What Would You Do with $500? Spending Responses to Gains, Losses, News and Loans[*]

Andreas Fuster[†]   Greg Kaplan[‡]   Basit Zafar[§]

February 9, 2018

Abstract

We use survey questions about spending in hypothetical scenarios to investigate features of propensities to consume that are useful for distinguishing between consumption theories. We find that (i) responses to unanticipated gains are vastly heterogeneous (either zero or substantially positive); (ii) responses to losses are much larger and more widespread than responses to gains; and (iii) even those with large responses to gains do not respond to news about future gains. These three findings suggest that limited access to disposable resources is an important determinant of spending behavior. We also find that (iv) households do not respond to the offer of a one-year interest-free loan, suggesting that this is not as a consequence of short-term credit constraints; and (v) people do cut spending in response to news about future losses, suggesting that neither is this a consequence of myopia. A calibrated two-asset life-cycle precautionary savings model can account for these features of propensities to consume, but cannot account for (vi) a positive extensive-margin size-effect for spending responses to gains, which suggests that non-convexities due to durability, salience or attention costs may also be important.

JEL Codes: D12, D14, E21.
Keywords: Consumption, Savings, Marginal Propensity to Consume, Survey.

1 Introduction

A large amount of research has been devoted to measuring marginal propensities to consume (MPCs). The majority of this work has focused on searching for observable characteristics that correlate with the heterogeneity in MPCs out of income shocks. However, this search has been largely fruitless: the only observable characteristic that has been robustly shown to correlate with MPCs is holdings of liquid wealth, and even then the explanatory power of wealth for MPC heterogeneity is weak.[1] In addition, most of the empirical work has focused on the consumption response to small, unanticipated one-time gains. Other than the cross-sectional correlation with liquid wealth, the limited variation in income changes has provided little in the way of evidence that is useful for evaluating theoretical models of consumption. In particular, it contains few findings that can help distinguish between alternative theories for why some households hold very little liquid wealth, and why some other households have large MPCs despite having moderate amounts of liquid wealth. Among the possible explanations for holding little liquid wealth are short-term credit constraints, low lifetime resources, access to illiquid investment opportunities, myopia, and financial illiteracy. Among the possible explanations for non-trivial MPCs among households with moderate liquid wealth are hand-to-mouth-like behavior due to committed expenditures or expense risk, and mental accounting.

In this paper, we use survey evidence on reported spending in various scenarios to generate new evidence that is useful for testing and refining existing models of consumption. Rather than focusing on correlates with observed heterogeneity as in the existing literature, we use within-person variation in consumption responses to different hypothetical treatments. In addition to MPCs out of unexpected gains of different amounts, we also elicit MPCs out of unexpected losses, news about future gains, news about future losses, and an interest-free loan. Thus, the advantage of our approach is that we generate variation in shocks (in terms of size, timing, and sign) that is otherwise very difficult to generate cleanly in natural settings. Moreover, using within-person variation generates results that are free from other confounds. For example, although it is possible to examine consumption responses to positive and negative income shocks in observational data, individuals who receive positive shocks are likely to differ along observable and unobservable dimensions from those who receive negative shocks, which limits the inferences that one can draw.

---

[1]Early work in this literature failed to find strong evidence for a correlation between liquid wealth and MPCs (see e.g. Johnson et al., 2006), but more recent work that uses larger samples and richer data routinely finds a significant correlation (see e.g. Fagereng et al., 2016; Baker, 2017; Aydin, 2018). However the R-squared measures remain very low.

Comparing spending responses across these treatments yields several insights about consumption behavior, which we demonstrate by implementing the treatments inside both simple theoretical models and calibrated life-cycle precautionary savings models with one or two assets.

Our first three findings describe a pattern of MPC behavior that suggests individuals act as if access to disposable resources is limited. First, as in the existing literature, we find a large amount of heterogeneity in consumption responses to small unexpected gains. Most people do not change their spending when given a $500 gain, but there is a set of people who spend a substantial fraction of the $500. Second, we find evidence of sign asymmetry. Spending responses to losses are larger and more widespread than spending responses to the same size gains.[2] Third, we find that very few respondents say that they would increase spending in response to news about a future gain, even those respondents who indicate that they would increase spending in response to an actual gain. These three findings are all consistent with a subset of the population acting as if they do not have access to disposable resources.

Our next two findings provide insights into the possible reasons why this group might act in this way. Fourth, we find that respondents do not increase their spending when offered a one-year interest-free loan, suggesting that short-term credit constraints are not a key factor in explaining high MPCs. Fifth, whereas very few respondents react to news about a future gain, the majority of respondents do react to news about a future loss, including those who react strongly to an immediate loss. This finding suggests that even low-wealth individuals are at least somewhat forward-looking and is evidence against extreme forms of myopia.

We then show that a calibrated two-asset life-cycle model, as in Kaplan and Violante (2014), is broadly consistent with the magnitude and distribution of MPCs, as well as with the five aforementioned findings from our hypothetical treatments. But the model is not consistent with our sixth finding. We find that as we increase the size of the windfall from $500 to $2,500 to $5,000, a larger fraction of respondents say that they would increase their spending. We refer to this as a positive extensive-margin size effect. Neither the two-asset model, nor any of the other baseline models, are consistent with this finding because none deliver meaningful predictions on the extensive margin, and all predict an intensive margin size effect in the opposite direction. In order to generate a meaningful extensive margin of MPCs, we extend the two-asset model by introducing a small "response cost" of modifying consumption in response to the treatments. This cost is intended to capture, in a parsimonious

---

[2]This asymmetric response to gains and losses is consistent with evidence from the the expiration of the 2013 payroll tax cuts (see Zafar et al., 2013, Bracha and Cooper, 2014, and Sahm et al., 2015).

way, the effects of salience, inertia or cognitive costs of changing consumption plans. We show that this modification improves the model *vis-a-vis* the MPC evidence, but still cannot generate a positive extensive-margin size effect. We suggest behavioral modifications that, when combined with the response cost, might improve the fit of the model in this regard.

One obvious explanation for why many households act as if their access to disposable resources is limited, is that they in fact possess very little liquid wealth. Indeed, this is the sole explanation in our calibrated models and almost all existing models of consumption behavior.[3] However, although we find a strong correlation in our data between liquid wealth and MPCs out of losses, we do not find a significant correlation between liquid wealth and MPCs out of gains. One possible reason is that in reality, liquid wealth is an imperfect proxy for disposable resources. For example, different households have different pre-committed expenditures, different expense risk and different access to informal credit, and hence consider themselves hand-to-mouth at different levels of liquid wealth. Another possible reason is that some behavioral phenomenon, such as mental accounting, lead households to act *as if* they are hand-to-mouth. In fact, Parker (2017) finds that spending responses to the 2008 stimulus payments are related with certain behavioral characteristics, such as impatience (but not with measures of self-control or procrastination). However, we are not aware of any existing single phenomenon that would lead high liquid wealth households to respond in the ways that we observe: high MPCs out of current gains and losses, low MPCs out of future gains, high MPCs out of future losses and low MPCs out of an interest-free loan.

The existing literature has followed one of two approaches to estimating MPCs. One strand of the literature uses what Parker and Souleles (2017) label the "revealed preference" approach, in which consumption is measured using data on actual housing expenditures. These data come either from household surveys or financial datasets—e.g. Consumer Expenditure Survey (Johnson et al., 2006), Kilts-Nielsen Consumer Panel (Parker, 2017), or banks and other financial service providers (Gelman et al., 2014; Baker, 2017; Ganong and Noel, 2017; Aydin, 2018) —or by backing out expenditures from administrative data on income and wealth (Fagereng et al., 2016). The revealed preference approach uses these data to estimate MPCs either by cleverly exploiting natural experiments that mimic unexpected changes in household budgets—e.g. fiscal stimulus payments (Parker et al., 2013), lottery winnings (Fagereng et al., 2016), or mortgage modifications (Ganong and Noel, 2017)—or by

---

[3]An important exception is Campbell and Hercowitz (2015) who propose a model in which some households have liquid wealth that has been earmarked for a foreseen large future expenditure.

extracting the transitory component of stochastic income fluctuations (Blundell et al., 2008). Another set of studies uses what Parker and Souleles (2017) label the "reported preference" approach, in which individuals are asked how their spending would respond in hypothetical or actual scenarios. A large part of the reported preference literature elicits qualitative spending responses using survey questions that follow Shapiro and Slemrod (2003). More recently, there has been a growing body of work, including ours, that elicit quantitative spending responses (Jappelli and Pistaferri, 2014; Graziani et al., 2016; Christelis et al., 2017). Using strategically-designed survey questions in conjunction with structural models has also been fruitfully applied to other questions related to household financial decisions (Ameriks et al., 2015, 2016).

This paper sits firmly in the "reported preference" approach. Our data come from a survey of 2,586 household heads from the NY Fed's Survey of Consumer Expectations, an online rotating panel of US household heads. We ask respondents to report how they would adjust their spending over the next quarter in response to receiving or losing dollar amounts ranging from $500 to $5,000, with the gain/loss occurring either now or in the future, or coming from a loan. Each respondent participates in two or more such treatments, allowing us to study within-person differences in responses. Ideally, we would compare actual spending data under these alternative scenarios rather than hypothetical spending data. The trade-off is that by using reported preferences rather than a revealed preferences, we have flexibility in designing treatments. We are not aware of any natural experiments that would allow us to compare actual spending data across scenarios in a controlled way.

Within the reported preference approach, our paper makes two main contributions. First, the variation that we generate across our scenarios is more extensive than has been implemented to date. This is important since this enables our setup to generate a much richer set of empirical results against which we can evaluate existing theoretical models of consumption behavior. More specifically, while previous studies have considered the size and sign effect, we are not aware of any study that has investigated response to news (about gains or losses) or loans. The closest paper to ours, fielded contemporaneously, is Christelis et al. (2017) who also use hypothetical scenarios (in a Dutch household survey) to study sign and size asymmetry. Our findings on these points are qualitatively similar (which is reassuring, given differences in the survey population, the design of the questions, and the size of the income shocks). Second, we contribute to this literature by using a survey instrument that is more precise, yet more flexible, than those that have been used in the existing literature. We discuss the advantages of our survey in Section 2. These advantages include a

4

relatively neutral wording of the question that does not prime respondents towards a non-zero response, a two-stage set-up which allows respondents to first think about whether they would change their spending at all and then by how much, an explicit mention of the spending horizon, and the ability to report an MPC outside the range of 0 to 1. We believe each of these features makes it easier for respondents to report their *true* MPCs.

An important underlying assumption when using reported preferences is that the responses have information content for what households would actually do in response to a (current or future) cash windfall or loan. Parker and Souleles (2017) provide a comprehensive analysis of this issue. They compare *reported* responses to hypothetical tax rebates with actual spending responses from past tax rebates and stimulus payments, and broadly conclude that the two approaches yield similar estimates.[4] In addition, Parker et al. (2013) added a similar question to Shapiro and Slemrod (2003) to the 2008 Consumer Expenditure Survey and found that respondents who said that they mostly spent their 2008 fiscal stimulus payment did in fact spend more. Thus, the reported preference approach has fared quite well when compared to the revealed approach, at least for current gains. Within the reported preference approach, studies have found a close correspondence between the ex ante MPC (that is, the MPC estimate based on how respondents say they will change their spending) and the ex post MPC (that is, the estimate based on what respondents say about how their spending changed) for one-time transfers (Shapiro and Slemrod, 2003; Sahm et al., 2010). Bunn et al. (2017) compare responses to retrospective survey questions (asking how spending adjusted in response to income being higher/lower than had been expected) to those from a survey featuring hypothetical scenarios similar to ours, and find that the sign asymmetry (which we also document) is present in both, although average MPCs are slightly smaller in the hypothetical scenarios. Similarly, for the payroll tax cut, Graziani et al. (2016) find that the ex-post reported MPC tends to be somewhat higher than the ex-ante MPC. Beyond consumption responses to income changes, other recent papers, mostly in the context of labor markets, have shown that the reported approach yields preference estimates that are similar to

---

[4]Parker and Souleles (2017) reach three main findings. First, reported spending is highly informative about actual spending in the sense that those that say they would mostly spend, do actually spend much more than those who say they would mostly save. Second, the average MPC implied by the reported response is similar to the MPC from the actual response. Third, for the reported responses, they find that MPCs are not correlated with liquid wealth, whereas for the actual responses they are. Interestingly, we find that in the one treatment of our survey that generates a large implied average MPC—the loss treatment—there is a strong correlation with liquid wealth. Thus, survey questions are able to detect a liquid wealth effect for losses, but we also do not find a correlation with liquid wealth for gains.

those from revealed choice (Mas and Pallais, 2017), and are predictive of real-world choices (Wiswall and Zafar, 2018). There is a growing consensus that the reported approach yields meaningful responses when the hypothetical scenarios presented to respondents are realistic and relevant for them, as is the case for the scenarios that we consider.

A third important contribution of our paper, relative to the existing literature, on reported preferences, is that we implement the hypothetical survey questions inside calibrated consumption models. Most of our insights stem from comparing the predictions of consumption theory with the elicited consumption responses, a step which the existing literature has largely avoided.

The remainder of the paper is structured as follows. Section 2 describes the survey instrument and the various treatments. Section 3 presents the results from the baseline gains treatment, and Section 4 analyzes the additional treatments (news, losses, and loans). Implications for theory are discussed in Section 5, and the last section concludes.

2   Data

2.1   NY Fed Survey of Consumer Expectations

Our data come from four modules added to the Survey of Consumer Expectations (SCE), which is a monthly survey fielded by the Federal Reserve Bank of New York. The SCE is an internet-based survey of a rotating panel of approximately 1,300 heads of household from across the US. The goal of the survey is to elicit expectations about a variety of economic variables, such as inflation and labor market conditions. Respondents participate in the panel for up to twelve months, with a roughly equal number rotating in and out of the panel each month. Respondents are invited to participate in at least one survey each month.

The survey is administered by the Demand Institute, a non-profit organization jointly operated by The Conference Board and Nielsen. The sampling frame for the SCE is based on that used for The Conference Board's Consumer Confidence Survey (CCS). Respondents to the CCS, itself based on a representative national sample drawn from mailing addresses, are invited to join the SCE internet panel. Each survey typically takes fifteen to twenty minutes to complete, and respondents receive $15 for completing a survey. The response rate for first-time invitees hovers around 55 percent, and for repeat respondents is around 80 percent.[5]

---

[5]See Armantier et al. (2016) for technical background information on the SCE, and www.

The four modules were added to the end of the monthly surveys in March 2016, May 2016, January 2017 and March 2017. Repeat and active panelists (i.e., those who were not participating in the SCE for the first time) were invited to participate in the modules. Because of the panel nature of the SCE some respondents answered multiple modules – those that were less than 12 months apart. In total we collected 9,061 responses to hypothetical spending questions from 2,586 panelists.[6]

Demographic and financial characteristics of respondents in the sample align well with corresponding characteristics of the US population. We report several of these characteristics in Table 1, along with their population counterpart from the 2015 American Community Survey or the 2013 Survey of Consumer Finances. For example, the average age of respondents in our sample is 50.4 years, of which 36% report annual household income of less than $50,000. The corresponding numbers in the US population are 51.1 years and 37%. 73% of respondents are homeowners, compared to a homeownership rate of 59% in the ACS. 75% of our respondents are white and non-Hispanic, compared to 69% of household heads in the ACS. One notable difference between our sample and the US population is in education. Households in our sample are on average more highly educated than the overall US population—56% of our respondents have at least a Bachelor's degree, compared with 31% of household heads in the ACS. We conjecture that this is partly due to differential internet access and computer literacy across education groups.[7]

2.2 Survey Instrument

Our baseline survey instrument asked respondents to report how they would change their spending behavior in response to an unexpected gain in resources. Respondents are first asked in what direction each of their spending, debt payments, and savings would change in response to the windfall. Next, respondents who say that they would change their (spending; debt payment; savings) are asked for the magnitude of the change. For example, the survey instrument for the $500 gain is as follows.

Respondents are first asked:

> Now consider a hypothetical situation where you unexpectedly receive a one-time payment of **$500** today.

---

newyorkfed.org/microeconomics/sce.html for additional information.

[6]There were a total of 9,086 scenarios submitted to these panelists, with 25 non-responses (corresponding to less than 0.3% of observations).

[7]We conduct our analysis unweighted; however, weighting to make the sample representative of the US population does not alter the qualitative patterns. This is illustrated in Appendix Table A-1, which provides a weighted version of our summary Table 3.

Table 1: Sample Characteristics

	Overall	March 2016	May 2016	January 2017	March 2017	U.S. Population
Sample Size	2586	1086	1087	1190	1180	
Demographics						
White/Non-Hispanic	0.75	0.75	0.75	0.75	0.76	0.69
Age	50.43	50.74	50.62	51.09	50.73	51.06
Education BA+	0.56	0.57	0.57	0.56	0.55	0.31
Married	0.64	0.65	0.64	0.63	0.65	0.50
Homeowner	0.73	0.73	0.73	0.73	0.73	0.59
Midwest	0.25	0.25	0.26	0.25	0.25	0.21
Northeast	0.20	0.19	0.18	0.19	0.21	0.18
South	0.33	0.34	0.34	0.33	0.32	0.38
West	0.22	0.22	0.22	0.23	0.22	0.24
Financial Characteristics						
Income <= 50k	0.36	0.36	0.37	0.36	0.36	0.37
Income 50k-100k	0.36	0.36	0.35	0.36	0.35	0.30
Income 100k+	0.28	0.28	0.28	0.28	0.29	0.31
Liquid Financial Assets $\geq$ 20k	0.50	0.48	0.51	0.51	0.51	0.35
Non-housing Debt > 20k	0.35	0.35	0.35	0.33	0.34	0.23
Net Worth > 200k	0.42	0.48	0.48	0.38	0.37	0.34

SCE respondents are equal-weighted.

For demographics, comparison is with the American Community Survey 2015; for financial characteristics, comparison is with the Survey of Consumer Finances 2013.

8

We would like to know whether this extra income would cause you to change your spending behavior in any way over the **next 3 months**.

Please select only one
- Over the next 3 months, I would **spend/donate more** than if I hadn't received the $500
- Over the next 3 months, I would **spend/donate the same** as if I hadn't received the $500
- Over the next 3 months, I would **spend/donate less** than if I hadn't received the $500

Please select only one
- Over the next 3 months, I would **pay off more debt (or borrow less)** than if I hadn't received the $500
- Over the next 3 months, I would **pay off the same amount of debt** as if I hadn't received the $500
- Over the next 3 months, I would **pay off less debt (or borrow more)** than if I hadn't received the $500

Please select only one
- Over the next 3 months, I would **save more** than if I hadn't received the $500
- Over the next 3 months, I would **save the same** as if I hadn't received the $500
- Over the next 3 months, I would **save less** than if I hadn't received the $500

Respondents are then asked by how much they would change their behavior for each category for which they do not select the middle option (spend/donate the same; pay off the same amount of debt; save the same). For example, a respondent who indicates that they would spend/donate more is asked the following question:

> You indicated that you would increase your spending/donations over the next 3 months following the receipt of the $500 payment.
> How much more would you spend/donate than if you hadn't received the $500?

The quantitative response to the increase or decrease in spending/donating forms the basis of our estimates of the marginal propensity to consume (MPC).[8] We refer to this baseline treatment for eliciting MPCs as the GAIN treatment:

GAIN: MPC over 1 quarter out of a one-time unexpected receipt of $Y, with Y={500; 2,500; 5,000}

Our survey instrument differs from those used in the existing literature on hypothetical consumption responses in several ways. The majority of this literature has based

---

[8]Note that the survey question distinguishes between paying down debt and saving. While paying down debt is a form of saving (and enters the same way in simple budget constraints), consumers may think of paying down debt as distinct from saving. Therefore, consistent with the approach used in the prior literature, we also make this distinction.

their survey instrument on the categorical response wording of Shapiro and Slemrod (2003), who focus on tax rebates. They ask respondents to choose between three uses of their tax rebate: (i) mostly spend; (ii) mostly save; or (iii) mostly pay off debt. Parker et al. (2013) added a Shapiro-Slemrod style question to the 2008 Consumer Expenditure Survey and found that respondents who said that they mostly spent their 2008 fiscal stimulus payment did in fact spend 75 cents more per dollar than those who said they mostly saved their stimulus payment.

More recently, the literature has started to employ survey questions that elicit direct quantitative responses for spending changes. For example, Jappelli and Pistaferri (2014), use the following question in the Survey of Household Income and Wealth (SHIW): "*Imagine you unexpectedly receive a reimbursement equal to the amount your household earns in a month. How much of it would you save and how much would you spend? Please give the percentage you would save and the percentage you would spend.*"[9] Whereas the Shapiro-Slemrod instrument asks a qualitative question and hence requires additional assumptions to be informative about the level of MPCs, the Japelli-Pistaferri (JP) instrument directly elicits a quantitative MPC. Similarly, Graziani et al. (2016) use a quantitative instrument to elicit consumption responses to the 2011 payroll tax cuts: "*Please indicate what share of the extra income [from the payroll tax cut] you are using or plan to use to save or invest, spend or donate, and pay down debts.*"

Christelis et al. (2017) use the following question to measure quantitative responses to hypothetical gains in an online survey of Dutch households: "*Imagine you unexpectedly receive a one-time bonus from the government equal to the amount of net income your household earns in (one-month / three months). In the next 12 months, how would you use this unexpected income transfer?*", with the respondent asked to allocate 100 points to saving; repaying debt; durable spending, and non-durable spending. They employ a similar wording for hypothetical losses, which are framed as one-time taxes. Finally, in a survey of British households, Bunn et al. (2017) ask respondents about the retrospective quantitative change in spending in response to unanticipated shocks to income over the past year. More specifically, they first ask households whether their income differed from what they expected a year ago, and if so, by how much. Next, they ask them how they adjusted their spending over the previous year in response to this unexpected change in income.

An advantage of eliciting a quantitative response is that it gives a direct measure of the individual MPC; this can then be aggregated up to yield the average

---

[9]The SHIW is administered to a sample of Italian households. The translation of the survey question from Italian to English is reproduced from Jappelli and Pistaferri (2014).

MPC, which is often the parameter of interest to policymakers. Although this elicitation approach may be more challenging for respondents to answer (as opposed to qualitative questions), it provides a much richer set of evidence to compare with theory.

We believe that our survey instrument is more precise than those in the existing literature. First, we explicitly state the size of the windfall, which we then vary, allowing us to measure potential size effects. Second, we start by asking respondents if they would change their spending at all, before asking the amount by which they would change their spending. This allows a more precise estimate of zero MPCs and does not prime respondents towards a non-zero response. We then ask only those respondents who say that they would actually change their spending behavior about how much they would spend. Third, our survey instrument is more explicit than most in the existing literature about the time horizon over which we are asking about spending responses (one quarter, in our case).[10] This is important because almost all economic models predict that any windfalls will ultimately be entirely spent over the respondents' remaining lifetime. So without explicitly stating a time horizon, it is difficult to make any comparisons with theory. Fourth, our elicitation strategy does not impose a household's MPC to be between 0 and 1. We leave open the possibility that an unexpected cash windfall may lead some respondents to increase their consumption by a larger amount than the windfall. This could occur if, for example, the respondent had been saving toward an expense and the windfall leads them to be alter the timing of the expense, as would be predicted by the model of Campbell and Hercowitz (2015).

2.3   Treatments

Differences in the survey instrument aside, our study advances the literature by also exposing respondents to a series of additional treatments beyond MPCs for income windfalls. These treatments are designed to elicit aspects of consumption behavior that are particularly useful for evaluating the predictions of theoretical models of consumption. In addition to the GAIN treatment, we conducted the following four treatments:

LOSS: MPC over 1 quarter out of a one-time unexpected loss of $500.

NEWS-GAIN: MPC over 1 quarter out of unexpected news about a one-time gain of $X, with X={500; 5,000}, 1 quarter from now.

---

[10]Bunn et al. (2017) and Christelis et al. (2017) specify time horizons of one year.

Table 2: Treatments and Survey Months.

	Mar-16	May-16	Jan-17	Mar-17
Gain				
$500 Gain	X [1085]		X [594]	
$2500 Gain		X [540]		
$5000 Gain	X [361]	X [1084]	X [595]	
$500 in 3 months	X [362]			
$5000 in 3 months			X [594]	
Loss				
$500 Loss	X [362]			X [1174]
$500 Loss in 3 months			X [594]	X [586]
$500 Loss in 2 years				X [589]
Loan				
$5000 Loan		X [541]		

Number of respondents given in square brackets. For Jan-17, half the sample got the $500 Gain and $500 News-Loss blocks, and the other half got the $5,000 Gain and News-Gain blocks.

NEWS-LOSS: MPC over 1 quarter out of unexpected news about a one-time loss of $500 Z quarters, with Z={1, 8}, from now.

LOAN: MPC over 1 quarter out of an unexpected interest-free loan of $5,000 to be repaid 1 year from now.

In each module, we exposed respondents to two possible treatments. The months in which the treatments were fielded are displayed in Table 2. For example, in the May 2016 module, all respondents were exposed to the $5,000 GAIN treatment and, in addition, were randomly assigned to either the $2,500 GAIN treatment or the $5,000 LOAN treatment. The order of the treatments within each survey was randomized. This design allows us to compare how the same respondent's spending behavior differs across alternative scenarios, thus providing a way to control for fixed unobserved individual characteristics.

In addition, some treatments were fielded in multiple months. For example, as shown in Table 2, the $5,000 GAIN treatment was fielded in March 2016, May 2016 and January 2017. This allows us to study whether the response distributions are consistent over time. Moreover, the panel structure of the survey ensures that some people appear in multiple surveys and, in some instances, in the same treatment in different months. This allows us to investigate whether individual respondents report stable spending responses. Finally, for some treatments we asked follow-up questions regarding the timing of spending adjustments (within the one-quarter horizon) and the composition of spending adjustments across different categories; these follow-up questions are discussed further below. The full texts of the survey instruments for each treatment are reproduced in Appendix A.

Another important advantage of exogenously varying the treatments is that we do not have to worry about (observable and unobservable) characteristics of the individuals confounding the effects across the different treatments. In observational data, positive and negative shocks are not randomly distributed and are usually systematically related with individual characteristics. For example, Bunn et al. (2017) find that households in their sample who experience positive shocks tend to be younger and hold more liquid assets (than those who experience negative shocks). Then, the extent to which a differential response to the positive and negative shocks is simply due to differences in preferences and characteristics of the two subsamples is not entirely clear. Likewise, the size of tax rebates usually tends to be a function of household income or size, which makes it hard to disentangle the size effect from underlying heterogeneity in characteristics and preferences of the different subsamples. Our approach bypasses these issues.

2.4 Summary Findings

Table 3 reports a summary of the MPCs implied by the responses to each treatment. We include this summary here without discussion in order to provide the reader with a concise overview of the findings. We will refer back to this table in the following sections as we discuss each treatment in turn. For each treatment, the table reports the total number of respondents (aggregated across multiple survey rounds for treatments that were conducted in more than one survey), the average MPC, the share of respondents with negative, zero and positive MPCs, and the average and median MPC conditional on being positive. When reporting average MPCs, we winsorize at the 2.5th and 97.5th percentiles.

Table 3: Summary of Findings

| | Count | Mean MPC | Share of Respondents with MPC | | | MPC \| MPC > 0 | |
			< 0	= 0	> 0	Mean	Median
GAIN							
$500	1638	0.08	0.06	0.75	0.19	0.54	0.50
$2500	540	0.11	0.06	0.66	0.27	0.43	0.40
$5000	1629	0.14	0.06	0.55	0.39	0.36	0.30
LOSS							
$500	1536	0.30	0.04	0.47	0.49	0.62	0.60
NEWS-GAIN							
$500 in 3 months	362	-0.00	0.08	0.86	0.06	0.41	0.50
$5000 in 3 months	594	0.04	0.05	0.81	0.14	0.31	0.30
NEWS-LOSS							
$500 in 3 months	975	0.31	0.02	0.46	0.51	0.61	0.55
$500 in 2 years	589	0.14	0.03	0.68	0.29	0.52	0.40
LOAN							
$5000	541	0.01	0.12	0.79	0.08	0.30	0.20

Note: Positive MPC corresponds to a negative change in spending for the loss treatments.

2.5 Order Effects

A potential concern with our survey design is that it may bias respondents toward stating that they would not adjust their spending so they can avoid the follow-up question of how much they would adjust their spending. If this were the case, we would expect that for a given treatment, we should see more non-zero spending responses if the treatment is shown to a respondent first rather than second. In Table 4, we test for this formally. We regress indicators for whether a respondent indicated they would change their spending (or increase their spending) in a given treatment on treatment-date fixed effects and an indicator for whether the respondent was exposed to this treatment first. We see that the estimated coefficients are close to zero and not statistically significant.

In addition, we can directly test for order effects by testing for the equality of distributions of MPCs depending on whether respondents saw a treatment first or second, for each treatment shown in Table 2. Out of the 14 treatments, there is 1 for which the null of equal distributions is rejected at $p < 0.05$, which is what one would expect based on random chance.[11]

---

[11]The treatment for which the distributions of MPCs between those respondents who see this

Table 4: Testing for Order Effects

	(1) MPC $\neq$ 0	(2) MPC $\neq$ 0
First Treatment Seen	0.010	0.003
	(0.007)	(0.007)
Treatment X Date FEs?	Yes	Yes
Avg. Y	0.37	0.31
Adj. R2	0.07	0.10
Obs.	9061	9061

Robust standard errors clustered by respondent in parentheses.

Significance: * < 0.1, ** < 0.05, *** < 0.01.

3   Baseline MPC Responses

In this section we consider responses to the GAIN treatment, in which we elicit the MPC out of a one time unexpected windfall of $500, $2500 or $5000. This treatment has been examined in the existing literature, both through surveys and choice data. In Section 4 we then compare the responses to the GAIN treatment with the four additional treatments that have been less well-studied.

3.1   Responses to Gains

The average reported quarterly MPC out of a $500 windfall is 8% (see Table 3). This small average MPC masks a large degree of heterogeneity across respondents. Three quarters of respondents say that they would not change their spending behavior at all, and hence have an MPC of zero, and an additional 6% report that they would reduce spending in response to the windfall. Only 19% of respondent say that they would increase their spending, but these households plan to spend a substantial fraction of the $500—the mean and median MPC conditional on a positive response are 54% and 50%, respectively. A more detailed breakdown of the distribution of MPCs is shown by the blue bars in Figure 1. For those respondents with a positive MPC the distribution is fairly evenly dispersed, although there is some evidence of bi-modality. Around 5% of households report that they would spend all of the $500 over the following quarter, while very few report spending more than 75% but less than 100% of the payment.

---

treatment first and those who see it second are significantly different is the $5,000 GAIN treatment in March 2016. The other treatment seen by these respondents was the $500 GAIN treatment. The fact that in May 2016 we do not see similar order effects for the $5,000 GAIN treatment when fielded together with the $2,500 GAIN treatment arguably makes it more likely that the difference in March happened by chance.
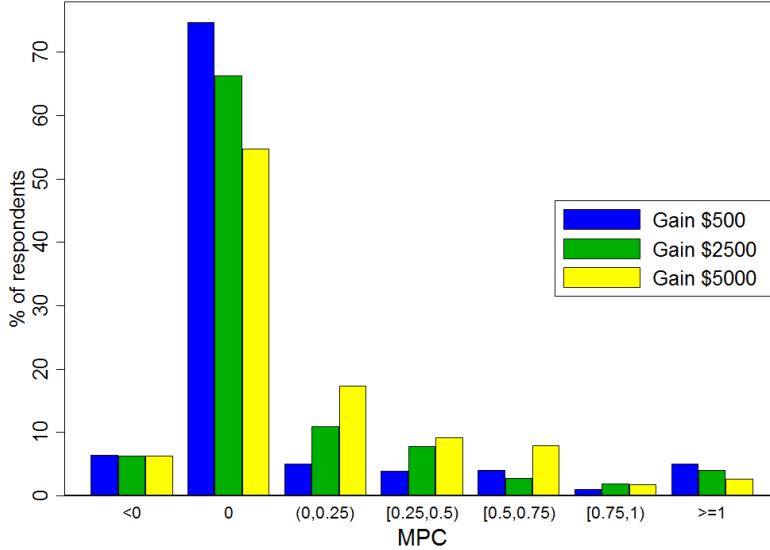
Figure 1: Histogram of MPCs for different Gain scenarios

Our average MPC is towards the lower end of the estimates found in the literature, for both hypothetical and actual gains of around this size. However, existing studies, like us, have found that majority of households respond little or not at all in response to an income windfall, but that a small sub-group of households (in our case, around one-fifth) spend a substantial fraction of the income windfall (see, for example, Bunn et al., 2017, and Christelis et al., 2017).

3.2   Effect of Windfall Size

As we increase the size of the windfall, a larger fraction of respondents say that they would increase their spending, but on average say they would spend a smaller fraction of the payment. For the $5,000 gain, 39% of respondents report a positive MPC, compared with 27% for the $2,500 gain and 19% for the $500 gain. Conditional on increasing spending, the median MPC is 30%, 40%, and 50% for the $5,000, $2,500 and $500 gains respectively. Overall the effect of the greater number of respondents with positive MPC dominates so that the average MPC increases slightly, from 8% to 11% to 14%, as the payment size increases.

This size effect in reported MPCs can be seen in Figure 1 by comparing the blue histogram ($500 windfall) with the green histogram ($2,500 windfall) and yellow histogram ($5,000) windfall. As the size of the windfall increases, the smaller mass of respondents with an MPC of zero is clearly evident, as is the larger mass of people with small, positive MPCs. Table A-2 in the Appendix shows that the difference in average MPCs (and the likelihood of reporting a positive MPC) across windfall

16

amounts are statistically significant and remain so once we condition on respondent fixed effects (for those that participate in more than one GAIN treatment) as well as date-by-treatment-order fixed effects.

In Section 5 we will show that this positive size effect, driven by the extensive margin, is difficult to reconcile with most standard models of optimal consumption behavior. Plausible explanations for why more people respond to larger windfalls include durables, salience, inattentiveness and cognitive costs of changing consumption plans. We discuss these possibilities in more detail in Section 5.

The size effect has not been studied much empirically, largely due to the fact that such variations are usually not observed in natural settings. We are aware of three other studies that investigate size effects, with little agreement. Bunn et al. (2017) find that for positive realized income shocks, MPCs increase in the size of the shock, in line with our results. Christelis et al. (2017) find similar overall MPC distributions for hypothetical positive shocks corresponding to one month or three months of income, though in line with our results, they find a smaller fraction of respondents that say they would not change their spending when the shock is larger. Fagereng et al. (2016), on the other hand, find that MPCs out of lottery winnings in Norway decline in the size of the amount won, which is consistent with our findings on the intensive margin (MPCs conditional on changing spending behavior), but not on the extensive margin (fraction of respondents who indicate they would change behavior).

3.3 Internal Consistency

Since the $500 gain and $5,000 gain treatments were each fielded in multiple survey waves, we can compare the distribution of responses across waves to examine the stability of responses over time. This comparison is shown in Figure 2a for the $500 windfall and Figure 2b for the $5,000 windfall. For both amounts, the distributions are very similar across the different survey waves.[12]

A subset of respondents were exposed to the same treatment in different survey waves (between two and six months apart). This makes it possible to examine how spending intentions in the same treatment vary within respondent over time. We examine within-person persistence in the $500 GAIN and $5000 GAIN treatments. Grouping respondents the into three MPC bins (≤ 0, $(0, 1)$, ≥ 1) reveals that 68% are

---

[12]For the two $500 gain MPC distributions (without binning values first as in the chart), a Kolmogorov-Smirnov (KS) test of equality of distributions gives a p-value of 0.95. Pairwise KS tests for the $5000 gain MPC distributions yield p-values of 0.71, 0.43, and 0.15 (where the smallest p-value is for the comparison between March 2016 and January 2017).
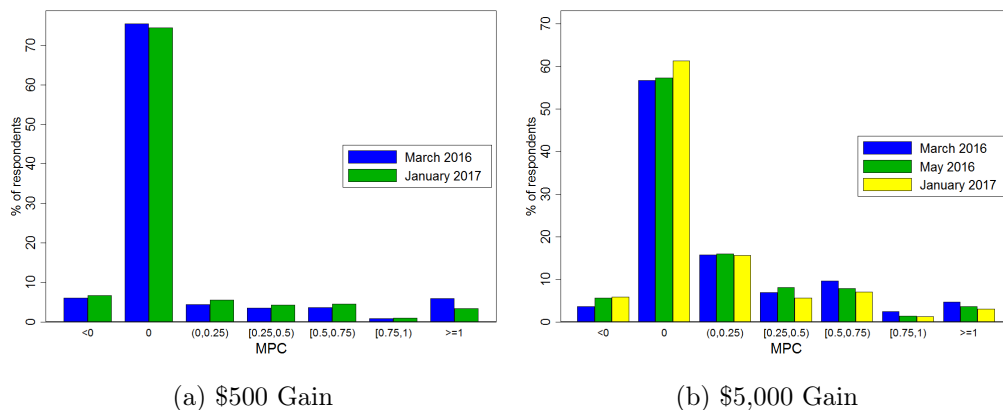
(a) $500 Gain   (b) $5,000 Gain

Figure 2: Consistency of responses across survey waves

in the same bin across different waves. This fraction is similar across both treatments. Since individual circumstances may change over time, we think that this degree of persistence is not a cause for concern.

3.4 Composition and Timing of Spending

We also asked respondents who indicated that they would adjust their spending in response to the treatment about how much of that additional spending would come from different spending categories.[13] The exact wording of seven possible spending categories can be found in Appendix A. We group the categories into non-durable spending ("traveling/vacation/eating out/other leisure activities"; "donations/gifts"; "general living expenses"), durables ("purchase of durables typically costing $1,000 or less"; "...typically costing more than $1,000"; "renovations or improvements to my home'; "pay for college/education/training for members of my household"), and "other". Table 5 shows average shares of spending responses for individuals with non-zero MPCs, in each of the three categories. For the three GAIN treatments, more of the adjustment comes from non-durable spending. However, as the size of the gain increases, the share that comes from durables increases (Christelis et al., 2017, find a similar result). This suggests that adjustment costs or other non-convexities may be important for understanding the positive size effect. We return to this possibility in Section 5.4.

In the May 2016 and January 2017 survey waves, we asked respondents who indicated they would increase their spending about the timing of spending within the following three month period. In May 2016, this was asked for all three treatments ($5000 GAIN, $2500 GAIN, $5000 LOAN) while in January 2017 it was asked for

---

[13]Spending composition was asked in all waves except March 2017.

18

Table 5: Average Spending Shares by Category

	Unweighted			MPC-Weighted			
	Nondur.	Dur.	Other	Nondur.	Dur.	Other	N
GAIN							
$500	0.71	0.24	0.05	0.65	0.28	0.06	410
$2500	0.63	0.33	0.04	0.55	0.39	0.06	179
$5000	0.60	0.36	0.04	0.53	0.41	0.06	830
LOSS							
$500	0.80	0.18	0.01	0.79	0.20	0.01	195
NEWS-GAIN							
$500 in 3 months	0.65	0.28	0.07	0.61	0.31	0.08	49
$5000 in 3 months	0.61	0.32	0.07	0.53	0.41	0.07	111
NEWS-LOSS							
$500 in 3 months	0.72	0.26	0.02	0.70	0.28	0.02	323
LOAN							
$5000	0.66	0.31	0.03	0.62	0.37	0.01	108

Nondurable and Durable definitions provided in text. N is the number of respondents with non-zero MPC for which the spending shares were elicited.
"Unweighted" means that all respondents with non-zero MPC are equal-weighted when calculating average shares. "MPC-weighted" means respondents are weighted by the absolute value of their MPC.

the $5000 GAIN treatment. The average shares of the spending increase happening in different time intervals are shown in Table A-3 in Appendix B (pooling the two $5000 gain waves). More than half of the increased spending (for those with MPC>0) occurs in the first month.

3.5 Individual Characteristics

Tables A-4 and A-5 in Appendix B show the average MPC and the share of respondents with MPC>0 in each treatment for different subgroups of respondents, defined by demographic characteristics (such as age or education), financial characteristics (such as income or liquid wealth), or preference parameters (discount rates). The definitions of the different groups are provided in Appendix A.2.

Focusing on the gain treatments for now, we see little systematic heterogeneity in spending responses. There is some evidence that those with lower discount factors (as measured from an incentivized choice experiment) tend to spend more out of their $5,000 windfall, although there are no significant differences for smaller gains. Similarly individuals with inconsistent time preferences (that is, those who exhibit a lower discount factor for choices that involve trade-offs today versus for choices involving trade-offs in the future) have a higher average MPC out of a $5,000 windfall;

for smaller gains, there are no differences. Also, those respondents who indicate in a qualitative question that they tend to spend rather than save do indeed have higher MPCs out of gains.[14]

4 Additional Treatments: News, Losses, Loans

The distribution of spending responses to a small unanticipated income windfall has been extensively studied in the existing literature. Relative to that literature, the results discussed in the previous section add only a few small new insights (such as the patterns of size asymmetry), while the main findings are consistent with prior studies. Rather, the role of the findings from the GAIN treatment is to act as a point of comparison for the more novel treatments that we discuss in this section. Unlike the GAIN treatment, there are few, if any, examples of behavioral studies that explore the NEWS, LOSS, NEWS-LOSS and LOAN treatments that we discuss in this section, necessitating a survey approach for these alternative treatments. In Section 5 we show that the responses to these additional treatments, and their comparison with the GAIN treatment provide a much richer set of findings for alternative theoretical models to confront.

4.1 News About Gains

In the NEWS treatment, we ask respondents how they would change their spending behavior over the next three months if they were to learn about a one-time windfall of either $500 or $5,000 that will be received in three months' time (see Appendix A for the exact wording of the survey instrument). In order to undertake within-person comparisons, these questions were asked only of respondents who also were exposed to the GAIN treatment of the same amount. A summary of our findings is that respondents do not react to news about a future windfall—even those respondents who say that they would react to the windfall if it were received immediately.

For the $500 news treatment, the average MPC is 0% and for the $5,000 treatment, the average MPC is only 4% (Table 3).[15] Moreover, 86% (81%) of respondents in the $500 ($5,000) treatment explicitly state that they would not change their

---

[14]The question is very similar to one in Parker (2017), who also finds that those who indicate that they are the "type of people who spend and enjoy today" have higher MPCs out of lump-sum payments.

[15]The zero average MPC for the $500 gain arises because a small fraction of respondents report small negative MPCs and small fraction report small positive MPCs, and these average to zero. See Table 3.

(a) Gain vs news of gain: $500

(b) Gain vs news of gain: $5,000

(c) Subset of respondents who react to im-
mediate gain: $500

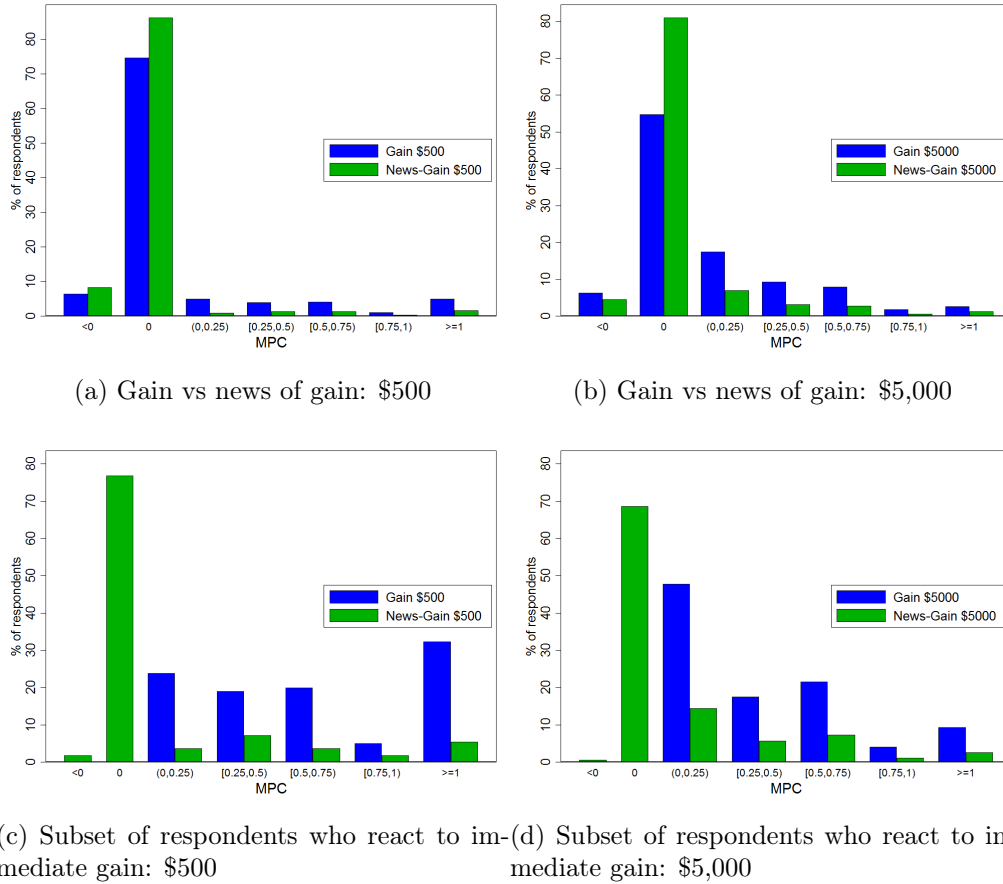(d) Subset of respondents who react to im-
mediate gain: $5,000

Figure 3: Spending response to news about future gains vs. response to gains today

spending over the quarter leading up to the payment in any way at all. Only 6%
(14%) of respondents say they would increase spending in response to the news, com-
pared with 19% (39%) for the immediate payment. The differences in these MPC
distributions between the GAIN treatment (blue histograms) and NEWS treatment
(green histograms) is displayed in Figures 3a and 3b. In both figures the additional
mass of respondents with a MPC of zero, and the much smaller fraction with a pos-
itive MPC, in the green histograms compared with the blue histograms, is clearly
evident.

We can find even stronger evidence for the absence of a spending response to the
news of a future windfall by examining the MPCs in the NEWS treatment for the
subset of households that say that they would indeed increase their spending in the
GAIN treatment. Focusing on the $5,000 windfall (where this subset is larger—195
out of 595 respondents), Figure 3d displays a histogram of the MPC for this subset.
The figure shows that the majority (nearly 70%) of respondents who would react to

an instantaneous windfall, would not react to a windfall in three months' time. And for those who do react, their spending response is typically much smaller than for the instantaneous windfall.

These findings are consistent with existing studies looking at the actual response to tax rebates, such as Johnson, Parker and Souleles (2006). The identification strategy in Johnson et al. exploits randomness in the timing of *when* households received their tax rebates, among a set of households who receive the rebate at some point during the observation period. As explained in Kaplan and Violante (2014), under reasonable assumptions about when households learned about their tax rebates, the estimated coefficients in the regression of consumption growth on the rebate received, should be interpreted as measuring the difference between the MPC out of a surprise tax rebate and an MPC out of an anticipated rebate, similarly to the difference between our GAIN and NEWS-GAIN treatments. The average coefficients of 20%-30% reported by Johnson et al. are thus indicative of a large difference between these two treatments. The analysis of consumption responses to different mortgage modification programs by Ganong and Noel (2017) is also consistent with the lack of a news effect on spending, although they study much larger amounts over a much longer time period than in our treatment.

4.2   Losses

We investigate the importance of sign asymmetry through a LOSS treatment, in which respondents were asked how they would change their spending in the event of an immediate unexpected loss of $500. We find that respondents are significantly more likely to react to a $500 loss than to a $500 gain, with an average MPC of 30% compared with with an average MPC of 8% for a $500 gain. The exact wording of the survey instrument can be found in Appendix A.

This sign asymmetry in the MPC is present on both the intensive and extensive margins. Whereas only 19% of respondents said they would increase spending under the GAIN treatment, 49% of respondents say they would decrease spending under the LOSS treatment.[16] Conditional on being positive, the median MPC is 60% for the loss compared with 50% for the gain. Figure 4a shows how the distribution of MPCs under the $500 LOSS treatment compares with the distribution under the GAIN treatment. The MPC distribution for the LOSS treatment (green histogram) is strongly suggestive of bi-modality—nearly 20% of respondents say that they would fully absorb the loss of $500 by reducing current spending. This sign asymmetry has

---

[16]In the LOSS treatment, a decrease in spending corresponds to a positive MPC.

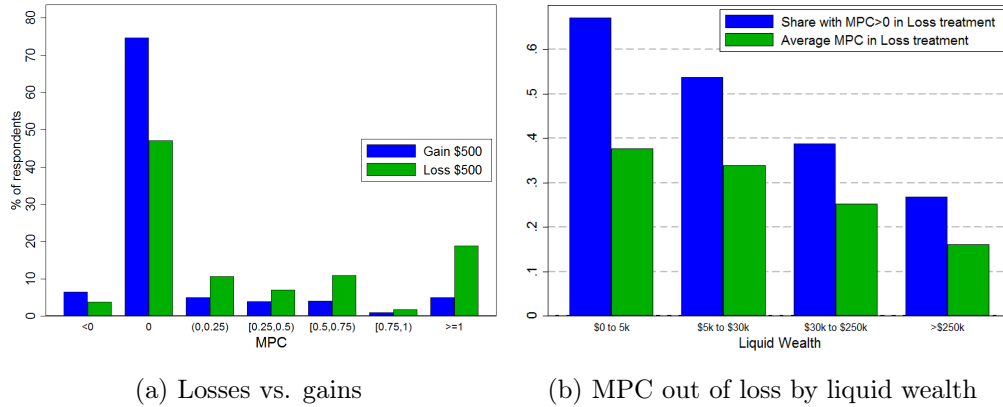(a) Losses vs. gains   (b) MPC out of loss by liquid wealth

Figure 4: Spending response to losses

also been documented by a relatively recent set of papers (Zafar et al., 2013; Sahm et al., 2015; Bracha and Cooper, 2014; Bunn et al., 2017; Christelis et al., 2017).

Because of the high average MPC and the bi-modality in responses, the LOSS treatment is a useful setting to compare the observable characteristics of individuals with high MPCs and low MPCs. As noted earlier, in Tables A-4 and A-5 in the Appendix we report how average MPCs for each treatment differ by many individual characteristics that one might *a priori* expect to be correlated with MPCs. Of these, a number of variables that proxy for being financially constrained (income, liquid assets, credit score) show a strong correlation with the MPC in the LOSS treatment. Figure 4b shows how the fraction of respondents with a positive MPC, and the average MPC, differ across liquid wealth categories (the cut-offs are chosen so that there are roughly the same number of respondents in each bin: 32%, 24%, 28% and 16% respectively). Both the fraction who say they would respond and the average MPC decline sharply with liquid wealth. Among respondents with less than $5,000 in liquid wealth, nearly 70% would reduce their spending, with an average MPC of 38%, whereas among respondents with more than $250,000 in liquid wealth, around 25% would reduce their spending, with an average MPC of around 17%.

The sign asymmetry of the average MPC also masks important heterogeneity in the extent and direction of sign asymmetry at the individual level. In Figure 5 we report the distribution of the *difference* in MPCs between the LOSS treatment and the GAIN treatment, separately for two groups of individuals – those who report a zero (or negative) MPC in the GAIN treatment (blue histogram), and those who report a positive MPC in the GAIN treatment. Of those respondents who do not react to the $500 windfall, Figure 5 shows that around half also do not react to the $500 loss. The remaining half do say that they would cut spending if faced with
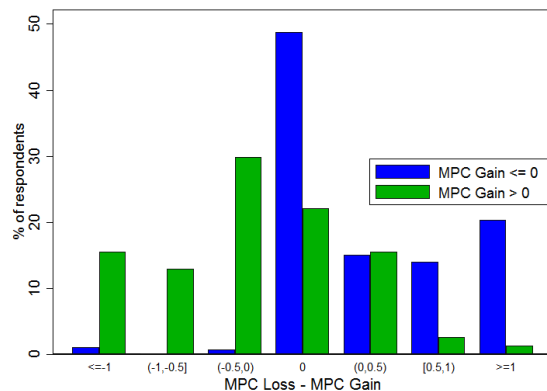
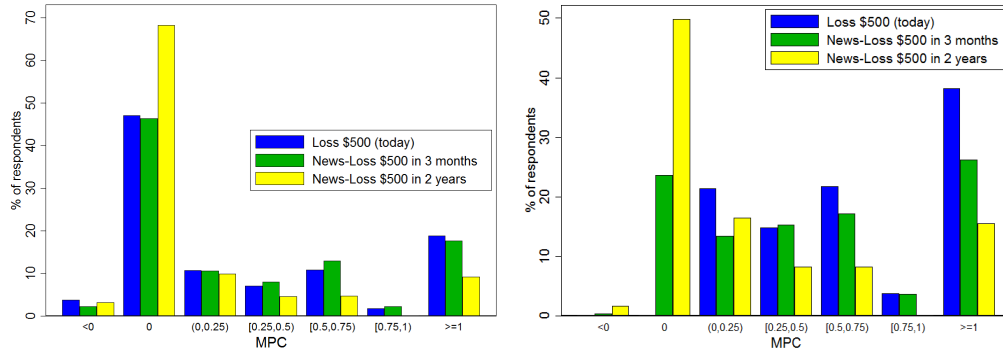Figure 5: Distribution of difference between MPCs out of losses and gains

a $500 loss, resulting in a larger average MPC for losses than gains. On the other hand, for those respondents who do react to the $500 windfall, Figure 5 shows that more than half of them react *less* to the loss than the gain; in fact, 39% would not cut their spending at all in response to the loss (not shown in the figure). However, since the latter group is much smaller than the former group (21% versus 79% of the sample that responds to these two treatments), the average MPC in the LOSS treatment is significantly larger than the MPC in the GAIN treatment.

What might explain this pattern of sign asymmetry? Recall that whereas the heterogeneity in MPCs in the LOSS treatment appears to be correlated with liquid wealth, there is much less evidence for such a correlation in the GAIN treatment. Since many of the households who react to the $500 gain do indeed have liquid wealth, it might not be surprising that they are able to smooth out the effect of the $500 loss. Why they then have a high MPC out of a windfall gain is an open question, and one we return to in Section 5.

4.3 News About Losses

The NEWS-LOSS treatment asks respondents how they would alter their spending behavior over the following three months if they were to immediately learn that they will suffer a $500 loss at a specified future date. Respondents are randomly assigned to two groups, one for which the loss is to occur in three months' time, the other for which the loss is to occur in two years' time. To allow within-person comparisons, all respondents exposed to the NEWS-LOSS treatment in March 2017 are also exposed to the LOSS treatment.

Figure 6a shows the distribution of spending responses to a $500 loss at different horizons. The response to the LOSS and 3 month NEWS-LOSS treatments are

24

(a) News about a \$500 loss at different hori-(b) News-loss effect for those respondents
zons   who have MPC>0 out of loss today

Figure 6: News-Loss

almost identical—the MPC distribution for a loss occurring in 3 months time is essentially the same as the distribution for an immediate loss. This lies in stark contrast to the comparison of the GAIN and NEWS-GAIN treatments in Section 4.1, where we found much smaller responses to news about a future windfall than to an immediate windfall. The implication of this finding is that even though many respondents say they would not smooth an immediate \$500 loss (suggesting a high MPC), they are willing to start preparing for a future loss of income by cutting spending today. Figure 6b shows the spending response to the NEWS-LOSS treatment for the subset of respondents who have a positive MPC in the LOSS treatment. The similarity of the \$500 NEWS-LOSS distribution with the MPC distribution for an immediate loss is evidence against the idea that high MPCs are driven by myopia, or even that that high MPCs are due to low liquid wealth which in turn is driven by myopia. Rather these findings suggest an element of rational, forward-looking behavior among individuals with high propensities to consume—they are willing to cut contemporaneous consumption in order to smooth out the effects of future anticipated losses.

Figure 6a also shows the distribution of MPCs out of an anticipated loss 2 years in the future. The MPC for a loss that far out is smaller than the MPC for a loss in three months' time, but even in this treatment almost one-third of people respond. Moreover, Figure 6b shows that around half of the households who say that they would cut consumption when faced with an immediate loss, also cut consumption in response to a loss in 2 years' time, albeit by a smaller amount. That so many high MPC households react to an anticipated loss 2 years in advance also implies that people are forward looking and that myopia alone cannot explain patterns of spending responses. In fact, 60% (32%) of respondents who cut spending in the
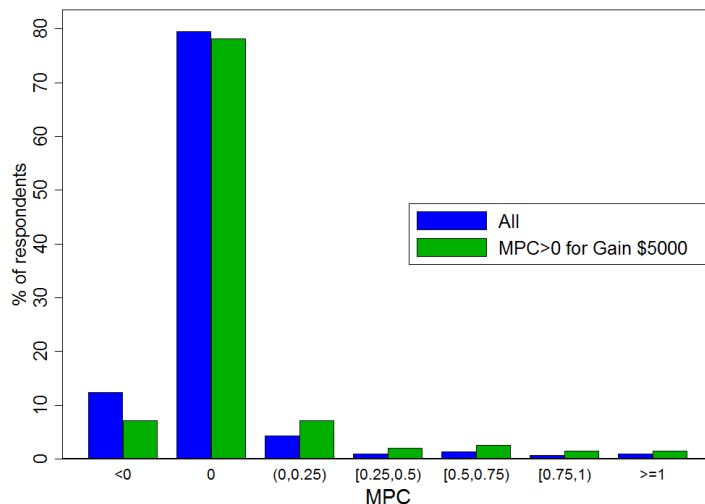
Figure 7: Response to Loan: all respondents and the subset who have MPC>0 for $5000 Gain

LOSS treatment report that they would cut spending by the same amount or more in the 3 months (2 year) NEWS-LOSS treatment.[17]

4.4 Loans

Binding borrowing constraints are often cited as a possible explanation for why some individuals have high MPCs out of small transitory gains and losses. To examine whether a loosening of borrowing constraints might have a large impact on MPCs, we included a LOAN treatment in the survey. In this treatment, respondents were asked how they would change their spending if offered an interest-free $5,000 loan to be repaid in one years's time. The distribution of MPCs for this treatment are displayed in Figure 7. The blue histogram shows that respondents react very little to the offer of a loan. Around 80% of respondents say that they would not change their spending at all, over 10% say that they would *reduce* spending (suggesting that they did not fully understand the concept of an interest free loan, or are worried about their ability to refrain from spending part of the loan in the time until repayment is due) and only 8% of respondents say that they would increase spending.

The green histogram in Figure 7 shows the MPC distribution for the LOAN treatment for the 36% of respondents who had a positive MPC in the $5,000 GAIN treatment (which was the other treatment these respondents were asked about). Among these respondents, the average MPC for the $5,000 windfall was 0.42, yet

---

[17]Among households that have an MPC≥ 1 in the LOSS treatment, 58% (34%) also have an MPC≥ 1 when the loss occurs 3 months (2 years) in the future.

when offered an interest-free loan for the same amount, nearly 80% of this group said that they would not change their spending at all, and the average MPC was only 0.04. That individuals who are known to have a large MPC out of an unanticipated windfall also have a zero MPC out of an interest free loan for the same amount is strong evidence that short-term borrowing constraints (shorter than the duration of the loan) are not a key reason for their high MPC. This does not rule out that longer-term borrowing constraints are important, but it is suggestive that high MPCs are associated with persistent low levels of disposable resources (i.e. longer than one year) rather than temporarily low levels.[18]

4.5 Takeaways

Before turning to theory, it is useful to summarize what we see as the broad qualitative takeaways from our survey treatments, with which we will confront existing theories. We will focus on six general findings:

1. Heterogeneity: Most people do not respond to gains, but there is a set of people who respond substantially.

2. Size effect: For bigger amounts, more people respond, i.e. a bigger average response driven by the extensive margin.

3. Sign asymmetry: More people respond, and by bigger amounts, to losses than to gains. Response to losses are correlated with liquid assets.

4. No response to news: People do not respond to news about future gains, even those with large responses to actual gains.

5. Response to news about losses: People do respond to news about future losses. Of those that respond to a loss, about half of them respond even when the loss is 2 years in the future.

6. No response to loans: People do not respond to loans, even those who do respond to gains.

5   Implications for Theory

What does economic theory predict for these treatments? In this section we view the MPCs for each treatment through the lens of alternative models. We

---

[18]In fact, in a trial at a large European retail bank where credit lines were randomly expanded, Aydin (2018) finds a MPC out of credit of about 0.20 after three months. One way to reconcile this with our findings is that the credit line expansions in his data had indefinite duration.

start by formally describing each treatment in a way that is amenable to theoretical and quantitative analysis. We then use the broad findings from Section 4.5 that summarize the survey results to evaluate a series of models. We start with the two simplest models of consumption behavior—the polar extremes of rule-of-thumb and permanent income behavior. We then consider richer models that incorporate precautionary savings, borrowing constraints, a life-cycle and assets with different degrees of liquidity. We end with a discussion of how various behavioral phenomena might further improve these models' ability to match the survey response data.

5.1 Definition of the Treatments

To organize ideas, it is useful to write an individual's budget constraint as

$$c_{it} + s_{it} = x_{it}$$
$$x_{i,t+1} = y_{i,t+1} + R(s_{it}) s_{it},$$

where x_{it} is cash on hand at the beginning of period t, c_{it} is the amount spent during period t and s_{it} is the amount saved in period t. We assume that interest is paid at the end of the period and that income $y_{i,t+1}$ is received at the beginning of the following period. Period $t+1$ cash-on-hand is thus given by period $t+1$ income plus savings from period t with accumulated interest. We allow the gross interest rate R to depend on the amount saved s_{it} to reflect the possibility that individuals face different interest rates on savings and borrowing. Consistent with the time horizon in our survey questions, we think of each time period as representing one quarter.

With this budget constraint, we can formally describe the five treatments. In the GAIN and LOSS treatments, the budget constraint unexpectedly becomes

$$c_{it} + s_{it} = x_{it} + \Delta$$
$$x_{i,t+1} = y_{i,t+1} + R(s_{it}) s_{it},$$

with $\Delta > 0$ in GAIN and $\Delta < 0$ in LOSS. In the NEWS-GAIN and NEWS-LOSS treatments, the budget constraint unexpectedly becomes

$$c_{it} + s_{it} = x_{it}$$
$$x_{i,t+1} = y_{i,t+1} + R(s_{it}) s_{it} + \Delta,$$

with $\Delta > 0$ in NEWS-GAIN and $\Delta < 0$ in NEWS-LOSS. In the LOAN treatment,

the budget constraint at time t unexpectedly becomes

$$c_{it} + s_{it} = x_{it} + \Delta$$

and the budget constraint at time $t + 4$ unexpectedly becomes

$$x_{i,t+4} = y_{i,t+4} + R\left(s_{i,t+3}\right) s_{it+3} - \Delta.$$

To ease notation, we label the treatments as follows: GAIN (G) , LOSS (L), NEWS-GAIN (NG), NEWS-LOSS (NL), and LOAN (LN). For each treatment $T \in \{G, L, NG, NL, LN\}$, we then define the MPC for an amount Δ as

$$MPC_{it}^T = \frac{c_{it}^\Delta - c_{it}}{\Delta},$$

where c_{it}^Δ is consumption under the treatment and c_{it} is consumption in the absence of the treatment.

5.2   Simple Benchmark Models

Before advancing to quantitatively plausible consumption-savings models, it is useful to clarify the predictions of three simple benchmark models of consumption behavior.

Rule-of-thumb consumers   Rule-of-thumb consumers consume all of their disposable income in every period. Hence they set $c_{it} = x_{it}$ and $s_{it} = 0$. This yields the following MPCs

$$MPC^G = MPC^L = MPC^{LN} = 1$$
$$MPC^{NG} = MPC^{NL} = 0$$

Rule-of-thumb behavior is thus not consistent with the substantial fraction of respondents who report not changing their consumption behavior in the GAIN and LOAN treatments, nor does it generate a size effect or sign asymmetry. Moreover, rule-of-thumb behavior is not consistent with the NEWS-LOSS responses, which suggest that people are at least somewhat forward looking, in contrast with the extreme myopia of rule–of-thumb consumers.

PIH consumers   Strict permanent income consumers have quadratic utility, face a fixed gross interest rate $R = 1 + r$ that is equal to the inverse of the discount rate, and face no constraints on borrowing other than a No-Ponzi condition that imposes

that they cannot die in debt. For such a household the optimal consumption policy is

$$c_{it} = \frac{R-1}{R}\left[x_{it} + \sum_{j=1}^{\infty} R^{-j} E_t y_{t+j}\right].$$

This gives the following MPCs

$$MPC^L = MPC^G = \frac{r}{1+r} \approx 0$$
$$MPC^{NL} = MPC^{NG} = \frac{r}{(1+r)^2} \approx 0$$
$$MPC^{LN} = \frac{r}{1+r} - \frac{r}{(1+r)^5} \approx 0,$$

where the approximations hold for low interest rates $r \approx 0$, which is true for the types of assets typically held for short-term consumption smoothing (e.g. cash, checking accounts). The strict PIH thus implies that households will have small responses in both the LOSS and GAIN treatments, and will not generate sign asymmetry nor a size effect. One of the starkest predictions of the PIH model is that the MPC out of gains should be essentially identical to the MPC out of news about future gains (with the only difference being the negligible effect of discounting). This prediction is not consistent with the finding from Section 4.1 that even among those respondents who reported substantial MPCs in the GAIN treatment, most reported low or zero MPCs in the NEWS-GAIN treatment.

Spender-saver model The spender-saver model is one in which the population is comprised of two groups of individuals – one group of rule-of-thumb consumers (the spenders), and another group of permanent income consumers (the savers) (Campbell and Mankiw, 1989). Assuming that a fraction α of the population are spenders and the remaining $1-\alpha$ are savers, and that the interest rate $r \approx 0$ then the MPCs for each of the five treatments are

$$MPC^G = MPC^L = MPC^{LN} = \alpha$$
$$MPC^{NG} = MPC^{NL} = 0.$$

The spender-saver model is thus able to generate large average MPCs, that are heterogeneous across individuals, as well as an average MPC out of news about future gains that is smaller than the MPC out of the actual gain. However, the model inherits from the rule-of-thumb and permanent income models the inability to generate meaningful sign asymmetry or size effects. Moreover the spender-saver model predicts no response to the NEWS-LOSS treatment, and predicts the same

size response to the LOAN treatment as to the GAIN and LOSS treatments. Both of these latter features are inconsistent with the survey responses.

5.3 Precautionary Savings Models

Modern workhorse models for understanding consumption behavior feature precautionary motives, due to either an occasionally binding borrowing constraint or the convexity of marginal utility. We start by examining an infinite-horizon precautionary-savings model with IID income risk, in which it is possible to develop several sharp theoretical predictions. We then examine the quantitative predictions of realistically calibrated life-cycle precautionary-savings models. With the exception of the size effect on the extensive margin of consumption responses, we find that a two-asset life-cycle model as in Kaplan and Violante (2014) can account for the observed MPC treatments. In Section 5.4, we then discuss how simple behavioral modifications to this model, in the spirit of salience or cognitive costs of changing consumption plans, can help reconcile the model with the evidence on the size effect.

5.3.1 Theoretical Predictions: Infinite-Horizon Model

Consider the following consumption-savings problem, expressed in recursive form

$$V\left(x\right) = \max_{c,s} u\left(c\right) + \beta E\left[V\left(x', y'\right)\right]$$

$$\text{subject to}$$

$$c + s = x$$
$$x' = Rs + y'$$
$$s \geq 0$$

The budget constraints are the same as previously described, except for the addition of the borrowing constraint $s \geq 0$. We assume that income y follows an IID process and that the individual has an increasing and convex utility function $u\left(c\right)$ with positive third derivative. The solution to this problem implies a value function $V\left(x\right)$ and an associated consumption policy function $c\left(x\right)$.

Under these conditions, both the value and policy functions are well-known to be strictly concave (see e.g. Carroll and Kimball, 1996; Carroll, 1997). Strict concavity of the consumption function implies that

$$MPC^L > MPC^G,$$

meaning that the consumption response to a windfall of a given size is bigger for a

loss than a gain, which is qualitatively consistent with the sign asymmetry reported in Section 4.2. But concavity also implies a negative size effect: the MPC in the GAIN treatment is smaller for larger windfalls. This is in contrast with the positive size effect reported in Section 3.2.

As the level of an individual's wealth increases, both the sign and size asymmetry in MPCs get weaker. In fact, one can show (see e.g. Benhabib et al., 2011) that for Constant Relative Risk Aversion (CRRA) utility functions, as $x \to \infty$, the consumption function approaches the linear function,

$$c\left(x\right) = \left[R\left(\beta R\right)^{-\frac{1}{\gamma}} - 1\right]x,$$

where γ is the coefficient of relative risk aversion. Thus for individuals with sufficiently high levels of wealth, there is neither any sign asymmetry nor a size effect. Moreover, when either βR is close to 1 or γ is close to 1, the MPC is approximately equal to $\beta^{-1} - 1$, as in the PIH. Thus both MPC^L and MPC^G are negligibly small.

In practice, this high wealth approximation of the consumption function tends to hold well except for individuals who are on, or very close to, the borrowing constraint $x = 0$.[19] For individuals who are borrowing constrained, the consumption function takes the simple form $c\left(y\right) = y$. It follows that $c\left(y - \Delta\right) = y - \Delta$ because if the borrowing constraint is binding at $x = y$ then it will also bind at $x = y - \Delta$. Hence borrowing constrained individuals respond to the LOSS treatment by cutting consumption by the amount of the loss, i.e. $MPC^L = 1$. Whether MPC^G is also equal to 1 depends on whether the borrowing constraint is also binding at the slightly higher level of wealth $x = y + \Delta$, which is less likely the larger is the size of Δ. This means that a borrowing constrained agent has $MPC^G = 1$ for small windfalls and $MPC^G < 1$ for larger windfalls.

The upshot of these results is that in terms of the simple GAIN and LOSS treatment, the simplest infinite-horizon precautionary savings model delivers MPCs that are qualitatively very similar to the even simpler spender-saver model. There is one group of individuals who have $MPC^G = MPC^L = 1$ and another who have $MPC^G = MPC^L \approx 0$. The key difference between the models is that in the precautionary savings model, the identity of the individuals in each group is endogenous and time-varying, whereas in the spender-saver model it is fixed exogenously.

---

[19]With sufficiently large transitory income risk, it is possible to generate consumption functions with substantial concavity (and hence high MPCs) even at moderate levels of liquid wealth. However, typically there are very few households at this part of the wealth domain in the ergodic distribution, since optimal savings decisions imply that households desire to save themselves away from the region where the consumption function is very concave.

We can also analyze the NEWS treatment separately for constrained and unconstrained agents. For unconstrained agents, the first-order condition for consumption is

$$u'(c) = \beta R E\left[V'\left(R(x-c) + \Delta + y'\right)\right],$$

where the Δ on the right-hand side reflects the future windfall at time $t+1$ that is learned about at time t. For small Δ, it is straightforward to show that the MPC out of news and the MPC out of an actual gains or losses are related by

$$MPC^{NG} = R^{-1}MPC^{G}$$
$$MPC^{NL} = R^{-1}MPC^{L}.$$

The MPC in the NEWS-GAIN treatment is less than the MPC in the GAIN treatment by a factor R and thus, for high and medium wealth individuals the MPCs in these two treatments are similar.[20] For constrained individuals, whose consumption function is $c(x) = x$, the MPC in the NEWS-GAIN treatment is 0, and thus for low wealth individuals the MPCs in the GAIN and NEWS-GAIN treatments can be very different.

The gap $MPC^{G} - MPC^{NG}$ is thus informative about whether an individual is hand-to-mouth (i.e. has the consumption function $c(x) = x$), since in this framework only hand-to-mouth agents exhibit a large difference between these two MPCs. The survey responses in Section 4.1 suggest that there are a substantial number of individuals for whom $MPC^{G} - MPC^{NG}$ is far from zero and hence may be hand-to-mouth. On the other hand, the informativeness of the corresponding gap for losses, $MPC^{L} - MPC^{NL}$ is more ambiguous. Constrained individuals have a large MPC out of the immediate loss but may or may not have a large MPC out of the news about a future loss, depending on the size of the multiplier on their borrowing constraint. We thus return to the NEWS-LOSS treatment in the context of the calibrated life-cycle model below.

In the context of the precautionary savings model with borrowing constraints, the LOAN treatment is informative about whether individuals are currently constrained and, if so, for how long they expect to be constrained. For unconstrained individuals with sufficient wealth that there is a low probability of still being constrained a year later, the loan has a negligible effect on their inter-temporal budget constraint, and

---

[20]Differentiating with respect to Δ and evaluating at $\Delta = 0$ defines the MPC out of news implicitly as $u''(c)\frac{\partial c}{\partial \Delta} = \beta R E\left[V''\left(R(x-c) + y'\right)\right]$. Differentiating with respect to x and evaluating at $\Delta = 0$ defines the MPC out of an immediate gain implicitly as $u''(c)\frac{\partial c}{\partial \Delta} = \beta R^2 E\left[V''\left(R(x-c) + y'\right)\right]$. Taking the ratio yields the result.

hence $MPC^{LN} \approx 0$.[21] For individuals who are currently borrowing constrained but expect to be unconstrained in the near future, the MPC in the LOAN treatment is similar to the MPC from the GAIN treatment, i.e. $MPC^{LN} \approx MPC^{G}$. However, for individuals who are constrained and expect to remain constrained for longer than the duration of the loan, the MPC in the LOAN treatment is approximately zero. In Section 4.4 we reported that almost no respondents indicated that they would increase spending when offered a one-year interest free loan, even those respondents who had large MPCs in the GAIN treatment and small MPCs in the NEWS-GAIN treatment. Viewed through the lens of the precautionary savings model, these responses are consistent with the presence of hand-to-mouth individuals who expect to remain in a low wealth state for a substantial period of time.

5.3.2 Quantitative Predictions: Life-cycle Model

The simple precautionary savings model described above lacks most of the features that existing work has shown are important for generating quantitatively realistic consumption behavior (see Kaplan and Violante (2010, 2014); henceforth KV10 and KV14). In this section we use a version of the life-cycle model of KV14 to assess the predictions of one-asset and two-asset life-cycle precautionary savings models for the five treatments. Appendix C contains a full description of both models and provides further details on the calibration, including plots of mean life-cycle profiles for income, consumption and wealth. We describe only the essential features in the main text.

In both models, the model period is one quarter and households live for $T^{work} + T^{ret} = 58 \times 4$ periods. During the first $T^{work} = 38 \times 4$ periods, households face a stochastic income process that consists of a deterministic age profile and AR(1) income shocks. During the last $T^{ret} = 20 \times 4$ periods households receive social security payments. Our calibration of the income process, and social security system are taken from KV14, and our calibration of the initial wealth distribution of newborn households is taken from KV10.[22] In the one-asset model households have CRRA preferences with a risk aversion coefficient of 2 and can save, but not borrow,

---

[21]The inter-temporal budget constraint is affected only to the extent that $\beta < 1$.

[22]Kaplan and Violante (2014) have permanent shocks, whereas we have AR(1) shocks with an autoregressive parameter of 0.97. We choose the size of the innovations to generate the same profile of variance of log labor income over the life-cycle. Our model does not include transitory shocks because there is no consensus on how to calibrate these in a quarterly model. We have experimented with transitory shocks of different sizes and provided that they are not too big, none of our substantive findings are affected. With very large, high-frequency, transitory shocks, households have an incentive to hold substantial liquid assets for precautionary reasons, and the model generates far fewer hand-to-mouth households than is suggested by US data. See Kaplan and Violante (2014) for further discussion.

Table 6: Average MPCs in Calibrated Life-Cycle Models

	Survey-based	(1) 1 asset model $r = 0\%$	(2) 1 asset model $r = 5\%$	(3) 1 asset model $r = 0\%$ Low wealth	(4) 2 asset model $r^a = 5\%$ $r^b = 0\%$
GAIN					
$500	0.08	0.05	0.06	0.53	0.28
$2500	0.11	0.04	0.05	0.29	0.12
$5000	0.14	0.03	0.04	0.22	0.06
LOSS					
$500	0.30	0.10	0.13	0.89	0.47
NEWS-GAIN					
$500 in 3 months	-0.00	0.03	0.03	0.02	0.04
$5000 in 3 months	0.04	0.02	0.03	0.01	-0.01
NEWS-LOSS					
$500 in 3 months	0.31	0.06	0.08	0.33	0.19
$500 in 2 years	0.14	0.02	0.02	0.01	0.02
LOAN					
$5000	0.01	0.00	0.01	0.16	0.01

in a single risk-free asset.

The average MPC for each of the five treatments in the one-asset model are displayed in the columns (1) and (2) of Table 6. Column (1) shows a version of the model with $r = 0\%$ p.a. (corresponding to the interest rate on liquid assets in the two-asset model) and column (2) shows a version of the model with $r = 5\%$ p.a., (corresponding to the interest rate on illiquid assets in the two-asset model). The discount factor β is calibrated in each version so that mean assets are equal to 3.5 times average annual labor income, which is approximately the level of mean household net worth in the United States (see KV14). Both versions are qualitatively consistent with many aspects of our survey results. They generate: the correct pattern of sign asymmetry; larger MPCs out of immediate gains than news about future gains; larger MPCs out of news about future losses than news about future gains; and a negligible MPC out of the loan. However, the models predict size effects in the wrong direction and, more importantly, the quantitative size of MPCs is substantially smaller than is suggested by the survey responses (e.g., the average MPC for a $500 loss is 0.1 in the model and 0.3 in our survey).

The key reason why the one-asset model produces such small average MPCs is that it generates very few hand-to-mouth households. Table 7 displays the average level of wealth and the fraction of hand-to-mouth households in each model. In

Table 7: Wealth Statistics in Calibrated Life-Cycle Models

	(1) 1 asset model $r = 0\%$	(2) 1 asset model $r = 5\%$	(3) 1 asset model $r = 0\%$ Low wealth	(4) 2 asset model $r^a = 5\%$ $r^b = 0\%$
Liquid Wealth				
Mean	3.50	3.50	0.05	0.05
Median	1.78	1.68	0.00	0.00
Illiquid Wealth				
Mean				3.45
Median				0.93
Fraction of Hand-to-Mouth Households (all)				
Poor HtM	0.03	0.04	0.80	0.15
Wealthy HtM				0.38
Fraction of Hand-to-Mouth Households (working-age)				
Poor HtM	0.03	0.04	0.52	0.13
Wealthy HtM				0.26

Notes: Wealth statistics are relative to average annual labor income.

the one-asset models calibrated to average net worth, less than 4% of households are hand-to-mouth, whereas Kaplan et al. (2014) report that up to one-third of households in United States can be characterized as hand-to-mouth.

A simple and popular way to modify the one-asset model so that it produces more low wealth households is to calibrate the discount factor β to match the average level of *liquid* wealth in the data, rather than the average level of *total* wealth. Column (3) of Table 6 displays the average MPCs in each of the five treatments when we set $r = 0\%$ and choose β so that average wealth is 0.05 times average labor income, which is the mean level of liquid wealth in the two-asset model described below. The MPCs in this liquid wealth calibration are much larger than in the net worth calibration and display similar patterns across treatments.

However, there are several drawbacks to the one-asset model with low liquid wealth. First, the MPCs are much larger than suggested by the survey responses (e.g., the average MPC for a $500 loss is 0.9 in the model and 0.3 in our survey), which is a consequence of the implausibly large fraction of hand-to-mouth households. Second, the very low discount factor required to generate such a low level of savings ($\beta = 0.83$ annual) causes the model to generate a substantial MPC in the LOAN treatment. Third, many would consider it unsatisfactory to calibrate a theory of consumption responses based on savings behavior by selectively ignoring the vast majority of observed savings.

To overcome these shortcomings, KV14 advocate the use of a two-asset precautionary savings model in which households hold both liquid and illiquid assets. The two-asset model that we consider is a simplified version of that in KV14. In the model, households can save in a liquid asset b that pays an interest rate $r^b = 0\%$ or an illiquid asset a that pays an annual interest rate $r^a = 5\%$. The asset is illiquid in the sense that in order to move funds between the liquid and illiquid accounts, households must pay a fixed cost κ. Illiquid assets also generate a flow of housing services, $d = \zeta a$ which enters utility according to the function

$$u\left(c, d\right) = \frac{\left(c^\phi d^{1-\phi}\right)^{1-\gamma} - 1}{1 - \gamma},$$

with the weight on housing services in consumption set to $1 - \phi = 0.15$. As with the one-asset model we choose the discount factor so that mean total assets are equal to 3.5 times average annual labor income.

Relative to the one-asset model, the two-asset model then has two additional parameters (κ, ζ). We choose these parameters to generate a fraction of poor hand-to-mouth households (households with zero net worth) and wealthy hand-to-mouth households (households with zero liquid assets but positive illiquid assets) that is broadly consistent with the evidence reported in Kaplan et al. (2014). Setting the fixed cost $\kappa = \$2,000$ and the service flow from housing $\zeta = 3\%$ generates 15% (13%) of all (working-age) households as poor hand-to-mouth and 38% (26%) of all (working-age) households as wealthy hand-to-mouth. Mean liquid asset holdings are $\$2,100$, which is equivalent to about 20% of mean quarterly labor income (see Table 7).

The average MPCs for each of the five treatments in the two-asset model are displayed in column (4) of Table 6. The two-asset is broadly consistent with many aspects of the survey response data. The two-asset model generates larger MPCs in the GAIN and LOSS treatments than the one-asset model, because of the larger fraction of hand-to-mouth households. Since many of these additional hand-to-mouth households are wealthy hand-to-mouth, they desire to smooth consumption over longer horizons and hence they respond to the 3 month NEWS-LOSS treatment, as in the data.[23] Unlike in the low-wealth calibration of the one-asset model, the large fraction of hand-to-mouth households does not arise because of extreme impatience,

---

[23]The model does not generate large MPCs out of anticipated losses two years out, while in the survey a surprisingly large number of respondents express a desire to cut their spending now rather than in the future. This finding may be related to other evidence showing that, in contrast to standard models of discounting, people tend to prefer incurring losses now rather than later (see e.g. Loewenstein and Prelec, 1991).

which explains why the two-asset model generates a low MPC out of the LOAN treatment whereas the low-wealth one asset model generates a substantial response in the LOAN treatment.

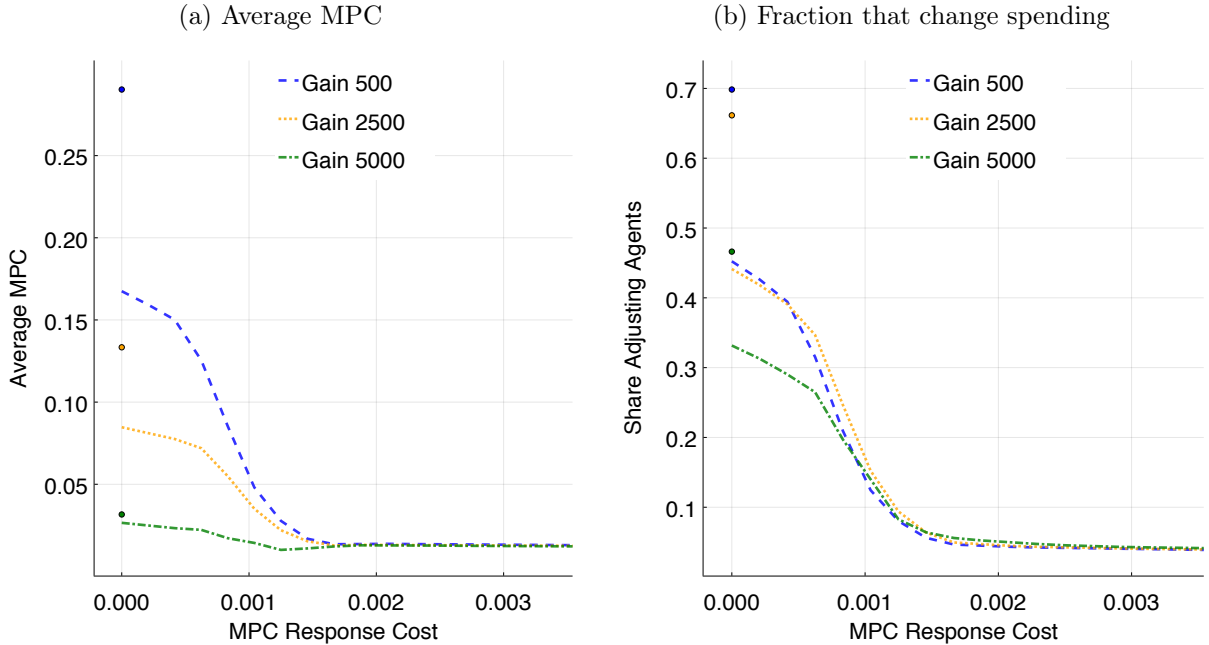## 5.4 Durables, Salience, and Behavioral Phenomena

An important aspect of the survey responses that neither the one-asset nor two-asset precautionary savings model is consistent with is the finding from Section 3.2 that MPCs increase with the size of the windfall. We found that the size effect is driven entirely by the extensive margin: as the size of the windfall was increased from $500 to $5000, the fraction of respondents saying that they would not change their spending decreased from 75% to 55%. This failure of the precautionary savings model is related to a more general failure of all of the models we have so far considered: whereas much of the action in the survey responses is on the extensive margin, the models only make predictions on the intensive margin because the MPC in these models is always non-zero.

One possible explanation for the extensive margin size effect is the presence of non-convex adjustment costs associated with the purchase of durable goods. In the presence of such goods, it is well known that consumption decisions have a threshold property, which can lead larger changes in individuals' budget sets to induce a larger fraction of individuals to change their consumption (see e.g. Berger and Vavra, 2015). This feature of durable spending is also emphasized by Broda and Parker (2014) as an interpretation of consumption responses to fiscal stimulus payments in 2008. Indeed, Table 5 showed that as the size of the windfall increases in the GAIN treatment, a larger fraction of expenditure is reported to be directed towards durable goods.

Another possible explanation for both the positive size effect and the large number of reported zero MPCs is salience or inertia. For many individuals, it may not be worth the effort to change their consumption plans when their budget constraint is changed by only a small amount. This may occur if there are cognitive costs involved with deviating from a status-quo or in deciding how to allocate a windfall across alternative spending possibilities. For small windfalls, the additional utility from optimally allocating the windfall between spending and savings might be too small to justify not saving the entire windfall.

A simple way to introduce salience and inertia into the precautionary savings model is to introduce a small "response cost". We model this cost by assuming that agents incur a fixed utility cost ψ in the period of the treatment if they choose a level

Figure 8: Two-asset precautionary savings model with response costs

(a) Average MPC

(b) Fraction that change spending

of consumption that differs from the level of consumption that was dictated by their optimal policy function in the absence of the treatment. The presence of this cost leads individuals with a small utility difference between responding to the treatment and not responding to the treatment to have an MPC of zero. Moreover, since this utility difference increases in the size of the windfall, this type of response cost has the potential to generate a positive size effect on the extensive margin.[24]

Figure 8a shows how the average MPC in the GAIN treatment in the two-asset model declines with the size of response cost ψ, for different size gains. Figure 8b shows the corresponding decline in the fraction of individuals with a positive MPC (extensive margin). These figures illustrate that addition of the response cost introduces a meaningful extensive margin of MPCs into the model and lowers MPCs to gains of all sizes. Importantly, the reduction in MPCs from the resource cost is smallest for the $5,000 gain and, for large enough responses costs, can overturn the negative size effect on the extensive margin. However, modeling salience in this way is not enough on its own to generate a positive extensive-margin size effect like that

---

[24]Formally, let $V^0(x)$ be the value for an individual with resources x in the absence of the treatment and let $c^0(x)$ be the associated consumption function. Let $V^1(x, \Delta)$ be the value for an individual with resources x in the GAIN treatment of size Δ and $\psi = 0$, and let $c^1(x, \Delta)$ be the associated consumption function. The value for an individual in the GAIN treatment who does *not* pay the response cost is given by $V^2(x, \Delta) = u(c_0(x)) + EV_1(R(x - c_0(x) + \Delta) + y)$. If $V^1(x, \Delta) - V^2(x, \Delta) > \psi$, the individual chooses to pay the response costs and sets $c = c^1(x, \Delta)$. If not, the individual does not pay the response cost, and sets $c = c^0(x)$.

observed in the survey responses. The reason is that the response cost lowers MPCs of all sizes and does not raise responses to larger windfalls.

A possible further modification that may be useful in this regard would be to combine the response cost with an additional behavioral phenomenon that raises MPCs to windfalls of all sizes. For example, the model of temptation and self-control in Gul and Pesendorfer (2001) and Kovacs and Moran (2017), or the model of quasi-hyperbolic discounting in Angeletos et al. (2001), typically generate higher MPCs to gains of all sizes than models that abstract from these features. Such a model combined with a response cost might be able to match the positive extensive margin size effect. We leave the quantitative assessment of these hybrid models to future work.

6 Conclusions

We have shown how carefully constructed survey questions about hypothetical treatments can be useful in distinguishing models of consumption behavior. Asking survey respondents how their spending would change in multiple different scenarios yielded six broad findings. First, there is a large amount of heterogeneity in consumption responses to small unexpected gains: most people do not react but there is a set of people who spend a substantial fraction of the windfall. Second, there is a positive extensive margin size effect: for bigger gains, more people respond. Third, we find evidence of sign asymmetry: spending responses to losses are larger and more widespread that spending responses to the same size gains. These responses to losses are correlated with holdings of liquid assets. Fourth, very few respondents increase their spending in response to news about future gains, even those respondents who indicate that they would increase spending in response to actual gains. Fifth, people generally do react to news about future losses. Sixth, almost no respondents indicated that they would increase spending when offered a one-year interest-free loan.

We then showed that many, but not all, of these aspects of consumption behavior are consistent with a two-asset life-cycle precautionary savings model. Several of the above findings are strongly suggestive of limited access to disposable resources being an important determinant of MPCs: higher MPCs out of losses than gains, the fact that MPCs out of losses are related to liquid wealth, and the very low MPCs out of news about future gains. The survey findings are also informative about the underlying reasons why many individuals act as if they have limited access to disposable resources: the substantial MPC out of news about future losses is evidence

against excessive impatience, myopia or extreme forms of present-bias; the lack of a spending response to the loan suggests that it is unlikely that short-term credit constraints play an important role.

The models that we considered all equate limited access to disposable resources with having little or no liquid wealth (i.e. being hand-to-mouth). A strict interpretation of these models thus suggests that liquid wealth, which is in principle easy to measure, should have strong explanatory power for cross-sectional heterogeneity in MPCs. However, with the exception of the LOSS treatment, we failed to find a significant correlation between MPCs and measures of hand-to-mouth behavior in terms of liquid wealth. This is true more generally in the literature on estimating MPCs: even in the more recent studies which have sufficient power to uncover a correlation between liquid wealth and MPCs, measured liquid wealth explains very little of the overall cross-sectional heterogeneity in MPCs (e.g. in terms of R-squared). One reason for this, we think, is that the relevant notion of *disposable* liquid wealth for determining MPCs, is actually very difficult to measure. It depends on the level, variation and timing of regular income; access to both informal and formal credit; and the level and variation in pre-committed and/or unforeseen expenses. Pre-Committed expenses are particularly difficult to measure and are likely to be correlated with income and liquid wealth in complicated ways that make make liquid wealth a very noisy proxy of disposable resources.

At least two important avenues for future work immediately arise from our findings. First, there is scope to develop theories of consumption which can explain why many households act as if they expect to be hand-to-mouth for long periods of time, and why many of these households have substantial measured liquid wealth. An obvious avenue is to introduce consumption commitments or other dimensions of heterogeneity in disposable resources into the model. Second, there is a need to systematically investigate the scope for behavioral phenomena to quantitatively account for the aspects of these survey responses where the two-asset precautionary savings model fails, most notably the increasing fraction of respondents who react to an unanticipated gain as the size of the gain increases.

Mental accounting, which is premised on the idea that resources are not fungible across different uses and that resources received in different circumstances may be used differently, is one such phenomenon. At this level of generality it is possible to reverse engineer a mental accounting story that is consistent with all of our findings. The downside of such an explanation, however, is that it is verging on tautological. The theory in this case is essentially a re-statement of the empirical findings. That being said, the basic ideas on which mental accounting are based, surely play a role

in understanding consumption behavior. Constructing a quantifiable and testable theory of mental accounting that can be confronted with the survey evidence alongside the other models that we have considered is an important open area for future research.

On the methodological front, we have demonstrated the usefulness of alternative treatments – gains, losses, news and loans – in distinguishing between different models of consumption behavior. Our hope is that future work will seek to identify similar experiments based on actual choice data that can complement our findings based on answers to hypothetical questions.

References

Ameriks, John, Joseph S. Briggs, Andrew Caplin, Matthew D. Shapiro, and Christopher Tonetti, "Long-Term-Care Utility and Late-in-Life Saving," Working Paper 20973, National Bureau of Economic Research February 2015.

_ , _ , _ , _ , and _ , "Late-in-Life Risks and the Under-Insurance Puzzle," Working Paper 22726, National Bureau of Economic Research October 2016.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The hyperbolic consumption model: Calibration, simulation, and empirical evaluation," *The Journal of Economic Perspectives*, 2001, *15* (3), 47–68.

Armantier, Olivier, Giorgio Topa, Wilbert van der Klaauw, and Basit Zafar, "An Overview of the Survey of Consumer Expectations," Staff Report No. 800, Federal Reserve Bank of New York 2016.

Aydin, Deniz, "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines," Working Paper, Washington University in St. Louis 2018.

Baker, Scott R., "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *Journal of Political Economy*, 2017, *forthcoming.*

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu, "The distribution of wealth and fiscal policy in economies with finitely lived agents," *Econometrica*, 2011, *79* (1), 123–157.

Berger, David and Joseph Vavra, "Consumption dynamics during recessions," *Econometrica*, 2015, *83* (1), 101–154.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption inequality and partial insurance," *The American Economic Review*, 2008, *98* (5), 1887–1921.

Bracha, Anat and Daniel Cooper, "Asymmetric responses to income changes: The payroll tax increase versus tax refund in 2013," *Economics Letters*, 2014, *124* (3), 534 – 538.

Broda, Christian and Jonathan A. Parker, "The economic stimulus payments of 2008 and the aggregate demand for consumption," *Journal of Monetary Economics*, 2014, *68*, S20–S36.

Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico, "The consumption response to positive and negative income shocks," *Journal of Monetary Economics*, 2017, *forthcoming.*

Campbell, Jeffrey R. and Zvi Hercowitz, "Liquidity Constraints of the Middle Class," Discussion Paper No. 2015-009, CentER, Tilburg University 2015.

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

Carroll, Christopher D., "Buffer-stock saving and the life cycle/permanent income hypothesis," *Quarterly Journal of Economics*, 1997, *112* (1), 1–55.

_ **and Miles S. Kimball**, "On the concavity of the consumption function," *Econometrica*, 1996, *64* (4), 981–992.

Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten van Rooij, "Asymmetric consumption effects of transitory income shocks," Working Paper No. 551, De Nederlandsche Bank 2017.

Fagereng, Andreas, Martin Blomhoff Holm, and Gisle James Natvik, "MPC heterogeneity and household balance sheets," 2016.

Ganong, Peter and Pascal Noel, "The Effect of Debt on Default and Consumption: Evidence from Housing Policy in the Great Recession," Working Paper, University of Chicago 2017.

Gelman, Michael, Kariv Kariv Shachar, Matthew Shapiro, Dan Silverman, and Steven Tadelis, "Harnessing Naturally Occuring Data to Measure the Response of Spending to Income," *Science*, 2014, *345* (6193), 212–215.

Graziani, Grant, Wilbert Van Der Klaauw, and Basit Zafar, "Workers' Spending Response to the 2011 Payroll Tax Cuts," *American Economic Journal: Economic Policy*, 2016, *8* (4), 124–159.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and self-control," *Econometrica*, 2001, *69* (6), 1403–1435.

Jappelli, Tullio and Luigi Pistaferri, "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 107–136.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, December 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L Violante, "How much consumption insurance beyond self-insurance?," *American Economic Journal: Macroeconomics*, 2010, *2* (4), 53–87.

_ **and Giovanni L. Violante**, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ , _ , **and Justin Weidner**, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014.

Kovacs, Agnes and Patrick Moran, "Temptation, Commitment, and Hand-to-Mouth Consumers," 2017.

Loewenstein, George and Drazen Prelec, "Negative Time Preference," *American Economic Review*, 1991, *81* (2), 347–352.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, December 2017, *107* (12), 3722–59.

Meier, Stephan and Charles Sprenger, "Present-Biased Preferences and Credit Card Borrowing," *American Economic Journal: Applied Economics*, January 2010, *2* (1), 193–210.

Parker, Jonathan A., "Why Don't Households Smooth Consumption? Evidence from a 25 million dollar experiment," *American Economic Journal: Macroeconomics*, 2017, *forthcoming.*

_ **and Nicholas S. Souleles**, "Revealed Preference vs. Reported Preference: Evidence from the propensity to spend tax rebates," Working Paper, MIT Sloan 2017.

_ , _ , **David S. Johnson, and Robert McClelland**, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–2553.

Sahm, Claudia, Matthew Shapiro, and Joel Slemrod, "Balance-Sheet Households and Fiscal Stimulus: Lessons from the Payroll Tax Cut and Its Expiration," Finance and Economics Discussion Series 2015-037, Federal Reserve Board 2015.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod, "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," *Tax Policy and the Economy*, 2010, *24* (1), 69–110.

Shapiro, Matthew D. and Joel Slemrod, "Consumer Response to Tax Rebates," *American Economic Review*, 2003, *93* (1), 381–396.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender*," *The Quarterly Journal of Economics*, 2018, *133* (1), 457–507.

Zafar, Basit, Max Livingston, and Wilbert van der Klaauw, "My Two (Per)cents: How Are American Workers Dealing with the Payroll Tax Hike?," *Liberty Street Economics*, 2013.

A   Further Details on Survey

In this Appendix, we provide the exact language and formatting (emphasis/underlining) used in the survey questions analyzed in the paper.

A.1   MPC Questions

For our MPC elicitation, we rely on questions across five different types of scenarios:

1. GAIN: windfall now ($500, $2,500, $5,000)

2. GAIN-NEWS: windfall in 3 months ($500, $5,000)

3. LOSS: occurring now ($500)

4. LOSS-NEWS: occurring later (500; either in 3 months or in 2 years)

5. LOAN ($5,000)

Parts in square brackets denote different variations depending on the gain amount or the timing (these were not shown to respondents). For scenario categories 2-5, we only provide the parts that differ from the GAIN scenarios. For the LOSS and LOSS-NEWS scenarios, there were minor changes to the wording across survey waves, as noted below. Example screenshots are provided at the end of this section.

<u>1. GAIN</u>

Now consider a hypothetical situation where you unexpectedly receive a one-time payment of **$500 [or 2,500, or 5,000]** today.

We would like to know whether this extra income would cause you to change your spending behavior in any way over the **next 3 months**.

Please select only one

• Over the next 3 months, I would **spend/donate** more than if I hadn't received the $500 [or 2,500, or 5,000]

• Over the next 3 months, I would **spend/donate** the same as if I hadn't received the $500 [or 2,500, or 5,000]

• Over the next 3 months, I would **spend/donate** less than if I hadn't received the $500 [or 2,500, or 5,000]

Please select only one

• Over the next 3 months, I would **pay off more debt (or borrow less)** than if I hadn't received the $500 [or 2,500, or 5,000]

• Over the next 3 months, I would **pay off the same amount of debt** as if I hadn't received the $500 [or 2,500, or 5,000]

• Over the next 3 months, I would **pay off less debt (or borrow more)** than if I hadn't received the $500 [or 2,500, or 5,000]

Please select only one

• Over the next 3 months, I would **save more** than if I hadn't received the $500 [or 2,500, or 5,000]

• Over the next 3 months, I would **save the same** as if I hadn't received the $500 [or 2,500, or 5,000]

• Over the next 3 months, I would **save less** than if I hadn't received the $500 [or 2,500, or 5,000]

[If selected "more" ["less"] spending]:
You indicated that you would increase [reduce] your spending/donations over the next 3 months following the receipt of the $500 [or $2,500, or $5,000] payment.

How much more [less] would you spend/donate than if you hadn't received the $500 [or $2,500, or $5,000]? [Enter value]

[Spending composition follow up (see screenshot below):]
And how much of these $(entered value) would you spend on each of the following: [And how much of this $(entered value) would come from a reduction in spending on each of the following:]
(Please note: The numbers need to add up to [entered value].)
Traveling / vacation / eating out / other leisure activities: $
Donation / gifts: $
General living expenses: $
Purchase of durables typically costing $1,000 or less (eg. electronics, sports equipment, clothing etc.): $
Purchase of durables typically costing more than $1,000 (such as a car, etc.): $
Renovations or improvements to my home: $
Pay for college / education / training for members of my household (including myself): $
Other (please specify:): $

[Spending timing follow up, asked for increases only (see screenshot below):]
You indicated that you would increase your spending/donations over the next 3 months by $(entered value) following the receipt of the $2,500 [or 5,000] payment. How would your spending change over time? I would increase my spending in...
(Please note: The numbers need to add up to [entered value].)
the next 2 weeks by $
the 2 weeks after that by $
the second month by $
the third month by $

48

2. GAIN-NEWS

Now consider a hypothetical situation where you learn that you will receive a guaranteed one-time payment of **$500 [or $5,000]** <u>exactly three months from today</u>.

We would like to know whether this extra income would cause you to change your spending behavior in any way over the **next 3 months (that is, before you receive the money)**.

Please select only one

• Over the next 3 months, I would **spend/donate more** than if I did not expect the guaranteed $500 [or 5,000] (in 3 months' time)

• Over the next 3 months, I would **spend/donate the same** as if I did not expect the guaranteed $500 [or 5,000]

• Over the next 3 months, I would **spend/donate less** than if I did not expect the guaranteed $500 [or 5,000] [questions on debt pay-down and saving]

[If selected "more" ["less"] spending]:
You indicated that you would increase [reduce] your spending/donations over the next 3 months after learning that you will receive a $500 [or $5,000] payment in 3 months.

How much more [less] would you spend/donate than if you did not expect to receive the $500 [or $5,000] in 3 months? [Enter value]
[Spending composition follow up]

3. LOSS

March 2016 wave:

Now consider a hypothetical situation in which you **unexpectedly** lose $500 today. Note that this is a one-time loss – it does not in any way affect your income going forward. You have simply found yourself suddenly to have $500 less than you previously had.

We would like to know whether this one-time $500 loss would cause you to change your spending behavior in any way over the **next 3 months.**
[The rest is identical to the March 2017 wave below]

March 2017 wave:

Now consider a hypothetical situation in which you **unexpectedly** lose $500 <u>**today**</u>. Note that this is a one-time loss – it does not in any way affect your income going forward. You have simply found yourself suddenly to have $500 less than you previously had.

We would like to know whether this one-time $500 loss would cause you to change your spending behavior in any way over the **next 3 months (that is, between now and June 2017).**

Please select only one

• Over the next 3 months, I would **spend/donate more** than if I hadn't lost $500

• Over the next 3 months, I would **spend/donate the same** as if I hadn't lost $500

• Over the next 3 months, I would **spend/donate less** than if I hadn't lost $500 [questions on debt pay down and saving]

[If selected "more" ["less"] spending]:
You indicated that you would increase [reduce] your spending/donations over the next 3 months following the one-time loss of $500.

How much more [less] would you spend/donate than if you hadn't lost $500? [Enter value]

[Spending composition follow up]

4. LOSS-NEWS

January 2017 wave:

Now consider a hypothetical situation where you learn today that you will lose **$500** <u>exactly</u> **three months** <u>from today</u>. Note that this is a one-time loss – it will not in any way affect your income otherwise.

We would like to know whether this one-time $500 loss would cause you to change your spending behavior in any way over the **next upcoming 3 months (that is, before you lose the $500).** [The rest is identical to the March 2017 wave below]

March 2017 wave:

Now consider a hypothetical situation where you learn today that you will lose **$500** <u>exactly</u> **three months** <u>from today (in June 2017)</u> [exactly **two years** from today (in March 2019)]. Note that this is a one-time loss – it will not in any way affect your income otherwise.

We would like to know whether this one-time $500 loss would cause you to change your spending behavior in any way over the **next upcoming 3 months (that is, between now and June 2017 – before you lose the $500).**

Please select only one

• Over the next 3 months, I would **spend/donate more** than if I did not expect the guaranteed $500 loss (in 3 months' time) [(in 2 years' time)]

• Over the next 3 months, I would **spend/donate the same** as if I did not expect the guaranteed $500 loss

• Over the next 3 months, I would **spend/donate less** than if I did not expect the guaranteed $500 loss [questions on debt pay down and saving]

[If selected "more" ["less"] spending]:
You indicated that you would increase [reduce] your spending/donations over the next 3 months after learning that you will lose $500 in 3 months [in 2 years].

How much more [less] would you spend/donate than if you did not expect the $500 loss in 3 months [2 years]? [Enter value]

[Spending composition follow up]

5. LOAN

Now consider a hypothetical situation where you are unexpectedly given **$5,000** today which you will have to return after 12 months. So this $5,000 is **a interest-free loan** that is to be repaid in a year's time.

We would like to know whether this $5,000 would cause you to change your spending behavior in any way over the **next 3 months.**

Please select only one
• Over the next 3 months, I would **spend/donate more** than if I hadn't received the $5,000 loan
• Over the next 3 months, I would **spend/donate the same** as if I hadn't received the $5,000 loan
• Over the next 3 months, I would **spend/donate less** than if I hadn't received the $5,000 loan [questions on debt pay down and saving]

[If selected "more" ["less"] spending]:
You indicated that you would increase [reduce] your spending/donations over the next 3 months following the receipt of the $5,000 loan.

How much more [less] would you spend/donate than if hadn't received the $5,000 loan? [Enter value]

[Spending composition follow up]

[Spending timing follow up]

nielsen

Now consider a hypothetical situation where you unexpectedly receive a one-time payment of **$500** today.

We would like to know whether this extra income would cause you to change your spending behavior in any way over the **next 3 months**.

Please select only one.

○ Over the next 3 months, I would **spend/donate** more than if I hadn't received the $500

○ Over the next 3 months, I would **spend/donate** the same as if I hadn't received the $500

○ Over the next 3 months, I would **spend/donate** less than if I hadn't received the $500

Please select only one.

○ Over the next 3 months, I would **pay off more debt (or borrow less)** than if I hadn't received the $500

○ Over the next 3 months, I would **pay off the same amount of debt** as if I hadn't received the $500

○ Over the next 3 months, I would **pay off less debt (or borrow more)** than if I hadn't received the $500

Please select only one.

○ Over the next 3 months, I would **save more** than if I hadn't received the $500

○ Over the next 3 months, I would **save the same** as if I hadn't received the $500

○ Over the next 3 months, I would **save less** than if I hadn't received the $500

| BACK | NEXT |

0%   25%   50%   75%   100%

(a) Screen 1

nielsen

You indicated that you would increase your spending/donations over the next 3 months following the receipt of the $500 payment.

How much more would you spend/donate than if you hadn't received the $500?

Please enter a number in the box below.

$☐

0%

(b) Screen 2, part 1

(c) Screen 2, Part 2 (spending composition follow-up – pops up after part 1 answered)

(d) Screen 3 (spending timing follow-up — only for some treatments, and only for increases)

A.2 Respondent Characteristics

Here, we provide the questions on which the cuts in Tables A-4 and A-5 below are based.

Age: the age of the respondent.

Bachelor's degree+: equals one if the respondent indicates that their highest completed education is a bachelor's degree, master's degree, doctoral degree, or professional degree.

Income: based on the following question:

> Which category represents the total combined pre-tax income of all members of your household (including you) during the past 12 months?
>
> Please include money from all jobs, net income from business, farm or rent, pensions, interest on savings or bonds, dividends, social security income, unemployment benefits, Food Stamps, workers compensation or disability benefits, child support, alimony, scholarships, fellowships, grants, inheritances and gifts, and any other money income received by members of your household who are 15 years of age or older.

[Respondents select from: Less than $10,000 / $10,000 to $19,999 / $20,000 to $29,999 / $30,000 to $39,999 / $40,000 to $49,999 / $50,000 to $59,999 / $60,000 to $74,999 /$75,000 to $99,999 / $100,000 to $149,999 / $150,000 to $199,999 / $200,000 or more]

Liquid Assets: based on the following two-part question:

> Which of the following do you or anyone in your family living with you have any money invested in?
>
> Please do NOT include any investments in retirement accounts (401k, 403b, 457, IRA, thrift savings plans etc.) or employer-sponsored pensions.

[Respondents select (possibly several) from: Checking accounts or cash / Savings accounts / Money market funds / CDs (Certificates of Deposit) / Government/Municipal Bonds or Treasury Bills / Stocks or bonds in publicly held corporations, stock or bond mutual funds, or investment trusts / None of the above]

   [If did not select "None of the above":]

If you added up all the money in these categories that you and your family members living with you have invested in, how much would it total?

[Respondents select from: Less than $500 / $500 to $999 / $1,000 to $1,999 / $2,000 to $4,999 / $5,000 to $9,999 / $10,000 to $19,999 / $20,000 to $29,999 / $30,000 to $49,999 / $50,000 to $99,999 / $100,000 to $249,999 / $250,000 to $499,999 / $500,000 to $749,999 / $750,000 to $999,999 / $1,000,000 or more]

Credit score: based on the question:

"What would you say is your credit score? [Respondents select from: Below 620 / 620-679 / 680-719 / 720-760 / above 760 / Don't know]

"Below Median" includes respondents in the first three categories (and we do not include in this cut the ones that answered "Don't know").

(Do not) have 2 months of funds: based on the response to the yes/no question (asked in February 2016) "In case of an unexpected decline in income or increase in expenses, do you [or your spouse/partner] have at least two months of covered expenses available in cash, bank accounts, or easily accessible funds?"

Tend to spend now / in future: based on the response to the question (asked in February 2016) "In general, are you the sort of person who would rather spend your money and enjoy it today or save more for the future? Where would you place yourself on the scale below?" [Respondents are asked to select on a scale from 1 to 5, where 1 is "Spend now" and 5 is "Save for the future". We classify as "Tend to spend now" those that select values 1-3.]

Time discounting questions: Respondents that were in the survey in February 2016 were asked the following:

In this final part of the survey, 5 respondents will be randomly chosen to win a significant amount of money, as explained below.

In completing this final part of the survey, you have to decide whether you want a smaller amount of money sooner, or a larger amount of money later. Specifically, in each of the following rows, please choose the option you prefer. For example, in row 1, you have a choice between $150 today versus $160 in a month from today.

If you are selected as a winner, one of these rows will be randomly picked and you will receive the money on the date indicated. (Note that "today" means within one business day.)

You will know immediately after the end of this survey whether you were chosen as a winner or not. If you are chosen as one of the winners, you would also be informed about the row for which you will be paid.x

They were then asked to make choices across three blocks of question (following Meier and Sprenger, 2010):

1. Money TODAY or IN A MONTH: they were presented with 6 rows, where in each row they chose between "$X guaranteed today" or "$160 guaranteed in a month", where X took the following values: 150, 140, 130, 120, 100, 80.

2. Money TODAY or IN 6 MONTHS: they were presented with 7 rows, where in each row they chose between "$X guaranteed today" or "$160 guaranteed in 6 months", where X took the following values: 150, 140, 130, 120, 100, 80, 60.

3. Money IN 6 MONTHS or IN 7 MONTHS: they were presented with 6 rows, where in each row they chose between "$X guaranteed in 6 months" or "$160 guaranteed in 7 months", where X took the following values: 150, 140, 130, 120, 100, 80.

We define a respondent to have a "Low 1-month Discount Factor" if in the first block, they prefer any of the smaller amounts today to the $160 in a month.

We define a respondent to exhibit "Inconsistent Time Discounting" if their implied discount factor in the first block is lower than in the third block (e.g. they prefer $150 today over $160 in a month, but prefer $160 in 7 months over $150 in 6 months).

B   Further Results

Table A-1: Summary of findings, weighting respondents to make average characteristics representative of US population

| | | Mean | Share of Respondents with MPC | | | MPC \| MPC > 0 | |
	Count	MPC	< 0	= 0	> 0	Mean	Median
GAIN							
$500	1638	0.07	0.08	0.74	0.18	0.53	0.50
$2500	540	0.09	0.09	0.69	0.22	0.43	0.40
$5000	1629	0.12	0.08	0.56	0.36	0.36	0.30
LOSS							
$500	1536	0.32	0.04	0.42	0.53	0.61	0.60
NEWS-GAIN							
$500 in 3 months	362	-0.02	0.11	0.86	0.04	0.43	0.50
$5000 in 3 months	594	0.04	0.05	0.82	0.13	0.30	0.30
NEWS-LOSS							
$500 in 3 months	975	0.31	0.03	0.43	0.54	0.58	0.50
$500 in 2 years	589	0.15	0.04	0.65	0.31	0.51	0.40
LOAN							
$5000	541	0.01	0.17	0.75	0.08	0.34	0.40

Note: Positive MPC corresponds to a negative change in spending for the loss treatments.

Table A-2: Size Effect for Gains: Within-Respondent Analysis

	(1) MPC	(2) MPC	(3) I(MPC > 0)	(4) I(MPC > 0)
Omitted category: Gain = $500				
Gain = $2500	0.033***	0.036**	0.088***	0.087***
	(0.012)	(0.017)	(0.020)	(0.027)
Gain = $5000	0.061***	0.064***	0.176***	0.193***
	(0.008)	(0.014)	(0.013)	(0.023)
Respondent FEs?	No	Yes	No	Yes
Date X Order FEs?	No	Yes	No	Yes
Avg. Y	0.11	0.12	0.28	0.30
Adj. R2	0.01	0.34	0.03	0.42
Obs.	4259	3094	4259	3094

Table reports regressions of a respondent's MPC (columns 1-2) or indicator for having a positive MPC (columns 3-4) on the size of the windfall in the three GAIN scenarios (for $500, $2,500, and $5,000). In columns (2) and (4), sample sizes are smaller because only respondents that participated in more than one GAIN scenario are included. Robust standard errors clustered by respondent in parentheses. Significance: * < 0.1, ** < 0.05, *** < 0.01.

Table A-3: Distribution of Timing

$5000 Gain	Mean	Median	25th Percentile	75th Percentile
Next Two Weeks	35.34	25.00	8.33	50.00
Next Two-Four Weeks	19.78	20.00	0.00	30.00
Second Month From Now	23.87	20.00	0.00	33.33
Third Month From Now	21.01	10.00	0.00	30.00
$2500 Gain	**Mean**	**Median**	**25th Percentile**	**75th Percentile**
Next Two Weeks	44.41	35.28	16.68	80.00
Next Two-Four Weeks	21.39	20.00	0.00	30.00
Second Month From Now	18.66	10.00	0.00	26.79
Third Month From Now	15.54	0.00	0.00	25.00
$5000 Loan	**Mean**	**Median**	**25th Percentile**	**75th Percentile**
Next Two Weeks	30.03	25.00	20.00	37.14
Next Two-Four Weeks	23.87	20.00	13.39	25.54
Second Month From Now	20.70	20.00	9.00	25.90
Third Month From Now	25.41	25.00	0.00	36.67

Table A-4: MPC Statistics (Average MPC)

	Gain			Loss	News-Gain		News-Loss		Loan
	$500 Gain	$2500 Gain	$5000 Gain	$500 Loss	$500 Gain in 3 months	$5000 Gain in 3 months	$500 Loss in 3 months	$500 Loss in 2 years	$5000 Loan
Full Sample	0.08 [1679]	0.11 [540]	0.14 [2040]	0.30 [1536]	-0.00 [362]	0.04 [594]	0.31 [1180]	0.14 [589]	0.01 [541]
Age ≤ 50	0.08 [829]	0.10 [265]	0.13 [977]	0.33 [741]	-0.00 [179]	0.05 [266]	0.34 [593]	0.17 [279]	0.01 [263]
Age > 50	0.08 [850]	0.13 [275]	0.15 [1063]	0.27 [795]	-0.01 [183]	0.04 [328]	0.28 [587]	0.12 [310]	0.02 [278]
P-value	0.72	0.14	0.06	0.00	0.76	0.25	0.02	0.03	0.12
< Bachelor's Degree	0.07 [717]	0.09 [224]	0.12 [885]	0.34 [704]	-0.02 [148]	0.03 [278]	0.32 [509]	0.17 [271]	0.01 [238]
Bachelor's Degree+	0.08 [961]	0.13 [316]	0.15 [1155]	0.27 [831]	0.01 [214]	0.05 [316]	0.30 [669]	0.12 [318]	0.02 [303]
P-value	0.39	0.04	0.01	0.00	0.09	0.08	0.31	0.07	0.42
Income ≤ $75k	0.08 [946]	0.12 [294]	0.14 [1137]	0.34 [850]	-0.01 [213]	0.04 [334]	0.34 [655]	0.16 [333]	0.01 [320]
Income > $75k	0.08 [712]	0.10 [241]	0.15 [880]	0.25 [669]	0.00 [147]	0.04 [257]	0.26 [510]	0.13 [250]	0.01 [213]
P-value	0.79	0.27	0.42	0.00	0.48	0.74	0.00	0.26	0.79
Liquid Assets ≤ $20k	0.08 [690]	0.12 [203]	0.14 [798]	0.37 [605]	0.01 [162]	0.05 [232]	0.37 [462]	0.16 [229]	0.02 [212]
Liquid Assets > $20k	0.09 [679]	0.11 [215]	0.14 [823]	0.23 [621]	-0.00 [142]	0.04 [244]	0.26 [497]	0.13 [229]	0.01 [217]
P-value	0.72	0.55	0.94	0.00	0.55	0.29	0.00	0.22	0.43
Below Median Credit Score	0.07 [691]	0.12 [218]	0.14 [834]	0.35 [611]	-0.01 [169]	0.05 [250]	0.36 [457]	0.17 [229]	0.01 [219]
Above Median Credit Score	0.10 [589]	0.12 [172]	0.15 [683]	0.24 [542]	0.01 [112]	0.04 [193]	0.27 [444]	0.11 [203]	0.01 [188]
P-value	0.13	0.82	0.51	0.00	0.48	0.53	0.00	0.06	0.58
Do not have 2 months of funds	0.08 [250]	0.13 [103]	0.14 [278]	0.47 [83]	0.00 [84]				0.01 [91]
Have 2 months of funds	0.10 [655]	0.12 [237]	0.15 [718]	0.26 [213]	-0.00 [224]				0.01 [261]
P-value	0.48	0.72	0.33	0.00	0.70				0.90
Tend to Spend Now	0.11 [582]	0.14 [214]	0.17 [615]	0.33 [195]	-0.00 [197]				0.01 [210]
Tend to Save in Future	0.07 [322]	0.08 [125]	0.12 [380]	0.29 [99]	0.00 [112]				0.02 [142]
P-value	0.07	0.04	0.01	0.37	0.62				0.59
Low 1-month Discount Factor	0.09 [362]	0.11 [214]	0.18 [526]	0.34 [124]	-0.01 [125]				0.02 [198]
High 1-month Discount Factor	0.09 [522]	0.13 [198]	0.13 [609]	0.30 [165]	0.01 [176]				0.00 [228]
P-value	0.86	0.58	0.00	0.47	0.21				0.08
Consistent Time Discounting	0.09 [714]	0.12 [265]	0.14 [790]	0.31 [236]	0.00 [242]				0.01 [287]
Inconsistent Time Discounting	0.09 [168]	0.14 [66]	0.20 [181]	0.34 [53]	-0.00 [58]				0.04 [57]
P-value	0.81	0.61	0.00	0.66	0.86				0.04

Note: Positive MPC corresponds to a negative change in spending for the loss treatments. Sample size reported in square brackets.

Table A-5: MPC Statistics (Proportion with MPC>0, in Percent)

	Gain			Loss	News-Gain		News-Loss		Loan
	$500 Gain	$2500 Gain	$5000 Gain	$500 Loss	$500 Gain in 3 months	$5000 Gain in 3 months	$500 Loss in 3 months	$500 Loss in 2 years	$5000 Loan
Full Sample	18.64 [1679]	27.41 [540]	36.23 [2040]	49.15 [1536]	5.52 [362]	14.48 [594]	48.14 [1180]	28.52 [589]	8.13 [541]
Age ≤ 50	19.30 [829]	27.17 [265]	36.13 [977]	53.58 [741]	7.26 [179]	15.79 [266]	52.45 [593]	29.75 [279]	6.84 [263]
Age > 50	18.00 [850]	27.64 [275]	36.31 [1063]	45.03 [795]	3.83 [183]	13.41 [328]	43.78 [587]	27.42 [310]	9.35 [278]
P-value	0.49	0.90	0.93	0.00	0.15	0.41	0.00	0.53	0.29
< Bachelor's Degree	18.13 [717]	23.21 [224]	30.73 [885]	56.11 [704]	4.05 [148]	12.95 [278]	52.65 [509]	36.16 [271]	7.56 [238]
Bachelor's Degree+	19.04 [961]	30.38 [316]	40.43 [1155]	43.20 [831]	6.54 [214]	15.82 [316]	44.54 [669]	22.01 [318]	8.58 [303]
P-value	0.64	0.07	0.00	0.00	0.31	0.32	0.01	0.00	0.67
Income ≤ $75k	19.98 [946]	29.25 [294]	36.15 [1137]	59.06 [850]	5.63 [213]	15.87 [334]	56.34 [655]	35.44 [333]	8.75 [320]
Income > $75k	17.13 [712]	25.31 [241]	36.36 [880]	36.77 [669]	5.44 [147]	12.84 [257]	38.04 [510]	19.60 [250]	7.04 [213]
P-value	0.14	0.31	0.92	0.00	0.94	0.30	0.00	0.00	0.48
Liquid Assets ≤ $20k	21.74 [690]	29.06 [203]	37.84 [798]	63.80 [605]	8.64 [162]	16.81 [232]	60.82 [462]	35.81 [229]	10.85 [212]
Liquid Assets > $20k	15.76 [679]	26.05 [215]	35.24 [823]	35.27 [621]	2.82 [142]	13.52 [244]	37.63 [497]	22.71 [229]	5.53 [217]
P-value	0.00	0.49	0.28	0.00	0.03	0.32	0.00	0.00	0.04
Below Median Credit Score	19.54 [691]	31.19 [218]	37.89 [834]	58.76 [611]	6.51 [169]	15.60 [250]	58.21 [457]	34.50 [229]	7.31 [219]
Above Median Credit Score	17.83 [589]	26.16 [172]	35.14 [683]	38.38 [542]	5.36 [112]	13.99 [193]	38.96 [444]	21.18 [203]	7.45 [188]
P-value	0.44	0.28	0.27	0.00	0.69	0.64	0.00	0.00	0.96
Do not have 2 months of funds	20.00 [250]	33.01 [103]	39.21 [278]	74.70 [83]	9.52 [84]				10.99 [91]
Have 2 months of funds	18.17 [655]	25.74 [237]	37.05 [718]	40.85 [213]	4.46 [224]				6.90 [261]
P-value	0.53	0.17	0.53	0.00	0.09				0.22
Tend to Spend Now	20.27 [582]	30.84 [214]	41.14 [615]	51.28 [195]	5.58 [197]				8.10 [210]
Tend to Save in Future	15.53 [322]	23.20 [125]	32.11 [380]	49.49 [99]	6.25 [112]				7.75 [142]
P-value	0.08	0.13	0.00	0.77	0.81				0.91
Low 1-month Discount Factor	18.51 [362]	25.23 [214]	41.63 [526]	53.23 [124]	4.00 [125]				12.12 [198]
High 1-month Discount Factor	18.39 [522]	30.81 [198]	34.48 [609]	48.48 [165]	7.39 [176]				4.39 [228]
P-value	0.96	0.21	0.01	0.43	0.22				0.00
Consistent Time Discounting	18.07 [714]	27.92 [265]	36.08 [790]	49.15 [236]	6.61 [242]				6.27 [287]
Inconsistent Time Discounting	20.24 [168]	28.79 [66]	44.20 [181]	56.60 [53]	3.45 [58]				15.79 [57]
P-value	0.51	0.89	0.04	0.33	0.36				0.01

Note: Positive MPC corresponds to a negative change in spending for the loss treatments.
Sample size reported in square brackets.

C Further Details on Model

This Appendix provides further details on the calibrated life-cycle models from Section 5.3. We first describe the two-asset model and then explain the parameter configuration such that it collapses to the one-asset model. The state variables for a working-age household of age $t \in \{1, \ldots, T^{work}\}$ are its liquid assets l, its illiquid assets a its labor productivity z. The value function in period t is given by

$$V_t(l, a, z) = \max \left\{ V_t^A(l, a, z), V_t^N(l, a, z) \right\}$$

where $V_t^A(l, a, z)$ is the value function if the household pays the adjustment cost in period t and $V_t^N(l, a, z)$ is the value function if the household does not pay the adjustment cost in period t.

The value function $V_t^A(l, a, z)$ is given by

$$V_t^A(l, a, z) = \max_{c, h, l', a'} \frac{\left(c^\phi d^{1-\phi}\right)^{1-\gamma} - 1}{1 - \gamma} + \beta \mathbb{E}_{z'} \left[V_{t+1}(l', a', z') | z \right]$$

$$\text{subject to}$$

$$c + h + R_l^{-1} l' + R_a^{-1} a' = l + a + y_t(z) - \tau(y_t(z)) - \kappa$$

$$l', a' \geq 0$$

$$h \geq -\zeta a$$

where total consumption of housing services $d = h + \zeta a$, where h is the amount of purchased housing services (at price 1) and ζa is imputed housing services from illiquid assets. The earnings function is given by $y_t(z) = e^{\chi_t + z}$ and the tax function is given by $\tau(y)$. Productivity risk z follows the AR(1) process

$$z' = \rho z + \eta_t$$

$$\eta_t \sim N(0, \sigma_{\eta,t})$$

$$z_0 \sim N(0, \sigma_{z,0}).$$

The problem for a retired household of age $t \in \{T^{work} + 1, \ldots, T\}$ is identical, except that the household faces no productivity risk, receives no labor income and instead receives a social security benefit $p(z_{T^{work}})$, which is a function of productivity in the last period of working life and is a function that returns social security benefits based on earnings in the last working age period.

The value function $V_t^N(l, a, z)$ is given by

$$V_t^N(l, a, z) = \max_{c,h,l'} \frac{\left(c^\phi d^{1-\phi}\right)^{1-\gamma} - 1}{1 - \gamma} + \beta \mathbb{E}_{z'}\left[V_{t+1}(l', a', z')|z\right]$$

subject to

$$c + h + R_l^{-1}l' = l + y_t(z) - \tau(y_t(z))$$
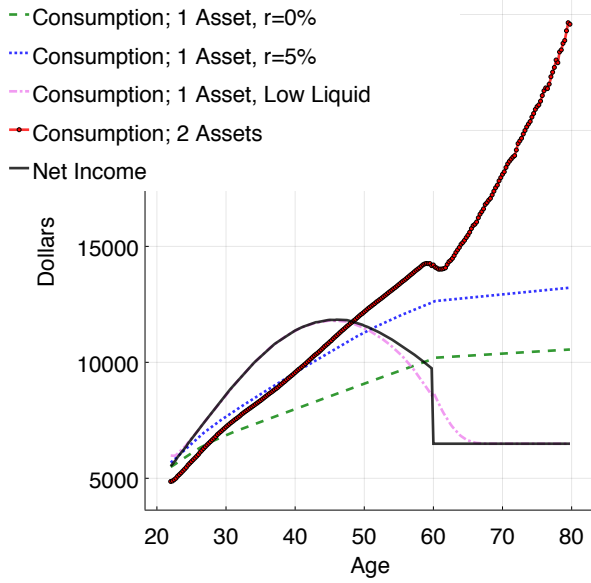
$$a' = R_a^{-1}a$$

$$l' \geq 0$$

$$h \geq -\zeta a$$

The parameter values in the two-asset model are as follows: $T^{work} = 38 \times 4$, $T^{ret} = 20 \times 4$, $\gamma = 2$, $\zeta = 0.03$, $\phi = 0.85$, $\kappa = \$2,000$, $\rho = 0.97$, $\sigma_{z,0}^2 = 0.18$, $R_l = 1$, $R_a = 1.05$ (annual). The social-security function, tax function, age profile χ_t and income risk $\sigma_{\eta,t}^2$ are the same as in Kaplan and Violante (2014). The distribution of initial wealth (a_0, l_0) is taken from Kaplan and Violante (2010), with a illiquid share of 81.4% as in Kaplan and Violante (2014). The discount factor $\beta = 0.9651$ to match a ratio of total wealth to mean quarterly labor income of 14.

The one-asset model is a special case of the two-asset model in which housing services do not enter the utility function ($\phi = 1$), there is an infinite cost of adjustment ($\kappa = \infty$) and all households start off life with zero illiquid assets $a_0 = 0$. We consider versions with $R \in \{1, 1.05\}$ annual. The discount factor is re-calibrated to generate the same ratio of total wealth to mean quarterly labor income as in the two-asset model.
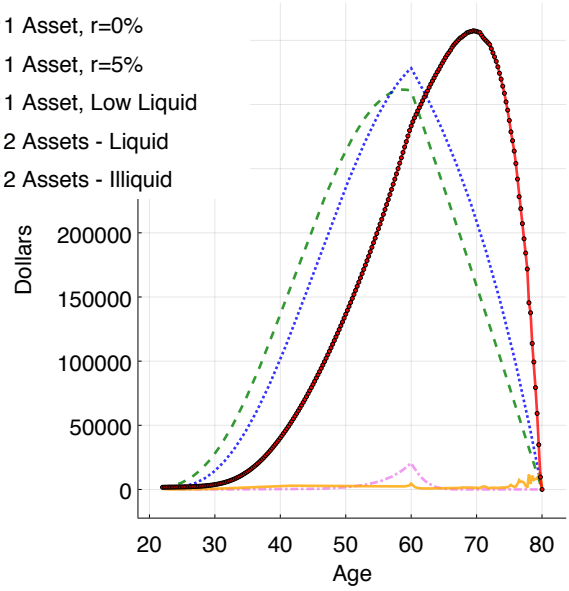
The life-cycle profiles of income and consumption are displayed in Figure 8a and the life-cycle profiles of wealth are shown in Figure 8b. Note that the increasing mean profile for consumption during retirement in the two-asset model and associated sharp decline in mean illiquid wealth are a consequence of the lack of a bequest motive in the model. Introducing a bequest motive would change this feature without having any meaningful affect for the results on MPCs in the main text. The age profile of the fraction of hand-to-mouth households in the model is displayed in Figure 8c.

Figure 8: Life-cycle profiles in precautionary savings models

(a) Consumption and income

(b) Wealth

(c) Hand-to-mouth households

63