American Economic Journal: Macroeconomics 2021, 13(4): 1–54
https://doi.org/10.1257/mac.20190211

MPC Heterogeneity and Household Balance Sheets[†]

By Andreas Fagereng, Martin B. Holm, and Gisle J. Natvik*

We use sizable lottery prizes in Norwegian administrative panel data to explore how transitory income shocks are spent and saved over time and how households' marginal propensities to consume (MPCs) vary with household characteristics and shock size. We find that spending peaks in the year of winning and gradually reverts to normal within five years. Controlling for all items on households' balance sheets and characteristics such as education and income, it is the amount won, age, and liquid assets that vary systematically with MPCs. Low-liquidity winners of the smallest prizes (around US$1,500) are estimated to spend all within the year of winning. The corresponding estimate for high-liquidity winners of large prizes (US$8,300–150,000) is slightly below one-half. While conventional models will struggle to account for such high MPC levels, we show that a two-asset life cycle model with a realistic earnings profile and a luxury bequest motive can account for both the time profile of consumption responses and their systematic covariation with observables. (*JEL* D12, D15, E21, G51, H24)

What determines how households adjust their expenditure and saving to transitory income shocks? This question is fundamental to several branches of economics. In particular, a growing literature articulates how statistics regarding the heterogeneity and dynamics of households' consumption responses to windfall income are essential to address long-standing macroeconomic questions about shock propagation and economic policy.[1] In this paper, we contribute by providing such statistics. Utilizing observed lottery prizes, we characterize (i) how transitory income shocks feed into

*Fagereng: BI Norwegian Business School (email: afagereng@gmail.com). Also affiliated with Statistics Norway; Holm: University of Oslo (email: martin.b.holm@outlook.com); Natvik: BI Norwegian Business School (email: gisle.j.natvik@bi.no). Virgiliu Midrigan was coeditor for this article. We thank three anonymous referees, Mario Alloza, Adrien Auclert, Paco Buera, Ådne Cappelen, Christopher Carroll, Thomas Crossley, Nicola Fuchs-Schündeln, Robert Hall, Elin Halvorsen, Steinar Holden, Ragnar Juelsrud, Greg Kaplan, Søren Leth-Petersen, Edwin Leuven, Davide Malacrino, Magne Mogstad, Benjamin Moll, Luigi Pistaferri, Martin Schneider, Terje Skjerpen, Kjetil Storesletten, Kjersti N. Torstensen, Nate Vellekoop, Gianluca Violante, and Matthew White, as well as participants at various seminars, workshops, and conferences, for helpful advice, discussions, and suggestions. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 851891) and the Norwegian Research Council (grant agreement No. 287720).

[†]Go to https://doi.org/10.1257/mac.20190211 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

[1] See, for instance, Auclert (2019); Auclert, Rognlie, and Straub (2018); Berger et al. (2018); and Kaplan, Moll, and Violante (2018).

Copyright American Economic Association; reproduced with permission.

consumption expenditure and savings over time, (ii) which household characteristics are systematically related to the magnitude of these responses, and (iii) how the responsiveness of consumption expenditure varies with the size of income shocks. Finally, we contrast these findings to the predictions from a two-asset life cycle model of the type that is widely applied in quantitative macroeconomic research of today.

Our contribution is rooted in how we deal with three econometric challenges. First, to credibly estimate households' marginal propensities to consume out of transitory income shocks (MPCs), one must observe exogenous income innovations.[2] Importantly, these innovations must be unanticipated and perceived by the recipients as transitory because theoretical predictions depend entirely on these properties (Modigliani 1954, Friedman 1957). Such exogenous shocks with a clear information structure are hard to come by in the data. We obtain transparent identification by using monetary rewards from betting activities in which most Norwegians participate. Second, the income shocks must be observed together with data on household-level consumption and balance sheets, which is a rare combination. We utilize detailed third-party-reported information on households' balance sheets to impute their total consumption *expenditure* from the budget constraint, leaving us with a measure of durable and nondurable consumption expenditure combined. Third, while average short-run consumption responses are interesting in themselves, in order to inform models, one really needs a better understanding of how income innovations are spent and saved over an extended period of time and of the determinants behind heterogeneity in these responses.[3] This requires panel data with rich information on household characteristics, in particular wealth and balance sheets since these play a central role in structural models of consumption dynamics; see, for instance, Kaplan and Violante (2014); Carroll et al. (2017); and Krueger, Mitman, and Perri (2016). We use data that cover the universe of Norwegian households for more than a decade. These data include a variety of household characteristics in addition to balance sheets. To the best of our knowledge, among the many carefully executed empirical MPC studies that exist, this is the first paper to meet all the data requirements above at once.[4] This leaves us with a thorough and transparent characterization of how transitory innovations to income affect household expenditure and saving.

Regarding (i) how transitory income shocks affect consumption and savings over time, we first establish that winners spend a substantial fraction of their prize within the first calendar year of receiving it. Our baseline estimate implies

[2] We use the term "MPC" to describe the fraction of an income shock that is spent over an extended period of time. In our application, MPC means the fraction spent within the calendar year of winning a lottery prize. This interpretation of the word *marginal* in MPC is admittedly somewhat misleading but widely adopted in the literature; see, for instance, Kaplan, Moll, and Violante (2018).

[3] Auclert, Rognlie, and Straub (2018) provide a detailed explanation of why dynamic consumption responses are useful to distinguish between alternative models of household behavior and necessary to address macroeconomic questions like the effect of fiscal policy.

[4] We here have in mind the existing evidence on *actual* household behavior. Surveys of existing studies are provided by Browning and Collado (2001), Jappelli and Pistaferri (2010), and Fuchs-Schündeln and Hassan (2016). An alternative route is to ask how households believe they would respond to hypothetical income shocks. Parker and Souleles (2019) discuss and compare the two approaches in the context of US tax rebates.

a within-year expenditure response around one-half. Of what is not spent, most is saved as deposits, which are thereafter gradually depleted to finance above-normal expenditure up to five years after winning. The dynamics of debt accumulation and saving in stocks, bonds, and mutual funds are quite different, as the responses occur mainly within the year of winning. Debt accumulation drops immediately and thereafter returns to normal. The flow of saving in stocks, bonds, and mutual funds jumps up and thereafter returns to normal. Our estimates imply that after 5 years, households have on average spent about 90 percent of their windfall.

Regarding (ii) which household characteristics systematically relate to MPC magnitude, we find that age and liquid assets stand out. When controlling for all items on households' balance sheets and a variety of other characteristics, it is only age and the liquid assets held before winning that systematically correlate with household-level MPCs. Estimates from a specification where consumption responses interact with a range of household observables imply that a one standard deviation increase in liquidity is associated with an MPC reduction of \$0.08 to the dollar won. Similar effects emerge when we group households along the liquidity strata and estimate consumption responses within each group. The magnitudes of age effects are also considerable. Moving from the youngest quartile (younger than 39 years) to the oldest quartile (older than 63 years), the MPC falls with \$0.15 to the dollar won. Relative to the existing literature, our contribution is to establish these effects in a setting where we observe all household balance sheet components and a variety of other household characteristics together.

Regarding (iii) how responses vary with shock size, we find a stark decreasing relationship. When we split our sample by the amount won, the within-year consumption response declines monotonically with prize size. Our estimates within the lowest prize quartile (less than about US\$2,000) imply that winners of small prizes tend to spend everything within the year of winning or even more by combining the money won with other financial means. In the highest prize size quartile (above US\$8,000), the within-year response lies below one-half.

Lastly, we benchmark our results against a relatively standard life cycle model. We show how our results point toward three model ingredients emphasized elsewhere in the established literature. First, a distinction between liquid and illiquid assets allows the model to match the finding that MPCs fall with liquid assets but not with total wealth. Second and third, we add a realistic lifetime profile of earnings and a luxury bequest motive in order for the model to capture how MPCs decline with age. The upshot is that when these three features are included, the model can fit the heterogeneity effects we detect in the data. It can also generate a time profile for the consumption response that is similar to our empirically estimated one. However, this relatively standard model cannot account for the high empirically estimated average MPC *level* without further extensions. We thus conclude that our empirical estimates are consistent with state-of-the-art economic models where liquidity and life cycle considerations cause cross-sectional MPC *variation*, but the high average MPC *level* we estimate cannot be explained by liquidity constraints alone.

As described above, we do not observe consumption expenditures directly but impute them from observed income flows and end-of-year balance sheets. It is well known that such an imputed measure may suffer from measurement errors

due to capital gains on financial assets and intra-year trading of these (Browning and Leth-Petersen 2003, Baker et al. 2018). To deal with these issues, we follow two alternative routes. First, we restrict attention to households whose financial balance sheets contain only deposits and debt, leaving us with approximately 60 percent of our original sample. Within this subgroup, concerns regarding measurement error in our imputed expenditure measure are reduced to a minimum. Second, we explore MPCs in more recent data where we can utilize a stockholder registry to estimate each household's capital gains stock by stock. All our results are confirmed in both exercises. Another concern, more specific to our strategy of using lottery prizes, could be that we do not observe the bets placed. However, among winners, there are no pre-trends in consumption expenditure or any other observables in the years prior to winning. Hence, any systematic increase in betting activity in the year of winning appears to be orthogonal to the household characteristics we consider when estimating heterogeneity effects. It follows that spending on lottery tickets cannot be driving the heterogeneity effects we estimate. Finally, an estimated average MPC across all households can be sensitive to outliers, in particular among high-prize winners. We establish a range from 0.35 to 0.71 for our baseline estimate, depending on how we deal with these outliers. However, apart from the MPC level, our main results are unaffected by outliers, as both the estimated dynamic paths of responses and the heterogeneity results are insensitive to how we deal with them.

A vast literature studies how income shocks affect consumption. One body of papers focuses on the US tax rebates in 2001 and 2008, and the bulk of existing evidence on actual income shocks and consumption stems from these quasi-experiments. Parker et al. (2013) consider total consumption expenditure like we do. Exploring the 2008 rebate episode, where transfers per adult ranged from US$300 to US$600, they find total consumption responses between 0.5 and 0.9 within 3 months of payment receipt.[5] Notably, both this study and those focusing on nondurable consumption compare households receiving a pre-announced rebate at different points in time, effectively identifying the effects of *anticipated* income shocks (Agarwal, Liu, and Souleles 2007; Johnson, Parker, and Souleles 2006; Parker et al. 2013; and Shapiro and Slemrod 2003, 2009). Relatedly, Hsieh (2003) and Kueng (2018) estimate consumption responses to large predetermined payments from the Alaska Permanent Fund. These responses are conceptually different from what we estimate since our estimates include both announcement and imbursement effects and hence should lie higher. More comparable to our context, Agarwal and Qian (2014) study a transfer episode in Singapore and find an average spending response around 80 percent of the stimulus received within 10 months after the transfer was announced. Overall, our baseline within-year MPC estimate is on the low side of the estimates provided by these studies, but once we restrict attention to prizes of more similar magnitudes to those used in existing studies, the estimates are similar.

[5] The 2007–2008 US tax rebate distributed US$300–600 to single individuals, US$600–1,100 to couples, and in addition gave US$300 for each child qualified for the tax credit. For details, see Parker et al. (2013).

The liquidity effects we find are consistent with previous findings in the literature, such as Misra and Surico (2014), who use the aforementioned survey evidence on the US tax rebates.[6] Our contribution here is to establish that liquidity is associated with MPC variation after controlling for all household balance sheet components and a variety of other household characteristics. Moreover, we find that illiquid households display markedly higher MPCs both in the short and in the medium run compared with liquid households.

Existing evidence on how MPCs vary with the magnitude of income shocks is limited. The literature has predominantly resorted to comparing MPC estimates across alternative studies that differ by the size of the shocks utilized, as in the survey of Browning and Crossley (2001). To the best of our knowledge, only Scholnick (2013) has previously provided direct evidence on how consumption responses vary with shock size like we do. He considers how consumption changes after final mortgage payments and finds that when these anticipated income changes are bigger, the marginal consumption response is smaller, consistent with our findings. Relative to this study, our value-added is to consider unanticipated income shocks with a far larger sample of treated observations (147 in Scholnick 2013, more than 20,000 in our data) over an extended sample period and with a wide range of household characteristics, including balance sheets. More recently and coincident with our work, Fuster, Kaplan, and Zafar (2020) and Christelis et al. (2019) address the issue with survey evidence on how respondents believe they would respond to hypothetical income shocks. The former find that MPCs increase with shock size; the latter find the opposite.

Baker et al. (2018) study a sample of higher-end individuals using retail investment data from a bank in Germany to quantify the measurement errors stemming from capital gains, trading fees, and within-year trading activities that are unobserved in annual registers of the type we use. They find that these errors are present yet quantitatively moderate and centered around zero. They also provide suggestive evidence that biases to estimated consumption responses to income shocks are moderate. Our results point in the same direction: our MPC estimates do not change notably when we exclude households where financial capital gains and intra-year trading activities create measurement errors.

Parker (2017) questions if empirically observed associations between liquidity and MPCs are *situational*, in the sense that an individual's MPC depends on how liquid he happens to be at the time of winning. The alternative is that liquidity correlates with unobserved household characteristics that raise consumption sensitivity, such as impatience, risk tolerance, or intertemporal elasticity of substitution, recently emphasized by Aguiar, Bils, and Boar (2020). In our setting, one would expect net wealth, education, and the portfolio share of risky assets to correlate

[6] Other examples are Leth-Petersen (2010), who studies the impact of a credit market reform on consumption in Denmark; Aydin (2015), who studies exogenously varying credit limits in a European retail bank; Baker (2018), who studies the interaction between household balance sheets, income, and consumption during the US Great Recession; Agarwal et al. (2015) and Gross, Notowidigdo, and Wang (2020), who study consumption dynamics around discontinuities in credit scores; and Cloyne, Ferreira, and Surico (2020), who explore micro responses to monetary policy shocks. Recently, Fuster, Kaplan, and Zafar (2020) and Christelis et al. (2019) provide survey evidence pointing in the same direction, when they ask how households would allocate hypothetical windfall gains.

with such unobserved characteristics. For instance, wealth accumulation is a channel through which patience affects MPCs in heterogeneous agent models (Carroll et al. 2017). It is therefore striking that when observables like wealth and education are controlled for together with liquidity and age, it is only the latter two that vary significantly with the consumption response to lottery prizes. In addition, liquidity and age remain significant even when we control for each household's historical co-movement between consumption and income, capturing the household-specific component of consumption sensitivity. These results support a situational interpretation of our findings.

Consumption and saving responses to lottery income have been studied before, most prominently by Imbens, Rubin, and Sacerdote (2001) and Kuhn et al. (2011). The former study considers 500 winners of large prizes in a Massachusetts lottery. Unlike the setting we explore, these prizes were paid out gradually, obscuring comparison with our estimates. The latter study considers a lottery in the Netherlands where households won €12,500. The Dutch findings stand out from ours and the tax rebate literature in that neither durable nor nondurable consumption responded by much.[7] More recently, Swedish lotteries have been used to identify income effects on health, labor supply, and portfolio choice, but not on consumption.[8]

From the perspective of standard economic theory, we are studying a well-defined income shock with implications that generalize to other sources of income variation. Still, the extent to which evidence from lotteries can be generalized to other income shocks is debatable. Ng (1965) and recently Crossley, Low, and Smith (2016) argue that households might gamble to "convexify" their feasibility set when discrete-type purchases are desired. This would imply that our estimates are upward biased, as some of the winners have gambled precisely because they have high MPCs. Here it is reassuring that our estimated spending responses align with the existing evidence on transfers. Moreover, participation in betting activities is widespread in Norway, partly because it is largely organized by the state-run entity Norsk Tipping, which redistributes their surplus to charitable purposes such as sports activities for children. According to Norsk Tipping, about 70 percent of the Norwegian adult population participated in one of their lotteries in 2012. Consistent with this observation, our descriptive statistics reveal that differences between winners and nonwinners are negligible. In particular, the consumption response to regular income variations and to the receipt of inheritance are almost identical for the two groups. In addition, while conceptually the gambling-to-convexify argument could explain (at least part of) high MPC levels, it seems less relevant for our main contribution, namely to explore the determinants of MPC heterogeneity. For all these reasons, it appears unlikely that this mechanism is driving our main results. Indeed, when we look at how our rich set

[7] While lottery prizes constitute unanticipated *transitory* income shocks, Fuchs-Schündeln (2008) studies an unanticipated *permanent* income shock, the German reunification. She finds results in line with a life cycle model of savings and consumption.

[8] Using Swedish lottery data, Cesarini et al. (2016) study effects of wealth on health and child development, Briggs et al. (2021) study effects on stock market participation, and Cesarini et al. (2017) study effects on labor supply. In the Appendix, we validate our empirical strategy by estimating earnings responses in our sample and comparing them to the findings of Cesarini et al. (2017).

of observables evolve before households win in the lottery, such as debt, wealth, and liquidity, none of these change in the immediate years before winning. These pre-trends support the identifying assumption that the timing of winning in a lottery is exogenous.

We believe our findings are most interesting when cast against incomplete markets models, as developed by Huggett (1993), Aiyagari (1994), and Carroll (1997). In these models, the main determinant of households' MPCs is net wealth. In contrast, our empirical findings indicate that net wealth is unimportant once liquidity is accounted for. While in conflict with a literal interpretation of buffer stock savings models, this finding supports extensions and modern interpretations of them. First, the approach of calibrating one-asset buffer stock models to data on liquid asset holdings rather than total wealth, as in for instance Carroll et al. (2017), is supported by our results. Second, the distinction between net wealth and liquid assets is explicit in recent two-asset frameworks. A cornerstone here is Kaplan and Violante (2014), who show how households might be rich yet behave in a hand-to-mouth fashion because their assets are illiquid. Norwegian households' balance sheets are dominated by housing, the prototypical illiquid asset, and we do indeed find that MPCs vary with liquid assets but not with housing wealth. The result that consumption responsiveness declines with shock magnitude also fits with what such buffer stock models predict. Still, even though we find that MPCs decline with liquidity and shock size, the responses remain high even among liquid winners of large prizes. Conventional models of nondurable consumption do not imply such magnitudes. Regarding the negative association between MPCs and age, this is at odds with a frictionless life cycle model with a flat earnings profile and no bequest. Instead, it points toward extensions emphasized in the more recent literature (De Nardi 2004 and De Nardi and Fella 2017). As our model exercise shows, a realistic earnings profile coupled with borrowing constraints will raise MPCs early in life, while luxury bequest motives will reduce MPCs later on.

The remainder of our paper is organized as follows. Section I presents the institutional setting and data. Section II provides our benchmark estimates of the MPC out of lottery earnings, including dynamic responses. Section III contains our results on how MPCs vary with household characteristics and the amount won. Section IV discusses robustness analyses. Section V compares our estimates to those from a specific model. Section VI concludes.

I. Institutional Background, Data, and Sample Selection

We base our study on Norwegian administrative data. Since Norway levies both income and wealth taxes, the data from the tax registry provide a complete and precise account of household income and balance sheets over time, down to the single asset category for all Norwegian households. From these records, we create an imputed measure of consumption using the household budget constraint. Moreover, as part of their yearly tax filings, Norwegian households must report received gifts and prizes above approximately US$1,100. Below, we describe the data sources, explain the consumption measure we construct, present the lottery data and summary statistics about our sample, and outline our empirical strategy.

A. *Administrative Tax and Income Records*

Our main data source is the register of tax returns from the Norwegian Tax Administration, which contains detailed information about all individuals' incomes and wealth, for the period 1993 to 2015 (Statistics Norway 2015d). We combine these data with data on educational achievement (Statistics Norway 2015e) and household identifiers from the historical event database FD-trygd (Statistics Norway 2015c) to aggregate all income and wealth information to the household level.[9] Every year, before taxes are filed in April, employers, banks, brokers, insurance companies, and any other financial intermediary send to both the individual and to the tax authority information on the value of the assets (and liabilities) owned by the individual and administered by the employer or intermediary as well as information on the income from these assets.[10]

The tax authority then prefills the tax form for the individual to amend and approve. These data have the advantage that there is no attrition from the original sample (apart from death or migration to another country) due to participants refusing to share their data. In Norway, these income and wealth records are in the public domain and pertain to all individuals.

B. *Measuring Consumption*

A challenge to most empirical studies of consumption is (a lack of) access to a precise longitudinal measure of household consumption expenditures (see Pistaferri 2015 for a recent summary of the literature on the measurement of consumption). Traditionally, studies have employed data on household consumption from surveys, as in Johnson, Parker, and Souleles (2006) or Parker et al. (2013) with the Consumer Expenditure Survey (CEX) in the United States or Jappelli and Pistaferri (2014) using the Survey on Household Income and Wealth (SHIW) in Italy. Surveys have the advantage that the researcher can obtain direct measures of self-reported consumption or the self-assessed marginal propensity to consume out of a hypothetical income shock as in the SHIW. However, as is well known, expenditure surveys and household surveys often suffer from small sample sizes and attrition and face considerable measurement errors that are potentially correlated with important observable and unobservable characteristics (Meyer, Mok, and Sullivan 2015). There is also an ongoing discussion about the reliability of self-reported marginal propensities to consume from hypothetical income shocks (Parker and Souleles 2019).

[9] In Norway, labor (and capital) income is taxed at the individual level, while a wealth tax is levied at the household level.

[10] These assets are for the most part assessed at market value. Housing values from the tax registries, however, are typically undervalued in Norway before 2010, when valuations for the purpose of wealth taxation were reassessed nationwide. We have therefore combined a variety of data sources to improve the valuation of housing. Transactions data (Statistics Norway 2015g), information on dwellings (Statistics Norway 2015b), and the population and housing census (Statistics Norway 2011) allow us to identify ownership of each single dwelling and its precise location. Following contemporary tax authority methodology, we estimate a hedonic model for the log price per square meter as a function of house characteristics (number of rooms, etc.), time and location indicator variables, and their interactions. The predicted values are then used to impute house values for each year. Detailed documentation of our estimated house prices is provided in Fagereng, Holm, and Torstensen (2021).

Instead of relying on consumption surveys, an alternative is to impute expenditure from income and wealth data in administrative tax records. We follow this approach. Equipped with the balance sheet data described above, we impute consumption for Norwegian households in a similar fashion as Browning and Leth-Petersen (2003) (for Denmark) and later Eika, Mogstad, and Vestad (2020) and Fagereng and Halvorsen (2017) for Norway.[11]

The imputation procedure starts from the accounting identity

$$Y = C + S, \tag{1}$$

which states that disposable income (Y) in each period must be either consumed (C) or saved (S). When combining this identity with balance sheet data, a number of issues must be dealt with to back out a consumption measure.

We start by excluding household-year observations that are known from the literature to cause measurement issues using the imputation procedure. First, we focus on stable households, excluding household-year observations involving a change in the number of adults (by divorce or marriage) to avoid the financial reshuffling of balance sheets that often takes place in these periods. Second, we exclude household-year observations in which members of the household are reported as a business owner or farmer, as both assets (private equity, business assets, or farm equipment) and income streams from these are not well measured in the data. We here also exclude households with extreme financial returns (for the subsample where we utilize data directly from the stockholder registry). Third, we leave out the household-year observations where a household moves or is involved in a housing market transaction. Housing transactions are observed in the data, so in principle, these are unproblematic. However, timing issues regarding the dates of the actual money transactions (purchase amount, debt uptake) relative to when the house sale was made can make it difficult to assign the correct money flow to the right calendar year. From this starting point, we now discuss issues that are of importance and relevance to our purpose of studying the consumption responses to lottery income.[12]

[11] Ziliak (1998) attempts to impute consumption using data from the Panel Study of Income Dynamics (PSID) in the United States. However, in the PSID, wealth is only reported in every fifth wave, making it necessary to also impute the yearly wealth data. Lately, several researchers have implemented the imputation method on Scandinavian countries where yearly data on both income and wealth are available. Browning and Leth-Petersen (2003) (and later Kreiner, Lassen, and Leth-Petersen 2015) implement this method using Danish register data; Eika, Mogstad, and Vestad (2020) and Fagereng and Halvorsen (2017) using Norwegian data; and Koijen, Van Nieuwerburgh, and Vestman (2015) and Kolsrud, Landais, and Spinnewijn (2020) using Swedish data. Other examples are Browning, Gørtz, and Leth-Petersen (2013); Leth-Petersen (2010); Autor et al. (2019); and Di Maggio, Kermani, and Majlesi (2020). Browning, Crossley, and Winter (2014) review the literature.

[12] After the imputation procedure described in this section, we also drop extreme observations of consumption, conditional upon the amount won, and winners of prizes above US$150,000. For each percentile of the prize size distribution, we exclude observations in the top and bottom 2.5 percentiles of the consumption distribution. By conditioning on prize size, we avoid systematically omitting low-prize winners with exceptionally low consumption and high-prize winners with exceptionally high consumption, which would bias MPC estimates downward. We return to the role of sample selection in Section IV. There we show that (i) our baseline MPC estimate is somewhat sensitive to how we deal with outliers in consumption and lottery prizes (it increases from 0.52 to 0.71 if we do not trim at all, while it drops to 0.35 if we trim unconditionally), and (ii) this does not affect our results regarding heterogeneity, which constitute our main finding. In a previously circulated version of this paper, we presented a

Disposable income observed in our data is defined as $Y_t = I_t - T_t + \sum_j^J d_{jt} + L_t$, where I_t is labor income, T_t is tax payments net of transfers, and d_{jt} is the capital income from each asset j held by the household during year t. For interest rate expenditure on debt, d is negative. For housing, indexed h, we impute $d_{ht} = \rho H_{t-1}$, with $\rho = 0.03$ and where H_{t-1} is beginning-of-year housing wealth.[13] The variable L_t is net income from any other source, such as inheritance or lottery prizes. Notably, all the components of Y_t except the implicit housing income (d_{ht}) are directly observed in the administrative tax records.

Consumption expenditure is imputed from the budget constraint, equation (1), where S_t consists of the period t income flow that is set aside and saved, often referred to as "active" saving. The challenge for consumption imputation is to calculate S_t and in particular to adjust wealth accumulation for unrealized capital gains. For illustration, assume that the household holds each asset over the entire year and then rebalances its portfolio at market prices at the end of the year. Then S_t is given by

$$S_t = W_t - W_{t-1} - \sum_{j=1}^{J} \left(p_{jt} - p_{jt-1}\right) a_{jt-1}, \tag{2}$$

where $W_t = \sum_{j=1}^{J} p_{jt} a_{jt}$ is end-of-year net wealth, while p_{jt} is the end-of-year price, and a_{jt} is the end-of-year stock of asset j. As the expression shows, we need to isolate capital gains and subtract them from the total wealth change. In the administrative tax records, we directly observe W_t and the value held within each specific asset class k, $w_{kt} = \sum_j^{J_k} p_{jt} a_{jt}$. In addition to housing, the classes are deposits, outstanding receivables,[14] debt, stocks, bonds, and mutual funds, held abroad and at home.

Our procedure is to use aggregate price indices p_{kt} to approximate $\sum_j^J \left(p_{jt} - p_{jt-1}\right) a_{jt-1} \equiv \sum_j^J w_{jt-1}\left((p_{jt}/p_{jt-1}) - 1\right)$ by $\sum_k^K w_{kt-1}\left((p_{kt}/p_{kt-1}) - 1\right)$. We approximate stock price changes with growth in the Oslo Stock Exchange (OSE), mutual fund prices with a weighted average of the OSE and the MSCI World Index, and bond prices with the Treasury bill rate. Hence, for these assets, we are assuming that each household holds the market portfolio. There are no capital gains on deposits and debt, so the imputed capital gains only apply to the risky share of the portfolio.

Under the assumptions above, we observe a measure of Y_t and S_t for each household. We then impute household consumption as $C_t = Y_t - S_t$. In Appendix Figure A1, we plot our imputed consumption per person against consumption per capita in the National Accounts. The two series track each other closely. The main difference is that the imputed consumption series is more volatile. This volatility stems from the fact that we exclude certain households, for instance, those who hold

baseline MPC estimate of about 1/3 because we did not condition on the amount won when trimming the consumption measure.

[13] We attribute to each homeowner's consumption expenditure a value of owner occupied housing services equal to 3 percent every year. This enables us to compare the consumption of renters (which includes rental payments) and homeowners. The value of these services is meant to represent the price the homeowner would have paid if s/he were to rent the same home on the market. Three percent is close to what Eika, Mogstad, and Vestad (2020) find as the rent-to-value for Norway, using data from National Accounts.

[14] Outstanding receivables consists of loans to family and friends, salary and maintenance payments that one is owed, and advances for services not yet received.

private equity, which makes our set of households vary somewhat from year to year, especially among the very wealthiest.

In short, the description above shows how our imputed consumption measure rests on two key assumptions. First, we assume there is no intra-year trading. Second, if a household owns stocks, bonds, or mutual funds, we assume it holds the market portfolio of the respective asset class. We now discuss scenarios in which the potential measurement errors that follow could be problematic for our purposes and how we deal with them.

First, we note that our interest lies in understanding dC_t/dl_t and its heterogeneity, where l_t is lottery income. In our main analysis, we will be controlling for individual-fixed effects in consumption levels. Hence, measurement error in C_t is only problematic insofar as it correlates with l_t, after controlling for individual fixed effects in C_t.

We face a potential problem when a lottery winner invests part of the prize in risky assets. If the acquired assets increase (or decrease) in value within the same year as the lottery win, our approach interprets the consequent wealth increase as "active" saving (S) and therefore subtracts it from income when imputing consumption. However, the capital gains from the newly acquired assets do not imply lower expenditure and should not be subtracted. As mean returns are positive over our sample period, this measurement error might bias our estimates of how lottery winnings affect consumption expenditure downward. Moreover, because the bias is positively correlated with unrealized returns, it is likely to be greater for households who buy riskier assets or for some other reason systematically obtain more extreme returns. To address this potential problem, we redo our main regressions on a sample of households who *never* hold risky assets (stocks, bonds, and mutual funds) in our sample period. For these households, measurement error in imputed consumption is of limited concern because capital income and expenditure flows (interest income and expenses) during the year are directly observed, together with end-of-year deposit and debt levels. As documented in Section IVA and Appendix A.8, restricting our attention to this subsample does not affect our results.

The assumption that each household holds the market portfolio within each asset class is obviously simplistic. For instance, Fagereng et al. (2020) document substantial heterogeneity in returns across households and within asset classes. We argue, however, that this source of measurement error is unlikely to drive our inference. First, there is little reason why it would correlate with lottery prizes. The main explanation would be heterogeneity in risk aversion, which might cause both greater gambling activity and higher MPCs, but there is little sign of such a relationship when we compare winners to nonwinners in Table 1 below or when we consider the predictability of prize size in Table 2. Second, when we drop all households with risky assets, we are left with a sample whose returns are directly observed and returns heterogeneity is unproblematic.

To be clear, our imputation procedure leaves us with a consumption measure that includes both durable and nondurable goods and services. This must be kept in mind when interpreting our estimates. If winners respond by increasing their end-of-year cash holdings and do not report these, this response will also be counted as consumption. However, this is likely to be a minor phenomenon, as cash holdings are

very low in Norway.[15] Unreported gifts and inter vivos transfers are also included in our measure of consumption expenditure. Hence, if parents transfer resources to offspring in response to a lottery winning but do not report it, then this transfer will be measured as consumption expenditure of the parents, while it may be saved by the offspring. In Section IVA, we show that our results are robust to including adult offspring responses in our estimated MPCs, indicating that this is indeed a minor issue.

C. *Gambling in Norway*

In Norway, only two entities are allowed to offer gambling services: Norsk Tipping (mainly lotteries and betting on sports events) and Norsk Rikstoto (horse racing). Both are fully state-owned companies, and all surpluses are earmarked charitable causes. According to Norsk Tipping, 70 percent of Norwegians above the age of 18 gambled in 2012 through their services.[16]

During our sample period between 1994 and 2006, gambling in Norway took place mainly through one of the more than 5,000 commissioned venues (about 1 per every 800 adult Norwegians), usually a kiosk or a local supermarket. Individuals filled out their betting forms and submitted them at one of the commissioned venues. In the event of a successful gamble, smaller prizes (less than NOK1,000, equivalent to about US$110) could be cashed out directly at any of these venues, whereas larger prizes were transferred directly to the winners' bank account within 14 days. All prizes included in our sample are paid out lump-sum within a few weeks. Income from gambling in Norway is generally tax-exempt, as is income from EU/EEA-area lotteries where the surplus primarily is given to charitable causes. However, Norwegian citizens are obliged to report lottery prizes exceeding NOK10,000 (about US$1,100) to the tax authority. Importantly, it is in the individuals' self-interest to report such windfall gains, as a sudden increase in wealth holdings from one year to another could raise suspicions of tax fraud and cause further investigation by the tax authority.[17] As the reporting requires display of a dated prize receipt, there is no scope for exaggerating such windfall gains or misreporting when the prize was won.

The data on lottery prizes include all games arranged by Norsk Tipping and Norsk Rikstoto and similar betting activities in other EEA countries. These data therefore cover a wide variety of games, such as scratch cards, bingo, horse racing, and sports betting. Our data do not include prizes won in card games or other casino games.

In 2007, the minimum threshold for reporting lottery prizes was increased to NOK100,000 (about US$11,000). To maintain the larger variation in windfall gains, we therefore limit our attention to the period 1994–2006. Moreover, we limit our sample to households who win only once. This is because we want to estimate responses to surprise income innovations, while for "serial winners," it is less clear whether yearly prize revenues can be considered as unexpected. In particular, we

[15] For instance, around 6 percent of M1 (deposits and cash) in 2008 were cash holdings (https://www.ssb.no/en/bank-og-finansmarked/statistikker/pengemengde).

[16] See Norsk Tipping Annual Report 2012. For details on gambling in Norway, see the Gaming and Foundation Authority.

[17] Norway also has a long tradition of public disclosure of tax filings, involving the public display of yearly information on income and wealth of individuals (Bø, Slemrod, and Thoresen 2015).



FIGURE 1. DISTRIBUTION OF LOTTERY PRIZES, 1994–2006

Notes: The figure shows the distribution of lottery prizes, denoted in US$, year-2000 prices. Each bin is US$1,000 wide, starting from US$1,000. The rightmost bar contains all prizes above US$60,000.

want to exclude systematic gamblers in horse racing and sports betting who might consider prizes as part of their regular income.

Figure 1 displays the distribution of lottery prizes in our sample. There is a clear peak for the smallest prize bin, which contains winners of US$1,100 to US$2,000. More than 20 percent of our prizes are of this magnitude. There is also substantial variation in the amount won, which will allow us to study how consumption responses vary with shock size.

D. *Descriptive Statistics*

Table 1 displays summary statistics for the sample of nonwinners and our sample of winners (measured in the year before winning) between 1994 and 2006. Age and education refer to the household head; all other variables are computed at the household level. Income after tax includes net transfers, capital income, labor income, and business income.

Winners and nonwinners are largely similar. Winners are slightly older, have somewhat fewer household members, and have slightly less education. The levels of income, consumption, and wealth are also similar. The small difference in mean net wealth that exists is primarily due to housing wealth. Regarding balance sheet composition, Table 1 reveals that a higher share of winners own risky assets (29 percent against 25 percent) and that their mean share of risky assets (stocks and mutual fund holdings relative to net wealth) is marginally higher than is the case for nonwinners. This pattern could suggest that households who win in lotteries are more risk tolerant than nonwinners, but the observed differences are small and do not suggest that winners exercise fundamentally different consumption behavior than nonwinners.

TABLE 1—SUMMARY STATISTICS, 1994–2006

	Nonwinners ($N = 2{,}980{,}347$)		Winners ($N = 23{,}728$)	
Variable	Mean	SD	Mean	SD
Age_t	49.67	(19.57)	50.61	(15.13)
Year_t	2,000.93	(4.26)	2,000.31	(3.34)
Household size_t	2.41	(1.39)	1.91	(1.14)
$\text{Number of children under 18}_t$	0.53	(0.94)	0.30	(0.71)
$\text{Years of education}_t$	12.99	(3.08)	12.55	(2.60)
$\text{Income after tax}_{t-1}$	25.98	(92.07)	24.11	(11.47)
Salary_{t-1}	22.15	(27.67)	23.49	(20.86)
Consumption_{t-1}	20.60	(15.84)	20.83	(13.76)
Lottery_t	.	.	9.24	(16.13)
Net wealth_{t-1}	79.17	(260.76)	77.26	(100.51)
Debt_{t-1}	36.07	(113.67)	30.89	(39.79)
$\text{Cars and boats}_{t-1}$	2.69	(7.49)	3.43	(6.03)
$\text{Housing wealth}_{t-1}$	95.16	(205.46)	93.60	(95.36)
$\text{Liquid assets}_{t-1}$	20.08	(164.06)	14.55	(25.51)
Deposits_{t-1}	15.93	(50.83)	12.30	(22.28)
Stocks_{t-1}	1.89	(145.17)	0.63	(4.72)
Bonds_{t-1}	1.06	(25.85)	0.62	(4.86)
$\text{Mutual funds}_{t-1}$	1.19	(12.05)	1.00	(4.45)
$\text{Risky share of balance sheet}_{t-1}$	0.07	(0.18)	0.08	(0.20)
$\text{Share of households owning risky assets}_{t-1}$	0.25	(0.44)	0.29	(0.46)
MPC-Income after tax	0.831	(0.001)	0.814	(0.014)
MPC-Inheritance	0.545	(0.007)	0.523	(0.049)

Notes: Nonwinners are defined as households that did not win a prize during the sample period, 1994 to 2006. In the table, each nonwinner is represented by a randomly selected year during the sample period. For winners, we display the year prior to winning ($t - 1$). Monetary amounts are CPI-adjusted (Statistics Norway 2015a) to the year 2000 and then converted to (thousands of) US$ using the mean exchange rate in the year 2000 (Norges Bank 2020). Liquid assets is equal to the sum of deposits, stocks, bonds, and mutual funds. Risky share of balance sheet is the fraction of liquid assets held in either stocks or mutual funds. Share of households owning risky assets is an indicator taking the value one if at least some fraction of liquid assets is invested in either stocks or mutual funds. MPC-Income after tax and MPC-Inheritance show the result from linear regressions of consumption on income after tax and inheritance, respectively. For MPC-Income after tax and MPC-Inheritance, the bracketed numbers are standard errors.

The final two rows of Table 1 display naïvely estimated marginal propensities to consume out of after-tax income and received inheritance. These estimates are not to be interpreted structurally. They simply are the resultant coefficients from regressing consumption on contemporaneous income and inheritance. Their purpose is to illuminate differences in consumption dynamics between the two groups. As we see, the estimates are approximately identical in the two groups.

E. *Empirical Strategy*

As explained above, we estimate the effects of lottery prizes on a sample of households who have won exactly once during our sample period. We utilize various regressions based on the specification

$$C_{i,t} = \beta_0 + \beta_1 \mathit{lottery}_{i,t} + \beta_2 X_{i,t-1} + \alpha_i + \tau_t + u_{i,t}, \tag{3}$$

where i is a household identifier, t represents calendar year, $C_{i,t}$ is household i's consumption in year t, $lottery_{i,t}$ is the amount won in year t, $X_{i,t-1}$ is a vector of controls, α_i is a household-fixed effect, and τ_t is a time-fixed effect.

To prevent lagged responses from contaminating our inference, we drop households in the years after they won.[18] Had we instead kept households after they won, our point estimates of β_1 would become downward biased if consumption responds persistently to income shocks. It follows that our identification is obtained by comparing households' consumption in the year of winning to their consumption in previous years. These individual responses are then weighted together to an average treatment effect across households winning different amounts. We are here leaning on two identifying assumptions: (i) the timing of winning is exogenous, and (ii) the amount won is exogenous.

Our estimate of β_1 represents an average increase in consumption expenditure per dollar won, consistent with how MPCs are estimated and interpreted elsewhere in the literature. Note, however, that the weights in this average increase with prize size. A point estimate of β_1 will therefore be pulled toward the MPCs of winners of relatively high prizes. In Appendix A.3, we derive the OLS weights to illustrate the intuition, and in Appendix A.4, we present a simple simulation exercise that illustrates this point. In what follows, we start with the linear specification in equation (3) and thereafter dissect the potential size effects together with the effects of various household characteristics.

F. *Internal Validity*

A shortcoming of our data is that we only observe how much households win, not how much they bet. Hence, one might worry that the households' lottery winnings are systematically related to other determinants of consumption. We therefore explore if observed household characteristics change in any systematic fashion in the years before winning and the extent to which they can predict the amount won.[19]

To conserve space, we here point to the pre-trends in the dynamic responses plotted in Figure 2, while further pre-trends are presented in Appendix A.6. Neither consumption, deposits, stocks, bonds, and mutual funds, nor debt evolve differently than normal in the years before winning. Appendix A.6 shows that the same holds for total income, net wealth, risky portfolio share, household size, and number of children. In short, all the pre-trends indicate that the timing of winning is exogenous.

Table 2 summarizes the predictability of prize size conditional upon winning, or the "intensive margin" of prize variation. Predictability along the intensive margin is useful to illuminate the extent to which our prizes can be considered exogenous. Column 1 focuses on lagged values of consumption and balance sheet

[18] Our sample includes 23,728 winners, as seen in Table 1, which amounts to 93,631 household-year observations when we also include observations from winners in the years prior to winning.

[19] We have also assessed our empirical strategy by estimating the effects of lottery income on labor earnings and comparing to Cesarini et al. (2017). That study observes the amounts bet together with the prize received and estimates labor supply effects in Sweden. The estimates we obtain in Norway using our lottery data and the same strategy for earnings as we use for consumption are similar to what Cesarini et al. (2017) find, but our estimates are less precise. See Appendix A.5 for details.

Table 2—Predictability of Lottery Prize Size

Dependent variable: Lottery Prize$_t$			
Regressors	(1)	(2)	(3)
Consumption$_{t-1}$	0.022		0.022
	(0.012)		(0.012)
Liquid assets$_{t-1}$	−0.006		−0.006
	(0.005)		(0.005)
Income$_{t-1}$	−0.035		−0.050
	(0.020)		(0.021)
Net Wealth$_{t-1}$	−0.000		0.000
	(0.002)		(0.002)
Debt$_{t-1}$	0.014		0.013
	(0.005)		(0.005)
Risky asset share$_{t-1}$	0.865		0.917
	(0.695)		(0.696)
Age$_t$		−0.019	−0.006
		(0.008)	(0.010)
Household size$_t$		0.110	0.385
		(0.481)	(0.650)
Household size$_t^2$		0.032	−0.003
		(0.100)	(0.137)
Number of children under 18$_t$		0.173	−0.053
		(0.246)	(0.356)
R^2	0.009	0.008	0.009
Partial R^2 of regressors	0.002	0.001	0.003
Observations	14,742	23,728	14,742

Notes: Each column represents a separately estimated regression of lottery prize among winners on predetermined characteristics. The sample sizes differ between columns because columns 1 and 3 require us to observe consumption in the year prior to winning. All regressions include time-fixed effects. Partial R^2 of regressors shows the increase in R^2 by adding the regressors to a specification with only time-fixed effects. Robust standard errors are in parentheses.

variables, column 2 closely follows Cesarini et al. (2017) by applying a similar vector of controls to the one they use in their study of Swedish lottery winners, and column 3 includes all controls together. Clearly, the predictive power of observable household characteristics for the amount won is low. All controls together explain hardly any of the variation in lottery prizes, as reflected by an R^2 below 1 percent. Some coefficients differ significantly from zero, in particular, those on debt (columns 1 and 3), income (column 3), and age (column 2). However, these associations with the amount won are small. A $1,000 increase in income predicts a $50 reduction in prize size, a $1,000 increase in debt predicts a $13 increase in prize size, and a $1,000 increase in consumption predicts a $20 increase in prize size.

Given the absence of visible pre-trends in observables and their lack of power in predicting the amount won, we find it unlikely that unobserved variables drive the MPC estimates that follow.

II. Consumption and Savings Responses to Lottery Prizes

This section presents results on how lottery income affects consumption expenditure and savings over time. We first estimate how consumption responds within the

TABLE 3—THE MPC OUT OF LOTTERY PRIZES

	(1)	(2)	(3)
Consumption response	0.587	0.524	0.520
(Observations = 93,631)	(0.013)	(0.017)	(0.017)
Time-fixed effects	Yes	Yes	Yes
Household-fixed effects	No	Yes	Yes
Additional controls	No	No	Yes

Notes: The estimated consumption response to a lottery prize from least-squares regressions using specification (3). Additional controls are age, age^2, age^3, age^4, household size, household size^2, number of children under 18, and $t-1$ values of after-tax income, liquid assets, debt, net wealth, and risky asset share. The sample size (93,631) now includes the 23,728 winners and all observations of winners in the years prior to winning. The standard errors in parentheses are robust and clustered at the household level.

year of winning, before providing time paths of responses for different balance sheet components as well as consumption.

Table 3 reports our estimates of the within-year consumption response to a lottery prize using specification (3). As we move horizontally across the table, we gradually add controls for individual-fixed effects and household characteristics. Note that when interpreting these results, it is key to recognize that our estimates of β_1 reflect weighted averages of individuals' within-year responses, where the weights increase with prize size as explained in Appendix A.3. These estimates are therefore to be considered a starting point, as our study's main contribution lies in estimating dynamic consumption responses and dissecting how responses vary with observables.

If prizes were perfectly random, additional controls beyond year-fixed effects would be superfluous. In Table 3, we see that when individual-fixed effects are added, the point estimate drops from 0.59 to 0.52. Partly, this is due to the fact that winners of higher prizes typically consumed a little more also before they won, as we saw in Table 2. The main takeaway from Table 3 is that the point estimates are unaffected when including further controls beyond fixed effects. Hence, for any omitted variable to drive our results, it must correlate with consumption and prize and be independent of the variables we observe and control for. In addition, any such joint correlation between our outcome variable and prize must lie far beyond the influence of the variables we observe. Because these observables span the main candidates that economic theory suggests for explaining consumption patterns, we regard it as reasonable to interpret the results in Table 3 as causal. Moreover, since additional controls have no impact on our point estimates, we include only time- and household-fixed effects going forward.[20]

The results above focus on within-year effects, but equally interesting is how consumption responds over time. As argued by Auclert, Rognlie, and Straub (2018), dynamic responses are informative for discriminating between economic

[20] In Table A1 in Appendix A.2, we report estimates from various other specifications than equation (3) that have been applied elsewhere in the literature, in particular dynamic specifications controlling for lagged consumption. These yield similar results to those in Table 3 and share the property that additional controls beyond time- and individual-fixed effects are superfluous.

models and likely to be key determinants of how aggregate impulses propagate. We therefore move on to estimate impulse responses over the five years after winning, utilizing the specification

$$(4) \quad Y_{i,t+k} = \beta_{0,k} + \beta_{1,k} lottery_{i,t} + \alpha_i + \tau_{t+k} + u_{i,t+k} \qquad \text{for } k = -4, -3, \ldots, 5.$$

The outcomes $Y_{i,t+k}$ that we consider are consumption and saving in various asset classes. The $\beta_{1,k}$ terms are the main coefficients of interest. For instance, when $Y_{i,t+k}$ is consumption, each $\beta_{1,k}$ represents the share of a lottery prize won in year t that is spent in year $t+k$. We also compute the cumulative responses as given by the sum of the $\beta_{1,k}$ values.

As previously, our sample consists of winners only. When we estimate the consumption response k years after winning, we include the consumption observations up to the year a household won. Hence $\beta_{1,k}$ is identified by comparing consumption k years after winning to consumption before winning. Moreover, we estimate equation (4) separately for each horizon k, including the years prior to winning to identify the fixed effects, instead of estimating the responses for all k jointly. The reason is that the latter would require we observe consumption for 11 straight years around winning, a restriction that greatly reduces our sample size since we do not impute consumption for households who move, dissolve, etc.

Figure 2 shows the dynamic responses of consumption, deposits, the sum of stocks, bonds, and mutual funds, and of debt. The top panel displays the flows and the bottom panel displays the cumulative effects. The estimated impact effect on consumption is the same as in column 2 of Table 3. For the balance sheet components, the estimated impact effects are 0.42 for saving in deposits, 0.06 for saving in stocks, bonds, and mutual funds, and about 0.07 for repayment of debt.

Beyond constituting quantitative moments to discipline economic models, the estimates provide four qualitative findings of particular interest. First, the 5-year cumulative response of consumption expenditure is almost 90 percent of the amount won, after which there is no observable effect of winning. This response contrasts with the textbook permanent income hypothesis, according to which a substantial share of a transitory income shock should be saved, also after five years. In Figure 2, we see how the remainder is spread across the balance sheet, primarily as deposits or repaid debt and to a smaller degree as stocks, bonds, and mutual funds.

Second, a substantial share of the prize-induced spending occurs immediately, as the consumption response drops from around 0.5 in the year of winning to around 0.2 in the following year. Thereafter, expenditure gradually reverts back to its pre-prize level. Deposits are used to support this consumption profile. They initially increase a great deal and are thereafter gradually depleted to finance the extranormal expenditure.[21] The sharp consumption decrease in the immediate year after winning might be due to durable goods purchases in the win-year, but the

[21] The pattern of deposits and stocks, bonds, and mutual funds is also consistent with recent evidence on savings from a surprise inheritance (Druedahl and Martinello 2020).

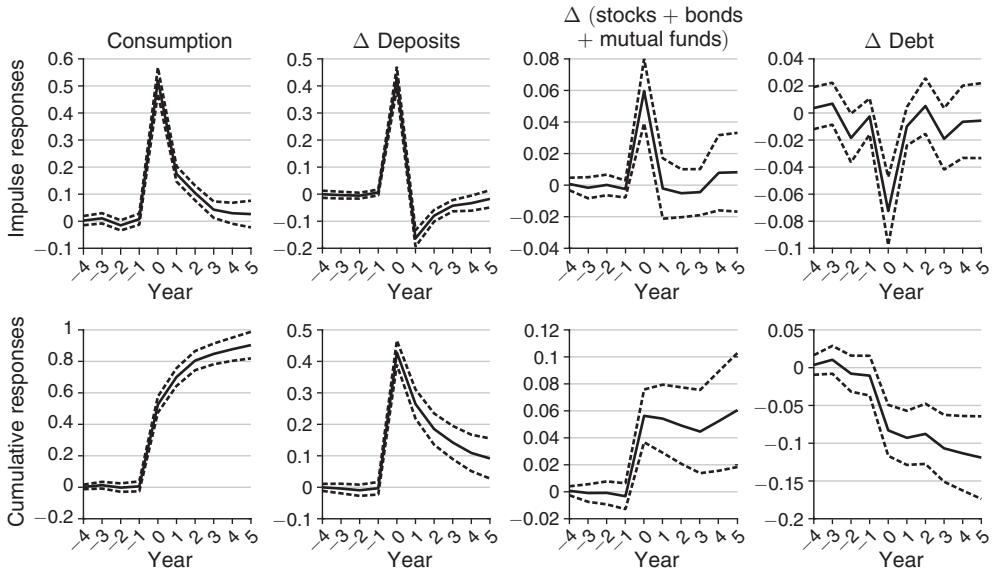


FIGURE 2. DYNAMIC HOUSEHOLD RESPONSES TO LOTTERY PRIZES

Notes: Each point is estimated as a separate regression of (4). Controls include time-fixed and household-fixed effects. Dotted lines represent 95 percent confidence intervals using robust standard errors clustered at the household level. The cumulative responses in the bottom panels are the sum of year-specific responses in the top panels. Standard errors of estimated cumulative effects are obtained through Monte Carlo simulations.

durables we observe, namely cars and boats, play a minor role here. When we zoom in on these durables, we find that in the years prior to winning, 11 percent of households purchase a car or boat by our measure, while in years of winning, the share rises to about 13 percent.[22] Among those 13 percent, the amount spent is approximately twice the amount won, on average. Our average estimate from specification (4) is that about 3 percent of lottery prizes are spent on car and boat purchases within the win-year.

Third, even though the consumption response drops rapidly from the year of winning to the next, it does not drop nearly as far as it would if households could be coarsely split into "savers" and "spenders." Such a distinction has been widely applied in macroeconomic models following the lead of Campbell and Mankiw (1989), in order to engineer a tight link between aggregate consumption and aggregate income. However, while spenders can generate a high average consumption response within the year of winning, their influence disappears after one period. Hence, the crude saver-spender model is inconsistent with Figure 2.[23] Auclert, Rognlie, and Straub (2018) discuss in more detail how such dynamic responses

[22] Purchases of cars and boats are not directly observed. Our data contain tax values of the vehicles a household owns. We consider a household as making a purchase whenever the value of its cars and boats increases. Hence, since we ignore transactions where households downscale their vehicle stock, our estimates along the extensive margin naturally underestimate the fraction making a purchase in a given year.

[23] Below, we find that small prizes typically are consumed entirely within the year of winning. Hence, for small windfalls, our results are more in line with a simplistic rule of thumb model.

can be used to distinguish between existing models. We also benchmark our results against a standard incomplete markets model in Section V.

Finally, we note that debt is repaid only within the year of winning, whereas deposits jump up and are thereafter gradually depleted. Given that debt typically comes with higher interest rates than deposits, the observed pattern is consistent with costs of altering amortization schedules.

When comparing our point estimates to existing evidence and models, one must bear two specific properties of our data in mind. First, we estimate the response of consumption *expenditure*, including durables as well as nondurables.[24] Generally, we would expect a greater MPC once durables are included. Second, we are studying responses within the entire year of winning, whereas theoretical models most commonly refer to a higher frequency. More specifically, if we assume that lottery prizes are uniformly distributed across the year, the average winner's within-year consumption response takes place over six months. In that perspective, our estimates can loosely be interpreted as a six-month MPC. In Appendix A.7, we combine our dynamic estimates with parametric assumptions to approximate MPCs at higher frequencies.[25]

Even after adjusting for time aggregation, our estimated consumption response is high relative to the predictions from standard models of household behavior. Broadly speaking, the class of models applied in macroeconomic research typically suggest a contemporaneous marginal propensity to consume out of transitory income shocks between 0.05 and 0.25 for nondurable consumption. For instance, the complete market infinitely lived household model suggests a nondurables MPC somewhere below 0.05, while the standard life cycle model suggests that the MPC is low for young households (smaller than 0.05) (see, for instance, discussion in Carroll 2001) but increases steadily with age. In the upper end of model-implied MPCs are the quarterly responses in Kaplan and Violante (2014), which lie around 0.25 for an unanticipated income shock of the same size as the 2001 US tax rebate. In contrast, our estimates are *not* large compared with the findings of existing empirical studies on transitory income shocks. For example, summarizing the findings from studies of US tax rebates, Parker et al. (2013) estimate that within a year, households spend between 0.50 and 0.90 per dollar received on total consumption.

[24] The aggregate share of durables in total household consumption expenditure lies around 15 percent in Norway over the sample period we study. We can isolate two components of durable consumption, cars and boats, and exclude them from our consumption expenditure measure. When we redo the estimates in Table 3 after this adjustment, the point estimates change by less than 3 percentage points. Note also that our measure of consumption expenditure includes imputed housing service flows from owner-occupied housing, but because we exclude movers from the sample, these housing service flows will not affect the estimated consumption responses to lottery prizes.

[25] In the existing literature, only Kaplan, Moll, and Violante (2018) report six-month MPCs. Their model predicts a nondurable consumption response around US$300 after a transfer of US$1,000.

III. MPC Heterogeneity

We now turn to addressing how MPCs vary with observable household characteristics and the size of income shocks. We first qualitatively establish which observables are associated with cross-sectional MPC variation, before we quantify their influence.

A. *Which Observables Matter for MPC Variation?*

We modify our benchmark specification to allow for interaction effects between prizes and explanatory variables that economic models suggest are important for MPCs. The specification is as follows:

$$(5) \qquad C_{i,t} = \beta_0 + \beta_1 lottery_{i,t} + \beta_2 lottery_{i,t} \times Z_{i,t-1} + \beta_3 Z_{i,t-1} + \alpha_i + \tau_t + u_{i,t},$$

where $Z_{i,t-1}$ contains variables we expect might correlate with MPCs. The interacting variables we consider are the amount won, liquid assets, income, net wealth, debt, education, the share of wealth held in risky assets, household size, and age. All these controls are lagged, except age, to avoid reverse causality. The term β_2 is our coefficient of interest, revealing whether the respective variable systematically varies with the consumption response to winning a prize of size $lottery_{i,t}$.

We estimate equation (5) both with each interaction term included in separate regressions (denoted "Univariate" in Table 4) and with all interactions included in a multivariate regression ("Multivariate" in Table 4). The latter estimates are our primary interest, as they indicate which factors affect MPCs directly, over and above their correlation with the other explanatory variables. To facilitate interpretation of the magnitudes involved, each column also contains in square brackets the regression coefficients when we standardize the interaction variable $Z_{i,t-1}$ by dividing by its cross-sectional standard deviation.

Among the candidates considered, prize size, liquid assets, and age stand out as the main observable factors associated with the magnitude of households' MPCs. All three are statistically significant both in separate regressions ("Univariate") as well as when every interaction term is included ("Multivariate"). Neither income, net wealth, debt, education, the portfolio share held in risky assets, nor household size has similar significant effects.

Column 2 ("Multivariate") further shows that when all variables are controlled for, a 1 standard deviation increase in prize size, liquidity, or age is associated with an MPC decrease of 0.7, 9.5, or 8 percentage points, respectively. As will become clearer when we zoom in on the importance of these variables below, this small effect of prize size is due to our estimator weighting winners of high prizes relatively more, whereas MPCs vary more with prize size among winners of small amounts.

As detailed in Section IB, imputed consumption is measured with error, where heterogeneous portfolios and intra-year trading of risky assets are the main concerns when studying responses to exogenous income shocks. We therefore redo our estimation after omitting all winners who have ever held stocks, bonds, or mutual funds, leaving us with about 40 percent of the original sample. While a selected

Table 4—The MPC out of Lottery Prizes: Interaction Effects

	Univariate		Multivariate		Multivariate, no-risky-assets	
	(1)		(2)		(3)	
$\text{Lottery}_t{}^2$	−0.001	[−0.007]	−0.001	[−0.007]	−0.001	[−0.007]
	(0.000)		(0.000)		(0.001)	
$\text{Lottery}_t \times \text{liquid assets}_{t-1}$	−0.003	[−0.109]	−0.003	[−0.095]	−0.003	[−0.109]
	(0.001)		(0.001)		(0.002)	
$\text{Lottery}_t \times \text{income}_{t-1}$	0.001	[0.016]	−0.002	[−0.026]	−0.005	[−0.057]
	(0.002)		(0.002)		(0.004)	
$\text{Lottery}_t \times \text{net wealth}_{t-1}$	−0.000	[−0.019]	0.000	[0.035]	0.000	[0.033]
	(0.000)		(0.000)		(0.000)	
$\text{Lottery}_t \times \text{debt}_{t-1}$	0.001	[0.061]	0.000	[0.022]	0.002	[0.083]
	(0.000)		(0.001)		(0.001)	
$\text{Lottery}_t \times \text{education}_t$	0.014	[0.037]	0.007	[0.019]	0.005	[0.013]
	(0.007)		(0.007)		(0.013)	
$\text{Lottery}_t \times \text{risky share}_{t-1}$	−0.006	[−0.001]	−0.046	[−0.009]	.	.
	(0.073)		(0.075)		.	.
$\text{Lottery}_t \times \text{household size}_t$	0.034	[0.037]	0.025	[0.028]	0.037	[0.041]
	(0.017)		(0.018)		(0.029)	
$\text{Lottery}_t \times \text{age}_t$	−0.005	[−0.082]	−0.005	[−0.076]	−0.005	[−0.071]
	(0.001)		(0.001)		(0.002)	
Observations	93,631	93,631	93,631	93,631	40,859	40,859

Notes: This table displays the estimated interaction terms from the estimation of equation (5). The univariate column presents the interaction coefficients when we include each interaction term with the lottery prize ($lottery_{i,t} \times Z_{i,t-1}$) individually in separate regressions (together with the remaining variables in equation (5)), while the multivariate columns present results when we include all interaction terms in the same regression. "No-risky-assets" refers to the sample where we keep only households that held no risky assets at any time during our sample period. Controls include the interacted variables, time-fixed effects, and household-fixed effects (see equation (5)). The sample sizes (93,631 or 40,859) include the winners and all observations of the winners in the years prior to winning. The standard errors in parentheses are robust and clustered at the household level. Regression coefficients where we standardize the interaction variable by its own cross-sectional standard deviation are reported in square brackets.

sample, it has the benefit that the main source of measurement error in imputed consumption is negligible. The third column of Table 4 shows that our results regarding MPC heterogeneity remain essentially the same within this subsample.

B. *Quantifying the Role of Shock Size, Liquidity, and Age*

We now zoom in on the influence of the three variables that stood out in Table 4. To this end, we group winners by the amount won, liquid asset holdings, and age. Stratification is conducted by quartile, and we estimate MPCs within strata.[26] There are two main reasons why we pursue this approach in addition to the interaction regressions above. First, we obtain directly interpretable estimates of each variable's co-movement with MPCs. Second, the approach will unveil nonlinearity in the importance of each covariate. Table 5 displays estimates when households are

[26] Ideally, we would have chosen an even more fine-grained stratification, but we are constrained by sample size when groups become too small.

Table 5—Heterogeneous Household Responses. Quartiles of Lottery Prize Size, Liquid Assets, and Age

	(1)	(2)	(3)	(4)
Panel A. Lottery prize size quartile				
Consumption	1.311	0.968	0.693	0.512
	(0.191)	(0.089)	(0.047)	(0.019)
Deposits	0.566	0.368	0.467	0.417
	(0.129)	(0.060)	(0.032)	(0.018)
Stocks, bonds, and mutual funds	0.021	0.081	0.007	0.060
	(0.056)	(0.037)	(0.015)	(0.009)
Debt	0.281	0.200	0.002	−0.080
	(0.147)	(0.068)	(0.038)	(0.011)
Observations	23,120	23,368	24,624	22,519
Panel B. Liquid asset quartile				
Consumption	0.621	0.528	0.466	0.455
	(0.041)	(0.047)	(0.038)	(0.049)
Deposits	0.323	0.361	0.479	0.599
	(0.035)	(0.041)	(0.035)	(0.037)
Stocks, bonds, and mutual funds	0.022	0.049	0.067	0.065
	(0.006)	(0.016)	(0.022)	(0.023)
Debt	−0.087	−0.134	−0.061	−0.011
	(0.025)	(0.033)	(0.022)	(0.018)
Observations	23,184	23,308	23,724	23,415
Panel C. Age quartile				
Consumption	0.565	0.555	0.523	0.441
	(0.037)	(0.038)	(0.033)	(0.031)
Deposits	0.344	0.385	0.409	0.539
	(0.035)	(0.028)	(0.032)	(0.031)
Stocks, bonds, and mutual funds	0.066	0.066	0.058	0.045
	(0.016)	(0.02)	(0.015)	(0.015)
Debt	−0.074	−0.109	−0.096	−0.017
	(0.027)	(0.022)	(0.020)	(0.012)
Observations	18,438	23,334	24,745	27,114

Notes: Each coefficient represents a separate regression of our baseline specification (3) within quartiles of the relevant variable. For lottery prize size, the cutoffs between the quartiles are US$2,070, US$5,200, and US$8,300, and the cutoffs between the age quartiles are 39, 51, and 63. The cutoffs between liquidity quartiles increase over time. In 1994 (2006), the cutoffs are US$720 (1,720), US$3,140 (7,160), and US$10,710 (21,600). The number of observations in each age quartile varies because the stratification is conducted at the time of winning and younger households are observed less frequently in the years before they won. Controls include time-fixed and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

grouped by each of the three covariates separately, while Figure 3 displays estimates when households are stratified by two out of the three dimensions at once.[27]

[27] As Table 5 reveals, quartiles differ by number of observations. The reason is that each winner is observed repeatedly, and some households are observed more frequently than others in our sample.

Shock Size.—Panel A of Table 5 presents estimates conditional on the amount won. The prize quartiles are US$1,100–2,070, US$2,070–5,200, US$5,200–8,300, and US$8,300–150,000.[28] In the top left corner, we see that winners of relatively small amounts tend to spend more than the prize itself. Moving down the table, we also see that the average debt response in this group is positive, suggesting that low-prize winners top their prize up with credit or lower debt repayment. However, all estimates in the lowest prize quartile come with relatively high standard errors, so the exact point estimates in this group should be interpreted with caution.[29]

Moving rightward in the prize distribution, we see that the consumption response declines monotonically with the amount won. The point estimate is approximately halved from the bottom to the top quartile. Still, the estimated response is remarkably high even in the top quartile, indicating that winners of more than US$8,300 spend on average about half their prize within the year of winning. The bottom row implies that only winners of relatively large amounts tend to cut debt. There is no clear monotonic pattern for the other asset classes.

A variety of theoretical mechanisms might explain why MPCs decrease with shock size. First, standard buffer stock saving models imply that the policy function for consumption is concave in wealth, so that MPCs are smaller for greater income innovations, either because of borrowing constraints (Carroll, Holm, and Kimball 2021) or risk (Carroll and Kimball 1996). Second, if purchases of high-return assets involve discrete transaction costs, then the rate of return on saving effectively increases with the amount won and thus motivates high-prize winners to save more. Third, several consumption decisions, such as purchases of durables, contain an element of discrete choice. In principle, this might explain why winners of small prizes seem to spend their entire income innovation, and even more by borrowing, whereas for larger prizes, such lumpy purchases could be less likely to dominate spending responses.[30] In Section V, we benchmark our results to a model where the first two mechanisms are present but not the third.

Liquidity.—Consider next the estimates conditional on liquid assets held at the end of the year before winning, presented in panel B of Table 5. Here the cutoffs defining quartiles are year-specific. The cutoffs in 1994 (2006) are US$720 (1,720), US$3,140 (7,160), and US$10,710 (21,600).

[28] When reading the estimates by size, two technical points about our estimator should be noted. First, within the lower quartile, the prize won (the "treatment") is low relative to all other factors that affect imputed consumption, rendering estimates imprecise. In higher quartiles, the treatment is stronger, and estimates are more precise. Second, our estimate in the highest prize quartile lies just below our pooled benchmark estimate in Table 3, whereas the three other within-quartile estimates are considerably higher than the benchmark. This reflects that because the prize distribution has a long right tail (see Figure 1), the linear specification behind Table 3 necessarily is drawn toward the MPC of high-prize winners. We illustrate this point in more detail in Appendix A.4.

[29] Median estimates from quantile regressions, previously referred to as LAD, are less sensitive to extreme consumption responses. In the lower size quartiles, this makes a substantial difference. The LAD estimates for consumption in the four quantiles are 0.65, 0.59, 0.54, and 0.48, respectively. If we drop owners of stocks, bonds, and mutual funds from our sample to minimize errors from capital gains, our OLS estimates across the size quartiles are 1.04, 0.92, 0.7, 0.49.

[30] On the other hand, discrete expenditure choice might in principle work in the opposite direction too, as the probability of actually making a discrete purchase increases with shock size.

Again, we see the negative relationship between MPCs and liquidity. The within-year consumption response is 0.62 in the low-liquidity quartile, gradually falling to 0.46 in the high-liquidity quartile. Among the three main savings vehicles considered, deposits track the consumption pattern most closely. The propensity to save in deposits increases from 0.33 among the least liquid to 0.59 among the most liquid. The propensity to save in stocks, bonds, and mutual funds also increases with liquidity but to a weaker extent than is the case for deposits. Debt stands out with an opposite pattern, as the least liquid winners tend to use more of their prize to repay debt than the most liquid do. One natural explanation is that households with high initial liquid asset levels were able to repay debt already before they won.

The association between liquidity and MPCs is consistent with previous empirical studies such as Misra and Surico (2014) and Leth-Petersen (2010) and with standard economic theory. However, the MPC is remarkably high even among the most liquid households, and this finding will prove hard to match with conventional buffer stock saving models, as we highlight in Section V.

Age.—Panel C in Table 5 displays estimates by age quartiles. The quartile cutoffs are at ages 39, 51, and 63. We see that the within-year expenditure response declines from 0.58 in the youngest quartile to 0.44 in the oldest quartile. Similarly, saving in deposits increases from 0.34 among the youngest to 0.53 among the oldest. There is no clear pattern for the response of saving in stocks, bonds, and mutual bonds, or debt repayment.

The tendency for MPCs to fall with age points toward mechanisms emphasized in the recent literature on life cycle models. Coupled with borrowing constraints, a realistic earnings profile will typically motivate high MPCs early in life. Among the old, it is well known that savings are remarkably high, which has led the structurally oriented literature to emphasize bequest as a luxury good (De Nardi 2004, De Nardi and Fella 2017) or precautionary savings for health expenditures (De Nardi, French, and Jones 2010). The finding that MPCs decrease with age even between the older age groups is consistent with a preference to increase the budget share of savings as one grows richer. We convey these points in more detail with a structural model in Section V.

Conditional Heterogeneity.—Prize size, liquidity, and age might be correlated, and unlike in Table 4, the estimates in Table 5 do not isolate the contribution from each variable alone. We now move on to explore how much the three variables matter conditional upon one another. To this end, we first stratify households by prize size and liquid assets, then by prize size and age, and finally age and liquid assets. Each time, we estimate an MPC within 16 (4×4) subgroups. By construction, the 16 different groups need not be equally large, as we group households by the quartile of the respective dimensions separately. The distributions across strata and the MPC estimates within them are displayed in Figure 3.[31]

[31] Point estimates with standard errors are provided in the Appendix (Figure A7). There we also provide the analysis when stratifying by Net illiquid and liquid assets (Figure A8).

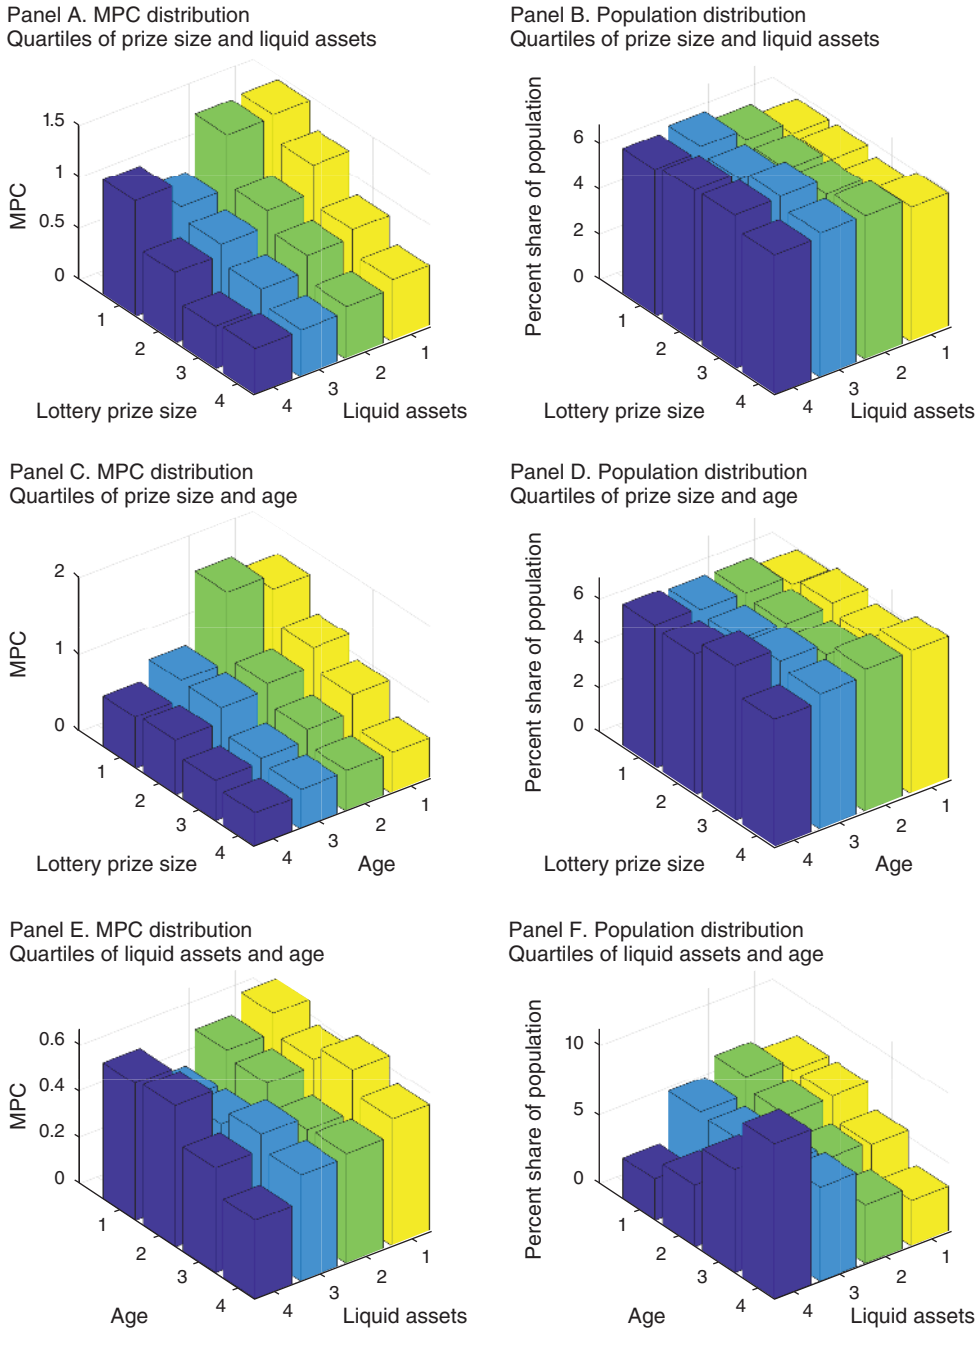


FIGURE 3. HETEROGENEOUS CONSUMPTION RESPONSES. QUARTILES OF LOTTERY PRIZE SIZE, LIQUID ASSETS, AND AGE

Panel A shows that within each liquidity quartile, the MPC decreases with the amount won and that within each size quartile, the MPC decreases with liquidity. The MPC is almost 1.5 among winners who are in the lowest liquidity quartile and the lowest prize size quartile simultaneously (northern corner of the diagram). In

contrast, the MPC is below 0.5 in the high-liquid-high-prize group (southern corner). Panel B shows that the strata are essentially equi-sized, consistent with the results in Table 2. Panel C displays MPC estimates across the joint distribution of prize size and age. Within all age groups considered, consumption responses decline with prize size. Within all prize quartiles, responses decline with age. Panel D reveals that the distribution is flat in the two dimensions, again consistent with Table 2. Panel E displays MPC estimates across the joint distribution of age and liquid asset holdings. Here there is no clear-cut pattern. We see why in panel F. Age and liquidity are correlated, preventing us from uncovering more than the average effects we already saw in Table 4.

IV. Robustness

We next extend our analysis in several dimensions. First, we provide a range of estimates for consumption and balance sheet responses by applying alternative assumptions in the consumption imputation procedure and by modifying our household definition to include adult offspring. Thereafter, we extend our analysis of MPC heterogeneity.

A. *Consumption Imputation and Household Definition*

As discussed in Section IA, the main challenge for accurate consumption imputation is to adjust the observed changes in net wealth for capital gains at the household level. As explained in Section IB, our baseline approach imputes capital gains by assuming that each household holds the market portfolio and that all transactions happen at the end of the year.[32] We here address the challenge posed by capital gains with two alternative strategies.

First, we drop all households who hold risky assets (stocks, bonds, or mutual funds), focusing on households with balance sheets consisting of bank deposits, debt, and housing wealth, where housing wealth is the only asset class experiencing capital gains. However, as explained in Section IB, capital gains on housing are unproblematic since our baseline sample excludes households who move so that errors from imputing capital gains should be minimal. Interest received and paid on the remaining two components are perfectly measured in our data, leaving us with a sample of households for which concerns with measurement error in our imputed measure of consumption will be minimal.

Second, we conduct our analysis on a different period, 2005–2011, for which we observe direct stockholding from the stockholder registry (Statistics Norway 2015f). This registry records individuals' holding of single stocks and therefore allows us to compute capital gains contingent on each individual's actual stock market portfolio rather than relying on the assumption that all individuals hold the market portfolio. The reason we did not focus on this subsample in our main

[32] For example, this implies that capital gains on stock wealth in year t are computed as stock wealth at the last day of year $t-1$ multiplied by the year-t capital gains rate at Oslo Stock Exchange.

Table 6—Robustness. Consumption and Balance Sheet Responses under Alternative Assumptions and Sample Restrictions

	Benchmark (1)	No risky assets (2)	Late sample (3)	Late sample incl. stockholder registry (4)	Incl. offspring responses (5)
Consumption	0.524	0.502	0.517	0.508	0.512
	(0.017)	(0.026)	(0.047)	(0.059)	(0.019)
Deposits	0.422	0.479	0.493	0.494	0.459
	(0.016)	(0.024)	(0.039)	(0.039)	(0.031)
Stocks, bonds, and mutual funds	0.058	.	0.041	0.040	0.063
	(0.008)	.	(0.017)	(0.017)	(0.010)
Debt	−0.073	−0.075	−0.077	−0.085	−0.124
	(0.010)	(0.014)	(0.027)	(0.027)	(0.042)
Observations	93,631	40,859	6,888	6,888	92,461

Notes: Each column represents a separate regression of equation (3). Controls include time-fixed and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

analysis is that in 2005, the reporting of lottery prizes changed from a minimum of NOK10,000 to a minimum of NOK100,000.

Table 6 presents consumption and balance sheet responses under alternative imputation procedures. Column 1 first shows our preferred specification from the main body of the paper. Next, in column 2, we present results for the sample of households who do not hold risky assets. We see that the consumption and balance sheet responses are almost identical in the two columns, with the exception that instead of acquiring risky assets, the sample of households with no risky assets save more in deposits.

Column 4 presents the results when we utilize information from the stockholder registry to compute capital gains. Note that because we can access the stockholder registry only after 2005, these estimates are not directly comparable to the baseline in column 1. We therefore also estimate household responses on this sample period using our baseline imputation method, reported in column 3. The main takeaway is that the results are similar for either imputation method. Overall, Table 6 implies that our results are not driven by biases due to unobserved capital gains.

Next, we explore the extent to which our estimated MPCs are biased by transfers to offspring. The motivation is that our estimated expenditure responses might include unreported transfers to adult offspring who subsequently save them. To investigate how important this channel is, we construct variables at the extended household level. We here include the offspring's consumption, deposits, stocks, bonds, mutual funds, and debt in the lottery winners' consumption and balance sheet variables. For example, this means that the new "household"-level consumption measure is the sum of expenditure by lottery winners and all their children above the age of 18.

Column 5 in Table 6 reports the results from including offspring responses in our analysis.[33] Compared to the baseline estimates in column 1, the consumption

[33] The number of observations decrease from column 1 to column 5 because we drop households where the offspring have missing consumption observations in the year the parents win.

TABLE 7—ROBUSTNESS. INTERACTION RESULTS UNDER ALTERNATIVE ASSUMPTIONS AND SAMPLE RESTRICTIONS

	Benchmark		Normalizing by permanent income		Controlling for individual-fixed MPC	Untrimmed sample		Unconditional 2.5 percent trim on consumption	
	(1)		(2)		(3)	(4)		(5)	
Lottery_t	0.524	0.887	0.563	0.680	0.870	0.712	0.716	0.353	0.662
	(0.017)	(0.092)	(0.024)	(0.174)	(0.093)	(0.017)	(0.091)	(0.011)	(0.068)
Lottery_t^2		−0.001		−0.000	−0.001		0.001		−0.001
		(0.000)		(0.007)	(0.000)		(0.000)		(0.000)
Lottery_t		−0.003		−0.042	−0.003		−0.002		−0.004
$\times$ liquid assets_{t-1}		(0.001)		(0.031)	(0.001)		(0.001)		(0.001)
$\text{Lottery}_t \times \text{income}_{t-1}$		−0.002		−0.020	−0.002		0.000		−0.000
		(0.002)		(0.104)	(0.002)		(0.001)		(0.001)
Lottery_t		0.000		0.008	0.000		−0.000		0.000
$\times$ net wealth_{t-1}		(0.000)		(0.006)	(0.000)		(0.000)		(0.000)
$\text{Lottery}_t \times \text{debt}_{t-1}$		0.000		−0.019	0.000		−0.001		0.000
		(0.001)		(0.014)	(0.001)		(0.000)		(0.000)
$\text{Lottery}_t \times \text{education}_t$		0.008		0.010	0.008		0.007		0.003
		(0.007)		(0.009)	(0.008)		(0.007)		(0.005)
Lottery_t		−0.046		−0.036	−0.033		0.023		−0.066
$\times$ risky share_{t-1}		(0.075)		(0.083)	(0.074)		(0.085)		(0.056)
Lottery_t		0.025		0.023	0.017		0.030		0.008
$\times$ household size_t		(0.018)		(0.020)	(0.019)		(0.015)		(0.011)
$\text{Lottery}_t \times \text{age}_t$		−0.005		−0.003	−0.005		−0.003		−0.003
		(0.001)		(0.020)	(0.001)		(0.001)		(0.001)
$\text{Lottery}_t \times \text{MPC}_i$					0.016				
					(0.007)				
Observations	93,631	93,631	61,190	61,190	87,968	120,139	120,139	94,113	94,113

Notes: Each column represents a separate regression of equation (3) or (5). Controls include interacted variables, time-fixed effects, and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

response is slightly lower, while the deposits and debt responses are somewhat stronger. Hence, the estimates suggest that lottery winners do transfer resources and that some of it is saved, but these effects are not quantitatively large.

B. *Heterogeneity*

We now scrutinize our finding that liquidity, age, and shock size are the main variables associated with MPC variation. We here focus on the results from variations of the interaction regression specified in equation (5), as this provides the most compact exposition of our findings. Results are reported in Table 7. For each alternative, the table reports MPC estimates and interaction effects from variations of specifications (3) and (5) next to each other.

From consumption theory, one could argue that all variables should be normalized by permanent income in specification (3); see, for instance, Carroll (forthcoming). Column 2 of Table 7 provides estimates under this modification. As detailed in Appendix A.9, permanent income is computed as a combination of age, time, and age-time fixed effects, together with a household's observed average income relative to its cohort. The latter aims to capture the permanent component of a household's

ability to earn, given age and time period.[34] Note that after normalization, the mean MPC estimate has the same quantitative interpretation as in our baseline, but this does not apply to the interaction coefficients.[35] Also, much of the variation in the interaction terms will now be driven by income variation.

Although not statistically significant, the results imply that after normalization, liquidity-to-permanent-income is the main household characteristic relevant for consumption responsiveness. The age coefficient is practically unaltered but less precisely estimated. Prize relative to income, in contrast, becomes unimportant after normalization. One reason is that when squared (as in the interaction term), variation in prize-to-permanent-income is dominated by variation in permanent income.

In column 2, we also note the coefficient on flow income normalized by permanent income. If the directly observed cross-sectional income variation is driven by income's permanent component, then theory predicts that flow income divided by permanent income matters negatively for MPCs (but not flow income alone). We return to this point in Section V. We see that the point estimate indeed is negative but statistically insignificant.

Next, we control for an estimate of each household's typical co-movement between consumption and labor income, denoted MPC_i in the table. Our objective here is to distinguish the influence of observed variables from individual-fixed effects in the marginal propensity to consume, for instance, due to persistent differences in patience. Had we observed households winning each year, we could have done this more directly, by controlling for an individual-fixed slope coefficient in the relation between consumption and prize.[36] Our approach is explained in Appendix A.10. In short, we (i) exclude the year in which a household wins and exclude households observed for less than five years; (ii) purge consumption and labor income of time-fixed effects; (iii) estimate an individual-specific MPC from labor income as $MPC_i = \text{cov}\left(\epsilon^c_{i,t}, \epsilon^y_{i,t}\right)/\text{var}\left(\epsilon^y_{i,t}\right)$, where $\epsilon^c_{i,t}$ and $\epsilon^y_{i,t}$ are the consumption and income residuals from step (ii).[37]

Column 3 in Table 7 shows the results when controlling for MPC_i. The estimated coefficient on MPC_i confirms the presence of unobserved persistent factors that raise both households' marginal propensity to consume out of income in general and out of lottery income in particular. Hence, nonsituational effects do seem relevant for understanding cross-sectional variation in MPCs. But more important for our purposes is the following finding: even after controlling for this persistent component of individual MPCs, the same three observables as before remain significant for cross-sectional MPC variation.

As emphasized in Section IE, we omitted extreme consumption observations from our baseline sample so as to prevent outliers from driving our inference. Column 4

[34] Age and cohort here refer to the households' main earner. Simply using average observed income, without controlling for time and age effects, would lead to estimates primarily driven by the age and time at which a household is observed.

[35] Both the left-hand-side variable (consumption) and prize are divided by income in the same way, whereas the interacting variables necessarily are divided by income squared. For instance, both lottery and liquidity are each divided by income.

[36] Note the difference between individual-fixed effects in consumption *level*, which we control for with α_i in equation (5), and individual-fixed effects in consumption *sensitivity*.

[37] The time-fixed effects in step (ii) are taken out through the regressions $c_{i,t} = \gamma^c_t + \epsilon^c_{i,t}$ and $y_{i,t} = \gamma^y_t + \epsilon^y_{i,t}$.

of Table 7 shows that if we instead include the full sample of one-time winners, the mean MPC increases to 0.71. If we instead unconditionally trim the top and bottom 2.5 percent consumption observations, the mean estimate falls mechanically (to 0.35) because of how lottery prizes raise the consumption of high-prize winners. Importantly though, our main findings regarding heterogeneity stand firm. Prize size, liquidity, and age matter irrespective of sample selection.

Rather than including the sum of deposits, stocks, bonds, and mutual funds, we have also considered each of these subcomponents separately. The results are provided in Table A4 in Appendix A.11. The estimated interaction coefficients are negative for each asset type, but bonds and mutual funds are not statistically significant. The estimate for deposits is the same as in Table 4, while the point estimate for stocks (which only 30 percent of our sample hold) is higher yet imprecise. As the deposit estimate implies, all our results remain essentially unchanged if we instead define liquid assets as deposits only.

V. Contrasting Our Findings to the Predictions of a Standard Model

We next benchmark our empirical finding against a buffer stock saving model of the kind that is widely applied in the literature. We also discuss possible extensions to bridge the gaps that open between our estimates and the model's implications.

A Consumption-Saving Model.—Households live T periods. They face an exogenous borrowing constraint; idiosyncratic income risk; and a life cycle income profile given by τ_t, where τ_t is the deterministic income component at age t. They can save in two assets: a zero-interest liquid asset b and an interest-bearing illiquid asset a, which is subject to transaction costs. Households maximize the sum of discounted utility from consumption c and bequests $a_T + b_T$:

$$\max_{\{c_t\}_0^T} \sum_{t=0}^{T} \beta^t \frac{c_t^{1-\gamma}}{1-\gamma} + \beta^T \psi_0 \frac{\left(a_T + b_T + \psi_1\right)^{1-\gamma}}{1-\gamma}$$

subject to

$$b_t = b_{t-1} + r a_{t-1} + e^{y_t} - d_t - \kappa \mathbf{1}_{a_t \neq a_{t-1}} - c_t, \qquad a_t = a_{t-1} + d_t,$$

$$y_t = \tau_t + p_t + \epsilon_t^T, \quad p_t = \rho p_{t-1} + \epsilon_t^P, \quad \epsilon_t^P \sim N\left(0, \sigma_P\right), \quad \epsilon_t^T \sim N\left(0, \sigma_T\right),$$

$$a_t \geq 0, \quad b_t \geq 0,$$

where β is the discount factor, ψ_0 and ψ_1 are parameters of the bequest function, r is the return on the illiquid asset, y is log labor income, d is the flow of deposits into the illiquid asset, κ is the adjustment cost for the illiquid asset, p is the persistent component of income, ϵ^T is a transitory income shock, ρ is the persistence

of income, and ϵ^P is an innovation to permanent income. Our framework nests many of the common models in the incomplete markets literature.[38]

The model is annual. Households start life at age 30 with a draw from the ergodic income distribution and end their life at age 75. We use income data from Norway for the period 1967 to 2014 to extract the systematic component of income by age, assuming that time-fixed effects are orthogonal to a time trend as in Ahn and Holm (2020). Households start working at age 30 and earn an average income of 0.2 that grows to 0.37 at age 60. When retired, households earn 66 percent of maximum earnings, approximating the Norwegian public pension system. We further set r to 0.0406, equal to the average excess return on illiquid assets in Norway between 2005 and 2015.[39] Next, after taking out year, age, and education effects from earnings, we estimate the permanent-transitory income process using GMM. The parameters we find are $\rho = 0.937$, $\sigma_P^2 = 0.024$, and $\sigma_T^2 = 0.009$.

We estimate the remaining five parameters (γ, β, ϕ_0, ψ_1, and κ) using GMM to minimize the distance between five of our previously estimated interaction coefficients and their counterparts implied by the model. The moments are equally weighted.[40] From our estimates, we use the standardized coefficients for prize size, liquidity, income, net wealth, and age reported in Table 4 above. The corresponding model-based coefficients are obtained by simulating 100,000 households.[41] We randomly select an age at which each household receives a lottery prize drawn from a log-normal distribution similar to the actual distribution in our lottery data.[42] For each household, we compute consumption after winning the lottery prize (treatment) and the counterfactual consumption that the household would have optimally chosen had it not won the lottery prize (control). We then regress consumption on lottery prize using the same interaction specification (5) as we applied to the actual data.

Our estimation procedure yields the following parameter values: $\gamma = 1.1$, $\beta = 0.902$, $\psi_0 = 479.66$, $\psi_1 = 8.46$, and $\kappa = 0.2037$. We note that the transaction cost parameter, κ, is approximately $2/3$ of the average annual wage at age 40, which seems relatively high. On the other hand, the model implies a frequency of housing transactions that lies close to what we observe in the data.[43]

Model versus Data.—Table 8 compares the standardized regression coefficients from our model to the corresponding targeted estimates from the data, displayed previously in Table 4. The model generates coefficients of the same signs as in the

[38] For example, $\kappa = 0$ is a one-asset model and $\kappa = 0$ with no income risk and $\psi_0 = 0$ is the canonical life cycle model.

[39] The return on illiquid asset is set equal to the return on nonfinancial wealth minus the return on financial wealth from Fagereng et al. (2020, table 3).

[40] Our GMM approach proceeds as follows. In each step, we pick a vector of parameters, solve the model, simulate the experiment, and compute the moments from the model and the error function to be minimized. We first use a global minimizer to search for potential candidate vectors. Next, we use a local minimizer to search for the best solution around the candidates from the global search. We use NLOPT (https://nlopt.readthedocs.io/en/latest/) for minimization, with controlled random search (CRS2_LM) as the global method and BOBYQA as the local method.

[41] The codes are available upon request and based on Fernández-Villaverde and Valencia (2018).

[42] Lottery prizes are drawn from a truncated log-normal distribution with a maximum of 3 times annual income at age 60, approximating the lottery prize size distribution observed in the data.

[43] In our model, 9.85 percent of households are involved in housing transactions each year, compared with an annual transaction frequency of 9 percent in the data.

TABLE 8—INTERACTION COEFFICIENTS IN DATA AND MODEL. STANDARDIZED COEFFICIENTS

	Data	Model
Shock size	−0.007	−0.013
	(0.002)	
Liquid assets$_{t-1}$	−0.095	−0.041
	(0.024)	
Income$_{t-1}$	−0.026	−0.037
	(0.023)	
Net wealth$_{t-1}$	0.035	0.041
	(0.023)	
Age$_t$	−0.076	−0.090
	(0.021)	

Note: This table shows estimated standardized interaction coefficients from the data (Table 4 above) and corresponding regression coefficients from the simulated model.

data. In terms of magnitudes, we note that the model-implied coefficient for liquidity is smaller than our empirical estimate, whereas the opposite applies to the coefficient on shock size. The model-implied age coefficient is similar to the empirical estimate.

There are two correlations in Table 8 that are particularly challenging to generate in models. One is the negative association between age and MPC. A simplistic life cycle model would imply a positive relationship because older households have shorter horizons. Our model instead aligns with the data in this dimension because we have a plausible life cycle earnings profile together with a bequest motive. The age-earnings profile motivates young households to consume a large fraction of windfall gains to smooth consumption over life, while the bequest motive induces old households to save a large share of windfall gains for their offspring.

The second challenge is to account for the low correlation between MPCs and net wealth. In a standard one-asset model, low net wealth implies that a household is constrained and thus has a high MPC. This no longer holds in a two-asset framework. Here some households have high net wealth but hold little of their wealth in liquid assets, which in turn gives rise to high MPCs. Such households are commonly referred to as "wealthy hand-to-mouth" agents. In this way, the two-asset framework breaks the negative association between net wealth and MPC, once liquid assets are controlled for as in the regressions behind Table 8.

Note that our model predicts that MPCs decrease with current income. The reason is that cross-sectional income variation is driven by income differences relative to a common trend, while MPCs in models like ours decrease with current relative to permanent income. If instead the cross-sectional income variation were due to permanent income differences, the negative association between MPCs and current income would disappear. Hence, a potential explanation why directly observed income is weakly associated with MPCs in our estimates is that the cross-sectional income variation is largely explained by permanent income differences.

We also explore whether our model can match the dynamic consumption responses estimated in Section II. Figure 4 plots the response from the actual data

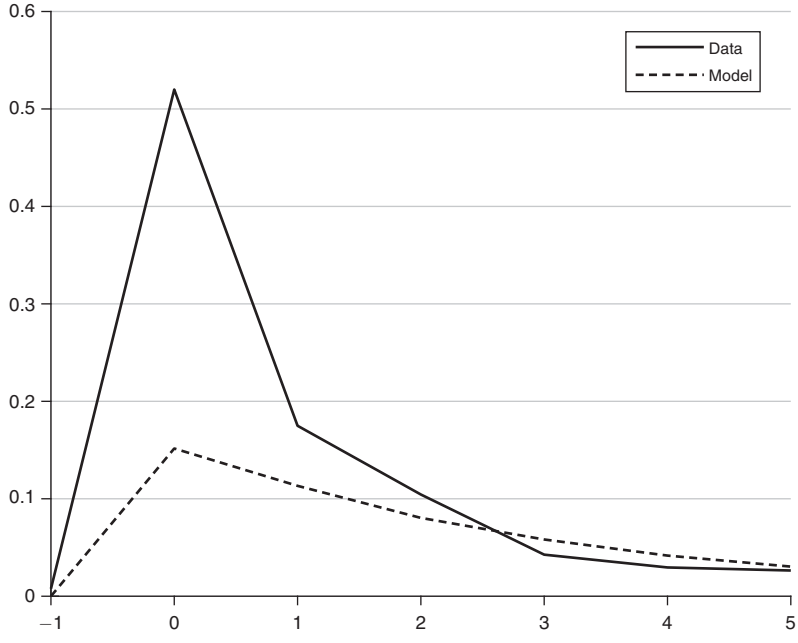


Figure 4. Dynamic Consumption Response in Model and Data

Note: The figure shows the estimated dynamic consumption response from the data (solid line) and from the simulated model (dashed line).

together with its counterpart from the simulated model. We see how our model fails to match the high *level* of our estimated within-year MPC. The MPC in year 0 is 0.15 in the model, significantly lower than our baseline empirical estimate of 0.52. The main reason why a model like ours will struggle to match the MPC level is that the lottery prizes are relatively large. Even in a framework with many constrained households, such as in our calibration, the actual lottery prizes are large enough to move most households sufficiently far away from the borrowing constraint that its influence on MPCs is limited.

Our model addresses nondurable consumption expenditure only. Our empirical approach, in contrast, combines durables and nondurables into one joint imputed consumption measure. In principle, this might explain why our model fails to match the estimated MPC level. Households could be smoothing their consumption of windfall gains by immediately purchasing durables, which they subsequently enjoy. However, a model with durable consumption goods will face two challenges in matching the empirical estimates. First, when the investment decision is modeled as a discrete choice due to the indivisible nature of durables, the probability of making a durable good purchase typically increases with liquid asset holdings. Hence, durable consumption choice will be a force toward positive correlation between liquidity and expenditure responses. Second, a model with durable consumption goods typically generates expenditure shifting between different time periods and high MPCs in the beginning. However, this need not

necessarily imply higher cumulated spending over a six-year period like in Figure 2. The presence of durable goods would raise the initial consumption response but also tend to lower expenditure responses thereafter. Hence, even with durables in the model, it would be challenging to match the sum of total spending over six years.

Similarly, our model ignores preference heterogeneity, which is a natural and realistic feature to include when assessing why MPCs differ across households. The route taken in the literature is typically to impose heterogeneity in patience or the intertemporal elasticity of substitution (IES). A recent example is Aguiar, Bils, and Boar (2020), where some households are rich, illiquid, and have high MPCs because their discount rate is low, while their IES is high. However, while such a model could successfully match the correlations we find between liquidity, prize size, and MPCs, it would still struggle to explain the high MPCs that we find among high-liquid households.

We believe that a reasonable summary of our model exercise is as follows. The main empirical patterns we find regarding MPC *heterogeneity* and response *dynamics* are consistent with models of precautionary saving under borrowing constraints, provided that liquid assets are distinguished from illiquid wealth and life cycle considerations are treated with some care. However, the base MPC *level* that we estimate lies well above what standard models typically imply. For a model to account for our findings, it must maintain the systematic heterogeneity and time-profile implied by the two-asset buffer stock framework but additionally contain a mechanism for high consumption responses.

VI. Conclusion

Applied macroeconomic research is increasingly emphasizing how micro-level heterogeneity matters for aggregate phenomena such as business cycles and policy transmission. A key ingredient in this research program is a detailed understanding of households' expenditure responses to income shocks and in particular, which household characteristics are associated with the cross-sectional variation in such responses. Different theories propose alternative factors, but including them all in structural models is infeasible. Our contribution lies here. We use detailed administrative data from tax and income records to identify the main variables that are systematically related to how households respond to unanticipated income shocks, as identified by lottery prizes. A weighted-average within-year propensity to spend out of a prize lies around one-half in our sample, but this estimate varies considerably with the amount won, predetermined household liquidity, and age. Winners of small to moderate prizes consume approximately everything within the calendar year of winning. The estimate is roughly halved when we move from the lowest to the highest quartiles of lottery prizes observed and falls by about one-fourth when we move from the least to the most liquid quartiles of households.

Importantly, the association between liquidity and MPCs is not matched by any observed household characteristic other than age. Neither income, net wealth, education, nor risky portfolio share correlates significantly with MPCs once liquid

wealth, age, and prize size are controlled for. In contrast, the liquidity, size, and age effects do not disappear when all these alternative explanatory variables are taken into account. While we cannot claim to control for all possible household characteristics that might correlate with both liquid wealth and consumption sensitivity, we do believe our combined evidence supports a situational interpretation of the estimated association between liquidity and MPCs. For instance, if heterogeneous impatience underlies the liquidity patterns we uncover, by causing both low holdings of liquid wealth and high willingness to consume windfall gains, we would expect net wealth to also correlate negatively with MPCs. However, it does not. Similarly, if heterogeneous risk aversion is driving the liquidity-MPC association, then we would expect the risky portfolio share to pick up some of this effect. It does not. Moreover, controlling for household-specific marginal propensities to consume out of regular labor income, estimated in the years households have not won prizes, does not alter our results either. This is not to say that nonsituational factors are irrelevant. To the contrary, our evidence shows that household-specific MPCs out of regular income do indeed matter. Our claim is that the effects of liquidity withstand as we control for a battery of such factors.

When it comes to shock magnitude, we find a clear pattern: MPCs decrease with the amount won. Qualitatively, this result fits well with standard consumption theory, where policy functions are concave due to borrowing constraints or risk. Quantitatively, however, such conventional precautionary savings motives imply a weak size effect for households who are some distance away from their relevant borrowing constraints. The MPCs we uncover decrease most sharply in size when prizes are low, consistent with the quantitative prediction from standard theory, but they decrease almost as markedly with size among high-liquid households as they do among low-liquid households. The latter result points toward mechanisms that are absent in the standard buffer stock models.

The age effect we estimate has the opposite sign of what the simplest life cycle model would imply. We find that MPCs decrease in age, whereas a straightforward argument based on households' time horizon would suggest the opposite. However, basic extensions of the life cycle model might explain why the age effect is negative. A realistic earnings profile together with borrowing constraints will tend to raise MPCs early in life, whereas nonhomothetic bequest-related preferences could explain the lower MPCs among the old. Both features are already common ingredients in state-of-the-art life cycle models. Indeed, our finding that older households respond relatively weakly to income shocks fits with the widespread observation of remarkably high saving rates among the elderly.

The base level of MPCs that we estimate, in particular the level among high-liquid winners, lies well above what standard models would typically imply. Hence, illiquidity alone does not underly the high average expenditure responses to income shocks that we find. On the other hand, as we show, the patterns we find regarding the time profile and heterogeneity of expenditure responses are consistent with models of precautionary savings under borrowing constraints, provided that liquid assets are distinguished from illiquid wealth as in the many recent studies and life cycle considerations are treated with some care. Hence, a framework that would successfully account for our findings is one that lifts the base MPC level without distorting the

systematic relationship between household characteristics and consumption dynamics implied by a standard model.

Appendix

A.1 *Comparing Imputed Consumption with the National Accounts*

In the paper, we use imputed consumption expenditures to represent consumption. In this appendix, we compare our measure of imputed consumption with consumption from the national accounts. Figure A1 shows the time series of average real imputed consumption per person in the national accounts and imputed from registry data. The two data sources have similar trends, but imputed consumption is more volatile than the national accounts. One reason why aggregate imputed consumption is more volatile is that the sample for imputing consumption may be different from year to year due to the sample restrictions imposed to impute consumption in Section IB. Overall, imputed consumption follows the same trajectory as consumption in the national accounts.

A.2 *Alternative Specifications*

Existing studies that estimate consumption responses to income shocks utilize a variety of slightly different econometric specifications. We here present estimates from the three main specifications considered in this literature, together with our benchmark specification (3) utilized in the main text. The alternative specifications are

$$\Delta C_{i,t} = \beta_0 + \beta_1 \mathit{lottery}_{i,t} + \beta_2 X_{i,t-1} + \alpha_i + \tau_t + u_{i,t}, \tag{6}$$

$$C_{i,t} = \beta_0 + \beta_1 \mathit{lottery}_{i,t} + \beta_2 X_{i,t-1} + \beta_3 C_{i,t-1} + \alpha_i + \tau_t + u_{i,t}, \tag{7}$$

where Δ is the one-year difference operator. Equation (6) is the difference estimator, while equation (7) is the dynamic estimator (including lagged consumption). Our coefficient of interest is still β_1.

A.3 *OLS Weights*

In this appendix, we derive how the distribution of lottery prizes affects our average estimated MPC. Assume that we have N observations of lottery prizes, $l_1, \ldots, l_N$, and consumption, $c_1, \ldots, c_N$. The estimated β on the full sample is defined as

$$\beta_1 = \frac{\sum_{i=1}^{N} l_i c_i}{\sum_{i=1}^{N} l_i^2}.$$

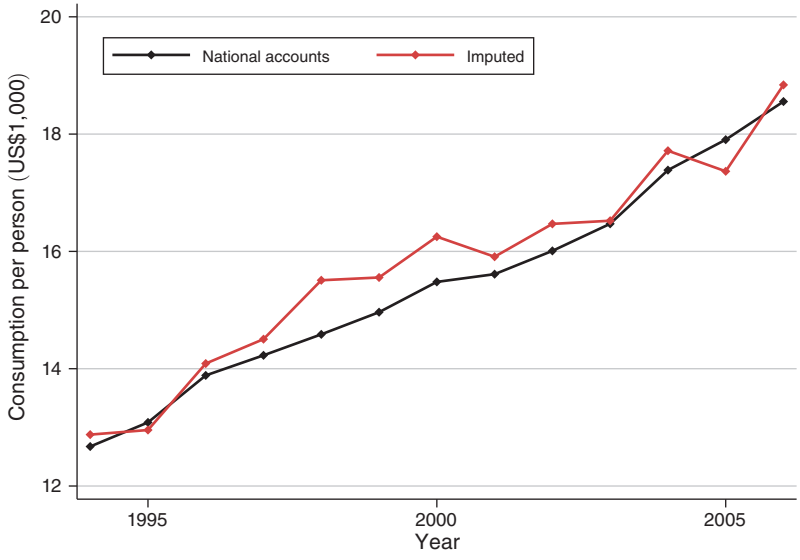


FIGURE A1. CONSUMPTION PER PERSON. IMPUTED AND NATIONAL ACCOUNTS. 1994–2006

Notes: The graphs show consumption per person in US$1,000 between 1994 and 2006. For the national accounts, we take household consumption at current prices, divide it by the population of Norway, CPI adjust it to the year 2000, and convert to US$1,000 using the year-2000 exchange rate. For imputed consumption, we take imputed household consumption + housing services at current prices, divide it by the size of the household, take the mean across years, CPI adjust to the year 2000, and convert to US$1,000 using the year-2000 exchange rate.

TABLE A1—THE MPC OUT OF LOTTERY PRIZES

	Specifications:		
	(1)	(2)	(3)
Levels (OLS)	0.587	0.524	0.520
(Observations = 93,631)	(0.013)	(0.017)	(0.017)
Differences (OLS)	0.504	0.504	0.493
(Observations = 59,909)	(0.018)	(0.025)	(0.025)
Levels (LAD)	0.609	0.486	.
(Observations = 93,631)	(0.015)	(0.018)	.
Differences (LAD)	0.467	0.460	.
(Observations = 59,909)	(0.015)	(0.024)	.
Dynamic (OLS-IV)[a]	0.506	.	0.503
(Observations = 59,909)	(0.021)	.	(0.020)
Time-fixed effects	Yes	Yes	Yes
Household-fixed effects	No	Yes	Yes
Additional controls	No	No	Yes

Notes: Each coefficient shows the estimated consumption response to a lottery prize. Levels correspond to the benchmark specification (3), differences correspond to specification (6), and dynamic corresponds to specification (7). Additional controls are age, age^2, age^3, age^4, household size, household size^2, number of children under 18, and $t-1$ values of after-tax income, liquid assets, debt, net wealth, and risky asset share. The sample size 93,631 (59,909) includes the 23,728 (14,743 if we restrict to also observing consumption in year $t-1$) winners and all observations of winners in the years prior to winning. LAD = Least Absolute Deviation estimator (Median). OLS-IV = Arellano-Bover/Blundell-Bond estimator (Arellano and Bover 1995, Blundell and Bond 1998). The standard errors in parentheses are robust and clustered at the household level (OLS).

[a]We do not include household-fixed effects in the OLS-IV estimator.

We can also define individual specific $\beta_{1,i}$ terms as

$$\beta_{1,i} = \frac{l_i c_i}{l_i^2}.$$

It is then straightforward to reformulate the complete sample β_1 as a function of individual sample βs in the following way:

$$\beta_1 = \sum_{i=1}^{N} \beta_{1,i} \omega_i,$$

where the weights are defined as $\omega_i = l_i^2 / \sum_{j=1}^{N} l_j^2$. That is, the weight of person i is defined as the share of the sum of squared lottery prizes that is explained by person i.

The weights are easier to interpret if we reformulate them as relative weights: the relative importance of $\beta_{1,i}$ compared with $\beta_{1,j}$ is then

$$\gamma_{ij} = \frac{\omega_i}{\omega_j} = \frac{l_i^2}{l_j^2}.$$

In short: if household i wins 100 times as much as household j, then it is 10,000 as important in determining the average β_1.

A.4 *Our Identification Strategy in Simulated Data*

To illustrate what we identify by our estimates, we show here how our identification strategy works on simulated data. We start by simulating lottery winners using the following specifications:

$$\begin{aligned} consumption_{it} &= 10\log\left(lottery_{it}/8 + 1\right) + \epsilon_{it}, \\ lottery_{it} &\sim \text{lognormal}\left(4{,}500\right), \\ \epsilon_{it} &\sim N\left(0, \sigma_\epsilon\right), \end{aligned}$$

where ϵ_{it} is measurement error on consumption. We discard lottery draws that are lower than US$1,100 or greater than US$150,000. In the simulation sample, we use 25,000 such lottery draws in addition to 75,000 observations with 0 in lottery prize to emulate the sample in the paper.

Figure A2 presents the results with low σ_ϵ. The blue line and red line represent the analytical marginal and average propensities to consume, respectively. The black dashed line shows what we estimate by our benchmark specification (equation (3)), and the black dotted line shows what we estimate when we split the lottery sample in quartiles.

There are two main takeaways from Figure A2. First, our results should be interpreted as the average propensity to consume. We estimate how large a fraction of the lottery prize the average lottery winner in our sample spends on consumption within

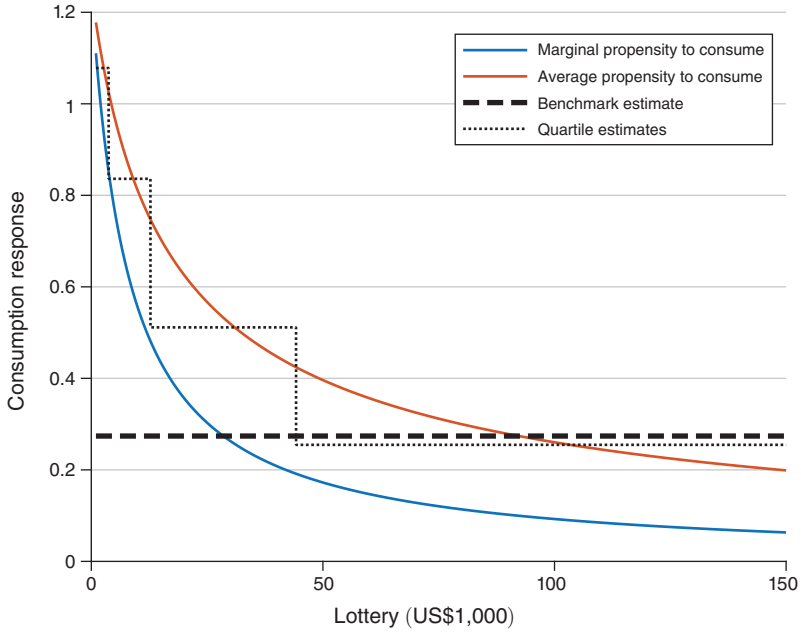


FIGURE A2. OUR IDENTIFICATION STRATEGY ON SIMULATED DATA

Notes: The figure shows how our identification strategy works on simulated data with the data generating process $consumption_{it} = 10\log(lottery_{it}/8 + 1)$. The blue and red lines depict analytical solutions. We draw lottery prizes from a truncated lognormal distribution with mean 4 and standard deviation 500. We include 25,000 lottery winners and 75,000 observations prior to winning in the simulation.

the year of winning. This is also the standard interpretation of MPC in the empirical literature. Second, our benchmark estimate is close to the estimate among big prize winners. We know from the derivations in Appendix A.3 that big winners acquire a greater weight in an OLS-regression. Our benchmark estimates should therefore be interpreted with caution, and we illustrate how the size of the lottery prize affects our consumption responses in Section IIIB.

A.5 *Dynamic Earnings Response to Lottery Prizes*

In this appendix, we reproduce the annual earnings responses to lottery prizes in Cesarini et al. (2017) using Norwegian data. The goal is to assess the reliability of our identification strategy. The main issue with our data is that we do not observe how much households spend on lottery tickets, only how much they win. We therefore rely on comparing winners with themselves in the years prior to winning to identify the effect of the lottery prize. The underlying assumptions are that households spend the same amount on lottery tickets each year and that the lottery prizes are random. For our identification strategy to be valid, a minimum requirement is that there is no effect of lottery prizes on earnings in periods prior to winning. Further, the labor supply response of winners should be of similar magnitude as that found using Swedish data in Cesarini et al. (2017).

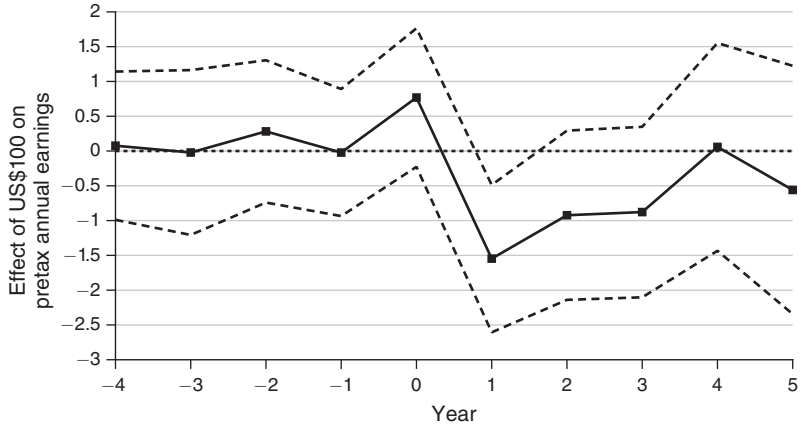


FIGURE A3. EFFECTS OF LOTTERY PRIZE ON HOUSEHOLDS' EARNINGS

Notes: Each dot represents a separate regression of gross household earnings on lottery prize. Controls include household-fixed effects and time-fixed effects. Dashed lines show 95 percent confidence intervals. The standard errors in parentheses are robust and clustered on the household level. Estimation method: OLS.

Figure A3 presents the results from estimating the response of pretax annual earnings of households to lottery prizes. There are no effects of future lottery prizes on current earnings. Furthermore, we find that after winning a lottery, households respond by reducing labor earnings by about 1.5 percent of the lottery prize. Compared with Cesarini et al.(2017), our estimated earnings response has about the same magnitude, although the response is less persistent.

A.6 *Pre-Trends*

Figure A4 shows how various household characteristics evolved in the years before households won the lottery prize. The paper's main body documents pre-trends for our outcome variables (consumption and savings), while we here document the pre-trends for a set of remaining variables. The vertical axis measures each respective variable's response per future dollar won. As the figure shows, no substantial movements are observed in the years before winning.

A.7 *Time Aggregation*

Since our data are yearly, our estimates constitute time-averaged responses to lottery prizes. For example, the within-year response is an average across households who won 0 to 12 months ago, the year 1 response is an average across households who won 0 to 24 months ago, and the year 2 response is an average across households who won 12 to 36 months ago. It is therefore not obvious how to map our estimates into structural models. A natural strategy is to impose an assumption on how prizes are distributed within a year and in addition impose a profile for how the response evolves over time. In this appendix, we use the information from our directly estimated dynamic consumption responses to infer MPCs at various time horizons.

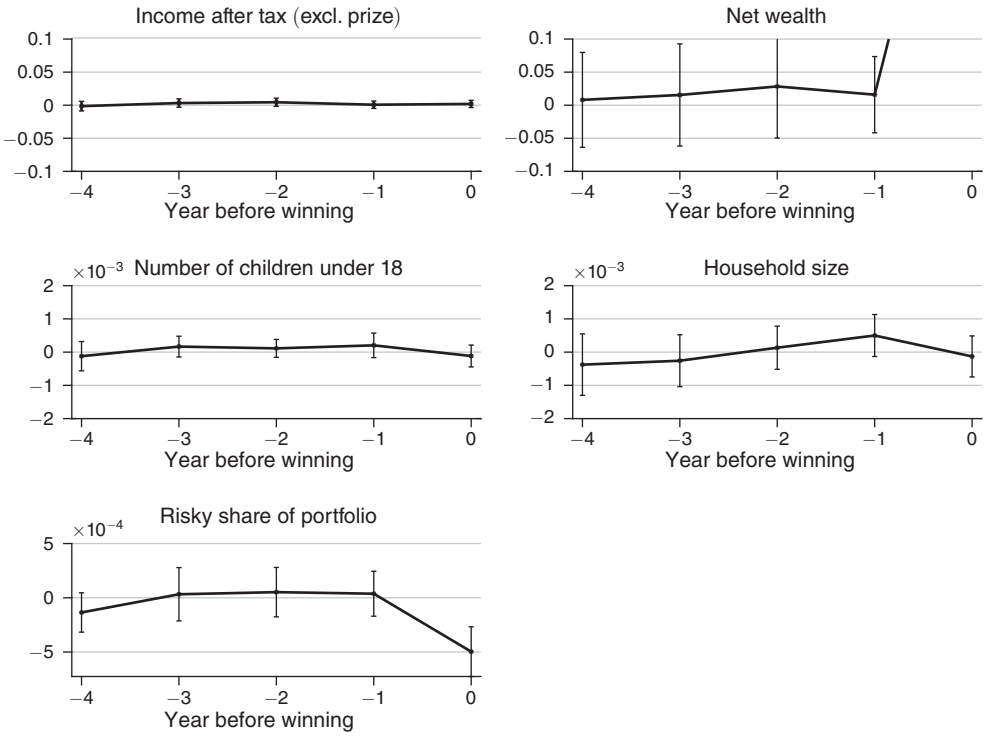


FIGURE A4. PRE-TRENDS

Notes: The figure shows responses of income, net wealth, risky share of portfolio, household size, and number of children to a lottery prize in the years before winning. Year 0 is the year of winning. The scale on the y-axis is in per unit of lottery prize.

We proceed in three steps. First, most games are weekly and played throughout the year, so we assume that prizes are distributed uniformly over a year. Second, we assume that the marginal consumption response as a function of time has the power form

$$mpc(t) = \theta_1 t^{\theta_2},$$

where θ_1 and θ_2 are parameters and t is years. Third, we search for the θ_1 and θ_2 that minimize the mean squared distance from our estimated dynamic consumption responses from Figure 2. In particular, we simulate $N = 50{,}000$ observations of within-year dates t_n from a uniform distribution. The model response that corresponds to the empirically estimated within-year response is then $1/N\sum_{n=1}^{N}\theta_1 t_n^{\theta_2}$. Similarly, the model response that corresponds to the empirical year 1 response is $1/N\sum_{n=1}^{N}\theta_1\left[(t_n+1)^{\theta_2} - t_n^{\theta_2}\right]$.

Figure A5 shows the dynamic consumption response and the fitted function. The outcome is $\theta_1 = 0.629$ and $\theta_2 = 0.214$. We can now use the estimated function to compute an MPC at any time horizon. For example, the model implies that the one-month MPC is 0.37, the quarterly MPC is 0.47, the half-year MPC is 0.54, and the

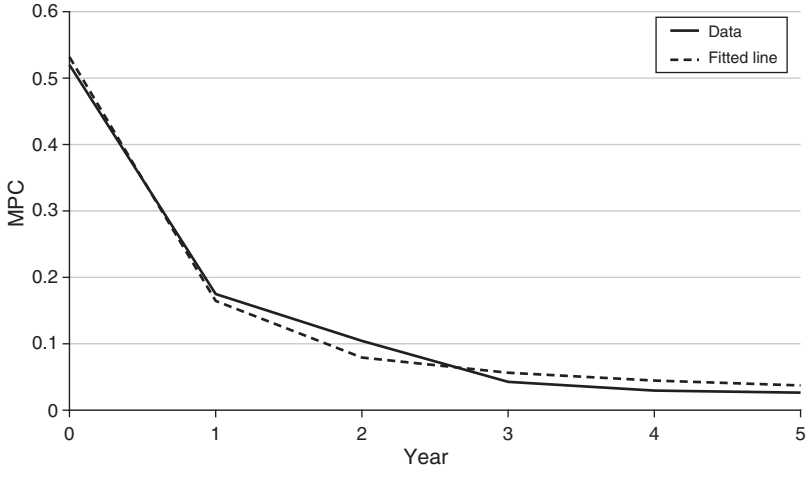


FIGURE A5. TIME AGGREGATION—DATA AND FITTED MODEL

Notes: The solid line is the estimated dynamic consumption responses from Figure 2. The dashed line is the fitted consumption response using the approach described above.

one-year MPC is 0.63. Similarly, we can use the model to calculate the consumption response within a specific year. For example, in the second year since winning (end of year 1 to end of year 2), the consumption response is 0.10.

A.8 *Additional Results on the "No-Risky-Assets" Sample*

As discussed in Sections IB and IV, our measure of imputed consumption contains measurement errors. There are, in particular, two assumptions that could introduce systematic errors in our marginal consumption responses. First, we assume that all households hold the market portfolio for every risky asset. And second, we assume that household portfolios are fixed over the year. Any deviations from these assumptions will introduce measurement errors in imputed consumption that will feed into our estimated consumption responses.

Measurement errors by themselves do not entail problems for our identification. If measurement errors are i.i.d., we only get higher standard errors. However, for our purpose, it does pose a potential problem since we show in Figure 2 that lottery winners tend to increase their holdings of risky assets after winning the lottery. While most of this increase in risky assets will be because lottery winners are net buyers, some of it will be capital gains throughout the year. Since we calculate capital gains only on the portfolio at the beginning of the year, the presence of capital gains within the year implies that we overestimate the active saving response of lottery winners and underestimate their consumption response.

To show that our assumptions of capital gains do not qualitatively affect our results, we presented the interaction regression in Table 4 also for the no-risky-asset sample. Here, we also present the dynamic consumption and balance sheet responses to lottery prizes in Figure A6. Both the patterns and the quantitative sizes of the responses are similar to the results using the main sample. The only exception is that

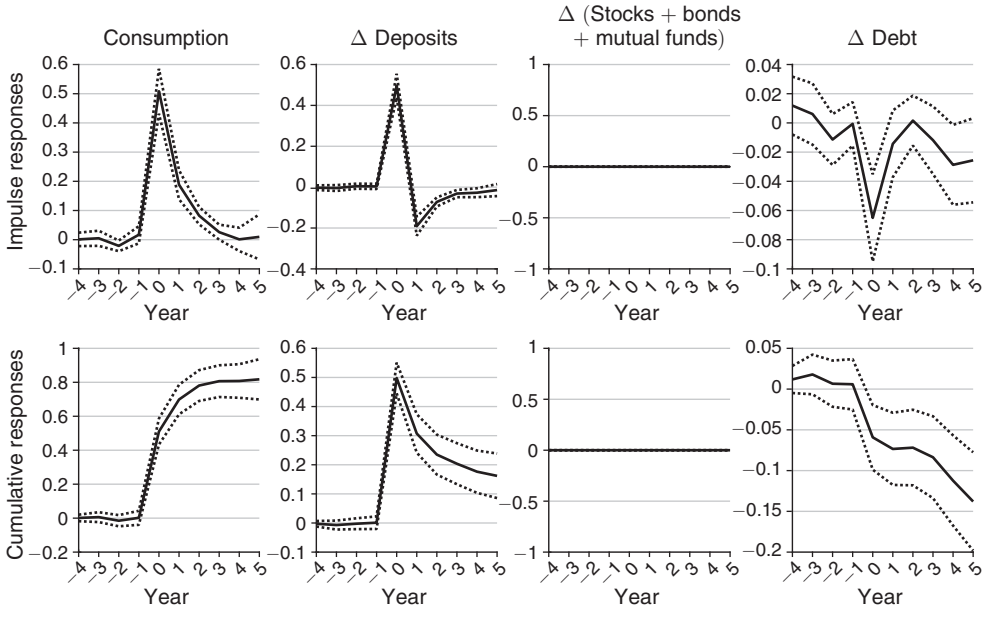


FIGURE A6. DYNAMIC HOUSEHOLD RESPONSES TO LOTTERY PRIZES. NO-RISKY-ASSETS SAMPLE

Notes: Each point is estimated as a separate regression of (4). Controls include time-fixed and household-fixed effects. Dotted lines represent 95 percent confidence intervals using robust standard errors clustered at the household level. The cumulative responses in the bottom panel are the sum of year-specific responses in the top panel. Standard errors of estimated cumulative effects are obtained through Monte Carlo simulations.

since we restrict on a sample of households that never hold any risky assets, there is no response of stocks, bonds, and mutual funds. Instead, the deposit response is a little bit larger than in the original sample.

A.9 *Normalizing by Permanent Income*

In some models of household behavior, it is not the level of different balance sheet variables but the level *relative* to permanent income that determines households' behavior. For example, in a standard Huggett model, one can collapse the model so that the relevant state variable is net wealth (liquid assets) relative to permanent income. In this appendix, we construct a measure of permanent income and use it to estimate consumption responses when we normalize by permanent income.

Assume that income follows an AR(1) of the following form in logs:

$$y_{i,j,a,t} = \bar{y}_{j,a,t-1} + \rho\left(y_{i,j,a,t-1} - \bar{y}_{j,a,t-1}\right) + \epsilon_{i,j,a,t},$$

where $y_{i,j,a,t}$ is the log income of household i, which is type j, age a, and in year t; $\rho \in [0,1)$ is the persistence; and $\epsilon_{i,j,a,t} \sim N(0,\sigma)$ is the error term. Then $\bar{y}_{j,a,t-1}$ is the long-run level at which income converges, i.e., the level of income you would have absent any shocks. We define this $\bar{y}_{j,a,t-1}$ as our measure of permanent income.

To construct this permanent income, we first restrict our sample to households between the ages of 25 and 60 to ensure that our households most likely work full-time. We construct permanent income in three steps:

(i) Regress the log of income net of taxes on a set of time, age, and time-age dummies:

$$y_{i,a,t} = \gamma_t + \alpha_a + \eta_{a,t} + \epsilon_{i,t},$$

where γ_t is the time-fixed effect, α_a is an age-fixed effect, and $\eta_{a,t}$ is the time-age-fixed effect.

(ii) Calculate the average of the residuals over all years (1993–2006) per household:

$$\bar{\epsilon}_i = \frac{1}{N}\sum_{1993}^{2006}\epsilon_{i,t},$$

where $N = \sum_{1993}^{2006}\mathbf{1}_{i \text{ exists}}$. Thus, $\bar{\epsilon}_i$ describes the type of the household.

(iii) Permanent income is then

$$\bar{y}_{i,a,t} = \gamma_t + \alpha_a + \eta_{a,t} + \bar{\epsilon}_i.$$

(iv) In the regressions below, we use the lagged version of the level of income to normalize in period t: $\exp(\bar{y}_{i,a,t-1})$.

We next reestimate Table 4, where we divide all dollar-denoted variables by lagged permanent income. In particular, both consumption (left-hand-side variable) and all right-hand-side variables are divided by permanent income. We define the variables that are not denominated in dollars, such as education, risky share of balance sheet, and age, as before. All monetary control variables (interacted terms alone) are also divided by permanent income.

Table A2 presents the results. We find no statistically significant effects. However, although the effects are not statistically significant, the signs and sizes of the effects are consistent with the finding in Table 4.[44]

A.10 *Co-movements between Consumption and Income*

In the data, some households' consumption tends to be systematically more responsive to variations in income, suggesting hand-to-mouth behavior or limited consumption smoothing. In theory, such behavior can be due to the presence of credit constraints but also to persistent household characteristics. For example, impatient households will consume almost all of their income in every period, while patient

[44] Note that we do not include the interaction term with income (defined as previous period income divided by permanent income) since this is close to just regressing on prize and suffer from almost perfect collinearity with the linear lottery term.

Table A2—The MPC out of Lottery Prizes. Interaction Effects. Normalizing by Permanent Income

	Univariate (1)	Multivariate (2)
$\text{Lottery}_t{}^2$	0.005	−0.000
	(0.008)	(0.007)
$\text{Lottery}_t \times \text{liquid assets}_{t-1}$	−0.029	−0.042
	(0.029)	(0.031)
$\text{Lottery}_t \times \text{income}_{t-1}$	−0.127	−0.020
	(0.098)	(0.104)
$\text{Lottery}_t \times \text{net wealth}_{t-1}$	0.005	0.008
	(0.006)	(0.006)
$\text{Lottery}_t \times \text{debt}_{t-1}$	−0.015	−0.019
	(0.014)	(0.014)
$\text{Lottery}_t \times \text{education}_t$	0.007	0.010
	(0.008)	(0.009)
$\text{Lottery}_t \times \text{risky share}_{t-1}$	0.003	−0.036
	(0.091)	(0.083)
$\text{Lottery}_t \times \text{household size}_t$	0.021	0.023
	(0.023)	(0.020)
$\text{Lottery}_t \times \text{age}_t$	−0.003	−0.003
	(0.003)	(0.003)
Observations	61,190	61,190

Notes: This table displays the estimated interaction terms from the estimation of equation (5). The univariate column presents the interaction coefficients when we include each interaction term with the lottery prize ($lottery_{i,t} \times Z_{i,t-1}$) individually in separate regressions (together with the remaining variables in equation (5)), while the multivariate column presents results when we include all interaction terms in the same regression. All variables that are denoted in US$ are divided by permanent income as described in Appendix A.9. Controls include the variables that are interacted with lottery prize, time-fixed effects, and household-fixed effects. The sample size (61,190) includes the winners and all observations of the winners in the years prior to winning, conditional on being able to estimate permanent income. The standard errors in parentheses are robust and clustered at the household level.

households will only consume a small fraction. The goal of this section is to include a measure of hand-to-mouth behavior interacted with the lottery prize to see if it can explain the heterogeneity in MPCs.

We proceed by constructing an estimate of the average marginal propensity to consume out of labor income at the household level in three steps:

(i) Exclude the year in which they win the lottery and keep only households with at least four observations.

(ii) Clean consumption and income of time-fixed effects by controlling for time dummies:

$$c_{i,t} = \gamma_t^c + \epsilon_{i,t}^c,$$

$$y_{i,t} = \gamma_t^y + \epsilon_{i,t}^y.$$

(iii) Estimate an MPC from labor income as $MPC_i = \text{cov}\left(\epsilon_{i,t}^c, \epsilon_{i,t}^y\right)/\text{var}\left(\epsilon_{i,t}^y\right)$.

TABLE A3—THE MPC OUT OF LOTTERY PRIZES. INTERACTION EFFECTS. CONTROLLING FOR MPC FROM INCOME

	Dependent variable: Consumption$_t$		
	(1)	(2)	(3)
Lottery$_t$	0.579	0.574	0.879
	(0.017)	(0.017)	(0.093)
Lottery$_t$ × MPC$_i$		0.009	0.015
		(0.008)	(0.007)
Lottery$_t{}^2$			−0.001
			(0.000)
Lottery$_t$ × liquid assets$_{t-1}$			−0.003
			(0.001)
Lottery$_t$ × income$_{t-1}$			−0.002
			(0.002)
Lottery$_t$ × net wealth$_{t-1}$			0.000
			(0.000)
Lottery$_t$ × debt$_{t-1}$			0.000
			(0.001)
Lottery$_t$ × education$_t$			0.008
			(0.007)
Lottery$_t$ × risky share$_{t-1}$			−0.034
			(0.074)
Lottery$_t$ × household size$_t$			0.017
			(0.019)
Lottery$_t$ × age$_t$			−0.005
			(0.001)
Observations	87,968	87,968	87,968

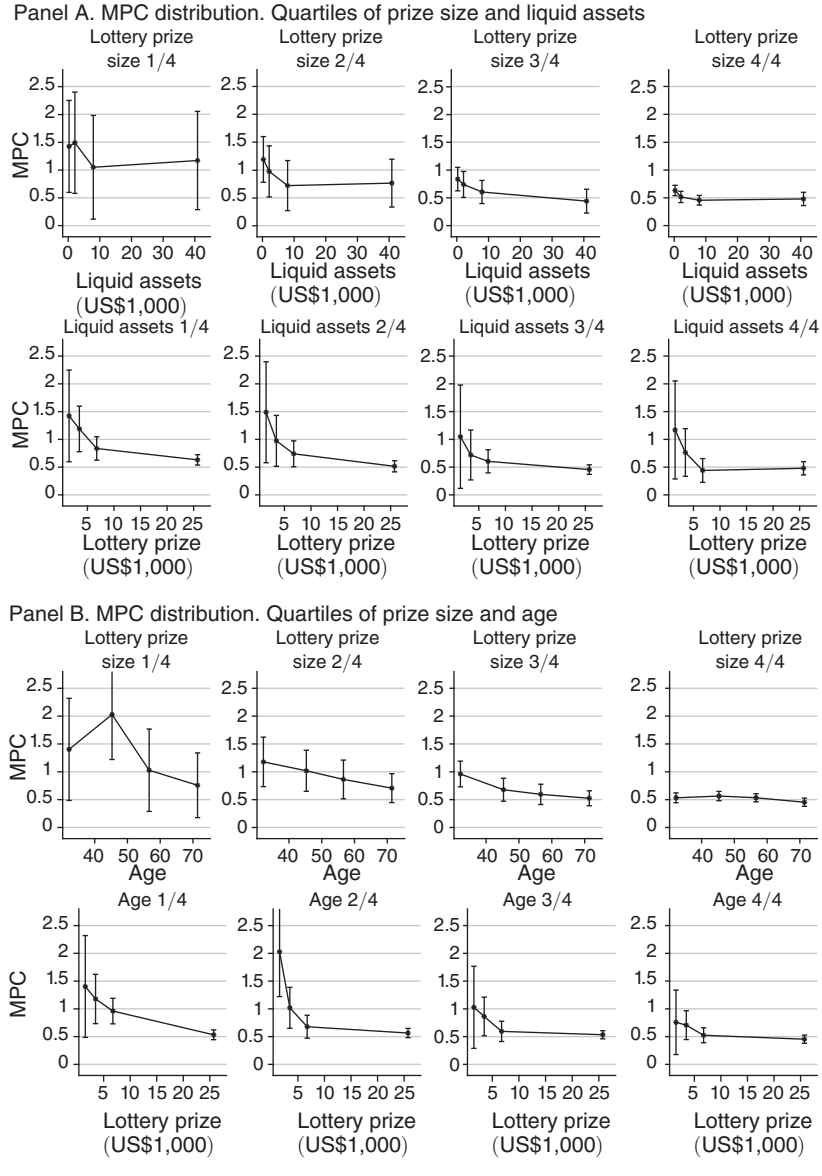
Notes: Each column represents a separate regression of (5). Controls include the variables that are interacted with lottery prize, time-fixed effects, and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

Then, MPC_i is a measure of the MPC of an individual household and meant to capture the extent to which the household behaves in a hand-to-mouth fashion.

The idea is now to include MPC_i in the interaction term regressions. By controlling for MPC_i at the same time as the balance sheet variables, the effect of MPC_i on the MPC out of lottery prizes can be interpreted as an innate characteristic since we have already controlled for the household's current financial position by including the level of liquid assets and other balance sheet variables in the regression.

Table A3 presents the results. The interaction effect of the MPC from labor income is positive, implying that households that let consumption vary more with labor income also have a higher MPC from lottery prizes. However, when we introduce the interaction effect with the MPC from labor income, all the other effects remain present. Thus, the effects of lottery size, liquid assets, and age are robust to introducing the estimated MPC from labor income in the regression.

A.11 *Additional Tables and Figures*



(*continued*)

FIGURE A7. HETEROGENEOUS CONSUMPTION RESPONSES. QUARTILES OF LOTTERY PRIZE SIZE, LIQUID ASSETS, AND AGE

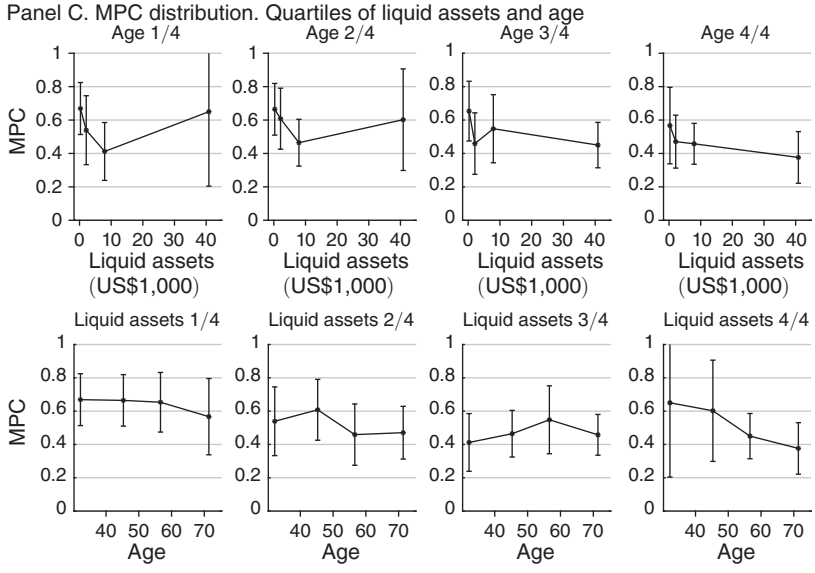


FIGURE A7. HETEROGENEOUS CONSUMPTION RESPONSES. QUARTILES OF LOTTERY PRIZE SIZE, LIQUID ASSETS, AND AGE (*continued*)

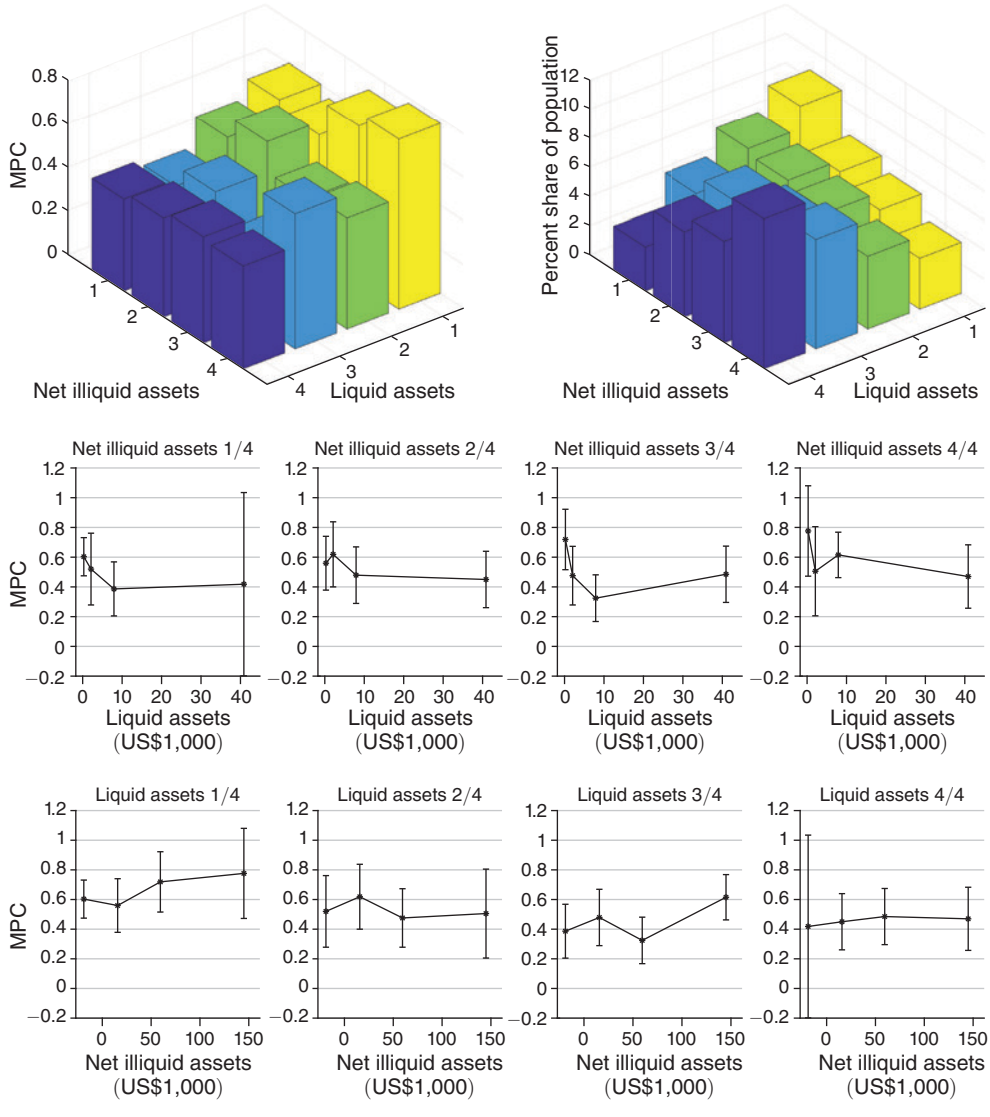


FIGURE A8. HETEROGENEOUS CONSUMPTION RESPONSES. NET ILLIQUID AND LIQUID ASSETS

Notes: Each bar/point is estimated as a separate regression of equation (3). Liquid assets = deposits + stocks + bonds + mutual funds. Net illiquid assets = net wealth − liquid assets = housing wealth − debt. Controls include time-fixed and household-fixed effects. The standard errors are robust and clustered at the household level. Total observations: 93,631.

Table A4—The MPC out of Lottery Prizes. Interaction Effects

	Dependent variable: Consumption$_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
Lottery$_t$	0.524	0.558	0.529	0.525	0.529	0.879
	(0.014)	(0.015)	(0.014)	(0.014)	(0.015)	(0.074)
Lottery$_t$ × deposits$_{t-1}$		−0.003				−0.001
		(0.001)				(0.000)
Lottery$_t$ × stocks$_{t-1}$			−0.012			−0.003
			(0.002)			(0.001)
Lottery$_t$ × bonds$_{t-1}$				−0.002		−0.011
				(0.003)		(0.002)
Lottery$_t$ × mutual funds$_{t-1}$					−0.005	−0.001
					(0.003)	(0.003)
Lottery$_t^2$						−0.005
						(0.003)
Lottery$_t$ × income$_{t-1}$						−0.002
						(0.002)
Lottery$_t$ × net wealth$_{t-1}$						0.000
						(0.000)
Lottery$_t$ × debt$_{t-1}$						0.000
						(0.000)
Lottery$_t$ × education$_t$						0.008
						(0.006)
Lottery$_t$ × risky share$_{t-1}$						0.021
						(0.072)
Lottery$_t$ × household size$_t$						0.025
						(0.015)
Lottery$_t$ × age$_t$						−0.005
						(0.001)
Observations	93,631	93,631	93,631	93,631	93,631	93,631

Notes: Each column represents a separate regression of (5). Controls include the variables that are interacted with lottery, time-fixed effects, and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

REFERENCES

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. 2015. "Regulating Consumer Financial Products: Evidence from Credit Cards." *Quarterly Journal of Economics* 130 (1): 111–64.

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

Agarwal, Sumit, and Wenlan Qian. 2014. "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore." *American Economic Review* 104 (12): 4205–30.

Aguiar, Mark A., Mark Bils, and Corina Boar. 2020. "Who Are the Hand-to-Mouth." NBER Working Paper 26643.

Ahn, S., and M.B. Holm. 2020. "Monetary Policy under Demographic Transitions." Unpublished.

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics* 109 (3): 659–84.

Arellano, Manuel, and Olympia Bover. 1995. "Another Look at the Instrumental Variable Estimation of Error-Components Models." *Journal of Econometrics* 68 (1): 29–51.

Auclert, Adrien. 2019. "Monetary Policy and the Redistribution Channel." *American Economic Review* 109 (6): 2333–67.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2018. "The Intertemporal Keynesian Cross." http://web.stanford.edu/~aauclert/ikc.pdf.

Autor, David, Andreas Kostøl, Magne Mogstad, and Bradley Setzler. 2019. "Disability Benefits, Consumption Insurance, and Household Labor Supply." *American Economic Review* 109 (7): 2613–54.

Aydin, D. 2015. "The Marginal Propensity to Consume Out of Liquidity: Evidence from Random Assignment of 54,522 Credit Lines." Unpublished.

Baker, Scott R. 2018. "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data." *Journal of Political Economy* 126 (4): 1504–57.

Baker, Scott R., Lorenz Kueng, Steffen Meyer, and Michaela Pagel. 2018. "Measurement Error in Imputed Consumption." NBER Working Paper 25078.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra. 2018. "House Prices and Consumer Spending." *Review of Economic Studies* 85 (3): 1502–42.

Blundell, Richard, and Stephen Bond. 1998. "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models." *Journal of Econometrics* 87 (1): 115–43.

Bø, Erlend E., Joel Slemrod, and Thor O. Thoresen. 2015. "Taxes on the Internet: Deterrence Effects of Public Disclosure." *American Economic Journal: Economic Policy* 7 (1): 36–62.

Briggs, J., D. Cesarini, E. Lindqvist, and R. Östling. 2021. "Windfall Gains and Stock Market Participation." *Journal of Financial Economics* 139 (1): 57–83.

Browning, Martin, and M. Dolores Collado. 2001. "The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates." *American Economic Review* 91 (3): 681–92.

Browning, Martin, and Thomas F. Crossley. 2001. "The Life cycle Model of Consumption and Saving." *Journal of Economic Perspectives* 15 (3): 3–22.

Browning, Martin, Thomas F. Crossley, and Joachim Winter. 2014. "The Measurement of Household Consumption Expenditures." *Annual Review of Economics* 6: 475–501.

Browning, Martin, Mette Gørtz, and Søren Leth-Petersen. 2013. "Housing Wealth and Consumption: A Micro Panel Study." *Economic Journal* 123 (568): 401–28.

Browning, Martin, and Søren Leth-Petersen. 2003. "Imputing Consumption from Income and Wealth Information." *Economic Journal* 113 (488): F282–F301.

Campbell, John Y., and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." In *NBER Macroeconomics Annual*, Vol. 4, edited by Olivier Jean Blanchard and Stanley Fischer, 185–216. Cambridge, MA: MIT Press.

Carroll, C. Forthcoming. "Theoretical Foundations of Buffer Stock Saving." *Quantitative Economics*. https://qeconomics.org/ojs/forth/354/354-1.pdf.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112 (1): 1–55.

Carroll, Christopher D. 2001. "A Theory of the Consumption Function, with and without Liquidity Constraints." *Journal of Economic Perspectives* 15 (3): 23–45.

Carroll, Christopher D., Martin B. Holm, and Miles S. Kimball. 2021. "Liquidity constraints and precautionary saving." *Journal of Economic Theory* 195: 105276.

Carroll, Christopher D., and Miles S. Kimball. 1996. "On the Concavity of the Consumption Function." *Econometrica* 64 (4): 981–92.

Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N. White. 2017. "The Distribution of Wealth and the Marginal Propensity to Consume." *Quantitative Economics* 8 (3): 977–1020.

Cesarini, David, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling. 2017. "The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries." *American Economic Review* 107 (12): 3917–46.

Cesarini, David, Erik Lindqvist, Robert Östling, and Björn Wallace. 2016. "Wealth, Health, and Child Development: Evidence from Administrative Data on Swedish Lottery Players." *Quarterly Journal of Economics* 131 (2): 687–738.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij. 2019. "Asymmetric Consumption Effects of Transitory Income Shocks." *Economic Journal* 129 (622): 2322–41.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary Policy When Households Have Debt: New Evidence on the Transmission Mechanism." *Review of Economic Studies* 87 (1): 102–29.

Crossley, Thomas F., Hamish Low, and Sarah Smith. 2016. "Do Consumers Gamble to Convexify?" *Journal of Economic Behavior and Organization* 131: 276–91.

De Nardi, Mariacristina. 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies* 71 (3): 743–68.

De Nardi, Mariacristina, and Giulio Fella. 2017. "Saving and Wealth Inequality." *Review of Economic Dynamics* 26: 280–300.

De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy* 118 (1): 39–75.

Di Maggio, M., A. Kermani, and K. Majlesi. 2020. "Stock Market Returns and Consumption." *Journal of Finance* 75 (6): 3175–219.

Druedahl, Jeppe, and Alessandro Martinello. 2020. "Long-Run Saving Dynamics: Evidence from Unexpected Inheritances." *Review of Economics and Statistics*, forthcoming.

Eika, Lasse, Magne Mogstad, and Ola L. Vestad. 2020. "What Can We Learn about Household Consumption Expenditure from Data on Income and Assets?" *Journal of Public Economics* 189: 104163.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and Persistence in Returns to Wealth." *Econometrica* 88 (1): 115–70.

Fagereng, Andreas, and Elin Halvorsen. 2017. "Imputing Consumption from Norwegian Income and Wealth Registry Data." *Journal of Economic and Social Measurement* 42 (1): 67–100.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. "Replication data for: MPC Heterogeneity and Household Balance Sheets." American Economic Association [publisher], —Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E121561V1.

Fagereng, Andreas, Martin Blomhoff Holm, and Kjersti Næss Torstensen. 2021. "Housing Wealth in Norway, 1993–2015." *Journal of Economics and Social Measurement* 45 (1): 65–81.

Fernández-Villaverde, Jesús, and David Zarruk Valencia. 2018. "A Practical Guide to Parallelization in Economics." NBER Working Paper 24561.

Friedman, Milton. 1957. *A Theory of the Consumption Function*. Vol. 81. Princeton, NJ: Princeton University Press.

Fuchs-Schündeln, Nicola. 2008. "The Response of Household Saving to the Large Shock of German Reunification." *American Economic Review* 98 (5): 1798–828.

Fuchs-Schündeln, N., and T.A. Hassan. 2016. "Chapter 12—Natural Experiments in Macroeconomics." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 923–1012. Amsterdam: North-Holland.

Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2020. "What Would You Do With $500? Spending Responses to Gains, Losses, News, and Loans." *Review of Economic Studies*, forthcoming.

Gross, Tal, Matthew J. Notowidigdo, and Jialan Wang. 2020. "The Marginal Propensity to Consume over the Business Cycle." *American Economic Journal: Macroeconomics* 12 (2): 351–84.

Hsieh, Chang-Tai. 2003. "Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund." *American Economic Review* 93 (1): 397–405.

Huggett, Mark. 1993. "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies." *Journal of Economic Dynamics and Control* 17 (5–6): 953–69.

Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote. 2001. "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players." *American Economic Review* 91 (4): 778–94.

Jappelli, Tullio, and Luigi Pistaferri. 2010. "The Consumption Response to Income Changes." *Annual Review of Economics* 2: 479–506.

Jappelli, Tullio, and Luigi Pistaferri. 2014. "Fiscal Policy and MPC Heterogeneity." *American Economic Journal: Macroeconomics* 6 (4): 107–36.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–610.

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

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica* 82 (4): 1199–239.

Koijen, Ralph, Stijn Van Nieuwerburgh, and Roine Vestman. 2015. "Chapter 11—Judging the Quality of Survey Data by Comparison with 'Truth' as Measured by Administrative Records: Evidence from Sweden." In *Improving the Measurement of Consumer Expenditures*, Vol. 74, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus, 308–46. Chicago: University of Chicago Press.

Kolsrud, Jonas, Camille Landais, and Johannes Spinnewijn. 2020. "The value of registry data for consumption analysis: An application to health shocks." *Journal of Public Economics* 189: 104088.

Kreiner, Claus T., David D. Lassen, and Søren Leth-Petersen. 2015. "Chapter 10—Measuring the Accuracy of Survey Responses Using Administrative Register Data: Evidence from Denmark." In *Improving the Measurement of Consumer Expenditures*, Vol. 74, edited by Christopher D. Carroll, Thomas F. Crossley, and John Sabelhaus, 289–307. Chicago: University of Chicago Press.

Krueger, D., K. Mitman, and F. Perri. 2016. "Chapter 11—Macroeconomics and Household Heterogeneity." In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 843–921. Amsterdam: North-Holland.

Kueng, Lorenz. 2018. "Excess Sensitivity of High-Income Consumers." *Quarterly Journal of Economics* 133 (4): 1693–751.

Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn. 2011. "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery." *American Economic Review* 101 (5): 2226–47.

Leth-Petersen, Søren. 2010. "Intertemporal Consumption and Credit Constraints: Does Total Expenditure Respond to an Exogenous Shock to Credit?" *American Economic Review* 100 (3): 1080–103.

Meyer, Bruce D., Wallace K.C. Mok, and James X. Sullivan. 2015. "Household Surveys in Crisis." *Journal of Economic Perspectives* 29 (4): 199–226.

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

Modigliani, Franco. 1954. "Utility Analysis and the Consumption Function: An Interpretation of Cross-Section Data." In *Post-Keynesians Economics*, edited by Kenneth K. Kurihara, 388–436. New Brunswick, NJ: Rutgers University Press.

Ng, Yew Kwang. 1965. "Why Do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure." *Journal of Political Economy* 73 (5): 530–5.

Norges Bank. 2020. "Exchange Rates." https://www.norges-bank.no/en/topics/Statistics/exchange_rates/?tab=currency&id=USD (accessed June 10, 2017).

Parker, Jonathan A. 2017. "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment." *American Economic Journal: Macroeconomics* 9 (4): 153–83.

Parker, Jonathan A., and Nicholas S. Souleles. 2019. "Reported Effects versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates." *American Economic Review: Insights* 1: 273–90.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6): 2530–53.

Pistaferri, L. 2015. "Household Consumption: Research Questions, Measurement Issues, and Data Collection Strategies." *Journal of Economic and Social Measurement* 40 (1–4): 123–49.

Scholnick, Barry. 2013. "Consumption Smoothing after the Final Mortgage Payment: Testing the Magnitude Hypothesis." *Review of Economics and Statistics* 95 (4): 1444–49.

Shapiro, Matthew D., and Joel Slemrod. 2003. "Consumer Response to Tax Rebates." *American Economic Review* 93 (1): 381–96.

Shapiro, Matthew D., and Joel Slemrod. 2009. "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review* 99 (2): 374–79.

Statistics Norway. 2011. "Population and Housing Census, 1980, 1990, 2001, 2011." https://www.ssb.no/en/omssb/tjenester-og-verktoy/data-til-forskning/befolkning (accessed June 10, 2017).

Statistics Norway. 2015a. "Consumer Price Index, 1990–2015." https://www.ssb.no/en/kpi (accessed June 10, 2017).

Statistics Norway. 2015b. "Dwellings, 1993–2015." https://www.ssb.no/en/bygg-bolig-og-eiendom/bolig-og-boforhold/statistikk/boliger (accessed June 10, 2017).

Statistics Norway. 2015c. "FD-trygd." https://www.ssb.no/en/sosiale-forhold-og-kriminalitet/artikler-og-publikasjoner/fd-trygd (accessed June 10, 2017).

Statistics Norway. 2015d. "Income and Wealth Tax Registry, 1993–2015." https://www.ssb.no/en/omssb/tjenester-og-verktoy/data-til-forskning/inntekt (accessed June 10, 2017).

Statistics Norway. 2015e. "National Education Database (NUDB)." https://www.ssb.no/en/omssb/tjenester-og-verktoy/data-til-forskning/utdanning (accessed June 10, 2017).

Statistics Norway. 2015f. "Shareholder Registry, 2004–2015." https://www.ssb.no/data-til-forskning/utlan-av-data-til-forskere/variabellister/aksjonaerregisteret (accessed June 10, 2017).

Statistics Norway. 2015g. "Transfer of Properties, 1990–2015." https://www.ssb.no/en/bygg-bolig-og-eiendom/eiendom/statistikk/eiendomsomsetning (accessed June 10, 2017).

Ziliak, James P. 1998. "Does the Choice of Consumption Measure Matter? An Application to the Permanent-Income Hypothesis." *Journal of Monetary Economics* 41 (1): 201–16.