

MPC heterogeneity and household balance sheets*

Andreas Fagereng[†], Martin B. Holm[‡] and Gisle J. Natvik[§]

First version: December 2015, this version: November 25, 2016

Abstract

Using Norwegian administrative data, we study how sizable lottery prizes affect household expenditure and savings. Expenditure responses (MPCs) spike in the year of winning, with a mean estimate of 0.35, and thereafter fall markedly. Controlling for all items on the household balance sheet and characteristics such as education and age, MPCs vary with the amount won and liquid assets only. Shock size matters: The MPC among the 25 percent winning least is twice as high as among the 25 percent winning most. Many households are wealthy, illiquid and have high MPCs, consistent with 2-asset models of consumer choice.

JEL Classification: D12, D14, D91, E21

Keywords: marginal propensity to consume, household expenditure response, household heterogeneity, income shocks

*We thank Mario Alloza, Adrien Auclert, Paco Buera, Thomas Crossley, Nicola Fuchs-Schündeln, Ragnar Juelsrud, Davide Malacrino, Ben Moll, Luigi Pistaferri, Terje Skjerpen, Kjetil Storesletten, Nate Vellekoop and Gianluca Violante, as well as participants at seminars, workshops and conferences, for helpful discussions and suggestions.

[†]Research Department, Statistics Norway, afagereng@gmail.com

[‡]Department of Economics, BI Norwegian Business School, martin.b.holm@bi.no

[§]Department of Economics, BI Norwegian Business School, gisle.j.natvik@bi.no

1 Introduction

How do households adjust their consumption in response to unanticipated and transitory income shocks? And what are the key determinants behind the magnitude of these responses? These are fundamental questions in macroeconomics. Their answers are key to address issues such as how distributional dynamics and inequality affect the economy, how fiscal and monetary policy can stimulate aggregate demand, and what are the welfare consequences of incomplete insurance markets. Consequently, these questions have received widespread attention both in the academic literature and among policy makers. In this article, we contribute by studying (i) how income surprises feed into consumption and savings over time, (ii) which household characteristics systematically relate to the magnitude of these responses, and (iii) how the size of income shocks matters for consumption responsiveness.

Despite the longstanding interest in the topic, evidence on the marginal propensity to consume (MPC) out of transitory income shocks and its determinants, is limited.¹ There are good reasons why. First, credible estimation of this specific MPC requires the researcher to observe an exogenous income shock. Moreover, it is not sufficient for the innovation to income to be exogenous, it must also be clear whether the innovation is anticipated or not, as theory gives very different predictions for the two (Modigliani and Brumberg, 1954; Friedman, 1957). For the same reason, it must be clear that the shock is transitory, and not persistent. Such exogenous shocks with a clear information structure are hard to come by in the data, with the notable exceptions of transfer schemes and lottery prizes. Second, the income shocks must also be observed together with reliable data on household consumption, which is a rare combination. Third, while mean short-run MPCs certainly are interesting in themselves, in order to inform theory one really needs a better understanding of the determinants behind MPC heterogeneity, and ideally how income innovations are spent over some time. This requires panel data with rich information on household characteristics, in particular data on wealth and balance sheets as these play a central role in contemporary structural models of consumption dynamics, see for instance Kaplan and Violante (2014), Carroll, Slacalek, Tokunaka, and White (2014) and Krueger, Mitman, and Perri (2016). These data requirements are rather formidable, leaving a detailed understanding of how households respond to unanticipated transitory income shocks still in the void.

With this study we aim to provide a thorough and transparent characterization of how transitory innovations to income affect household expenditure and saving, dealing with all of the aforementioned issues. To this end, we utilize detailed and precise third-party reported information on household balance sheets, as well as several other characteristics, covering the universe of Norwegian households for more than a decade. From these data we construct an imputed measure of consumption expenditure utilizing the household budget constraint. By construction, this measure does not distinguish durable from non-durable consumption expenditure, but we are able to isolate

¹See for instance surveys by Browning and Collado (2001), Jappelli and Pistaferri (2010) and Fuchs-Schündeln and Hassan (2016).

purchases of boats and cars and study them separately. Importantly, the data also include prizes from betting activities where a majority of the Norwegian population participate. As we will argue, these prizes constitute unexpected transitory income shocks.

Our data allow us to explore and document an array of features regarding households' MPCs. We start by establishing that winners spend a large fraction of the prize money within the first year of receiving it. The marginal propensity to consume out of a lottery prize lies around 0.35. Of the remainder, most is saved in liquid assets, primarily deposits. Five years after winning, households typically have spent about 60 percent of their windfall gain. However, these estimates mask considerable heterogeneity and our main contribution is to dissect this.

First, we provide the unconditional correlations between estimated MPCs and households' income and balance sheet components. These moments are essential in the increasingly utilized "sufficient statistics" approach to analyzing macroeconomic shock propagation, see for instance Berger, Guerrieri, Lorenzoni, and Vavra (2015), Auclert and Rognlie (2016) and Auclert (2016), but empirical estimates have up until now been close to non-existent. Our estimates reveal that housing and income are uncorrelated with MPCs, whereas liquid assets, particularly deposits, have a clear negative correlation with MPCs.

Second, to better understand what the raw correlations actually capture, we estimate MPCs allowing for interaction effects with a host of different control variables, including income and balance sheet components. When controlling for the stock of liquid assets prior to winning in the lottery, other characteristics hardly matter at all for consumption sensitivity. Liquidity stands out as the household characteristic most strongly related to MPC magnitude. Moreover, illiquid households display markedly higher MPCs both in the short and in the long run. Within the impact year, households in the lowest quartile of the liquidity distribution spend about twice as much of their windfall gain as households in the highest liquidity quartile. These strong effects of liquidity are consistent with previous findings in the literature. Examples are Misra and Surico (2014) who use survey evidence on the U.S. tax rebates, Leth-Petersen (2010) who studies the impact of a credit market reform on consumption in Denmark, Aydin (2015) who studies exogenously varying credit limits in a European retail bank, Baker (2014) who studies the interaction between household balance sheets, income and consumption during the U.S. Great Recession, and Agarwal, Chomsisengphet, Mahoney, and Stroebel (2015) and Gross, Notowidigdo, and Wang (2016) who study consumption dynamics around discontinuities in credit scores.

One might reasonably question whether the association between liquidity and MPCs simply reflects correlation with omitted variables. In the words of Parker (2015), the question is if the association uncovered is *situational*, in the sense that an individual's MPC depends on how liquid he happens to be at the time of winning. The alternative explanation is that liquidity happens to correlate with unobserved household characteristics that raise consumption sensitivity. From theory, the leading candidates are impatience and risk tolerance. However, we expect wealth, education and the share of risky assets in a household's portfolio to correlate with these unobservable characteristics. It is therefore striking that when wealth, education and the risky share (and several

other variables) are all controlled for together with liquidity, it is only the latter which significantly influences the consumption response to lottery prizes. This supports a *situational* interpretation of our findings.

Third, beyond household characteristics, we ask if consumption and spending responses vary with the amount won. The role of shock magnitude has so far been largely overlooked by the empirical literature, mostly because there has been little variation in the magnitude of the income shocks studied. However, this question is of great interest. It is directly relevant for transfer policies that aim to boost demand and it can inform the development of structural models on consumption choice. Regarding the latter, both liquidity constraints and discrete consumption and savings choice imply that larger prizes are associated with weaker consumption responses. Our findings qualitatively align with this prediction. When we group lottery winners by the amount won, the within-year consumption response to a marginal prize change is monotonically declining in prize size. It is about 0.73 in the lowest size quartile (USD 1,100 - USD 2,150), while it is 0.28 in the highest size quartile (above USD 8,926).²

Notably, our within-year MPC estimate for the lowest size quartile is of similar magnitude as provided by the literature studying U.S. tax rebates in 2001 and 2008. The bulk of existing evidence on exogenous 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, in which transfers per adult were between USD 300 and 600, they find total consumption responses in the range of 0.5 to 0.9 within three months of payment receipt.³ 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 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 \(2015\)](#) estimate the consumption response to large predetermined payments from the Alaska Permanent Fund. These responses are conceptually different from what we are estimating, which includes both announcement and imbursement effects. More comparable to our estimates, [Agarwal and Qian \(2014\)](#) study a transfer episode in Singapore, where natives received between USD 78 and USD 702, and find an average spending response around 80 percent of the stimulus received within ten months after the transfer was announced.

Finally, we characterize how MPCs vary across the joint distribution of prize size and liquidity. We find that MPCs decrease with both characteristics: Within any quartile of the distribution of lottery prizes, the consumption response declines with liquidity. And within any quartile of the distribution of liquidity, the consumption response declines with prize size. The former pattern corroborates that liquidity constraints are indeed important, the second pattern indicates that

²Throughout the article values are in 2000-prices and converted using the average exchange rate during the year 2000 (NOK/USD = 0.114).

³The 2007-2008 U.S. tax rebate distributed USD 300-600 to single individuals, USD 600-1,200 to couples, and in addition gave USD 300 for each child qualified for the tax credit. For details, see [Parker et al. \(2013\)](#).

discrete consumption choice is also playing a separate role. Notably, the importance of discrete choice is consistent with [Agarwal and Qian \(2014\)](#), who find that the aforementioned Singaporean fiscal transfer primarily was spent on objects such as apparel, travel and small durable goods.

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, but unlike the setting we study, these prizes were paid out gradually, obscuring comparison with our estimates. The latter study considers a lottery in the Netherlands where households received prizes of 12,500 euros. The Dutch findings stand out from ours and the tax rebate literature, as neither durable nor non-durable consumption responded by much.⁴ More recently, Swedish lotteries have been used to identify income effects on health, labor supply, and portfolio choice, but not on consumption.⁵

The extent to which evidence from lotteries generalizes to other income shocks, is debatable. [Ng \(1965\)](#) and recently [Crossley, Low, and Smith \(2016\)](#), argue that households might gamble to “convexify” their feasibility set when discrete-type purchases are desired. This would imply that our estimates are upward biased, as some of the winners have gambled precisely because they have a high MPC. Here it is reassuring that our estimated spending responses align well with the existing evidence on transfer policies. Moreover, gambling is widespread in Norway. According to the Norwegian state owned gambling entity Norsk Tipping, about 70 percent of the Norwegian adult population participated in one of their lotteries in 2012.⁶ Consistently with this observation, our descriptive statistics reveal only minor systematic differences between winners and non-winners, primarily that winners tend to hold riskier wealth portfolios. In addition, while conceptually the gambling-to-convexify argument might explain high MPC levels, it seems less relevant for our main contribution, namely to establish determinants of MPC heterogeneity. For all these reasons, while we do not claim that households never gamble to convexify, it seems unlikely that this mechanism is driving our main results.

Our findings are perhaps most interesting when cast against incomplete markets models, as developed by [Huggett \(1993\)](#), [Aiyagari \(1994\)](#), and [Carroll \(1997\)](#). In these models, households face uninsurable idiosyncratic risk and a borrowing constraint. As a result, households acquire a buffer stock of capital in order to prevent the constraint from becoming binding. The main determinant of households’ MPC is then their net wealth level. In contrast, our empirical findings indicate that wealth is unimportant, once liquidity is controlled for. While in conflict with the implication of classic buffer stock savings models, this finding is very much in spirit with a close

⁴While 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.

⁵Using data from Sweden [Cesarini, Lindqvist, Östling, and Wallace \(2016\)](#) study the effects of wealth on health and child development, [Briggs, Cesarini, Lindqvist, and Östling \(2015\)](#) study the effect on stock market participation and [Cesarini, Lindqvist, Notowidigdo, and Östling \(2015\)](#) the effect on labor supply.

⁶See [Norsk Tipping Annual Report 2012](#).

extension to it, namely two-asset frameworks distinguishing liquid from illiquid assets. The prime example here is the model by [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 indeed, we do find that greater housing wealth does not reduce spending responses after liquidity is controlled for, as it is liquid wealth that matters. Our finding that consumption responsiveness declines with the size of the lottery prize is also consistent with what such a two-asset model predicts.

The remainder of this article is organized as follows. Section 2 presents the institutional setting and the data. Section 3 provides the 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 concludes.

2 Institutional background, data and sample selection

We base our study on Norwegian data collected for administrative purposes. Norway levies both income and wealth taxes, and the data from 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 imputed measures of consumption using the household budget constraint. Moreover, as part of the tax return, Norwegian households must report larger lottery wins, above approximately USD 1,100, in their yearly tax report. Below we describe the data sources in some detail, explain the consumption measure we construct, and discuss the lottery sample in more detail.

2.1 Administrative tax and income records

Our main data source is the register of tax returns from the Tax Authority, which contains detailed information about all individuals' incomes and wealth, for the period 1993 to 2014.⁷ We combine these data with family identifiers from the population register to aggregate all income and wealth information at the family level.⁸ 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 (electronically) information on the value of the assets owned by the individual and administered by the employer or the intermediary, as well as information on the income on these assets.⁹ The tax authority then pre-fills the tax form for the individual for approval. Further,

⁷The quality of this data is similar to that in the Swedish data studied by [Calvet, Campbell, and Sodini \(2007\)](#). Until 2007, Sweden like Norway collected taxes on both individual income and wealth. In 2007, however, Sweden abandoned the wealth tax, leaving Norway as the only Scandinavian country with the arrangement of collecting detailed wealth information for the purpose of collecting a wealth tax.

⁸In Norway, labor (and capital) income is taxed at the individual level, while wealth tax is levied at the household level. For further institutional details see [Fagereng, Gottlieb, and Guiso \(2015\)](#).

⁹An important illiquid asset in the household portfolio is real estate wealth. Housing values from the tax registries are typically undervalued in Norway before 2010, when values for the purpose of wealth taxes were

these data have the advantage that there is no attrition from the original sample (apart from death or migration to another country) due to refusal by participants to consent to data sharing. In Norway, these records are in the public domain. Also, our income and wealth data pertain to all individuals, and not only to jobs covered by social security or individuals who respond to wealth and income surveys.

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.¹⁰ 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 US, 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](#)).

An alternative to using consumption surveys, is to impute consumption from income and wealth data in administrative tax records. We follow this approach. Equipped with the data on household balance sheets discussed above, we impute consumption for Norwegian households as in [Fagereng and Halvorsen \(2015\)](#) (for Norway) and [Browning and Leth–Petersen \(2003\)](#) (for Denmark).¹¹

Underlying the imputation procedure is the simple accounting relation of the household budget constraint:

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

stating that disposable income (Y) in each period must be either consumed (C) or saved (S). When

reassessed nation wide. Hence, for the exercises here in which we split the sample on net wealth we have utilized real estate transaction data (from the national portal for land and GIS data in Norway, Ambita Infoland) to backtrack housing values (with municipality wide price-indices) of households also prior to the 2010-reevaluation.

¹⁰[Pistaferri \(2015\)](#) provides a recent summary and discussion of the literature on the measurement of consumption.

¹¹[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, [Brinch, Eika, and Mogstad \(2015\)](#) and [Fagereng and Halvorsen \(2015\)](#) using Norwegian data, and [Kojien, Van Nieuwerburgh, and Vestman \(2015\)](#) using Swedish data. Other examples are [Browning, Gørtz, and Leth-Petersen \(2013\)](#), [Leth-Petersen \(2010\)](#), and [Autor, Kostøl, and Mogstad \(2015\)](#). [Browning, Crossley, and Winter \(2014\)](#) provide a recent review of this literature.

combining this equation with balance sheet data, a number of issues must be dealt with to obtain a consumption measure. Fagereng and Halvorsen (2015) provide a detailed discussion of these. Below we raise those issues that are most important and most relevant for our purpose of studying the consumption responses to lottery income.

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_j is the realized capital income from each asset j that the household held during the year. For interest rate expenditure on debt, d necessarily is negative. L_t is income from any other source, such as bequests or lottery prizes. Notably, all the components of Y_t are directly observed in the administrative tax records.

Consumption expenditure is imputed from the budget constraint 1, where S_t is "active", or net, savings. The challenge for consumption imputation is to calculate S_t , and in particular to adjust wealth accumulation for unrealized capital gains. To simplify, we assume that the household holds each asset over the entire year, and then rebalances its portfolio at market prices at the end of the year. S_t is then given by

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

where $W_t = \sum_j^J p_{jt} a_{jt}$ is end-of-year wealth, while p_{jt} is the end-of-year price and a_{jt} is the end-of-year stock of asset j . If a liability, a_{jt} is negative. As the expression shows, we need to isolate capital gains and subtract them from the total wealth change. From the administrative tax records, we directly observe W_t and the value held within each specific asset class k , $w_{kt} = \sum_j^{J_k} p_{jt} a_{jt}$; real estate, deposits, 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 price changes in stock prices with 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.

For housing, we do observe which asset the individual holds, as given by its address. However, the value of each house cannot be directly observed as a home is only infrequently transacted. Our approach here is to employ estimated market values of each specific home in our sample. To this end, data on all housing transactions in Norway over the sample period are combined with a set of key house characteristics, such as location, age, size, and number of balconies and bathrooms. REF DOCUMENT. As documented there, the estimated values fare well when compared to realized values in the observed transactions data. The R -squared of the estimates is 0.8XX. Compared to the existing literature, the home values we use are of high quality.

Finally, we observe deposits and debt directly, together with interest income and expenditure. Hence, capital gains are unproblematic for these assets.

In Appendix Figure A.1 we plot the distribution of our consumption measure for the years of our estimation period.

In sum, the description above reveals how our imputed consumption measure rests on two key assumptions. One, we assume there is no intra-year trading. Two, we assume each household holds the market portfolio of stocks, bonds and mutual funds. In addition, home values are necessarily estimated. We now discuss when the potential measurement errors that follow might be problematic for our purposes, and how we deal with these issues.

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

Intra-year trading involves the realization of capital gains or losses. These are observed and therefore unproblematic. However, we do face a potential problem when a lottery winner invests part of his prize in risky assets, say stocks, and does not realize the returns on this investment within the year. We will interpret the consequent wealth increase as "active" savings (S), and therefore subtract it from income when imputing consumption. However, the unrealized capital gains from purchase until the end of the year does not imply lower expenditure and hence should not be subtracted. As mean returns are positive, this measurement error is likely to cause a downward bias in our estimates of how lottery winnings affect consumption expenditure. 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 higher returns. To deal with this issue, we conduct all our analyses on various samples and report results in the online appendix. In particular we redo all our regressions on samples that exclude the 30 percent of the population who hold risky assets (stocks, bonds and mutual funds). We also control for household characteristics such as education and observed risk taking.

When housing wealth increases due to revaluation, this is directly observed and unproblematic for imputed consumption. Should a household transact in the housing market, this is also observed, together with sale and repurchase prices, so in principle this is unproblematic too. However, a home purchase is an extraordinary expenditure decision, involving large costs in several dimensions. To prevent households involved in home purchases from driving our results, we exclude all these household-year observations from our benchmark sample. For similar reasons we exclude household-years where there is change in the number of adults (e.g. by divorce or marriage), where there are extreme returns from financial markets, or where the household is a business owner or a farmer. Results without these sample restriction are reported in the online appendix. When we do include these observations, the conclusions from our main analysis remain qualitatively unchanged, but the precision of estimates falls considerably. This is to be expected, and similar to what Autor et al. (2015) report.

2.3 Gambling in Norway

In Norway, only two entities are allowed to offer gambling services: Norsk Tipping (mainly lotteries and sports betting) and Norsk Rikstoto (horse racing). Both of these entities 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 Norwegians above the age of 18 gambled in 2012 through their services.¹²

In the "pre-internet era" gambling in Norway took place mainly through one of the more than 5,000 commissioned sites (about one per every 800 adult Norwegian), usually a kiosk or a local super market. An individual filled out his or her betting forms and submitted them at one of the commissioned sites. Smaller prizes (less than around NOK 1,000, about 110 USD) were possible to cash out directly at any of these sites, whereas larger amounts were paid directly to the individuals bank account within a few weeks. Income from gambling in Norway is generally tax exempt, as is income from EU/EEA-area lotteries where the surplus is primarily directed to charity. However, Norwegian citizens are obliged to report lottery prizes to the tax authority if the total prize amount exceeds 10,000 NOK (about 1,100 USD). 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 suspicion of tax fraud and cause further investigation by the Tax Authority.¹³ As the reporting is done by display of the prize receipt, there is no scope for exaggerating such windfall gains. In 2007 the minimal reporting requirement was raised to 100,000 NOK (about 11,000 USD).

2.4 Descriptive statistics

The data on lottery prizes include all games arranged by Norsk Tipping and Norsk Rikstoto, and similar games in other EEA-countries. The data on lottery prizes therefore include a wide variety of games, such as scratch cards, bingo, horse racing and sports gambles. Note, however, that our lottery data does not include prizes won in card games such as poker or blackjack, nor does it include prizes won in other casino games.

As explained above, the threshold for reported lottery prizes was increased in 2007, from about 1,100 USD to 11,000 USD. To maintain the larger variation in the windfall gain, we therefore limit our attention to the period 1994-2006. Moreover, and importantly, we include only households who win once in our sample. The reason is that we want to study the responses to surprise income innovations, while for serial winners it is unclear whether yearly lottery winnings are best considered as unexpected. In particular, we want to exclude systematic gamblers in horse racing and sports betting, for whom prizes are part of their "regular" income.

Figure 1 displays the distribution of lottery prize in the sample we later use for our econometric analysis. We see a clear peak for the smallest prize bin, consisting of winners of 1,100 to 2,000

¹²See [Norsk Tipping Annual Report 2012](#).

¹³Norway 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](#)).

USD. 20 percent of our prizes are of this magnitude. Importantly, there is substantial variation in the size of lottery prizes, which allows us to study the role of the prize size.

[FIGURE 1 ABOUT HERE]

Table 1 displays basic summary statistics for winners and non-winners over the sample period. For the sample of winners, all characteristics are measured in the year before they won. For non-winners, we have drawn a random year during the sample period to represent their observations.

[TABLE 1 ABOUT HERE]

As we see from Table 1, winners and non-winners are similar, but there are some noteworthy differences. Winners are somewhat older than the rest of the population,¹⁴ have slightly smaller families, and somewhat lower education. The levels of income, consumption, and wealth are very similar. Any difference in mean wealth is primarily due to housing wealth. Regarding balance sheet composition, the last two rows of Table 1 reveal that a slightly higher share of winners own risky assets, and that their mean share of risky assets (stocks and mutual fund holdings relative to net wealth) is higher than among non-winners. Intuitively, this is not surprising as one expects households who play in lotteries to be more risk tolerant than non-players.¹⁵

To identify the effects of lottery prizes below, we will only rely on variation in prize sizes *within* the sample of winners. Still, for the purpose of external validity, we will use a sample of winners that are matched with non-winners, using propensity score matching. In the last column of Table 1 we report summary statistics on this matched sample. The matching is based on age, salary, debt, risky share and deposits, and the propensity score regression estimates are reported in Appendix Table A.1. As we see, these households are closer to the general population along most of the observed dimensions. Notably, all our results that follow hold irrespective of whether the samples are matched or not.

Given that our identification will come from size-variation among winners, one might worry that high-prize winners systematically differ from winners of low prizes. For instance, a shortcoming of our data is that we only observe how much households win, but not how much they have gambled. To a considerable extent, we ameliorate this problem by dropping serial winners from our sample, but it might still be true that households who win higher prizes are systematically different from those who win less. For instance, different betting games have different prize structures and are likely to attract different types of players. It could also be that winners of higher prizes bet more. We therefore study the relation between pre-determined characteristics and the amount won within

¹⁴To represent age, we use the age of the household head.

¹⁵In Table 1, the mean levels of each risky asset component (stocks and mutual funds) are higher among non-winners, while their mean shares of total wealth are higher among the winners. The reason for this discrepancy is that among non-winners, risky assets are heavily concentrated among the wealthiest households, whereas among the non-winners risky wealth is more evenly spread out, as indicated by the participation rates in the last row of Table 1.

the winner-sample.

Columns 1 and 2 of Table A.2 show that our main variable of interest, consumption, in the period before winning is uncorrelated with prize size. In Columns 2 and 3 we closely follow Cesarini et al. (2015), using a similar vector of controls for our Norwegian sample as they use in their sample of Swedish lottery winners. Even though we find small significant effects of age (that are likely related to the prize structure of lotteries played by younger and older individuals), we find that the vector of controls has little predictive power. Assuringly, the characteristics explain hardly any of the variation in the lottery prize, as measured by the R^2 . In Column 4 we further add lagged consumption growth, in case a particular trajectory of consumption growth could affect the gambling propensity of a household. We find no such pattern. Overall, it seems unlikely that systematic assignment of prize sizes will be a main driver of our results.

3 Consumption and savings responses to lottery prizes

Existing studies that estimate consumption responses to income shocks utilize a variety of slightly different econometric specifications. Our starting point is the three main specifications considered in this literature:

$$C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \beta_2 \text{lottery}_{i,t} + \alpha_i + \tau_t + u_{i,t} \quad (3)$$

$$\Delta C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \beta_2 \text{lottery}_{i,t} + \alpha_i + \tau_t + u_{i,t} \quad (4)$$

$$C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \beta_2 \text{lottery}_{i,t} + \beta_3 C_{i,t-1} + \alpha_i + \tau_t + u_{i,t} \quad (5)$$

Here i is a household identifier, t represents a year, C is the level of household consumption, $X_{i,t}$ is a vector of controls, $\text{lottery}_{i,t}$ is the lottery prize, and Δ is the one-year difference operator. α_i and τ_t denote household fixed effects and year fixed effects respectively, while $u_{i,t}$ is a normally distributed, mean zero error term. The coefficient of interest is always β_2 .

When interpreting the results that follow, it is key to recognize that it is variation in lottery prize size, $\text{lottery}_{i,t}$, that identifies β_2 . Moreover, we are including only lottery winners in our sample. Hence, our estimates of β_2 reflect the within-year consumption response to a *marginal increase* in lottery prize. At this point we are assuming this marginal effect to be independent of the amount won. We relax this assumption later on when studying the role of prize size.

[TABLE 2 ABOUT HERE]

Table 2 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 therefore less sensitive to outliers. As we move horizontally in Table 2, we gradually add controls for individual and year fixed effects, household characteristics, and lagged household income. The estimated responses lie between 0.3 and 0.4. Notably, once time and individual fixed effects are controlled for, the point estimates vary only across specifications (3) to (5), and hardly change as we add further controls.

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, often 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 that typically is less than 0.05, while the standard life-cycle model suggests that the MPC is very low for young households (less 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, Imbens et al. (2001) find a marginal propensity to save out of lottery prizes of 0.16, which implies a marginal propensity to consume of more than 0.80. In the tax rebate literature, Johnson et al. (2006) and Parker et al. (2013) find that households spend between 0.50 and 0.90 per dollar received on total consumption.

Importantly, when comparing our point estimates to existing evidence and models, one must bear two specific properties of our data in mind. First, our consumption measure includes both durables and non-durables.¹⁶ From theory we should expect a greater MPC once durables are included. Second, we are studying the increase in consumption that comes within the entire year of winning in the lottery. Again this should raise the expenditure response considerably above the instantaneous MPC typically emphasized in the literature. More specifically, if we assume that lottery prizes are approximately uniformly spread across the year, the average winner’s consumption response takes place over 6 months. In that perspective, our estimates can loosely be interpreted as a 6-month MPC.¹⁷

For the remainder of this article, we concentrate on the level specification given in equation 3. The reason is that among the alternatives, the level specifications (either OLS or LAD) allow us to use the largest sample of lottery winners because we do not need two consecutive consumption observations. This is of great value when we later explore dynamic responses and narrowly defined subgroups of the population. Moreover, as seen from Table 2, the point estimates of β_2 from the

¹⁶We can isolate two components of durable consumption and exclude them from the consumption expenditure measure. When we redo the estimates in Table 2 with this measure, we get very similar results, with less than 1 percentage point in difference different. These results are available upon request.

¹⁷In existing literature, only Kaplan, Moll, and Violante (2016) report 6-month MPCs. Their model-based prediction is a non-durable consumption response around 0.3 after a transfer of USD 1,000.

two level specifications represent a middle ground among the alternative estimates in Table 2. As the OLS estimates have the clearest interpretation, we will primarily focus on these. However, as we eventually move on to studying narrowly defined household groups with fewer observations, the OLS estimates become highly sensitive to outliers, and we will focus on the LAD estimates instead.

In addition to investigating the within-year consumption response of lottery winners, we consider saving and portfolio decisions by estimating the following equation:

$$\Delta Z_{i,t} = \beta_0 + \beta_1 X_{i,t} + \beta_2 \text{lottery}_{i,t} + \alpha_i + \tau_t + u_{i,t}, \quad (6)$$

where $\Delta Z_{i,t}$ is the change in balance sheet component Z from $t - 1$ to t . Results are reported in Table 3.

[TABLE 3 ABOUT HERE]

To ease comparison, column 1 in Table 3 restates the consumption response of 0.34, while the next report mean estimates for the responses of deposits, debt and the sum of stocks, bonds, and mutual funds. Our results suggest that within the year of winning, the mean household saves about 51 percent of the lottery prize in deposits, 7 percent in stocks, bonds and mutual funds, while about 12 percent of the lottery prize is used to repay debt.

3.1 Dynamic consumption 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 here we move on to estimating responses over the next 5 years. In order to assess the yearly dynamics and the cumulative consumption response, we estimate the following equation:

$$C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \sum_{k=0}^5 \beta_{2,k} \text{lottery}_{i,t-k} + \alpha_i + \tau_t + u_{i,t} \quad (7)$$

where $X_{i,t}$ is the vector of controls from specification VII in Table 2. The $\beta_{2,k}$'s are the main coefficients of interest. Each $\beta_{2,k}$ represents the share of a lottery prize won in year $t - k$ that is spent on consumption expenditures in year t . We estimate the cumulative responses as given by the sum of the $\beta_{2,k}$'s.

In addition to the dynamic consumption responses, we are interested in the dynamics of the balance sheet. We therefore estimate an equation similar to (7) for debt repayment and the accumulation of liquid assets:

$$\Delta Z_{i,t} = \beta_0 + \beta_1 X_{i,t} + \sum_{k=0}^5 \beta_{2,k} \text{lottery}_{i,t-k} + \alpha_i + \tau_t + u_{i,t} \quad (8)$$

where $\Delta Z_{i,t}$ is the change in balance sheet component Z from period $t - 1$ to t .

[FIGURE 2 ABOUT HERE]

Figure 2 shows the dynamic responses of consumption, car and boat purchases, deposits, the sum of stocks, bonds and mutual funds, and debt. The estimated within-year effects are the same as we found above. The bottom row in the figure plots the cumulative household responses to a marginal prize increase. Three findings stand out as particularly interesting. First, after five years, about 60 percent of the last dollar won is spent.¹⁸ The magnitude of this response contrasts with the textbook permanent income hypothesis, by which a substantial share of a temporary 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, either as deposits (30 percent), as stocks, bonds or mutual funds (4 percent) or as a reduction in debt (10 percent).¹⁹

Second, a substantial share of the prize induced spending occurs immediately, as the consumption response drops from 0.34 to 0.11 from the year of winning to the year thereafter. The expenditure boost falls gradually from year 1 to year 5 after winning. Deposits are used to support this consumption profile. They initially increase a great deal and then are gradually depleted to finance the spending profile after winning. The sharp decrease in consumption in the immediate year after winning is a further sign of lack of consumption smoothing. Part of the sharp movements may be due to durables purchases in the win-year, but the durables which we can observe, namely cars and boats, play a minor role here. About 2.5 percent of the lottery prize is spent on car and boat purchases in the win-year and about 4 percent of the lottery prize is spent on car and boat purchases five years after winning.

Third, the dynamics of deposits, debt, and stocks, bonds, and mutual funds are consistent with the existence of fixed adjustment costs in debt and asset markets. Debt is discretely repaid upon impact, but thereafter evolves with the same amortization profile as before winning. Similarly, upon impact winners discretely save a fraction in stocks, bonds and mutual funds, and thereafter deplete these negligibly. The movements after winning are small and statistically insignificant. Deposits, in contrast, continue to move down by a considerable amount in the years after winning. The discrete movements in debt and in stocks, bonds, and mutual funds are consistent with the existence of fixed costs of adjusting these balance sheet components, whereas the smoother decline in deposits align with a much smaller role of adjustment costs here.

¹⁸Note that 0.6 is the cumulative *marginal effect* of the last dollar won, and not the share of the entire prize that is spent after 5 years. The two are identical only if the marginal effect is independent of lottery prize size. We later show that the marginal effect decreases with prize size, and hence the cumulative responses to the entire prize typically are considerably above 0.6.

¹⁹Note that because the impulse response of each savings and spending object is estimated separately, the sum of responses needs not sum to 100 percent.

4 Household characteristics and MPC heterogeneity

The aggregate estimates above constitute a natural starting point, but mask potential heterogeneity in how households respond to income shocks. We here turn to this heterogeneity, and address which household characteristics, if any, are associated with cross-sectional variation in MPCs. Our approach is motivated by macroeconomic models with incomplete markets, which are used to address an increasingly wide range of issues, such as the likely effects of fiscal stimulus, monetary policy shocks, house price changes, and economic inequality. This literature highlights five main variables as likely determinants of household level consumption dynamics. In addition to age, these are income, net wealth (see e.g. [Huggett, 1993](#), [Aiyagari, 1994](#), and [Carroll, 1997](#)), liquid assets, ([Kaplan and Violante, 2014](#)) and debt ([Mian, Rao, and Sufi, 2013](#)).

4.1 Covariances for the “sufficient statistics” approach

The cross-sectional co-variances between individual MPCs and various household characteristics have recently been emphasized as key to understand macroeconomic phenomena. For instance, [Auclert \(2016\)](#), calculates a “sufficient statistic” for how strongly redistribution propagates monetary policy shocks. Key to this statistic is the covariance between individuals’ MPCs and their balance sheet exposure to interest rate changes, coined “unhedged interest rate exposure”. A similar sufficient statistics approach is used by [Berger et al. \(2015\)](#) in assessing how house price changes affect the economy, where the key covariance is between MPCs and housing wealth. Likewise, [Auclert and Rognlie \(2016\)](#) emphasize the cross-sectional covariance between MPCs and income in addressing how inequality shocks might influence aggregate demand. However, for all these covariances direct empirical evidence is in short supply. We therefore start our analysis of heterogeneity by displaying the unconditional correlations between MPCs and household income and balance sheet components. We construct the correlations by first sorting households into percentiles of the variables in question and then estimate the marginal consumption responses to lottery prizes within each percentile. Using the estimated consumption responses and their respective standard errors, we simulate the correlations 100,000 times. The reported correlations and standard errors in [Table 4](#) are the means and standard errors from the samples of simulated correlations. [Table 4](#) presents the correlations.

[TABLE 4 ABOUT HERE]

The marginal propensity to consume out of a lottery prize is virtually unrelated to income. The correlation is insignificant and small, particularly compared to the correlations between MPCs and liquidity. In the second row liquid assets refer to the sum of deposits, stocks, bonds and mutual funds. The correlation is -0.13 . The third row sharpens the definition of liquidity further, by computing the MPC-correlation with deposits only. Here the correlation is about the same, -0.13 . Thereafter, debt and housing are uncorrelated with consumption responses. The correlation for net

worth has the same sign as implied by conventional single-asset buffer stock theory, but the pattern is a bit weaker than for deposits and liquid assets. Finally, the bottom row implies a sizable correlation between MPCs and unhedged interest rate exposure, -0.15 . The latter is crudely measured as the sum of financial wealth net of debt since nearly all mortgage debt in Norway is flexible rate mortgages. Notably, [Auclert \(2016\)](#) finds a correlation of smaller magnitude, -0.06 , based on Italian survey data from [Jappelli and Pistaferri \(2014\)](#).

In light of the work discussed above, these correlations suggest that redistribution between different income groups plays a negligible role for aggregate demand and the propagation of aggregate shocks to the economy. What matters is instead how impulses are distributed across households with different balance sheets. For instance, the high correlation between MPCs and uncovered interest rate exposure suggests that redistribution might play a key role in the propagation of interest rate shocks in Norway. Notably, the prevalence of flexible rate mortgages in Norway is key to this pattern, and hence the implication cannot be generalized to countries dominated by fixed rate mortgages.

4.2 Which household characteristics matter?

The household characteristics in [Table 4](#) are most likely correlated both with each other and with other household attributes. Hence, the correlations displayed say little about which household characteristic that actually matters for MPC variation. To shed light on this issue, we modify our estimating equation and allow for interaction effects between prizes and a multitude of explanatory variables that might affect MPCs. The specification is as follows:

$$C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \beta_2 \text{lottery}_{i,t} + \beta_3 \text{lottery}_{i,t} * Z_{i,t-1} + \beta_4 Z_{i,t-1} + \alpha_i + \tau_t + u_{i,t} \quad (9)$$

where $Z_{i,t-1}$ contains variables we expect might be correlated with MPCs. All these variables are lagged, except age, to avoid problems of reverse causality. β_3 is the coefficient of interest, revealing whether the variable systematically relates to the consumption response to winning a prize of size $\text{lottery}_{i,t}$.

We estimate equation (9) both with each variable of interest and its interaction with the lottery prize separately, and a joint specification where all variables of interest and their interactions with the lottery prize are included. The results of the latter regression indicates which factors that have a separate effect on MPCs, over and above any correlation with the other explanatory variables. The characteristics we consider are liquid assets, income, net wealth, debt, the share of wealth held in risky assets, education, and age. We first group the liquid assets into one, before we distinguish between stocks, bonds, mutual funds, and deposits.

[TABLE 5 ABOUT HERE]

[Table 5](#) presents the estimated interaction coefficients. Among the candidates considered, liquid

assets stand out as the most important determinant of household level MPCs. Higher holdings of liquid assets are associated with lower MPCs. The estimated effect is substantial. For example, given the results from specification I, the within-year MPC is 0.39 for a household with no liquid assets and 0.19 for a household with USD 100,000 in liquid assets at the beginning of the year. The estimated interaction effect is significant at the 0.1 percent level, also when the other candidate interaction effects are controlled for in the last column. Neither net wealth, income, debt, the portfolio share held in risky assets, education, nor age, have similar effects.

[TABLE 6 ABOUT HERE]

Rather than including just the sum of deposits, stocks, bonds, and mutual funds, Table 6 studies them separately. As we see from the table, the interaction with each of them are important for MPCs. It is only stocks that are not statistically significant at the five percent significance level. As we proceed, we will focus on the interaction with deposits, as this is the main component of the liquid portfolio. However, just as Table 6 suggests, all the results below also hold if we utilized the entire basket of liquid assets instead.

Our results suggest that the focus on asset liquidity in the recent incomplete market literature, as exemplified by Kaplan and Violante (2014), is well-grounded. Yet, with what degree of confidence should we interpret the estimated liquidity effects as *situational*, meaning that an individual’s MPC depends on how liquid he happens to be at any given point in time? The alternative explanation is that liquidity happens to correlate with some unobserved household characteristic such as patience or risk aversion, which also affects the propensity to consume out of income increases. Parker (2015) argues in this direction when exploring survey evidence on how households responded to the 2008 U.S. stimulus payments. We certainly cannot exclude the possibility that such unobserved behavioral characteristics affect our estimated liquidity effects. But if they were driving our estimates entirely, we expect significant interaction effects with either net wealth, education or the risky portfolio share. We do not find such interaction effects and the fact that the association between liquidity and MPCs holds also when interactions with all those three variables are included, lends support to a situational interpretation of our results.

4.3 Deposits and the dynamics of consumption and savings.

To shed further light on the role of liquidity, we move on to a less parametric approach than the interaction specification. We divide the sample of households by how much deposits they hold and then display consumption and savings dynamics across the deposit distribution.

We split all households into four equal-sized groups (within-year quartiles), given by their holdings of deposits in the year prior to winning. We denote these categories as low deposits, low-mid deposits, high-mid deposits, and high deposits.²⁰ We first investigate how the within-year

²⁰The cut-offs of deposit quartiles vary with the year won. Mean (median) deposits in each quartile are: USD 404 (270), USD 2,844 (2,385), USD 10,390 (9,140), USD 52,347 (36,967).

MPC varies with the level of deposits by estimating the following equation:

$$C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \sum_{j=1}^4 \beta_{2,j} lottery_{i,t} I_{j,t-1} + \alpha_i + \tau_t + u_{i,t}, \quad (10)$$

where $I_{j,t-1}$ equals 1 if the household belongs to deposits quartile j at time $t - 1$, and is zero otherwise. Results are displayed in Table 7.

[TABLE 7 ABOUT HERE]

Again we see the negative relationship between MPCs and liquidity. The within-year consumption response is 0.44 in low-deposit quartile, and gradually falls to 0.22 in the most liquid quartile. Moreover, among the three main savings vehicles considered, it is deposits that most closely reflects the consumption pattern. The marginal propensity to save in deposits increases from 0.4 in the least liquid quartile, to 0.73 among the most liquid. The propensity to save in stocks, bonds and mutual funds, is also increasing in liquidity, but to a weaker extent than deposits savings. Debt stands out with an opposite pattern, as the least liquid winners tend to spend more of their prize on debt repayment than what the most liquid do. This stands to reason, as the households with high initial deposit levels are already choosing not to repay debt.

We next explore how the dynamics of the households responses vary with the initial level of deposits. The equation we estimate is:

$$C_{i,t} = \beta_0 + \beta_1 X_{i,t} + \sum_{j=1}^4 \sum_{k=-1}^3 \beta_{2,k,j} lottery_{i,t-k} I_{j,t-k-1} + \alpha_i + \tau_t + u_{i,t} \quad (11)$$

where $I_{j,t-k-1}$ indicates if the household belongs to deposit quartile j at time $t - k - 1$.

[FIGURE 3 ABOUT HERE]

The first row of Figure 3 shows how the dynamic consumption response varies between deposit quartiles. Interestingly, while the initial responses decline with deposits, as we saw in Table 7, there is little variation in the three subsequent years. The highest deposit quartile has the smoothest consumption profile. In short, the main difference in consumption responses to lottery prizes occurs within the year of winning.

The remaining rows of Figure 3 display how the dynamics of the savings responses vary with initial deposit levels, distinguishing between saving in deposits, debt, and stocks, bonds and mutual funds. The dynamics of deposits closely follows the consumption response. For all liquidity quartiles there is a sharp increase in the year of winning, consistently with Table 7, thereafter some of these savings are depleted for a year or two. In stocks, bonds and mutual funds, the three lower quartiles conduct most of their savings within the year of winning, while the highest liquidity group has a slightly smoother profile; increasing their holdings of stocks, bonds and mutual funds also in the year after winning. Debt repayment is a one-off event for all liquidity levels, and as we saw in

Table 7 it decreases with initial liquidity.

4.4 Wealthy hand-to-mouth households

All the above results indicate an association between liquidity and marginal propensities to consume and save. By implication, there may in principle exist households who are wealthy, yet behave in a hand-to-mouth fashion by letting consumption track swings in disposable income, as emphasized in the structural model of [Kaplan, Violante, and Weidner \(2014\)](#). We here investigate further the prevalence of such households.

To this end, we first sort households into quartiles along two dimensions: (i) liquid assets (deposits + stocks + bonds + mutual funds), and (ii) net illiquid assets (housing wealth - debt).²¹ We then combine the two, leaving us with a total of 16 liquid-illiquid asset groups. We then estimate the consumption response to lottery prizes within each of these 16 groups. The results are visualized in Figure 4.

[FIGURE 4 ABOUT HERE]

The right-hand side plot in Figure 4 shows the population density within each of the illiquid-liquid asset categories, among our prize winners. As one might expect, there are hardly any households who are in both the highest liquid asset quartile and the lowest net illiquid asset quartile, located in the figure’s north-east corner. There is, however, a considerable fraction in the south-west 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 about 24 percent of all the winners in our sample. Almost 18 percent of the population are in the lowest liquidity quartile but above the bottom illiquid wealth quartile.²²

In Figure 4’s left-hand side plot, we see how higher liquidity is associated with lower MPCs within illiquidity quartiles. Among households with the least illiquid wealth, the negative association between liquidity and MPCs is visible only between the two highest liquidity levels. Among households with higher liquidity holdings, the relationship between liquidity and MPC holds throughout. Hence, there are households with both high levels of illiquid wealth and high MPCs in our sample, and their high MPCs do coincide with low holdings of liquid wealth. Both findings are consistent with structural models that emphasize the distinction between liquid and illiquid assets in explaining household consumption.

²¹For expositional purposes we here use the entire sum of liquid assets, rather than just deposits. Results using deposits instead are qualitatively identical, and are available from the authors upon request.

²²[Kaplan et al. \(2014\)](#) define hand-to-mouth (HtM) households as households with liquid reserves less than half the monthly wage (wages are usually paid monthly in Norway) and wealthy as households with positive net illiquid assets. When following this approach we find that the universe of households in Norway contains 9.5 % Wealthy-HtM, 13 % Poor-HtM and 77.4 % Non-HtM. Their respective MPCs (with standard errors in parentheses) are W-Htm: 0.462 (0.036); P-Htm: 0.428 (0.034), and N-Htm: 0.313 (0.011).

5 Size of income shocks and MPCs

The lottery prizes we study are larger and vary more between individuals than the income shocks typically studied in the existing empirical literature. As reviewed in the introduction, the US tax rebates in 2001 and 2008 were between USD 600 and USD 1,200 (see Johnson et al., 2006 and Parker et al., 2013), and the tax rebate in Singapore in 2011 was below USD 700 (Agarwal and Qian, 2014)), whereas the prizes we consider vary from approximately USD 1,100 to USD 1,000,000. Our data therefore allows us to provide novel insight into how consumption responses depend on the size of income shocks.

Before turning to the evidence, we summarize the main reasons why the size of the prize may 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 will choose to save more of the income innovation. Second, if purchases of high-return assets involves discrete transaction costs, then the rate of return on saving a received prize will effectively increase with the amount won. Third, several consumption decisions contain an element of discrete choice, in particular this is so for durable goods, but also for certain non-durables like vacations. Winners of small prizes might therefore choose to spend their entire income innovation, and possibly even more if they can borrow, to make a discrete purchase. The larger is the prize won, the less likely are such lumpy purchases to dominate spending decisions. For all these reasons, we expect the MPC to decline with prize magnitude, at least for the lottery prizes that are present in our sample ($>$ USD 1,100).

We investigate the effect of prize size by dividing the winner population into four groups according to the amount won. The four groups are "small size" (USD 1,100-2,150), "small-mid size" (USD 2,150-5,332), "large-mid size" (USD 5,332-8,926), and "large size" ($>$ USD 8,926). We then estimate consumption responses for each of these four groups separately. Results using both OLS and LAD are provided in Table 8.²³

[TABLE 8 ABOUT HERE]

The LAD estimates are included here because discrete choice type mechanisms are likely to cause high MPC outliers among winners of small prizes. This is exactly what we see in Table 8 when we compare the OLS and the LAD estimates in the small size group. The OLS estimated MPC is 1 while the LAD estimate is just above 0.7. Thereafter, the two estimators give similar results, and the main pattern is clear: MPCs do decline considerably with the amount won, reaching approximately 0.3 in the highest prize size quartile.

Upon reading Table 8, correct interpretation is essential. We are in effect estimating a piecewise linear regression, and hence each estimate represents the mean (OLS) or median (LAD) consumption response to a *marginal* increase in prize within each group. Hence, when the winners in the

²³As Table 8 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.

highest quartile have a mean marginal propensity to consume of 0.316, that does *not* necessarily imply that they spend 31.6 percent of their entire prize. To find that share, we must impose additional assumptions on how the estimated marginal responses relate to the inframarginal responses. One approach here is to assume that the marginal consumption response estimated in the lower size quintiles also applies to the inframarginal portion of the prize won by high-prize winners. Our OLS estimates then imply that a high-prize winner spends 100 percent of his first USD 2,150, 68.7 percent out of the next USD 3,182, 58.8 percent of the prize money between USD 5,332 and 8,926, and then finally 31.6 percent of the prize above USD 8,926. Based on this back-of-the-envelope calculation, Table 8 indicates that winners at the 25th percentile of the prize size distribution spend the entire amount won within the year (73 percent according to the LAD estimate), winners of the median prize size spend 82 percent (66 percent according to the LAD-estimate), and winners at the 75th percentile spend 72 percent (59 percent according to the LAD estimate).

The final row of Table 8 reports 5-year cumulative marginal effects. These estimates represent how much households spend out of the last dollar won, over a 5-year period including the year of winning. Unsurprisingly, the pattern here is the same as it was for the impact year effects. Moreover, following the same procedure as above we can use these estimates to calculate how much households spend out of their entire prize over the first 5 years after winning. For winners at the 25th percentile and below, the implied estimate is identical to the cumulative marginal response, 86 percent. Winners of a median sized prize are estimated to spend 73 percent out of the amount won, and winners at the 75th percentile are estimated to spend 65 percent of the amount won over the first 5 years after winning.

Above we sketched out two main explanations why lottery size might matter: Liquidity constraints and discrete savings or expenditure choice. Now, if the former explanation drives the size effects in Table 8, we expect to see a particularly strong size effect among the least liquid households. On the other hand, if the size effects are due to discrete choice alone, we expect the size effect to hold across the deposit distribution.

In order to illuminate the relative importance of these channels, we here divide the population into 16 groups, differentiated in two dimensions. First, we distinguish between the 4 deposit quartiles (deposits in the year before winning). Second, we distinguish between the 4 size quartiles. Within each of these 16 groups we then estimate the MPC. We here use the median estimator (LAD) as the sample sizes necessarily become rather small and hence sensitive to outliers.

By construction, the 16 different groups need not be equally large. However, according to the right-hand panel in Figure 5, they are essentially equi-sized, consistently with the results of Table A.2 (random assignment). The main results are visualized in the left panel of Figure 5 and in Table 9, which reports standard errors together with the same point estimates as displayed in the figure.

[FIGURE 5 ABOUT HERE]

The pattern is stark: Both size and liquidity matter. Within each deposit quartile, the MPC

decreases with prize size. Within each size quartile, the MPC decreases with deposits. Consistently with the importance of liquidity constraints, the magnitude by which the MPC drops when moving from “small” to “small-mid” prizes is markedly greater in the lowest deposit quartile than for more liquid households. Yet, prize size does matter across the deposit distribution, even among the most liquid, indicating that discrete choice is key to understand how households respond to income shocks.

6 Conclusion

We use detailed administrative data from tax and income records of Norwegian households to document the main determinants of how households respond to unanticipated income shocks, as identified through lottery prizes. The average within-year propensity to consume out of a marginal increase in lottery prize lies around 0.35, but this estimate varies considerably with predetermined household liquidity and the amount won. The estimate is approximately halved when we move from the least to the most liquid quartiles of households, and it falls by approximately two thirds when we move from the lowest to the highest lottery prizes observed. Among winners of prizes in the lower half of the prize size distribution, the *marginal* propensity to consume is above 0.6, and in a back-of-the-envelope calculation, the winner of a median sized lottery prize is found to consume more than 65 percent of the *entire* prize received, within the year of winning.

Liquidity is in our study measured as cash and deposits as well as stocks, bonds, and mutual funds. Importantly the association between these assets and MPCs is not matched by any other household characteristics. In particular, neither income, net wealth, education nor risky portfolio share correlate significantly with MPCs once liquidity is controlled for. In contrast, the liquidity effects do not disappear when all these alternative explanatory variables are taken into account. While we cannot say that we control for all possible household characteristics that might be correlated with both liquid wealth and consumption sensitivity, we do believe our estimates support a situational interpretation of the estimated association between liquidity and MPCs. For instance, if heterogeneous impatience underlies the pattern we uncover, by causing both low holdings of liquidity and high willingness to consume windfall gains, then we expect measures of net wealth and education to also correlate with MPC variability. However, they do not. Similarly, if heterogeneous risk aversion is driving the liquidity-MPC association, then we expect the risky portfolio share to pick up some of this effect. However, it does not.

In terms of the magnitude of income shocks, we find a clear pattern: greater prizes are associated with lower MPCs. Two underlying explanations seem important here: First, credit constraints no longer bind when prizes are above a certain size, and second, discrete consumption choice can motivate households to spend far larger fractions of low than of high prizes. By separating between liquidity and prize size, we find both mechanisms to be relevant.

Our findings have direct implications for transfer policies that aim to stimulate aggregate de-

mand. For instance, our estimates provide no support for the popular suggestion that in order for transfers to stimulate the economy more effectively, they should be targeted to low-income households. Instead, such policies should be directed to the least liquid. Of course, such a systematic policy will face its own problems by rewarding households for being illiquid, but that is beyond the scope of our study to explore.

More importantly than the direct policy implications, our findings provide guidance for the development of structural models. In particular, the distinction between liquid and illiquid assets, emphasized in several recent contributions following [Kaplan and Violante \(2014\)](#), seems highly relevant for understanding how households respond to income shocks. Net wealth, as emphasized by traditional models of consumer choice, seems less relevant than liquid wealth for MPCs, as several households are wealthy, yet let their consumption expenditure respond strongly to lottery prizes. Finally, our results indicate that discrete purchases are an important driver of micro level consumption dynamics.

References

- 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,” Mimeo 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. R. (2014): “Debt and the consumption response to household income shocks,” SSRN Research Paper No. 2541142.
- BERGER, D., V. GUERRIERI, G. LORENZONI, AND J. VAVRA (2015): “House Prices and Consumer Spending,” NBER Working paper No. 21667.
- 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.
- BRINCH, C., L. EIKA, AND M. MOGSTAD (2015): “What can we learn about household consumption from information on income and wealth?” Mimeo Statistics Norway.
- BROWNING, M. AND M. D. COLLADO (2001): “The response of expenditures to anticipated income changes: panel data estimates,” *The 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,” *The Economic Journal*, 123, 401–428.
- BROWNING, M. AND S. LETH-PETERSEN (2003): “Imputing consumption from income and wealth information,” *The Economic Journal*, 113, F282–F301.

- CALVET, L. E., J. Y. CAMPBELL, AND P. SODINI (2007): “Down or Out: Assessing the Welfare Costs of Household Investment Mistakes,” *Journal of Political Economy*, 115, 707–747.
- CARROLL, C. D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *The Quarterly Journal of Economics*, 112, 1–55.
- (2001): “A theory of the consumption function, with and without liquidity constraints,” *The Journal of Economic Perspectives*, 15, 23–45.
- CARROLL, C. D., J. SLACALEK, K. TOKUOKA, AND M. N. WHITE (2014): “The distribution of wealth and the marginal propensity to consume,” ECB Working paper No. 1655.
- CESARINI, D., E. LINDQVIST, M. J. NOTOWIDIGDO, AND R. ÖSTLING (2015): “The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries,” NBER Working paper No. 21762.
- 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*, forthcoming.
- CROSSLEY, T. F., H. LOW, AND S. SMITH (2016): “Do Consumers Gamble to Convexify?” *Journal of Economic Behavior and Organization*, forthcoming.
- FAGERENG, A., C. GOTTLIEB, AND L. GUISO (2015): “Asset Market Participation and Portfolio Choice over the Life-Cycle,” *The Journal of Finance*, forthcoming.
- FAGERENG, A. AND E. HALVORSEN (2015): “Imputing consumption from Norwegian income and wealth registry data,” Statistics Norway Discussion Paper No. 831.
- 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,” NBER Working paper No. 21228.
- GROSS, T., M. J. NOTOWIDIGDO, AND J. WANG (2016): “The Marginal Propensity to Consume Over the Business Cycle,” NBER Working paper No. 22518.
- HSIEH, C.-T. (2003): “Do Consumers React to Anticipated Income Shocks? Evidence from the Alaska Permanent Fund,” *American Economic Review*, 93, pp. 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 (2016): “Monetary policy according to HANK,” NBER Working paper No. 21897.

- 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.
- 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,” NBER Working paper No. 22319.
- KUENG, L. (2015): “Explaining Consumption Excess Sensitivity with Near-Rationality: Evidence from Large Predetermined Payments,” NBER Working paper No. 21772.
- 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,” *The 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?” *The 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.
- MIAN, A., K. RAO, AND A. SUFI (2013): “Household Balance Sheets, Consumption, and the Economic Slump,” *Quarterly Journal of Economics*, 128, 1687–1726.
- 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. (2015): “Why Don’t Households Smooth Consumption? Evidence from a 25 Million Dollar Experiment,” NBER Working paper No. 21369.
- 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.

Tables

Table 1: Summary statistics, 1994-2006

Variable	Non-winners		Winners		Matched Winners	
	(N = 4,093,070)		(N = 20,578)		(N = 18,520)	
	Mean	S.D.	Mean	S.D.	Mean	S.D.
Age _t	46.51	(20.03)	51.70	(15.08)	49.79	(14.24)
Year _t	2000.54	(4.34)	2000.48	(3.36)	2000.43	(3.38)
Family Size _t	2.41	(1.41)	1.86	(1.11)	1.89	(1.13)
No. of Children under 18 _t	0.52	(0.93)	0.27	(0.69)	0.30	(0.72)
Years of Education _t	13.22	(3.12)	12.47	(2.57)	12.49	(2.58)
Income After Tax _{t-1}	24.57	(52.48)	23.72	(11.51)	23.66	(11.14)
Salary _{t-1}	21.64	(25.56)	22.63	(21.05)	22.80	(20.01)
Consumption _{t-1}	20.63	(15.92)	20.66	(13.41)	20.66	(13.37)
Lottery _t	.	.	8.63	(15.24)	8.65	(15.35)
Net Wealth _{t-1}	50.34	(108.29)	47.38	(86.51)	46.50	(87.08)
Debt _{t-1}	35.81	(58.62)	30.14	(39.01)	32.11	(39.89)
Car & Boat _{t-1}	2.63	(7.02)	3.41	(5.91)	3.43	(5.89)
Housing Wealth _{t-1}	68.59	(102.34)	62.55	(78.63)	63.35	(78.53)
Deposits _{t-1}	14.39	(31.30)	12.67	(22.36)	13.35	(23.34)
Stocks _{t-1}	1.13	(13.51)	0.65	(4.61)	0.48	(3.87)
Bonds _{t-1}	0.88	(9.19)	0.61	(4.45)	0.56	(4.30)
Mutual Funds _{t-1}	1.17	(7.47)	1.04	(4.59)	0.86	(4.09)
Risky Share of Balance Sheet _{t-1}	0.07	(0.19)	0.08	(0.20)	0.06	(0.17)
Share of Households Owning Risky Assets _{t-1}	0.25	(0.44)	0.30	(0.46)	0.27	(0.45)

Notes: The sample of "Non-winners" contains one randomly selected observation from the period we observe the non-winning households. Summary statistics for winners are beginning of year values of the year they win. All variables except age, family size, children, education length, risky share of balance sheet and share of households owning risky assets are in thousands of 2000 USD.

Table 2: The marginal propensity to consume out of lottery prizes

	Specifications:			
	I	II	III	IV
Levels (OLS) (N = 266,263)	0.385*** (0.008)	0.340*** (0.009)	0.335*** (0.009)	0.336*** (0.009)
Differences (OLS) (N = 173,341)	0.349*** (0.013)	0.377*** (0.015)	0.377*** (0.015)	0.377*** (0.015)
Levels (LAD) (N = 266,263)	0.449*** (0.011)	0.308*** (0.009)	0.301*** (0.010)	0.312*** (0.010)
Differences (LAD) (N = 173,341)	0.386*** (0.013)	0.306*** (0.013)	0.306*** (0.013)	0.309*** (0.014)
Dynamic (OLS-IV) ¹ (N = 173,341)		0.265*** (0.014)	0.265*** (0.014)	0.266*** (0.014)
Time-fixed effects	Yes	Yes	Yes	Yes
Household-fixed effects	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
Lagged Income	No	No	No	Yes

Notes: Results are obtained using the matched sample of winners. Controls include age, age², family size, family size² and no. of children under 18. LAD = Least Absolute Deviation estimator (Median). OLS-IV = Arellano-Bover/Blundell-Bond estimator (Arellano and Bover, 1995; Blundell and Bond, 1998). Standard errors are in parentheses and clustered at the household levels in the OLS specifications. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.

¹⁾ The dynamic-IV estimation does not include household-fixed effects in any specification.

Table 3: Marginal propensities to save and spend out of lottery prizes

Consumption	Deposits	Stock, bonds & mutual funds	Debt
0.336*** (0.009)	0.507*** (0.013)	0.071*** (0.009)	-0.118*** (0.009)

Notes: Results are obtained using the matched sample of winners. Each column represents a separate regression. Controls include household-fixed effects, time-fixed effects, income_{t-1} , age, age^2 , family size, family size^2 , and no. of children under 18. Estimation method: OLS. Standard errors are in parentheses and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively. N = 266,263.

Table 4: MPC heterogeneity. Unconditional correlations

X	Corr(MPC_t, X_{t-1})	S.E.
Income	0.0181	(0.0994)
Liquid assets	-0.1370	(0.0237)
Deposits	-0.1311	(0.0195)
Debt	0.0191	(0.0330)
Housing	-0.0201	(0.0366)
Net wealth	-0.0943	(0.0311)
Unhedged interest rate exposure	-0.1519	(0.0333)

Notes: Each estimate is constructed by sorting the population into percentiles on the X-variable and estimating equation (3) in each percentile. Estimation method: OLS. Liquid assets = deposits + stocks + bonds + mutual funds; Unhedged interest rate exposure = liquid assets - debt. Standard errors are generated from Monte Carlo simulations.

Table 5: The marginal propensity to consume out of lottery prizes, cross-term regressions

Dependent variable: Consumption _t									
	I	II	III	IV	V	VI	VII	VIII	IX
Lottery _t	0.390*** (0.010)	0.379*** (0.021)	0.354*** (0.010)	0.335*** (0.011)	0.334*** (0.012)	0.338*** (0.009)	0.422*** (0.030)	0.281*** (0.098)	0.254*** (0.099)
Lottery _t *liquid assets _{t-1}	-0.0020*** (0.0004)								-0.0017*** (0.0004)
Lottery _t *income _{t-1}		-0.0019* (0.0008)							-0.0016 (0.0010)
Lottery _t *net wealth _{t-1}			-0.0002 (0.0001)						0.0000 (0.0001)
Lottery _t *debt _{t-1}				0.0001 (0.0002)					-0.0001 (0.0002)
Lottery _t *risky share _{t-1}					-0.0421 (0.0439)				-0.0249 (0.0440)
Lottery _t *education _t						0.0002 (0.0033)			0.0004 (0.0038)
Lottery _t *age _t							-0.0017** (0.0005)	0.0041 (0.0039)	0.0048 (0.0039)
Lottery _t *age _t ²								-0.0551 (0.0367)	-0.0587 (0.0365)

Notes: Each column represents a separate regression. All squared variables are divided by thousand. Controls include all variables which are cross-termed with lottery, household-fixed effects, time-fixed effects, income_{t-1}, age, age², family size, family size², and no of children under 18. Estimation method: OLS. Standard errors are in parentheses and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively. N = 281,131.

Table 6: The marginal propensity to consume out of lottery prizes, cross-term regressions, continued

Dependent variable: Consumption _t					
	I	II	III	IV	V
Lottery _t	0.379*** (0.010)	0.337*** (0.009)	0.339*** (0.009)	0.347*** (0.009)	0.348*** (0.099)
Lottery _t *deposits _{t-1}	-0.0018*** (0.0004)				-0.0011* (0.0005)
Lottery _t *stocks _{t-1}		-0.0045* (0.0023)			-0.0025 (0.0024)
Lottery _t *bonds _{t-1}			-0.0061*** (0.0016)		-0.0032* (0.0016)
Lottery _t *mutual funds _{t-1}				-0.0103*** (0.0030)	-0.0074* (0.0037)
Lottery _t *income _{t-1}					-0.0016 (0.0009)
Lottery _t *net wealth _{t-1}					0.0000 (0.0001)
Lottery _t *debt _{t-1}					-0.0001 (0.0002)
Lottery _t *risk share _{t-1}					0.0424 (0.0594)
Lottery _t *education _t					0.0003 (0.0038)
Lottery _t *age _t					0.0050 (0.0040)
Lottery _t *age _t ²					-0.0609 (0.0366)

Notes: Each column represents a separate regression. All squared variables are divided by thousand. Controls include all variables which are cross-termed with lottery, household-fixed effects, time-fixed effects, income_{t-1}, age, age², family size, family size², and no. of children under 18. Estimation method: OLS. Standard errors are in parentheses and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively. N = 281,131.

Table 7: Heterogeneous household responses. Quartiles of deposits

Dependent variable	Deposit quartile			
	Low	Low-mid	High-mid	High
Consumption	0.436*** (0.021)	0.416*** (0.020)	0.336*** (0.022)	0.224*** (0.021)
Deposits	0.398*** (0.024)	0.437*** (0.030)	0.530*** (0.025)	0.727*** (0.038)
Stocks, bonds & mutual funds	0.031*** (0.009)	0.056*** (0.014)	0.086*** (0.021)	0.081** (0.028)
Debt	-0.149*** (0.019)	-0.123*** (0.022)	-0.087*** (0.016)	-0.027 (0.018)
N	70,542	70,450	70,643	69,496

Notes: Each coefficient represents a separate regression. Controls include household-fixed effects, time-fixed effects, income_{t-1} , age, age^2 , family size, family size², and no. of children under 18. Estimation method: OLS. Standard errors are in parentheses and clustered at the household level. *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.

Table 8: Heterogeneous household responses. Quartiles of lottery prize sizes

Dependent variable:	Lottery prize size quartiles:			
	Small	Small-mid	Large-mid	Large
Consumption (OLS)	1.007*** (0.111)	0.687*** (0.051)	0.588*** (0.027)	0.316*** (0.009)
Consumption (LAD)	0.731*** (0.053)	0.619*** (0.028)	0.485*** (0.016)	0.281*** (0.010)
5 Year Cumulative Marginal Consumption Response (LAD)	0.862*** (0.1421)	0.640*** (0.0639)	0.541*** (0.0329)	0.404*** (0.0121)
N	68,373	65,309	67,070	61,339

Notes: Each coefficient represents a separate regression. The groups are: small (USD 1,100-2,150), small-mid (USD 2,150-5,332), large-mid (USD 5,332-8,926), and large (> USD 8,926). Controls include household-fixed effects, time-fixed effects, $income_{t-1}$, age, age^2 , family size, $family\ size^2$, and no. of children under 18. Estimation methods: OLS and LAD. Standard errors are in parentheses and clustered at the household level for OLS. Standard errors for the cumulative marginal consumption response are calculated using bootstrapping (100,000 replications). *, **, *** denote significance at the 5, 1 and 0.1 percent level, respectively.


Table 9: Heterogeneous consumption responses. Quartiles of deposits and lottery prize sizes

		Deposit quartiles			
		Low	Low-mid	High-mid	High
Lottery size quartiles	Small	1.047*** (0.109)	0.745*** (0.087)	0.720*** (0.105)	0.490*** (0.140)
	Small-mid	0.762*** (0.053)	0.640*** (0.050)	0.559*** (0.054)	0.437*** (0.083)
	Large-mid	0.663*** (0.029)	0.546*** (0.042)	0.390*** (0.035)	0.386*** (0.039)
	Large	0.354*** (0.032)	0.325*** (0.025)	0.242*** (0.030)	0.216*** (0.034)

Notes: Each coefficient represents a separate regression. Controls include household-fixed effects, time-fixed effects, age, age², family size, family size² and no. of children under 18. Estimation method: LAD. Standard errors are in parentheses and clustered at the household level. *, **, *** denote significance at the 5, 1, and 0.1 percent level, respectively.


Figures

Figure 1: Distribution of lottery prizes, 1994-2006




Notes: Lottery prizes are in thousands of 2000 USD. Each bin is 1,000 USD wide, starting from USD 1,100.

Figure 2: Dynamic household responses to lottery prizes




Notes: Dotted lines represent 95 percent confidence interval. Standard errors are clustered at the household level. Standard errors of estimated cumulative effects are obtained from monte carlo simulations. Controls include household-fixed effects, time-fixed effects, $income_{t-1}$, age, age^2 , family size, $family\ size^2$ and no. of children under 18. Estimation method: OLS. $N = 180,119$.

Figure 3: Heterogeneous household responses to lottery prizes. Quartiles of deposits




Notes: Results are obtained using the matched sample of winners. Dotted lines represent 95 percent confidence interval. Standard errors are clustered at the household level. Controls include household-fixed effects, time-fixed effects, $income_{t-1}$, age, age^2 , family size, $family\ size^2$ and no. of children under 18. Estimation method: OLS.

Figure 4: Heterogeneous consumption responses. Quartiles of liquid and net illiquid assets



Notes: Controls include time-fixed effects, $income_{t-1}$, age, age^2 , family size, $family\ size^2$ and no. of children under 18. Estimation method: OLS. Total N: 266,263.


Figure 5: Heterogeneous consumption responses. Quartiles of deposits and lottery size



Notes: Controls include household-fixed effects, time-fixed effects, $income_{t-1}$, age, age^2 , family size, $family\ size^2$ and no. of children under 18. Estimation method: LAD. Total N: 210,783.

A Appendix

Figure A.1: Consumption dispersion



Notes: The figure plots the kernel density of the consumption measure, valued in 2000 USD for selected time periods.

Table A.1: Propensity score regression

	Dependent variable: Non-winner _t
Age _t	-0.0057*** (0.0001)
Salary _{t-1}	-0.0023*** (0.0001)
Debt _{t-1}	0.0008*** (0.0001)
Risk share _{t-1}	-0.1079*** (0.0119)
Deposits _{t-1}	0.0022*** (0.0001)

Notes: The sample contains all lottery winners and a random selection of year observations of non-winners (one year observation for each non-winner household). Standard errors in parentheses. Estimation method: probit. *, **, *** denote significance at the 5, 1, and 0.1 percent level, respectively. N = 4,114,154.

Table A.2: Random assignment of lottery prizes

Regressors	Dependent variable: lottery prize size _t			
	I	II	III	IV
Consumption _{t-1}	0.0222 (0.0127)	0.0220 (0.0129)	0.0226 (0.0137)	0.0290 (0.0194)
Age _t		-0.138* (0.0632)	-0.137* (0.0641)	-0.161* (0.0813)
Age _t ²		0.0013* (0.0006)	0.0013* (0.0006)	0.0014* (0.0007)
Family size _t		0.166 (0.809)	0.167 (0.809)	-0.421 (1.361)
Family size _t ²		0.0107 (0.167)	0.0106 (0.167)	0.0934 (0.302)
No. of children under 18 _t		0.411 (0.379)	0.414 (0.381)	1.203* (0.536)
Born in Norway _t		-0.222 (0.756)	-0.221 (0.756)	-0.045 (0.875)
Couple _t		-0.569 (0.445)	-0.559 (0.463)	-0.366 (0.609)
College-graduated _t		-0.449 (0.370)	-0.448 (0.370)	-0.417 (0.463)
Net income _{t-1}			-0.0019 (0.0171)	-0.0152 (0.0213)
Consumption growth (percentage) _{t-1}				-0.0020 (0.0121)
F-statistics for joint significance of regressors	3.09	1.80	1.63	1.43
[p-value]	[0.08]	[0.06]	[0.09]	[0.15]
Partial R-squared of regressors	0.0014	0.0014	0.0014	0.0029
N	13,064	13,064	13,064	9,235

Notes: Each column represents a separately estimated regression of lottery prize among winners on predetermined characteristics. All regressions include time-fixed effects. Standard errors are in parentheses. *, **, *** denote significance at the 5, 1, and 0.1 percent level, respectively.