CESifo WORKING PAPERS

7134
2018

Original Version: June 2018
This Version: June 2019

MPC Heterogeneity and Household Balance Sheets

Andreas Fagereng, Martin B. Holm, Gisle J. Natvik

CESifo

Impressum:

CESifo Working Papers
ISSN 2364-1428 (electronic version)
Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH
The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute
Poschingerstr. 5, 81679 Munich, Germany
Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de
Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl
www.cesifo-group.org/wp

An electronic version of the paper may be downloaded
· from the SSRN website: www.SSRN.com
· from the RePEc website: www.RePEc.org
· from the CESifo website: www.CESifo-group.org/wp

CESifo *Working Paper No. 7134*
Category 6: Fiscal Policy, Macroeconomics and Growth

MPC Heterogeneity and Household Balance Sheets

Abstract

We use sizeable 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 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 USD 1,500) are estimated to spend all within the year of winning. The corresponding estimate for high-liquidity winners of large prizes (USD 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 co-variation with observables.

JEL-Codes: D120, D140, D910, E210.

Keywords: marginal propensity to consume, household heterogeneity, income shocks.

Andreas Fagereng
Statistics Norway
Oslo / Norway
andreas.fagereng@ssb.no

Martin B. Holm
University of Oslo
Oslo / Norway
martin.b.holm@outlook.com

Gisle J. Natvik
BI Norwegian Business School
Oslo / Norway
gisle.j.natvik@bi.no

May, 2019
We thank 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 discussions and suggestions.

1 Introduction

What determines how households adjust their consumption expenditure 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 longstanding macroeconomic questions about shock propagation and economic policy.[1] In this paper, we contribute by providing such statistics. With direct evidence on observed windfall gains, we characterize (i) how transitory income shocks feed into 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.

Our contribution is rooted in how we deal with three econometric challenges. These hurdles explain why existing evidence on the determinants of households' marginal propensity to consume (MPC) out of transitory income shocks is incomplete, despite longstanding interest in the subject.[2,3] First, credible estimation of MPCs requires that the researcher observes exogenous income innovations. Moreover, it is not sufficient that the innovations are exogenous, it must also be clear whether they are anticipated or unanticipated, as theory gives different predictions for the two (Modigliani and Brumberg, 1954; Friedman, 1957). For the same reason, one must know that the shock was perceived by the recipient as transitory, not persistent. Such exogenous shocks with a clear information structure are hard to come by in the data. We use sizable prizes derived from gambling activities in which a majority of Norwegians participate. Second, the income shocks must be observed together with data on household consumption, which is a rare combination. We utilize detailed and precise third-party reported information on households' balance sheets to impute their total consumption *expenditure* from the budget constraint. By construction, this is a measure of durable and non-durable 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 the determinants behind MPC *heterogeneity*, and ideally how income innovations are spent over an extended period of time. This requires panel data with rich information on household characteristics, in particular data on wealth and balance sheets since these play a central role in structural models of consumption dynamics, see for instance Kaplan and Violante (2014); Carroll, Slacalek, Tokuoka, and White (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

[1] See for instance Auclert (2016); Auclert, Rognlie, and Straub (2018); Berger, Guerrieri, Lorenzoni, and Vavra (2018); and Kaplan, Moll, and Violante (2018).

[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 et al. (2018).

[3] See for instance surveys by Browning and Collado (2001); Jappelli and Pistaferri (2010); and Fuchs-Schündeln and Hassan (2016).

knowledge, among the many carefully executed empirical MPC-studies that exist, this is the first paper to meet all these formidable data requirements.[4]

Our data allow us to explore and document an array of features regarding households' MPCs, defined as their propensity to spend out of a transitory income shock. We start by establishing that winners spend a substantial fraction of their prize within the first calendar year of receiving it. Our baseline estimate implies a within-year expenditure response around one half, with a range from 0.35 to 0.7 depending on how we deal with outliers among high-prize winners. While this point estimate is sensitive to outliers, the systematic patterns we find regarding dynamic consumption responses and MPC heterogeneity are robust. These findings constitute our contribution. Of what is not spent, most is saved as deposits. Deposits are thereafter gradually depleted to finance above-normal consumption expenditure up to five years after winning. The dynamics of saving in stocks, bonds, and mutual funds, or debt repayment, are quite different from this, as the responses occur mainly within the year of winning. Our estimates imply that after five years, households have spent about 90 percent of their windfall.[5] Moreover, behind these estimates lies considerable heterogeneity which we dissect in detail.

We explore how MPCs vary with a host of household characteristics, income and balance sheet components, and prize size. By allowing for interaction effects between lottery income and each of these variables, we single out those that are most important. The stock of liquid assets held prior to winning, age, and the amount won stand out. Once controlling for these three variables, other household characteristics play hardly any role for consumption sensitivity.

Quantitatively, the association between MPCs and liquidity is considerable. Estimates from the interaction model imply that a one standard deviation increase in liquidity is associated with an MPC reduction of 8 cents to the dollar won. Similar effects emerge when we abandon the linear interaction structure and estimate consumption responses within liquidity strata instead. These effects are consistent with previous findings in the literature, such as Misra and Surico (2014).[6] Our contribution here is to establish that liquidity is associated with MPC variation after controlling for household balance sheets 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.

[4]What we here have in mind is evidence on actual consumption behavior. An alternative route is to ask how households believe they would respond to hypothetical income shocks. Parker and Souleles (2017) discuss and compare the two approaches in the context of U.S. tax rebates.

[5]Auclert et al. (2018) detail how the dynamic consumption responses that we estimate can be used to distinguish between alternative models of household behavior.

[6]Misra and Surico (2014) use survey evidence on the U.S. tax rebates. Other examples are Leth-Petersen (2010) who studies the impact of a credit market reform on consumption in Denmark; Aydın (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 U.S. Great Recession; Agarwal, Chomsisengphet, Mahoney, and Stroebel (2015) and Gross, Notowidigdo, and Wang (2016) who study consumption dynamics around discontinuities in credit scores; Cloyne, Ferreira, and Surico (2018) who explore micro responses to monetary policy shocks. Recently, Fuster, Kaplan, and Zafar (2018) and Christelis, Georgarakos, Jappelli, Pistaferri, and van Rooij (2018) provide survey evidence pointing in the same direction, when they ask how households would allocate hypothetical windfall gains.

The consumption responses also decline with age and prize size. Moving from the youngest quartile (younger than 39 years) to the oldest quartile (older than 63 years), the MPC falls with 15 cents to the dollar won. When we split our sample by the amount won, we find that the within-year consumption response declines monotonically with prize size. Our estimates within the lowest prize quartile (less than about USD $2,000$) imply that winners of small prizes tend to spend everything within the year of winning, and even more by combining the money won with other sources of financing. In the highest prize size quartile (above USD $8,000$), the within-year response lies below one half.[7] This evidence is novel, as the role of shock magnitude has largely been overlooked by the empirical literature.[8]

Importantly, as part of our study we restrict attention to households who possess only deposits and debt on their financial balance sheet. Within this sub-sample, concerns regarding measurement error in our imputed expenditure measure are reduced to a minimum.[9] Our results are confirmed in this sub-sample. Another concern, more specific to our strategy of using lottery prizes, could be that we do not observe the bets placed; money spent on lottery tickets is part of our consumption expenditure measure. 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 of betting activity in the year of winning must be orthogonal to the household characteristics we consider. It follows that spending on lottery tickets cannot be driving the heterogeneity effects that we estimate.

Our baseline within-year MPC estimate is substantial, but still on the low side of estimates provided by the literature studying U.S. tax rebates in 2001 and 2008. The bulk of the existing evidence on actual income shocks and consumption stems from these quasi-experiments. Within this literature, Parker, Souleles, Johnson, and McClelland (2013) consider total consumption expenditure like we do. Exploring the 2008 rebate episode, where transfers per adult ranged from USD 300 to USD 600, they find total consumption responses in the range of 0.5 to 0.9 within three months of payment receipt.[10] However, both this study and those focusing on non-durable consumption compare households receiving a pre-announced rebate at different points in time, effectively identifying

[7] Throughout, we report values in fixed year 2000-prices and convert them using the NOK/USD exchange rate from the year 2000 (NOK/USD = 0.114).

[8] The only other evidence we are aware of on size effects, is in the concurrent studies of Fuster et al. (2018) and Christelis et al. (2018). Both use survey evidence on *hypothetical* income shocks. The former finds that MPCs increase with shock size, the latter that MPCs fall with shock size.

[9] It is well-known that an imputed measure of consumption expenditure may suffer from certain measurement issues related to capital gains on financial assets, which are especially important among the richest (Browning and Leth-Petersen, 2003). Baker, Kueng, Meyer, and Pagel (2018) study a sample of higher-end individuals using retail investment data from a bank in Germany to highlight the potential measurement errors stemming from trading fees and within-year trading activities that one does not observe in the annual data. They find that these errors are present, yet tend to be quantitatively moderate and centered around zero on average. In Appendix A.8, we redo our analysis on a sample of households for which these concerns are irrelevant, namely households with only deposits and debt on their financial balance sheet. Measurement of capital gains will here be virtually absent as interest income and expenses during the year are directly observed.

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

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 the consumption response to large predetermined payments from the Alaska Permanent Fund. These responses are conceptually different 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 ten months after the transfer was announced.

Parker (2017) questions if empirically observed associations between liquidity and MPCs are to be interpreted as *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, such as impatience or risk tolerance, that raise consumption sensitivity. In our setting one would expect net wealth, education and the portfolio share of risky assets to correlate 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. This statistic will arguably capture the household-specific component of consumption sensitivity. These results support a situational interpretation of our findings.

Consumption and savings responses to lottery income have been studied before, most prominently by Imbens, Rubin, and Sacerdote (2001) and Kuhn, Kooreman, Soetevent, and Kapteyn (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 Euros 12,500. The Dutch findings stand out from ours and the tax rebate literature, in that neither durable nor non-durable consumption responded by much.[11] More recently, Swedish lotteries have been used to identify income effects on health, labor supply, and portfolio choice, but not on consumption.[12]

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

[11] 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.

[12] Using Swedish lottery data, Cesarini, Lindqvist, Östling, and Wallace (2016) study effects of wealth on health and child development; Briggs, Cesarini, Lindqvist, and Östling (2015) study effects on stock market participation; and Cesarini, Lindqvist, Notowidigdo, and Östling (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).

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 a high MPC. 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 who redistribute 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 non-winners 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 seems unlikely that this mechanism is driving our main results. Indeed, when we look at how our rich set 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). Here households face uninsurable idiosyncratic risk and a borrowing constraint. The main determinant of households' MPC then is their net wealth level. 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 well motivated by our results. Second, the distinction between net wealth and liquid assets is explicit in recent two-asset frameworks. The prime example here is the model of Kaplan and Violante (2014), where 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 non-durable 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): a realistic earnings profile coupled with borrowing constraints will tend to generate high MPCs early in life, while luxury bequest motives might explain low MPCs late in life.

We substantiate the above points through a structural life-cycle model with a realistic earnings profile, a distinction between liquid and illiquid assets, and a luxury bequest motive. This model can be matched to the heterogeneity effects we document empirically. It also generates a time profile for the consumption response that is similar to our empirically estimated one. Still, it cannot account for the high empirically estimated average MPC *level* without further extensions. In a final exercise, we therefore add an additional property to the framework: on top of their rational consumption decision implied by the life-cycle model, households consume a random fraction of whatever they won. Because this additional feature is orthogonal to the optimal consumption decision, it lifts households' expenditure response level without distorting the model-generated time profile and systematic relationships between MPCs and observables. Supported by this model exercises, we conclude that our empirical estimates are consistent with the hypothesis that liquidity and age cause cross-sectional MPC *variation*, while the high average MPC *level* we estimate cannot be explained by liquidity-constraints alone.

The remainder of our paper is organized as follows. Section 2 presents the institutional setting and data. Section 3 provides our benchmark estimates of the MPC out of lottery earnings, including dynamic responses. Sections 4 and 5 contain our main contributions, as we characterize how MPCs vary with household characteristics and the amount won. Section 6 discusses robustness analyses. Section 7 compares our estimates to those from a specific model. Section 8 concludes.

2 Institutional background, data and sample selection

We base our study on Norwegian administrative data. Norway levies both income and wealth taxes, and the data in tax returns are third-party reported. Hence, the tax registry data 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 USD 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.

2.1 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. We combine these data with household identifiers from the population register to aggregate all income and wealth information at the household level.[13] 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 owned by the

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

individual and administered by the employer or intermediary, as well as information on the income from these assets.[14]

The tax authority then pre-fills 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.

2.2 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 of household consumption, as in Johnson et al. (2006) or Parker et al. (2013) with the Consumer Expenditure Survey (CEX) in the U.S., 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 in general 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 an ongoing discussion about the reliability of self-reported marginal propensities to consume from hypothetical income shocks (Parker and Souleles, 2017).

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 (2017) and Fagereng and Halvorsen (2017) for Norway.[15]

[14] These assets are for the most 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 from the Norwegian Mapping Authority (Kartverket), the land registration, and the population census 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 (2018b).

[15] Ziliak (1998) attempts to impute consumption using data from the Panel Study of Income Dynamics (PSID) in the US. 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 et al. (2017) and Fagereng and Halvorsen (2017) using Norwegian data, and Koijen, Van Nieuwerburgh, and Vestman (2015) and Kolsrud, Landais, and Spinnewijn (2017) using Swedish data.

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 take 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. 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 makes it difficult to assign the correct money flow to the right year.[16] 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.[17]

Disposable income observed in our data is $Y_t = I_t - T_t + \sum_j^J d_{jt} + L_t$, 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.[18] L_t

Other examples are Browning, Gørtz, and Leth-Petersen (2013); Leth-Petersen (2010); Autor, Kostøl, and Mogstad (2015); and Di Maggio, Kermani, and Majlesi (2018). Browning, Crossley, and Winter (2014) review the literature.

[16] Each year about 8 percent of Norwegian households complete a housing transaction according to our data. There is no significant change in this fraction after winning the lottery.

[17] 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 USD 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 6. There we show that; (i) while 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 baseline MPC estimate of about 1/3 because we did not condition on the amount won when trimming the consumption measure.

[18] We attribute to each homeowner's consumption expenditure a value of owner occupied housing services equal to 3% every year. This enables us to compare the consumption of renters (which includes rental payments) and home owners. The value of these services is meant to represent the price the home owner would have paid if s/he were to rent the same home on the market. Three percent is close to what Eika

is income from any other source, such as bequests 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" savings. 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 re-balances its portfolio at market prices at the end of the year. S_t is then given by

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

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,[19] debt, stocks, bonds, and mutual funds, held abroad and at home. However, we do not observe each individual's exact portfolio composition within these asset classes.

Our procedure is to use aggregate price indices p_{kt} to approximate $\sum_{j}^{J} (p_{jt} - p_{jt-1})\, a_{jt-1} \equiv \sum_{j}^{J} w_{jt-1} \left(\frac{p_{jt}}{p_{jt-1}} - 1 \right)$ by $\sum_{k}^{K} w_{kt-1} \left(\frac{p_{kt}}{p_{kt-1}} - 1 \right)$. 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 A.1 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.

In short, the description above shows how our imputed consumption measure rests on two key assumptions. One, we assume there is no intra-year trading. Two, 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

et al. (2017) find as the rent-to-value for Norway, using data from National Accounts.

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

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" savings (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, 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 the main text and in 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, Guiso, Malacrino, and Pistaferri (2018a) document substantial heterogeneity in (un)realized returns across households, 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 substantially 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 non-winners in Table 1 below, or when we consider the predictability of prize size in Table 2. Furthermore, when we drop all households with risky assets, we are left with a sample whose returns are directly observed and returns heterogeneity is unproblematic.

Finally, gifts and bequests are included in our measure of consumption expenditure. Conceptually, one might question whether this is appropriate. Bequests may arguably be classified as savings. However, for our purpose of estimating MPCs, that imputation choice seems innocuous as observed bequests are a marginal phenomenon among winners.[20]

2.3 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. These restrictions, however, do not mean that Norwegians are prevented from gambling. According to Norsk Tipping, 70 percent of

[20]In our sample period, it was compulsory to report inheritance to the tax authorities, and accumulated inheritance above USD 34,000 was subject to taxation.

Norwegians above the age of 18 gambled in 2012 through their services.[21]

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 one per every 800 adult Norwegians), usually a kiosk or a local super market. 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 NOK 1,000, equivalent to about USD 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 NOK 10,000 (about USD 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.[22] 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. In 2007, the minimal reporting requirement was raised to NOK 100,000 (about USD 11,000).

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.

As explained above, the threshold for reporting lottery prizes was increased in 2007. 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 are to be considered as unexpected. In particular, we 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 USD 1,100 to 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.

2.4 Descriptive statistics

Table 1 displays summary statistics for a random sample of non-winners and our sample of winners between 1994 and 2006. For winners, all characteristics are measured in the year before they won.

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

[22] 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.

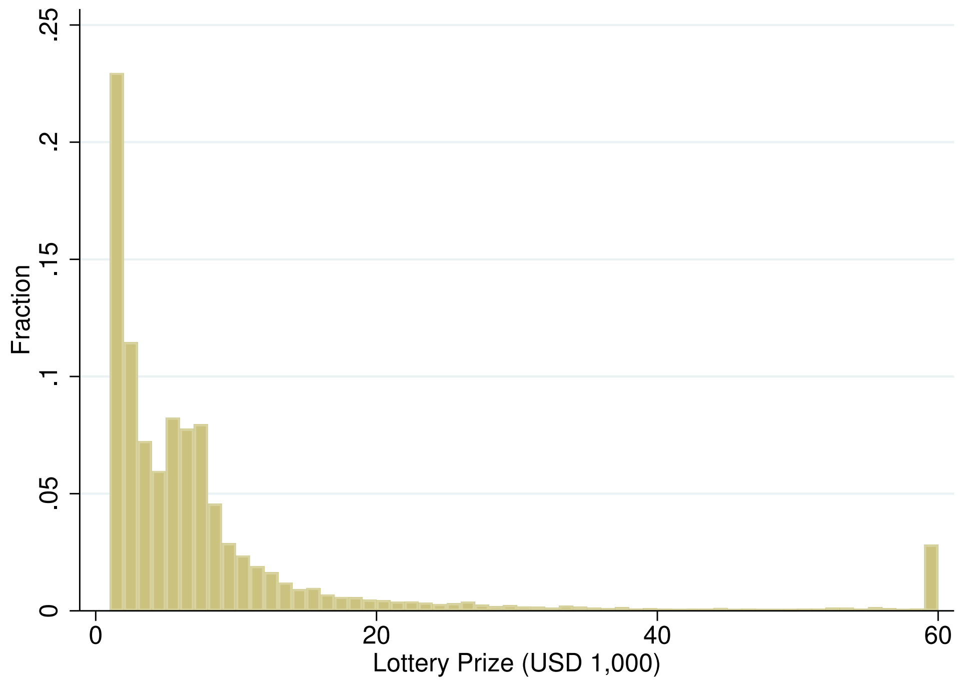

Notes: The figure shows the distribution of lottery prizes, denoted in USD, year-2000 prices. Each bin is USD 1,000 wide, starting from USD 1,000. The rightmost bar contains all prizes above USD 60,000.

For each non-winner, we have drawn a random year during the sample period to represent their observations. 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.

We see that winners and non-winners 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 does exist, is primarily due to housing wealth. Regarding balance sheet composition, Table 1 reveals that a higher share of winners own risky assets (29% against 25%), 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 non-winners. This pattern could suggest that households who win in lotteries are more risk tolerant than non-winners, but the observed differences are small and far from any levels that would raise the concern that winners exercise fundamentally different consumption behavior than others.

The final two rows of Table 1 display naively estimated marginal propensities to consume out of after-tax income or received inheritance. These estimates are not to be interpreted structurally. They simply are the resultant coefficients from regressing consumption on contemporaneous income

Table 1: Summary statistics, 1994-2006.

	Non-winners (N = 2,999,823)		Winners (N = 23,728)	
Variable	Mean	St. Dev.	Mean	St. Dev.
Age_t	49.69	(19.54)	50.61	(15.13)
Year_t	2000.92	(4.26)	2000.31	(3.34)
Household size_t	2.41	(1.39)	1.91	(1.14)
No. of children under 18_t	0.53	(0.94)	0.30	(0.71)
Years of education_t	12.99	(3.08)	12.55	(2.60)
Income after tax_{t-1}	25.97	(91.79)	24.11	(11.47)
Salary_{t-1}	22.16	(27.64)	23.49	(20.86)
Consumption_{t-1}	20.63	(15.86)	20.83	(13.76)
Lottery_t	.	.	9.24	(16.13)
Net wealth_{t-1}	79.20	(260.10)	77.26	(100.51)
Debt_{t-1}	36.04	(113.37)	30.89	(39.79)
Cars & boats_{t-1}	2.70	(7.48)	3.43	(6.03)
Housing wealth_{t-1}	95.18	(204.99)	93.60	(95.36)
Liquid assets_{t-1}	20.05	(163.57)	14.55	(25.51)
Deposits_{t-1}	15.92	(50.71)	12.30	(22.28)
Stocks_{t-1}	1.89	(144.73)	0.63	(4.72)
Bonds_{t-1}	1.05	(25.77)	0.62	(4.86)
Mutual funds_{t-1}	1.19	(12.02)	1.00	(4.45)
Risky share of balance sheet_{t-1}	0.07	(0.18)	0.08	(0.20)
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: Non-winners are defined as households that did not win a prize during the years 1994 to 2006. Among non-winners, we randomly select one household-year observation in our sample. For winners, we select the year prior to winning. Monetary amounts are CPI-adjusted to the year 2000 and then converted to (thousands of) USD using the mean exchange rate in the year 2000. 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 shows 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.

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.

2.5 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

$$C_{i,t} = \beta_0 + \beta_1 lottery_{i,t} + \beta_2 X_{i,t} + \alpha_i + \tau_t + u_{i,t}, \tag{3}$$

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, $X_{i,t}$ 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. 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.[23] 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.2, we derive the OLS-weights, and in Appendix A.3 we present a simple simulation exercise that illustrates this point clearly. 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.

2.6 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.[24]

[23] When we estimate dynamic responses, i.e. the effect of $lottery_{i,t}$ on $C_{i,t+j}$, we drop household-year observations from t to $t+j-1$ and after $t+j$.

[24] 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

Table 2: Predictability of lottery prize size

Dependent variable: Lottery Prize$_t$			
Regressors	(1)	(2)	(3)
Consumption$_{t-1}$	0.022		0.021
	(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$_{t-1}$	0.864		0.916
	(0.695)		(0.696)
Age$_t$		-0.019*	-0.006
		(0.008)	(0.010)
Household size$_t$		0.114	0.384
		(0.481)	(0.650)
Household size$_t^2$		0.031	-0.003
		(0.100)	(0.137)
No. of children under 18$_t$		0.175	-0.052
		(0.246)	(0.356)
R-squared	0.009	0.008	0.009
Partial R-squared of regressors	0.005	0.004	0.006
N	14,743	23,728	14,743

Notes: Each column represents a separately estimated regression of lottery prize among winners on predetermined characteristics. All regressions include time-fixed effects. Partial R-squared of regressors shows the increase in R-squared by adding the regressors to a specification with only time-fixed effects. Robust standard errors are in parentheses. *, **, *** denote significance at the 5, 1, and 0.1 percent level, respectively.

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.5. Neither consumption, deposits, stocks, bonds and mutual funds, nor debt evolve differently than normal in the years before winning. Appendix A.5 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

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.4 for details.

extent to which our prizes can be considered exogenous. Column (1) focuses on lagged values of consumption and balance sheet 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 and income. However, these associations with the amount won are small. A one-thousand-dollar increase in income predicts a fifty-dollar reduction in prize size, a one-thousand-dollar increase in debt predicts a thirteen-dollar increase in prize size, and a one-thousand-dollar increase in consumption predicts a twenty-dollar 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.

3 Consumption and savings responses to lottery prizes

3.1 Static responses

Existing studies that estimate consumption responses to income shocks utilize a variety of slightly different econometric specifications. To set the stage for our analysis of MPC heterogeneity, we first consider average estimates from the three main specifications considered in this literature. In addition to (3), the specifications are

$$\Delta C_{i,t} = \beta_0 + \beta_1 lottery_{i,t} + \beta_2 X_{i,t} + \alpha_i + \tau_t + u_{i,t} \tag{4}$$

$$C_{i,t} = \beta_0 + \beta_1 lottery_{i,t} + \beta_2 X_{i,t} + \beta_3 C_{i,t-1} + \alpha_i + \tau_t + u_{i,t} \tag{5}$$

where Δ is the one-year difference operator. Equation (4) is the difference estimator, while equation (5) is the dynamic estimator (including lagged consumption). Our coefficient of interest is still β_1.

When interpreting the results that follow, 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.2. These estimates are to be considered a starting point, as our study's contribution lies in estimating dynamic consumption responses and in dissecting how the average responses vary with the amount won and observable characteristics.

Table 3 shows our estimates of the within-year consumption response to a lottery prize using specifications (3) to (5), in the table referred to as "Levels," "Differences," and "Dynamic" respectively. "Dynamic" is estimated using the instrumental variable method proposed by Arellano and Bover (1995) and Blundell and Bond (1998). In addition, we distinguish between OLS and LAD estimates. The latter is based on the least absolute deviation estimator, or the median estimate as opposed to the mean estimate which the OLS provides. The LAD estimator is less sensitive to outliers. The first column in Table 3 controls for time-fixed effects only, which is necessary due

Table 3: The MPC out of lottery prizes.

	Specifications:		
	(1)	(2)	(3)
Levels (OLS)	0.587	0.523	0.520
(N = 93,627)	(0.013)	(0.014)	(0.014)
Differences (OLS)	0.504	0.504	0.501
(N = 59,906)	(0.018)	(0.021)	(0.021)
Levels (LAD)	0.609	0.485	.
(N = 93,627)	(0.015)	(0.017)	.
Differences (LAD)	0.467	0.460	.
(N = 59,906)	(0.015)	(0.024)	.
Dynamic (OLS-IV)[1]	.	0.506	0.504
(N = 59,906)	.	(0.021)	(0.021)
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 corresponds to the benchmark specification (3), differences corresponds to specification (4), and dynamic corresponds to specification (5). Additional controls are: age, age^2, age^3, age^4, household size, household size^2, no. of children under 18, and after tax income$_{t-1}$. 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). [1]We do not include household-fixed effects in the OLS-IV estimator.

to common trends in consumption and prize size. As we move horizontally across the table, we gradually add controls for individual-fixed effects and household characteristics.

If prizes were perfectly random, additional controls beyond year-fixed effects would be superfluous. Instead, we see that when individual-fixed effects are added, the point estimates in the levels specifications drop somewhat. The OLS estimate falls from 0.59 to 0.52, whereas the LAD estimate falls from 0.61 to 0.49. In part, these drops reflect that winners of higher prizes typically consume a little more also in years when they do not win, as we saw in Table 2. Consistent with the latter observation, the specifications in differences and with lagged consumption take the historical individual-fixed association between consumption and prize size into account by construction, so they do not change when individual-fixed effects are added.[25]

An important take-away 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 be far beyond the influence of the variables we observe. Because these observables span the main

[25]Adding individual fixed effects in specifications (4) and (5) might in principle have affected the estimates in Table 3 if consumption *growth* (not levels) were systematically correlated with prize size, but no such correlation is present, as reflected by the stability of the point estimates.

candidates suggested by economic theory for understanding consumption, we regard it as reasonable to interpret the results in Table 3 as causal.

Given that a lottery prize represents a transitory income shock to the household, the point estimates might seem high. Most models of household behavior suggest a substantially smaller instantaneous marginal propensity to consume out of transitory income shocks, typically in the range between 0.05 and 0.25 for non-durable consumption. For instance, the complete market infinitely-lived household model suggests a non-durables 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 Carroll 2001, p. 26), 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 U.S. tax rebate. In contrast, our estimates are *not* large compared to what existing empirical articles on transitory income shocks typically find. For example, in the literature on U.S. tax rebates, Johnson et al. (2006) and Parker et al. (2013) find that within a year households spend between 0.50 and 0.90 per dollar received on total consumption.

Importantly, when comparing our point estimates with existing evidence and models, one must bear two specific properties of our data in mind. First, we are estimating the response of consumption *expenditure*, including durables as well as non-durables.[26] Generally, we would expect a greater MPC once durables are included. Second, we are studying responses within the entire year of winning. Again, this will raise the expenditure response above the MPCs in structural models of a higher frequency. More specifically, if we assume that lottery prizes are uniformly spread 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 (Appendix A.6).[27] We return to time-aggregation at the end of Section 3.2, after estimating dynamic responses.

For the remainder of this paper, we concentrate on OLS estimates from the level specification in equation (3), as this has the clearest interpretation and allows us the largest sample by not requiring winners to be observed in consecutive years. Sample size will be valuable when we later explore dynamic responses and narrowly defined subgroups of the population. Furthermore, since the inclusion of additional controls beyond time- and household-fixed effects had no impact on point estimates, we include only time- and household-fixed effects going forward.

In addition to investigating the within-year MPC, we estimate responses of different balance

[26] The aggregate share of durables in total household consumption expenditure lies around 15% 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 point. These results are available upon request.

[27] In the existing literature, only Kaplan et al. (2018) report six-month MPCs. Their model predicts a non-durable consumption response around USD 300 after a transfer of USD 1,000.

Table 4: The propensities to save and spend out of lottery prizes.

	Dependent variable:				
	(1) Consumption	(2) Cars & boats	(3) Deposits	(4) Stocks, bonds & mutual funds	(5) Debt
Lottery_t	0.523	0.029	0.422	0.058	-0.073
	(0.014)	(0.003)	(0.013)	(0.007)	(0.008)
N	93,627	93,627	93,627	93,627	93,627

Notes: Each column represents a separate regression of equation (6). The dependent variables are consumption, change in cars and boats, change in deposits, change in stocks, bonds, and mutual funds, and change in debt. A negative number for the debt response means repayment of debt. Controls include time-fixed and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level. Behind the estimates for cars and boats lie a 2 percentage point increase in the share of households that make a purchase.

sheet components and car and boat purchases, utilizing the specification

$$\Delta Z_{i,t} = \beta_0 + \beta_1 lottery_{i,t} + \alpha_i + \tau_t + u_{i,t}, \tag{6}$$

where $\Delta Z_{i,t}$ is the change in balance sheet component Z from $t-1$ to t. Results are reported in Table 4.

To facilitate comparison, Column (1) in Table 4 restates the consumption response of 0.52. The next columns report the responses of cars and boat purchases, deposits, the sum of stocks, bonds and mutual funds, and debt. Note that purchases of cars and boats are included in the consumption response of Column (1). Note also that because we run separate regressions for each object of interest, the estimates of consumption (Column (1)) and saving (Columns (3) - (5)) need not sum to one. Our results suggest that within the year of winning, the mean response is to save about 42 percent of the lottery prize in deposits, 6 percent in stocks, bonds and mutual funds, while about 7 percent of the lottery prize is used to repay debt.

For cars and boats, we observe a small within-year expenditure increase of around 3 percent of the amount won. Naturally, this number is an average across the majority of households who do not purchase any vehicle at all, and those who cross the extensive margin and buy a car or boat. 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.[28] Among those 13 percent, the amount spent is approximately twice the amount won, on average. The latter is a simple illustration of how discrete spending choice can easily make the response of expenditure greater than the prize.

[28]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.

3.2 Dynamic responses

Obviously, how strongly consumption and savings respond to an income shock depends on the time frame one has in mind. The results above focus on within-year effects, but of equal interest is how consumption responds over time. For instance, recent work by Auclert et al. (2018) emphasizes that the intertemporal consumption response is key to understanding how certain aggregate impulses such as tax changes propagate to the macro economy. Therefore, we here move on to estimate impulse responses over the five years after winning. We utilize the following specification:

$$C_{i,t+k} = \beta_0 + \beta_{1,k} lottery_{i,t} + \alpha_i + \tau_{t+k} + u_{i,t+k} \text{ for } k = -4, -3, ..., 5 \tag{7}$$

The $\beta_{1,k}$'s are the main coefficients of interest. 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 $\beta_{1,k}$'s. Importantly, when we estimate the consumption responses in the years after winning, we include the consumption observations up to the year a household won in the sample. As always, our sample consists of winners only. Hence $\beta_{1,k}$ is identified by comparing consumption k years after winning to consumption before winning. Moreover, we estimate equation (7) separately for each year, instead of estimating the responses for all k jointly. The reason is that the latter would require we observe consumption for 10 straight years around winning, a restriction that greatly reduces our sample size since we do not impute consumption for households who transact in the housing market, dissolve, etc.

In addition to the intertemporal MPCs, we study the dynamic responses of the balance sheet. We therefore estimate an equation similar to (7) for financial asset accumulation:

$$\Delta Z_{i,t+k} = \beta_0 + \beta_{1,k} lottery_{i,t} + \alpha_i + \tau_{t+k} + u_{i,t+k} \text{ for } k = -4, -3, ..., 5, \tag{8}$$

where $\Delta Z_{i,t+k}$ is the change in balance sheet component Z from period t to $t+k$.

Figure 2 shows the dynamic responses of consumption, deposits, the sum of stocks, bonds and mutual funds, and debt. The top panel displays the flows and the bottom panel displays the cumulative effects. The estimated within-year responses are the same as those discussed above. Beyond constituting quantitative moments to discipline economic models of consumption dynamics, the estimates provide four qualitative findings of particular interest. First, the five-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 after 5 years 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.

Figure 2: Dynamic household responses to lottery prizes.

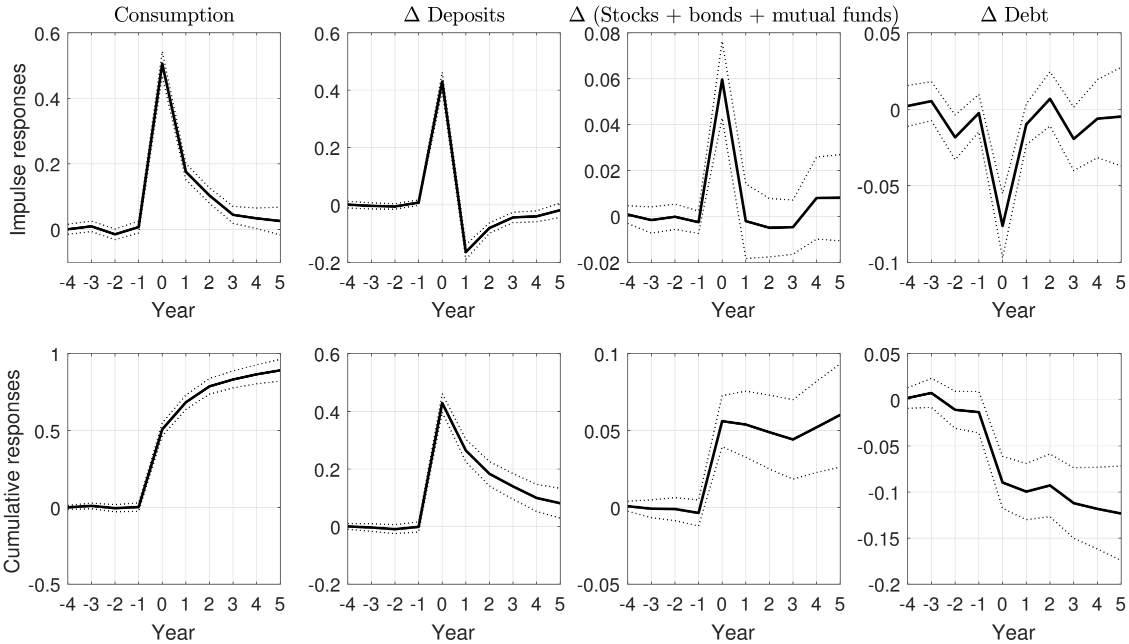

Notes: Each point is estimated as a separate regression of (7) or (8). Controls include time-fixed and household-fixed effects. Dotted lines represent 95 percent confidence intervals using robust standard errors clustered at the household level. The cumulative responses in the bottom panel are the sum of year-specific responses in the top panel. Standard errors of estimated cumulative effects are obtained through Monte Carlo simulations.

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 extra expenditure.[29] The sharp consumption decrease in the immediate year after winning might be due to durable goods purchases in the win-year, but the durables we observe, namely cars and boats, play a minor role here. As reported above, about 3 percent of the lottery prize is spent on car and boat purchases in the win-year.

Third, even though the consumption response drops rapidly from year one to year two, it does not drop nearly as far as it would if households could be coarsely split in to "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.[30] Auclert et al. (2018) discuss in more detail how such dynamic responses can be used to distinguish between existing consumption-savings models, and in Section 7 we relate them to models of discrete durable goods purchases.

[29] 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, 2017).

[30] 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.

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 models where there are costs of altering amortization schedules.

We may use the dynamic responses to deal with time aggregation, and to convert our estimates into MPCs at different frequencies. This conversion allows closer comparison to structural models and other empirical studies, but it requires assumptions. We first assume that prizes are uniformly distributed over the year of winning. Next, motivated by the patterns in Figure 2, we assume that average MPCs obey the functional form $mpc(t) = \theta_1 t^{\theta_2}$, where θ_1 and θ_2 are parameters to be determined, and t is time since winning. We select θ_1 and θ_2 so as to minimize the distance between the directly estimated MPCs in Figure 2 and the functional form imposed on $mpc(t)$. The outcome is $\theta_1 = 0.629$ and $\theta_2 = 0.214$ (when t is measured in years), with a resultant time-profile that fits Figure 2 closely. This exercise implies a quarterly MPC of 0.37, a six-month MPC of 0.54, and a one-year MPC of 0.63. Further details are provided in Appendix A.6.

4 Shock size, household characteristics, and MPCs

While the above estimates constitute a natural starting point, they mask heterogeneity in how households respond to income shocks. We here turn to this heterogeneity, and ask which observable variables are associated with cross-sectional variation in MPCs. Our approach is largely motivated by buffer-stock saving models. We proceed in three steps. First, we estimate unconditional covariances between MPCs and different covariates suggested in the literature. Second, we use saturated regressions where we can control for many plausible explanations of MPC heterogeneity at once. The ambition here is primarily qualitative, namely to assess which explanatory variables that matter individually. Third, in Section 5 we focus on the variables that stood out in step two, with the ambition to establish their quantitative influence.

4.1 Unconditional covariances

The cross-sectional covariances between individual MPCs and various household characteristics have recently been emphasized as key to understanding different macroeconomic phenomena. Notably, the issue here is not really causality, but simply correlation. For instance, Auclert (2016), calculates a "sufficient statistic" for the role of redistribution in the transmission of monetary policy shocks to aggregate consumption. Key to this statistic is the covariance between individuals' MPC and their balance sheet exposure to interest rate changes, coined "unhedged interest rate exposure" (URE). Similarly, the covariance between individuals' MPC and their balance sheet exposure to inflation, the net nominal position (NNP), is important to determine the role of redistribution in inflation transmission. Another example is Berger et al. (2018), who address how house price changes are likely to affect individual-level consumption, and derive a statistic emphasizing house-

Table 5: Unconditional covariances between MPCs and household characteristics.

Z_{t-1}	Cov(MPC$_{it}$,$Z_{it-1}/\bar{C}_{t-1}$)	s.e.
Income$_{t-1}$	-0.002	(0.006)
Liquid Assets$_{t-1}$	-0.051	(0.014)
Deposits$_{t-1}$	-0.049	(0.013)
Debt$_{t-1}$	0.045	(0.029)
Housing wealth$_{t-1}$	0.023	(0.059)
Net wealth$_{t-1}$	-0.088	(0.048)
Age$_t$	-0.054	(0.010)
URE$_{t-1}$	-0.109	(0.040)
NNP$_{t-1}$	-0.084	(0.029)

Notes: We estimate these covariances by (1) dividing our sample into equal-sized bins based on the predetermined characteristic; (2) estimate equation (3) within each bin i; and (3) simulate 10,000 covariances using the estimated consumption responses, their respective standard error, and the mean value of the predetermined characteristic within each bin. $\bar{C}_{t-1}$ is the cross-sectional mean level of consumption in the year before winning. We redo the procedure for all bin sizes between 10 and 75, giving us a total of 660,000 covariances. The reported covariances are the mean and standard deviations in the sample of 660,000 covariances. Liquid assets = deposits + stocks + bonds + mutual funds; Unhedged interest rate exposure (URE) = income after tax - consumption + liquid assets - debt; Net nominal position (NNP) = deposits + bonds - debt.

holds' MPC out of transitory income and their housing wealth. By implication, the cross-sectional covariance between MPCs and housing wealth will be key to how house price shocks affect aggregates. A third example is Auclert and Rognlie (2016), who address how inequality might influence aggregate demand, and then emphasize the cross-sectional covariance between MPCs and income.

Direct empirical evidence on such covariances is in short supply. We therefore display the unconditional covariances between MPCs and the key household characteristics that we observe. To this end, we first rank and group households by the variable in question. Next, we estimate the consumption responses to lottery prizes within each group using equation (3). With these estimates, their standard errors, and the mean value of the relevant variable within the bin at hand, we simulate and compute each covariance 10,000 times. We do this for 66 different stratifications using 10 to 75 bins, giving us a sample of 660,000 covariances for each variable. Table 5 reports the average covariance and the sample standard error generated by this bootstrapping procedure. Note that each covariance is computed after scaling the characteristic by average observed consumption.

In the first line of Table 5, we see that the marginal propensity to consume out of a lottery prize is virtually unrelated to income. The covariance is insignificant by any conventional measure. In contrast, there is a visible and significant negative covariance between MPCs and liquid assets, deposits, net wealth, NNP,[31] and in particular between MPCs and our crude measure of URE. The latter is simply calculated as income minus consumption plus liquid assets minus debt, an approximation that is motivated by the fact that nearly all mortgage debt in Norway has a flexible interest rate.[32] For debt, the covariance is positive, but with a relatively high standard error. Housing wealth is approximately unrelated to MPCs. Age co-varies negatively with MPCs.

4.2 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:

$$C_{i,t} = \beta_0 + \beta_1 lottery_{i,t} + \beta_2 lottery_{i,t} * Z_{i,t-1} + \beta_3 Z_{i,t-1} + \alpha_i + \tau_t + u_{i,t} \tag{9}$$

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. 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 (9) both with each interaction term included in separate regressions (denoted Univariate in Table 6) and with all interactions included in a multivariate regression (Multivariate in Table 6). The latter estimates are of our primary interest, as they indicate which factors that affect MPCs directly, over and above their correlation with the other explanatory variables. Table 6 presents the estimated interaction coefficients. As the objective here is to single out which variables that matter qualitatively, we mark statistical significance by asterisks. To facilitate interpretation of the magnitudes involved, each column also contains standardized coefficients in square brackets. These show the effect of a one standard deviation increase in the interaction variable in question.

Among the candidates considered, prize size, liquid assets, and age stand out as the main observable factors associated with the magnitude of households' MPCs. As the second column shows, all three are statistically significant also when every interaction term is included. Neither income, net wealth, debt, education, the portfolio share held in risky assets, nor household size, have similar significant effects.

[31] We compute the net nominal position as the sum of deposits and bonds minus liabilities.

[32] Auclert (2016) estimate a similar covariance between URE and MPC of −0.06, based on Italian survey data from Jappelli and Pistaferri (2014).

Table 6: The MPC out of lottery prizes: Interaction effects.

	(1) Univariate		(2) Multivariate		(3) Multivariate, no-risky-assets	
Lottery_t^2	-0.001*** (0.000)	[-0.008]	-0.001*** (0.000)	[-0.007]	-0.001** (0.000)	[-0.007]
Lottery_t*liquid assets_{t-1}	-0.003*** (0.001)	[-0.093]	-0.003*** (0.001)	[-0.081]	-0.003** (0.001)	[-0.108]
Lottery_t*income_{t-1}	0.001 (0.001)	[0.015]	-0.002 (0.002)	[-0.025]	-0.005 (0.003)	[-0.057]
Lottery_t*net wealth_{t-1}	-0.000 (0.000)	[-0.016]	0.000 (0.000)	[0.031]	0.000 (0.000)	[0.033]
Lottery_t*debt_{t-1}	0.001*** (0.000)	[0.059]	0.000 (0.000)	[0.022]	0.002* (0.001)	[0.083]
Lottery_t*education_t	0.015** (0.005)	[0.037]	0.008 (0.006)	[0.019]	0.006 (0.010)	[0.015]
Lottery_t*risky share_{t-1}	-0.006 (0.060)	[-0.001]	-0.045 (0.062)	[-0.009]	. .	. .
Lottery_t*household size_t	0.034* (0.014)	[0.038]	0.026 (0.015)	[0.029]	0.038 (0.023)	[0.042]
Lottery_t*age_t	-0.005*** (0.001)	[-0.081]	-0.005*** (0.001)	[-0.076]	-0.005** (0.002)	[-0.070]
Observations	93,627	93,627	93,627	93,627	40,852	40,852

Notes: This table displays the estimated interaction terms from the estimation of Equation (9). The univariate column presents the interaction coefficients when we include each interaction term with the lottery prize ($lottery_{i,t} * Z_{i,t-1}$) individually in separate regressions (together with the remaining variables of Equation (9)), while the multivariate columns present results when we include all interaction terms in the same regression. 'no-risky-assets' refers to the sample where 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 (9)). The standard errors in parentheses are robust and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively. Standardized regression coefficients are reported in squared brackets.

Column 2 shows that when all variables are controlled for, one standard deviation increases in prize size, liquidity and age are each associated with MPC decreases of 0.7, 8 and 8 percentage points, respectively. As will become clear below, the quantitatively small effect of prize size is due to the fact that the estimator is overweighting winners of high prizes, whereas MPCs vary more strongly with prize size among winners of small amounts.

As detailed in Section 2.2, 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 sample, it has the benefit that the main source of measurement error in imputed consumption is absent. The third column of Table 6 shows that within this sub-sample, our main results regarding heterogeneity remain the same.

Rather than including the sum of deposits, stocks, bonds, and mutual funds, we have also considered each of these sub-components separately. These results are provided in the Table A.4 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 6, while the point estimate for stocks (which only 30 percent of our sample hold) is higher, yet imprecise. As we proceed, we will focus on the impact of the entire basket of liquid assets. As the deposit estimate implies, our results would remain largely unchanged if we instead had focused on deposits only.

5 Dissecting the role of shock size, liquidity, and age

We now turn to further dissecting and quantifying the influence of the three variables that stood out in Section 4.2. To this end we group winners by amount won, liquid assets, and age. Stratification is done by quartile of the respective variable.[33]

5.1 Shock size

Table 7 displays how the responses of consumption, deposits, stocks, bonds and mutual funds, and debt vary with prize size. Each estimate is obtained within the respective quartile. The quartiles are USD 1,100-2,070, USD 2,070-5,200, USD 5,200-8,300, and USD 8,300-150,000.[34]

Given that our paper is the first to provide direct evidence on this issue, we first summarize the main reasons why the magnitude of an income shock might matter for households' consumption response. First, if a household initially is credit constrained, a sufficiently large income shock will eventually relax the constraint, at which point the household chooses to save more of the income innovation. More generally, standard precautionary savings models imply that the policy function for consumption is concave in wealth, either because of borrowing constraints (Carroll and Kimball, 2001; Holm, 2018) or risk (Carroll and Kimball, 1996). This concavity in wealth implies that the MPC is smaller for greater income innovations. Second, if purchases of high-return

[33]Ideally, we would have chosen an even more fine-grained stratification here. However, to obtain statistical power we need to include enough observations in each group.

[34]As Table 7 reveals, each size quartile has a different number of observations. The reason is that each winner is observed repeatedly, and some households are present longer in our sample than others.

Table 7: Heterogeneous household responses. Quartiles of lottery prize.

	Lottery prize size quartile			
Dependent variable	(1)	(2)	(3)	(4)
Consumption	1.354	0.957	0.700	0.511
	(0.161)	(0.075)	(0.040)	(0.015)
Deposits	0.569	0.371	0.483	0.417
	(0.109)	(0.050)	(0.027)	(0.014)
Stocks, bonds & mutual funds	0.022	0.082	0.007	0.060
	(0.047)	(0.031)	(0.013)	(0.007)
Debt	0.292	0.203	0.010	-0.081
	(0.124)	(0.058)	(0.032)	(0.009)
Observations	23,109	23,366	24,633	22,519

Notes: Each coefficient represents a separate regression of equation (3) or (6) within a prize size quartile. The groups are: 1 (USD 1,100 - 2,070), 2 (USD 2,070 - 5,200), 3 (USD 5,200 - 8,300), and 4 (USD 8,300 - 150,000). Controls include time-fixed and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

assets involve discrete transaction costs, then the rate of return on savings for a received prize will effectively increase with the amount won and motivate high-prize winners to save more. Third, several consumption decisions contain an element of discrete choice. In particular, this is true for durable goods purchases, but also for several non-durables like vacations. Winners of small prizes might therefore choose to spend their entire income innovation, and possibly even more if they can deplete savings or borrow, to make a discrete purchase. The larger the amount won, the less likely such lumpy purchases will dominate spending decisions. On the other hand, discrete choice could in principle work in the opposite direction too, as the probability that a household actually makes a purchase increases with shock size.

In the top left corner of Table 7 we see that winners of relatively small amounts tend to spend all they win, or even more than the prize itself. The mean estimate is 1.35 to the dollar won. This is likely to reflect discrete expenditure choice, as explained above.[35] Moving down the table, we also see that the average debt response in this group is positive, suggesting that several of the low-prize winners top their prize up with credit or lower debt repayment. Surprisingly though, the estimated deposit coefficient is high too, and the sum of coefficients (minus debt) exceeds unity. This, together with the fact that all estimates in the lowest prize quartile come with relatively high standard errors, suggests that the exact point estimates in this group should be interpreted with caution.[36]

[35] Relatedly, Aaronson, Agarwal, and French (2012) find that minimum wage hikes trigger spending responses in excess of income responses, and evaluate how structural models with durable goods can account for this pattern.

[36] 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.

Moving rightward in Table 7, we see that the consumption response declines with the amount won. Still, the point estimate in the top quartile is remarkably high, indicating that even winners of more than USD 8,300 spend on average half their prize within the year of winning. There is no clear monotonic pattern for deposits or for stocks, bonds, and mutual funds. But for debt, the bottom row implies that it is only among winners of relatively large amounts, the highest prize quartile, that the average response is to cut debt.

Upon reading Table 7, 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, making the estimates imprecise. In higher quartiles, the treatment is stronger and the estimates are more reliable. Second, our estimate in the highest size 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 pulled toward the MPC of high-prize winners. We illustrate this point in more detail in Appendix A.3.

5.2 Liquidity

Table 8 displays the within-year estimates across quartiles of liquid assets held in the year before winning. The cutoffs for liquidity quartiles vary with the year of winning. The cutoffs in 1994 (2006) are: USD 720 (1,720), USD 3,140 (7,160), and USD 10,710 (21,600).

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, it is deposits that most closely track the consumption pattern. 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. This stands to reason, as households with high initial liquid asset levels were already able to repay debt before they won.

All the above results indicate an association between liquidity and marginal propensities to consume and save, consistent with previous empirical studies such as Misra and Surico (2014) and Leth-Petersen (2010), and with standard economic theory. An interesting question then, is the extent to which this association differs from that between illiquid wealth and MPCs, as suggested by the recent wave of two-asset structural models following Kaplan and Violante (2014). To explore this issue, we sort households not only by liquid assets, but also group them into quartiles by their net illiquid assets (housing wealth minus debt). We then combine the two stratifications, leaving us with a total of 16 liquid-illiquid asset groups, and estimate the consumption response to lottery prizes within each of these 16 groups. The results are visualized in Figure 3.

The right-hand-side bar diagram in Figure 3 shows the density of prize winners within each of

Figure 3: Heterogeneous consumption responses. Net illiquid and liquid assets.

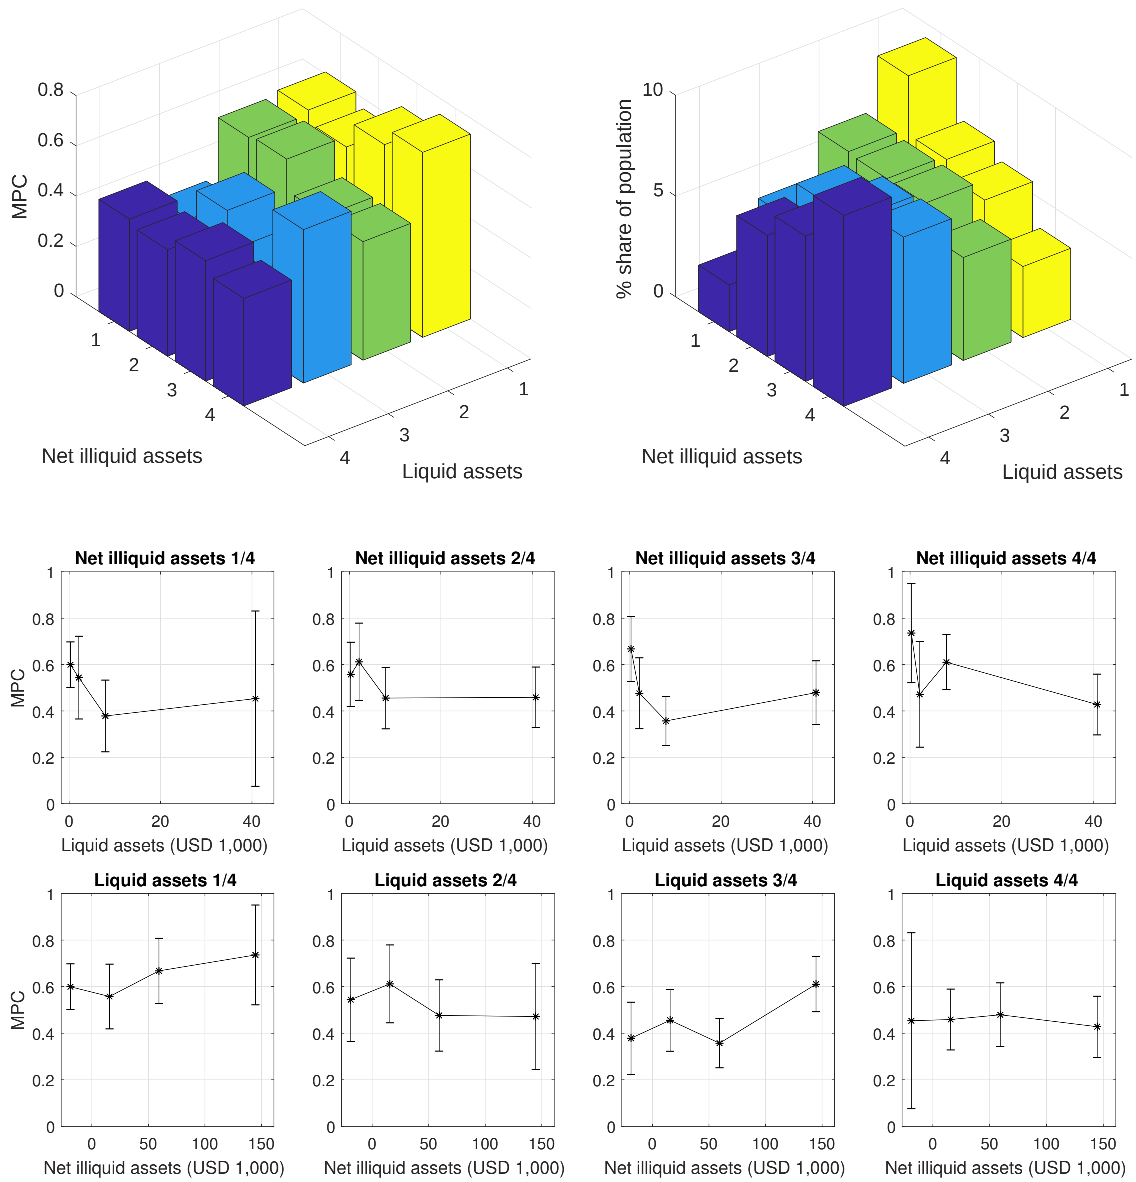

Notes: Each bar/point is estimated as a separate regression of equation (3). Liquid assets = deposits + stocks + bonds + mutual funds. Net illiquid assets = net wealth - liquid assets = housing wealth - debt. Controls include time-fixed and household-fixed effects. The standard errors are robust and clustered at the household level. Total N: 93,627.

Table 8: Heterogeneous household responses. Quartiles of liquid assets.

	Liquid asset quartile			
Dependent variable	(1)	(2)	(3)	(4)
Consumption	0.617	0.521	0.458	0.459
	(0.033)	(0.036)	(0.030)	(0.040)
Deposits	0.327	0.376	0.475	0.594
	(0.028)	(0.032)	(0.027)	(0.029)
Stocks, bonds & mutual funds	0.021	0.047	0.066	0.068
	(0.005)	(0.011)	(0.017)	(0.018)
Debt	-0.087	-0.128	-0.072	-0.010
	(0.020)	(0.025)	(0.018)	(0.015)
Observations	23,401	23,409	23,406	23,411

Notes: Each coefficient represents a separate regression of equation (3) or (6) within a liquid asset quartile. The cutoffs between liquidity quartiles increase over time. In 1994 (2006), the cutoffs are: USD 720 (1,720), USD 3,140 (7,160), and USD 10,710 (21,600). Controls include time-fixed and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level.

the liquid-illiquid asset categories. The shape of this distribution is hardly surprising. It spikes in the two corners where households hold high levels of both liquid and illiquid wealth or low levels of both liquid and illiquid wealth. Few households are simultaneously in the highest liquid asset quartile and the lowest net illiquid asset quartile, to the north-west of the figure. A somewhat higher fraction are in the south-east corner, with high illiquid wealth but low holdings of liquid assets. The four groups that simultaneously are in the upper two illiquid wealth quartiles and in the bottom two liquid asset quartiles, constitute more than 20 percent of all the winners in our sample.[37]

Figure 3's left-hand-side bar diagram shows how MPCs vary in the two dimensions. When viewing the diagram, one must bear in mind how the right-hand plot implies small samples, and hence imprecise estimates, in several strata. We see the same overall pattern reflected in Table 6, i.e., that MPCs tend to fall with liquid, but not with illiquid, wealth. Also, there is a tendency for households with high levels of both illiquid and liquid wealth to have relatively high MPCs. Still, the MPCs among households with high liquidity are above what a model like that in Kaplan et al. (2014) would imply. Hence, our estimates are consistent with the hypothesis that liquidity matters for MPC variation, but inconsistent with the hypothesis that illiquidity alone explains the high average MPC in our sample.

Below the two bar diagrams, estimates conditioned on both liquid and illiquid assets are presented in a complementary way. Standard errors are here displayed vertically, while the horizontal

[37] Kaplan, Violante, and Weidner (2014) define hand-to-mouth (HtM) households as households with liquid reserves less than half their monthly wage and wealthy as households with positive net illiquid assets. With this stratification, our distribution of households features 13.9 % Wealthy-HtM, 9.5 % Poor-HtM and 76.6 % Non-HtM. Their respective MPCs (with standard errors in parentheses) are W-Htm: 0.675 (0.052); P-HtM: 0.565 (0.049), and N-HtM: 0.500 (0.017).

axes show how much the strata differ by mean holdings of liquid and illiquid assets. Again we see the overall tendency for responses to fall with liquid, but not with illiquid assets, but we also see how the within-group estimates are too uncertain to justify more fine-grained conclusions.

5.3 Age

Table 9: Heterogeneous household responses. Quartiles of age.

	Age quartile			
Dependent variable	(1)	(2)	(3)	(4)
Consumption	0.578	0.527	0.532	0.436
	(0.028)	(0.030)	(0.028)	(0.026)
Deposits	0.342	0.390	0.413	0.537
	(0.028)	(0.022)	(0.027)	(0.026)
Stocks, bonds & mutual funds	0.066	0.070	0.053	0.046
	(0.013)	(0.017)	(0.011)	(0.013)
Debt	-0.072	-0.111	-0.091	-0.019
	(0.021)	(0.018)	(0.017)	(0.009)
Observations	18,416	23,364	24,719	27,128

Notes: Each coefficient represents a separate regression of equation (3) or (6) within an age quartile. The cutoffs between the quartiles are: 39, 51, and 63. Controls include time-fixed and household-fixed effects. The number of observations in each quartile varies because the stratification is conducted by age at the time of winning, and younger households are observed less frequently in the years before they won. The standard errors in parentheses are robust and clustered at the household level.

Table 9 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, 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 in the saving in stocks, bonds and mutual bonds, or debt repayment.

The tendency for MPCs to fall with age is at odds with a frictionless life-cycle model with a flat earnings profile and no bequests. Instead, it points toward extensions emphasized in the recent literature. Coupled with borrowing constraints, a realistic earnings profile would 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). The finding that MPCs decrease with age even among the older age groups might be caused by such a preference to increase the budget share of savings as one grows richer. We illustrate this point with a structural model in Section 7.

5.4 Conditional effects

Given that shock size, liquidity and age each matter for MPCs, a key quantitative question is how much they matter conditional on one another. We therefore proceed with a similar approach as we used to distinguish between the influence of liquid and illiquid assets in Figure 3 above. That is, we first stratify households by prize size and liquid assets, then by liquid assets and age, and finally by age and prize size. The results are displayed in Figures 4 and 5, and in Figure A.10.

Figure 4's right-hand side bar diagram displays the joint distribution of the amount won and liquidity in the year before winning. MPC estimates within each subgroup are shown in the diagram to the left. By construction, the 16 different groups need not be equally large, as we have grouped households by quartile in the liquidity and prize size distributions separately. Nevertheless, the right-hand panel in the figure shows that the strata are essentially equi-sized, consistent with the results in Table 2.

The pattern is stark: Within each liquidity quartile, the MPC decreases with the amount won. Within each size quartile, the MPC decreases with liquidity. The consumption expenditure response is almost 1.5 among winners who are in the lowest liquidity quartile and the lowest prize size quartile simultaneously (north-east corner). In contrast, in the jointly high-liquid-high-prize group (south-west corner), the MPC is below 0.5.

The plots below the bar diagrams add to this picture by displaying the magnitudes with which the groups differ along the horizontal axes and standard errors around point estimates. The upper plots show how MPCs vary with prize size within each of the four liquidity quartiles. The lower plots show the converse.

In Section 5.1 we explained how prize size might matter if liquidity constraints are present. Now, if this explanation drives the size effects in Table 7, we would expect to see a particularly strong size effect among the least liquid households. Consistent with the effects of liquidity constraints, the estimates indicate that liquid assets matter less among high-prize winners, and that prize size matters more for MPCs at lower liquidity levels. Yet, we also see that prize size co-varies with MPCs across the distribution of initial liquidity, even among the most liquid. Hence, the presence of liquidity constraints alone cannot entirely explain the role of shock size.

Figure 5 displays estimates across the joint distribution of prize size and age. The upper-right panel reveals a flat distribution across the two dimensions, again consistent with Table 2. The upper-left diagram and the bottom plot show that within all age groups considered, consumption responses decline with prize size. The association between MPCs and prize size is nonlinear, falling more sharply with the amount won at lower prizes. As we consider the other dimension, how MPCs vary with age within prize groups, we again see a declining relationship.

The patterns in Figures 4 and 5 are similar. This similarity is related to the fact that liquid wealth and age are highly correlated: Low-liquidity households tend to be relatively young, and high-liquidity households tend to be relatively old. Figure A.10 in the appendix displays the age-liquidity distribution and MPC estimates by age-liquidity quartiles. Age and liquidity are highly

Figure 4: Heterogeneous consumption responses. Lottery prize size and liquid assets.

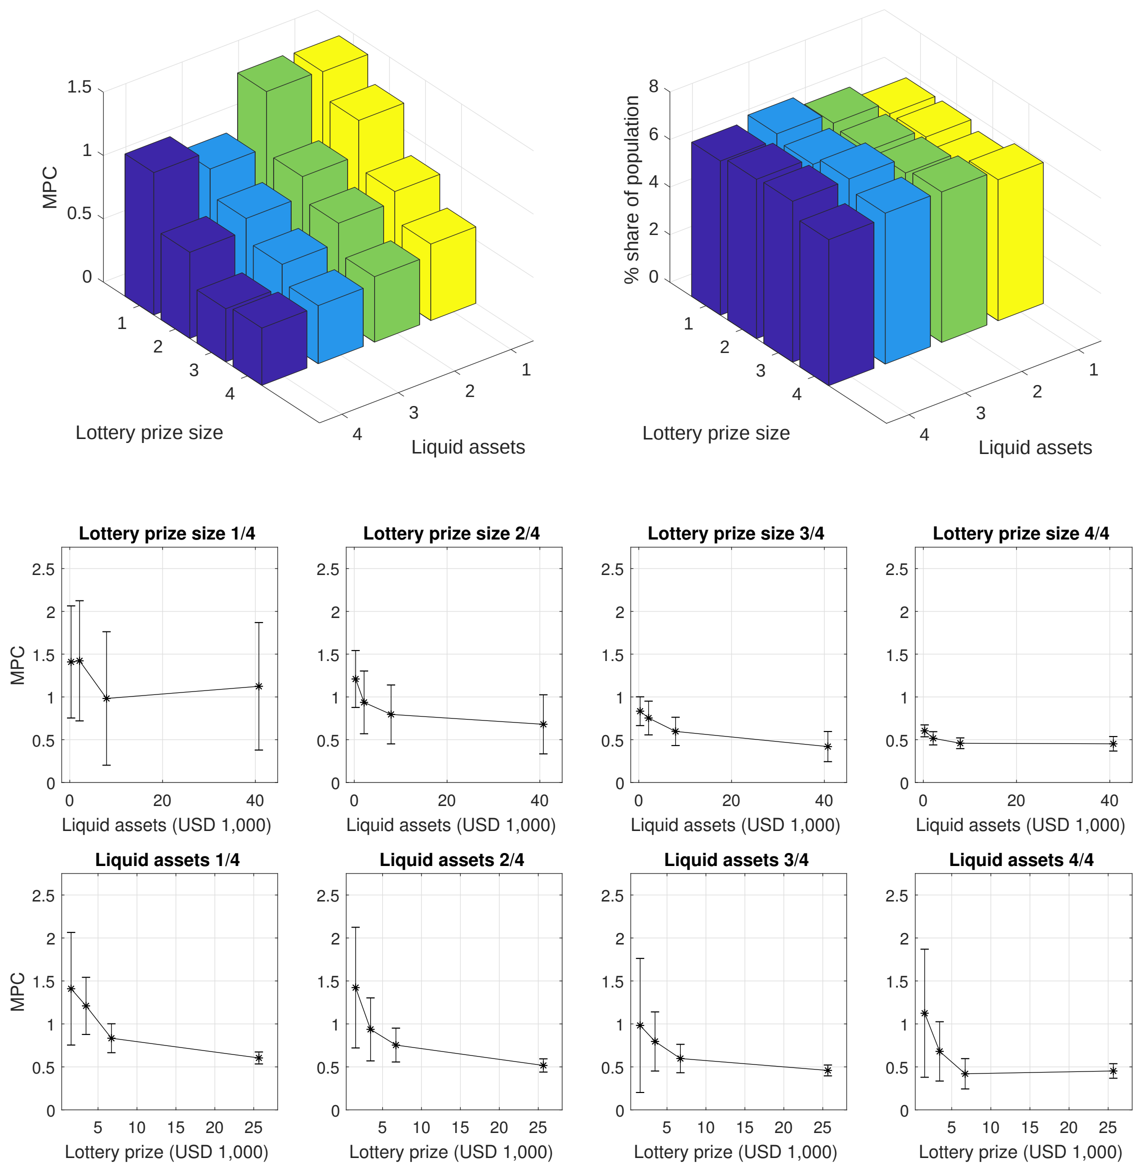

Notes: Each bar/point is estimated as a separate regression of equation (3). Controls include time-fixed and household-fixed effects. Standard errors are robust and clustered at the household-level. Total N: 93,627.

Figure 5: Heterogeneous consumption responses. Lottery prize size and age.

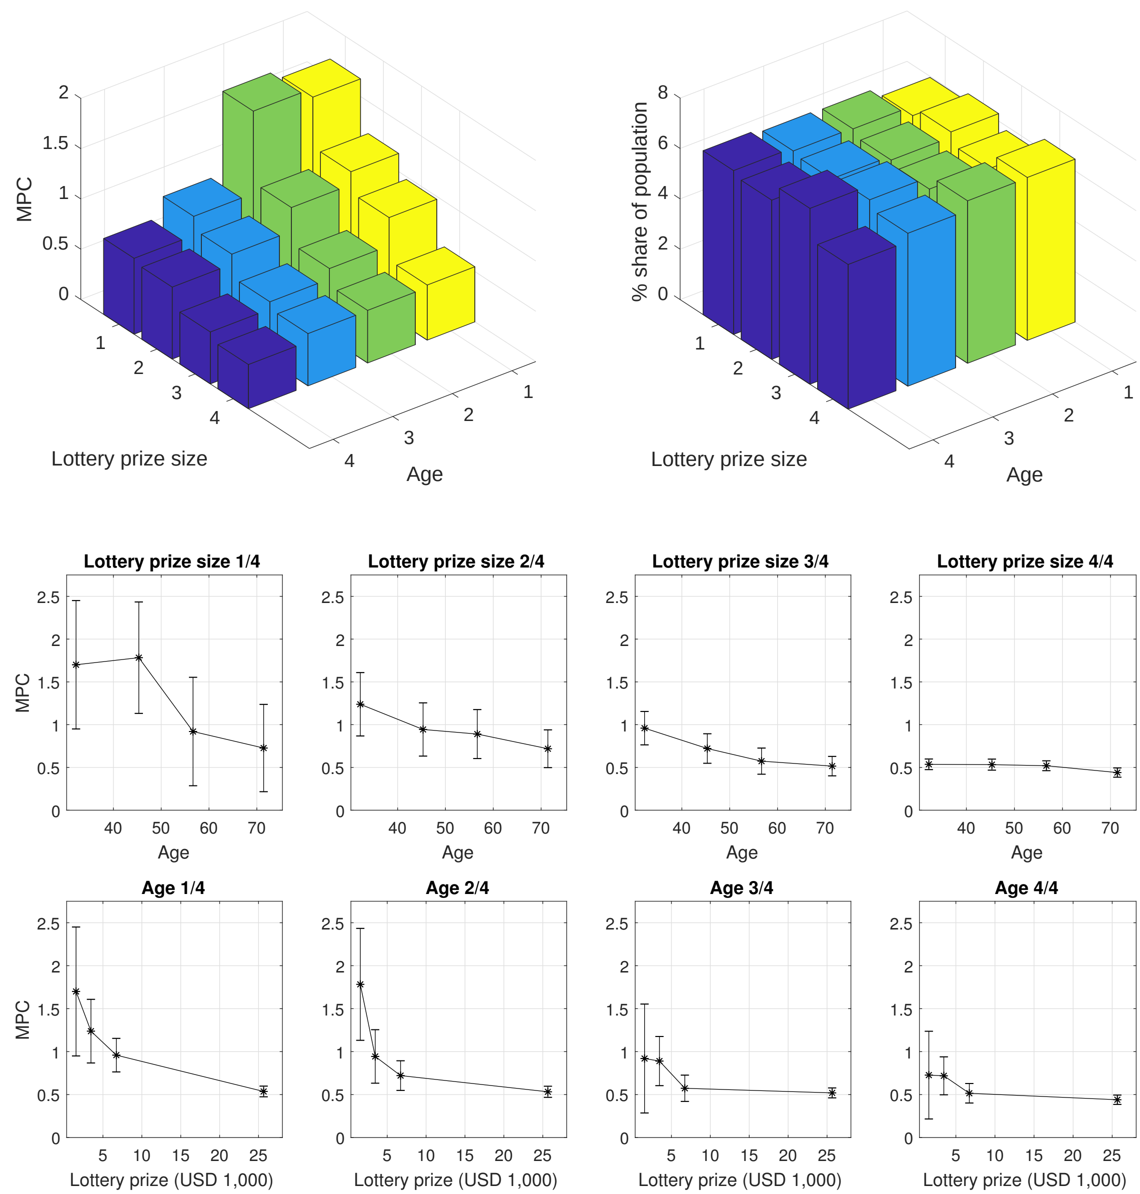

Notes: Each bar/point is estimated as a separate regression of equation (3). Controls include time-fixed and household-fixed effects. Standard errors are robust and clustered at the household-level. Total N: 93,627.

correlated, blurring more precise patterns than the average effects in Table 6. Figure A.10 visualizes how MPCs on average decline with age and liquidity separately, and more strongly so for liquidity. In addition, within age quartiles there is a tendency for the same non-linearity in liquid assets as we saw in Figures 3 and 4.

6 Robustness

We conduct extensive robustness tests, all documented in the Appendix. The upshot is that our main results regarding liquidity, age, and shock size hold irrespective of sample selection criteria and variable definitions. We here discuss the robustness checks we believe to be of most interest. We primarily focus on the results from variants of the interaction regression reported in Table 10, as this provides the most compact exposition of our findings. For each alternative, the table reports MPC estimates and interaction effects from variants of specifications (3) and (9) 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 (2004). The second main column of Table 10 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.[38] 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.[39] 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. The likely reason is that when squared as in the interaction term, variation in prize-to-permanent-income is swamped by permanent income.

Next, we control for an estimate of each household's typical co-movement between consumption and labor income, denoted MPC_i in the table. The motivation 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.[40] Our approach is explained in Appendix A.10. In short, we (i)

[38] 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.

[39] 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.

[40] Note the difference between individual-fixed effects in consumption *level*, which we control for with α_i

Table 10: Robustness. Interaction results under alternative assumptions and sample restrictions.

	(1) Benchmark		(2) Normalizing by permanent income		(3) Controlling for individual-fixed MPC	(4) Untrimmed sample		(5) Unconditional 2.5% trim on consumption	
Lottery_t	0.523*** (0.014)	0.883*** (0.075)	0.563*** (0.019)	0.648*** (0.115)	0.865*** (0.077)	0.712*** (0.014)	0.718*** (0.070)	0.353*** (0.009)	0.661*** (0.056)
Lottery_t^2		-0.001*** (0.000)		0.000 (0.005)	-0.001** (0.000)		0.001*** (0.000)		-0.001*** (0.000)
Lottery_t*liquid assets_{t-1}		-0.003*** (0.001)		-0.043 (0.024)	-0.003*** (0.001)		-0.002*** (0.001)		-0.004*** (0.001)
Lottery_t*income_{t-1}		-0.002 (0.002)			-0.002 (0.002)		0.000 (0.001)		-0.000 (0.001)
Lottery_t*net wealth_{t-1}		0.000 (0.000)		0.008 (0.005)	0.000 (0.000)		-0.000 (0.000)		0.000 (0.000)
Lottery_t*debt_{t-1}		0.000 (0.000)		-0.021 (0.011)	0.000 (0.000)		-0.001 (0.000)		0.000 (0.000)
Lottery_t*education_t		0.008 (0.006)		0.011 (0.007)	0.008 (0.006)		0.007 (0.005)		0.003 (0.004)
Lottery_t*risky share_{t-1}		-0.045 (0.062)		-0.034 (0.065)	-0.033 (0.061)		0.023 (0.063)		-0.065 (0.047)
Lottery_t*household size_t		0.026 (0.015)		0.024 (0.016)	0.018 (0.016)		0.030* (0.012)		0.008 (0.009)
Lottery_t*age_t		-0.005*** (0.001)		-0.003 (0.002)	-0.005*** (0.001)		-0.004** (0.001)		-0.003*** (0.001)
Lottery_t*MPC_i					0.016** (0.006)				
Observations	93,627	93,627	61,182	61,182	87,518	120,139	120,139	94,113	94,113

Notes: Each column represents a separate regression of equation (3) or (9). Controls include interacted variables, time-fixed effects, and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.

exclude the year in which a household wins and exclude households observed for less than five years; (ii) clean consumption and labor income of time-fixed effects; (iii) estimate an individual-specific MPC from labor income as $\text{MPC}_i = \frac{cov(\epsilon^c_{i,t},\epsilon^y_{i,t})}{var(\epsilon^y_{i,t})}$, where $\epsilon^c_{i,t}$ and $\epsilon^y_{i,t}$ are the consumption and income residuals from step (ii).[41]

Column (3) in Table 10 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, non-situational effects do seem relevant for understanding cross-sectional variation in MPCs. But more interestingly for our purposes, we see the following: 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 2.5, we omitted extreme consumption observations from our baseline sample so as to prevent outliers from driving our inference. Column four of Table 10 shows that if we instead do 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.

7 Which structural models are consistent with our results?

Our results imply that liquidity and age affect MPCs. In this section, we therefore explore the extent to which a two-asset life-cycle model can match our main empirical findings. We first describe the general framework (which also nests other simpler models). Then we compare our empirical results to model-generated counterparts, aiming to illuminate how the different economic mechanisms in this relatively standard framework contribute to explaining the patterns of our econometric estimates. As our exercises will show, our model will largely account for the estimated MPC heterogeneity and time profile, but struggle to account for the MPC level. We therefore conclude with a discussion of how the model could be extended to match the estimated MPC levels, while still being consistent with our empirical results on heterogeneity and dynamics.

7.1 A general consumption-saving model

A continuum of ex ante identical households with a finite (deterministic) lifetime T maximize their discounted utility from consumption. They face idiosyncratic income risk, a life-cycle income

in equation (9), and individual-fixed effects in consumption *sensitivity*.

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

profile, a borrowing constraint, and they can save in two assets: a liquid asset with no interest and a high return illiquid asset with transaction costs. The household problem is as follows:

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

$$\text{subject to}$$

$$b_t = b_{t-1} + r_t a_{t-1} + w_t y_t - d_t - \kappa \mathbb{1}_{a_t \neq a_{t-1}} - c_t$$

$$a_t = a_{t-1} + d_t$$

$$y_t = p_t + y$$

$$p_t = \rho p_{t-1} + \epsilon_t$$

$$\epsilon_t \sim N(0, \sigma)$$

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

Here c is consumption, b is the liquid asset, a is the illiquid asset, β is the discount factor, r is the return on the illiquid asset, w is the wage, y is the idiosyncratic income component, d is deposits into the illiquid assets, κ is the adjustment cost for the illiqud asset, p is the persistent income component, ρ is the persistence of the persistent income component, and ϵ is the income shock (note that ϵ is *not* the shock we use to study model-implied responses to lottery prizes).[42] Note that many simpler models are nested in the framework. For example, $\kappa = \infty$ and $T = \infty$ is the one-asset model, and $\kappa = 0$ is the standard life-cycle model. $B(\cdot)$ is the utility function over end-of-life net worth $(a_t + b_T)$, given by

$$B(a+b) = \psi_0 \frac{(a+b)^{1-\psi_1}}{1-\psi_1}.$$

Note that we allow for different curvature of the utility function and the bequest function ($\psi_1 \neq \gamma$).

The model is annual and we assume that households start life at age 21 with a draw from the ergodic income distribution. They live 55 years, until age 75. The wage profile is calibrated to match the wage profile in Kaplan and Violante (2014), where households start at age 21 with 0.12 in income that grows deterministically to 0.3 at age 60. After retirement (age 60), income is 0.12 again. The persistence of idiosyncratic income (ρ) is set to 0.91 and the standard deviation of income innovations (σ) is 0.12. The return on the illiquid asset (r) is 0.02. The preference parameters γ and ψ_1, are set to 2 and 1, respectively, implying that bequest is a luxury good. The remaining three parameters are set through a matching exercise explained below.

[42]The presence of ϵ means that households make decisions under income uncertainty. However, lottery prizes are not part of this income uncertainty. Lottery prizes are introduced as an unanticipated one-time shock, and households respond under the expectation that they will never win again.

7.2 Model-based versus empirical estimates

We select the values of β, ψ_0, and κ so as to minimize the difference between model-based and empirically estimated interaction effects for shock size, liquid asset holdings, income, net wealth, and age in Table 6, column 2. To facilitate comparability between model and empirical estimates, we use the standardized coefficients from Table 6. The model-based coefficients are obtained by simulating 100,000 households.[43] For each household, we randomly select an age at which they receive a lottery prize of a random size.[44] For each household, we compute its consumption after winning the lottery prize (treatment) and the counterfactual consumption that the household would have chosen had it not won the lottery prize (control). We then regress consumption on the lottery prize, with and without the interaction terms. Since our model is annual, our regression coefficient on lottery prize is an estimate of the annual marginal propensity to consume, analogously to our empirical estimate presented above. Our matching procedure yields the following values for the three remaining model parameters: $\beta = 0.997$, $\psi_0 = 10$, and $\kappa = 0.12$.

Table 11 presents the standardized regression coefficients from our baseline model (column 2) together with their empirical counterparts from Table 6 (column 1). Our main interest lies in the multivariate regression results, shown in the right-hand-side section of the table. Qualitatively, we see that the matched model generates the same correlations as in the data. The coefficients on size, liquid assets, income, and age are negative, while the coefficient on net wealth is positive, both in model and data. Quantitatively, the baseline two-asset model overstates the effects of size, income, net wealth, and age, while it understates the importance of liquid assets for consumption responses. In the left section of Table 11, we also present results from univariate regressions. Here, each coefficient estimate stems from a separate regression with only one interaction term, as in the empirical estimates of Table A.4. Here, the model generates qualitatively similar coefficients as in the data for all variables except income.

Taking a step back, there are two correlations that are challenging to generate from models. First, the correlation between age and MPC is negative. In contrast, a standard life-cycle model will typically imply a positive correlation because older households have a shorter horizon. The key reason why our two-asset model aligns with the data here, is the combination of a plausible life-cycle earnings profile and a luxury bequest motive. The earnings profile implies that young households typically choose to be constrained and have high MPCs since they expect higher income in the future. A luxury bequest motive induces old households to save for their offspring. They therefore consume less out of positive income and wealth innovations .

The second challenge for models is the non-negative correlation between net wealth and MPCs. The key component that allows our model to generate this correlation is the two-asset structure. In conventional one-asset buffer stock models, low net wealth is a predictor of high MPCs, unlike what we find in the data. This is also true in our model when we estimate the correlation between

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

[44] For simplicity, we use a uniform distribution of four different lottery sizes.

Table 11: Correlations between MPCs and other covariates. Model vs. data. Standardized coefficients.

	(1) Univariate regression			(2) Multivariate regression		
	Data	Two-asset Model baseline	Two-asset Model w/randomness	Data	Two-asset Model baseline	Two-asset Model w/randomness
Shock size	-0.008	-0.041	-0.035	-0.007	-0.027	-0.021
Liquid assets$_{t-1}$	-0.093	-0.074	-0.085	-0.081	-0.018	-0.040
Income$_{t-1}$	0.015	-0.135	-0.068	-0.025	-0.144	-0.077
Net wealth$_{t-1}$	-0.016	-0.049	-0.050	0.031	0.101	0.102
Age$_t$	-0.081	-0.073	-0.079	-0.076	-0.134	-0.137

Notes: This table displays the estimated standardized interaction terms from Table 6 together with the coefficients of a regression of the MPC in the model on model observables. The univariate column presents the interaction coefficients when we estimate each coefficient separately, while the multivariate column presents results when we estimate all coefficients jointly. All coefficients estimated from the model simulation are significant at the 0.1 % level.

net wealth and MPCs in a univariate regression. However, the correlation between MPCs and net wealth in the univariate model is determined by other observables that are correlated with wealth, such as liquid assets, income, and age. Hence, the correlation switches sign once the other variables are controlled for.

Figure 6 presents our second experiment where we compare the dynamic consumption responses we estimated empirically (as in Figure 2) to counterparts from our model. As one might expect, our two-asset model is unable to match the *level* of our estimated within-year MPC. The baseline two-asset model produces a within-year MPC slightly above 0.15, substantially lower than our empirical estimate of 0.5. Notably though, in terms of the qualitative time profile, the model captures the persistent nature of the consumption response.

7.3 Discussion

An upshot from our baseline model is that it can capture our estimated MPC patterns on dynamics and heterogeneity, but it misses our estimated MPC level. Informally stated, the model gets the systematic patterns right, but misses a ”base” response of around 0.35 in the year of winning. To substantiate this point further, we extend our model by adding a random component to households’

Figure 6: Dynamic consumption response. Model vs. data.

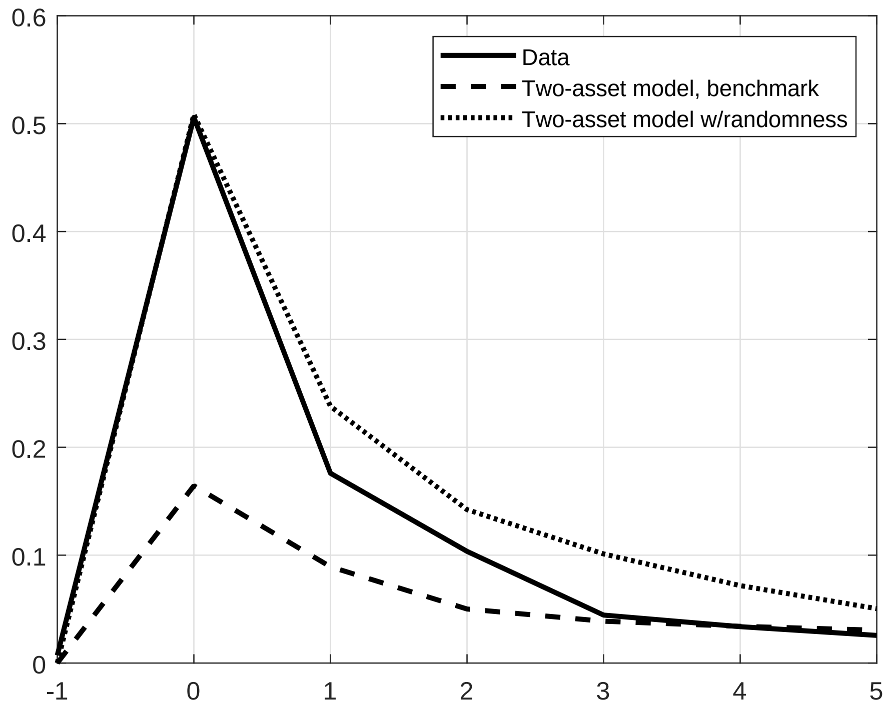

Notes: The figure shows the estimated dynamic consumption responses from the data (solid line), two-asset baseline model (dashed line), and the two-asset model with randomness (dotted line).

consumption response to a lottery prize.

We introduce the random component as follows: after winning, households first compute their optimal consumption plan as in the two-asset benchmark model. Thereafter they draw a random number determining how much extra to spend out of their prize. This random share comes on top of their model-implied consumption response. It is drawn from a uniform distribution over 0 to 0.7, where the support is chosen to add an average of 0.35 to households' response. Each period after winning, households are subject to the same procedure. Note that the random expenditure share always applies to the remaining lottery prize that has not yet been spent.

We choose this randomization approach so as to lift the MPC level without distorting the model-implied systematic patterns that result from households' optimal choice. Hence, in Table 11 we see that the systematic relationships between observable characteristics and MPCs are very similar to what they were in the baseline model without randomization. Moving to Figure 6, we see how the model by construction now matches the within-year response level. In addition it matches the time-profile fairly well, but exaggerates the subsequent responses somewhat.

Our randomization procedure is a pure engineering device. One possible interpretation is that after winning and choosing an optimal consumption plan, households receive additional 'temptation shocks' of arbitrary magnitudes. Mechanically, this randomizaion procedure has an important, no-table feature that allows our extended model to match the estimates: it combines rational behavior with an element that is orthogonal.

An alternative, seemingly similar, modeling strategy to obtain high MPCs, could to introduce "near-rational" behavior as emphasized by Cochrane (1989) and Kueng (2018). The idea there is that for certain households, the utility loss from consuming an arbitrary share of an income shock, instead of optimizing, is negligible. Hence, these households could have arbitrarily high MPCs. While this mechanism is very similar to our randomization procedure, there is an important difference. The utility loss incurred by deviating from optimal consumption decreases with household liquidity. Near-rationality therefore introduces a positive correlation between MPCs and liquid asset holdings. Hence, if we extended our baseline model with this mechanism instead of our chosen randomization approach, we would not be able account for the estimated systematic relationships between MPCs and household balance sheets.

Because our expenditure measure includes durable as well as non-durable consumption, another natural model extension to generate a high expenditure response might be to include purchases of durable consumption goods. However, there are two problems that such a model will face in matching our data.[45] First, the typical model-implied probability of making a durable consumption purchase is increasing in the holdings of liquid assets. Hence, discrete purchases of durable goods would tend to cause a positive correlation between liquid assets and MPCs, just like near-rationality does. Second, a model with durable consumption goods will typically generate expenditure shifting across time, without necessarily increasing the sum of total spending over a five-year period like in Figure 2. Hence, although the addition of durable goods would raise the initial consumption response, it would typically reduce expenditure thereafter as part of the first-period increase follows from planned durable purchases which are shifted forward in time.

We believe that a reasonable summary of our model results is as follows. The main patterns we find regarding MPC *heterogeneity* and response *dynamics* are consistent with models of precautionary savings 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 withhold the systematic behavior implied by the two-asset buffer-stock framework, and in addition contain a mechanism for high consumption responses that does not distort the patterns following from conventional models.

8 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 response to income

[45]To conserve on space, we do not show results for such a model. These results are available upon request.

shocks, and in particular which household characteristics that 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 correlate 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 support 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 non-situational factors are irrelevant. To the contrary, our evidence shows that household-specific MPCs out of regular income do indeed matter. But the fact 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 precautionary savings model.

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 non-homothetic 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 wide-spread observation of remarkably high saving rates among the elderly.

The base level of MPCs that we estimate, in particular the level among highly 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. We illustrate this by adding a random component to household consumption responses in a two-asset life-cycle model. What constitutes the micro foundations behind such a level shifter, remains an open question for future research.

References

Aaronson, D., S. Agarwal, and E. French (2012): "The Spending and Debt Response to Minimum Wage Hikes," *American Economic Review*, 102, 3111–39.

Agarwal, S., S. Chomsisengphet, N. Mahoney, and J. Stroebel (2015): "Regulating consumer financial products: Evidence from credit cards," *Quarterly Journal of Economics*, 130, 111–164.

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

Agarwal, S. and W. Qian (2014): "Consumption and Debt Response to Unanticipated Income Shocks: Evidence from a Natural Experiment in Singapore," *American Economic Review*, 104, 4205–4230.

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

Arellano, M. and O. Bover (1995): "Another look at the instrumental variable estimation of error-components models," *Journal of Econometrics*, 68, 29–51.

Auclert, A. (2016): "Monetary policy and the redistribution channel," MIT Working paper.

Auclert, A. and M. Rognlie (2016): "Inequality and Aggregate Demand," Working paper, Stanford University.

Auclert, A., M. Rognlie, and L. Straub (2018): "The Intertemporal Keynesian Cross," Working paper, Stanford University.

Autor, D., A. R. Kostøl, and M. Mogstad (2015): "Disability Benefits, Consumption Insurance, and Household Labor Supply," MIT Working paper, July.

Aydin, D. (2015): "The Marginal Propensity to Consume Out of Liquidity: Evidence from Random Assignment of 54,522 Credit Lines," Stanford University Working paper.

Baker, S., L. Kueng, S. Meyer, and M. Pagel (2018): "Measurement Error in Imputed Consumption," Mimeo.

Baker, S. R. (2018): "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *Journal of Political Economy*, 126, 1504–1557.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2018): "House Prices and Consumer Spending," *Review of Economic Studies*, 85, 1502–1542.

Blundell, R. and S. Bond (1998): "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics*, 87, 115–143.

Bø, E. E., J. Slemrod, and T. O. Thoresen (2015): "Taxes on the internet: Deterrence effects of public disclosure," *American Economic Journal: Economic Policy*, 7, 36–62.

Briggs, J. S., D. Cesarini, E. Lindqvist, and R. Östling (2015): "Windfall Gains and Stock Market Participation," NBER Working paper No. 21673.

Browning, M. and M. D. Collado (2001): "The response of expenditures to anticipated income changes: panel data estimates," *American Economic Review*, 91, 681–692.

Browning, M., T. F. Crossley, and J. Winter (2014): "The Measurement of Household Consumption Expenditures," *Annual Review of Economics*, 6, 475–501.

Browning, M., M. Gørtz, and S. Leth-Petersen (2013): "Housing Wealth and Consumption: A Micro Panel Study," *Economic Journal*, 123, 401–428.

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

Browning, M. and S. Leth−Petersen (2003): "Imputing consumption from income and wealth information," *Economic Journal*, 113, F282–F301.

Campbell, J. Y. and N. G. Mankiw (1989): "Consumption, income, and interest rates: Reinterpreting the time series evidence," *NBER macroeconomics annual*, 4, 185–216.

Carroll, C. (2004): "Theoretical foundations of buffer stock saving," NBER Working paper No. 10867.

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

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

——— (2001): "A theory of the consumption function, with and without liquidity constraints," *Journal of Economic Perspectives*, 15, 23–45.

Carroll, C. D. and M. S. Kimball (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64, 981–992.

——— (2001): "Liquidity constraints and precautionary saving," NBER Working paper No. 8496.

Cesarini, D., E. Lindqvist, M. J. Notowidigdo, and R. Östling (2017): "The effect of wealth on individual and household labor supply: evidence from Swedish lotteries," *American Economic Review*, 107, 3917–46.

Cesarini, D., E. Lindqvist, R. Östling, and B. Wallace (2016): "Wealth, health, and child development: Evidence from administrative data on Swedish lottery players," *Quarterly Journal of Economics*, 131, 687–738.

Christelis, D., D. Georgarakos, T. Jappelli, L. Pistaferri, and M. van Rooij (2018): "Asymmetric Consumption Effects of Transitory Income Shocks," *forthcoming in Economic Journal.*

Cloyne, J., C. Ferreira, and P. Surico (2018): "Monetary policy when households have debt: new evidence on the transmission mechanism," *forthcoming in Review of Economic Studies.*

Cochrane, J. (1989): "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, 79, 319–37.

Crossley, T. F., H. Low, and S. Smith (2016): "Do consumers gamble to convexify?" *Journal of Economic Behavior & Organization*, 131, 276–291.

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

De Nardi, M. and G. Fella (2017): "Saving and Wealth Inequality," *Review of Economic Dynamics*, 26, 280–300.

Di Maggio, M., A. Kermani, and K. Majlesi (2018): "Stock Market Returns and Consumption," Working paper.

Druedahl, J. and A. Martinello (2017): "Long-Run Saving Dynamics: Evidence from Unexpected Inheritances," Working paper.

Eika, L., M. Mogstad, and O. Vestad (2017): "What can we learn about household consumption from information on income and wealth," Working paper.

Fagereng, A., L. Guiso, D. Malacrino, and L. Pistaferri (2018a): "Heterogeneity and Persistence in Returns to Wealth," CESifo Working Paper No. 7107.

Fagereng, A. and E. Halvorsen (2017): "Imputing consumption from Norwegian income and wealth registry data," *Journal of Economic and Social Measurement*, 42, 67–100.

Fagereng, A., M. B. Holm, and K. N. Torstensen (2018b): "Housing Wealth in Norway, 1993 - 2015," Statistics Norway Reports.

Fernández-Villaverde, J. and D. Z. Valencia (2018): "A Practical Guide to Parallelization in Economics," Tech. rep., National Bureau of Economic Research.

Friedman, M. (1957): *A Theory of the Consumption Function*, Princeton University Press.

Fuchs-Schündeln, N. (2008): "The Response of Household Saving to the Large Shock of German Reunification," *American Economic Review*, 98, 1798–1828.

Fuchs-Schündeln, N. and T. A. Hassan (2016): "Natural experiments in macroeconomics," in *Handbook of Macroeconomics*, Elsevier, vol. 2, 923–1012.

Fuster, A., G. Kaplan, and B. Zafar (2018): "What Would You Do With $500? Spending Responses to Gains, Losses, News and Loans," NBER Working Papers 24386.

Gross, T., M. J. Notowidigdo, and J. Wang (2016): "The Marginal Propensity to Consume Over the Business Cycle," NBER Working paper No. 22518.

Holm, M. B. (2018): "Consumption with liquidity constraints: An analytical characterization," *Economics Letters*, 167, 40 – 42.

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

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

Imbens, G. W., D. B. Rubin, and B. 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, 778–794.

Jappelli, T. and L. Pistaferri (2010): "The consumption response to income changes," *Annual Review of Economics*, 2, 479–506.

——— (2014): "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 6, 107–136.

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

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

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

Kaplan, G., G. L. Violante, and J. Weidner (2014): "The wealthy hand-to-mouth," *Brookings papers on economic activity*, 1, 77–153.

Koijen, R., S. Van Nieuwerburgh, and R. Vestman (2015): "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*, ed. by C. Carroll, T. Crossley, and J. Sabelhaus, University of Chicago Press, chap. 11, 308–346.

Kolsrud, J., C. Landais, and J. Spinnewijn (2017): "Studying Consumption Patterns using Registry Data: Lessons From Swedish Administrative Data," CEPR Discussion Papers 12402, C.E.P.R. Discussion Papers.

Kreiner, C. T., D. D. Lassen, and S. Leth-Petersen (2015): "Measuring the Accuracy of Survey Responses using Administrative Register Data: Evidence from Denmark," in *Improving the Measurement of Consumer Expenditures*, ed. by C. Carroll, T. Crossley, and J. Sabelhaus, University of Chicago Press, chap. 10, 289–307.

Krueger, D., K. Mitman, and F. Perri (2016): "Macroeconomics and household heterogeneity," in *Handbook of Macroeconomics*, Elsevier, vol. 2, 843–921.

Kueng, L. (2018): "Excess Sensitivity of High-Income Consumers*," *The Quarterly Journal of Economics*, 133, 1693–1751.

Kuhn, P., P. Kooreman, A. Soetevent, and A. Kapteyn (2011): "The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery," *American Economic Review*, 101, 2226–2247.

Leth-Petersen, S. (2010): "Intertemporal consumption and credit constraints: Does total expenditure respond to an exogenous shock to credit?" *American Economic Review*, 100, 1080–1103.

Meyer, B. D., W. K. Mok, and J. X. Sullivan (2015): "Household Surveys in Crisis," *Journal of Economic Perspectives*, 29, 199–226.

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

Modigliani, F. and R. Brumberg (1954): "Utility analysis and the consumption function: An interpretation of cross-section data," in *Post-Keynesians Economics*, ed. by K. Kurihara, New Brunswick, NJ: Rutgers University Press, 388–436.

Ng, Y.-K. (1965): "Why do People Buy Lottery Tickets? Choices Involving Risk and the Indivisibility of Expenditure," *Journal of Political Economy*, 73, 530–535.

Parker, J. (2017): "Why Don't Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment," *American Economic Journal: Macroeconomics*, 9, 153–183.

Parker, J. A. and N. S. Souleles (2017): "Reported Preference vs. Revealed Preference: Evidence from the propensity to spend tax rebates," NBER Working paper No. 23920.

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

Pistaferri, L. (2015): "Household consumption: Research questions, measurement issues, and data collection strategies," *Journal of Economic and Social Measurement*, 40, 123–149.

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

——— (2009): "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review*, 99, 374–379.

Ziliak, J. P. (1998): "Does the choice of consumption measure matter? An application to the permanent-income hypothesis," *Journal of Monetary Economics*, 41, 201–216.

A Appendix

A.1 Imputed consumption versus 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 A.1 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. Overall, imputed consumption follows the same trajectory as consumption in the national accounts.

Figure A.1: Consumption per person. Imputed and national accounts. 1994 - 2006.

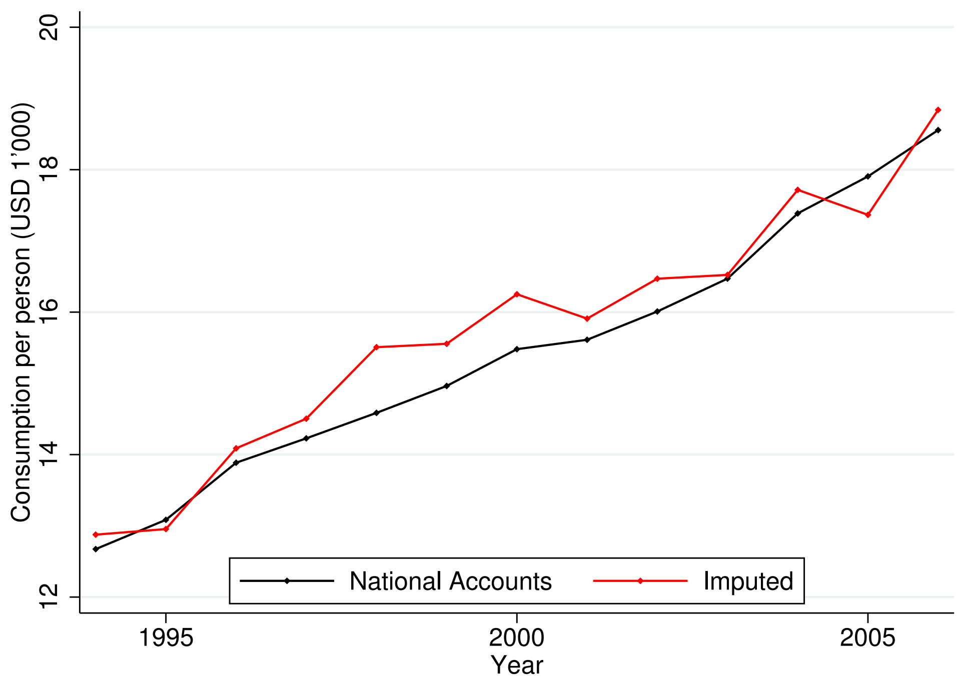

Notes: The graphs show consumption per person in USD 1,000 between 1994 and 2006. For the national accounts, we take household consumption at current prices, divide it by the population of Norway, CPI adjust it to the year 2000, and convert to USD 1,000 using the year-2000 exchange rate. For imputed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

A.2 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, $l_1, ..., l_N$, and consumption, $c_1, ..., c_N$.

The estimated β on the full sample is defined as

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

We can also define individual specific $\beta_{1,i}$'s 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 β's in the following way

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

where the weights are defined as $\omega_i = \frac{l_i^2}{\sum_{j=1}^{N} l_j^2}$. 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.3 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

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

where ϵ_{it} is measurement error on consumption. We discard lottery draws that are lower than USD 1,100 or greater than USD 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 A.2: Our identification strategy on simulated data

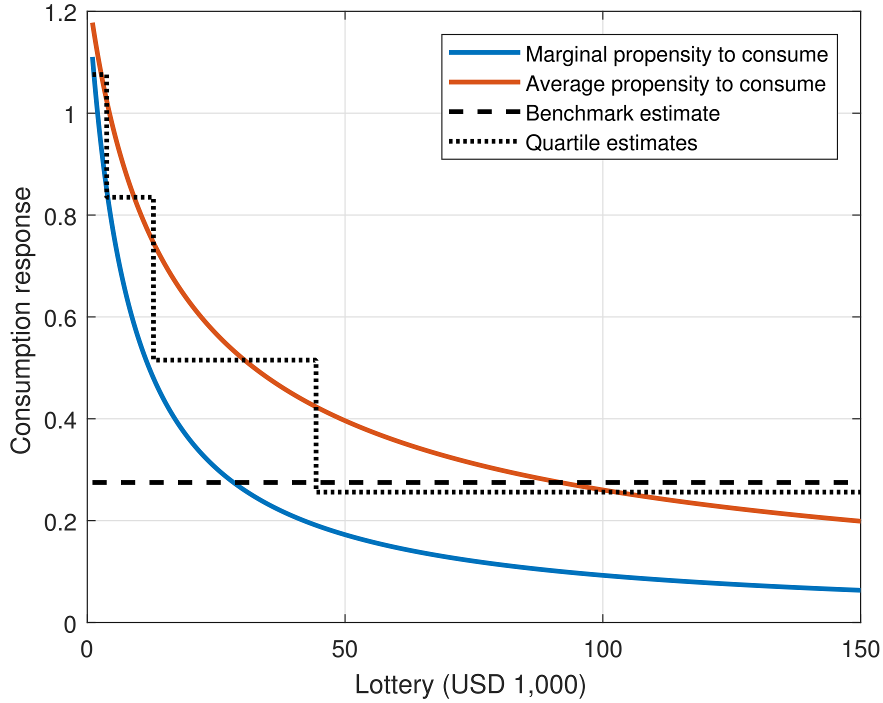

Notes: The figure shows how our identification strategy works on simulated data with the data generating process $\text{consumption}_{it} = 10\log(\text{lottery}_{it}/8+1)$. The blue and red lines depict analytical solutions. We draw lottery prizes from a truncated lognormal distribution with mean 4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

Figure A.2 presents the results with low σ_ϵ. The blue line and red line represents 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 A.2. 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 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.2 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 5.1.

A.4 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 play for, 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 all players play for more or less the same amount 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 A.3: Effects of lottery prize on households' earnings

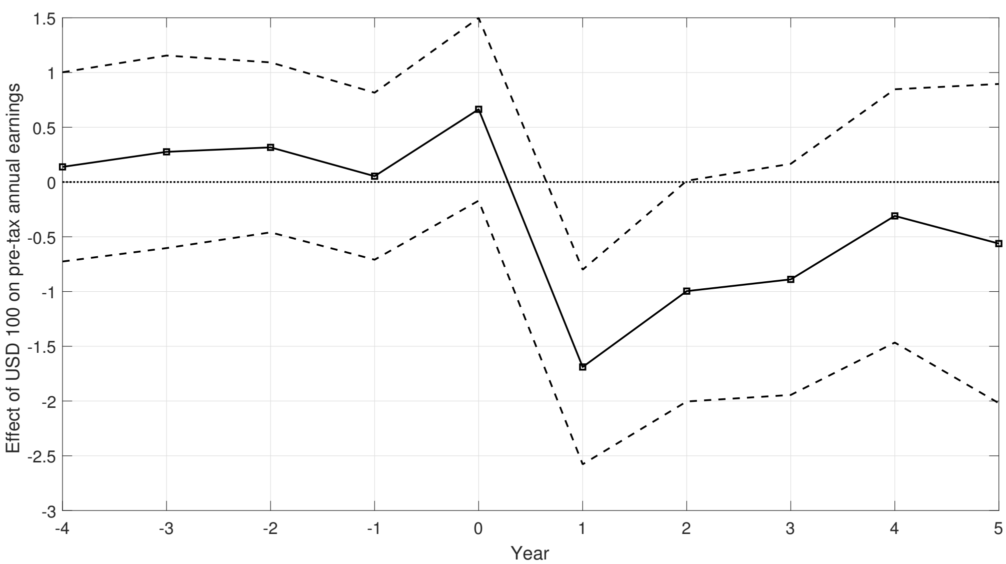

Notes: Each dot represents a separate regression of gross household earnings on lottery prize. Controls include household-fixed effects and time-fixed effects. Dashed lines show 95 percent confidence intervals. The standard errors in parentheses are robust and clustered on the household level. Estimation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

Figure A.3 presents the results from estimating the response of pre-tax 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.5 Pre-trends

Figure A.4 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.

Figure A.4: Pre-trends

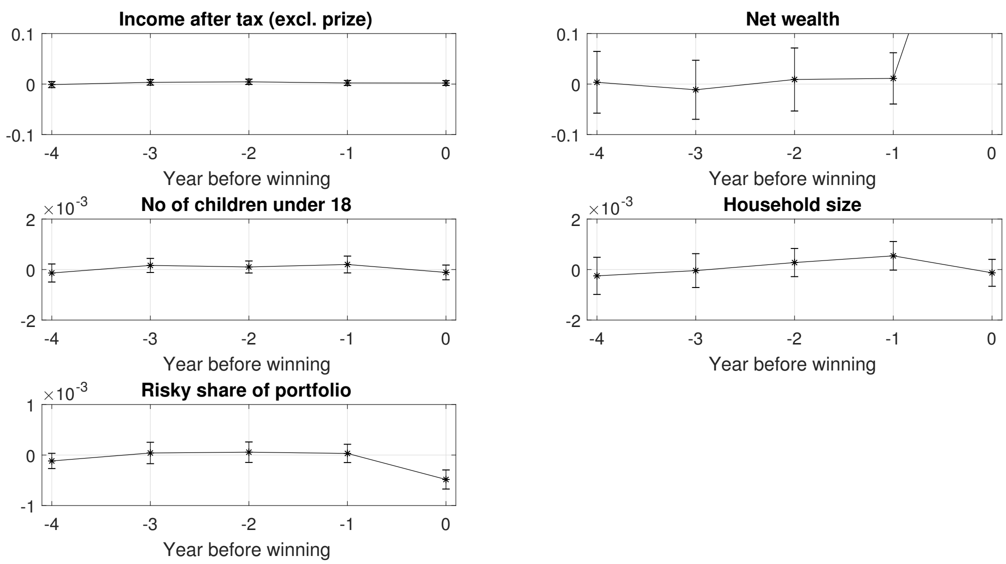

Notes: The figure shows responses of income, net wealth, risky share of portfolio, household size, and number of children to a lottery prize in the years before winning. Year 0 is the year of winning. The scale on the y-axis is in per unit of lottery prize.

A.6 Time aggregation

Since our data is 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.

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 correspond 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 A.5: Time aggregation - data and fitted model

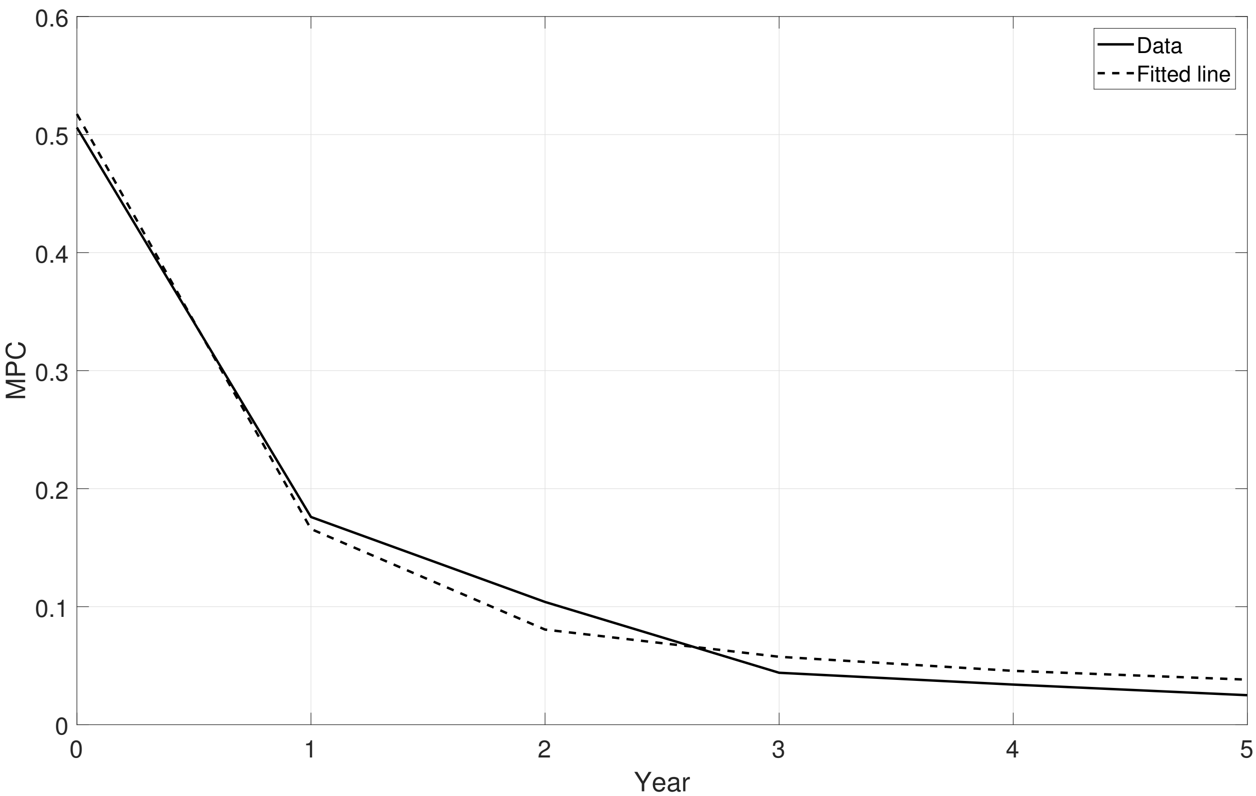

Notes: The solid line is the estimated dynamic consumption responses from Figure 2. The dashed line is the fitted consumption response using the approach described above.

Figure A.5 shows the dynamic consumption response and the fitted function. The outcome is $\theta_1 = 0.6285$ and $\theta_2 = 0.2142$. We can now use the estimated function to calculate the 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 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.7 Heterogeneous dynamic responses to lottery prizes

In the body of the paper, we focus on heterogeneity of the within year responses to lottery prizes. It is also important to understand how the dynamics of household responses are heterogeneous. In this section, we therefore present the results for heterogeneous dynamic responses to lottery prizes. To construct these results, we first divide the population of lottery winners into quartiles of liquid assets or age in the year before they win the lottery prize. We then estimate the dynamic responses to lottery prizes within each quartile.

Figure A.6: Dynamic household responses to lottery prizes. Quartiles of liquid assets.

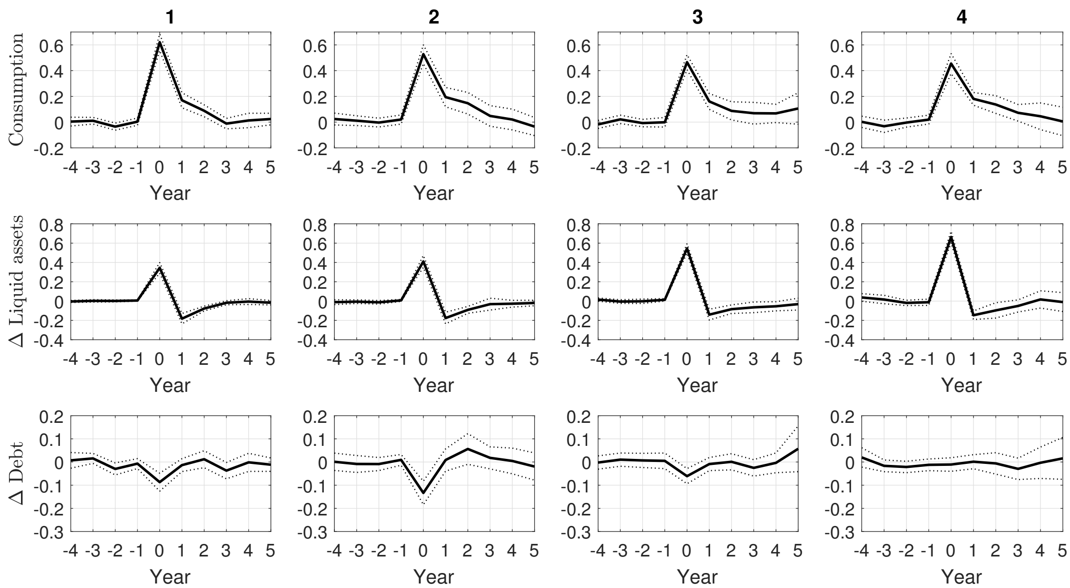

Notes: Each point is estimated as a separate regression of (7) or (8) within each liquid asset quartile. The cutoffs between the liquid asset quartiles vary with the year won. In 1994 (2006), the cutoffs are: USD 720 (1,720), USD 3,140 (7,160), and USD 10,710 (21,600). Controls include time-fixed and household-fixed effects. Dotted lines represent 95 percent confidence intervals using robust standard errors clustered at the household level.

Figure A.7: Dynamic household responses to lottery prizes. Quartiles of age.

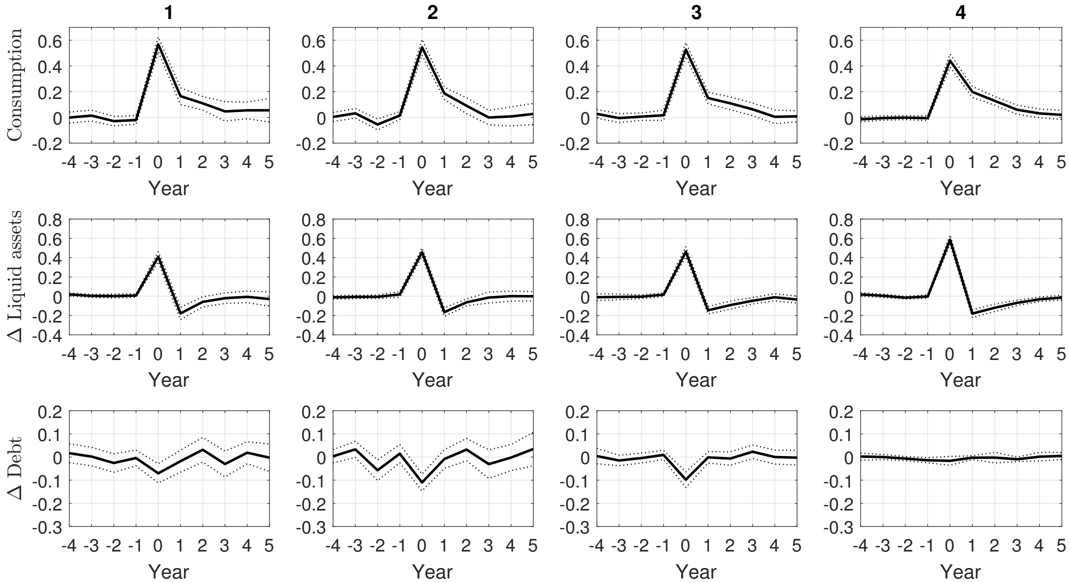

Notes: Each point is estimated as a separate regression of (7) or (8) within each liquid asset quartile. The cutoffs between the quartiles are: 39, 51, and 63. Controls include time-fixed and household-fixed effects. Dotted lines represent 95 percent confidence intervals using robust standard errors clustered at the household level.

Figure A.6 presents the result for quartiles of liquid assets. The year zero response is the same as that presented in Table 8: the consumption response declines with liquid assets, while the liquid assets response increases with the prior period level of liquid assets (liquid assets = deposits + SBM in Table 8). The interesting feature is that households with more liquid assets tend to smooth out the lottery prize over more years. Households in the bottom quartile have no consumption response from year 3 and forward, while those in the top quartile have a positive consumption response also in period 3 and 4. Furthermore, since the within year response in quartile 4 is lower than for the bottom quartile, the path of consumption is less front-loaded.

We see the same smoothing pattern also for the dynamic age effects in Figure A.7. Young households tend to consume more in the year they won, but then cut back on consumption relatively fast. Old household, on the other hand, consume less in the year they won, but smooth consumption over more years.

A.8 Sample selection

One concern with our results is that our measure of imputed consumption might contain 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.

Now, 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 Table 4 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 test whether risky asset purchases might bias our results, we re-estimate Table 6 with a sample of lottery winners that never owned any risky assets (stocks, bonds, or mutual funds) in the sample between 1993 and 2006. Table A.1 presents the results. The table is qualitatively and quantitatively almost identical to the results in Table 6 in the body of the paper with our benchmark sample. The results still point to lottery prize size, liquid assets, and age as the main culprits in explaining the size of MPCs. Moreover, our benchmark MPC is a little bit lower (0.500 vs. 0.521).

Further, we present two additional results from the paper using the no-risky-assets sample. Figure A.8 shows the dynamic responses of consumption and balance sheet components to lottery prizes, the same as Figure 2 in the body of the paper. Both the patterns and the quantitative sizes of the responses are similar to the results using the main sample. The only exception is that 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.

Figure A.9 shows the heterogeneity in consumption responses across the lottery prize size and liquid asset dimension using the no-risky-assets sample, similar to Figure 4 in the body of the paper. Again, the results in the no-risky-assets sample is consistent with the results in the whole sample. MPC responses decrease in both the level of liquid assets and the size of the lottery prize. Given the similarities in the results, we argue that measurement errors in imputed consumption are unlikely to be an important driver of our main results.

Table A.1: The MPC out of lottery prizes. Interaction effects. No-risky-assets sample.

	Dependent variable: Consumption_t									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Lottery_t	0.502***	0.613***	0.541***	0.423***	0.506***	0.448***	0.457***	0.417***	0.841***	0.862***
	(0.021)	(0.033)	(0.023)	(0.049)	(0.026)	(0.025)	(0.026)	(0.043)	(0.075)	(0.115)
Lottery_t^2		-0.001**								-0.001**
		(0.000)								(0.000)
Lottery_t*liquid assets_{t-1}			-0.004***							-0.003**
			(0.001)							(0.001)
Lottery_t*income_{t-1}				0.004						-0.005
				(0.002)						(0.003)
Lottery_t*net wealth_{t-1}					-0.000					0.000
					(0.000)					(0.000)
Lottery_t*debt_{t-1}						0.002***				0.002*
						(0.001)				(0.001)
Lottery_t*education_t							0.025**			0.006
							(0.009)			(0.010)
Lottery_t*household size_t								0.046*		0.038
								(0.021)		(0.023)
Lottery_t*age_t									-0.006***	-0.005**
									(0.001)	(0.002)
Observations	40,852	40,852	40,852	40,852	40,852	40,852	40,852	40,852	40,852	40,852

Notes: Each column represents a separate regression of (9). Controls include the variables that are in interacted with lottery prize, time-fixed effects, and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.

Figure A.8: Dynamic household responses to lottery prizes. No-risky-assets sample.

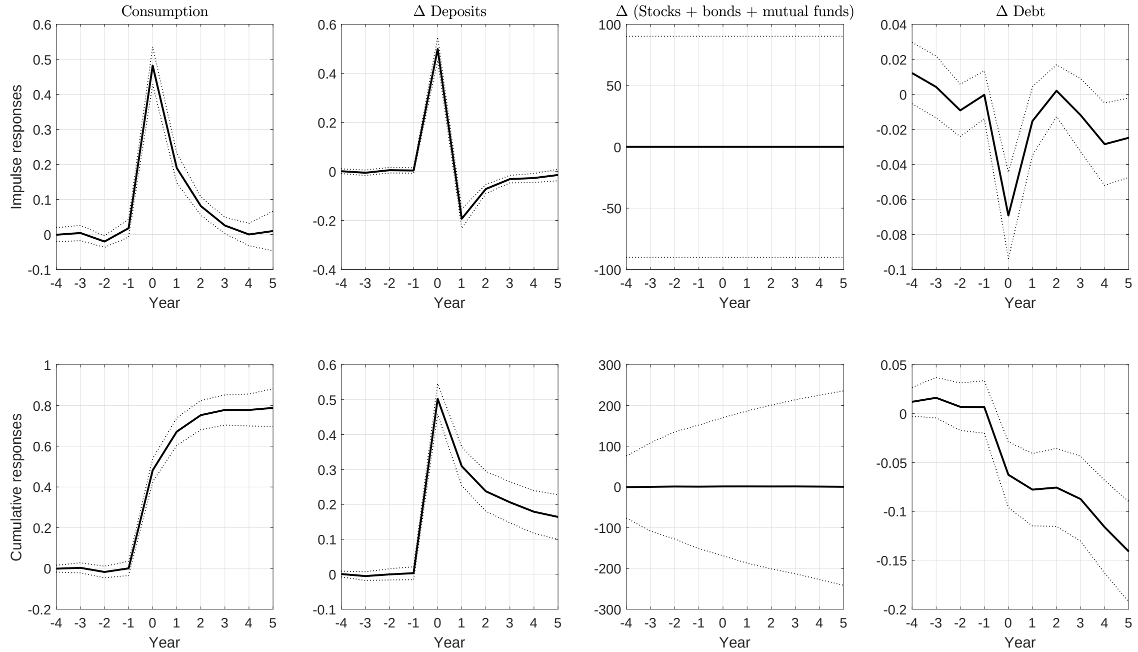

Notes: Each point is estimated as a separate regression of (7) or (8). Controls include time-fixed and household-fixed effects. Dotted lines represent 95 percent confidence intervals using robust standard errors clustered at the household level. The cumulative responses in the bottom panel are the sum of year-specific responses in the top panel. Standard errors of estimated cumulative effects are obtained through Monte Carlo simulations.

Figure A.9: Heterogeneous consumption responses. Lottery prize size and liquid assets. No-risky-assets sample.

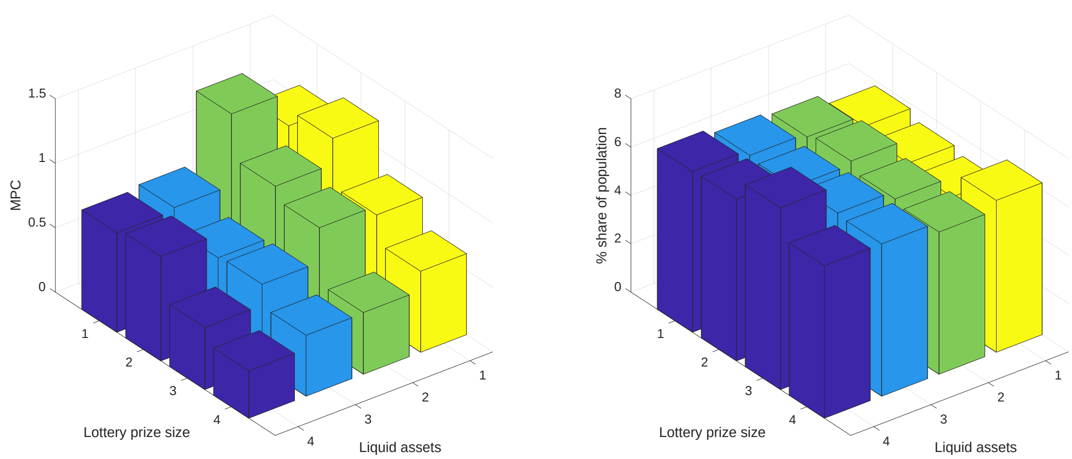

Notes: Each bar/point is estimated as a separate regression of equation (3). Controls include time-fixed and household-fixed effects. Total N: 40,852.

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 income that is cleaned of transitory shocks, mimicking the permanent income variable in standard macro models.

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

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

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 to, 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. This is to ensure that our households are most likely to be working full-time. We construct permanent income by three steps:

1. 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.

2. 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 $N = \sum_{1993}^{2006} 1_{i \text{ exists}}$. $\bar{\epsilon_i}$ thus describes the type of the household.

3. Permanent income is then

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

4. 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 re-estimate Table 6 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,

Table A.2: The MPC out of lottery prizes. Interaction effects. Normalizing by permanent income.

	Dependent variable: $Consumption_t$									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$Lottery_t$	0.563*** (0.019)	0.536*** (0.027)	0.578*** (0.020)	0.547*** (0.022)	0.578*** (0.027)	0.542*** (0.027)	0.563*** (0.021)	0.524*** (0.037)	0.675*** (0.092)	0.648*** (0.115)
$Lottery_t^2$		0.005 (0.006)								0.000 (0.005)
$Lottery_t$*liquid $assets_{t-1}$			-0.029 (0.023)							-0.043 (0.024)
$Lottery_t$*net $wealth_{t-1}$				0.005 (0.005)						0.008 (0.005)
$Lottery_t$*$debt_{t-1}$					-0.015 (0.011)					-0.021 (0.011)
$Lottery_t$*$education_t$						0.008 (0.007)				0.011 (0.007)
$Lottery_t$*risky $share_{t-1}$							0.004 (0.072)			-0.034 (0.065)
$Lottery_t$*household $size_t$								0.021 (0.018)		0.024 (0.016)
$Lottery_t$*age_t									-0.003 (0.002)	-0.003 (0.002)
Observations	61,182	61,182	61,182	61,182	61,182	61,182	61,182	61,182	61,182	61,182

Notes: Each column represents a separate regression of (9). All variables that are denoted in USD are divided by permanent income as described in Appendix A.9. Controls include the variables that are in interacted with lottery prize, time-fixed effects, and household-fixed effects. The standard errors in parentheses are robust and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.

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 A.2 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 6.[46]

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 households will only consume a small fraction. The goal of this section is to include such 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:

1. Exclude the year in which they win the lottery and keep only households with at least 4 observations.

2. Clean consumption and income of time-fixed effects by controlling for time dummies:

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

3. Estimate an MPC from labor income as $\text{MPC}_i = \frac{cov(\epsilon_{i,t}^c, \epsilon_{i,t}^y)}{var(\epsilon_{i,t}^y)}$.

MPC_i is then 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.

[46] 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 A.3: The MPC out of lottery prizes. Interaction effects. Controlling for MPC from income.

	Dependent variable: Consumption_t		
	(1)	(2)	(3)
Lottery_t	0.519***	0.514***	0.865***
	(0.015)	(0.015)	(0.077)
Lottery_t*MPC_i		0.009	0.016**
		(0.006)	(0.006)
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
			(0.000)
Lottery_t*education_t			0.008
			(0.006)
Lottery_t*risky share_{t-1}			-0.033
			(0.061)
Lottery_t*household size_t			0.018
			(0.016)
Lottery_t*age_t			-0.005***
			(0.001)
Observations	87,969	87,969	87,969

Notes: Each column represents a separate regression of (9). 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. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.

Table A.3 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

Figure A.10: Heterogeneous consumption responses. Liquid assets and age.

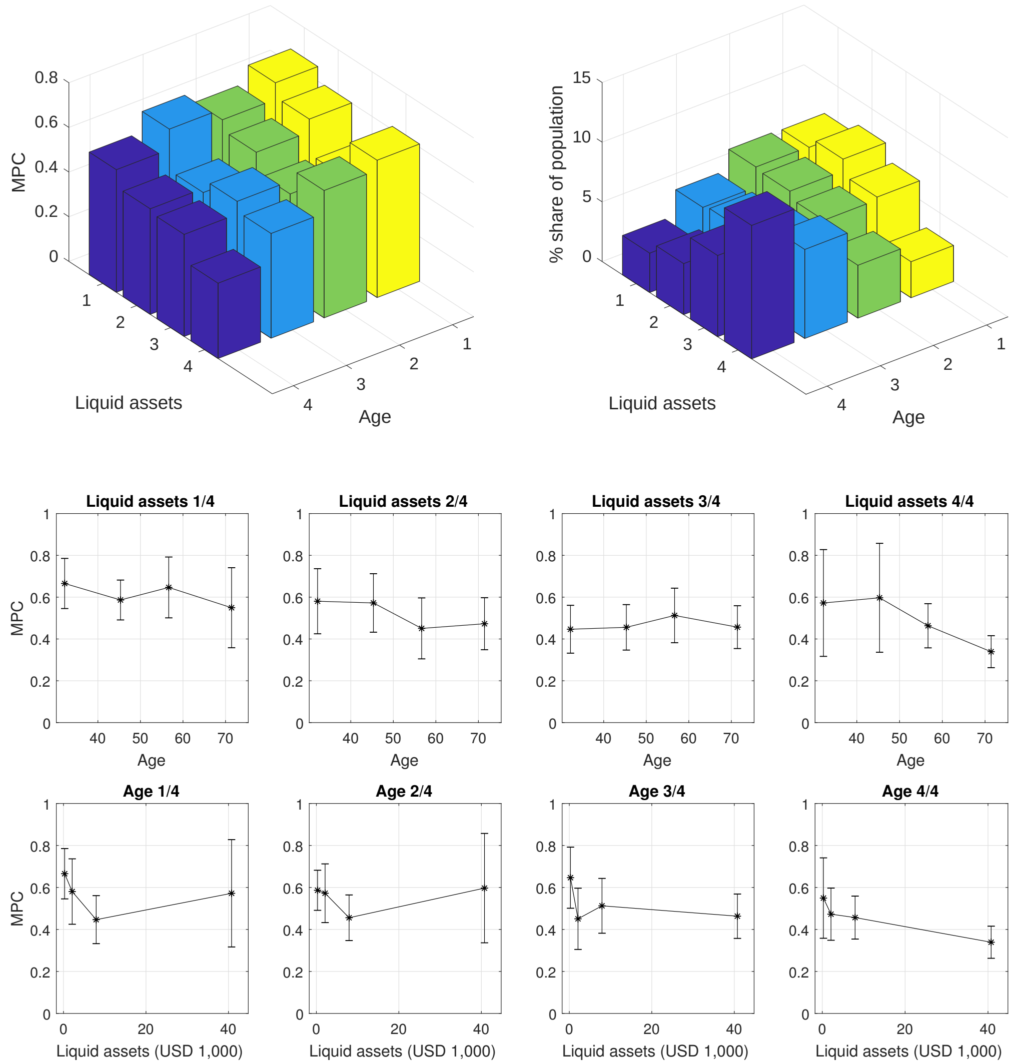

Notes: Each bar/point is estimated as a separate regression of equation (3). Controls include time-fixed and household-fixed effects. The standard errors are robust and clustered at the household-level. Total N: 93,627.

Table A.4: The MPC out of lottery prizes. Interaction effects.

	Dependent variable: Consumption_t					
	(1)	(2)	(3)	(4)	(5)	(6)
Lottery_t	0.523***	0.558***	0.529***	0.525***	0.529***	0.874***
	(0.014)	(0.015)	(0.014)	(0.014)	(0.015)	(0.074)
Lottery_t*deposits_{t-1}		-0.003***				-0.003**
		(0.001)				(0.001)
Lottery_t*stocks_{t-1}			-0.012***			-0.011***
			(0.002)			(0.002)
Lottery_t*bonds_{t-1}				-0.002		-0.001
				(0.003)		(0.003)
Lottery_t*mutual funds_{t-1}					-0.005	-0.005
					(0.003)	(0.003)
Lottery_t^2						-0.001***
						(0.000)
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}						0.023
						(0.072)
Lottery_t*household size_t						0.025
						(0.015)
Lottery_t*age_t						-0.005***
						(0.001)
Observations	93,627	93,627	93,627	93,627	93,627	93,627

Notes: Each column represents a separate regression of (9) 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. *, **, *** denote significance at the 5, 1, and 0.1 percent level, respectively.