPC comes in the first month, compared with 39% and 36% in the other two models. Finally, mental accounting also predicts that the seasonal pay group of households accumulate more wealth than the flat pay group, something none of the other models predict. This is consistent with our data, where seasonal pay households carry higher deposit balances (and the model quantitatively matches the magnitude of this difference).

The differences in Figure 17 reflect the different ways the models generate high MPCs. Heterogeneous discounting and present bias both generate MPCs through the classic buffer-stock mechanism - hand to mouth households spend money when it arrives because it frees them from constraints imposed

[21] For the MPC-SB, mental accounting makes the liquidity profile too shallow: we discuss this at length below.

by lack of liquidity and the desire to maintain precautionary buffers. Both work by amplifying this mechanism by making more households hand to mouth. The mental accounting model does still feature the buffer-stock mechanism: this is why it also generates higher MPCs at the bottom of the liquidity distribution. But it also features an entirely different mechanism: the marginal utility kink. Many households sit exactly at the kink: they would consume more than their current income if doing so was unpenalized, but given the penalty on marginal utility from dis-saving, they choose not to, and set $c = y$. For these households, an increase in their income - whether anticipated or not - shifts their reference point upwards, and they increase consumption up to the point they would have chosen before had there been no penalty. This mechanism comes purely from the arrival of cash resetting the household's monthly spending target (the reference point), and it therefore it generates similar effects whether the cash is anticipated or not. The intuition is simple: households do not like dis-saving for psychological reasons. When some extra cash arrives, they can spend more without dis-saving, so they do. Because the shift in reference point last only during the month that the income arrives, the spending response is very front-loaded.[22]

Our empirical evidence: that seasonal bonus and lottery MPCs are very similar, and that both remain substantial across the income and liquidity distributions, strongly suggests the need for some mechanism linking spending to current income, as in the mental accounting model. There are a number of similar modelling approaches which generate a similar response. For example, the 'splurge' mechanism of C. D. Carroll, Crawley, Du, et al. (2026), where households spend a fixed fraction of their labour income on arrival, is effectively a simplified reduced form for the behaviour generated by the mental accounting model. Their mechanism applies universally, whereas the mental accounting model produces 'splurging' only for households at or near to their reference point. Another example is the 'rule of thumb' consumer (Campbell and Mankiw (1989)), where some households simply spend a constant fraction of their current income. This too generates similar spending responses out of anticipated and unanticipated income, and can in principle be assumed to apply to households carrying liquidity. Both 'splurge' and 'rule of thumb' also produce front-loaded spending responses, and responses from liquid households. All of these mechanisms can coexist in a model which also

[22]This assumption about the reference point might seem arbitrary, but given that most households in our sample are paid monthly, it seems most natural to assume that budgeting is done month to month (this is what much popular finance advice recommends), so the reference point should depend on income in the current month.

has the buffer stock mechanism generating high MPCs at the bottom of the liquidity distribution. We have focussed here on mental accounting because it provides a clear psychological rationale for the behaviour we observe which is consistent with other evidence, and it produces richer predictions about the wealth distribution and the distribution of MPCs, but depending on the context all of these 'spend on arrival' modelling devices are viable ways of matching the data.

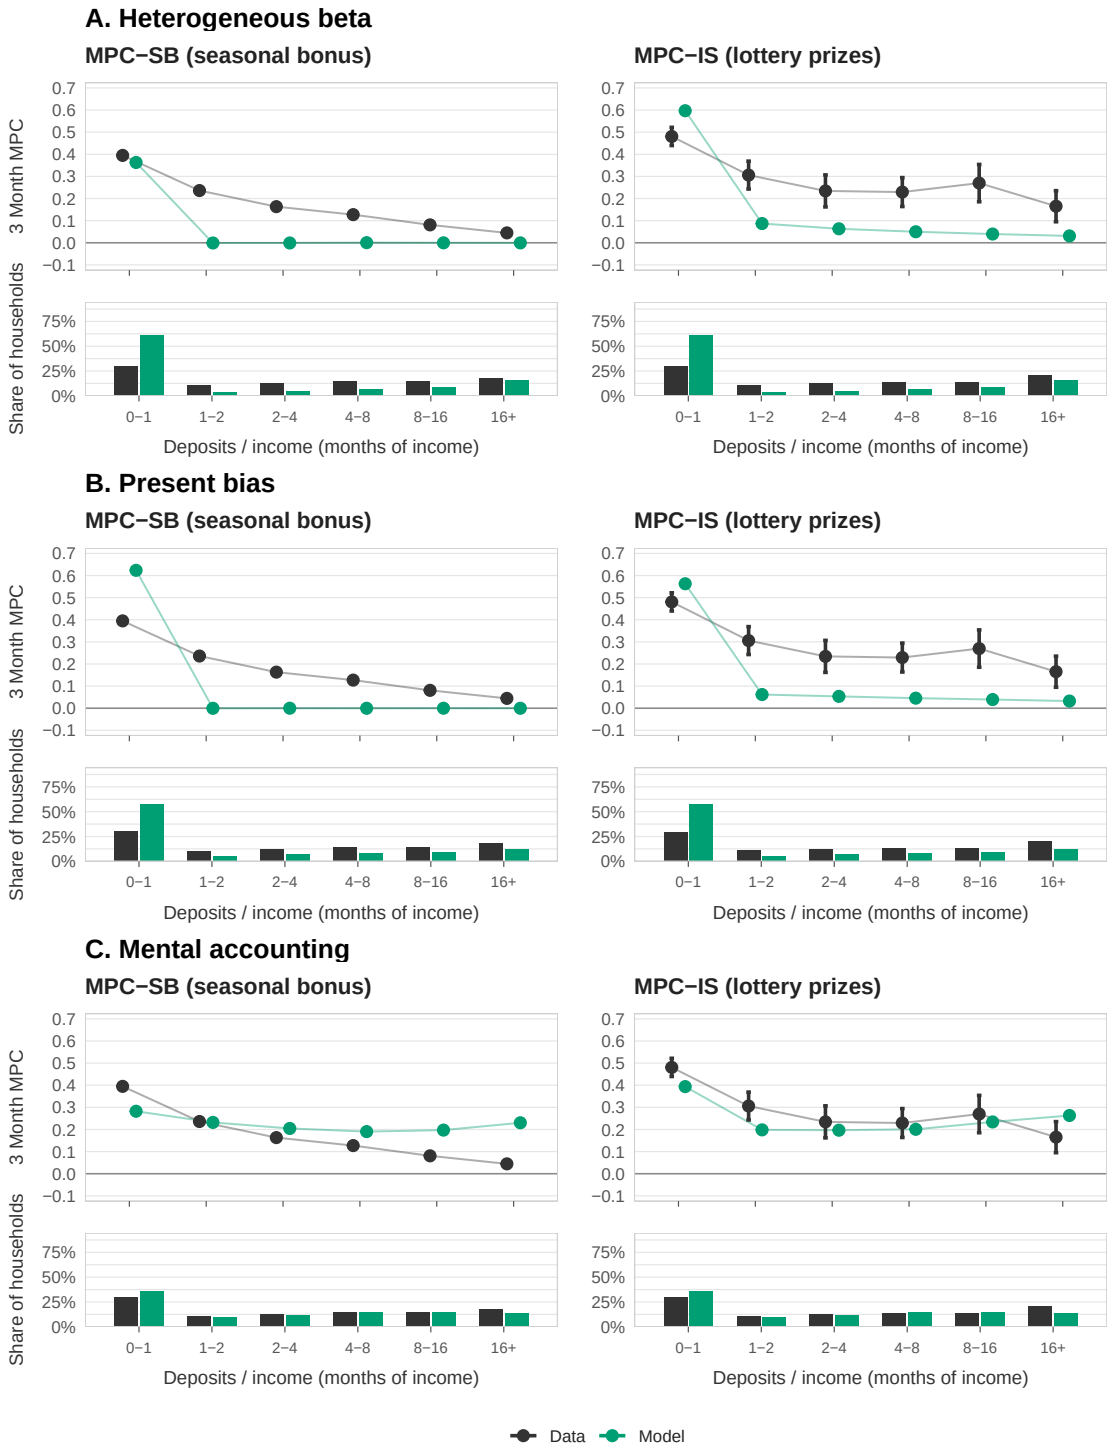


Figure 17: Distribution of MPCs and asset distribution compared across models.

The one substantial empirical failure of the mental accounting model is that it significantly under-predicts the gradient of the MPC-SB with respect to liquidity: it generates MPCs which are too small for hand to mouth households, and too high for liquid households. The other two models miss in the opposite direction: they make the liquidity gradient too steep. The reason that seasonal bonus MPCs remain large up to the top of the liquidity distribution in the mental accounting model is because a large fraction of the wealthiest households remain near their reference point: their spending is close to their income. When the seasonal bonus arrives, the reference point reset induces them to spend more whilst it is unpenalized. In our data, the most liquid seasonal bonus recipients are on average older, and spend less relative to their income than less liquid households (see Figure 10). These liquid households therefore seem most likely to be savers. If the model reproduced this fact: that high liquidity households are also savers, then it would generate much lower MPCs for them, because these households would sit below the reference point in all periods and would therefore not respond to receiving cash. But the model cannot predict that the highest liquidity households are savers: the requirement of a stationary wealth distribution means that the highest liquidity households must be dis-saving on average, so many of them end up at or close to the reference point.

There are therefore essentially three routes to restoring the relationship between the MPC-SB and liquidity. The first is any mechanism which can make a large fraction of the most liquid households also be net savers. A life-cycle with income increasing over the career and a motive to save for retirement is one way to do this: the most liquid households are those near the end of their working life, and they are also saving for retirement. The opposite route is to put a large fraction of the most liquid households beyond their reference point in all periods. A two-asset model where the wealthiest benefit from a large amount of income from the illiquid asset puts the consumption of these households well beyond their labour income. If the reference point for consumption is set by labour income, then these households are beyond the marginal utility kink in all periods, and shifts in that reference point do not shift their consumption. A third option would be to add ex ante heterogeneity into the model. Graham and McDowall (2026) do this: rather than making all households mental accounters, they give a small group a very high value of ψ, whilst the rest have $\psi = 0$. The dis-saving aversion generated by $\psi > 0$ discourages saving to some extent (accumulating assets is worth less if drawing them down is penalized). Therefore households sort, and the top of the liquidity distribution contains mostly households with $\psi = 0$. Quantitatively, this sorting force is not very strong, so it

requires a very high value of ψ for the mental accounting households. Another (perhaps a little contrived) approach would therefore be to have correlated ex ante heterogeneity: a household type which is both impatient and averse to dis-saving, and a household type which is patient and not averse to dis-saving. This combination is not unreasonable: the former group could be interpreted a less sophisticated households which plan less, but follow a simple budgeting heuristic, the latter group are more long-termist households who budget fully rationally rather than via a heuristic. Appendix A.13 presents results from some of these alternative versions of the mental accounting model.

7.5 Understanding the mental accounting model

In this section, we discuss mental accounting in more detail. We first discuss how the model relates to other empirical evidence on consumption, and how plausible its core behavioural mechanism is. Then, we explain in greater detail the core mechanism through which mental accounting generates higher MPCs across the liquidity distribution.

7.5.1 Other evidence on mental accounting

Mental accounting is a priori an attractive framework for a number of reasons. Firstly, it helps to resolve a number of consumption 'puzzles'. Empirical research has generally found that households do not spend in anticipation of receiving income windfalls, even when they are not credit constrained (S. R. Baker et al. (2025)). In a survey, Fuster, Kaplan, and Zafar (2021) find that households with high MPCs do not want to increase their spending when offered a one-year interest-free loan, and do not want to spend out of a future income gain. This can be rationalized by mental accounting, with households psychologically not wanting to spend out of 'future income'. Baugh et al. (2021) find evidence of 'asymmetric consumption smoothing': households increase spending (fail to smooth) after getting tax refunds, but they keep spending the same (they smooth) when faced with tax payments. This behaviour can be explained by mental accounting: tax refunds are treated like windfalls, and households are happy to spend freely out of them, while tax payments are mentally categorized as outside of the normal income-consumption process, so they do not impact consumption choices. Mental accounting is one explanation for why 'wealth effects' are generally found to be very small: a mental accounter might classify capital gains along with savings rather than as income. The MPC out of dividends has been found to be higher than the MPC out of capital gains (M. Baker, Nagel, and

Wurgler (2006)), suggesting non-fungibility. Gimeno-Ribes (2026) shows how a model with mental accounting can match a wide range of findings on consumption behaviour.

Secondly, a body of lab and field experiments have provided evidence that people do treat money from different sources as non-fungible (Arkes et al. (1994), Hodge and Mason (1995), Milkman and Beshears (2009)). Most of this evidence focuses on the effects of the labelling and structure of different types of current income, rather than the distinction between current, past, and future income, but it does help to establish the plausibility of the underlying behavioural mechanisms.

Finally, mental accounting is consistent with common budgeting practices and popular financial advice. Trying to keep your spending within your income month by month (rather than on average over the life-cycle) is a natural and intuitive way to budget. Much of the advice in popular finance books and on the internet recommends exactly this: to avoid spending beyond your income month by month. For example, the personal finance influencer Dave Ramsey recommends 'zero-based budgeting', that each month, all expenses and saving should be covered by income. Closely related is the notion of 'paying yourself first', common financial advice according to which you should set aside some fraction of your income each paycheck and earmark it as savings - often by literally storing it in a different bank account. Personal finance advice often instructs people to have separate emergency and long-term savings pots, and to avoid touching these unless absolutely necessary. This is essentially advising people to behave as mental accounters. In their study of budgeting behaviours, Zhang et al. (2022) find that the majority of households do budget, and that this is true across the income distribution. They find that 'the desire to avoid overspending one's income is by far the predominant reason why respondents budget'. Furthermore, they find a large asymmetry: whilst more than 85 percent of respondents would adjust their spending or budget in response to over-spending, only 30 percent would do so in response to under-spending. We interpret all of this as supporting the notion that a large proportion of households compare spending with **current** income, and suffer some psychological cost when the former exceeds the latter.

7.5.2 How mental accounting generates high MPCs

Recall that the mental accounting model puts a kink in the utility function where consumption equals current income:

$$v(c) = \begin{cases} u(c) & \text{if } c \leq y \\ u(c) - \psi\left[u(c) - u(y)\right] & \text{if } c > y \end{cases}$$

The consequence of the kink is that the Euler equation holds with equality only away from current income. Below the kink, consumption satisfies the usual condition; above it, the same condition holds with $(1-\psi)u'(c)$ in place of $u'(c)$. In between, marginal utility jumps down discretely, and there is a range of values of the marginal value of saving that no consumption choice can equate to it. This range of values is:

$$(1-\psi)\,u'(y) \;\leq\; \beta(1+r)\,\mathbb{E}\left[v'(c')\right] \;\leq\; u'(y)$$

In this range, the household would want to consume more than their income based on unpenalized marginal utility, but does not want to consume more based on penalized marginal utility. Therefore, they will consume their income exactly: they behave as a hand to mouth household and have an MPC of one out of earnings. If they receive an income windfall, the reference point moves with it: the ceiling on unpenalized spending rises euro for euro. For a small windfall the household remains at the kink and spends it in full. For a windfall as large as our seasonal bonuses, the ceiling typically rises above the household's unpenalized target, so the penalty becomes inactive and the household consumes its unpenalized optimum out of the higher cash. Either way, for such a household, the bonus month is the one month of the cycle in which spending carries no penalty, and that is what drives the results below. Figure 18 illustrates the mechanism in a 2-period model.

Let c^* denote the consumption the household would choose if the penalty were removed. The right-hand inequality holds when $c^* \geq y$; the left holds when $c^* \leq (1-\psi)^{-1/\sigma} y$. Equivalently, a household consumes exactly its current income whenever current earnings fall in the window:

$$(1-\psi)^{1/\sigma}\, c^* \;\leq\; y \;\leq\; c^*$$

The important feature of this window is that as y increases, so does target consumption, c^* (as long as income is persistent - in our calibrations and in reality income is highly persistent). This means

the hand to mouth window remains even for high income households. The likelihood of y falling in the hand to mouth region does fall with assets: households with high assets relative to current income will have a high target for c^* relative to their income, and they are therefore less likely to fall in the window. However, this mechanism can generate hand to mouth behaviour at high wealth levels. In our calibrations, the mental accounting model generates MPCs which do fall as assets increase, but slowly: they remain substantial in the top quintile of the wealth distribution.

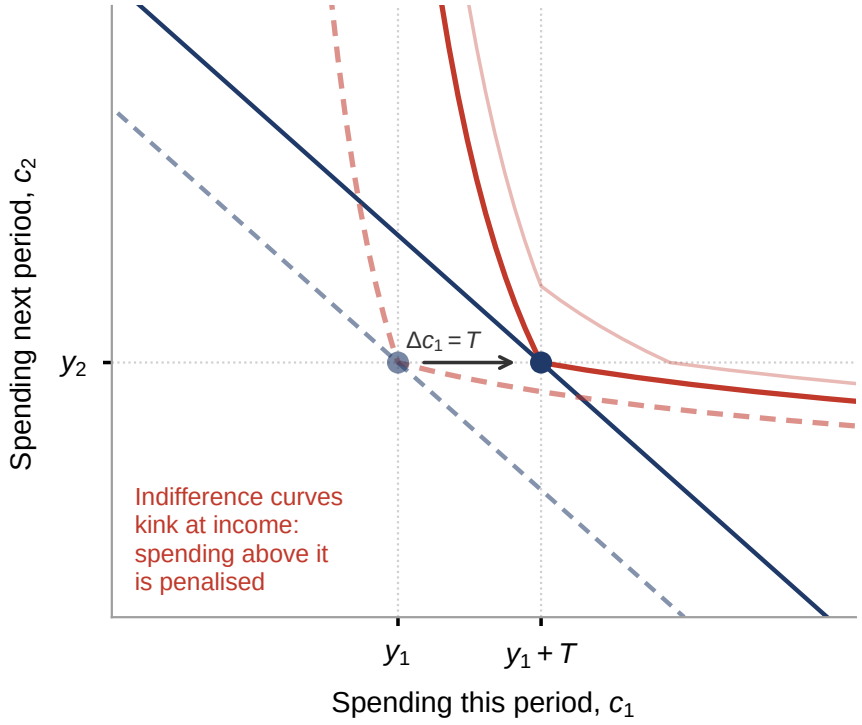


Figure 18: A two period example of how mental accounting impacts MPCs. When an income windfall is received, the household may remain at the kink, consuming all of the windfall.

The utility kink does not make hand-to-mouth status universal: being at the kink requires the current earnings draw to land inside the window. The frequency with which this occurs depends on the width of the window - which depends on ψ - and on the dispersion of income realizations. The width of the window is set by the parameters only, so it does not narrow as households grow richer — wealth shifts the window without shrinking it. What limits the hand-to-mouth share is therefore how far the window has drifted above ordinary earnings realizations, which is governed by the ratio of wealth to

income. Because income is highly persistent, high-wealth households tend also to have high income, so this ratio rises only slowly as wealth increases.

How does mental accounting relate to our theoretical results? Write $v'(c) = \mu\, u'(c)$, so $\mu = 1$ when the marginal unit of consumption is funded from current income and $\mu = 1-\psi$ when it is funded from savings. The Euler inequality becomes $\mu_t u'(c_t) \geq \beta R\, \mathbb{E}[\mu_{t+1} u'(c_{t+1})]$, so the effective one-month discount factor is

$$\tilde{\beta}_t \;=\; \beta\, \frac{\mu_{t+1}}{\mu_t}.$$

The aggregate version of Lemma 1 therefore becomes:

$$\bar{C}_{t+1} \;\geq\; (R\tilde{\beta}_t)^{1/\sigma} \bar{C}_t$$

The path of $\tilde{\beta}_t$ depends on where the household is relative to the utility function kink. Consider a household that draws on its savings in non-bonus months, so that $y < (1-\psi)^{1/\sigma} c^*$ and $\mu = 1-\psi$ outside the bonus month. Receiving the bonus temporarily shifts the reference point and therefore raises their marginal utility of an additional unit of consumption. If T is the month they receive the bonus, then $\mu_T = 1$ and $\mu_{T+1} = 1-\psi$, so that $\tilde{\beta}_T = \beta(1-\psi)$. Receiving the bonus multiplies the floor on $\bar{C}_{T+1}/\bar{C}_T$ by $(1-\psi)^{1/\sigma}$ — equivalently, it multiplies the ceiling on bonus-month consumption by $\kappa \equiv (1-\psi)^{-1/\sigma}$, which is 1.21 at $\psi = 0.25$. This directly relaxes both the restriction on the shape of the MPC time path, allowing it to be short and sharp, and the restriction on the size of the MPC-SB, allowing it to be significantly larger. Crucially, however, chaining $\tilde{\beta}_t$ over the six-month cycle starting at T gives $\prod_{s=0}^{5} \tilde{\beta}_{T+s} = \beta(1-\psi) \cdot \beta^4 \cdot \frac{\beta}{(1-\psi)} = \beta^6$, where the final factor is $\beta/(1-\psi)$ because the month before the next bonus is penalized while the bonus month itself is not, so the household is correspondingly more patient on the margin running into the bonus. The penalty for crossing into the savings account is refunded on crossing back out, and the long-run discount factor governing asset accumulation is therefore unchanged. This means that while mental accounting households occasionally discount at a much higher rate between adjacent periods, their long-run desire to accumulate assets remains strong. For any given β, mental accounting households do accumulate somewhat lower asset buffers, because they anticipate that future dis-saving will be penalized by the kink, but this is quantitatively second-order. How much mental accounting actually

relaxes the restriction of Lemma 1 depends on the distribution of households relative to the kink, and what fraction of total consumption is accounted for by these households, so we will rely on simulations to quantify it.

Figure 19 shows simulation results for mental accounting models with three values of ψ - 0.15, 0.25, and 0.35 - across the full range of possible discount factor values. The top left panel shows our theoretical bound for the buffer stock model, and it shows that the realized MPC-SB values in the mental accounting models can exceed that bound. The top right panel shows clearly the core advantage of the mental accounting model: it massively relaxes the trade-off between asset accumulation and the MPC-SB. The line for $\psi = 0.25$ almost passes through the empirical point, showing clearly that this model will be able to match both moments. Another attractive feature of the mental accounting model is apparent in the bottom left panel: whilst it increases both the MPC-SB and the MPC-IS, the effect on the former is much larger, meaning the model can reproduce the empirical relationship between the two quantities.

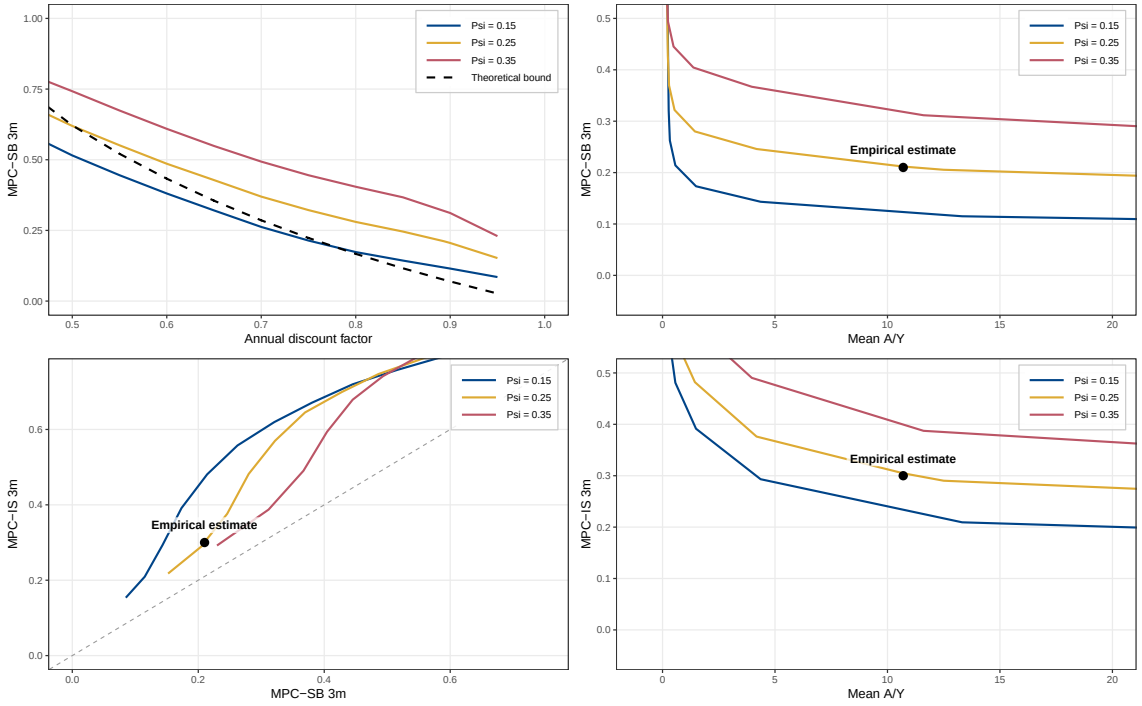


Figure 19: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with mental accounting at different values of psi.

The general idea of mental accounting is quite flexible. The way we model it here, there are two mental accounts: current income, which includes both seasonal bonuses and lottery winnings, and assets, which in this model come only from accrued past income (borrowing is not possible due to the constraint, so future income is inaccessible). The mental accounting in this framing is primarily about a division between 'current' and 'previously accrued' income into different brackets, rather than depending on the source or labelling of the income, which was more the emphasis in earlier treatments. It may well be the case that both seasonal bonuses and lottery winnings are bracketed into different mental accounts to normal salary income and to each other, and these different mental accounts may have different "ψ parameters". We do not explore this possibility here: although mental accounting might in theory lead to different MPCs out of every different kind of income windfall depending on the source and labelling, allowing for this seems unlikely to provide the basis for tractable macroeconomic models which can be applied across a range of settings. In addition, our estimated MPCs are actually quite similar for both seasonal bonuses and lottery winnings, and the mental accounting model with a single ψ value can match this fact.

8 Implications

The first main message of this paper is that MPCs out of fully anticipated and recurring payments are roughly as large as MPCs out of income shocks. The second message is that this fact should be interpreted through models linking spending to income mainly through some mechanism other than borrowing constraints if we want to match a realistic liquidity distribution and the evidence on heterogeneity of MPCs and front-loading. In this section, we discuss the implications of these two points. In Appendix A.14, we show results from simulating the HANK model of Auclert, Rognlie, and Straub (2025) with mental accounting added to the household block, providing a quantification of some of the points here in a baseline model.

One approach to rationalizing evidence of deviations from textbook consumption smoothing has been to point out that the welfare stakes are actually very small. Kueng (2018) derives a sufficient statistic for the welfare loss from failing to smooth an anticipated payment, expressed as a money-metric percentage of lifetime wealth, and shows that it is small: Loss $\approx \text{MPC}^2\ (w^*/c^{pih})^2\ (\sigma/2)\ (T-1)/T^2$, where T is the number of periods over which the payment could be spread and c^{pih} is permanent-

income consumption over one such period. The loss is second order because it expands around the PIH benchmark, where utility is already optimized, and so the first-order term vanishes by the envelope theorem. Using our three-month (quarterly) MPC-SB of 0.21 the relevant T is the number of quarters over which the bonuses can be smoothed, which is 2. Using $\sigma = 1.5$, and the ratio of the average seasonal pay to average three-month consumption in our data, which is 0.37, Kueng's expression implies an average welfare loss of about 0.11% of lifetime (human and financial) wealth — equivalently, a permanent reduction in consumption of the same 0.11%, or less than five euros per month. This is similar to the value calculated by Kueng: while our bonuses are somewhat larger relative to consumption, and the horizon for smoothing shorter (both raising the loss), our MPC estimates are smaller.

Despite the small welfare costs, this failure to smooth can have large consequences at the macroeconomic level: it is well known that MPCs are the crucial determinant of the effectiveness of fiscal stimulus. Indeed, Auclert, Rognlie, and Straub (2024) show that the matrix of intertemporal MPCs is the sufficient statistic for the fiscal multiplier under deficit finance, and it is the whole path that matters, not the impact response alone — so the short, sharp profile documented here is itself macroeconomically relevant, not just the level of the MPC. Therefore, finding alternative models which do match consumption behaviour remains important, despite the small welfare costs. In fact, the small welfare costs of deviating from the optimal consumption plan can help rationalize behavioural models: if the costs of deviation are small, then even small psychological frictions will be enough to produce meaningfully different behaviour. Here there is a direct analogy with menu costs, where small frictions at the firm level can produce large effects at the macro level. Indeed, the reasoning is very similar: a firm's profit is flat in its own price at the optimum, so arbitrarily small menu costs rationalize not adjusting prices (Akerlof and Yellen (1985), Mankiw (1985)); likewise a household's objective is flat in its own consumption path at the optimum, so arbitrarily small psychological costs can rationalize not smoothing. Important aggregate effects arise because the small individual-level deviations are correlated.[23] We therefore turn to considering the policy

[23] Kueng is very much aware of this point, p1748: "While the failure of the standard theory documented herein is not economically significant for individual households, it has potentially important implications for policy makers, since these deviations (or mistakes) are correlated across households and therefore do not disappear in the aggregate. For macroeconomic policies, such near-rational alternatives might therefore be the more relevant behavior than the one predicted by the standard consumption models."

implications of our findings.

We showed above that our first core message - that anticipation does not matter much - can be reproduced by models with heterogeneous discounting, present bias, or mental accounting. All of these models share certain policy implications. Fiscal policy is powerful, and takes effect when money hits accounts, not when it is announced. Announcing some stimulus in the future has little impact on current spending.

Our second point - that 'spend on arrival' behavioural models match the overall evidence much better than models generating MPCs only through hand to mouth households - has further, distinctive implications. Firstly, policies affecting cash-flows for liquid households do affect their consumption. For example, changes in mortgage payments affect consumption even of relatively wealthy and liquid households. Fiscal stimulus can still boost consumption even when households are already flush with liquidity (for example during the pandemic). The share of households who are hand to mouth is not the single crucial statistic for the effectiveness of stimulus policies as it is in most standard models. Taken literally, the mental accounting model we presented in the previous section also suggests that a policy redistributing from liquid to illiquid households would have much less impact than in standard models. But since this conclusion comes from the one important prediction of the model which is not consistent with our data (the flat profile of MPCs with respect to liquidity), we do not think this is a robust take-away.

A second distinctive implication is that changes to households' financial position coming via their main current income stream, to which their budgeting is anchored, have a greater impact than changes coming from other sources. So for example, changes in wealth due to fluctuations in the value of retirement or investment accounts have less impact on consumption than equivalent changes in income (this point is discussed in more detail by Gimeno-Ribes (2026)). This is consistent with a body of empirical evidence suggesting that MPCs out of wealth are lower than those out of income. Two-asset models, where drawing on wealth gains incurs costs, are one way to interpret this evidence, and this is very plausible with for example housing wealth, which is usually costly to draw on. But for wealth gains coming from stock portfolios - which are really quite liquid in many settings - mental accounting provides a complementary explanation for small wealth MPCs. A related implication is that debt relief policies will affect consumption more if they target the payment stream rather than the

principal. In standard models, providing debt relief by reducing the principal owed has nearly the same impact as reducing the payment stream by a present value equivalent amount: the impact on lifetime resources is the same. With mental accounting, reducing the payment stream has a greater effect on spending. This is consistent with the evidence of Ganong and Noel (2020), who found no impact on consumption from principal reductions, but significant effects from maturity extensions. The authors show that this is primarily due to borrowing constraints, but under mental accounting this pattern could also exist for liquid households.

A third, related, implication concerns the effectiveness of policies which extend credit or otherwise relax the borrowing constraint for constrained households. If MPCs are entirely driven by constrained households, extending credit to these households can have a powerful stimulative effect. Under mental accounting, the response is much weaker: spending is constrained by current income as well as by access to credit. This is consistent with the survey evidence of Fuster, Kaplan, and Zafar (2021) that households with high MPCs do not respond to the offer of a one year interest-free loan. Credit isn't completely ineffective under mental accounting - this would be inconsistent with empirical evidence (Aydin (2022) estimates a quarterly MPC of 0.11 out of credit expansions in an RCT) - but it is less powerful.

Our findings also have implications for intertemporal MPCs (iMPCs). Table A.12 in Appendix A.8 reports the full profile of MPCs out of lottery winnings from 4 months before to 12 months after winning for each of the five quintiles of prize size, and for prize sizes matched to the seasonal bonuses. Auclert, Rognlie, and Straub (2024) show that iMPCs are the crucial sufficient statistic for the impact of deficit-financed fiscal transfers, and we therefore hope our estimates will be useful for calibrating models to income shocks of different sizes. The iMPC profile we estimate is very front-loaded, coming mostly within the first month (with the exception of the top prize bin). This profile is inconsistent with the smooth decay generated by any buffer stock variant, but is re-produced by the mental accounting model. Since HANK models are often calibrated to the impact MPC, they will overstate the iMPCs beyond impact, therefore overstating the overall multiplier.

What about 'anticipatory' iMPCs: the response of consumption to fully anticipated income? Our core result - that anticipation does not matter for the MPC - suggests that these are small. Figure 3 showed that the spending of bonus and non-bonus groups does not diverge until the month when

bonuses arrived, consistent with evidence from other settings that households wait for income to arrive before spending it.

9 Conclusion

How much does anticipation matter for MPCs? Our evidence suggests it does not matter very much: consumption responds similarly to fully anticipated, recurring bonuses as it does to winning the lottery. Consumption responds to both types of income receipt for households across the income and liquidity distributions. And this response is front-loaded and short-lived.

We showed that standard models cannot match these facts. The Euler inequality restricts the rate at which consumption can fall in these models, implying that MPCs must be quite persistent, and out of recurring, anticipated income, they must be small. Simulations across a range of model variants, including two-asset models, showed the generality of this conclusion.

After this, we considered two ways of modifying standard models to match our findings. The first - represented by a model with heterogeneous discounting and a model with present bias - is to make a large mass of households hand to mouth, and generate the MPC from these. Because the only households with positive MPCs in these models are constrained, anticipation does not matter. The second strategy is to add some behavioural feature to the model which links desire to consume with current income, specifically a mental accounting friction where households prefer to consume from their current income than their accrued savings. Mental accounting can narrow the gap between anticipated and unanticipated MPCs without hand to mouth households. Overall, we find this second strategy to be more consistent with the evidence. The first type of models implies that too many households are hand to mouth, that MPCs are too small for liquid households, and that MPCs are not as front-loaded as in the data.

In the model with mental accounting, current income matters most for consumption. Therefore policies operating directly on current flows of income matter more than policies shifting stocks of wealth or expectations about the future.

References

Aguiar, Mark, Mark Bils, and Corina Boar (2025). "Who are the Hand-to-Mouth?" In: *Review of Economic Studies* 92.3, pp. 1293–1340.

Akerlof, George A and Janet L Yellen (1985). "Can small deviations from rationality make significant differences to economic equilibria?" In: *The American Economic Review* 75.4, pp. 708–720.

Angeletos, George-Marios et al. (2001). "The hyperbolic consumption model: Calibration, simulation, and empirical evaluation". In: *Journal of Economic perspectives* 15.3, pp. 47–68.

Arkes, Hal R et al. (1994). "The psychology of windfall gains". In: *Organizational behavior and human decision processes* 59.3, pp. 331–347.

Auclert, Adrien (2019). "Monetary policy and the redistribution channel". In: *American Economic Review* 109.6, pp. 2333–2367.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2024). "The intertemporal keynesian cross". In: *Journal of Political Economy* 132.12, pp. 4068–4121.

– (2025). "Fiscal and monetary policy with heterogeneous agents". In: *Annual Review of Economics* 17.1, pp. 539–562.

Aydin, Deniz (2022). "Consumption response to credit expansions: Evidence from experimental assignment of 45,307 credit lines". In: *American Economic Review* 112.1, pp. 1–40.

Baker, Malcolm, Stefan Nagel, and Jeffrey Wurgler (2006). *The effect of dividends on consumption*.

Baker, Scott R et al. (2025). "Anticipatory Spending". In.

Barsky, Robert B and Jeffrey A Miron (1989). "The seasonal cycle and the business cycle". In: *Journal of Political Economy* 97.3, pp. 503–534.

Baugh, Brian et al. (2021). "Asymmetric consumption smoothing". In: *American Economic Review* 111.1, pp. 192–230.

Beaulieu, J Joseph and Jeffrey A Miron (1992). “A cross country comparison of seasonal cycles and business cycles”. In: *The Economic Journal* 102.413, pp. 772–788.

Bergholt, Drago, Francesco Furlanetto, and Lorenzo Mori (2025). *Consumption inequality, household risks, and the business cycle*.

Best, Michael Carlos et al. (2020). “Estimating the elasticity of intertemporal substitution using mortgage notches”. In: *The Review of Economic Studies* 87.2, pp. 656–690.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel (2025). “Five facts about MPCs: Evidence from a randomized experiment”. In: *American Economic Review* 115.1, pp. 1–42.

Bordalo, Pedro et al. (2020). “Overreaction in macroeconomic expectations”. In: *American Economic Review* 110.9, pp. 2748–2782.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess (2024). “Revisiting event-study designs: robust and efficient estimation”. In: *Review of economic studies* 91.6, pp. 3253–3285.

Browning, Martin and M Dolores Collado (2001). “The response of expenditures to anticipated income changes: panel data estimates”. In: *American Economic Review* 91.3, pp. 681–692.

Campbell, John Y and N Gregory Mankiw (1989). “Consumption, income, and interest rates: Reinterpreting the time series evidence”. In: *NBER macroeconomics annual* 4, pp. 185–216.

Carroll, Christopher et al. (2017). “The distribution of wealth and the marginal propensity to consume”. In: *Quantitative Economics* 8.3, pp. 977–1020.

Carroll, Christopher D (1997). “Buffer-stock saving and the life cycle/permanent income hypothesis”. In: *The Quarterly journal of economics* 112.1, pp. 1–55.

Carroll, Christopher D, Edmund Crawley, William Du, et al. (2026). “Welfare and spending effects of consumption stimulus policies”. In: *Quantitative Economics* 17.3, pp. 741–790.

Carroll, Christopher D, Edmund Crawley, Jiri Slacalek, et al. (2020). “Sticky expectations and consumption dynamics”. In: *American economic journal: macroeconomics* 12.3, pp. 40–76.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva (2024). *The how and why of household reactions to income shocks*. Tech. rep. National Bureau of Economic Research.

Crawley, Edmund and Alexandros Theloudis (2024). "Income shocks and their transmission into consumption". In: *arXiv preprint arXiv:2404.12214*.

De Nardi, Mariacristina, Eric French, and John B Jones (2010). "Why do the elderly save? The role of medical expenses". In: *Journal of political economy* 118.1, pp. 39–75.

Deaton, Angus (1991). "Saving and Liquidity Constraints". In: *Econometrica: Journal of the Econometric Society*, pp. 1221–1248.

Debortoli, Davide and Jordi Galí (2024). "Idiosyncratic income risk and aggregate fluctuations". In: *American Economic Journal: Macroeconomics* 16.4, pp. 279–310.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996). "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach". In: *Econometrica* 64.5, pp. 1001–1044. DOI: 10.2307/2171954.

European Central Bank (Dec. 2024). *Study on the Payment Attitudes of Consumers in the Euro Area (SPACE) – 2024*. Tech. rep. Frankfurt am Main: European Central Bank. DOI: 10.2866/9860161. URL: https://www.ecb.europa.eu/stats/ecb_surveys/space/shared/pdf/ecb.space2024~19d46f0f17.en.pdf (visited on 09/04/2026).

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik (2021). "MPC heterogeneity and household balance sheets". In: *American Economic Journal: Macroeconomics* 13.4, pp. 1–54.

Friedman, Milton (1957). "The permanent income hypothesis". In: *A theory of the consumption function*. Princeton University Press, pp. 20–37.

Fulford, Scott and David Low (2024). "Expense shocks matter". In: *Consumer Financial Protection Bureau Office of Research Working Paper* 24-08.

Fuster, Andreas, Greg Kaplan, and Basit Zafar (2021). "What would you do with 500 dollar? Spending responses to gains, losses, news, and loans". In: *The Review of Economic Studies* 88.4, pp. 1760–1795.

Ganong, Peter and Pascal Noel (2019). “Consumer spending during unemployment: Positive and normative implications”. In: *American economic review* 109.7, pp. 2383–2424.

– (2020). “Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession”. In: *American Economic Review* 110.10, pp. 3100–3138.

Gelman, Michael (2021). “What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics”. In: *Journal of Monetary Economics* 117, pp. 521–542.

Gerard, François and Joana Naritomi (2021). “Job displacement insurance and (the lack of) consumption-smoothing”. In: *American Economic Review* 111.3, pp. 899–942.

Gimeno-Ribes, Jaime (2026). *A Model of Consumption with Mental Accounting and Heterogeneous Agents*. Tech. rep.

Gómez Yáñez, José Antonio (2025). *Anuario del Juego en España 2025: El juego en España y su aportación social*. Spanish. Ed. by María Gómez González. Madrid: CEJUEGO (Consejo Empresarial del Juego). ISBN: 978-84-09-76892-9.

Graham, James and Robert A McDowall (2026). “Mental accounts and consumption sensitivity across the distribution of liquid assets”. In: *American Economic Journal: Macroeconomics* 18.1, pp. 1–33.

Guvenen, Fatih et al. (2021). “What do data on millions of US workers reveal about lifecycle earnings dynamics?” In: *Econometrica* 89.5, pp. 2303–2339.

Hall, Robert E (1978). “Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence”. In: *Journal of political economy* 86.6, pp. 971–987.

Hodge, Sharon K and Charlotte H Mason (1995). “Work versus windfall: An exploration of saving on subsequent purchase”. In: *Marketing Letters* 6.2, pp. 91–100.

Hori, Masahiro and Satoshi Shimizutani (2012). “Do households smooth expenditure over anticipated income changes? Evidence from bonus payments to public employees in Japan”. In: *Journal of the Japanese and International Economies* 26.3, pp. 405–433.

Hsieh, Chang-Tai (2003). “Do consumers react to anticipated income changes? Evidence from the Alaska permanent fund”. In: *American Economic Review* 93.1, pp. 397–405.

Jappelli, Tullio and Luigi Pistaferri (2010). “The consumption response to income changes”. In: *Annu. Rev. Econ.* 2.1, pp. 479–506.

– (2017). *The economics of consumption: theory and evidence*. Oxford University Press.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles (2006). “Household expenditure and the income tax rebates of 2001”. In: *American economic review* 96.5, pp. 1589–1610.

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante (2018). “Monetary policy according to HANK”. In: *American Economic Review* 108.3, pp. 697–743.

Kaplan, Greg and Giovanni L Violante (2014). “A model of the consumption response to fiscal stimulus payments”. In: *Econometrica* 82.4, pp. 1199–1239.

– (2022). “The marginal propensity to consume in heterogeneous agent models”. In: *Annual Review of Economics* 14.1, pp. 747–775.

Kőszegi, Botond and Matthew Rabin (2006). “A model of reference-dependent preferences”. In: *The Quarterly Journal of Economics* 121.4, pp. 1133–1165.

Krusell, Per and Anthony A Smith Jr (1998). “Income and wealth heterogeneity in the macroeconomy”. In: *Journal of political Economy* 106.5, pp. 867–896.

Kueng, Lorenz (2018). “Excess sensitivity of high-income consumers”. In: *The Quarterly Journal of Economics* 133.4, pp. 1693–1751.

Laibson, David (1997). “Golden eggs and hyperbolic discounting”. In: *The Quarterly Journal of Economics* 112.2, pp. 443–478.

Mankiw, N Gregory (1985). “Small menu costs and large business cycles: A macroeconomic model of monopoly”. In: *The Quarterly Journal of Economics* 100.2, pp. 529–538.

Mankiw, N Gregory and Ricardo Reis (2002). “Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve”. In: *The Quarterly Journal of Economics* 117.4, pp. 1295–1328.

Mastrobuoni, Giovanni and Matthew Weinberg (2009). “Heterogeneity in intra-monthly consumption patterns, self-control, and savings at retirement”. In: *American Economic Journal: Economic Policy* 1.2, pp. 163–189.

Milkman, Katherine L and John Beshears (2009). "Mental accounting and small windfalls: Evidence from an online grocer". In: *Journal of Economic Behavior & Organization* 71.2, pp. 384–394.

Miranda-Pinto, Jorge et al. (2025). "A model of expenditure shocks". In: *Journal of Monetary Economics*, p. 103807.

O'donoghue, Ted and Matthew Rabin (1999). "Doing it now or later". In: *American economic review* 89.1, pp. 103–124.

Olafsson, Arna and Michaela Pagel (2018). "The liquid hand-to-mouth: Evidence from personal finance management software". In: *The Review of Financial Studies* 31.11, pp. 4398–4446.

Parker, Jonathan A (1999). "The reaction of household consumption to predictable changes in social security taxes". In: *American Economic Review* 89.4, pp. 959–973.

– (2017). "Why don't households smooth consumption? Evidence from a $25 million experiment". In: *American Economic Journal: Macroeconomics* 9.4, pp. 153–183.

Parker, Jonathan A et al. (2013). "Consumer spending and the economic stimulus payments of 2008". In: *American economic review* 103.6, pp. 2530–2553.

Rodríguez Molina, Isabel and Álvaro Esandi Erro (2025). "Trends and developments in payments by Spanish consumers". In: *Banco de Espana Article* 1.2025, Q3.

Sant'Anna, Pedro HC and Jun Zhao (2020). "Doubly robust difference-in-differences estimators". In: *Journal of econometrics* 219.1, pp. 101–122.

Shapiro, Jesse M (2005). "Is there a daily discount rate? Evidence from the food stamp nutrition cycle". In: *Journal of public Economics* 89.2-3, pp. 303–325.

Shefrin, Hersh M and Richard H Thaler (1988). "The behavioral life-cycle hypothesis". In: *Economic inquiry* 26.4, pp. 609–643.

Sokolova, Anna (2023). "Marginal propensity to consume and unemployment: A meta-analysis". In: *Review of Economic Dynamics* 51, pp. 813–846.

Souleles, Nicholas S (1999). "The response of household consumption to income tax refunds". In: *American economic review* 89.4, pp. 947–958.

Stephens Jr, Melvin (2003). ""3rd of tha month": do social security recipients smooth consumption between checks?" In: *American Economic Review* 93.1, pp. 406–422.

Stephens Jr, Melvin and Takashi Unayama (2011). "The consumption response to seasonal income: Evidence from Japanese public pension benefits". In: *American Economic Journal: Applied Economics* 3.4, pp. 86–118.

Sun, Liyang and Sarah Abraham (2021). "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects". In: *Journal of econometrics* 225.2, pp. 175–199.

Thaler, Richard (1985). "Mental accounting and consumer choice". In: *Marketing science* 4.3, pp. 199–214.

Wen, Yi (2002). "The business cycle effects of Christmas". In: *Journal of Monetary Economics* 49.6, pp. 1289–1314.

Zhang, C Yiwei et al. (2022). "How consumers budget". In: *Journal of Economic Behavior & Organization* 204, pp. 69–88.

A Appendices

A.1 Appendix: Data Construction

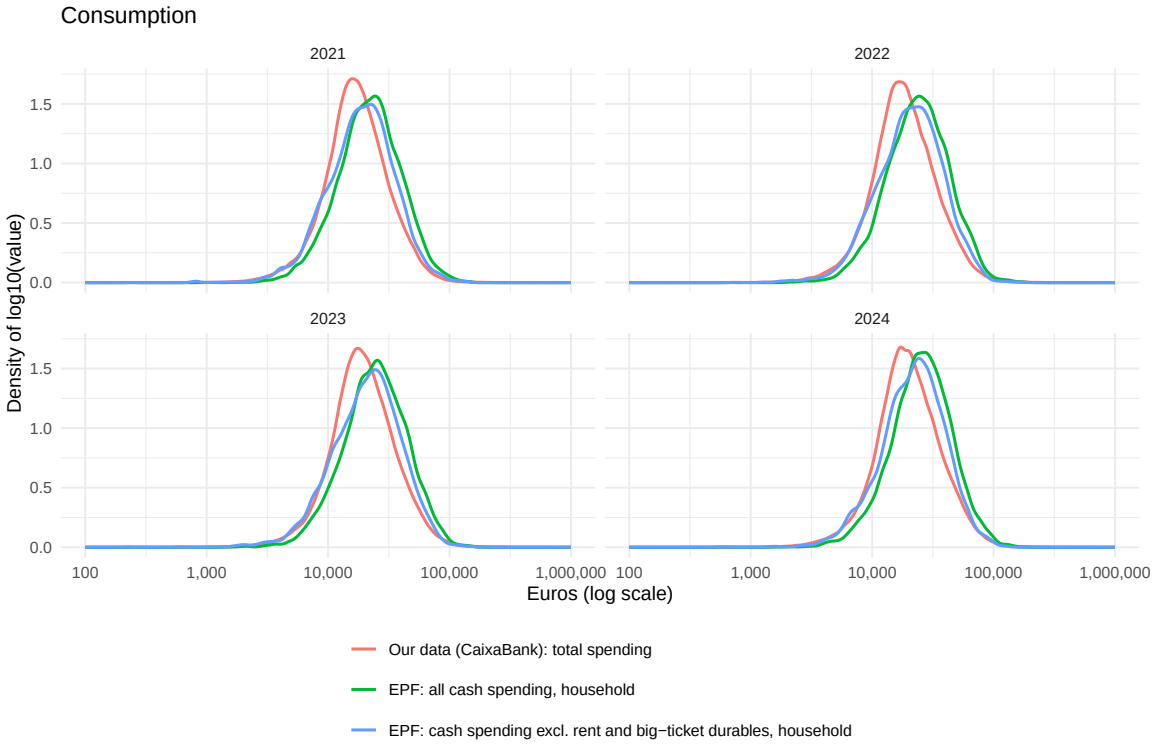


Figure A.1: Distribution of spending in our dataset compared with the distribution in the official data.

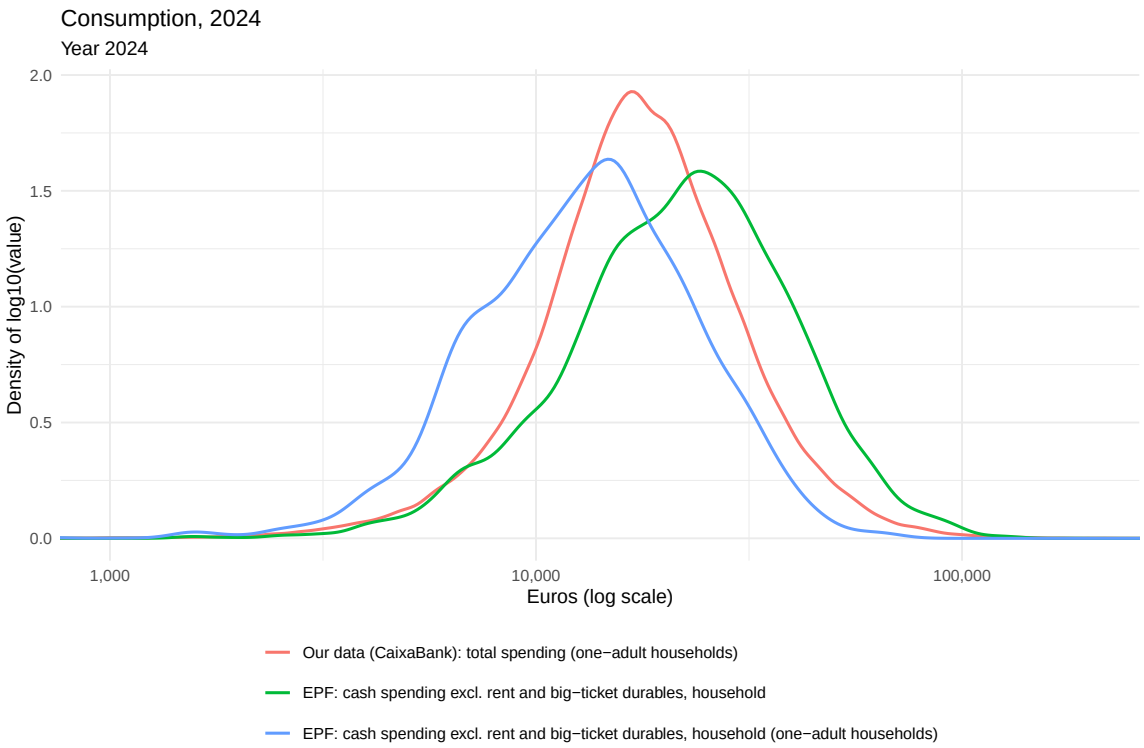


Figure A.2: Distribution of spending in our dataset compared with the distribution in the official data: total household versus one-adult households.

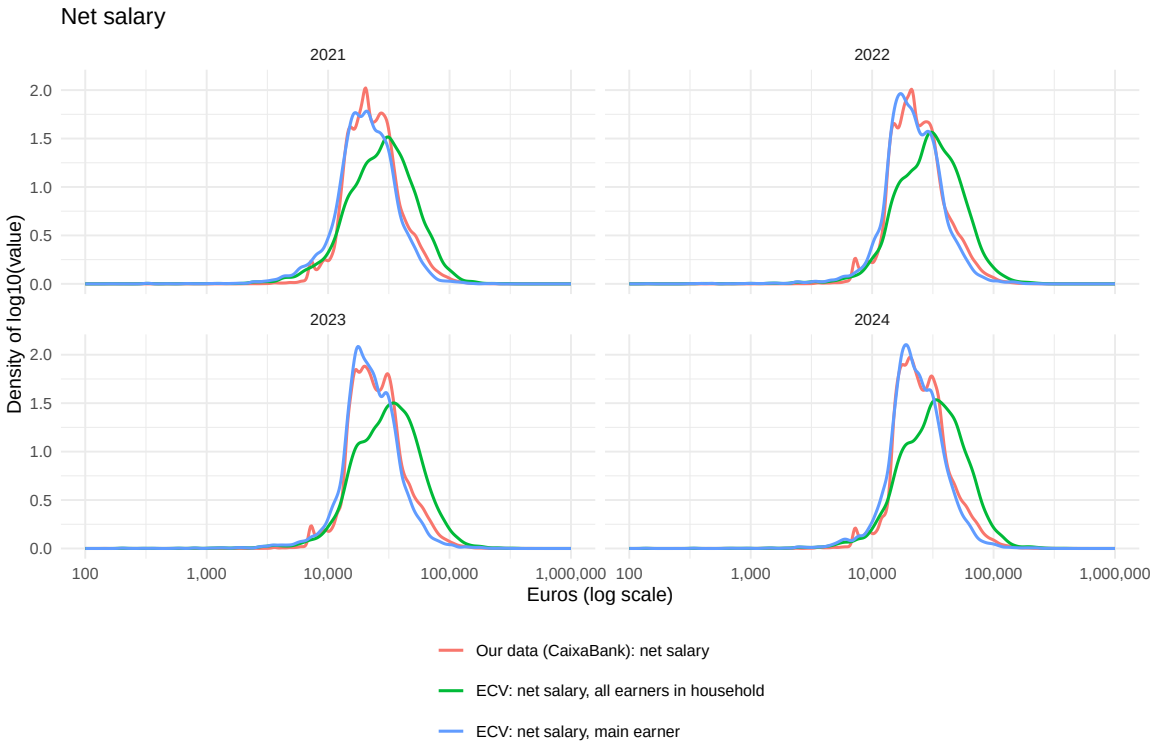


Figure A.3: Distribution of salary income in our dataset compared with the distribution in the official data.

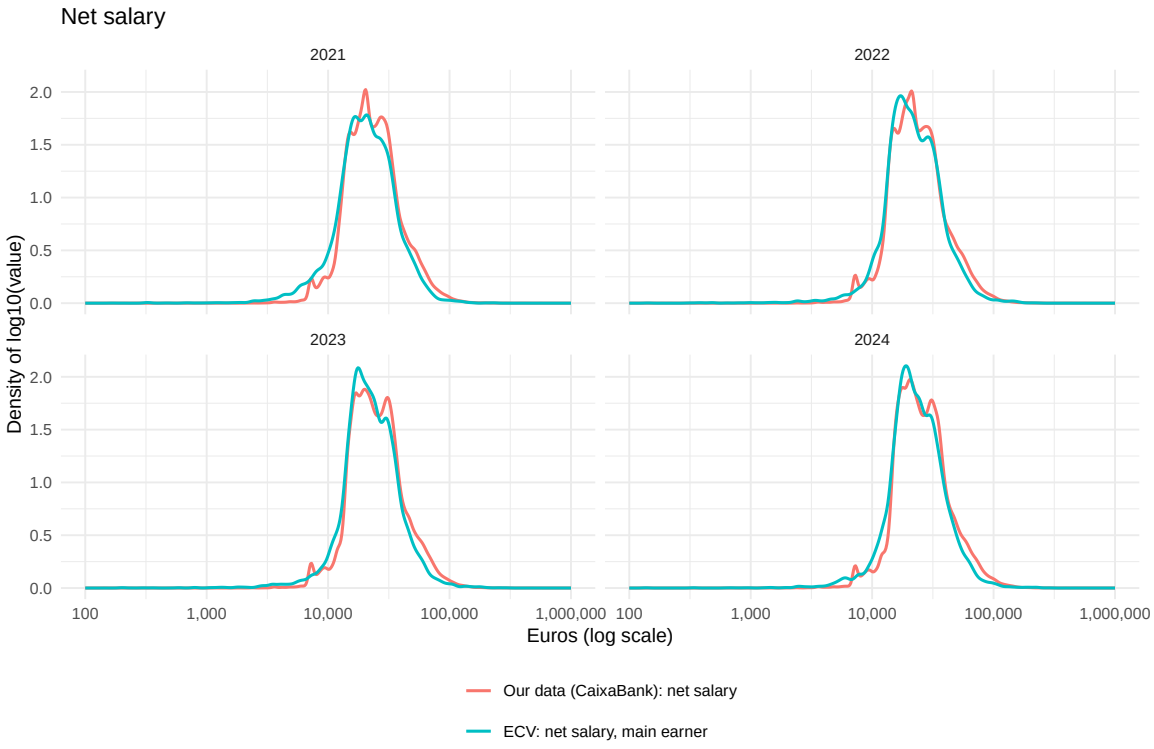


Figure A.4: Distribution of salary income in our dataset compared with the distribution in the official data.

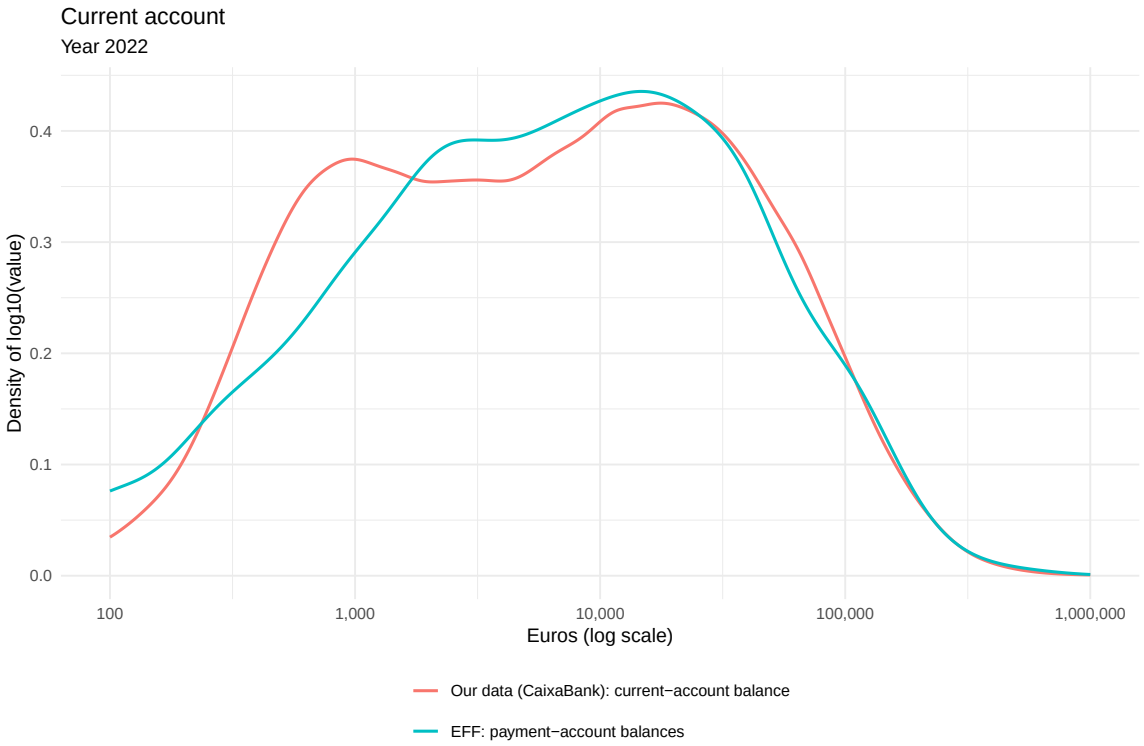


Figure A.5: Distribution of checking deposits in our dataset compared with the distribution in the official data.

Table A.1: Composition of household cash spending, 2021–2025

	2021	2022	2023	2024	2025	Euros, 2025
Total cash spending	**100.0**	**100.0**	**100.0**	**100.0**	**100.0**	**30,205**
Services	**43.3**	**44.6**	**46.4**	**47.8**	**47.6**	**14,388**
Restaurants and accommodation	10.6	12.5	13.5	13.3	12.7	3,838
Rent paid	5.6	5.6	6.4	6.4	6.6	1,988
Recreation, culture and package holidays	2.9	3.6	3.9	4.5	4.7	1,406
Insurance and financial services	4.3	3.9	3.8	4.1	4.1	1,225
Vehicle maintenance, parking and tolls	3.5	3.5	3.7	3.3	3.4	1,014
Water, refuse and community fees	3.2	3.0	2.9	3.1	3.2	959
Health services	2.0	1.8	2.0	3.3	3.0	914
Communication services	3.6	3.1	2.9	3.0	2.9	875
Education	2.5	2.4	2.4	2.5	2.7	805
Personal care and other services	2.1	2.2	2.0	1.8	2.0	592
Transport services	1.2	1.5	1.5	1.5	1.6	471
Domestic services and repairs	0.9	0.8	0.7	0.7	0.6	190
Housing maintenance and repairs	0.8	0.5	0.5	0.3	0.2	71
Transfers (pocket money, remittances)	0.2	0.2	0.1	0.0	0.1	40
Non-durable goods	**36.9**	**36.6**	**35.6**	**34.4**	**34.9**	**10,544**
Food and non-alcoholic drinks	19.2	18.3	18.7	18.4	18.8	5,690
Vehicle fuel	5.4	6.3	5.7	5.0	4.7	1,415
Household energy	4.5	4.9	4.2	3.8	4.0	1,193
Household and personal care products	3.0	3.0	3.1	3.0	3.1	942
Alcohol and tobacco	2.1	1.8	1.8	1.6	1.5	445
Pharmaceutical and medical products	1.2	1.1	1.0	1.1	1.3	394
Other (newspapers, stationery, garden, pets)	1.1	1.0	1.0	1.0	1.0	315
Housing maintenance materials	0.3	0.3	0.2	0.5	0.5	151
Durable goods	**10.8**	**10.1**	**9.3**	**8.8**	**8.9**	**2,696**
Vehicles (cars, motorbikes, bicycles)	5.4	5.4	4.8	5.2	5.5	1,669
ICT and audiovisual equipment	1.4	1.2	1.0	1.1	1.1	317
Furniture and furnishings	1.3	1.1	1.1	0.9	0.9	265
Major household appliances	0.8	0.7	0.7	0.6	0.6	181
Other (jewellery, large recreational items)	0.4	0.4	0.4	0.4	0.5	140
Therapeutic appliances (glasses, hearing aids)	1.6	1.3	1.4	0.4	0.4	124
Semi-durable goods	**9.0**	**8.7**	**8.7**	**9.0**	**8.5**	**2,577**
Clothing and footwear	5.6	5.4	5.5	5.7	5.5	1,673
Other (vehicle parts, small tools, personal effects)	1.1	1.2	1.2	1.3	1.2	375
Recreational goods, toys and books	1.4	1.4	1.2	1.3	1.1	330
Household textiles and small appliances	1.0	0.8	0.8	0.7	0.7	199
Memo: big-ticket durables						
Big-ticket durables, share of cash spending	7.5	7.2	6.5	6.9	7.2	2,169
Vehicles, share of cash spending	5.4	5.4	4.8	5.2	5.5	1,669
Big-ticket durables, share of durables	69.3	71.5	70.4	78.3	80.4	

Table A.2: From survey cash spending to national accounts consumption

	2021	2022	2023	2024
EPF monetary (cash) spending	22,525	24,713	25,731	26,510
+ imputed rent of owner-occupiers	6,456	6,489	6,653	7,306
+ other non-monetary (in kind, own production)	186	206	232	227
= EPF total consumption	**29,167**	**31,409**	**32,617**	**34,044**
+ gap to national accounts	7,187	8,465	9,386	10,310
= National accounts household consumption	**36,354**	**39,873**	**42,003**	**44,354**
Cash spending, share of national accounts (per cent)	62.0	62.0	61.3	59.8

Table A.3: Wealth statistics for Spain: calculated by the authors using the 2022 wave of the EFF survey (Banco de España).

	Euros							Ratio to monthly income		
	Total net wealth	Housing wealth	Financial wealth	Illiquid wealth	Liquid wealth	Current account	Income	Net wealth	Liquid wealth	Current account
Mean	315,613	238,554	72,225	263,263	46,296	27,562	38,069	125.42	15.06	9.36
p5	-2,872	0	60	0	0	21	7,200	-2.31	0.00	0.02
p10	390	0	308	0	100	200	9,973	0.30	0.07	0.12
p15	4,406	0	802	0	400	486	12,000	2.70	0.23	0.25
p20	17,956	0	1,469	4,000	872	870	14,000	9.52	0.46	0.42
p25	37,724	40,000	2,218	26,042	1,500	1,289	16,538	17.50	0.75	0.63
p30	57,291	64,967	3,520	47,400	2,036	2,000	18,867	24.77	1.09	0.90
p35	78,042	86,164	5,646	63,736	3,000	2,665	21,174	32.30	1.57	1.19
p40	95,905	105,000	8,000	81,538	5,000	3,564	23,777	39.05	2.25	1.60
p45	118,887	121,773	10,938	99,185	7,000	5,000	25,940	45.75	3.02	2.18
p50	143,031	147,280	15,005	120,000	10,000	7,000	28,939	53.46	3.97	2.85
p75	320,556	280,000	60,000	264,299	40,391	28,400	47,195	118.65	13.30	9.70
p90	669,915	547,936	159,575	557,189	112,527	77,000	75,989	231.20	35.18	22.13
p95	1,063,692	798,080	274,038	877,312	200,000	120,000	98,237	361.99	59.29	36.24

A.2 Appendix: Robustness to Covariates

Table A.4: Robustness: parallel trends conditional on covariates

Event	N treated	N control	Covariates	Baseline	Income effect (T)	MPC 1m	MPC 3m
1	119221	130046	Salary, Age	T-2	1830.935	0.077	0.259
2	87198	76592	Salary, Age	T-3	1876.856	0.125	0.185
3	79187	63213	Deposits, Salary, Age	T-3	1886.116	0.034	0.196
4	79502	65566	Deposits, Salary, Age	T-3	1913.392	0.115	0.170
5	76760	77385	Deposits, Salary, Age	T-3	1885.544	0.036	0.176
6	78892	84995	Deposits, Salary, Age	T-3	1831.620	0.115	0.181
7	77133	87173	Deposits, Salary, Age	T-3	1922.681	0.024	0.175
8	81595	96059	Deposits, Salary, Age	T-3	1877.934	0.150	0.251
9	81313	95955	Deposits, Salary, Age	T-3	2058.170	0.012	0.165
10	81692	98815	Deposits, Salary, Age	T-3	2005.048	0.116	0.206
11	81240	96345	Deposits, Salary, Age	T-3	2012.539	0.055	0.178
12	82267	100734	Deposits, Salary, Age	T-3	2080.310	0.104	0.169
Average					1931.762	0.080	0.192

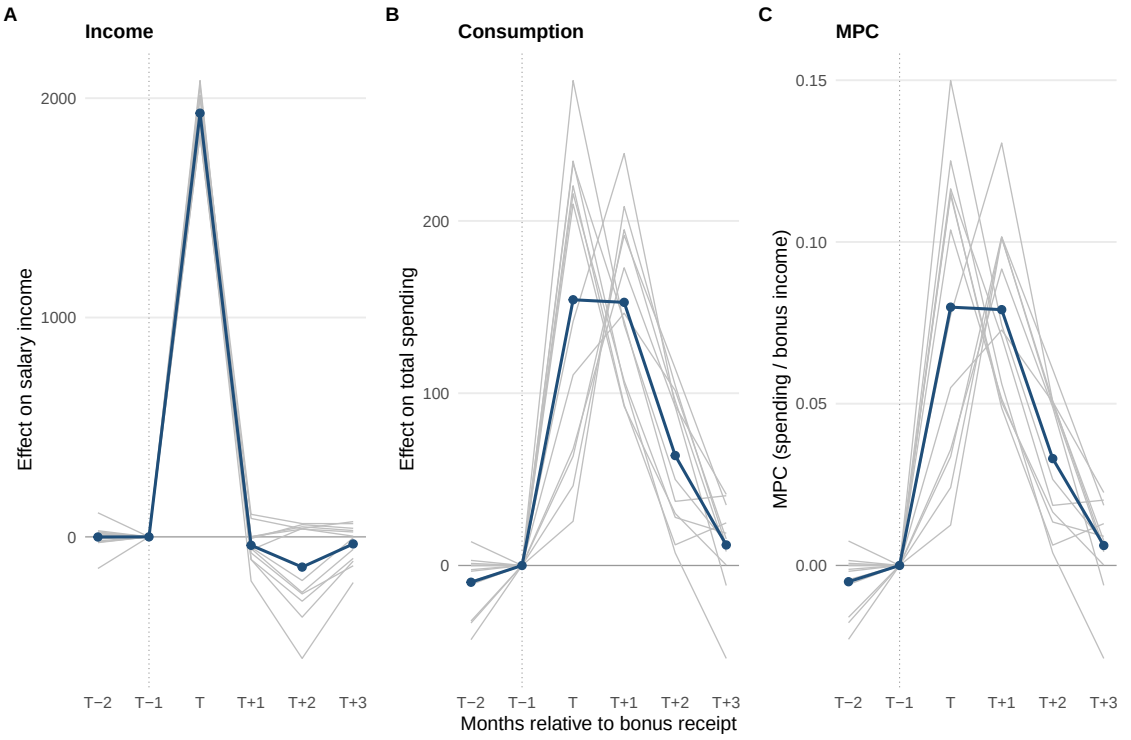


Figure A.6: Robustness check: Callaway & Sant'Anna (2021) with the Sant'Anna & Zhao (2020) doubly robust estimator, run separately for each bonus event. Average path in blue.

A.3 Appendix: Additional Graphs

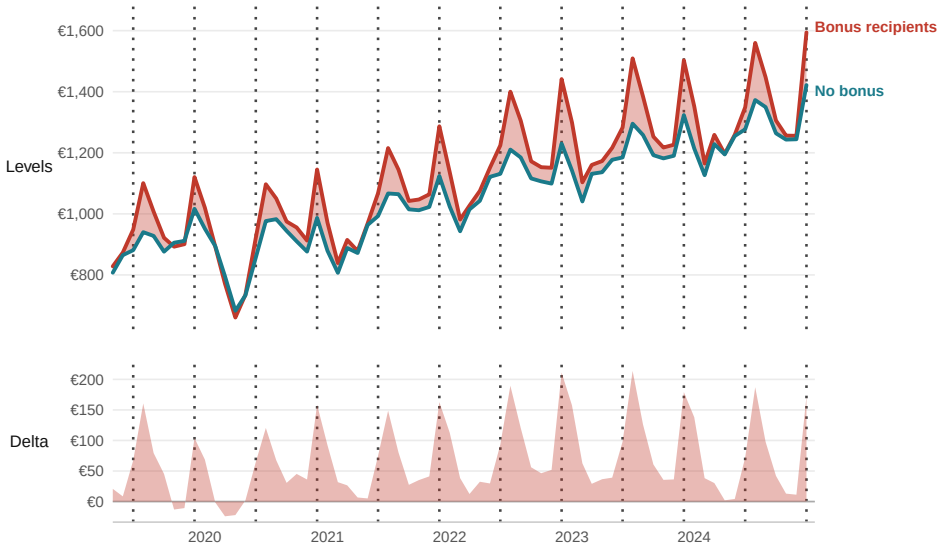


Figure A.7: Monthly card spending across the bonus and non-bonus groups (upper panel), and the delta between them (lower panel).

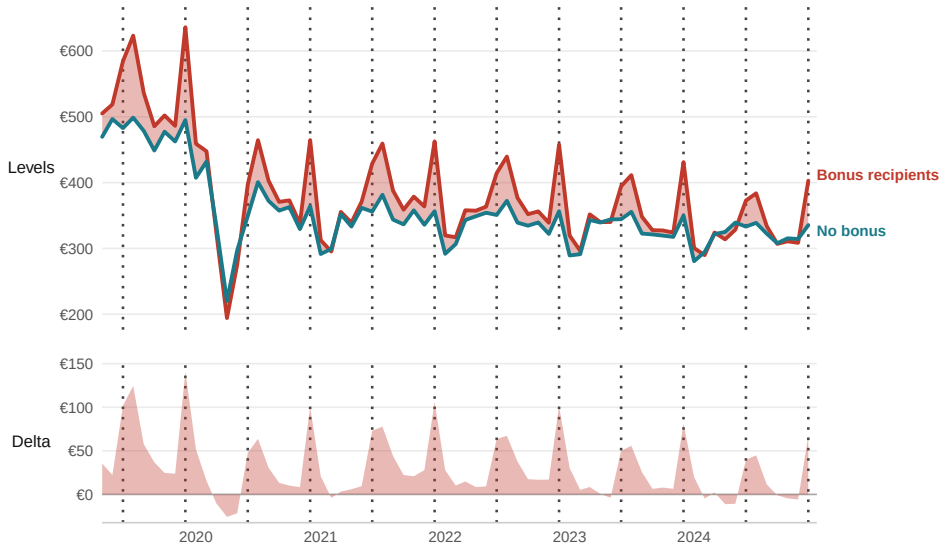


Figure A.8: Monthly cash withdrawals across the bonus and non-bonus groups (upper panel), and the delta between them (lower panel).

A.4 Appendix: More Results from Main Regressions

Table A.5: Seasonal Bonus MPC Estimates: all data pooled

	Estimate	Lower (95%)	Upper (95%)
1 Month	0.0916	0.0899	0.0934
3 Month	0.2075	0.2034	0.2117

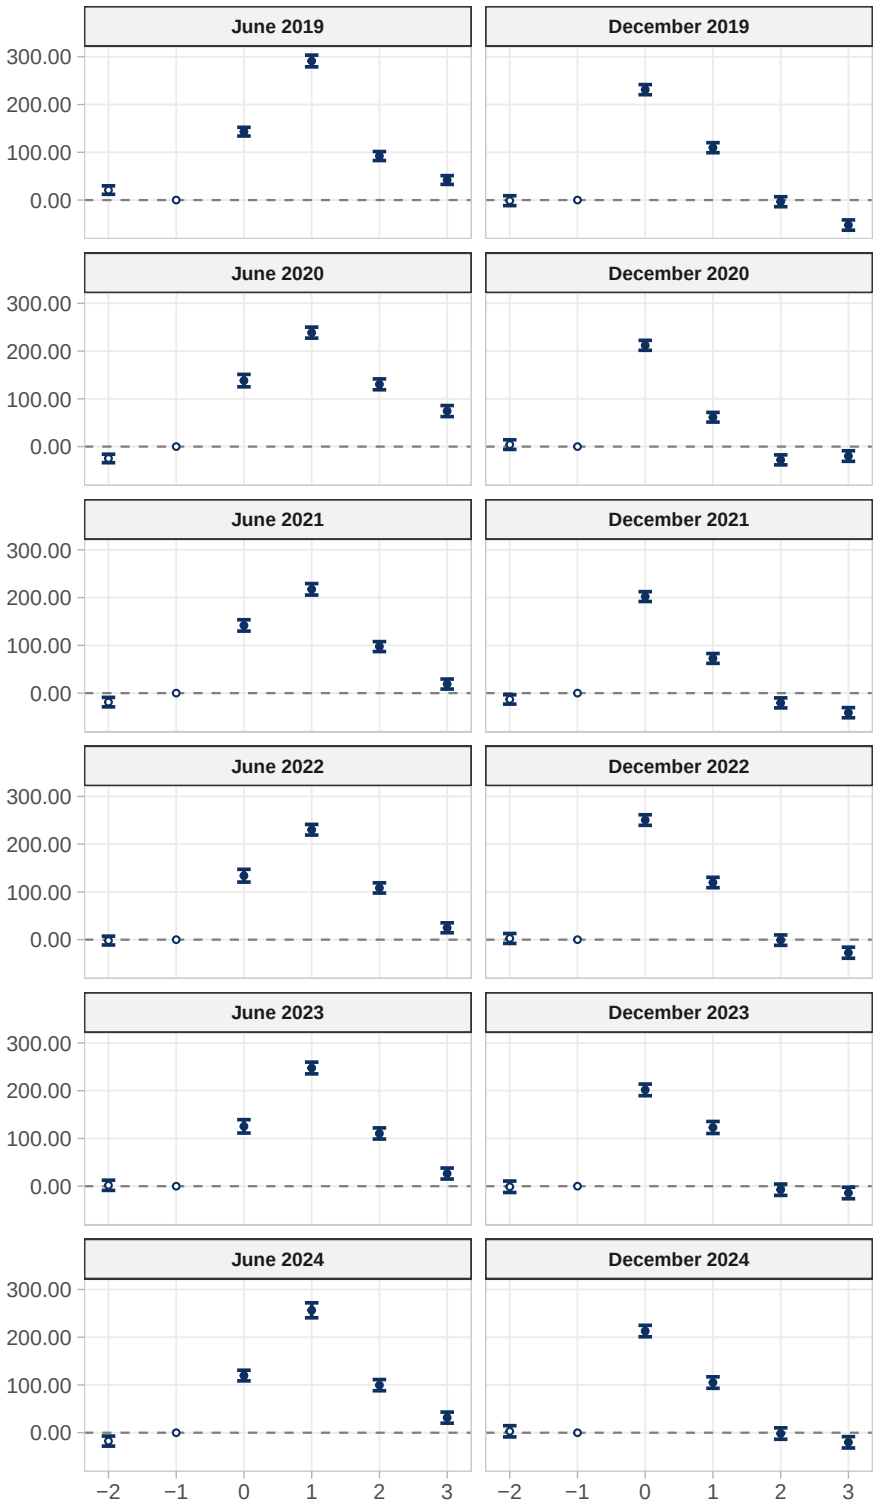


Figure A.9: Event-study plots for all 12 bonus events, using main sample of employed households. 95% confidence intervals shown, based on standard errors clustered at household level.

Table A.6: Stacked Regression Results

	Log salary	Log total spending
T-2	−0.001	0.002
	(0.006)	(0.006)
T	0.648***	0.100***
	(0.005)	(0.009)
T+1	−0.019***	0.097***
	(0.006)	(0.009)
T+2	−0.028***	0.027***
	(0.009)	(0.007)
T+3	0.011*	0.003
	(0.006)	(0.005)
Num.Obs.	16 520 400	16 520 400
R2	0.920	0.754
R2 Adj.	0.904	0.705

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

Standard errors are clustered at household by event and date level.
Period T is the month of bonus receipt. T-1 is the base period.

Table A.7: Stacked Regression Results: Levels

	Salary	Total spending
T-2	−1.713*	−3.330**
	(0.969)	(1.519)
T	1916.563***	175.644***
	(1.069)	(1.718)
T+1	−53.251***	173.989***
	(1.005)	(1.764)
T+2	−117.362***	48.078***
	(1.102)	(1.620)
T+3	−0.646	3.320**
	(1.100)	(1.633)
Num.Obs.	16 520 400	16 520 400
R2	0.853	0.690
R2 Adj.	0.823	0.628

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

Standard errors are clustered at household by event and date level.
Period T is the month of bonus receipt. T-1 is the base period.

Table A.8: Stacked Regression Results: June and December

	(1) Jun, C	(2) Jun, Y	(3) Dec, C	(4) Dec, Y
T-2	−5.643***	−4.361***	−0.986	0.969
	(2.043)	(1.469)	(2.251)	(1.261)
T	133.436***	1909.043***	218.405***	1924.182***
	(2.530)	(1.503)	(2.320)	(1.521)
T+1	247.883***	−33.236***	99.127***	−73.527***
	(2.650)	(1.473)	(2.324)	(1.367)
T+2	105.485***	33.902***	−10.081***	−270.605***
	(2.274)	(1.142)	(2.306)	(1.885)
T+3	35.133***	39.841***	−28.909***	−41.662***
	(2.249)	(1.248)	(2.369)	(1.818)
Num.Obs.	8 316 300	8 316 300	8 204 100	8 204 100
R2	0.671	0.876	0.711	0.833
R2 Adj.	0.606	0.851	0.653	0.799

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

Dependent variable is log spending.

Standard errors are clustered at household by event level.

Period T is the month of bonus receipt. T-1 is the base period.

Table A.9: Summary Statistics With Additional Detail

	Non-bonus HH (N = 125,889)					Bonus HH (N = 107,386)				
	Mean	SD	P10	Median	P90	Mean	SD	P10	Median	P90
Age	44.46	9.92	31.00	45.00	58.00	48.30	9.34	35.00	49.00	60.00
Number of adults	1.35	0.69	1.00	1.00	2.00	1.43	0.75	1.00	1.00	2.00
Monthly salary	2,176	1,795	990	1,612	3,975	2,491	1,495	1,221	2,108	4,291
Bonus month salary	2,162	1,659	1,000	1,617	3,926	4,156	2,012	2,346	3,728	6,637
Non-bonus month salary	2,164	1,627	1,000	1,637	3,955	2,159	1,064	1,198	1,930	3,458
Total spending	1,808	1,752	540	1,370	3,474	1,927	1,535	634	1,591	3,548
Card spending	1,136	1,148	189	845	2,365	1,202	1,029	228	968	2,421
Cash withdrawals	330	571	0	120	920	355	576	0	160	960
Checking deposits	17,394	43,601	246	3,085	46,218	24,426	43,493	575	8,393	66,208
Other CBK assets	33,118	76,420	0	2,547	104,048	33,763	65,812	0	4,702	102,976
Checking deposits / income	7.93	28.05	0.17	1.73	20.58	10.66	25.93	0.31	3.71	27.92
Borrowing from CBK	23,234	53,734	0	0	84,113	20,233	42,513	0	0	77,088

A.5 Appendix: Additional Results on Heterogeneity of Seasonal Bonus MPCs

In the paper, we showed results on the heterogeneity of seasonal bonus MPCs with respect to covariates, with the cut-offs chosen based on the distribution of all the data (not just seasonal bonus

recipients winners), with the objective being to compare across seasonal bonus and lottery MPCs with the same cut-offs. Here, we show results on the heterogeneity of seasonal bonus MPCs across the quintiles of the distribution of covariates among seasonal bonus recipients only.

The first element of heterogeneity we focus on is over time. Figure A.10 plots our estimates of the 3 month MPC in each of the event windows separately. Across years, the estimates are fairly stable, but they are consistently higher in June than December. It is evident in Figure A.9 that the spending responses in June tend to be spread over June and July, and are generally higher in July. In contrast, the December spending response is generally concentrated in December itself. This is likely the result of when bonuses are paid: June bonuses are normally paid out towards the end of the month, whilst in December they are generally paid halfway through the month. Bearing this in mind, in both cases the spending response is quite concentrated in the first month after bonus receipt.

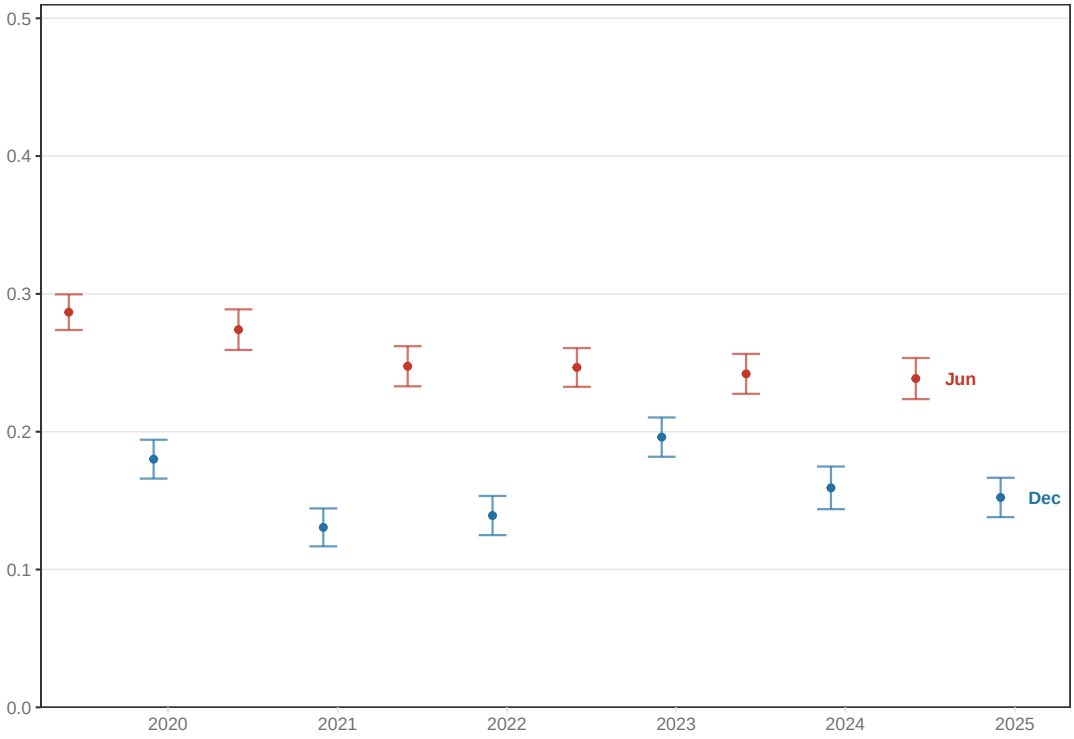


Figure A.10: Three month MPCs over time. 95% confidence intervals for the sum of the three coefficients are shown.

Next we exploit the cross-sectional variation in our data to examine how MPCs vary across groups.

Our bank account data includes three crucial covariates which we focus on: the age of the main account holder in the household, the salary of the household, and the current account balance of the household (which we refer to as 'cash' as it provides us with a proxy for the cash on hand of the household). Within each event window we calculate the average value for households of each of these covariates. [24]

We first show how the MPC varies stratifying by one covariate at a time: Figure A.11 shows the three month MPC within each quintile of each covariate. These MPCs are calculated restricting both treatment and control households to fall within the relevant quintile, so they are identified under the somewhat weaker assumption of conditional parallel trends. Consistent with much of the previous literature, the MPC is found to be steeply decreasing with respect to cash, somewhat decreasing in income, and mildly decreasing with age. MPCs are greater than 0.1 in all five quintiles of all three covariates.

[24] Because the event windows are six months with one bonus month, the average salary over the window reflects the annualized salary, incorporating the bonus appropriately.

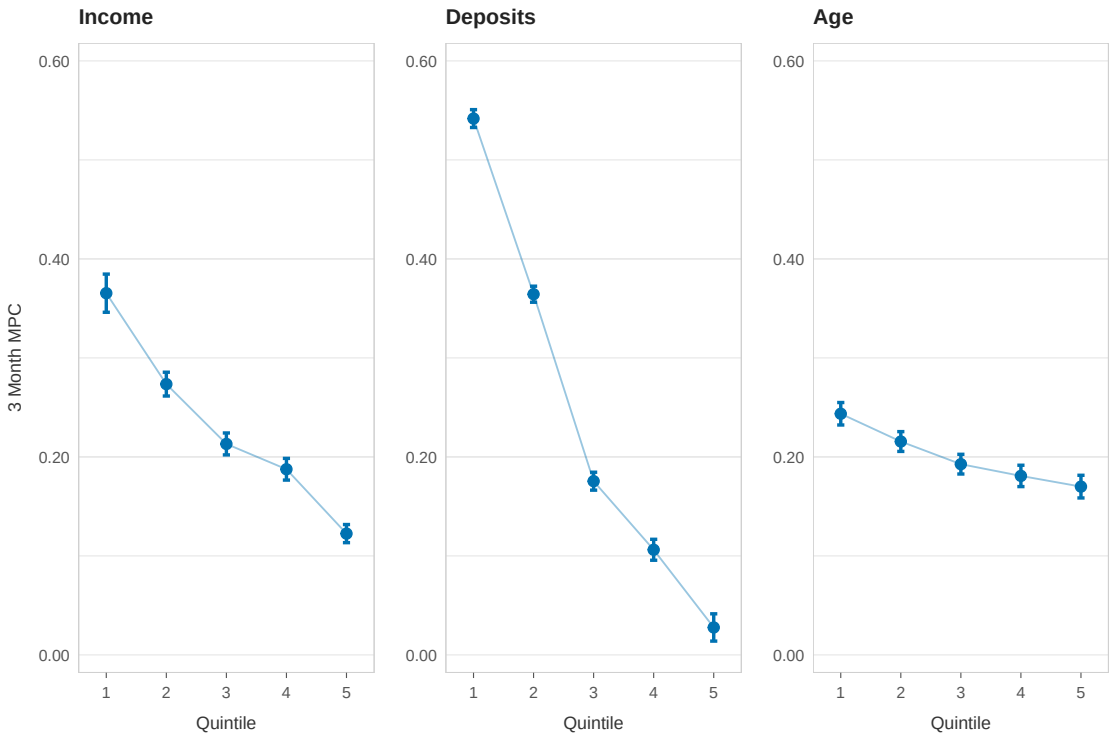


Figure A.11: Three month MPCs estimated within each quintile of income, cash, and age. 95% confidence intervals shown, based on standard errors clustered at household by event level.

Since our dataset tracks roughly 250,000 households, we can analyse MPCs within more granular cells. We sort households into groups based on the 5x5 intersection of quintiles of each pair of covariates: for example one cell is those households which are in the first quintile for income and the first quintile for cash. Since the covariates are correlated, the number of households within each cell varies a lot. The smallest cell is households in the first quintile for cash and the fifth quintile for income - this cell still contains 1,371 households.

The payoff of this is demonstrated by Figure A.12, which shows a separate panel for each cash quintile, and within each panel plots the MPCs for households in that cash quintile, split by which income quintile they fall into. [25] The figure makes it clear that for households above the bottom quintile for cash, income is not a very important driver of MPC heterogeneity. For low cash households, income

[25] To be clear: the stratification is by the intersection of the unconditional quintiles of both variables. We are not looking at the quintiles of the cash distribution within each income quintile. So the left-most point on the graph represents the 3 month MPC for households in the first unconditional income quintile and the first unconditional cash quintile.

does seem to matter for MPCs, although we cannot rule out that this reflects a correlation between income and cash within the bottom cash quintile. The overall message is that cash balance seems to matter a lot for MPCs, independently of income.

The other important result is that 22 of the 25 cells (which include 92% of households) have MPCs greater than 0.1. Households in the top quintile for income and cash have an MPC of 0.07 (and significantly different from zero). Figure A.13 shows a similar graph stratifying by income and age quintiles. MPCs are decreasing with income within each age quintile, but as per the preceding discussion, this likely largely reflects the correlation of income with cash. In this partition of the data, MPCs are greater than 0.1 in all cells.

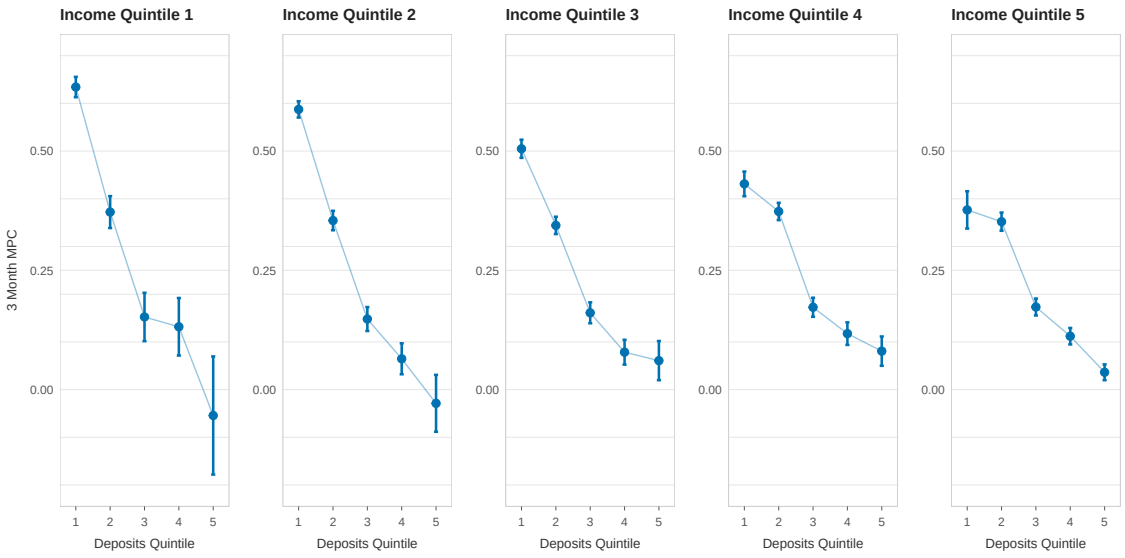


Figure A.12: Three month MPCs estimated within each of the 25 cells created from stratifying the data by income quintile and cash quintile. 95% confidence intervals shown, based on standard errors clustered at household by event level.

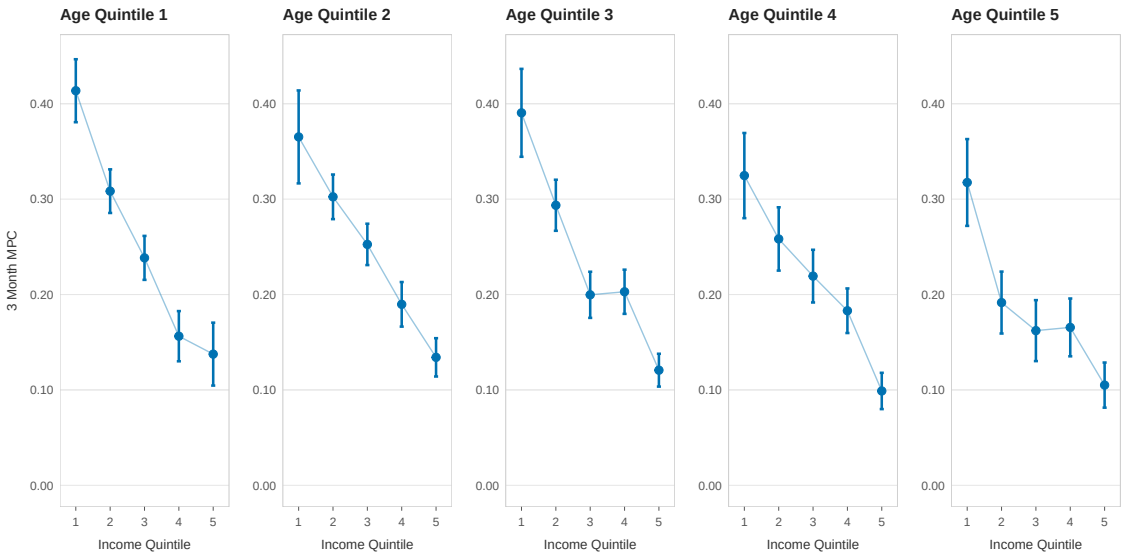


Figure A.13: Three month MPCs estimated within each of the 25 cells created from stratifying the data by age quintile and income quintile. 95% confidence intervals shown, based on standard errors clustered at household by event level.

A.6 Appendix: Additional Results on Heterogeneity of Lottery Win MPCs

In the paper, we showed results on the heterogeneity of lottery MPCs with respect to covariates, with the cut-offs chosen based on the distribution of all the data (not just lottery winners), with the objective being to compare across lottery and seasonal bonus MPCs with the same cut-offs. Here, we show results on the heterogeneity of lottery MPCs across the quintiles of the distribution of covariates amongst lottery winners only. We form quintiles based on the value of each covariate in the six months before winning the lottery, and estimate MPCs within each quintile for each covariate. For imputing the counter-factual, we use only controls which lie within the same quintile. We perform this exercise with the set of seasonal-bonus-matched prizes, and also with all prizes below 10,000 euros.

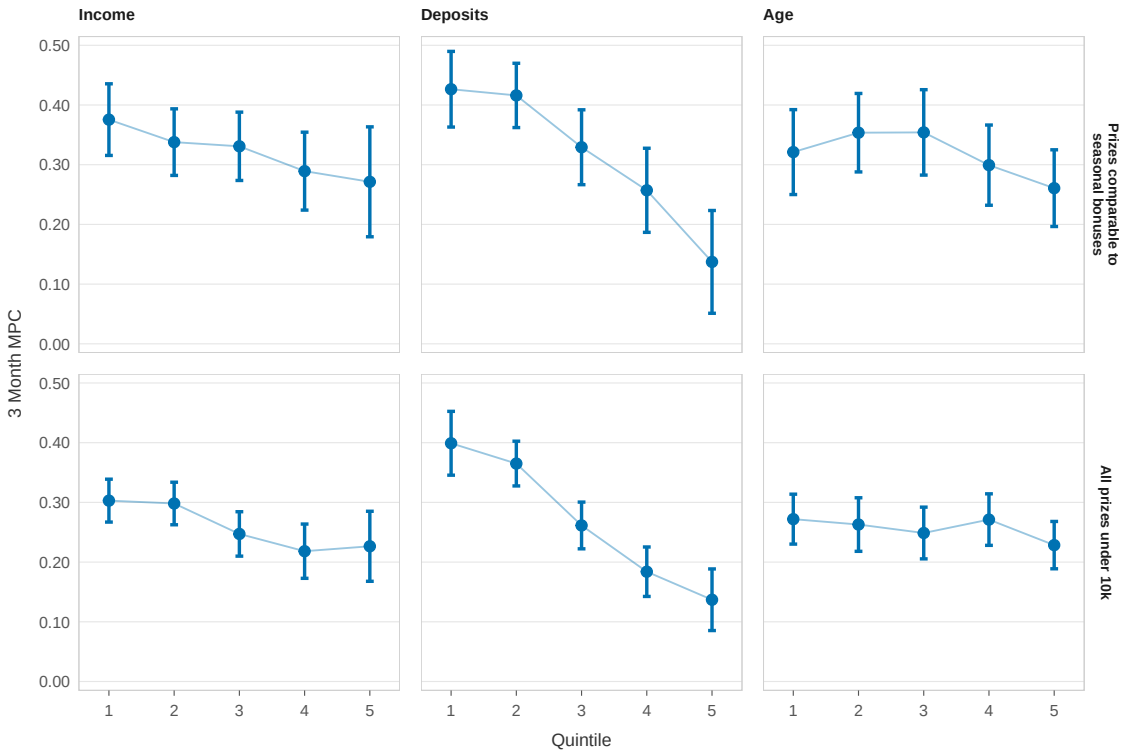


Figure A.14: MPC event studies for lottery winnings of comparable size to seasonal bonuses, heterogeneity with respect to key covariates.

Figure A.14 summarizes our results, plotting the three month (cumulative) MPC across the quintiles of each variable for the two samples. The patterns are similar to our results from the seasonal bonuses, but tempered. MPCs decline slightly with income, and more so with deposits. The bottom quintile of deposit holders have three month MPCs of 0.43, falling to 0.14 in the top quintile - a lesser gradient than for seasonal bonus MPCs, but still substantial. Having already restricted ourselves to prizes in the same size range of the seasonal bonuses, our sample size is too small to facilitate estimating MPCs across the full 5x5 grid of deposits and income as in the previous section.

A.7 Appendix: Distribution of Main Covariates

Table A.10: Summary Statistics: Key Covariate Quintiles

Quintile	Mean	Median	SD
Age			
Quintile 1	33	33	3
Quintile 2	42	42	2
Quintile 3	48	48	1
Quintile 4	53	53	2
Quintile 5	60	59	4
Income			
Quintile 1	1,163	1,258	260
Quintile 2	1,623	1,621	108
Quintile 3	2,070	2,058	160
Quintile 4	2,725	2,702	220
Quintile 5	4,816	4,228	2,009
Cash deposits			
Quintile 1	424	425	216
Quintile 2	1,596	1,489	564
Quintile 3	5,536	5,237	1,920
Quintile 4	17,478	16,578	5,609
Quintile 5	79,529	56,531	72,495

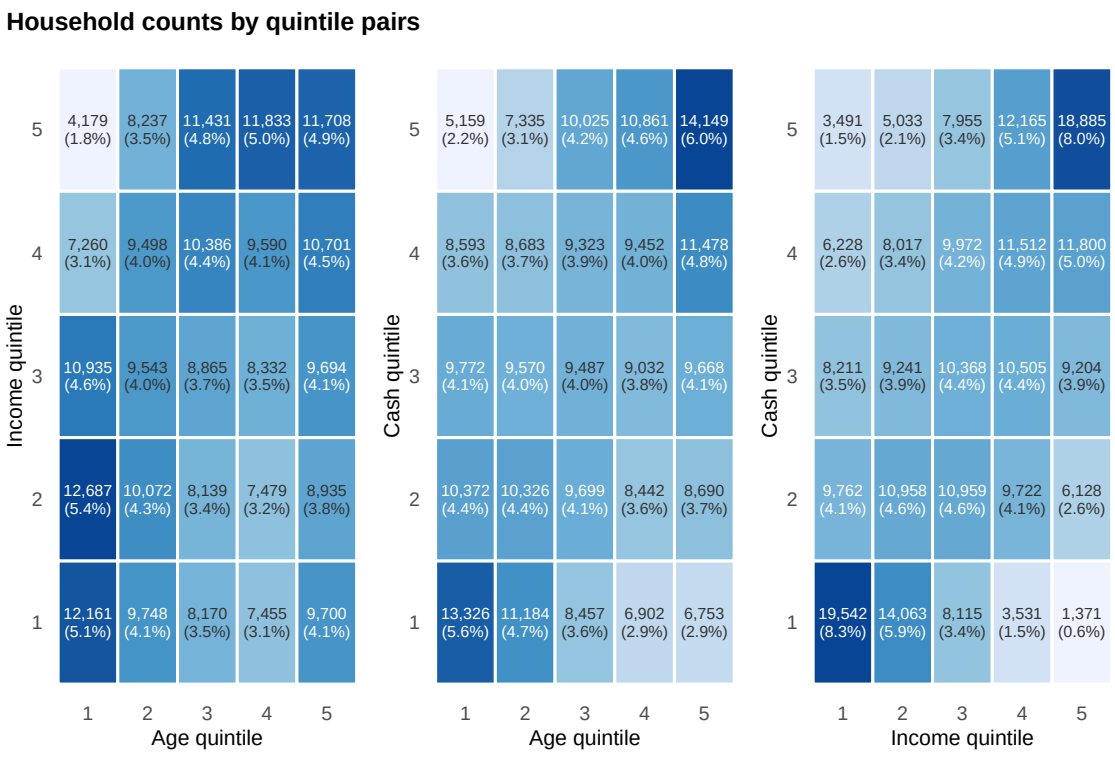


Figure A.15: Number of households within each cell when stratifying by quintiles of two variables simultaneously.

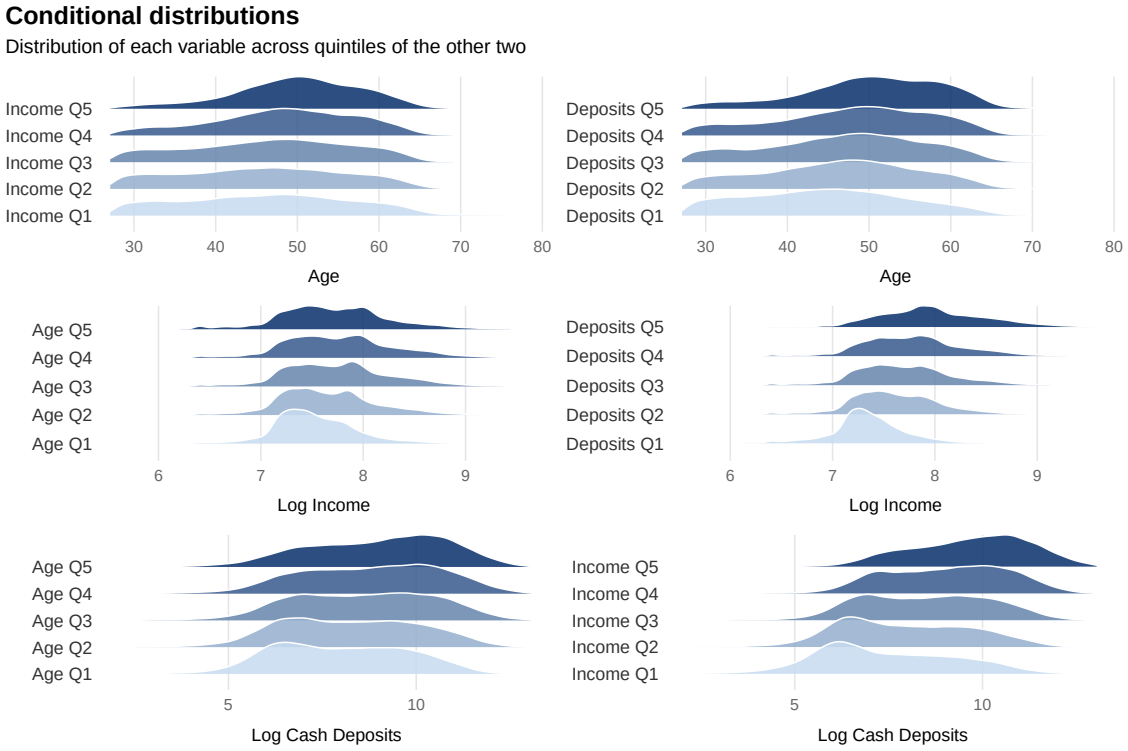


Figure A.16: Number of households within each cell when stratifying by quintiles of two variables simultaneously.

A.8 Appendix: Additional Material on Lotteries

Table A.11: Decomposition of lottery MPC by category

Category	Effect	S.E.	Share (%)	Cumulative (%)
Cash withdrawals	0.1064***	(0.0400)	37.2	37.2
Insurance bills	0.0266	(0.0283)	9.3	46.5
Restaurants	0.0225***	(0.0079)	7.9	54.3
Culture and shows	0.0208**	(0.0094)	7.2	61.6
Other retail	0.0183**	(0.0091)	6.4	68.0
Other bills	0.0109	(0.0129)	3.8	71.8
Food, drinks, tobacco	0.0069	(0.0099)	2.4	74.2
Gas stations	0.0059	(0.0037)	2.0	76.2
Travel agencies	0.0049	(0.0086)	1.7	77.9
Other essential goods	0.0043	(0.0056)	1.5	79.4
Home appliances, IT	0.0015	(0.0051)	0.5	79.9
Hotels	0.0012	(0.0086)	0.4	80.4
Pharmacies	0.0001	(0.0022)	0.0	80.4
Fashion	-0.0004	(0.0065)	-0.1	80.3
Furniture	-0.0013	(0.0052)	-0.5	79.8
Utility bills	-0.0027	(0.0037)	-0.9	78.9
Transport	-0.0033	(0.0097)	-1.1	77.7
Residual	0.0637		22.3	
TOTAL (Total spending)	0.2864***	(0.0676)	100.0	

Notes: Sum of coefficients on the bonus indicator at t, t+1 and t+2. Standard errors clustered by household × event. * p<0.10, ** p<0.05, *** p<0.01.

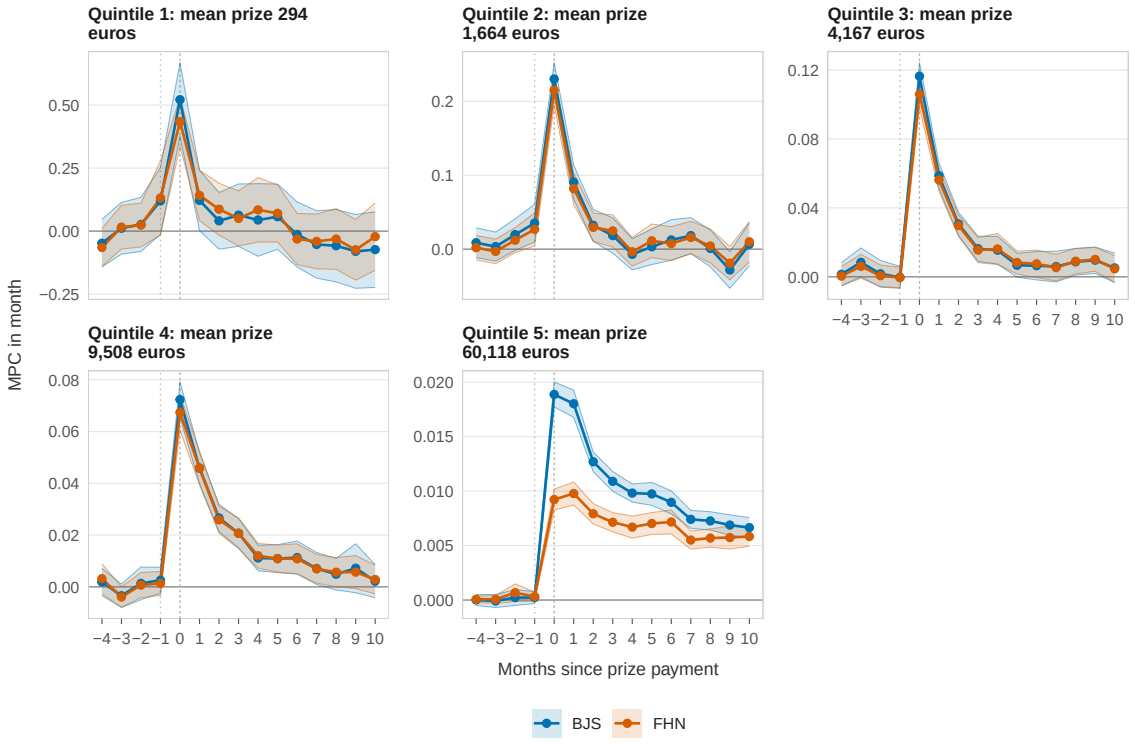


Figure A.17: MPC event studies for lottery winnings, facetted by win quintile, comparing BJS estimator with FHN method.

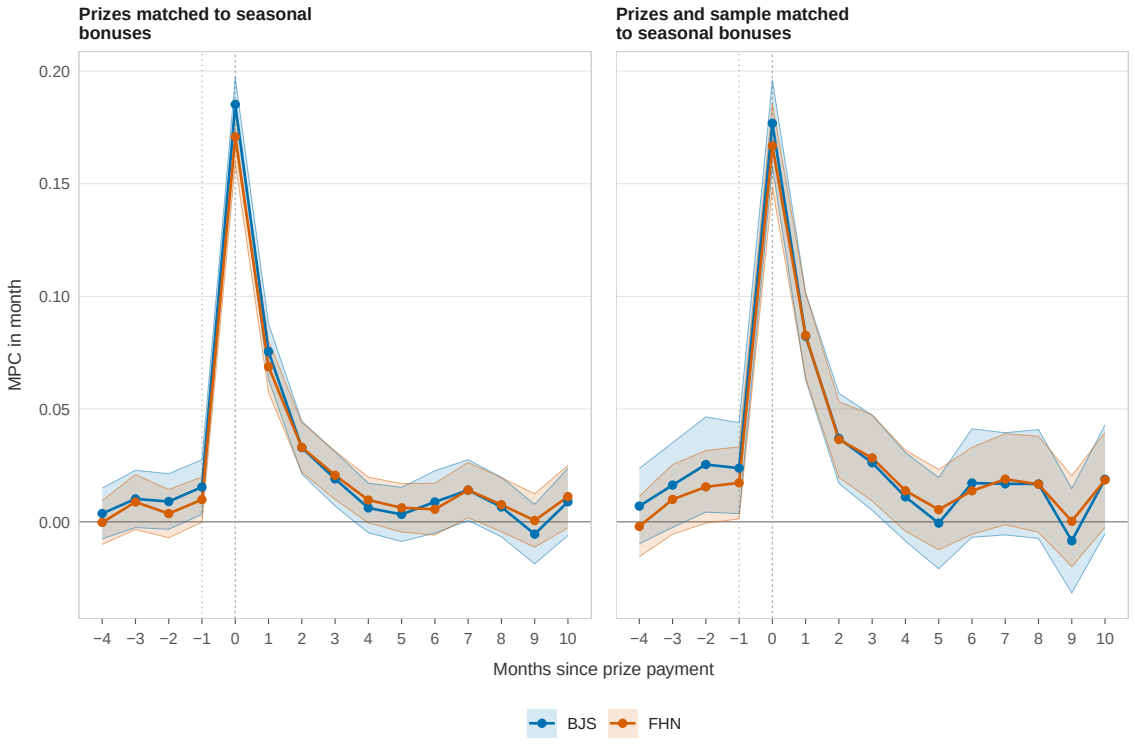


Figure A.18: MPC event studies for lottery winnings of size comparable to seasonal bonus, comparing BJS estimator with FHN method.

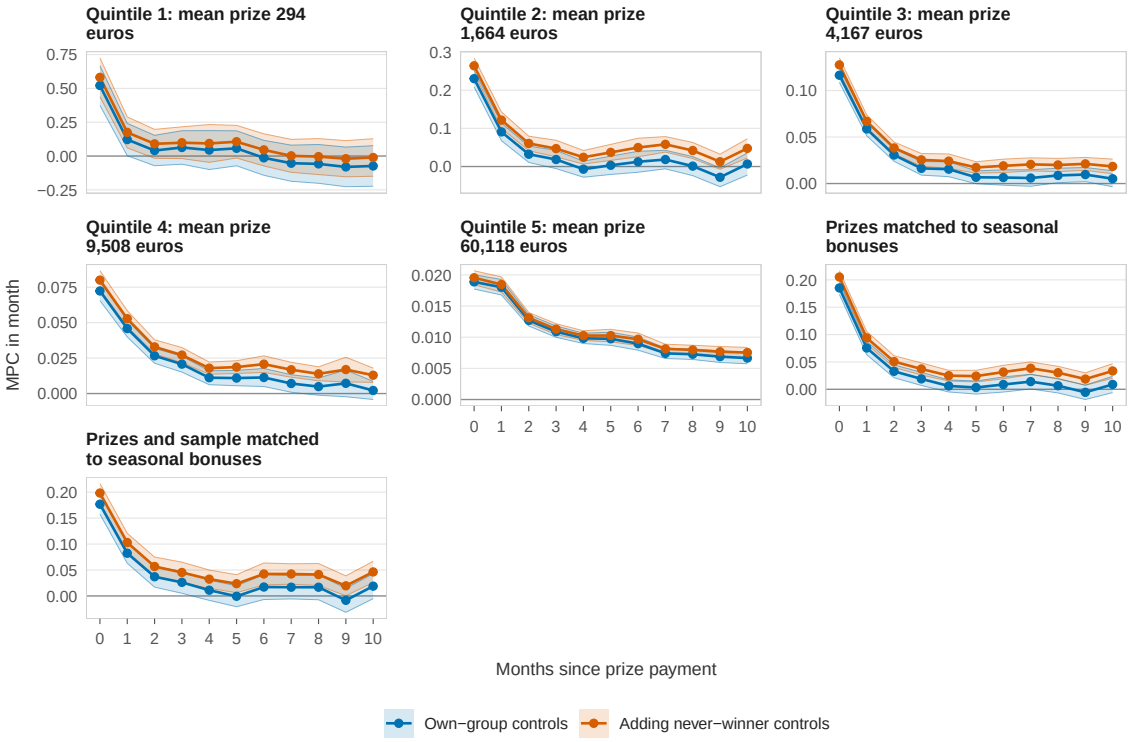


Figure A.19: MPC event studies for lottery winnings, using not-yet-winners as controls versus using a selection of characteristic-matched never-winners.

Table A.12: Full event study results.

	Prize quintile					Matched to seasonal bonuses	
Month	1	2	3	4	5	Prizes	Prizes and sample
-4	-0.048	0.009	0.001	0.002	0.000	0.004	0.007
	(0.048)	(0.010)	(0.003)	(0.003)	(0.000)	(0.006)	(0.009)
-3	0.011	0.003	0.008	-0.003	-0.000	0.010	0.016
	(0.052)	(0.010)	(0.004)	(0.002)	(0.000)	(0.006)	(0.010)
-2	0.027	0.020	0.002	0.001	0.000	0.009	0.025
	(0.055)	(0.011)	(0.004)	(0.003)	(0.000)	(0.006)	(0.011)
-1	0.120	0.035	-0.000	0.003	0.000	0.015	0.024
	(0.068)	(0.013)	(0.003)	(0.003)	(0.000)	(0.006)	(0.010)
0	0.521	0.230	0.116	0.072	0.019	0.185	0.177
	(0.075)	(0.011)	(0.004)	(0.003)	(0.001)	(0.006)	(0.010)
1	0.122	0.091	0.059	0.046	0.018	0.075	0.082
	(0.061)	(0.012)	(0.004)	(0.003)	(0.001)	(0.006)	(0.010)
2	0.041	0.032	0.031	0.027	0.013	0.033	0.037
	(0.057)	(0.011)	(0.003)	(0.003)	(0.000)	(0.006)	(0.010)
3	0.063	0.018	0.016	0.021	0.011	0.019	0.026
	(0.063)	(0.012)	(0.004)	(0.003)	(0.000)	(0.006)	(0.011)
4	0.044	-0.007	0.015	0.011	0.010	0.006	0.011
	(0.073)	(0.011)	(0.004)	(0.003)	(0.000)	(0.006)	(0.010)
5	0.057	0.003	0.007	0.011	0.010	0.003	-0.001
	(0.066)	(0.012)	(0.004)	(0.003)	(0.001)	(0.006)	(0.010)
6	-0.013	0.012	0.006	0.011	0.009	0.009	0.017
	(0.066)	(0.014)	(0.004)	(0.003)	(0.001)	(0.007)	(0.012)
7	-0.053	0.018	0.006	0.007	0.007	0.014	0.017
	(0.068)	(0.012)	(0.005)	(0.003)	(0.000)	(0.007)	(0.012)
8	-0.058	0.001	0.009	0.005	0.007	0.007	0.017
	(0.073)	(0.013)	(0.004)	(0.003)	(0.000)	(0.007)	(0.012)
9	-0.080	-0.028	0.010	0.007	0.007	-0.005	-0.008
	(0.075)	(0.013)	(0.004)	(0.005)	(0.000)	(0.007)	(0.012)
10	-0.073	0.007	0.005	0.002	0.007	0.009	0.019
	(0.076)	(0.015)	(0.004)	(0.003)	(0.000)	(0.008)	(0.012)
11	0.034	-0.027	0.004	0.006	0.006	-0.013	0.002
	(0.083)	(0.013)	(0.004)	(0.004)	(0.000)	(0.006)	(0.011)
12	-0.088	-0.010	-0.000	0.009	0.006	-0.001	0.001
	(0.078)	(0.014)	(0.005)	(0.004)	(0.000)	(0.008)	(0.012)

Note:
MPC in each month relative to payment; standard errors in parentheses. Month -1 is the anticipation month; months -4 to -2 are pre-trend coefficients relative to months -5 and earlier.

Table A.13: Lottery cummulative MPCs: 1, 3, and 5 months.

Group	1 month	3 months	5 months
Quintile 1: mean prize 294 euros	0.521 (0.075)	0.727 (0.145)	0.833 (0.221)
Quintile 2: mean prize 1,664 euros	0.230 (0.011)	0.346 (0.026)	0.365 (0.040)
Quintile 3: mean prize 4,167 euros	0.116 (0.004)	0.206 (0.008)	0.237 (0.012)
Quintile 4: mean prize 9,508 euros	0.072 (0.003)	0.145 (0.007)	0.177 (0.011)
Quintile 5: mean prize 60,118 euros	0.019 (0.001)	0.049 (0.001)	0.070 (0.002)
Prizes matched to seasonal bonuses	0.185 (0.006)	0.292 (0.014)	0.321 (0.021)
Prizes and sample matched to seasonal bonuses	0.177 (0.010)	0.296 (0.023)	0.340 (0.035)

Table A.14: Characteristics of households who receive seasonal bonuses compared with lottery winners.

	Lottery winners		Seasonal bonus recipients		Difference	
	Mean	Median	Mean	Median	Mean	Median
Monthly income (euros)	2,321	2,008	2,497	2,251	-175	-243
Checking deposits (euros)	24,995	7,162	25,036	9,174	-41	-2,012
Age (years)	59.6	60.0	48.9	49.3	10.7	10.7
Spending / income	1.06	0.87	0.79	0.74	0.27	0.13
Bonus or prize (euros)	2,266	2,276	2,140	1,929	126	347
Observations	8,746		799,518			
Households	8,746		193,847			

A.9 Appendix: Alternate re-weighting of seasonal bonus MPC

In Section 5, we showed results from re-weighting the seasonal bonus MPC to the lottery population distribution based on the quintiles of deposits, age, and spending relative to income. Here we show the robustness of that exercise with respect to different re-weighting techniques. Table A.15 shows the results from re-weighting by: (i) quintiles of the three covariates, as in the main paper, (ii) giving more granularity to covariates more predictive of the MPCs: deciles of deposits, terciles of age, and quintiles of spending relative to income, (iii) propensity score, estimated with a natural spline with 5 degrees of freedom for each covariate. The different methods do not make much difference: the share of the difference explained is close to 50%.

Table A.15: Results from re-weighting seasonal bonus MPC with alternative weighting schemes.

Method	Change from reweighting	Share explained
Main: quintiles of deposits, age and spending / income	0.0464	0.4979
Deposit deciles x age bands x spending / income quintiles	0.0503	0.5394
Propensity score	0.0425	0.4555

A.10 Appendix: Proofs

A.10.1 Proof of average consumption identity:

Start from the household budget constraint, for household i in month t:

$$C_{i,t} = R\, a_{i,t} - a_{i,t+1} + Y_{i,t}.$$

Average across households. Everything is linear and R is common, so the cross-sectional mean passes through exactly:

$$\bar{C}_t = R\, \bar{a}_t - \bar{a}_{t+1} + \bar{Y}_t.$$

Multiply by R^{-s} at $t = T + s$ and sum over one full cycle, $s = 0, \ldots, 5$, anchoring the sum at the bonus month:

$$\sum_{s=0}^{5} R^{-s}\, \bar{C}_{T+s} = \sum_{s=0}^{5} \left(R^{1-s}\, \bar{a}_{T+s} - R^{-s}\, \bar{a}_{T+s+1} \right) + \sum_{s=0}^{5} R^{-s}\, \bar{Y}_{T+s}.$$

The asset sum mostly cancels: only the two endpoints survive. Imposing periodic stationarity, $\bar{a}_{T+6} = \bar{a}_T$, the asset sum is therefore:

$$R\, \bar{a}_T - R^{-5}\, \bar{a}_{T+6} = \bar{a}_T\, R\big(1 - R^{-6}\big).$$

Now divide through by the 6-month present-value factor $S_6^{PV} \equiv \sum_{s=0}^{5} R^{-s} = \big(1 - R^{-6}\big)/\big(1 - R^{-1}\big)$.

The coefficient on assets collapses exactly:

$$\frac{R\big(1 - R^{-6}\big)}{S_6^{PV}} = R\big(1 - R^{-6}\big) \cdot \frac{1 - R^{-1}}{1 - R^{-6}} = R\big(1 - R^{-1}\big) = R - 1 = r.$$

The income side is w^* in every month of the cycle, plus an extra w^* in the bonus month itself, which is undiscounted because the sum is anchored there. So $\sum_{s=0}^{5} R^{-s}\,\bar{Y}_{T+s} = w^*\big(1 + S_6^{PV}\big)$, and dividing by S_6^{PV} gives exactly one, since w^* is defined by present-value neutrality against a flat unit stream, $w^* = S_6^{PV}/\big(1 + S_6^{PV}\big)$.

Collecting terms, the present-value-weighted average of mean consumption across the cycle is

$$\widetilde{C} \;\equiv\; \frac{\sum_{s=0}^{5} R^{-s}\,\bar{C}_{T+s}}{S_6^{PV}} \;=\; 1 + r\,\bar{a}_T,$$

where $\bar{a}_T$ denotes mean assets entering the bonus month. In words: in a periodic steady state the household is not running its buffer down from one cycle to the next, so over a full cycle it consumes its income plus the interest the buffer earns, and no more.

A.10.2 Proof of bound on $\bar{C}^b_{T-1}$

The previous section derived:

$$\sum_{s=0}^{5} R^{-s}\,\bar{C}_{T+s} \;=\; \Big(\sum_{s=0}^{5} R^{-s}\Big)\big(1 + r\,\bar{a}_T\big).$$

The identity fixes a present-value-weighted average of the six monthly means. To convert it into an upper bound on the single month $\bar{C}_{T-1}$, each month must be bounded *below* by a multiple of $\bar{C}_{T-1}$.

For $s = 0, \ldots, 4$, chain Lemma 1 forward from month $T-1$ to month $T+s$.

$$\bar{C}_{T+s} \;\geq\; R_\beta^{s+1}\,\bar{C}_{T-1}, \qquad s = 0, \ldots, 4.$$

Chaining Lemma 1 six steps would give $\bar{C}_{T+5} \geq R_\beta^6\,\bar{C}_{T-1}$. But month $T+5$ occupies the same phase of the cycle as month $T-1$, so periodic stationarity supplies a tighter restriction:

$$\bar{C}_{T+5} \;=\; \bar{C}_{T-1}.$$

Substituting these steps into the right hand side of the identity,

$$\sum_{s=0}^{5} R^{-s}\bar{C}_{T+s} \;=\; \Big(\sum_{s=0}^{5} R^{-s}\Big)\big(1 + r\,\bar{a}_T\big) \;\geq\; \bar{C}_{T-1}\left[\sum_{s=0}^{4} R^{-s}R_\beta^{s+1} \;+\; R^{-5}\right].$$

The bracket is strictly positive, so it can be divided through.

$$\boxed{\bar{C}_{T-1} \;\leq\; \Phi\,\big(1 + r\,\bar{a}_T\big), \qquad \Phi \;\equiv\; \frac{\displaystyle\sum_{s=0}^{5} R^{-s}}{\displaystyle\sum_{s=0}^{4} R^{-s}R_\beta^{s+1} \;+\; R^{-5}}}$$

using only Lemma 1, the averaged budget constraint, and periodic stationarity. □

Combining with Theorem 1 gives the bound in the form used in the text,

$$\text{MPC-SB}(3m) \;\leq\; \Phi \cdot \frac{1 + r\,\bar{a}_T}{w^*} \sum_{h=0}^{2} \Big(R_\beta^{-(5-h)} - 1\Big),$$

where w^* is the bonus size.

A.10.3 How growth affects Theorem 1

Theorem 1 is stated for a stationary environment. Suppose instead consumption carries a deterministic trend — aggregate productivity growth, or the rising portion of the life-cycle profile — so that $c_t = G^t\hat{c}_t$ with $\hat{c}_t$ periodic-stationary and $G = 1 + f_g$ the gross monthly growth factor.

The Euler inequality holds regardless of the trend, and in detrended terms reads $\hat{c}_t^{-\sigma} \geq \beta R\,G^{-\sigma}\,E_t[\hat{c}_{t+1}^{-\sigma}]$, so Lemma 1 and Theorem 1 apply to $\hat{c}$ with R_β replaced by R_β/G: over five months the trend rises by G^5, giving the household more room for a temporary spike before it must fall back toward a rising path. Quantitatively:

$$(R_\beta/G)^{-5} - 1 \;\approx\; \big(R_\beta^{-5} - 1\big) \;+\; 5\,f_g,$$

The bound on the de-trended first difference rises by about $5\,f_g \times \bar{c}^{b}_{T-1}/w^*$.

However, the MPC-SB uses a difference-in-differences comparison, removing any common trends. Write the seasonal consumption gap between the groups at phase j as $\Delta_j = \hat{c}^b_j - \hat{c}^f_j$. Consumption in the bonus month is $G^t \hat{c}^g_T$, in the month before is $G^{t-1}\hat{c}^g_{T-1}$, and the bonus grows with income, $w^*_t = G^t w^*_0$. So:

$$\text{MPC-SB} \;=\; \frac{\left(G^t \hat{c}^b_T - G^{t-1}\hat{c}^b_{T-1}\right) - \left(G^t \hat{c}^f_T - G^{t-1}\hat{c}^f_{T-1}\right)}{G^t w^*_0} \;=\; \frac{G^t\Delta_T - G^{t-1}\Delta_{T-1}}{G^t w^*_0} \;=\; \underbrace{\frac{\Delta_T - \Delta_{T-1}}{w^*_0}}_{\text{no-growth MPC-SB}} + \frac{f_g}{1+f_g}\cdot\frac{\Delta_{T-1}}{w^*_0}.$$

Therefore, assuming the growth rate over time is the same across the two groups (which holds in the data), the growth factor mostly cancels. The only residual left is $f_g/(1+f_g)$ times Δ_{T-1}/w^*_0, the consumption gap between the two groups in the month before the bonus. If the consumption gap across groups in $T-1$ is small, which it is both in the data and the model (by the present value normalization), then this term is very small. Even if the gap is substantial, being multiplied by $f_g/(1+f_g)$, a very small number, makes it negligible.

A.10.4 How Theorem 1 changes with seasonal preferences

Let the per-period utility function be $\gamma_{j(t)}u(c_t)$ with $\gamma_j > 0$ a deterministic periodic taste-shifter, and write $q_j \equiv \gamma_j^{1/\sigma}$. The shifter is assumed to be common to both pay groups.

Define taste-adjusted consumption $\tilde{c}_t \equiv c_t/q_{j(t)}$. Then $\gamma_j\, u'(c_t) = \gamma_j\left(q_j\tilde{c}_t\right)^{-\sigma} = \tilde{c}_t^{-\sigma} = u'(\tilde{c}_t)$, so the household's Euler inequality now holds in terms of $\tilde{c}$, and Lemma 1 is:

$$\bar{\tilde{C}}_{t+1} \;\geq\; R_\beta\, \bar{\tilde{C}}_t, \qquad R_\beta = (\beta R)^{1/\sigma}.$$

Since γ_j is deterministic it factors out of the cross-sectional mean, so $\bar{\tilde{C}}_{T+s} = \bar{C}_{T+s}/q_{j(T+s)}$.

Recall that by definition:

$$\text{MPC-SB}(T) = \frac{\left(\bar{C}^{\text{b}}_T - \bar{C}^{\text{b}}_{T-1}\right) - \left(\bar{C}^{\text{f}}_T - \bar{C}^{\text{f}}_{T-1}\right)}{w^*}$$

We can bound the first term in a very similar way to above.

By Lemma 1, $\tilde{\bar{C}}^b_{T+5} \geq R^5_\beta \tilde{\bar{C}}^b_T$, implying $\bar{C}^{\mathrm{b}}_{T+5} \geq R^5_\beta \frac{q_{j(T+5)}}{q_{j(T)}} \bar{C}^{\mathrm{b}}_T$.

Now, if the taste shifter is on a six month cycle too ($j \in 1, ..., 6$) then periodic stationarity over the bonus cycle still holds, so $\bar{C}^b_{T+5} = \bar{C}^b_{T-1}$. The bonus group need not follow the deterministic cycle of the flat pay group, indeed in general it will not, but it does still follow a regular six month cycle.

Combining these, $\bar{C}^{\mathrm{b}}_{T-1} \geq R^5_\beta \frac{q_{j(T+5)}}{q_{j(T)}} \bar{C}^{\mathrm{b}}_T$, and inverting this, $\bar{C}^{\mathrm{b}}_T \leq R^{-5}_\beta \frac{q_{j(T)}}{q_{j(T+5)}} \bar{C}^{\mathrm{b}}_{T-1}$.

Therefore, $\bar{C}^{\mathrm{b}}_T - \bar{C}^{\mathrm{b}}_{T-1} \leq (R^{-5}_\beta \frac{q_{j(T)}}{q_{j(T+5)}} - 1)\bar{C}^b_{T-1}$.

We do not directly observe the ratio $\frac{q_{j(T)}}{q_{j(T+5)}}$, but we can apply Lemma 1 to the flat pay group: $\tilde{\bar{C}}^f_T \geq R_\beta \tilde{\bar{C}}^f_{T-1}$, which implies $\bar{C}^{\mathrm{f}}_T \geq R_\beta \frac{q_{j(T)}}{q_{j(T-1)}} \bar{C}^{\mathrm{f}}_{T-1}$ and then $R^{-1}_\beta \frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} \geq \frac{q_{j(T)}}{q_{j(T-1)}}$. In the six month preference cycle, $q_{j(T+5)} = q_{j(T-1)}$, so $R^{-1}_\beta \frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} \geq \frac{q_{j(T)}}{q_{j(T+5)}}$.

So, finally: $\bar{C}^{\mathrm{b}}_T - \bar{C}^{\mathrm{b}}_{T-1} \leq (R^{-6}_\beta \frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} - 1)\bar{C}^b_{T-1}$

The second term $\bar{C}^{\mathrm{f}}_T - \bar{C}^{\mathrm{f}}_{T-1}$ is no longer zero. Write

$$\bar{C}^{\mathrm{f}}_T - \bar{C}^{\mathrm{f}}_{T-1} = (\frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} - 1)\bar{C}^{\mathrm{f}}_{T-1}$$

Then the whole bound is

$$\text{MPC-SB}(T) \leq \frac{\bar{C}^b_{T-1}}{w^*}(R^{-6}_\beta \frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} - 1) - \frac{\bar{C}^f_{T-1}}{w^*}(\frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} - 1)$$

We can then exploit the fact that $\bar{C}^{\mathrm{f}}_{T-1} \geq \bar{C}^{\mathrm{b}}_{T-1}$ (the bonus group runs its seasonal reserve down through the ordinary months and so holds least cash in hand in the month before the bonus, whereas the flat group's resources are the same every month) implying

$$\text{MPC-SB}(T) \leq \frac{\bar{C}^b_{T-1}}{w^*}(R^{-6}_\beta \frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} - 1) - \frac{\bar{C}^b_{T-1}}{w^*}(\frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}} - 1)$$

Which simplifies to

$$\text{MPC-SB}(T) \leq \frac{\bar{C}^b_{T-1}}{w^*}\frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}}(R^{-6}_\beta - 1)$$

This is looser than the original bound for two reasons: there is an extra power of R^{-1}_β within the brackets, and there is the multiplier $\frac{\bar{C}^{\mathrm{f}}_T}{\bar{C}^{\mathrm{f}}_{T-1}}$. In the data, control group spending is 4% higher in June

than May, and 6% higher in December than November, so this is a number in the range 1.04 to 1.06. Together, these loosen the bound modestly.

If instead preferences follow a twelve month cycle ($j \in 1, ..., 12$), the result is somewhat weaker. The crucial stage which forces bonus group consumption to return to where it started after six months can no longer be used ($\bar{C}^b_{T+5} = \bar{C}^b_{T-1}$), and instead one can use $\bar{C}^b_{T+11} = \bar{C}^b_{T-1}$, which leads to a significant loosening:

$$\text{MPC-SB}(T) \leq \frac{\bar{C}^b_{T-1}}{w^*} \frac{\bar{C}^{\text{f}}_T}{\bar{C}^{\text{f}}_{T-1}} (R_\beta^{-12} - 1)$$

We run simulations using a full twelve month seasonal preference cycle, calibrated to the seasonality of the non-bonus group. It makes no material difference to the model predictions; in fact for any given β it produces lower MPC-SBs, because the seasonal preference cycle forces households to save more, making them less constrained.

A.10.5 MPC-SB bounds for alternative sigma values

Table A.16: Upper bound on the MPC-SB at different discount rates, sigma = 1

β (annual)	Φ	MPC-SB(1m), A/Y = 10	MPC-SB(3m), A/Y = 10	MPC-SB(1m), A/Y = 60	MPC-SB(3m), A/Y = 60
0.975	1.001	0.003	0.007	0.003	0.007
0.950	1.007	0.016	0.038	0.017	0.041
0.900	1.018	0.044	0.105	0.047	0.113
0.850	1.030	0.075	0.179	0.081	0.193
0.800	1.043	0.109	0.261	0.118	0.282
0.750	1.057	0.148	0.352	0.160	0.381
0.700	1.071	0.192	0.454	0.207	0.491
0.650	1.088	0.241	0.570	0.261	0.616
0.600	1.105	0.298	0.702	0.322	0.759
0.550	1.124	0.363	0.854	0.393	0.923
0.500	1.145	0.440	1.031	0.476	1.115

Table A.17: Upper bound on the MPC-SB at different discount rates, sigma = 10

β (annual)	Φ	MPC-SB(1m), A/Y = 10	MPC-SB(3m), A/Y = 10	MPC-SB(1m), A/Y = 60	MPC-SB(3m), A/Y = 60
0.975	1.000	0.000	0.001	0.000	0.001
0.950	1.001	0.002	0.004	0.002	0.004
0.900	1.002	0.004	0.010	0.005	0.011
0.850	1.003	0.007	0.017	0.008	0.018
0.800	1.004	0.010	0.024	0.011	0.026
0.750	1.006	0.013	0.032	0.014	0.035
0.700	1.007	0.017	0.040	0.018	0.044
0.650	1.009	0.021	0.050	0.022	0.054
0.600	1.010	0.025	0.059	0.027	0.064
0.550	1.012	0.029	0.070	0.032	0.076
0.500	1.014	0.034	0.082	0.037	0.089

A.11 Appendix: Additional Simulation Results

A.11.1 Model results from two buffer stock simulations

Table A.18: Buffer Stock Model Simulation Results

Statistic	Data	Wealth / income target	MPC-IS target
MPCs			
MPC-SB (3 months)	0.21	0.01	0.05
MPC-IS (3 months)	0.29	0.14	0.29
MPC-IS, 1m / 3m		0.35	0.36
Hand-to-mouth			
Share of MPC-SB from HtM	55%	100%	100%
Share of MPC-IS from HtM	46%	75%	95%
Households with < 1 month	30% / 29%	40%	84%
Calibration			
Wealth / income	9.8	10.0	1.6
Parameters		$\beta = 0.88$	$\beta = 0.80$

Note:
Share of MPC from HtM means the share of the total MPC accounted for by households with assets < 1 months of income. Data for households with < 1 months is for the seasonal bonus sample / lottery sample in our bank data. In the Spanish Survey of Household Finances the share of households with any earnings who hold deposits less than 1 times their monthly earnings is 29.8 percent (falling to 27.1 percent for all liquid assets). Data for wealth / income is mean deposits divided by mean salary income in our bank data, for the seasonal pay group. In the Spanish Survey of Household Finances the ratio of mean deposits to mean earnings is 10.9.

A.11.2 Simulation plots using median income

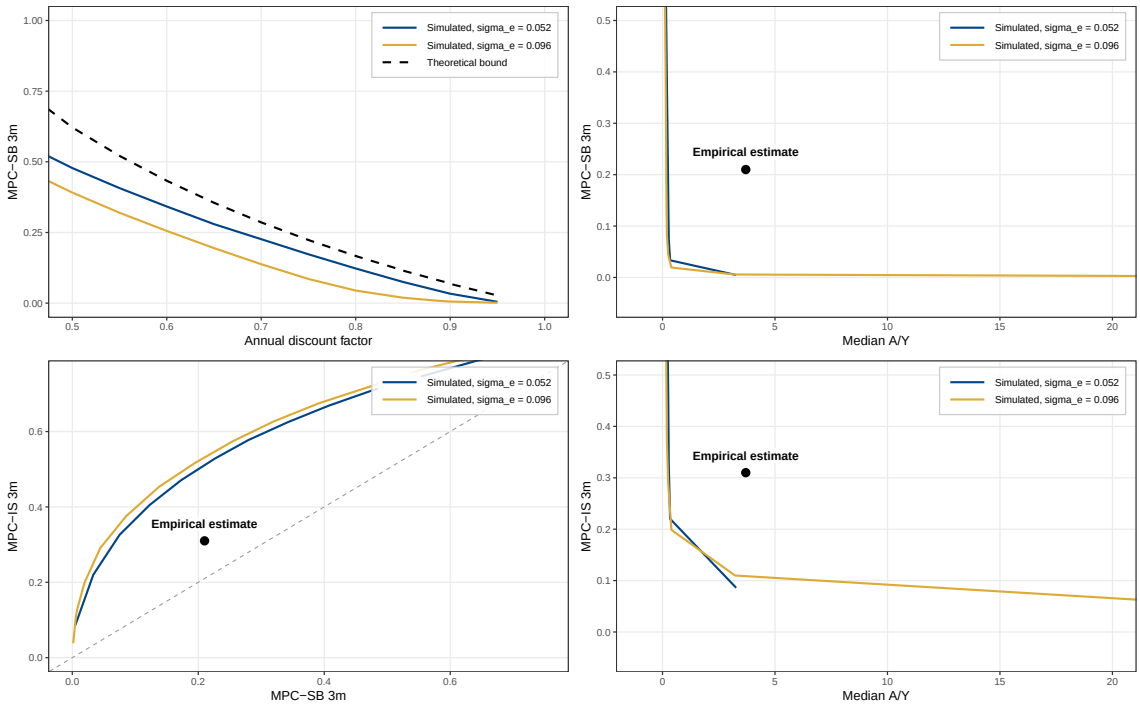


Figure A.20: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with high and low income variance.

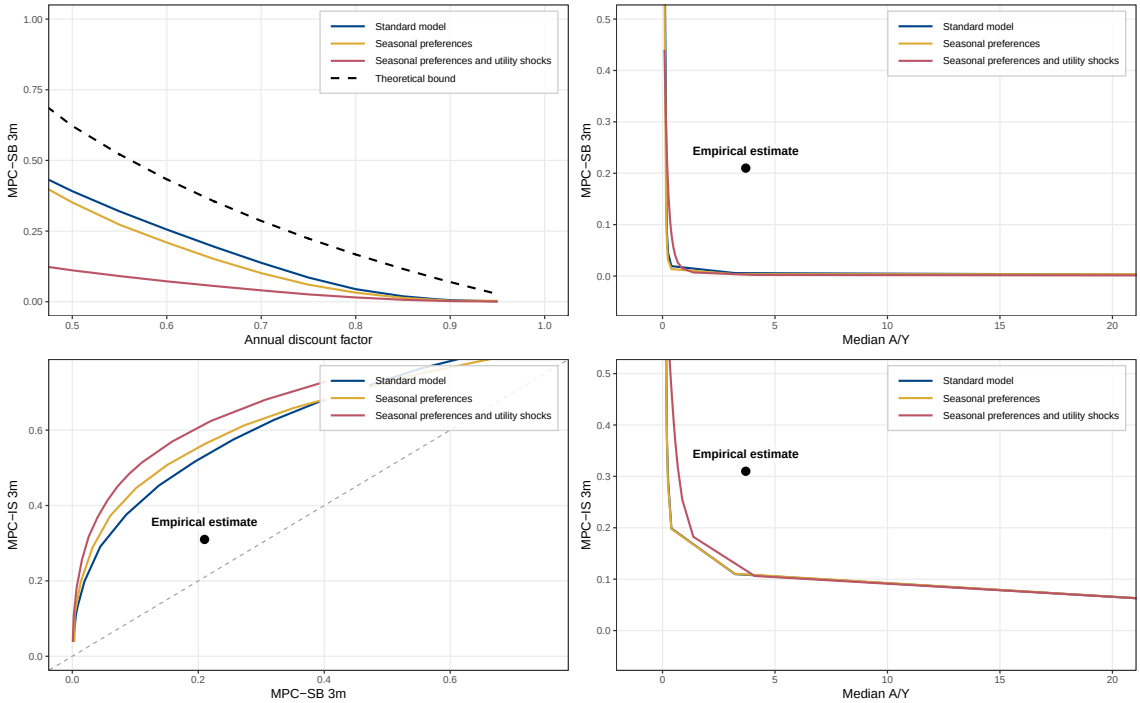


Figure A.21: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with seasonal preferences and utility shocks added.

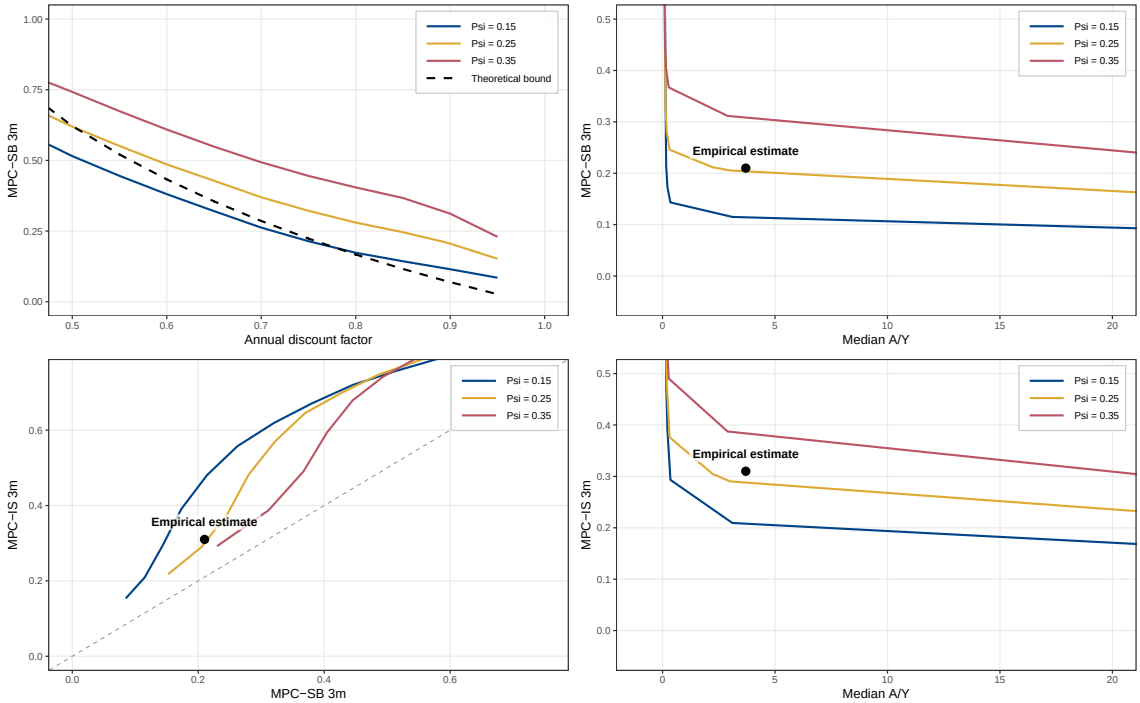


Figure A.22: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with mental accounting at different values of psi.

A.11.3 Simulation with GKOS income process

The income process in our baseline simulations is a simple Gaussian AR(1). Guvenen et al. (2021), henceforth GKOS, use Social Security records to show that earnings changes are heavily fat-tailed and negatively skewed: most workers' earnings barely move from one year to the next, but a minority experience very large, mostly negative, jumps (job loss, business failure, a health shock). A Gaussian AR(1) cannot reproduce this pattern: its innovations are symmetric, so risk can only be raised symmetrically for all households. The income process has been shown to matter for MPCs in buffer-stock models (**commault2025heterogeneity**), so we examine what adopting a GKOS-style process does in our application. We use the persistent-plus-transitory specification of their Model 5. Pre-bonus income is $y_t = w\,\kappa_j\,z_t$ with

$$\log z_t = p_t + \mathrm{tr}_t - c,$$

where c is a constant that normalises $E[z_t] = 1$ (as for the AR(1) baseline). The persistent component is an AR(1),

$$p_{t+1} = \rho\, p_t + \eta_{t+1},$$

whose innovation is a mixture of a rare draw (component 1) and a normal shock (component 2):

$$\eta \sim \pi_p\, N\left(\mu_1^p, (\sigma_1^p)^2\right) + (1 - \pi_p)\, N\left(\mu_2^p, (\sigma_2^p)^2\right).$$

The rare draw has small probability, a large negative mean and a standard deviation well above that of the normal shock; the normal shock has a slightly positive mean, which offsets the rare draw and produces the negative skew. The transitory component is i.i.d. over time and is also a Gaussian mixture,

$$\mathrm{tr}_t \sim \pi_{tr}\, N\left(\mu_1^{tr}, (\sigma_1^{tr})^2\right) + (1 - \pi_{tr})\, N\left(\mu_2^{tr}, (\sigma_2^{tr})^2\right),$$

whose rare component has a very large variance and a slightly positive mean. The two components are independent.

We calibrate the monthly process to preserve the following features of the GKOS (Table IV, Model 5) estimates: - The rare-event probabilities. We divide the annual probabilities by 12, so the expected number of rare events per year is unchanged. - The shape of each mixture: we keep the ratios $\mu_1 : \sigma_1 : \sigma_2$ within each component. - Mean zero: μ_2^p and μ_2^{tr} are chosen so that η and tr each have mean zero.

We then match two moments to our baseline AR(1): - Persistence: $\rho = 0.9922$ monthly (0.91 annual), the same as the baseline. This is deliberately not GKOS's own persistence, which is higher. - Overall dispersion: a common factor is applied to $\mu_1, \sigma_1, \sigma_2$ of both components so that the stationary standard deviation of $\log z$ equals 0.77, the raw cross-sectional dispersion of log income in our data and the same total variance as the baseline AR(1).

Holding persistence and total variance fixed, the calibration isolates the effect of the shape of the

income process. Note that the resulting shocks are somewhat larger than GKOS's own estimates (the rare persistent draw is −0.63 in logs rather than −0.52; unrescaled, their estimates imply $sd(\log z) \approx 0.61$), so this is a GKOS-shaped process at our level of risk rather than their estimated process.

Our calibration parameters for the monthly income process are:

	π	μ_1	σ_1	μ_2	σ_2
Persistent innovation, η	0.01467	−0.6266	0.1351	0.0093	0.0550
Transitory shock, tr	0.00367	0.1695	0.9639	−0.0006	0.0696

Figure 14 shows the results from simulating the model with this income process across the range of discount factors.

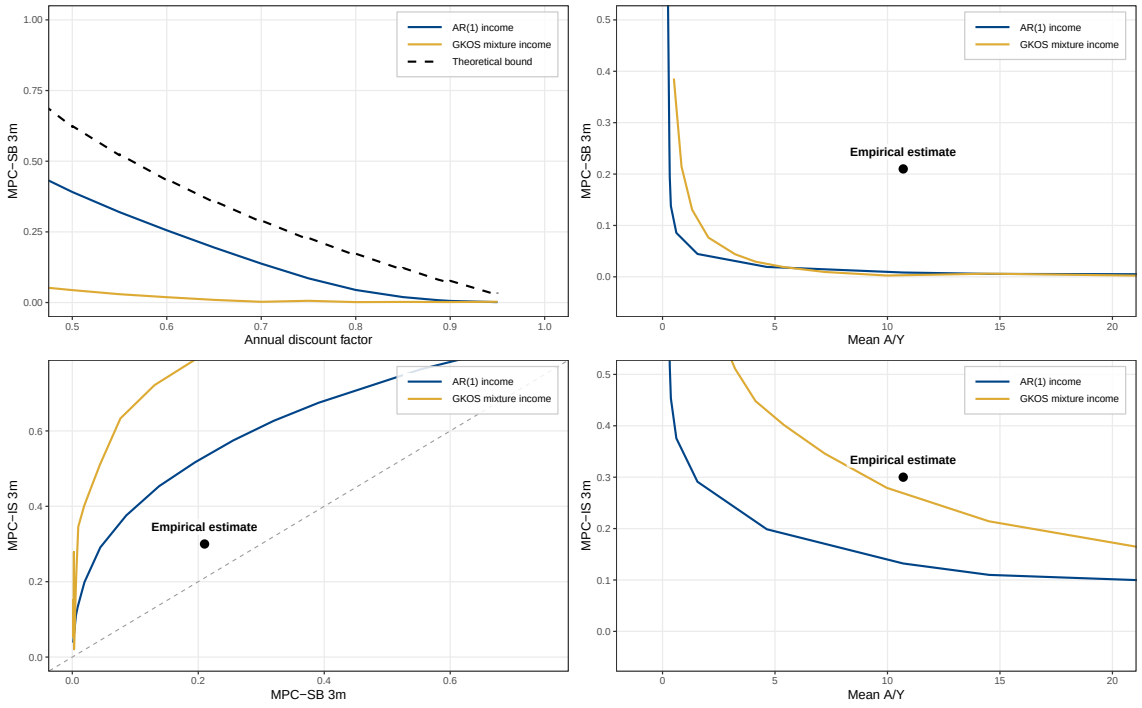


Figure A.23: Model statistics from simulations over a range of discount factors from 0.1 to 0.975, with AR(1) and Guvenen et al. (2021) income processes.

Adopting the GKOS income leads to much more wealth accumulation, and consequently much lower MPC-SBs, at each given β. It also raises the MPC-IS very significantly at any given level of assets. The MPC-SB remains very small except at very low levels of wealth, and the ratio of the MPC-SB to MPC-IS shifts much further away from its empirical value. Therefore, if anything adopting a more realistic income process pushes the model predictions even further from our findings.

A.11.4 Full results table for heterogeneous discounting, present bias, and mental accounting

Here we show the full results table for heterogeneous discounting, present bias, and mental accounting, showing how the results vary as the tuning parameters vary.

Table A.20: All model runs, each calibrated to a seasonal-pay asset-income ratio of 10 months.

		Assets, months (share)		MPC-SB, 3m					MPC-IS, 3m					
Parameter	Annual β	<0.5	≥10	Agg.	Q1	Q5	<0.5	≥10	Agg.	1m/3m	Q1	Q5	<0.5	≥10
Exponential baseline: no behavioural friction														
none	0.884	41%	24%	0.01	0.12	0.00	0.07	0.00	0.14	0.35	0.33	0.04	0.26	0.04
Mental accounting: strength ψ of the dissaving penalty														
$\psi = 0.05$	0.885	40%	24%	0.04	0.15	0.03	0.10	0.03	0.16	0.49	0.35	0.06	0.28	0.07
$\psi = 0.10$	0.886	40%	24%	0.08	0.20	0.06	0.14	0.07	0.19	0.60	0.35	0.10	0.30	0.10
$\psi = 0.15$	0.888	39%	25%	0.12	0.24	0.10	0.19	0.10	0.23	0.68	0.39	0.13	0.32	0.13
$\psi = 0.20$	0.889	39%	25%	0.17	0.28	0.14	0.23	0.14	0.26	0.74	0.43	0.16	0.36	0.16
$\psi = 0.25$	**0.891**	**38%**	**25%**	**0.21**	**0.34**	**0.18**	**0.29**	**0.19**	**0.30**	**0.80**	**0.43**	**0.20**	**0.40**	**0.20**
$\psi = 0.30$	0.892	37%	26%	0.27	0.39	0.23	0.35	0.23	0.35	0.84	0.48	0.24	0.44	0.24
$\psi = 0.40$	0.895	35%	26%	0.38	0.52	0.32	0.48	0.33	0.45	0.90	0.59	0.33	0.52	0.33
Naive present bias: extra one-month discount β_0 on the future														
$\beta_0 = 0.99$	0.886	42%	24%	0.01	0.17	0.00	0.10	0.00	0.14	0.35	0.34	0.04	0.27	0.04
$\beta_0 = 0.95$	0.893	46%	23%	0.05	0.36	0.00	0.28	0.00	0.17	0.35	0.36	0.04	0.32	0.04
$\beta_0 = 0.90$	0.901	50%	23%	0.10	0.59	0.00	0.49	0.00	0.22	0.35	0.45	0.03	0.40	0.04
$\beta_0 = 0.85$	0.909	55%	22%	0.16	0.71	0.00	0.54	0.00	0.28	0.35	0.54	0.03	0.47	0.03
$\beta_0 = 0.80$	**0.916**	**57%**	**21%**	**0.21**	**0.79**	**0.00**	**0.65**	**0.00**	**0.34**	**0.36**	**0.59**	**0.03**	**0.56**	**0.03**
$\beta_0 = 0.75$	0.924	60%	21%	0.26	0.84	0.00	0.77	0.00	0.41	0.36	0.69	0.03	0.66	0.03
$\beta_0 = 0.70$	0.931	62%	20%	0.30	0.89	0.00	0.84	0.00	0.48	0.37	0.77	0.03	0.74	0.03
Heterogeneous discount factors: 50/50 split of patient and impatient households														
$\{\beta_L, \beta_H\} = \{0.85, 0.902\}$		45%	23%	0.01	0.13	0.00	0.07	0.00	0.15	0.35	0.34	0.04	0.27	0.04
$\{\beta_L, \beta_H\} = \{0.80, 0.910\}$		54%	21%	0.02	0.17	0.00	0.10	0.00	0.19	0.36	0.37	0.03	0.31	0.03
$\{\beta_L, \beta_H\} = \{0.75, 0.912\}$		59%	20%	0.05	0.22	0.00	0.15	0.00	0.23	0.36	0.40	0.03	0.36	0.03
$\{\beta_L, \beta_H\} = \{0.70, 0.912\}$		61%	20%	0.07	0.29	0.00	0.20	0.00	0.27	0.37	0.45	0.03	0.41	0.03
$\{\beta_L, \beta_H\} = \{0.65, 0.912\}$		62%	20%	0.10	0.35	0.00	0.27	0.00	0.30	0.38	0.50	0.03	0.46	0.03
$\{\beta_L, \beta_H\} = \{0.60, 0.912\}$		62%	20%	0.13	0.42	0.00	0.35	0.00	0.33	0.38	0.55	0.03	0.51	0.03
$\{\beta_L, \beta_H\} = \{0.55, 0.912\}$		62%	20%	0.16	0.49	0.00	0.42	0.00	0.36	0.39	0.59	0.03	0.55	0.03
$\{\beta_L, \beta_H\} = \{0.50, 0.912\}$		**62%**	**20%**	**0.20**	**0.56**	**0.00**	**0.49**	**0.00**	**0.38**	**0.39**	**0.63**	**0.03**	**0.59**	**0.03**
$\{\beta_L, \beta_H\} = \{0.40, 0.912\}$		63%	20%	0.28	0.73	0.00	0.63	0.00	0.43	0.41	0.71	0.03	0.66	0.03
$\{\beta_L, \beta_H\} = \{0.30, 0.913\}$		63%	20%	0.37	0.84	0.00	0.72	0.00	0.47	0.43	0.78	0.03	0.72	0.03
$\{\beta_L, \beta_H\} = \{0.20, 0.913\}$		63%	20%	0.44	0.88	0.00	0.81	0.00	0.50	0.46	0.85	0.03	0.78	0.03

Notes: Each run is calibrated to a seasonal-pay A/Y between 9.9 and 10.1 months by choosing the annual discount factor. Bold: in each model class, the run whose aggregate MPC-SB is closest to 0.21. Assets are in months of average monthly income; the first two columns are the shares of seasonal-pay households with assets below 0.5 and at least 10 months, and the last two columns of each MPC block are the MPC of households in those ranges.

A.11.5 Results for Stone-Geary Seasonal Preferences

A.12 Appendix: Browning and Collado (2001)

One previous paper has studied the same Spanish seasonal bonus setting as us: Browning and Collado (2001). They find the opposite result to us: they find no evidence that spending patterns are affected by receipt of seasonal bonuses. In this section, we try to reconcile our findings with theirs. Our overall conclusion is that we believe two factors mostly explain the difference: lack of power, and some misclassification of bonus status.

To compare our results, we accesed the same public microdata from the Encuesta Continua de Presupuestos Familiares (ECPF) for 1985–1995. We were able to closely replicate the results in the paper. With this in hand, we identified two potential reasons for the discrepancy between our findings and theirs.

Most importantly, their dataset is small: 1874 bonus households and 467 non-bonus households. Making assumptions about how our estimated spending response is distributed over the weeks within each month, we estimate that - under the alternative hypothesis that our MPC of 0.21 is true - their main test on total expenditure would correctly reject the null only 56% of the time.

A second potential explanation relates to categorization into bonus and non-bonus households. Comparing the bonus and non-bonus groups formed using their classification method, we find some evidence of misclassification. Firstly, bonus group pay does not increase as much as we would expect in the bonus quarters, and there is also some increase in the non-bonus group in these quarters. Secondly, there are some cases where bonus households either experience only small increases in a bonus quarter, or experience large increases in non-bonus quarters. Thirdly, the survey used to collect the data asked respondents to report their earnings over the previous three months, with the potential for recall errors and misreporting.

A.13 Appendix: Alternate Versions of the Mental Accounting Model

A.14 Appendix: Mental Accounting in a Baseline HANK Model

In this appendix we take the one-asset heterogeneous agent New Keynesian (HANK) model of Auclert, Rognlie, and Straub (2025), and test how its policy implications change when we add mental

accounting to the consumption block. Auclert, Rognlie, and Straub (2025) provide a lot of useful materials on their Github (https://github.com/shade-econ/annual-review) and codes for solving all of their models, making this exercise possible. The mental accounting model simulations are not intended as serious quantitative estimates for policy purposes - the goal here is to highlight some of the potential policy implications of mental accounting using an existing baseline model as a laboratory.

We leave most features of the baseline model unchanged, changing only the household block. Time is quarterly. We add mental accounting to the utility function exactly as above: marginal utility falls by a factor $1 - \psi$ when consumption exceeds current (after-tax) labour income.[26] Adding mental accounting changes the MPC and the amount of wealth accumulation, so we re-calibrate the model.

We compare three versions of the heterogeneous agent model. The first (HA) is the baseline model with no mental accounting, using the calibration of Auclert, Rognlie, and Straub (2025): households have one of two discount factors which they randomly switch between with some small probability each period. The two discount factors and the share of households with each are jointly chosen to match wealth of 20 times quarterly output, an impact (first quarter) MPC of 0.20 out of an income shock proportional to income, and the Lorenz curve of total wealth in the 2019 SCF.

Our first alternative calibration (MA1) exploits the fact that a model with mental accounting can simultaneously match total wealth of 20 times quarterly income and an MPC of 0.20 without heterogeneous discounting. We choose ψ and a single discount factor to jointly match the wealth and MPC target. This requires a large ψ of 0.44. This is more a proof of concept than necessarily a realistic calibration.

We also test a mental accounting model which retains the heterogeneous discounting of the baseline model (MA2). This model sets $\psi = 0.25$, the value we chose above to match our seasonal bonus MPCs. We then follow the same calibration strategy as Auclert, Rognlie, and Straub (2025), choosing both discount factors and the share of impatient households to match the same assets, MPC, and Lorenz curve targets. With $\psi = 0.25$, the model is unable to hit all targets simultaneously, so we match the first two and choose the calibration which gets closest to hitting the Lorenz curve.[27] The

[26] We solve the household problem with the endogenous grid method on 1,600 asset points. For dynamics we replace the kink by a smooth transition within 0.02% of income.

[27] Closer fits exist, but only as patient households approach the point at which their wealth would grow

details of each calibration are in the table below.

	HA (baseline)	MA1 (one β)	MA2 (two β)
β (annual)	0.69 / 0.995	0.966	0.90 / 0.998
Patient share	0.49	–	0.35
ψ	0	0.44	0.25
Targets	Wealth, MPC, Lorenz	Wealth, MPC	Wealth, MPC, Lorenz
Impact MPC (proportional)	0.20	0.20	0.20
Impact MPC (flat, $\bar{y}$ (1m))	0.18	0.46	0.35

The MA2 calibration can simultaneously match MPCs and wealth without a very impatient group: the share of impatient households rises from 0.51 to 0.65, but their discount factor rises from 0.69 to 0.90. The last row of the table calculates the MPC out of a flat windfall where all households receive 1 month's worth of average income - an experiment comparable to our headline lottery experiment with prize-size matched to seasonal bonuses. This differs from the proportional MPC which the models are calibrated to in two ways: firstly it is a larger transfer, and secondly, it is of equal amount to all households rather than being proportional to income. For the baseline model (HA), these changes reduce the MPC. Whilst shifting to a flat (non-proportional) windfall increases the MPC, because hand-to-mouth households get a larger proportion of the total money transferred, the increase in the size of the transfer reduces MPCs by more. In the mental accounting models, MPCs increase with transfer size over some range: whilst a small transfer may not move many households close to their reference point, a larger one does so for more of them.

A.14.1 Spending is concentrated when income arrives

Our first policy experiment is the tax cut considered by Auclert, Rognlie, and Straub (2025): the government cuts taxes by 1% of quarterly output. This tax cut decays by a factor of 0.9 per quarter and is financed by issuing debt. Debt subsequently mean-reverts, requiring higher taxes in the future. The taxes in the model are proportional taxes on labour income and dividends, so the tax cut raises each household's income in proportion to their labour income (as in the experiment to which MPCs

without bound; we stop where they are as close to that point as in the baseline calibration.

are calibrated). The model is solved for the output path that clears the asset market. The table below reports the first-year multiplier: the cumulative additional output over the first four quarters, divided by the cumulative tax cut over the same period (which is 3.44% of quarterly output).

As we discussed in the main text, mental accounting models make MPCs much more front-loaded: most of the consumption response to a transfer occurs within the three months after receipt. This shape is much more consistent with our evidence, and with other recent papers (for example Boehm, Fize, and Jaravel (2025), Graham and McDowall (2026)). The next table shows this quantitatively for our model variants. In the baseline HA model, the MPC in the quarter after impact is 46% of its impact value. For the mental accounting variants MA1 and MA2, it is 9% and 19% respectively. For comparison, we also show the results from the two-agent New-Keynesian model (TANK) which is also discussed by Auclert, Rognlie, and Straub (2025).

Auclert, Rognlie, and Straub (2024) show that it is the entire path of iMPCs which matters for the fiscal multiplier. Since each model is calibrated to the same impact MPC, the models where iMPCs fall faster (MA and TANK) have lower transfer multipliers: about 50% lower for MA1 and 20% lower for MA2. In this respect, the mental accounting models lie somewhere between the HANK and the TANK models.

	HA	MA1	MA2	TANK
Impact MPC	0.20	0.20	0.20	0.20
MPC one quarter later / impact MPC	0.46	0.09	0.19	0.02
Fiscal multiplier (first year)	0.60	0.32	0.47	0.20
Multiplier / impact MPC	3.0	1.6	2.4	1.0
Monetary easing: output gain (%)	2.18	2.00	2.15	1.90
Seasonal MPC (data: 0.21)	0.04	0.26	0.28	-

We also estimated the impact of mental accounting on the monetary policy experiment of Auclert, Rognlie, and Straub (2025): the central bank cuts the (annual) interest rate by one percentage point, with the cut then reverting by a factor of 0.9 per quarter. The government passes the lower interest bill onto households through lower taxes. This monetary expansion therefore affects households through a number of channels: a direct effect from lower returns on wealth now and into the near

future, a capital gain at date 0 as the stock market is revalued at a lower discount rate, and higher labour income via both higher output from the direct effect and lower taxes from the government. The table reports the average output response to the rate cut over the first four quarters, as a percentage of steady state output.

Adding mental accounting alters a number of these transmission channels. The labour income channel is weakened, for the same reason that the fiscal multiplier falls: the iMPCs are lower beyond the first quarter. The direct impact of the rate cut, on the other hand, is higher. But this is not because of mental accounting per se, it is because calibrating the mental accounting model yields much higher discount factors, and more patient households are more responsive to changes in interest rates. For a fixed discount factor, mental accounting would make households less responsive to the direct effects of monetary policy, because households at their reference point do not respond to changes in interest rates over some range. Overall, the impact of the 1 percentage point rate cut is similar in the mental accounting models as in the HA baseline.

A.14.2 Income flows versus balance sheets

The next policy experiment departs from the applications in Auclert, Rognlie, and Straub (2025) and focuses on comparing the impact of policies targeting income flows with policies targeting balance sheet quantities. In standard models, income and assets are fungible, and policies targeting one or the other only have different effects because of borrowing constraints. But with mental accounting, policies targeting income flows may shift the reference point, therefore having a greater impact. We demonstrate this by comparing the impact of a debt-relief policy with a payment-relief policy of equivalent cost (in present value terms). This model has a borrowing constraint at zero, so there are no net-debtor households. We therefore simulate debt relief as an addition to the stock of assets. Payment relief is modelled as a transfer, spread over 4 quarters.[28] We assume that debt payments are deducted from the income households budget out of, so the payment relief increases households' reference point for the 4 quarters.

[28] In our experiment, payment relief is equal to 1% of average quarterly income to each household for 4 quarters, and the principal relief is the present value equivalent added to the asset stock in the first quarter. Both are financed by the government issuing debt which is repaid gradually through income taxes in the same way as in the tax cut experiment.

The table below shows the impact of these two policies as the additional dollars of spending per (present-value) dollars of principal or payment relief. In the baseline model (HA), these two policies have essentially the same impact - principal relief is slightly more stimulating because it is front-loaded. In both mental accounting models, on the other hand, payment relief stimulates spending by more than twice as much as principal relief.[29]

	HA	MA1	MA2
Debt relief: first-year spending per dollar of present value			
Payment relief	0.44	0.27	0.42
Principal relief	0.47	0.10	0.21
Debt relief with GE effects: first-year output per dollar			
Payment relief	0.63	0.31	0.54
Principal relief	0.64	0.08	0.23
Looser borrowing limit: spending on impact per dollar			
MPC from credit	0.24	0.02	0.06
MPC from income	0.29	0.19	0.27

Finally, we investigate the impact on consumption of relaxing the borrowing constraint. In the baseline (HA) model, borrowing constraints are the core driver of high MPCs. Extending credit therefore stimulates a lot of additional spending upon impact. We model this as a change in the borrowing limit: from zero, it is relaxed to allow all households to borrow up to 1% of average quarterly income at the same interest rate they can save at.[30] In the baseline HA model, households spend 0.24 dollars per dollar of credit extended. In the mental accounting models, this policy is much less effective: stimulating only 0.02 or 0.06 dollars of additional spending for MA1 and MA2 respectively. The households in the mental accounting model do not have high MPCs primarily

[29]Principal relief is less stimulative in both mental accounting models largely because the households in these calibrations are more patient. But the gap between the efficacy of principal and payment relief is driven by the mental accounting mechanism. In MA1 the general-equilibrium effect is below the direct effect because the small initial impulse, spent mostly on impact, generates little feedback through incomes, while the taxes that repay the debt still reduce spending.

[30]To be clear: this is 1% of all households' average income, so the borrowing limit extension is the same in dollar terms for all households.

because they are constrained, it is rather because of their dis-saving aversion. So they do not spend heavily out of extended credit.

The final row of the table gives the MPC when the same 1% of average quarterly income is given to all households as income.[31] In the baseline HA calibration, the MPC out of credit is very similar to the MPC out of income: 0.24 versus 0.29. In the mental accounting models, on the other hand, the MPC out of credit is far lower than the MPC out of income: 0.02/0.06 compared with 0.19/0.27.

A.14.3 What are the implications of these exercises?

The core point is that the mechanism through which we generate high MPCs in our model can alter its policy implications. Mental accounting models, calibrated in the same manner as the baseline HA model, predict lower fiscal multipliers, and smaller responses to principal relief and to credit, relative to payment relief or cash transfers. Our calibrations here are not designed to give a serious quantitative assessment of the impact of different policies, but instead to show the broad qualitative implications of the mental accounting mechanism for high MPCs. An important avenue for future work is to figure out how to build serious quantitative models for policy with some sort of 'spend on arrival' mechanism, like mental accounting, so that they can reflect the type of behaviour this paper documents empirically.

[31] These MPCs are not the 0.20 the models were calibrated to because this experiment gives an equal amount to all households of 1% of average quarterly income, rather than giving each household 1% of their own income.