Savings and Consumption Responses to Persistent Income Shocks

Brian E. Higgins* Terry O'Malley† Fang Yao‡

March 26, 2026

Abstract

We propose a novel identification strategy to identify the impact of news about future income on consumption today—intertemporal marginal propensities to consume (iMPCs)—which are important determinants of the extent of Ricardian non-equivalence. We implement the identification strategy using an unexpected and persistent shock to disposable income arising from a shock to mortgage contracts in Ireland. We estimate that today households consume 0.8 percent of the income that will be received in more than ten years time—and show that this is smaller than the response in typical one- and two-asset heterogeneous agent models. The equivalent iMPC of high liquid-asset households is considerably larger than the iMPC of low liquid-asset households. We also find that the total response of consumption to our shock is high on average (91 percent of income received today is consumed today), and that it is consistent with the average response of a standard consumption-savings model with an income shock that has the same duration in our data (17 years). In contrast to mixed results in the literature, we find that low liquid asset households have larger consumption responses than high liquid asset households, which are of similar magnitude to those in standard quantitative models.

*IIES. Email: brian.higgins@iies.su.se.

†Unaffiliated. Email: terryomalley89@gmail.com.

‡Central Bank of Ireland. Email: fang.yao@centralbank.ie.

We thank Adrien Auclert, Luigi Bocola, Doireann Fitzgerald, Peter Ganong, Robert Kelly, Pete Klenow, Per Krusell, Philip Lane, Kieran Larkin, Reamonn Lydon, Neale Mahoney, Greg Mankiw, Fergal McCann, Kurt Mitman, Arash Nekoki, Jonathan Parker, Monika Piazzesi, Luigi Pistaferri, Ciaran Rogers, Martin Schneider, Amit Seru, Chris Tonetti, and Christian Wolf, as well as many participants at the Central Bank of Ireland (Seminar and Winter Workshop), Esade (Macro Meetings), Ghent (Empirical Macro), the IIES, Jonkoping, SED (Copenhagen), and Stanford (Labor-Public Workshop, Macroeconomics Lunch and GSB Household Finance Workshop). This project was started while O'Malley was affiliated with the Central Bank of Ireland. Higgins is grateful to Fergal Mc Cann for hosting him at the Central Bank of Ireland and thanks the George P. Shultz Dissertation Support Fund at Stanford Institute for Economic Policy Research for financial support.

Recent progress has developed macroeconomic models that are consistent with the empirical evidence that consumers have high marginal propensities to consume out of transitory income (Jappelli and Pistaferri, 2010, 2017; Kaplan and Violante, 2022). This responsiveness of consumers also suggests that the stimulative effects of fiscal policy may be large. Yet, the stimulative effect of fiscal policy does not only depend on how consumers react to income that is received today, but also how they respond to the expected future taxes—in other words, on the degree of Ricardian non-equivalence.

Indeed, with sufficient empirical evidence on how consumption responds to income at all past and future dates, one could dispense with a fully specified model to perform counterfactuals on fiscal policy. That is because the matrix of intertemporal marginal propensities to consume ("iMPCs")—the aggregate response of consumption in period t to income shocks in in past and future periods s $\partial C_t/\partial Y_s$—is a sufficient statistic for the transmission of fiscal policy to consumption (Auclert, Rognlie and Straub, 2023). Yet, as Auclert, Rognlie and Straub (2023) note, there is practically no evidence on how consumers respond to future income shocks, and in its absence we must rely on the structure of our models to fill in this missing information.

In this paper we propose a novel empirical strategy to estimate forward looking iMPCs using persistent income shocks and implement it in a unique setting in Ireland. We estimate that consumers spend 0.8 percent of income that will be received 15 to 25 years in the future—about half as much as what models predict. We also find that the equivalent iMPC of high liquid-asset households is ten times larger (0.013) than low liquid-asset households (0.001).

The challenge to identifying forward-looking iMPCs arises because there are few instances in which the ideal experiment exists. Such an ideal experiment would credibly induce random variation in income in a single period in the future. The experiment would need to have sufficient power to detect changes in consumption today, a challenge heightened by the fact that most models predict that the magnitude would be very small.

Our identification strategy leverages persistent income shocks. Consider a consumer who knows she will receive income $d\tau$ every period for k periods. And suppose that an otherwise identical consumer receives the same income stream for $k+1$ periods. Since they both receive the same stream of payments until the final period, the difference between their consumption today is only driven by knowledge about the additional income in period $k+1$. Thus by comparing the *total* response of the consumers, to the sequence of income shocks, we identify the *marginal* response of consumption today to the new about the income received in period $k+1$, i.e. the (k+1)-th element of the iMPC matrix.

Even with this identification strategy in hand, it remains challenging to find a shock that

provides variation in the duration of the shock. We use a novel shock to mortgages that occurred in Ireland in the wake of the financial crisis. The shock— depicted in Figure 1—resulted in a quasi-*permanent* divergence in the payments due on two ex-ante similar mortgage contracts. But, since mortgage contracts have a fixed maturity, the shock was in effect of fixed *duration*—and the duration varied across consumers based on the remaining time-to-maturity at the time of the shock. Another advantage of the shock is its large magnitude, and thus allows us to test for the importance of liquidity constraints in a manner that is robust to Cochrane (1989)'s critique that non-optimal responses to small shocks imply trivial utility costs.

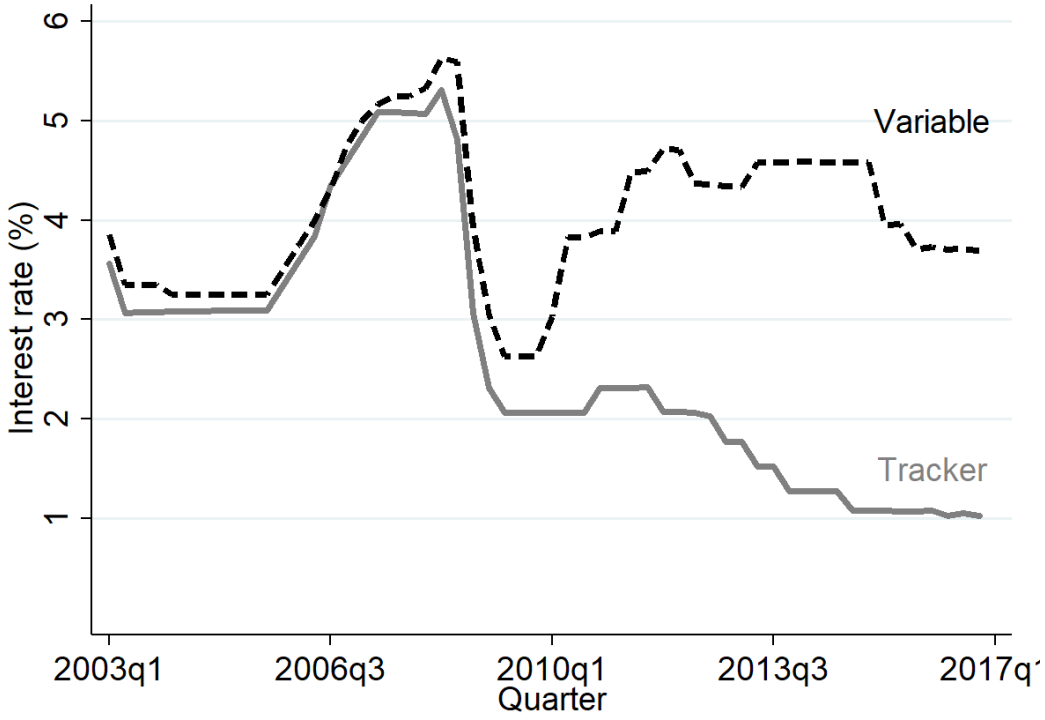


Figure 1: Divergence of the *Tracker* and *Variable* mortgage interest rates

Notes: This figure plots the mean interest rate of Tracker and Variable rate mortgages in our data. The Tracker rate is contracted to follow the ECB's rate plus a fixed spread. The Variable rate is the so-called Standard Variable Rate that banks offer to customers. It is not indexed to any specific rate and typically follows the funding costs of banks. *Source:* Mortgage Dataset, Central Bank of Ireland.

The shock comes from the divergence of the two main types of adjustable rate mortgages in Ireland in 2010. In the early 2000s, Irish banks offered two comparable adjustable-rate mortgage contracts: the so-called standard *variable* rate was set at the discretion of each bank and typically reflected the cost of funding due to competitive pressures. The second product was known as a *tracker* mortgage, because it was contractually indexed to the ECB rate plus a fixed spread, and thus tracked the ECB's rate. As Figure 1 shows, these products were highly comparable: the

interest rates closely followed each other until the financial crisis, after which they diverged due to increased funding costs for Irish banks. Banks stopped offering *new* tracker mortgages before the rates diverged, so the difference only affected existing mortgages.

A challenge to using mortgage interest rates to estimate consumption responses is that they typically generate both substitution and income effects. However, an attractive feature of our setting is that the divergence in interest rates only affected existing loans, so the interest rate on *new* lending and savings remained the same for both groups. There remains a possible substitution effect coming from the existing mortgage: the higher variable rate holders may want to repay their mortgage faster than the trackers. Conveniently, we can test this using data on repayments and find no difference between the two groups. The remaining difference is that tracker borrowers received a large and positive stream of disposable income relative to comparable variable rate borrowers.

We estimate the propensity to save using linked bank account and mortgage data, and have have three empirical findings. To be comparable to the literature on consumption responses we discuss our results in terms of consumption responses.[1] In addition to the forward looking iMPCs, our shock also delivers estimates of the total response of consumption to persistent shocks, and heterogeneity across the distribution of liquid assets.

Our first empirical finding is that the total response of consumption $dc_t/d\tau_t$ is very high on average: households consume 91 percent of the income that is received in a given period. While this is much higher than the typical estimate in the literature on transitory shocks, it is qualitatively consistent with the view consumption should move almost one-for-one with highly persistent shocks. To quantitatively evaluate this intuition, we compare our estimate to the consumption response of a standard incomplete markets consumption-savings model in which infinitely lived consumers are subject to incomes shocks and a borrowing constraint. The average household in the data had a mortgage with 17 years (68 quarters) of remaining maturity when the shock occurred. We thus feed a 68 quarter income shock into our model and show that we can exactly match the average response, an MPC of 91 percent. At least for the average response to a persistent shock, we conclude that the standard model performs well.

Second, we estimate the consumption response to shocks of different duration. We identify variation in the duration in the shock using the time-to-maturity of the mortgage at the time that the divergence occurred. We find that the the total consumption response is increasing in the duration of the shock. The estimates imply that households consume 0.8 percent of the income that

[1]Since we do not observe the components of consumption, we do not distinguish between consumption and expenditures, where the difference arises from durables (Laibson, Maxted and Moll, 2022). In the language of Laibson, Maxted and Moll (2022), our consumption responses are the marginal propensity to spend (MPX).

will be received 15-24 years in the future. This is about half as large as standard consumption-savings models predict. We also find that the equivalent iMPC of high liquid-asset households is ten times larger (0.013) than low liquid-asset households (0.001). This is consistent with the presence of binding liquidity constraints: even if low liquid-asset households wanted to consume more from the income they expect to receive in the future, borrowing constraints may prevent them from doing so.

Third, we evaluate the importance of liquidity constraints by testing whether households with low liquid assets have higher MPCs than those with high liquid assets—as is tradition in the literature since Zeldes (1989). We find that this is the case; the MPC for households in the lowest liquid asset quartile is 100 percent and monotonically decreases to 82 percent for households in the highest liquid asset quartile. The quantitative model predicts a similar consumption response, with an MPC of 102 percent for the lowest quartile and of 72 percent for the highest quartile. Again, that the model and data responses are within 10 percentage points suggests that the standard model with liquidity constraints fits the data well in response to persistent shocks. The empirical variation across the liquid asset distribution is of similar magnitude to the estimates for transitory shocks: we estimate a spread across liquid asset quartiles of 18 percentage points, while Fagereng, Holm and Natvik (2021) estimate a 16 percentage point spread in response to transitory shocks. We consider this to be new evidence that liquidity constraints can generate large, heterogeneous consumption responses even for persistent shocks.

An advantage of our persistent shock is that it is also large. By 2014 Q1, the median tracker household had lower payments that cumulatively exceeded €12,000 (the median income at origination is €66,000). The size of our shock may explain why we find strong evidence in favor of liquidity constraints, whereas the prior literature has found mixed results in response to tax rebates (Johnson, Parker and Souleles, 2001; Parker, Souleles, Johnson and McClelland, 2013). Kueng (2018) suggests that the excess sensitivity of unconstrained households may be explained by the fact that welfare losses from sub-optimal behavior are small when the income shock is also small. In the case of our larger, persistent shock, the welfare losses from mis-optimizing are larger, and this may rationalize why we do not estimate excess sensitivity for non-liquidity constrained households.

We undertake a number of tests to ensure that our results are robust and not affected by selection.[2] While we find some evidence that the tracker and variable rate groups differ on observables,

[2]It is worth noting that an inherent advantage of our shock is that is occurs in a population of interest, adjustable rate mortgage holders, that is representative of over 80 percent of the Irish mortgage market. This compares favorably to adjustable rate mortgages in the US, which represent between 16-38 percent of the mortgage market (Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru and Yao, 2017). That being said, there remains a reasonable concern that despite being highly substitutable products, there may have been systematic selection into either the

we find that controlling for these observable characteristics does not change our results. We flexibly control for income, age, mortgage size, location, year of origination and mortgage size, and allow the effect of each control to vary over time. In our preferred specification, the average MPC is exactly the same, 91.4 percent, with and without controls, and the inclusion of controls alters the coefficient by no more than 1 percentage point across our four specifications.

To rule out selection on unobserved heterogeneity we follow Byrne, Kelly and O'Toole (2021) and use an ancillary survey of Irish mortgage holders conducted by the Central Bank of Ireland in 2012 and 2013. This is the period after the interest rate shock occurred and during the highest levels of unemployment in Ireland following the financial crisis. We show that at this time, there is no difference in the probability of being unemployed or in the income change experienced by tracker and variable rate mortgage holders. It is reassuring that, at least for these important predictors of consumption changes, there is no evidence of selection on unobserved heterogeneity.

Lastly, we check the robustness of our measurement of assets. We show that that savings held in bank checking and savings accounts capture the vast majority of financial assets of Irish households. Using the ECB's Household Finance and Consumption Survey, we show that bank deposits represent 78 percent of financial assets for Irish mortgage holders. Furthermore, this share is stable across the liquid asset distribution. This suggests that unobservable savings are modest on average and are unlikely to bias our estimates of MPC heterogeneity.

Contribution to the literature. Our paper contributes to a literature on the role of consumer heterogeneity in macroeconomic models. Since *transitory* shocks map directly into a the definition of a marginal propensity to consumed, much of the empirical evidence has focused on responses to *transitory* income shocks (Jappelli and Pistaferri, 2017; Kaplan and Violante, 2022).[3] Our contribution is to show that *persistent* shocks can be particularly informative about *forward-looking* iMPCs, which Auclert, Rognlie and Straub (2023) show are important determinants of fiscal multipliers. The only evidence on forward-looking iMPCs come from self-reported responses to surveys and are measured at short horizons (e.g. up to one year ahead) (Colarieti, Mei and Stantcheva, 2025; Eichenbaum, Guerreiro and Obradovic, 2026). We complement their evidence with estimates from observed behaviour and for income received at further horizons in the future (up to 80 quarters)—a time horizon that is especially applicable for deficit financed fiscal stimulus.

tracker or variable product.

[3]This literature has leveraged expected changes such as tax rebates (Parker, 1999; Shapiro and Slemrod, 2003; Johnson, Parker and Souleles, 2001; Parker, Souleles, Johnson and McClelland, 2013) and unexpected changes such as dividends (Bodkin, 1959; Agarwal and Qian, 2014), lottery wins (Fagereng et al., 2021; Golosov et al., 2023), unemployment (Gruber, 1997; Browning and Crossley, 2001; Ganong and Noel, 2019) and firm-wide income changes (Ganong, Jones, Noel, Farrell, Greig and Wheat, 2023). Pilossoph, Lewis and Melcangi (2022) estimate latent heterogeneiety in the MPC using the tax rebates data.

We also show that persistent shocks are also informative about the role of liquidity constraints in driving consumption behaviour. While the literature consistently finds empirical MPCs that are higher on average, there is mixed evidence on the key prediction of models with liquidity constraints: that households with less liquid assets consume more out an income shock. For example, US tax rebates studies have found different covariances with liquid assets (Parker, 1999; Johnson, Parker and Souleles, 2001), and recent experimental evidence finds no covariance with liquid assets (Boehm, Fize and Jaravel, 2025). Cochrane (1989) argued that small shocks, like those in much of the empirical literature, may not be informative about sensitivity, because low stakes windfall have small utility costs to near-rational behaviour. Persistent shocks have the potential to overcome this challenge as even modest shocks may accumulate into large stakes windfalls. Consistent with this argument, we find strong evidence in favor that low-liquid asset households having higher MPCs.

These empirical results are informative for models of studying the impact of fiscal policy (Kaplan, Moll and Violante, 2018; Nakamura and Steinsson, 2018; Auclert, Rognlie and Straub, 2023; Wolf, 2023). As Auclert, Rognlie and Straub (2023) show, the matrix of iMPCs is a sufficient statistic for the transmission of fiscal policy to consumption; our estimates provide the first empirical discipline on the forward-looking elements of this matrix. Wolf (2025) shows that the same objects are important for studying monetary policy. Orchard, Ramey and Wieland (2025) re-evaluate micro-evidence fiscal stimulus payments using macro counteractions.

Secondly, we contribute to a literature testing the permanent income hypothesis based on whether there is *excess sensitivity* to expected and or *excess smoothness* to unexpected income shocks (Flavin, 1981, 1993).[4] We find evidence in favour of the permanent income hypothesis, and against excess smoothness to *permanent shocks*, in the sense that consumption does move almost one-for-one with a highly persistent income shock, and can be rationalized with a reasonable calibration of a quantitative model. Imbens, Rubin and Sacerdote (2001) also quasi-experimentally identify persistent income shocks from lotteries.[5]

[4]See also the textbook treatments on excess smoothness and excess sensitivity in Jappelli and Pistaferri (2017) and Romer (2019).

[5]Baker (2018) identifies persistent income shocks using firm shocks as instruments. He finds an elasticity of consumption to income of 0.4. This lower elasticity may be reflective of the fact that his income shocks, while persistent, are mean reverting by around 40 percent after six quarters. Beyond our focus on consumption, there is a also a recent literature studying the impact of persistent shocks on labor supply Zator (2022), debt reduction Cookson, Gilje and Heimer (2022), and business formation Cookson, Bellon, Gilje and Heimer (2021). There is also a rich literature using covariance restrictions (Hall and Mishkin, 1982; Blundell, Pistaferri and Preston, 2008) and subjective expectations (Hayashi, 1985; Pistaferri, 2001) to estimate consumption responses to both transitory and persistent income shocks. Similarly, while our focus is on unexpected persistent shocks, our paper relates to those studying income changes that were expected in advance, such as at retirement (Bernheim, Skinner and Weinberg, 2001; Aguiar and Hurst, 2005, 2007).

Third, we contribute to the literature studying the interaction of housing, mortgages and consumption (Berger, Guerrieri, Lorenzoni and Vavra, 2018; Beraja, Fuster, Hurst and Vavra, 2019; Berger, Milbradt, Tourre and Vavra, 2021; Guren, McKay, Nakamura and Steinsson, 2021; Wong, 2021; Kinnerud, 2025; Hedlund, Larkin, Mitman and Ozkan, 2025). Gupta (2019) first proposed a similar identification strategy—leveraging differing indexes of adjustable-rate mortgages—to estimate the impact of default on neighbourhood spillovers. (Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru and Yao, 2017) also use mortgage-rate variation to identify the spending response to changes in mortgage payments.[6]

Lastly, we contribute to a growing literature using Irish microdata to study a variety of mortgages and household finance topics (Byrne, Kelly and O'Toole, 2021; O'Malley, 2021; Higgins, 2024; Acharya, Bergant, Crosignani, Eisert and McCann, 2022; Byrne, Devine, King, McCarthy and Palmer, 2022). Ireland is an especially interesting setting considering that it had one of the most severe economic and banking crises in response to the global finance crisis (e.g. see Lane (2012) and Honohan (2019)). Among this literature, we most directly relate to Byrne, Kelly and O'Toole (2021), who use the same interest rate shock to identify the impact of changes in mortgage payments on mortgage default.

1 Identification of Forward-Looking iMPCs

In this section we describe how persistent income shocks of varying length can identify forward-looking intertemporal marginal propensities to consume (iMPCs). Consider the class of models in which individual consumption can be written as a function of unearned income τ_t at all future dates $c_t\left(\{\tau_t\}_0^\infty\right)$—the one- and two-asset models in Section 6 fall into this class of models. To first order, the *total* response of consumption to a sequence of income shocks of length k is the sum of the marginal propensities to consume at each time horizon multiplied by the size of the shock at each horizon

$$dc_0(k) \equiv dc_0(\{d\tau_s\}_0^k) = \sum_{s=0}^{k} \frac{\partial c_0}{\partial \tau_s} d\tau_s. \tag{1}$$

Comparing consumers receiving different length shocks $dc_0(k+1)$ and $dc_0(k)$ yields the marginal propensity to consume out of income at the date $k+1$

$$\frac{dc_0(k+1)}{d\tau} - \frac{dc_0(k)}{d\tau} = \frac{\partial c_0}{\partial \tau_{k+1}}. \tag{2}$$

[6]Druedahl et al. (2021) study high-frequency identification of mortgage payment shocks.

where for simplicity we have assumed constant sequence of $d\tau_s = d\tau$.[7] This is the forward-looking iMPC: the response of consumption today to income in a future period $k+1$. The key insight is that while estimating the *partial* response $\frac{\partial c_0}{\partial \tau_{k+1}}$ directly would require an experiment that varies income in a single future period, the differencing strategy above identifies it from the *total* responses to shocks of different lengths.

Expectations. The first-order approximation implicitly assumes the consumer has perfect foresight over the sequence of shocks. In practice there may be uncertainty over whether a future shock will realize. If the income will be received with probability $p_{k+1} \in [0, 1]$, then the total response can be still be decomposed into the iMPC and the probability of the shock

$$\frac{dc_0(k+1)}{d\tau} - \frac{dc_0(k)}{d\tau} = p_{k+1}\frac{\partial c_0}{\partial \tau_{k+1}}. \tag{4}$$

In that case, going from the total response of consumption to the iMPC requires knowledge over the probability of the shock realizing. Notice that once we're equipped with an empirical estimate of the left-hand side (the difference in total responses), then assuming the shock was known with certainty, $p_{k+1} = 1$, identifies an lower bound of the iMPC. This is useful because detecting a difference in consumption for a shock at a given horizon implies that the MPC for that horizon is non-zero. With knowledge on the probability of the shock realizing, then we can recover the underlying iMPC.

2 Background on Our Persistent Shock

It is challenging to identify a shock to income with clear heterogeneity in its duration. In this section we describe a unique setting in Ireland in which such a shock occurred. As described in relation to Figure 1, our shock arising from the divergence indexing of two mortgage contracts in Ireland in the aftermath of the financial crisis— this is similar to an identification strategy first used by Gupta (2019) to study mortgages in the US.

In the early 2000's, Irish banks offered two adjustable rate mortgage products: the *variable* rate, and the *tracker* rate; as well as a short-term fixed mortgage (typically fixed for up to five years). As figure 2a illustrates, the majority of mortgages issued at this time were one of these two adjustable rate mortgages (Kelly, Lyons and O'Toole, 2015). In the years while both the tracker

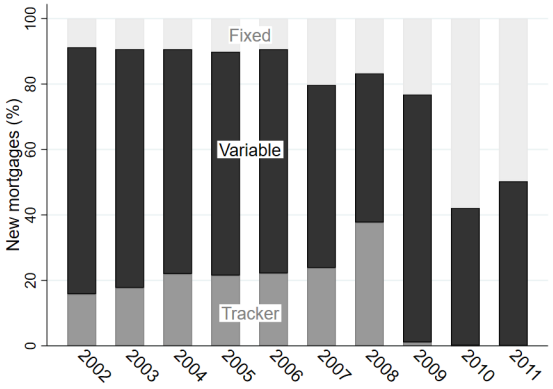
[7] The identification relies on the same time path of the shocks, and not a constant stream. In such a case, the total response should be normalized by the income received at the future date

$$\frac{dc_0(k+1)}{d\tau_{k+1}} - \frac{dc_0(k)}{d\tau_{k+1}} = \frac{\partial c_0}{\partial \tau_{k+1}}. \tag{3}$$

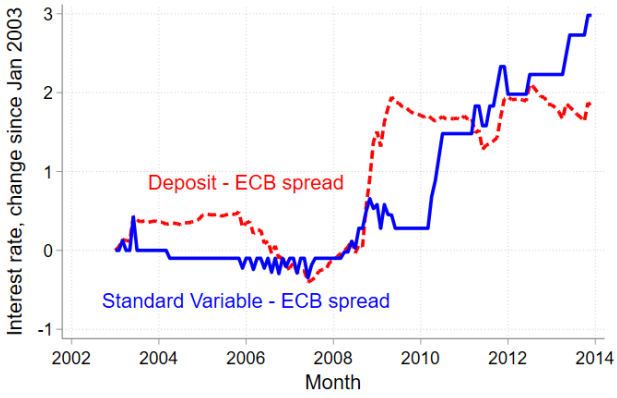
and variable products were offered to new customers their interest rates closely tracked each other. The interest rate on the variable product is set at the discretion of the banks, yet competition between banks meant that it is closely related to the funding costs of banks. Thus, changes in the ECB's main refinancing rate are typically passed on to consumers. The tracker rate on the other hand is contractually indexed to the ECB's main refinancing rate plus a fixed spread that is typically less than 100 basis points. The tracker rate was a relatively new product, having been introduced by one of the banks in 2001 and the remaining banks followed suit shortly thereafter (Byrne, Kelly and O'Toole, 2021).

Figure 2: Variable and Tracker mortgages rates and issuance over time

(a) Share of new mortgages issued by type



(b) Change in spread since Jan 2003 (%)



Note: Left plot shows the share of new mortgages issued by interest rate contract type. Source: Mortgage Data, Central Bank of Ireland. Right plot shows the change in two spreads, (i) deposit rate minus ECB rate and (ii) standard variable rate minus ECB rate. To focus on the change over time, we normalize both spreads to zero in January 2013 Source: Central Bank of Ireland, Retail Interest Rates Statistics (Tables B.1.1. and B.1.2).

Source of the Divergence. The divergence in Figure 1 occurred because, in the aftermath of the global financial crisis, the funding costs of Irish banks increased substantially while the ECB's rate fell, leading to a wedge between the index of each contract. Given this wedge, the tracker product became a loss-making product and banks stopped issuing *new* tracker mortgages as shown in Figure 2a.

There is no direct data available on the exact funding costs of Irish banks. That said Figure 2b compares a component of the funding costs (the interest rate on deposits) with the gap in the mortgage products. It shows that as the banks' cost of deposit funding rose relative to the ECB rate, the spread between the variable rate and the ECB rate increased by a similar magnitude. Between 2007 and 2012, both spreads increased by about 2 percentage points. To the extent that bond or interbank funding was still available to Irish banks at this time, the cost of these funds also increased substantially. Lane (2015) reports that spreads on Irish banks' credit default swaps

increased from approximately 0 percentage points in 2007 to between 2 and 6 percentage points in 2009 and 2010.

Income and Substitution Effects. Changes in interest rates usually have both substitution and income effects, but we consider our shock to primarily induce income effect. The primary reason is that the divergence only affected pre-existing mortgages, because tracker mortgages were no longer offered after the divergence occurred. Thus, both groups faced the same interest rates on *new* mortgages and on savings products, and in the absence of different prices there is no difference in substitution effects. On the other hand, the lower interest rate on tracker mortgages directly maps to lower required payments and therefore an increase in disposable income.

There is one caveat, that there remains a possible a substitution effect on the existing loan. In particular, the variable group have a higher interest rate on the pre-existing mortgage and this may induced them to repay their mortgage at a faster rate than the variable rate group. Conveniently, this is an effect that we can directly test for because it induces faster *repayments* rather than higher savings in cash. We test for such differential repayments in Table 11 and find no difference between the groups.

Even in the presence of substitution effects from interest rates, the variation in duration in mortgages could still be informative about the impact of news of future income and interest rates. In the same spirit as Kaplan and Violante (2014) used a quantitative model to purge tax rebates of anticipation effects, one could use a quantitative to model discipline or purge the substitution effects of mortgage interest rates changes.

Duration and Expectations. The shock occurred unexpectedly and was highly persistent. Despite banks making considerable losses on the tracker product, they issued tracker mortgages until two years before the divergence occurred, indicating that they did not anticipate such a large divergence while issuing the mortgages.[8]

As for the persistence of the shock, Figure 1 shows that once a gap of about 2 percentage points opened up, it remained throughout and beyond the sample period.[9] The fact that the shock turned out to be persistent ex-post does not mean that it was known to be persistent when it first occurred. But there is reason to think that households could reasonably forecast that the shock would be highly persistent. Firstly, it could be reasonably deduced that the cause of the divergence was due to the divergence between ECB rates and bank funding costs, and given that the Irish government providing a bailout for the Irish banks it was evident that the high funding cost was not transitory.

[8]The Irish government provided a bailout for the Irish banks, which ultimately led to a bailout of the Irish government itself in the years following (Honohan, 2019).

[9]The spread remained until mid-2022 when the ECB increased its benchmark rate in response to rising inflation.

Secondly, the closest financial product for the yield curve of bank funding costs is the yield curve on Irish government bond yields. At the time of the shock, a similar sized spread opened up between Irish and safe (German) bonds. Crucially, this spread also opened up at long horizon (20 year) maturities indicating that financial markets were pricing in a quasi-permanent spread in funding costs.

Taken together, we think the evidence points a shock that might best be approximated by a highly persistent regime-switching process. In the one regime there is no risk premia between Ireland and the rest of the world, while in another regime there is a gap. While in the high spread regime, the *effective* duration of the shock is governed by the duration of mortgages.

Substitutability of the Products and Selection. While the tracker and variable rate products were close substitutes, one might still be concerned that there is unobservable selection into either rate type, which could occur either because of screening by banks or self-selection by households. In terms of screening by banks, there are some features of the Irish market that mitigate the risk that screening is another source of selection. In general, Irish banks do not tailor mortgage terms or price discriminate based on individual risk. For instance, 94 percent of households with a variable rate in 2014 Q1 had the exact same rate of 4.58 percent.[10] Based on reports at the time, banks sometimes restricted access to tracker mortgages based on certain mortgage characteristics, though importantly these are characteristics that we observe in our data and control for, such as the size of the loan. For example, the Irish Times (2004) reported that some banks only offered tracker rates to loans over €150,000 or €250,000.[11] Other than these observable mortgage characteristics, the Irish Times (2003) reported that tracker mortgages were accessible to all customers, "At most lenders it is possible for standard variable-rate customers to switch to their ECB tracker mortgage. "It's a case of writing in and asking for it," says Mr Ferguson [of mortgage brokers Ferguson and Associates]."

In terms of self-selection by households, there are a number of reasons to think that households considered these products to be close substitutes (and in the next section we provide empirical evidence against self-selection). Since both contracts are adjustable rate mortgages, they are clearly closer substitutes than fixed rate mortgages. This mitigates risks associated with comparing adjustable and fixed rate consumers, who may differ in their risk aversion, or in their expectations

[10] Banks even refer to the variable rate product as their "Standard Variable Rate", miplying that it is offered as standard to most customers.

[11] Irish Times (2004) reported "First Active's tracker mortgage offers several different rates depending on the value of the loan in relation to the property price, known as the loan-to-value ratio (LTV)... But in all cases, customers will only qualify for the tracker rate if they borrow €250,000 or more." "Meanwhile, Permanent TSB has two rates according to the size of the loan: the more you borrow, the less interest you will be charged, but the lowest amount on which it will advance a tracker mortgage rate is €150,000. "

of or exposure to future interest rates. Discussions in newspapers at the time described trackers as "a good alternative to the standard variable rate" (Irish Times, 2003). These reports also noted that some banks offered lower interest on variable rates while other banks did so on tracker rates, which implies that neither product type was clearly dominated by the other.[12] And even if either product was slightly cheaper, there is wider evidence that households rarely choose the minimum cost mortgage product (Robles-Garcia, 2020; Benetton, 2021), which means that it is likely that comparable customers chose both contracts for reasons other than the structure of the contract. All that being said, it is natural to be concerned that there is selection into tracker status based on unobservable characteristics, such as financial sophistication. With this in mind, in the next section we use ancillary survey data to test whether tracker holders were better educated or more exposed to income shocks, and find little evidence that they are.

3 Data

This section describes the datasets used, compares the characteristics of tracker and variable rate borrowers, and outlines the construction of our shock variable.

3.1 Data description

We use data held by the Central Bank of Ireland. We create a household level dataset by merging a loan-level mortgage dataset (the "Mortgage Data") with a dataset of checking and savings bank accounts (the "Bank Account Data"). The data were collected by regulators for separate purposes and were not intended to be linked. We therefore only have unique customer identifiers for one Irish bank. Our final sample has 10351 households and Table 1 summarizes the main variables of this linked dataset (the "Merged Mortgage-Bank Dataset").

Mortgage data. At origination, we observe information about the loan, including the origination balance, house price and the originating LTV, and about the borrower, including their income, age and post code. We also observe quarterly updates about the mortgage, including the current interest rate type, interest rate, outstanding balance, and days past due.

Bank account data. For each account, we observe the quarterly average account balance (computed daily over the past quarter), the account balance on the final day of each half year, and interest accrued.[13] We observe whether the account was a checking or savings account, and current

[12] "three of the six tracker mortgage providers currently have more competitive rates on these mortgages than they do on their standard variable rate, while a further two offer better deals if LTV ratios are below 60 per cent... "The ICS tracker is not particularly competitive," says Ms Wellband. "It's the same as the standard variable rate. I don't see why they didn't shade it a little bit."" (Irish Times, 2003).

[13] Every six months, banks reported the average account balance over the past 3 months and past 6

Table 1: Summary Statistics

	Mean	P25	P50	P75
Household information (at 2012Q4)				
Borrower Age	46	39	45	51
Number of Household Members	1.83	1	2	2
Dublin (%)	49			
Bank account information (at 2012Q4)				
No of Deposit Accounts	1.76	1	1	2
Quarterly Average Account Balance	8581	779	2776	10156
Mortgage account information (at 2012Q4)				
No of Liens	1.72	1	1	2
Outstanding Balance	149956	61974	127036	213287
Current Interest Rate (%)	4.19	4.68	4.68	4.68
Current Loan-to-Value	0.78	0.32	0.68	1.12
Quarterly Mortgage Payments	3214	1908	2813	4025
Tracker Rate (%)	20			
SVR Rate (%)	80			
Quarters to Maturity	59	31	58	86
Primary Dwelling Home (%)	83			
Mortgage account information (at origination)				
Borrower Age at Origination	37.14	30	35	43
Year of Origination	2002.69	2000	2003	2006
Income at Origination	73794	48000	65743	90964
Loan-to-Value at Origination	0.58	0.38	0.6	0.81

Note: Table presents the summary statistics for our main sample of 10351 households. We group the variables into four categories: households level information; bank account information that is observable quarterly; mortgage account information that is observable quarterly; and mortgage account information that is only observable at origination. For information that varies over time, we summarize the data at a single snapshot (2012 Q4). P25, P50 and P75 are the 25th, 50th and 75th percentiles respectively. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland.

information about the account holder (age, post code).[14]

Unit of observation. We construct our data at the household level. We observe the unique customer identifier for each account holder, which we use to merge the mortgage and bank account data. We classify households as any unique customers who hold a bank or mortgage account together. Once we identify households, we merge every bank account and mortgage held by every customer in the household, including single and joint accounts. For example, if we see two

months. We compute the average balance for the missing quarter using Average Balance in Quarter_{t-1} = $2 \times$ Average Balance in Half-Year_t − Average Balance in Quarter_t.

[14]Irish banks call the checking account a *transactional* account and the savings account a *deposit* account.

customers with a joint mortgage, then this household would include all of these customer's bank accounts and mortgages.

Timing. The mortgage data was first collected in 2012 Q1 and has been updated quarterly ever since. The first collection included information at origination which we use to reconstruct the data back to each mortgage's origination date. We are missing mortgages that were originated and repaid or refinanced before 2012. The bank account data was collected on an ad-hoc basis regulators between 2011 Q3 and 2014 Q4. Therefore, our regressions are limited to this period.

Sample restrictions. We do our best to restrict the sample to households who are using actively using their bank accounts. We do this by excluding: (i) households who have exactly zero balance across all bank accounts; and (ii) households whose bank balances (on all bank accounts) are exactly the same in two consecutive two quarters. We make four further restrictions. First, we restrict to households with at least one mortgage issued between 2000 and 2008, which excludes years when tracker mortgages were not issued. Second, we restrict to mortgage holders with at least one tracker or variable rate mortgage, and exclude those with only a fixed rate mortgage. Third, we exclude households who defaulted on their mortgage before or during our sample period, because our disposable income shock does not apply to households who are not paying their mortgage. Fourth, we drop households in the top percentile of balances and changes in balances, corresponding to balances greater than €152k and quarterly changes in excess of 200 log points.[15] To ensure that our sample is consistent across specifications, we also exclude a small number of households for whom we do not observe a full set of controls.

3.2 Measurement of Housing Savings

In this section we evaluate the share of assets that are captured by our measure of checking and savings account balances, because one might be concerned that we are missing savings held in other accounts. To do so, we use the ECB's Household Finance and Consumption Survey, a representative survey of household balance sheets in the euro area. We use the second wave that was undertaken in Ireland between March 2013 - September 2013 and aligns with the period in which we observe our bank account data.[16] We focus on two measures of household assets: real assets, which includes housing, business assets, vehicles and valuables; and financial assets, which includes non-self employment private business, sight accounts, savings accounts, mutual funds, bonds, shares, managed accounts, voluntary pension plans. Together these sum to total assets (excluding public and occupational pension plans).

[15] In robustness, we show that our results are robust the inclusion of these groups.

[16] While this is the second wave of the survey, it is the first wave conducted in Ireland.

Table 2 shows that bank deposits make up 78 percent of financial assets for the sample of households who own a home with a mortgage. This reassuringly suggests that our measure of savings is capturing the vast majority of financial assets, and therefore we are not missing a large fraction of savings. Households may also save in real assets and for completeness table 2 reports that deposits represent a small fraction of real assets. This leaves open the possibility that missing savings may occur in real assets, such as second homes purchases. That said, since real assets are mostly illiquid, it is likely that such savings would show up in financial assets before they are invested in illiquid assets.

Table 2: Deposits as a Share of Assets

	Deposits	Share of Financial	Share of Real
Housing Tenure Status			
Own - outright	35188	84	12
Own - w/mortgage	13222	78	5
Rent/other	7884	90	184
Total	19691	83	54

Note: Table reports mean deposits, and deposits as a share financial and real assets. Deposits are included in financial assets and not real assets. The sample includes 5,419 observations and estimates are weighted using survey weights. Source: Household Finance and Consumption Survey, European Central Bank.

3.3 Sample Balance and Selection

To evaluate the balance between the tracker and variable rate groups, Table 3 compares summary statistics for both groups. We restrict the comparison to a reasonably short window —in 2006 and 2007— to avoid changes driven by large increase in house prices and mortgage balances over the entire sample period. The groups are comparable in terms of location (as measured by the percentage in the capital city Dublin) and age, which is statistically significantly different but not economically so at less than one year. Tracker mortgage holders earn aout 18 percent more. They also bought more expensive houses, and had correspondingly larger initial balances and payments. Despite these average differences, Figure 3 plots the distributions of income, age, mortgage balance and interest rate for both groups and shows that there is considerable overlap across these distributions. This means that we are able to find comparable households in both groups. Given these differences in observables, our identifying assumption will rely on exogeneity *conditional* on observable controls, and we will include a rich set of controls to ensure that we are comparing within groups of similar households.

Though we can use controls to compare tracker and variable rate borrowers with similar observ-

Table 3: Summary Statistics By Interest Rate Type

	Variable (N=2952)		Tracker (N=1096)			
	Mean	Std. Dev.	Mean	Std. Dev.	Difference	p-value
Age	37.5	9.97	36.7	9.56	−0.83	0.02
Dublin (%)	48	49.98	47.35	49	−0.95	0.59
Mortgage Balance	117	107.30	218	154.05	101.46	0.00
Annual Gross income	82	373.82	977	62.75	14.30	0.05
Interest Rate (%)	4.6	0.56	4.5	0.56	−0.14	0.00
Quarterly Mortgage Payment	1987	1712.84	3474	2370.69	1486.58	0.00
Loan-to-Value Ratio (%)	51	26.90	64	23.51	12.42	0.00
House Price	277	410	400.48	314.78	122.90	0.00

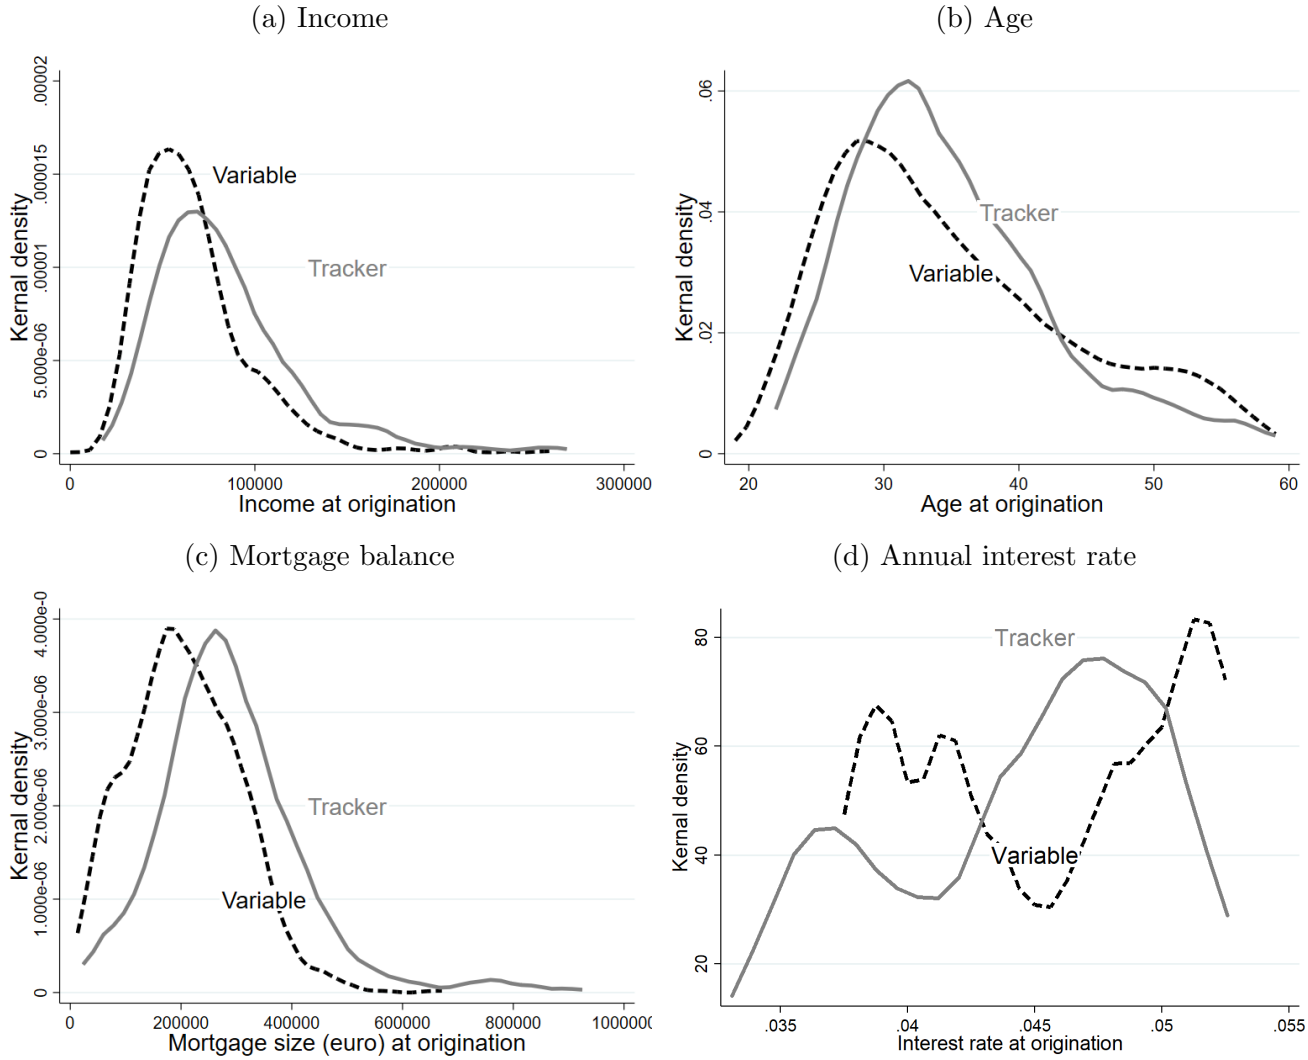
Notes: Table presents comparison of borrower and loan characteristics of variable and tracker groups. Given that contract types may differ across liens, we restrict attention to first lien. We also restrict attention to loans issues between 2006 and 2007, to avoid differences in incomes and house prices driven by variation over time. Mortgage balance, income and house price are shown in 1000s of EUR. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland.

able characteristics, it is reasonable to be concerned that these households differ on unobservable variables that will also influence the savings and consumption behavior of these groups. In the case of mortgages, a natural concern is that financially sophisticated borrowers are self-selecting into either product or that banks are screening based on repayment risk. Fortunately we can test both of these hypotheses using an ad-hoc survey of mortgage borrowers undertaken by the Central Bank of Ireland, which we merge the survey to our mortgage data. The survey took place between May 2012 and February 2013, around the peak of unemployment in Ireland following the financial crisis. The survey asked respondents about their income and employment status at the time of the survey, allowing us to test whether either group experienced bigger incomes changes or more unemployment.[17].

Table 4 reports the results of regressions of (i) the log income change and (ii) unemployment status on the interest rate group and a number of controls. The income change is measured between each household's origination date and the survey date. Columns 3 and 6 show that there is no difference between the two groups conditional on controls: both estimates are economically close to zero and statistically insignificant. While column 1 shows that tracker households experience larger income shocks, this is not statistically significant after the inclusion of origin year fixed effects. Since most tracker mortgages were issued later in the boom, this suggests that households later cohorts of borrowers were more likely to receive negative income shocks, but, reassuringly

[17]The survey has been used by numerous studies at the Central Bank of Ireland. See McCarthy and McQuinn (2013), McCarthy (2014), and Byrne, Kelly and O'Toole (2021) for more information.

Figure 3: Distribution of Variable and Tracker Mortgage Borrowers



Note: Figures show kernal densities for income, age, mortgage balance and interest rates at origination, split by tracker and variable rate group. The sample includes all mortgages for new house purchases originated in 2006 and 2007. For presentation purposes, the plots are truncated to include only ages up to 60, incomes up to €500,000 and mortgage balances up to €1m. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland.

for our purposes, these differences do not differ across the two groups conditional on origin year. Table A1 sequentially adds the controls and reports that adding origin year and income are the biggest contributors to reduction in the coefficient.

Table A2 tests whether education status predicts the interest rate group, which is a test of whether financially sophisticated households self-selected into the tracker group. Unconditionally, tertiary education positively predicts tracker group status, however the sign is reversed when including controls. Indeed, both results are statistically insignificant and thus suggest that education is not an important predictor. Taken together, the survey provides reassuring evidence that the

Table 4: Survey of Mortgage Holders After Interest Rate Shock

	Income Change			Unemployed		
	(1)	(2)	(3)	(4)	(5)	(6)
Tracker	-0.254***	-0.129	0.029	-0.029	0.001	0.005
	(0.0683)	(0.0774)	(0.0611)	(0.0215)	(0.0244)	(0.0256)
Observations	616	616	593	626	626	593
Adjusted R^2	0.020	0.044	0.464	0.001	0.026	0.006
Origin year and bank FE		Yes	Yes		Yes	Yes
Controls			Yes			Yes

Note: This table uses a survey of mortgage holders that is linked to our mortgage data. Columns 1-3 test whether tracker mortgage holders were more likely to experience changes in income between mortgage origination and the survey date. Columns 4-6 test whether tracker mortgage holders were more likely to be unemployed at the time of the survey. Controls include age, income at origination, LTV, and whether the household was a first-time buyer. Source: Survey of Mortgage Holders 2012-13 and Mortgage Dataset, Central Bank of Ireland.

tracker group is not selected on unobservable financial literacy or income risk.

3.4 Construction of the shock

The divergence in interest rates resulted in tracker rate mortgage holders paying lower amounts compared to variable rate mortgage holders with the same mortgage terms. We therefore construct our shock so that it can be interpreted as a positive shock to disposable income for the tracker mortgage group. In the robustness section we will show that our results are not sensitive to constructing an analogous disposable income shock that is negative for the variable rate group.

We observe the actual payments of tracker households. To compute their counterfactual payment, we need to know what interest rate they would have had, if they were in the variable rate group. We use the average interest rate of the variable rate group as the counterfactual interest rate for the tracker borrower. The average interest rate, as opposed to an individual specific risk-adjusted interest rate, is a very good approximation of the counterfactual interest rate because there is almost no variation in interest rates within the variable rate group. For example, in 2014 Q1 94 percent of these households had the same rate of 4.58 percent. With this counterfactual interest rate, we compute the counterfactual payment using the annuity formula

$$\text{payment}_t = \text{balance}_t \times r \times \frac{(1+r)^{n(t)}}{(1+r)^{n(t)} - 1}, \tag{5}$$

where t is the current quarter, r is the *quarterly* interest rate and n is the number of quarters to

maturity[18].

In our analysis, it is important to consistently compare stocks and flows of savings. Bank account balances measure the stock of savings while changes in bank account balances measure the flow of savings. Since our shock occurs over multiple periods, we also define its flow and stock components.

The payment savings in a given quarter t —flow payment savings— are the difference between the counterfactual (variable rate) payment and the actual (tracker rate) payment

$$\begin{aligned} m_t^{flow} \equiv \text{payment savings shock}_t = &\ \text{payment}_t^{\text{variable}} - \text{payment}_t^{\text{tracker}} && \text{if tracker} \\ = &\ 0 && \text{if variable.} \end{aligned} \tag{6}$$

The stock of payment savings at time t is the cumulative payment savings received up to time t since the mortgage originated

$$m_t^{stock} = \sum_{j=0}^{t} m_j^{flow}, \tag{7}$$

where $j = 0$ is the origination date of the mortgage[19]. The change in the stock over one quarter is equal to the flow $\Delta m_{t+1,t}^{stock} = m_{t+1}^{flow}$. We estimate the savings response over longer horizons, thus it is useful to note that the change in the stock, $\Delta m_{t+k,t}^{stock} = \sum_{j=0}^{t+k} m_{t+k}^{flow} - \sum_{j=0}^{t} m_j^{flow}$, is equal to the sum of the flows between these dates

$$\Delta m_{t+k,t}^{stock} = \sum_{j=t+1}^{t+k} m_j^{flow}. \tag{8}$$

We calculate the payment savings for each lien held by the household and use the total payment savings for the household. This accounts for borrowers who have different interest rates on separate liens.

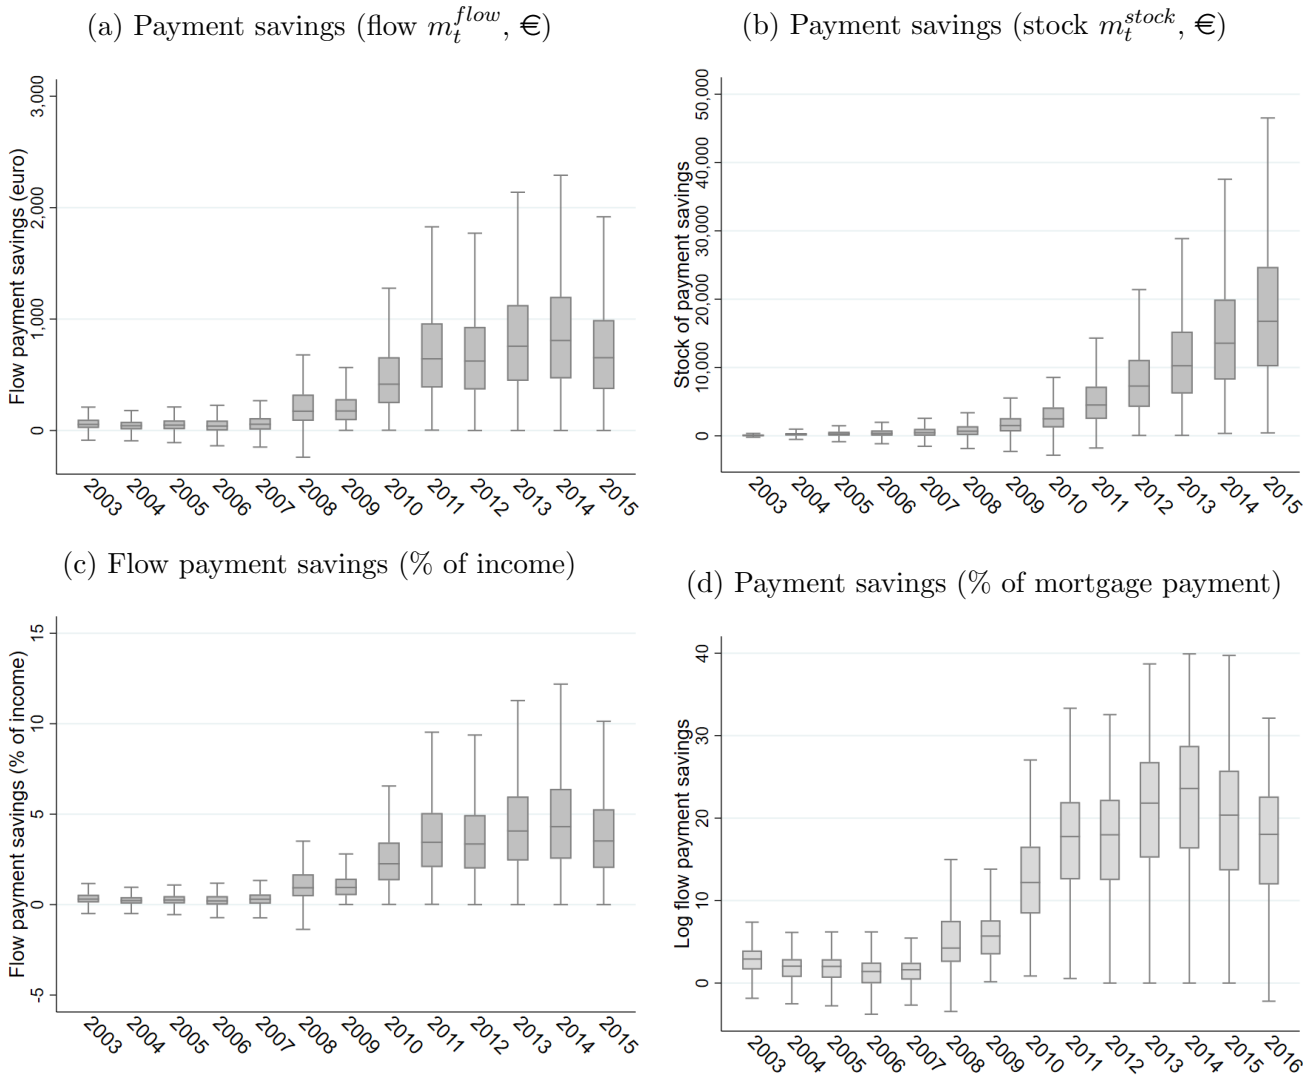
Figure 4 plots the size of the payment shock by year. The flow payment sayings, m_t^{flow} (panel a) were very close to zero in the years leading up to 2010, after which a large difference opens up. In 2011, the interquartile range of this differences is €361-908 per quarter. In line with interpreting

[18]Irish banks typically adjust the interest rate monthly. However, since our data is quarterly, we compute the paymetns that a quarterly frequency. There will however be some small approximation error arising from interest rates changes that occur in months that are not at the quarter end.

[19]While we ignore the accumulation of interest in this definition. Since interest will affect both bank balances and payment savings in the similar manner, the impact on propensities to save/consume is minimal. We account for interest when converting propensities to save at different horizons.

this as a positive disposable income shock for the tracker mortgage holders, the payment savings are positive.

Figure 4: Box Plot of the Size of Payment Savings



Note: Figure shows the distribution of our payment savings shock over time via box plots. For each year the box reports the median, 25th and 75th percentile, and whiskers show the upper and lower adjacent values. Panel (a) reports the flow payment savings, Panel (b) reports the stock of payment savings, and Panel (c) reports the flow payment savings as a percentage of income (as measured at origination). Panel (d) reports the flow payment savings as a percentage of the counterfactual mortgage payment on the first lien. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland.

Panel b plots the stock of payment savings in each year. Given that the flow payment savings are approximately flat after 2011, the stock of savings increases linearly in the years following. By 2014 Q1, the median tracker household saved €12,206. To give a sense of how large these savings are, it is worth noting that the cumulative size of the shock by this point is larger that the mean

bank balance —€8,581 — reported in Table 1.

Panel c plots the flow payment savings as a percentage of the household's income at origination. For the median household, the payment savings in 2014 was equivalent to about to 4 percent of their income. Panel d plots the flow payment savings as a percentage of the household's counterfactual mortgage payment. From 2010 onwards, the majority of tracker households saved at least 10 percent relative to their counterfactual payment and, in 2014, the median tracker household saved 23 percent.

4 Empirical strategy

To estimate the savings response to a change in mortgage payments, we estimate the following equation

$$\Delta b_{i,t+j} \quad = \quad \beta_0 + \beta_1 \Delta m^{stock}_{i,t+j} + \sum_t \gamma_t (X_i \times \tau_t) + u_{i,t}, \tag{9}$$

where $\Delta b_{i,t+j}$ is the change in bank balance of household i in between quarter t and $t+j$, $\Delta m^{stock}_{i,t+j}$ is the change in the stock of payment savings between quarter t and $t+j$, X_i is a vector of time-invariant characteristics of households i, and τ_t are quarter fixed effects that are interacted with the time-invariant characteristics.

We only include household characteristics that were known at the time of the shock and could otherwise affect changes in consumption, including geographic, lifecycle and leverage factors. Controlling for household characteristics that vary after the shock, such as income or total outstanding balance, would be, in the language of Angrist and Pischke (2009), "bad controls" because they may be directly affected by the shock. For instance the shock may induce households to work more or households may use the income to pay down their mortgage. The characteristics X_i include the age of the household at the time of the shock, the age the household when they took out the mortgage, the income of the household at origination, the postcode of the households primary residence (as measured by the collateral on the mortgage), the year of origination of the first lien and the total outstanding balance as at 2008Q4. To allow for non-linearities, we group continuous control into between five and tens bins, and estimate a separate effect for each bin. To allow for the effect of these fixed characteristics to vary over time, we interact them with quarter fixed effects τ_t, and the estimated coefficients are thus indexed by time, γ_t. For example, this allows for income shocks to have a different effect on savings for high and low income households, and for those affects to be different in 2010 than in 2014.

Since equation (9) is in differences, we are estimating the relationship between the *flow* of bank

savings and *flow* of mortgage savings between t and $t+j$. The coefficient β_1 estimates the proportion of the payment savings that are saved in bank accounts, which is the marginal propensity to save. Our primary specification is in euros because this identifies the marginal propensity to save that is directly comparable to our model. We also report results for all specifications in logs.[20]

Our identifying assumption is that the error term is uncorrelated with the payment difference conditional on our controls

$$E[\Delta m_{it}^{stock} \cdot u_{it} | X_{it}, \tau_t] = 0. \tag{10}$$

This is a stronger identifying assumption than unconditional exogeneity. If our shock was truly randomly distributed across borrowers we could rely on the weaker assumption that the error is unconditionally exogenous $E[\Delta m_{it}^{stock} \cdot u_{it}] = 0$, however Section 3 showed that there is some evidence of selection between the two borrower groups. Under this assumption, we estimate the causal impact of payment reductions on savings. Our results are stable to adding control variables which we interpret as evidence that selection is not a major concern.

Equation (9) is also general enough to estimate changes at different horizons. For $j = 1$ it encompasses the typical equation used in the literature to estimate one period MPCs from transitory income shocks (Fagereng, Holm and Natvik, 2021). For $j > 1$ it can also estimate longer differences, which may be important given that our shock is persistent. We also run this in regression in levels, in which case we include individual with fixed effects and estimate the following equation

$$b_{i,t+1} \quad = \quad \beta_0 + \beta_1 m_{i,t}^{stock} + \sum_t \gamma_t (X_i \times \tau_t) + \alpha_i + u_{it}. \tag{11}$$

As noted by (Fagereng, Holm and Natvik, 2021), in the presence of heterogeneous responses our estimates are a weighted average of the marginal propensity to save. This is because OLS weights the coefficient on the regression by the square of the size of the shock, resulting in an estimate that is biased towards those who received the largest change in income. Since theory predicts that there will be heterogeneous responses, we investigate these by performing sub-sample analysis along the dimensions of suggested by the theory.

[20] To account for zeros, we define the log of variable x to be $log(1 + x)$.

5 Empirical results

5.1 Average saving response

One period time horizon. Table 5 reports our main estimates of the propensity to save. The columns report estimates in levels (column 1), log levels (column 2), level differences (column 3) and log differences (column 4). Column 5-8 report estimates in the same order while including a rich set of controls. We prefer the specification in column 7, which corresponds to equation 9 with our full set of controls, because it is in euros and thus directly corresponds to a model implied MPC. This specification estimates that for every dollar of income households save 0.087 dollars or 8.7 percent. This implies an MPC of 91 percent. The estimate is statistically significant at the one percent level. Alternatively, we could test the hypothesis that the MPC is zero and that the MPS is one. The final row of the table tests this hypothesis with a t-test that $\beta = 1$ and reports that the probability is below 0.001 for all specifications.

Table 5: The Marginal Propensity to Save

	(1) Savings	(2) Log Savings	(3) Δ Savings	(4) Δ Log Savings	(5) Savings	(6) Log Savings	(7) Δ Savings	(8) Δ Log Savings
Payment Savings	0.083***				0.077***			
	(0.0210)				(0.0242)			
Log Payment Savings		0.067***				0.076**		
		(0.0241)				(0.0326)		
D.Payment Savings			0.086***				0.086***	
			(0.0221)				(0.0219)	
D.Log Payment Savings				0.056***				0.059**
				(0.0154)				(0.0217)
Observations	144914	144914	134563	134563	144914	144914	134563	134563
Adjusted R^2	0.902	0.907	0.001	0.001	0.903	0.907	0.002	0.002
Individual FE	Yes	Yes			Yes	Yes		
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE					Yes	Yes	Yes	Yes
Prob($\beta = 1$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports estimates from regression equations 9 (differences) and 11 (levels). The dependent variable is a household's bank account balance ("Savings"), $b_{i,t}$. The independent variable is stock of payment savings $m_{i,t}^{stock}$. Columns 1, 2, 5, and 6 are level specifications. Columns 3, 4, 7, and 8 are difference specifications and the independent and dependent variables are first-differenced. Variable names indicate euro and log specifications. Controls are current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, deciles of total outstanding balance as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered by quarter in the differenced specifications, or double clustered by quarter and individual in the levels specifications. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$.

The magnitude and statistical significance of the estimate is practically the same as the same specification without controls (0.086) shown in column 3. This suggests that selection on observables does not meaningfully alter our estimate of the MPS. The estimate is also insensitive to the

specific model chosen. The range of estimates from the levels and level differences specifications, in the odd numbered columns, imply an MPS in the range of 7.7 to 8.6 percent, and therefore an implied MPC in the range of 91-92 percent. The specifications using logs, in even numbered columns, imply a range for the elasticity of savings with respect to income of 5.6-7.6 percent. This range also implies a small savings response and large consumption response.

Multiple period time horizon. Table 6 examines savings response over multiple horizons. Recall that we first observe bank account savings on 2011Q3, which is over one year since the gap in payments opened up. The left most column shows the differences in savings at this date, and we estimate that the tracker group have saved 2 percent of the payment savings that they've received up to this point in time. This difference in savings is not statistically significant. That the gap is small at the first time we observe savings is reassuring as it suggests that the differences in savings by the end of the sample are not simply driven by differences in savings at the beginning of the sample.

Table 6: The Average Marginal Propensity to Save at Different Horizons

	Δ Savings						
	(1) At 2011Q3	(2) 2011Q4	(3) 2012Q1	(4) Q2	(5) Q3	(6) Q4	(7) 2013Q1
Δ Payment Savings 2011Q3-	0.023	0.058	0.093	0.135*	0.144**	0.108*	0.094*
	(0.0519)	(0.1263)	(0.0564)	(0.0574)	(0.0528)	(0.0459)	(0.0434)
Observations	10351	10351	10351	10351	10351	10351	10351
Adjusted R^2	0.064	0.004	-0.001	0.002	0.001	-0.001	0.000
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

	Δ Savings						
	(8) Q2	(9) Q3	(10) Q4	(11) 2014Q1	(12) Q2	(13) Q3	(14) Q4
Δ Payment Savings 2011Q3-	0.083*	0.081*	0.075*	0.074**	0.074**	0.084**	0.093***
	(0.0387)	(0.0353)	(0.0294)	(0.0268)	(0.0247)	(0.0283)	(0.0266)
Observations	10351	10351	10351	10351	10351	10351	10351
Adjusted R^2	0.003	0.004	0.004	0.002	0.003	0.003	0.004
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table reports estimates from regression equation 9 (differences). The dependent variable is the change in household i's bank account balance ("Savings"), $b_{i,t+j}$ between 2011Q3 and the various horizons specified in the titles. The independent variable is the change in household i's stock of payment savings $m^{stock}_{i,t+j}$, at the same horizon as the dependent variable. All variables are in euros. Controls include current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, deciles of total outstanding balance as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered by quarter. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$.

The remaining columns report the share of payment savings that are saved between various horizons, Δm^{stock}_{t+j}. Each regressions starts at $t = 2011Q3$, the first quarter that we observe the savings data, and j is the number of quarters to the quarter listed in the column header. The final column reports the share of payment savings that households have saved by 2014Q4, which is thirteen quarters after we first observe savings balances. The estimates suggest that by the third quarter, there is a statistically significant difference in savings of 0.135. This implies that households have saved 14 percent of the payment savings that they have received up to this point in time. By the end of the sample, the difference in savings is 0.093, which implies that households have saved 9 percent of the payment savings that they have received up to this point in time.[21] Overall the main takeaway is that, on average, households save a very small fraction of the payment savings received over this time period. Therefore the implied consumption response is large, between 86 and 91 percent.

At first glance, our shock might seem well-placed to estimate intertemporal MPCS ("iMPCs") which Auclert, Rognlie and Straub (2023) show are important for disciplining the dynamics of consumption. However, once the gap in flow payments begins, there is little variation flows over time; they are highly auto-correlated. It is therefore infeasible to estimate the impact of any given period's shock while controlling for shocks that occurred in neighboring periods. Relatedly, we are unable to estimate event study specifications because we do not observe a pre-period. Our estimates remain valid under our exogeneity assumption, and the small difference shown in column 1 reassuringly suggests that both groups are similar that at the beginning of the shock.

5.2 Heterogeneity by liquid assets

Table 7 reports the savings response for each quartile of the bank balance distribution. We group households based on their bank balances in 2013Q4, which is the first date at which we observe their savings. The savings response is monotonically increasing as bank balances increase. We estimate that households in the upper quartile of the balance distribution saved 18.4 percent of the payment savings (column 4), the third quartile saved 7 percent (column 3), the second quartile saved 5.2 percent (column 2) and the bottom quartile did not save any of the payment savings (-0.001, column 1). The results are statistically significant at the 1 percent level for the upper quartile and at the 5 percent level for the second quartile. The results are not statistically significantly different from zero for the lowest or third quartile, however, the response of the lowest quartile is a precisely estimated zero, the coefficient is -0.001 and the standard error is 0.01, which implies that all but

[21] It is worth noting that even thought the propensity to save declines between the third and the final quarter, the total savings balance is still increasing. This is because both the denominator (the stock of payment savings) is increasing in each period.

the third quartile are precisely estimated. The final row strongly rejects the hypothesis that the MPC of each group is zero using a t-test of $\beta = 1$.

Table 7: The Marginal Propensity to Save by Liquid Asset Quartile

	Δ Savings				Δ Log Savings			
	(1) Lower	(2) Quartile 2	(3) Quartile 3	(4) Upper	(5) Lower	(6) Quartile 2	(7) Quartile 3	(8) Upper
D.Payment Savings	-0.001	0.052**	0.070	0.184***				
	(0.0106)	(0.0219)	(0.0416)	(0.0542)				
D.Log Payment Savings					-0.023	0.099	0.122***	0.066***
					(0.0596)	(0.0598)	(0.0350)	(0.0091)
Observations	35828	31057	32006	35672	35828	31057	32006	35672
Adjusted R^2	0.000	0.009	0.000	0.006	0.015	0.005	0.003	0.009
Individual FE								
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prob($\beta = 1$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports estimates from regression equation 9 where the sample is split by quartiles of the bank balance distribution. The dependent variable is the change in household i's bank account balance ("Savings"), $b_{i,t+1}$. The independent variable is the change in the stock of payment savings $m^{stock}_{i,t+1}$. Variable names indicate euro and log specifications. Controls are current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, deciles of total outstanding balance as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered by quarter. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*\quad p < 0.1$, $**\quad p < 0.05$, $***\quad p < 0.01$.

The magnitude of the differences in the MPC across the liquid asset distribution is large when compared against the response to one-time shocks. For instance, Fagereng, Holm and Natvik (2021) estimate that the MPC out of a one-time lottery payment is 62 percent for the bottom quartile and 46 percent for the top quartile, which at a 16 percentage point spread is slightly smaller than our estimated 18.4 percentage point spread. This is surprisingly large if your prior was that differences in MPCs arising from liquidity constraints should be smaller for persistent shocks than for one-time shocks.[22]

We acknowledge that the grouping of our bank balance quartiles is based on bank balances that we observe several years after the shock initially occurred. In principle, households could switch between quartiles if the savings response was non-monotonic over these quartiles. That said, our estimates are monotonic. We estimate that those who initially had the least savings continue to save the least and this implies that the ordering of households in the bank balance distribution does not change in response to the shock.

[22]See for instance Campbell (2023)'s video discussing the redistribution channel of monetary policy [13 minutes onwards]. He argues that the redistribution channel of monetary policy only works for transitory changes in interest rates because for permanent changes every household will have the same MPC of one. Our results suggest that large differences in MPCs remain even when the shock is highly persistent.

Columns 5-8 report the results in log differences, which similarly show that the elasticity of savings with respect to the payment shock is larger for households with higher initial balances. The elasticities are not, however, monotonic with the upper quartile reporting a lower elasticity than the third quartile. It is possible that this is driven by the fact that higher quartile households start from a higher based level of savings. For completeness, Table B2 in the appendix reports the estimates in levels rather than level differences. The estimates are similar with the highest quartile saving 18 percent and the lowest quartile reducing savings by one percent.

5.3 Heterogeneity by the duration of the income shock

Table 8 reports how the savings response varies by the duration of the shock. The year-to-maturity of the mortgage, at the time of the shock (2010Q1, is our measure of the maximum potential duration of the income shock. The coefficient of interest is the interaction term which estimates how the savings response changes with the potential duration of the shock. To isolate the impact of news about income, and not differences in realised incomes, I restrict to those with duration greater than five years in column 2 and ten years in column 3.

Table 8: The Marginal Propensity to Save by Year-to-Maturity

	Savings		
	(1)	(2) YTM > 5	(3) YTM > 10
Payment Savings	0.077***	0.271**	0.243**
	(0.0252)	(0.1106)	(0.1132)
Payment Savings x Year to Maturity		-0.008*	-0.006
		(0.0043)	(0.0044)
Observations	144914	137928	116858
Adjusted R^2	0.903	0.903	0.904
Individual FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Controls × Quarter FE	Yes	Yes	Yes
Prob($\beta < 0$)		0.038	0.079

Note: This table reports estimates from regression equation 9 including an interaction term. Standard errors are reported in parentheses. Standard errors are clustered by individual. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*$ $p < 0.1$, $**$ $p < 0.05$, $***$ $p < 0.01$.

Column (2) reports a clear negative relationship: the savings response falls by 0.008 (80 basis points) for each year of longer duration. The coefficient is statistically significant at the ten percent level. Given that theoretical models predict that that magnitude of the coefficient is likely to be small, it is also worth commenting on the precision of the estimate, in addition to testing its

statistical significance relative to zero. The 90 percent confidence interval lies between (-0.0151, -0.0009) which gives a tight range of less than 2 percentage points. Restricting to households with longer duration (column 3, greater than ten years) does not meaningfully change the economic magnitude nor the precision of the estimate (the standard error is practically unchanged despite no longer being statistically significant).

The estimated coefficient on the interaction can also be interpreted as evidence that households expected the shock to have a long duration. Since we estimate that the *potential* duration of income shock is associated with households choices, this suggests that households indeed had large enough expectation that the shock would realise that they reacted to the news. While we will require assumptions on expectations to map total savings responses into marginal propensities to consume, the reduced form estimates can be taken as model free evidence that the duration matters for consumption.

We did not comment on the magnitude, nor do we place much weight on the economic significance, of the coefficient on the payment savings variable once the interaction is included. Taken literally, it could be interpreted as the saving response for someone with zero years-to-maturity, i.e. the savings response to a transitory income shock. However, practically no one in our data has such short duration, so it is an extrapolation outside the range of the data and which relies on the linearity assumption.[23]

Heterogeneity by Liquid Assets. Table 9 explores heterogeneity by initial liquid assets. There is a clear difference in the interaction term: households with above median liquid assets respond about ten times more to duration than households with low liquid assets. The negative coefficient implies that households respond to longer duration by lowering their savings. The fact that low liquid-asset households do not lower their savings as much as higher liquid asset households is consistent with the presence of borrowing constraints. Even if these households wanted to respond to the news of future income, a borrowing constraint could prevent them from doing so.

5.4 Mapping Empirical Estimates to iMPCs

Our empirical estimates correspond to the the change in bank balances for a change in income, $[da_t - (1+r)da_{t-1}]/dy$, which through the budget constraint can be mapped to the total response of consumption

$$\frac{da_t - (1+r)da_{t-1}}{dy} = 1 - \frac{dc_t}{dy}. \tag{12}$$

[23]Recall that tracker mortgages were first issued in the early 2000s and the standard mortgage length is at least twenty years, so at the time of the shock that the shock realized most holders had duration of more than ten years.

Table 9: The Marginal Propensity to Save by Year-to-Maturity and Liquid Assets

	Savings	
	(1) Below Median	(2) Above Median
Payment Savings	0.038	0.441**
	(0.0431)	(0.1844)
Payment Savings x Year to Maturity	-0.001	-0.013*
	(0.0016)	(0.0073)
Observations	69468	68460
Adjusted R^2	0.4863	0.8788
Individual FE	Yes	Yes
Quarter FE	Yes	Yes
Controls × Quarter FE	Yes	Yes
Prob($\beta < 0$)	0.3414	0.0403

Note: This table reports estimates from regression equation 9 including an interaction term. Standard errors are reported in parentheses. Standard errors are clustered by individual. The sample is restricted to households with more than 10 years-to-maturity upon impact. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*\quad p < 0.1$, $**\quad p < 0.05$, $***\quad p < 0.01$.

The interaction term tell us how how this object depends on the duration of the shock, k. It tell us that the total saving response $[da_t(k) - (1+r)da_{t-1}(k)]/dy$ falls by 0.006 as duration k increases. This implies that the total consumption response dc_t/dy increases by 0.006 as duration k increases. From Equation 2 we know that this change corresponds to the marginal propensity to consume out of income received at horizon k.

Our interaction term therefore directly maps into an implied iMPC (of opposite sign to the coefficient), as reported in Table 10. However, which at what horizon does it identify the iMPC? We interpret our estimate as an *average* iMPC for large k approximately in the range $(15, 24)$ years in the future. It is an average because the linearity assumption implies that we estimate a single iMPC in table 8.[24] Since we restrict to our estimates to households with duration above five and ten (in columns 2 and 3 respectively), we are estimating for k above this threshold. We say approximately in this range because the interquartile range of year-to-maturity in this sample is 15 to 24 years. Many models suggest that iMPCs are similar at large horizons, which will mitigate the quantitative importance of this averaging and approximation. When we compare model and data, we will compute a similarly averaged iMPC.

A second source of averaging occurs from the fact that we estimate consumption responds 1

[24]Linearity is a natural starting point to maximise power, but we relax this constraint in on-going work. Since theory puts restrictions on the time-path and sum of iMPCs, non-linear least squares is a natural avenue to estimate how the iMPC changes as k grows.

Table 10: Implied iMPCs

	iMPCs
Baseline Estimates	
Average $\overline{iMPC}_2^{15,24}$	0.008
Heterogeneity by Liquid Assets	
Average $\overline{iMPC}_2^{15,24}$ (below median)	0.001
Average $\overline{iMPC}_2^{15,24}$ (above median)	0.013

Notes: The table reports implied iMPCs from the linear estimates in Table 8 (top panel) and Table 9 (bottom panel). The implied iMPC is the (average) response of consumption at period 2 in response to news of income in periods 15 to 24.

to 3 years after shock occurred.[25] Taking the average of these, we will refer to our estimate of the iMPC at year 2 to a shock realized at between years $(15, 24)$. We test in the model how sensitive results are to this averaging.

5.5 Heterogeneity by savings account type

Table 11 reports the total savings response for split samples based on accounts held. In addition to providing information about how households savings, this is a robustness test of whether missing savings held in alternative accounts. Column (1) reports the savings response for everyone in our sample, column (2) is for those who have a checking account, column (3) is for those who have a savings account. These are overlapping samples as most savings account holders also hold a checking account at our bank. We report the total savings response for these holders of these account types, so column (3) reports the savings in both checking and savings accounts as a fraction of the income shock for the sample of savings account holders.

For households with a savings account, the average propensity to save is 1.6 percent. This is not statistically significant, however the standard error is small (0.05) suggesting that it is estimated with reasonable precision. One limitation of our data is that we only observe savings within a single institution, therefore one could be concerned that our estimates of a small savings responses are missing savings in other institutions. To test this, we look at households in column 3 for whom we observe the savings (and the vast majority of whom we also observe the checking account). That the savings response is *smaller* for this group, reassuringly suggests that most savings occurs in the checking account. Thus, if we observed the savings accounts for households in other banks,

[25] Recall that we miss the on-impact effect of the shock which is date-0.

Table 11: The Marginal Propensity to Save, by Account Types

	Savings				Log Savings			
	(1) Total	(2) Checking a/c	(3) Savings a/c	(4) Mortage prepay	(5) Total	(6) Checking a/c	(7) Savings a/c	(8) Mortage prepay
D.Payment Savings	0.086***	0.095***	0.016	0.045				
	(0.0219)	(0.0233)	(0.0519)	(0.1939)				
D.Log Payment Savings					0.059**	0.059**	0.068**	-0.135
					(0.0217)	(0.0226)	(0.0298)	(0.1363)
Observations	134563	122720	37531	144914	134563	122720	37531	144144
Adjusted R^2	0.002	0.001	0.005	0.001	0.002	0.003	-0.002	0.058
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prob($\beta = 1$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Note: This table reports estimates from regression equation 9, for different types of savings account types. In the first three columns, the dependent variable is the change in household i's bank account balance ("Savings"), $b_{i,t+1}$. The samples are split by whether we observe a checking or savings account: column 1 includes households with either account; column 2 includes only households with a checking account; and column 3 includes households with only a savings account. In column 4 the dependent variable is the change in mortgage prepayment. For all four columns, the independent variable is the change in the stock of payment savings $m^{stock}_{i,t+1}$. Variables are in euros. Controls include current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, deciles of total outstanding balance as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered by quarter. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*\quad p < 0.1$, $**\quad p < 0.05$, $***\quad p < 0.01$.

we might not estimate a much larger savings propensity.

Another avenue for savings that is missing in our main estimates is the reduction of debt, for instance by prepaying mortgage debt. Column (4) reports the fraction of the income shock that is used to *voluntarily* prepay mortgages. There is a *mechanical* reduction in the outstanding balance for tracker mortgages, that we exclude from this measure of prepayment. Since mortgage payments based on the annuity formula and fixed over time, lower interest rates imply a fast payment of the principle. Since we want to know the MPC out of the disposable income arising from reductions in payments, our mortgage prepayment measure only includes additional voluntary prepayment and not this mechanical balance reduction.

We estimate that 4.5 percent of payment savings are used to pre-pay mortgage debt, though the estimate is not statistically significant. While the point estimate suggests that our average savings response could be about 50 percent higher, this does not change our conclusion that the average MPC very high. This implies that we are not missing a large amount of savings via reduction in mortgage balances. Combined, these results suggest that the MPC is indeed high, as reported in our main estimate. There is little evidence that we are missing an economically meaningful amount of savings for households for whom we do not observe savings accounts, nor are households savings by pre-paying mortgage debt.

5.6 Robustness

Reversing the sign of the shock. Our identification comes from a relative change in mortgage payments between our tracker and variable rate mortgage holders. We have referred to the tracker group as the treated group and constructed the shock as a positive disposable income shock for this group. In theory, at least in the case of a simple difference-in-difference estimation, our results should be the same (with the sign changed) if we switch the definition of which groups is treated. In practice, the results are likely to differ because we use regressions to adjust for covariates and our treatment variable is continuous in intensity. Both of these features may alter the underlying weights of observations in the regressions (for example, Fagereng et al. (2021) show that observations who receive bigger shocks will receive greater weight).

Table B3 reports the results when we reverse the treatment group, so the variable rate group are treated and the tracker rate group are controls. We continue to define the payment savings as a positive number for the treated group and zero for the untreated group. It therefore should be interpreted as the *increase in payments* for the variable rate group. Based on specifications in level and level differences, we estimate that households reduce their savings by 5.8-9.9 dollars in response to a 100 dollar increase in payments, and our benchmark estimate of 8.6 percent result lies within this range (with the sign reversed).

Including defaulters in the sample. Our primary sample excludes households who default on their mortgage during or before the start of our sample, where default is defined as being more than 90 days behind on scheduled mortgage payments. We prefer to exclude these households, because there is no payment difference between them while they are not making payments. To ensure that this sample selection is not an important driver of our results, Table B4 reports the savings response while including the group of households who default during our sample. Based on our preferred specification, in column 7, households in this sample save 5.7 percent of the payment savings. It is not too surprising that the estimate is closer to zero because we are including households in the tracker and variable group who are in fact closer to each other because neither is making payments. While the estimate in this sample is lower than our main estimate of 8.6 percent, it reinforces our conclusion that the savings response is close to zero.

Households balance sheets across the liquid asset distribution. In section 3 we showed that deposits represent 78 percent of financial assets for the average household. However, our estimates of savings across the liquid asset distribution could be biased if some quartiles were more likely to hold savings in other financial assets, such as stocks or bonds. To test this, Table B6 reports deposits as a share of financial assets for each quartile of liquid assets. Reassuringly, there are minimal differences in the share of deposits across the liquid asset distribution, with each

quartile holding between 76 and 79 percent of their financial assets in deposits.

6 Consumption-Savings Model

In this section we study a benchmark incomplete markets consumption-savings model and evaluate its ability to match our estimates.

6.1 Model Description

Infinitely lived households choose consumption c and savings a to maximize the present discounted value of present discounted value of utility $\mathbf{E_0}\left[\sum_{t=0}^{\infty}\beta^t u(c_t)\right]$, where u has the constant relative risk aversion form, $\frac{c_t^{1-\sigma}}{1-\sigma}$. Endowments follow an AR(1) process in logs, $\ln e_t = \rho_e \ln e_{t-1} + \epsilon_t$, which we approximate with n_e idiosyncratic states that have a probability $\mathcal{P}(e,e')$ of transitioning between e and e'. Households can save in a risk free asset at a constant interest rate r and face a borrowing constraint $a \geq \underline{a}$. We model our income shock using an exogenous sequence of lump-sum transfers or taxes $\{\tau_s\}$ over which households have perfect foresight.

The value function of a household in state (e,a) at time t is therefore

$$\begin{aligned} V_t(e,a) = \max_{c,a'} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \beta \sum_{e'} V_{t+1}(e',a')\mathcal{P}(e,e') \right\} \\ c + a' = (1+r)a + e + \tau_t \\ a' \geq \underline{a}. \end{aligned} \tag{13}$$

We purposefully keep the model as simple as possible to evaluate the estimated responses. It incorporates heterogeneity in bank balances via the endogenous distribution of assets a and can incorporate heterogeneity in the length of the income shock via the sequence $\{\tau_s\}$. The model nests the representative-agent permanent income case when $n_e = 1$ and $\underline{a} = -\infty$.

Denote by $c_t^*(e,a;\boldsymbol{\tau})$ and $a_t^*(e,a;\boldsymbol{\tau})$ the policy functions that solve the Bellman equation (13). Here we make is explicit that these depend on the idiosyncratic state of the household (e,a) as well as the sequence of transfers $\boldsymbol{\tau}$.We also denote by $D_t(e,\mathcal{A}) \equiv Pr(e_t = e, a_t \in \mathcal{A})$ the measure of households in state e that have savings in a set $\mathcal{A}$ at the start of date t. This distribution has a standard law of motion $D_{t+1}(e',\mathcal{A}') = \sum_e D_t\left(e', a_t^{*-1}(e,\mathcal{A})\right)\mathcal{P}(e,e')$.

6.2 Consumption Responses

In our experiments we compute the response of consumption to a perfect foresight path of constant transfers $\boldsymbol{\tau}$ and compare this to steady state consumption when there are no transfers $\tau_{ss} = 0$.

$$mpc_t(e, a; \boldsymbol{\tau}) \equiv \frac{dc_t(e, a; \boldsymbol{\tau})}{d\tau} = \frac{[c_t^*(e, a; \boldsymbol{\tau}) - c_{ss}^*(e, a, \mathbf{0})]}{\tau}. \tag{14}$$

Note that since our shock is persistent, our definition of the *mpc* is the total change in consumption at time t divided by the total change in the transfer τ. This is different to the usual definition as the partial derivative of consumption with respect to income today $\frac{\partial c_t(e,a)}{\partial \tau_t}$ (e.g. Auclert, Rognlie and Straub (2023)). When the shock is constant, our definition corresponds to the sum of the partial derivatives of consumption with respect to income at every future period

$$\frac{dc_t(e, a; \boldsymbol{\tau})}{d\tau} = \sum_{s=t}^{t+T} \frac{\partial c_t(e, a)}{\partial \tau_s}, \tag{15}$$

where T is the length of the sequence of transfers.[26] The definition $mpc_t(e, a; \boldsymbol{\tau})$ also highlights that the consumption response depends on the length of the sequence of transfers $\boldsymbol{\tau}$ as well as on the individual states (e, a). We can also aggregate the consumption response across consumers to give an average consumption $C_t(\boldsymbol{\tau}) = \sum_e \int_a c_t^*(e, a; \boldsymbol{\tau}) D_t(e, a)$ and the average MPC

$$MPC_t(\boldsymbol{\tau}) = [C_t(\boldsymbol{\tau}) - C_{ss}(\mathbf{0})] / \tau. \tag{16}$$

To be consistent with our empirical estimates, we define the average savings response MPS_t as the cumulative savings relative to the total transfer received up to that point in time

$$MPS_t(\boldsymbol{\tau}) = [A_t(\boldsymbol{\tau}) - A_{ss}(0)] / (\sum_{s=0}^{t} \tau_s). \tag{17}$$

Since savings are a stock, this reports the average savings of the cumulative shock and is thus consistent with how we estimate the empirical savings response.

[26] In the case of non-constant transfers, the consumption response is

$$dc_t(e, a; \boldsymbol{\tau}) = \sum_{s=t}^{\infty} \frac{\partial c_t(e, a)}{\partial \tau_s} d\tau_s.$$

6.3 Model Experiment

Our model experiment considers the response to a perfect foresight path to transfers $\{\tau_s\}_{s\geq 0}$.This formulation implicitly assumes that, at the time at which it occurred, households perfectly anticipated the length of the shock. While there was naturally some uncertainty about how long the spread between the two interest rates would last, it is reasonable for households to have expected that it was very persistent.[27] We approximate the path of the empirical shock, as shown in Figure 4, with a sequence where the size of the transfer is the same in every period. For our benchmark experiment, the length of the shock is 68, which matches the average length of the shock in our data. Figure 5 illustrates this, showing both the flow transfers and the cumulative ("stock") transfers received at each date.

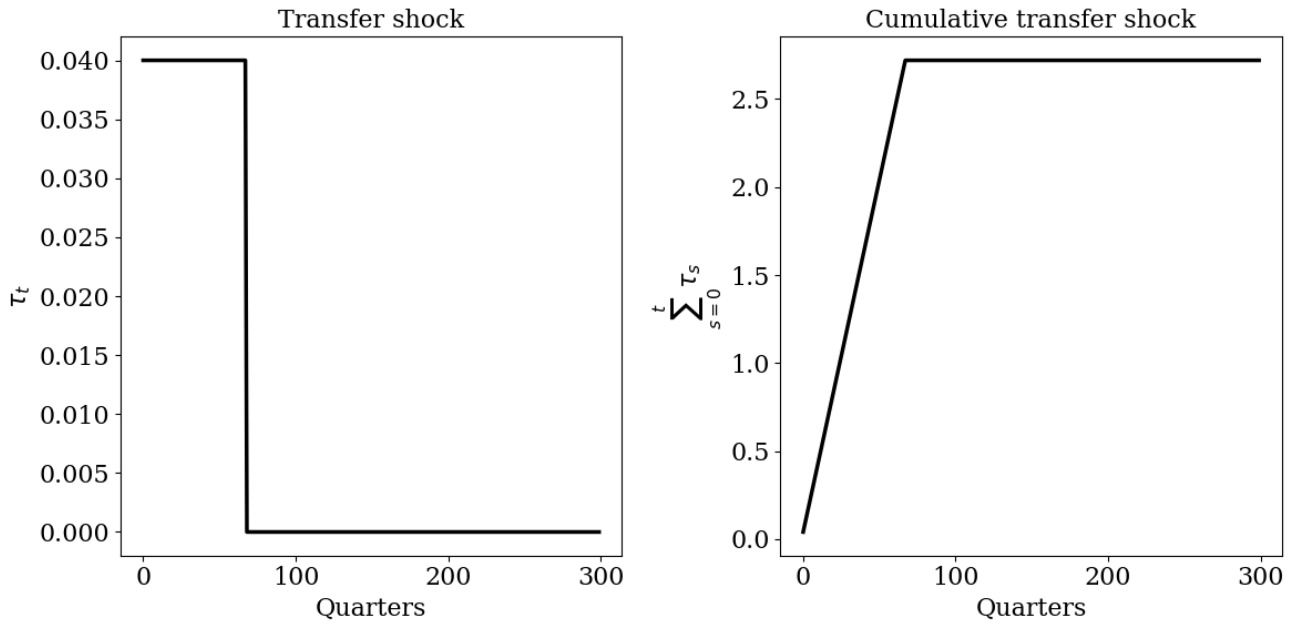


Figure 5: Transfer Shock in the Benchmark Model Experiment

Note: Figure reports the transfer shock that is the benchmark model experiment. The left panel shows the per period shock, which is the same size in every period until the 68th period. The right panel shows the cumulative shock, which is the sum of the per period shocks.

6.4 Model Quantification

Table 12 reports the parameter values that we use to quantify the model. The model is calibrated to a quarterly frequency, which is the same frequency of our data. We set the inverse elasticity of

[27]The perfect foresight assumption is also an implicit assumption in any macro model that makes a first order approximation (Auclert, Bardoczy, Rognlie and Straub, 2021).

intertemporal substitution to 2 and the real interest rate to 0.01, as is standard in the literature (Auclert, Bardoczy, Rognlie and Straub, 2021) and is consistent with the low real interest rates on savings in the data during this time. In the absence of an estimate for the persistence of income shocks in Ireland, we follow the same procedure as de Ferra, Mitman and Romei (2020). We use an estimate of the persistence from the US of 0.966 (Auclert, Bardoczy, Rognlie and Straub, 2021) and estimate the variance of log residual earnings σ^2_{data} in Ireland using the second wave (2014) of the ECB's Household Finance and Consumption Survey (HFCS). With these two parameters we can compute the variance the income shock, $\sigma^2_\epsilon = \sigma^2_{data} \times (1 - \rho^2_\epsilon) = 0.44 \times 0.066 = 0.029$ (Krueger, Mitman and Perri, 2016). Since our model is non-linearly, the size of the transfer shock that we feed in could affect the consumption response. We set the size of the transfer shock τ to 0.05 to match the average size of the payment shock relative to incomes.

Table 12: Model Quantification

Parameter	**Description**	**Value**	**Source**
σ	Inverse Elasticity of Intertemporal Substitution	2	Auclert et al. (2021)
r	Real Interest Rate	0.01	Auclert et al. (2021)
ρ_e	Persistence of Income Risk	0.966	Auclert et al. (2021)
σ^2_e	Variance of Income Risk	0.029	HFCS (2014)
τ	Transfer Shock Size	0.05	Internal Estimation
β	Discount Factor	0.97	Internal Estimation

Note: Table reports the calibration of the model parameters. The model is calibrated to a quarterly frequency. We use a standard calibration of an incomplete markets model with idiosyncratic income risk, following Auclert, Bardoczy, Rognlie and Straub (2021). We calibrate the size of the transfer shock τ using the average size of the payment shock relative to income the data. We estimate the discount factor β to match the average MPC in the data. We estimate the variance of income risk using the second wave (2014) of the ECB's Household Finance and Consumption Survey (HFCS) for Ireland.

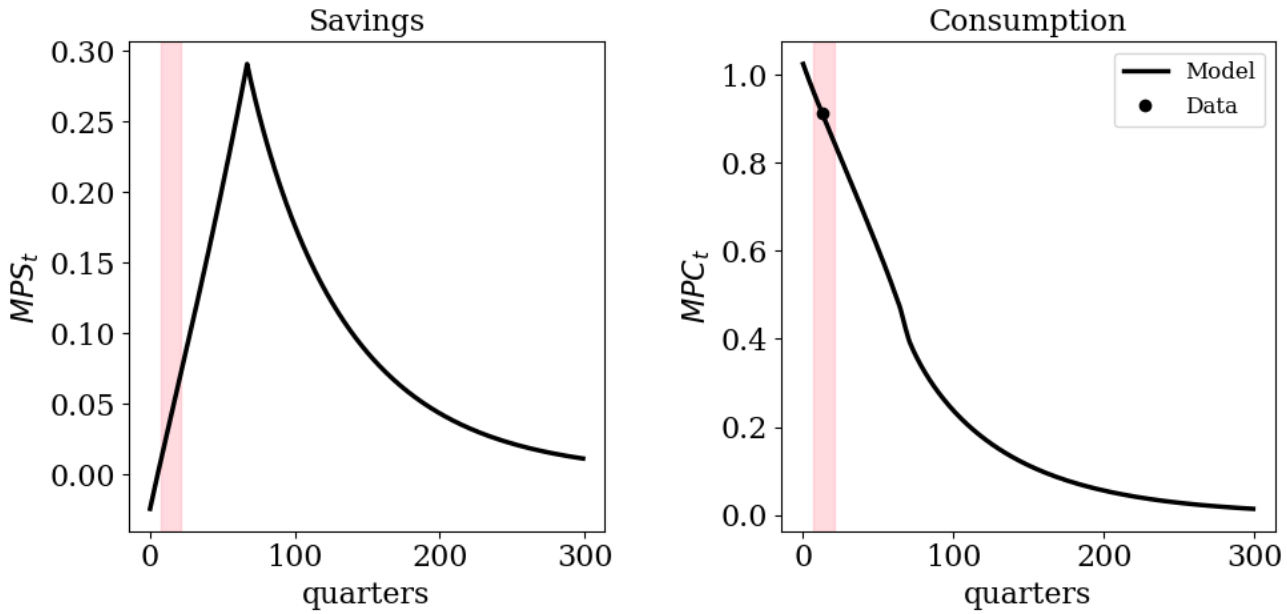
We estimate the discount factor to match the average MPC in the data. This involves finding a single parameter to match a single moment and therefore we can exactly match the MPC in the data (0.91) with a discount factor $\beta = 0.986$. Alternatively, we could have used a standard value of the discount factor and then evaluated the fit of the model based on the discrepancy between the model and data's average MPC. This approach yields and similar conclusion about model fit, since our estimated discount factor $\beta = 0.986$ is remarkably close the standard value $\beta = 0.982$ used by Auclert, Bardoczy, Rognlie and Straub (2021) in their calibration of a Krussel-Smith economy[28]. We consider this evidence that our empirical estimates are consistent with a standard calibration of a consumption-savings model. Targeting the average MPC better allows us to evaluate the

[28]The average MPC of our model with the discount factor of Auclert, Bardoczy, Rognlie and Straub (2021) is 1.0. This is higher than our MPC of 0.91, but still represents a good fit relative to the large differences between model and data for transitory shocks.

fit of the model using the heterogeneous responses. Otherwise, we might be concerned that any mismatch on the heterogeneous responses was affected by a mismatch on the average MPC.

6.5 Results

Figure 6 plots the response of savings (in the left panel) and consumption (in the right panel) to our benchmark 69 quarter transfer shock. The line shows the model's simulated response and the dot is our preferred estimate from Table 5. The model is simulated for more periods than we observe in the data, so the pink shaded area highlights period during which the model and data overlap. Households consume the vast majority of the shock; they spend 95 percent of the transfer on impact and 91 percent during the pink shaded period. Even though households spend the majority of the shock, they also save to smooth the increase in consumption over time. Savings rise until the end of the transfer shock period and then fall back to their steady state level. At the end of the sequence of transfers, household have saved 23 percent of the *cumulative* transfer, which enables consumption to remain above its steady state level after the shock has stopped. Indeed, in the quarter that the shock ends, consumption falls from consuming 51 percent to 49 percent of the flow transfer.



(a) Savings Response

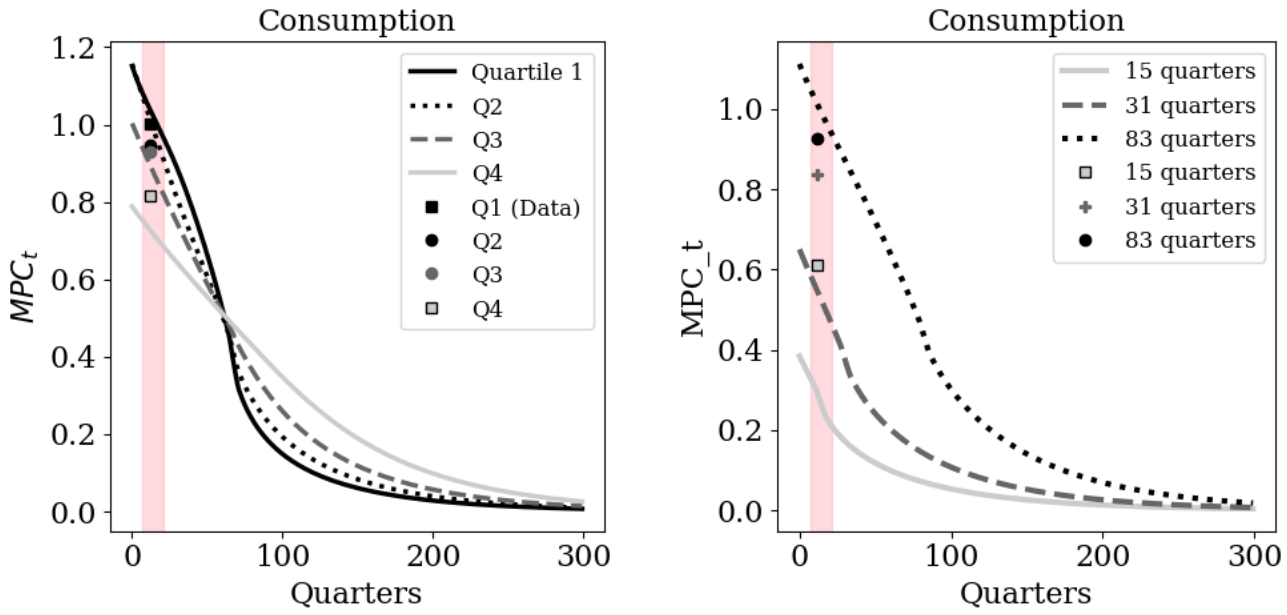
(b) Consumption Response

Figure 6: Savings and Consumption Responses to the Benchmark Model Experiment

Note: Figure reports the average savings and consumption responses to the benchmark model experiment. The line is the model's simulated path. The dot is the data estimate. The gray shaded area is the time period over which the model and data overlap. The sequence of transfers is reported in Figure 5.

Figure 6 shows that the model exactly matches the average consumption response during the period when the model and data overlap. As noted in the quantification, while this is a targeted moment, the reasonable value of the discount factor suggests that the model is not too far from the data. Next we consider the heterogeneous responses.

Figure 7 plots the heterogeneous responses of consumption. The left panel (7a) plots the response of households split by quartiles of liquid assets in response to the benchmark shock length. The right panel (7b) plots the average response of households for different shock lengths. As in the previous figure, the lines are model outputs, the dots are data estimates from section 5, and the pink shaded area is the period of overlap between the model and data.



(a) By Quartile of Liquid Asset Balances

(b) By Persistence of the Shock

Figure 7: Heterogeneous Consumption Response, Model Versus Data

Note: Figures reports the heterogeneous consumption responses to the transfer shock. The left plots consumption responses by liquid assets under the benchmark experiment. The right plots the average consumption responses for three shocks that differ in their persistence. In each plot the lines are the model's simulated path and the dots are the data estimates. The light pink shaded areas are the time period over which the model and data overlap.

The left panel (7a) shows that model fits the heterogeneous response by liquid assets well. Table 13 reports the model and data consumption responses, and column (3) reports the difference between the model and data for each quartile. The model's response varies between 0.72 for those with the most liquid assets to 1.03 for those with the least liquid assets. Like we seen in the data, the model MPC is monotonically decreasing as liquid assets increase. Three of quartiles are within 5 percentage points of the data estimate and the remaining quartile is within ten percentage points.

We consider this a good fit. The only feature of the model that generates this heterogeneity in consumption response is the presence of liquidity constraints. In the presence of liquidity constraints, households who receive negative shocks and approach the borrowing constraint are not able to borrow to sustain their consumption. Additional income thus alleviates the constraint for these households with low liquid assets but not for those with higher liquid assets. And this results in a larger consumption response for those with low liquid assets. Therefore, the fact that the model fits this form of heterogeneity well is evidence in favor of the theory that liquidity constraints are important for the consumption response to shocks. While it is well known that liquidity constraints can generate larger responses for households with lower liquid assets, our results provide novel theoretical and empirical evidence that liquidity constraints are quantitatively important even for highly persistent shocks.

Table 13: Heterogeneous Consumption Response, Model Versus Data

	(1) **Data MPC**	(2) **Model MPC**	(3) **Difference**
Panel A: Heterogeneity By Liquid Assets			
Highest Quartile	1.001	1.026	-0.025
3rd Quartile	0.948	0.995	-0.047
2nd Quartile	0.93	0.882	0.048
Lowest Quartile	0.816	0.72	0.096
Panel B: Heterogeneity By Shock Persistence			
Most Persistent	0.92	1.0	-0.07
Middle	0.84	0.53	0.31
Least Persistent	0.61	0.27	0.34

Note: Table reports the average MPC in the data (column 1) and model (column 2), and the difference between them (column 3). Estimates are reported for three groups of households based on the persistence of the shock. The data estimates are from Tables 7 and 8.

The right panel (7b) shows the model fits the heterogeneous response by the length of the shock less well. The model and data have the same covariance between the length of the shock and the consumption response. More persistent shocks have larger consumption responses, as predicted by the permanent income hypothesis. For the most persistent shock, the model's consumption response is only 6 percentage points higher than the empirical estimate, which is perhaps unsurprising given that the model matches the average response and this shock lasts only 13 more quarters. However, the model performs less well for the less persistent shocks. Table 13 reports that the model's response is 34 percentage points lower than the empirical estimate of 0.61 for the least persistent shock of 15 quarters. The model's response is also 31 percentage points under the empirical estimates for the moderately persistent shock (31 quarters). Taken together, these

estimates suggest that the model predicts too much variation in the consumption response by the length of the shock. While the covariance suggest that the persistence of the shock does indeed predicts the size of the consumption response, the model under-performs quantitatively. This result that the model performs worst for the least persistent shocks is consistent with the prior literature (Parker, Souleles, Johnson and McClelland, 2013; Fagereng, Holm and Natvik, 2021) that finds that the standard model's consumption response is too small in the case of one-time shocks.

6.6 iMPCs

Finally, we compare the model's average iMPCs, computed using the same differencing procedure as in Section 5 in Table 14. The iMPC in the model is more than twice as high as in the data. Since the forward-looking iMPC captures how much consumption today responds to anticipated future income—and thus governs the extent of Ricardian non-equivalence—these estimates suggest that households are less forward-looking than the standard model predicts. On-going work continues to quantify the importance of this discrepancy.

Table 14: Implied iMPCs

	iMPCs	
	Data	**Model**
Average $\overline{iMPC}_2^{15,24}$	0.008	0.023

Notes: The implied iMPC is the (average) response of consumption at period 2 in response to news of income in periods 15 to 24 (annually).

7 Conclusion

In conclusion, we use a natural experiment in Ireland to provide novel identification of persistent income shocks with varying degrees of persistence. Our findings provide valuable insights into household behavior in the face of long-lasting changes in disposable income. We find that households exhibit a relatively high average marginal propensity to consume when confronted with such persistent shocks.

Importantly, our results highlight a noteworthy contrast between persistent shocks and temporary income shocks. In the case of persistent shocks, our findings show that the standard consumption-savings model aligns well with the empirical evidence. The model's predictions regarding the MPC closely match the observed consumption responses, suggesting its relevance for

understanding household behavior in the presence of prolonged income fluctuations. This stands in stark contrast to temporary income shocks, where significant discrepancies exist between the standard model's predictions and empirical estimates of the MPC.

Furthermore, our analysis uncovers heterogeneity in consumption responses. Households with limited liquid assets display higher MPCs. Additionally, we find that the length of the shock's persistence positively correlates with the MPC, with households facing longer remaining mortgage maturities exhibiting higher consumption propensities.

By shedding light on the marginal propensity to save and consume in the context of persistent income shocks, our findings contribute to a better understanding of household behavior and its implications for macroeconomic dynamics. These insights underscore the importance of considering individual heterogeneity in macroeconomic models and policy analysis. We hope that our findings provide valuable guidance for policymakers and contribute to the ongoing dialogue surrounding macroeconomic stabilization measures.

References

Acharya, Viral V., Katharina Bergant, Matteo Crosignani, Tim Eisert, and Fergal McCann, "The Anatomy of the Transmission of Macroprudential Policies," *The Journal of Finance*, 2022, *77* (5), 2533–2575.

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

Aguiar, Mark and Erik Hurst, "Consumption versus Expenditures," *Journal of Political Economy*, 2005, *113* (5), 919–948.

_ **and** _ , "Life-Cycle Prices and Production," *American Economic Review*, 2007, *97* (5), 1533–1559.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, 1st ed., Princeton, NJ: Princeton University Press, 2009.

Auclert, Adrien, Bence Bardoczy, Matthew Rognlie, and Ludwig Straub, "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 2021, *89(5)*, 2375–2408.

_ **, Matthew Rognlie, and Ludwig Straub**, "The Intertemporal Keynesian Cross," *Journal of Political Economy*, 2023, *forthcoming.*

Baker, Scott R., "Debt and the Response to Household Income Shocks: Validation and Application of Linked Financial Account Data," *Journal of Political Economy*, 2018, *126* (4), 1504–1557.

Benetton, Matteo, "Leverage Regulation and Market Structure: A Structural Model of the UK Mortgage Market," *Journal of Finance*, 2021, *76* (6), 2997–3053.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, "Regional Heterogeneity and the Refinancing Channel of Monetary Policy," *The Quarterly Journal of Economics*, 2019, *134* (1), 109–183.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra, "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy," *American Economic Review*, September 2021, *111* (9), 2829–2878.

_ , **Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra**, "House Prices and Consumer Spending," *The Review of Economic Studies*, July 2018, *85* (3), 1502–1542.

Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg, "What Accounts for the Variation in Retirement Wealth Among U.S. Households?," *American Economic Review*, 2001, *91* (4), 832–857.

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

Bodkin, Ronald G., "Windfall Income and Consumption," *American Economic Review*, 1959, *49* (4), 602–614.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel, "Five Facts about MPCs: Evidence from a Randomized Experiment," *American Economic Review*, January 2025, *115* (1), 1–42.

Browning, Martin and Thomas F. Crossley, "Unemployment Insurance Benefit Levels and Consumption Changes," *Journal of Public Economics*, 2001, *80* (1), 1–23.

Byrne, David, Robert Kelly, and Conor O'Toole, "How does monetary policy pass-through affect mortgage default? Evidence from the Irish mortgage market," *Journal of Money, Credit and Banking*, 2021, *54 (7)*, 2081–2101.

Byrne, Shane, Kenneth Devine, Michael King, Yvonne McCarthy, and Christopher Palmer, "The Last Mile of Monetary Policy: Inattention, Reminders, and the Refinancing Channel," *NBER Working Paper*, 2022, (31043).

Campbell, John Y., "John Campbell on Mortgage Choice and Monetary Policy," *Princeton Bendheim Center for Finance, Youtube*, 2023.

Cochrane, John H., "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *American Economic Review*, June 1989, *79* (2), 319–337.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The How and Why of Household Reactions to Income Shocks," 2025, (32191).

Cookson, Anthony, Aymeric Bellon, Erik P. Gilje, and Rawley Z. Heimer, "Personal Wealth, Self-Employment, and Business Ownership," *Review of Financial Studies*, 2021, *34* (8), 3935–3975.

_ , **Erik P. Gilje, and Rawley Z. Heimer**, "Shale Shocked: Cash Windfalls and Household Debt Repayment," *Journal of Financial Economics*, 2022, *146* (3), 905–931.

de Ferra, Sergio, Kurt Mitman, and Federica Romei, "Household Heterogeneity and the Transmission of Foreign Shocks," *Journal of International Economics*, 2020, *124*.

Di Maggio, Marco, Amir Kermani, Benjamin J. Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao, "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging," *American Economic Review*, 2017, *107 (11)*, 3550–88.

Druedahl, Jeppe, Emil Bjerre Jensen, and Søren Leth-Petersen, "The Intertemporal Marginal Propensity to Consume out of Future Persistent Cash-Flows: Evidence from Transaction Data," *Working Paper*, 2021.

Eichenbaum, Martin S., Joao Guerreiro, and Jana Obradovic, "Ricardian Non-Equivalence," NBER Working Paper 34691 January 2026.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik, "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 2021, *13*, 1–54.

Flavin, Marjorie A., "The Adjustment of Consumption to Changing Expectations About Future Income," *Journal of Political Economy*, October 1981, *89* (5), 974–1009.

_ , "The Excess Smoothness of Consumption: Identification and Interpretation," *The Review of Economic Studies*, July 1993, *60* (3), 651–666.

Ganong, Peter and Pascal Noel, "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, 2019, *109* (7), 2383–2424.

_ **, Damon Jones, Pascal Noel, Diana Farrell, Fiona Greig, and Chris Wheat**, "Wealth, Race, and Consumption Smoothing of Typical Labor Income Shocks," *Working Paper*, 2023.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky, "How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income," *Quarterly Journal of Economics*, 2023, *forthcoming*.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 1997, *87* (1), 192–205.

Gupta, Arpit, "Foreclosure Contagion and the Neighborhood Spillover Effects of Mortgage Defaults," *Journal of Finance*, October 2019, *74* (5), 2249–2301.

Guren, Adam M., Alisdair McKay, Emi Nakamura, and Jon Steinsson, "Housing Wealth Effects: The Long View," *The Review of Economic Studies*, March 2021, *88* (2), 669–707.

Hall, Robert E and Frederic S Mishkin, "The sensitivity of consumption to transitory income: Estimates from panel data on households," *Econometrica: Journal of the Econometric Society*, 1982, *50* (2), 461–481.

Hayashi, Fumio, "The Effect of Liquidity Constraints on Consumption: A Cross-Sectional Analysis," *Quarterly Journal of Economics*, 1985, *100* (1), 183–206.

Hedlund, Aaron, Kieran Larkin, Kurt Mitman, and Serdar Ozkan, "Mortgage Market Structure and the Transmission of Monetary Policy During the Great Inflation," June 2025, (2025-016).

Higgins, Brian E., "Mortgage Borrowing Limits and House Prices: Evidence from a Policy Change in Ireland," *Working Paper*, 2024.

Honohan, Patrick, *Currency, Credit and Crisis: Central Banking in Ireland and Europe*, Cambridge University Press, 5 2019.

Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote, "Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players," *American Economic Review*, 2001, *91* (4), 778–794.

Irish Times, The, "Versatile tracker mortgages are worth a visit," 2003.

_ , "Targeting a tracker mortgage could pay dividends," 2004.

Jappelli, Tullio and Luigi Pistaferri, "The Consumption Response to Income Changes," *Annual Review of Economics*, 2010, *2*, p.479–509.

_ **and** _ , *The Economics of Consumption: Theory and Evidence*, Oxford University Press, 11 2017.

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

Kaplan, Greg and Giovanni L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82(4)*, 1199–1239.

_ **and** _ , "The Marginal Propensity to Consume in Heterogeneous Agent Models," *Annual Reviews of Economics*, 2022, *14*, 747–775.

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

Kelly, Robert, Paul Lyons, and Conor O'Toole, "Mortgage Interest Rate Types in Ireland," *Economic Letters, Central Bank of Ireland*, 2015, *2015* (9).

Kinnerud, Karin, "The Effects of Monetary Policy Through Housing and Mortgage Choices on Aggregate Demand," *Quantitative Economics*, 2025, *16* (2), 659–703.

Krueger, Dirk, Kurt Mitman, and Fabrizio Perri, "Macroeconomics and Heterogeneity," *Handbook of Macroeconomics*, 2016, *Volume 2*, 843–921.

Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers," *The Quarterly Journal of Economics*, 06 2018, *133* (4), 1693–1751.

Laibson, David, Peter Maxted, and Benjamin Moll, "A Simple Mapping from MPCs to MPXs," *Working Paper*, 2022.

Lane, Philip R., "The European Sovereign Debt Crisis," *Journal of Economic Perspectives*, 2012, *26* (3), 49–68.

_ , "The Funding of The Irish Domestic Banking System During the Boom," *Journal of Statistical and Social Inquiry Society of Ireland*, 2015.

McCarthy, Yvonne, "Dis-Entangling the Mortgage Arrears Crisis: The Role of the Labour Market, Income Volatility and Housing Equity," *Research Technical Papers, Central Bank of Ireland*, 2014, *2014* (2).

_ **and Kieran McQuinn**, "Credit Conditions in a Boom and Bust Property Market," *Research Technical Papers, Central Bank of Ireland*, 2013, *2013* (8).

Nakamura, Emi and Jón Steinsson, "Identification in Macroeconomics," *Journal of Economic Perspectives*, 2018, *32* (3), 59–86.

O'Malley, Terry, "The Impact of Repossession Risk on Mortgage Default," *The Journal of Finance*, 2021, *76* (2), 623–650.

Orchard, Jacob D., Valerie A. Ramey, and Johannes F. Wieland, "Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates," *The Quarterly Journal of Economics*, August 2025, *140* (3), 2001–2052.

Parker, Jonathan A., "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes," *American Economic Review*, 1999, *89* (4), 959–973.

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

Pilossoph, Laura, Daniel Lewis, and Davide Melcangi, "Latent Heterogeneity in the Marginal Propensity to Consume," 2022.

Pistaferri, Luigi, "Superior information, income shocks, and the permanent income hypothesis," *Review of Economics and Statistics*, 2001, *83* (3), 465–476.

Robles-Garcia, Claudia, "Competition and Incentives in Mortgage Markets: The Role of Brokers," *Working Paper*, 2020.

Romer, David, *Advanced Macroeconomics*, 5 ed., McGraw-Hill Education, 2019.

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

Wolf, Christian, "The Missing Intercept: A Demand Equivalence Approach," *American Economic Review*, 2023, *113* (8), 2232–2269.

Wolf, Christian K., "Interest Rate Cuts versus Stimulus Payments: An Equivalence Result," *Journal of Political Economy*, April 2025, *133* (4), 1235–1275.

Wong, Arlene, "Refinancing and the Transmission of Monetary Policy to Consumption," 2021. Working paper.

Zator, Michal, "Working More to Pay the Mortgage: Household Debt, Interest Rates and Labor Supply," *Working Paper*, 2022.

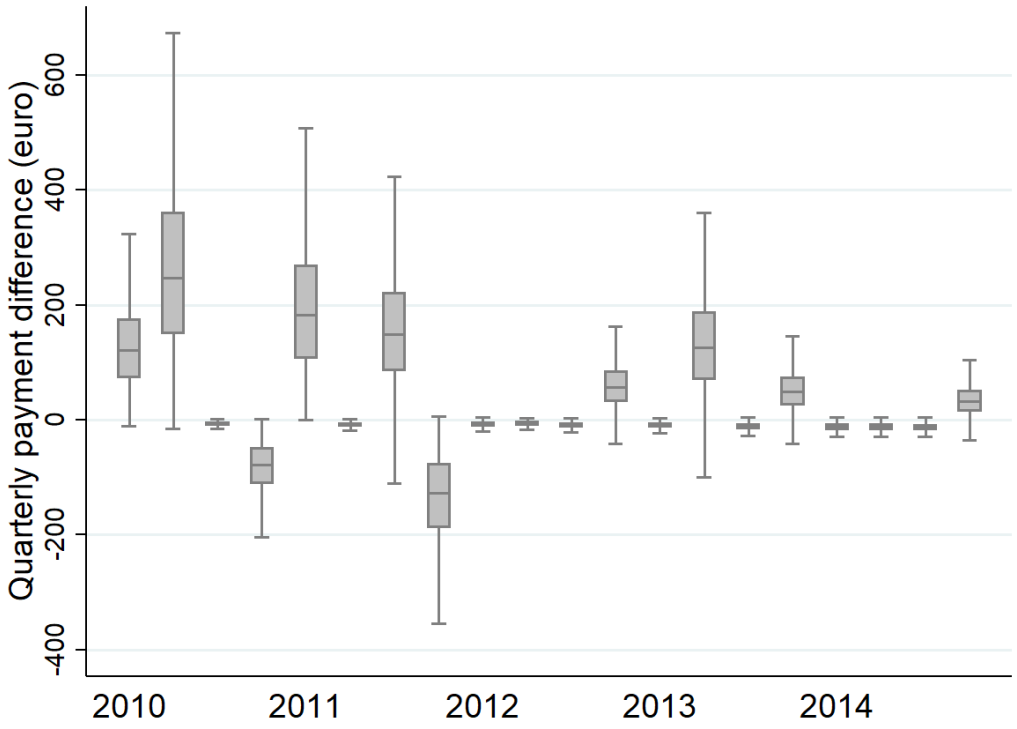
Zeldes, Stephen P., "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 1989, *97* (2), 305–346.

For Online Publication

A Appendix to Section 3: Data

Figure A1: Box plot Change in (Flow) Payment Savings, Quarterly 2011Q2 - 2015Q4

(a)



Note: Plot shows the change in the payment pavings at a quarterly frequency. Payment savings are in euros.

Table A1 presents the results of sequentially adding variables to a survey of mortgage holders post-shock, in order to investigate which controls have the biggest impact on the outcome variable. The table includes seven columns, each corresponding to a different specification with additional controls. The first column presents the results of the baseline regression, which includes only the treatment variable and no controls. The second column adds bank fixed effects, while the third column adds the variable for the origin year of the mortgage. The table shows that controlling for origin year and initial income of the borrower matter the most for explaining differences in income changes for tracker versus variable rate holders. This is intuitive because income shocks after the financial crisis did not occur uniform across the income distribution. However, once controlling for the ex ante differences in income, there are small and statistically insignificant differences in income changes (2.8 percent) between tracker and variable rate borrowers. This suggests that

our assumption that the tracker and variable rate borrowers are as good as randomly assigned conditional on observables is a reasonable assumption.

Table A1: Exposure of Tracker Holders to Ex-Post Income Changes

	Income Change						
	(1) None	(2) Bank	(3) Origin year	(4) Age	(5) Income	(6) LTV	(7) First time buyer
Tracker	-0.254***	-0.233**	-0.129	-0.132	0.034	0.028	0.028
	(0.0683)	(0.0723)	(0.0774)	(0.0795)	(0.0607)	(0.0609)	(0.0610)
Observations	616	616	616	597	597	595	595
Adjusted R^2	0.020	0.026	0.044	0.063	0.463	0.464	0.463
Origin year and bank FE		Yes	Yes	Yes	Yes	Yes	Yes
Controls			Yes	Yes	Yes	Yes	Yes

Note: This table reports the same regression as (column 3) of Table 4, which tests whether tracker mortgage holders were more likely to experience changes in income between mortgage origination and the survey data. The columns report the results of sequentially adding fixed effects and controls, with column 7 corresponding to main specifiation. Source: Survey of Mortgage Holders 2012-13 and Mortgage Dataset, Central Bank of Ireland.

Table A2 reports the estimated coefficients and standard errors for the effect of education on the outcome variable in a survey of mortgage holders post-shock. The table includes three variables, each corresponds to a different level of education: primary, secondary, and higher. The estimated coefficients represent the change in the outcome variable associated with each level of education, relative to the omitted category (primary). The table shows that the results are not statistically significant, indicating that there is little evidence that tracker borrowers have better education than non-tracker borrowers. Column one shows that tracker borrowers are more likely to have a third-level higher degree. However, controlling for the year and bank reverses this association. This suggests that there may be differences in the education of the types of borrowers who borrowed at different points in the boom, and that tracker mortgages were also more likely to be sold in later years (since it was a new product). That said, once we control for these associations, there is no remaining association between those who took tracker mortgages and their education status.

Table A2: Ex-Ante Selection by Education Status

	Tracker		
	(1)	(2)	(3)
Primary	0.000	0.000	0.000
	(.)	(.)	(.)
Secondary	-0.045	-0.088	-0.135
	(0.0772)	(0.0687)	(0.0700)
Higher	0.097	-0.028	-0.094
	(0.0675)	(0.0604)	(0.0640)
Observations	605	605	593
Adjusted R^2	0.011	0.234	0.252
Origin year and bank FE		Yes	Yes
Controls			Yes

Note: Table reports a linear probability model using educational attainment to predict whether a households is a tracker (versus variable) rate mortgage holder. Column 1 includes no controls, column 2 include bank and year fixed effects, and column 3 additionally includes controls for age, income, LTV and first-time buyer status. Source: Survey of Mortgage Holders 2012-13 and Mortgage Dataset, Central Bank of Ireland.

B Appendix to Section 5: Empirical results

B.1 Heterogeneity in levels

Table B1: The Marginal Propensity to Save out of payment shocks, by checking and savings account types

	Savings				Log Savings			
	(1) Total	(2) Checking a/c	(3) Savings a/c	(4) Mortage prepay	(5) Total	(6) Checking a/c	(7) Savings a/c	(8) Mortage prepay
Payment Savings	0.077***	0.088***	0.019	-0.027				
	(0.0242)	(0.0262)	(0.0504)	(0.0374)				
Log Payment Savings					0.076**	0.077	0.102*	0.107
					(0.0326)	(.)	(0.0543)	(0.1761)
Observations	144914	132160	40418	144914	144914	132160	40418	144301
Adjusted R^2	0.903	0.893	0.929	0.203	0.907	0.890	0.958	0.677
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prob($\beta = 1$)	0.000	0.000	0.000	0.000	0.000	.	0.000	0.000

Note: Controls include current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, total outstanding balance as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered at the level of the fixed effects.

Table B2: The Marginal Propensity to Save out of payment shocks, by quartiles of the balance distribution

	Savings				Log Savings			
	(1) Lower	(2) Quartile 2	(3) Quartile 3	(4) Upper	(5) Lower	(6) Quartile 2	(7) Quartile 3	(8) Upper
Payment Savings	-0.007	0.044**	0.038	0.180**				
	(0.0073)	(0.0209)	(0.0439)	(0.0702)				
Log Payment Savings					0.018	0.156***	0.138**	0.085*
					(0.0518)	(0.0583)	(0.0674)	(0.0437)
Observations	38584	33446	34468	38416	38584	33446	34468	38416
Adjusted R^2	0.506	0.413	0.471	0.854	0.742	0.521	0.563	0.756
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Prob($\beta = 1$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: This table reports estimates from regression equation 11 where the sample is split by quartiles of the bank balance distribution. The dependent variable is household i's bank account balance ("Savings"), $b_{i,t}$. The independent variable is household i's stock of payment savings $m_{i,t}^{stock}$. Variable names indicate euro and log specifications. Controls are current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, deciles of total outstanding balance as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses and are double clustered by quarter and individual. Source: Merged Bank-Mortgage Dataset, Central Bank of Ireland. $*\quad p < 0.1$, $**\quad p < 0.05$, $***\quad p < 0.01$.

B.2 Negative shock

In the main paper, the shock is defined as a positive payment saving for tracker mortgage holders and zero savings for variable rate mortgage holders. However, we can also define the shock as a payment shock for variable rate mortgage holders and a zero shock for tracker mortgage holders. The payment difference is still defined as a positive number, so that we can take logs, and it represents the *increase* in the payments for the variable rate group. In this case we expect that it would be associated with a decline in savings. Table B3 reports the results for this specification and as expected the coefficients are negative. The coefficient is a similar magnitude to the main specification, with a savings decline of 6 percent of the payment difference (column 3). The implied MPC is also high at 94 percent, implying that even if we think of the shock as increasing the mortgage payments for variable rate holders, most of the response comes from consumption rather than savings.

Table B3: The Marginal Propensity to Save out of payment shocks, negative shock

	(1) Savings	(2) Log Savings	(3) Δ Savings	(4) Δ Log Savings	(5) Savings	(6) Log Savings	(7) Δ Savings	(8) Δ Log Savings
Cumulative Payment Difference	-0.063				-0.099***			
	(.)				(0.0221)			
Log Cumulative Payment Difference		-0.039**				-0.037**		
		(0.0156)				(0.0152)		
D.Cumulative Payment Difference			-0.058				-0.083*	
			(0.0397)				(0.0452)	
D.Log Cumulative Payment Difference				0.003				0.003
				(0.0111)				(0.0101)
Observations	144914	144845	134563	134493	144848	144779	134502	134432
Adjusted R^2	0.902	0.907	0.001	0.001	0.903	0.907	0.002	0.002
Individual FE	Yes	Yes			Yes	Yes		
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE					Yes	Yes	Yes	Yes
Prob($\beta = 1$)	.	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Controls include current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, total outstanding balance as of 2008 Q4, loan to income as of 2008 Q4, loan to value as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered at the level of the fixed effects.

B.3 Defaulter sample

The main sample excludes households who default before or during our sample, because the concept of payment savings does not apply to households who are not making payments. However, we still include households who we know default at a point in time after our sample, because the payment savings apply up until they stop paying. For transparency, we also consider the marginal propensity to save out of payment shocks for an extended sample of households that includes those who default during our sample. Table B4 reports the results for this sample. The results are similar to the main sample, with a savings response of 6 percent (column 3). This response is lower than the savings response in the main sample and is consistent with the idea that there should be minimal differences in savings response for households who have no difference in payment savings, which is the case when neither variable rate nor tracker rate defaulters are making payments.

Table B4: The Marginal Propensity to Save out of payment shocks, including sample of defaulters

	(1) Savings	(2) Log Savings	(3) Δ Savings	(4) Δ Log Savings	(5) Savings	(6) Log Savings	(7) Δ Savings	(8) Δ Log Savings
Cumulative Payment Difference	0.059*				0.060*			
	(0.0230)				(0.0253)			
Log Cumulative Payment Difference		0.077*				0.092*		
		(0.0278)				(0.0332)		
D.Cumulative Payment Difference			0.057*				0.057*	
			(0.0234)				(0.0244)	
D.Log Cumulative Payment Difference				0.059**				0.061*
				(0.0187)				(0.0243)
Observations	151228	151213	140426	140411	151162	151147	140365	140350
Adjusted R^2	0.903	0.906	0.001	0.001	0.903	0.905	0.002	0.002
Individual FE	Yes	Yes			Yes	Yes		
Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls × Quarter FE					Yes	Yes	Yes	Yes
Prob($\beta = 1$)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000

Note: Controls include current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, total outstanding balance as of 2008 Q4, loan to income as of 2008 Q4, loan to value as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered at the level of the fixed effects.

B.4 Differences at the start of the sample and over the entire sample, i.e. the long differences

Table B5 reports the savings response to a payment shock at the start of the sample and over the entire sample. The table includes two sets of results, each corresponding to a different time period: the start of the sample and the entire sample. The table shows that at the start of the sample, there is a small and statistically insignificant savings response of 2 percent. However, over the entire sample, there is a large savings response of 9 percent. Since there is no difference in savings at the start, this suggests that the savings of both groups were similar at the beginning. But over time, the tracker group saved about ten percent of the payment difference. The results over the entire sample are statistically significant at the 1 percent level.

Table B5: The Marginal Propensity to Save out of payment shocks, at start of sample and long differences

	Savings		Log Savings		Δ Savings 2011-14		Δ Log Savings 2011-14	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cumulative Payment Difference	0.086	0.023						
	(0.0441)	(0.0519)						
Cumulative Payment Difference			0.019***	0.019***				
			(0.0050)	(0.0055)				
Δ Cumulative Payment Difference 2011-14					0.091***	0.093***		
					(0.0234)	(0.0266)		
Δ Cumulative Payment Difference 2011-14							0.072**	0.080**
							(0.0234)	(0.0265)
Observations	10351	10351	10351	10351	10351	10351	10351	10351
Adjusted R^2	0.000	0.064	0.001	0.061	0.002	0.004	0.001	-0.002
Controls × Quarter FE		Yes		Yes		Yes		Yes

Note: Controls include current age, age at origination (both 5 year bins), income deciles, post code and year of origination of first lien, total outstanding balance as of 2008 Q4, loan to income as of 2008 Q4, loan to value as of 2008 Q4, all of which are interacted with quarter dummies. Standard errors are reported in parentheses. Standard errors are clustered at the level of the fixed effects.

B.5 Deposits as a Share of Assets, by Liquid Asset Quartile

Table B6: Deposits as a Share of Assets, by Liquid Asset Quartile

	Deposits	Share of Financial	Share of Real
Quartiles of Deposits (within Own w/ Mortgage)			
1	146	78	0
2	1694	79	1
3	6497	79	3
4	45068	76	15
Total	13222	78	5

Note: Table reports deposits as a share financial and real assets. Sample is only households who own a home with a mortgage. Figures are weighted using survey weights. Source: Household Finance and Consumption Survey, European Central Bank.

C Appendix to Section 6: Model

C.1 Derivations

Equation (15) comes from applying the implicit function theorem to the consumption function $c(e, a; \tau)$

$$dc_t(e, a; \boldsymbol{\tau}) = \sum_{s=t}^{\infty} \frac{\partial c_t(e, a)}{\partial \tau_s} d\tau_s \tag{18}$$

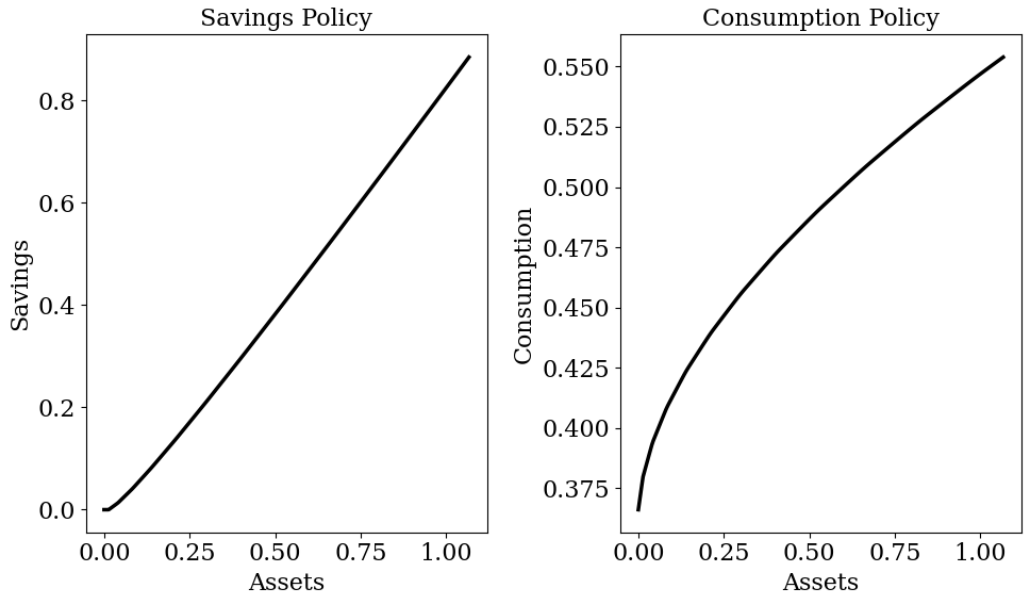
Imposing that the shock is constant $d\tau_s = d\tau$ and lasts for k periods gives equation (15)

$$\frac{dc_t(e, a; \boldsymbol{\tau})}{d\tau} = \sum_{s=t}^{k} \frac{\partial c_t(e, a)}{\partial \tau_s}. \tag{19}$$

C.2 Appendix for Quantitative Results

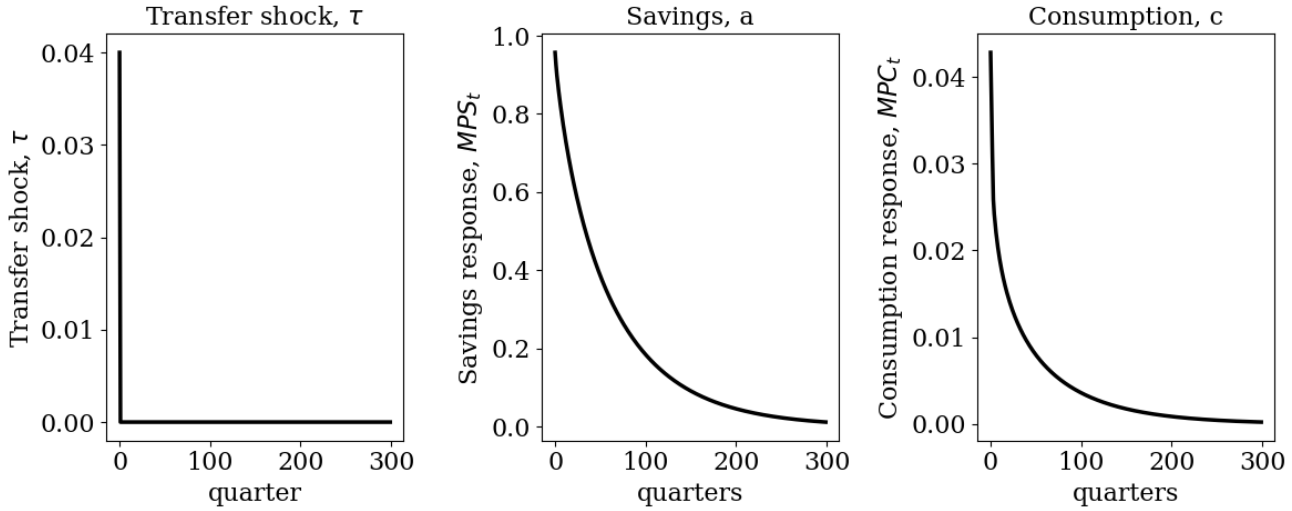
This section presents additional results for the quantitative model. Figure C1 plots the policy functions for the model. Figure C2 plots the response of the model to a one time income shock. The relevant results are discussed in the main body.

Figure C1: Model's Policy Functions



Note: Plots show the policy functions for savings and consumption under the benchmark calibration that is reported in the main text.

Figure C2: Model Response to One Time Income Shock



Note: Plots show the response of the model to a one time income shock. The left panel shows the path of the transfer shock, i.e. the income shock. The middle panel shows the response of savings and the right panel shows the response of consumption. The MPC on impact 0.043.

C.3 Comparing to Models of One-Time Shocks

To provide some context for our results, we discuss our model's response to a one-time shock as shown in Figure C2. The average MPC is 4.3 percent, which is similar to the 4.6 percent that Kaplan and Violante (2022) report in the benchmark calibration of a one-asset model. This is an order of magnitude lower than the empirical estimates of MPCs, such as Fagereng, Holm and Natvik (2021) who estimate an MPC of 53 percent using lottery winners. It is in this context that we consider the fit of the model's response to persistent shocks to be a success. The model can fit the average response with reasonable parameters values, and the heterogeneous responses are within ten percentage points for the liquid assets. As we've acknowledged, the 34 percentage point difference for less persistent shocks shows the limitation of the one-asset model for transitory shocks. Overall though, the model's response to these persistent shocks is much closer to the empirical estimates than the model's response to a one-time shock.

The literature on one-time shocks has also found mixed results regarding the importance of liquidity constraints. Like our approach, the literature has tested whether households with low liquid assets have higher consumption responses as suggested by a model, like ours, with liquidity constraints. Parker, Souleles, Johnson and McClelland (2013) find that such households have *lower* MPCs in response to the United States' 2008 fiscal stimulus checks, while Johnson, Parker and Souleles (2001) and find that they have *higher* MPCS when studying the 2001 stimulus checks. Likewise Kueng (2018), while studying the Alaska Permanent Fund, finds that there is little clear

variation across households who differ in liquid assets and that high-income households with large liquid assets display excess sensitivity with MPCs above 50 percent. On the other hand, Fagereng, Holm and Natvik (2021) find that households with low liquid assets have higher MPCs in response to lottery winnings in Norway. Our results suggest that liquidity constraints are quantitatively important for the consumption response to persistent shocks.

One potential explanation for these discrepancies is that the cumulative size of our shock is much larger than the size of the shocks in the previous literature. As shown in Figure 4, the median household received over €12,000 by 2014Q1, compared with between $600 and $1,200 per couple in the 2008 stimulus checks. Kueng (2018) shows that the welfare loss from mis-spending a transfer is increasing in the size of the transfer relative to consumption. In the case of smaller transfers, this can rationalize high MPCs for richer households, because the welfare loss is small compared to following an optimization problem. As shocks get larger, like in our setting, the incentive to optimally save is higher, and therefore we can better detect the impact of liquidity constraints in the higher MPCs of poorer households.