How Do Interest Rates Affect Consumption? The Household Debt Channel*

Angus Foulis, Bank of England
Jonathon Hazell, LSE
Atif Mian, Princeton
Belinda Tracey, Bank of England

September 11, 2025

Preliminary draft.
Please do not quote without the permission of the author(s).

Abstract

We combine the staggered expiry of mortgage deals in the United Kingdom, with administrative data on mortgages and consumption, to construct six million household-level natural experiments. We estimate the household debt channel—the causal effect of interest-rate changes on the cash-on-hand and consumption of indebted households. Our estimates are large: a 1 pp rate cut raises cash-on-hand by 1,800 pounds (1.4%), and consumption by 950 pounds (3%), implying a consumption response of 0.6% of GDP. Rate cuts matter due to a household financial accelerator—lower rates increase house prices and collateral values, leading to lower mortgage spreads, and higher cash-on-hand and consumption. Using plausibly exogenous variation in the sensitivity of house prices to interest rates, we show that the financial accelerator accounts for half of the effect of rate cuts.

*We thank Amil Mumssen, Verónica Bäcker-Peral and Alexander Prior for excellent research assistance; and Matias Bayas-Erazo, Matthias Doepke, Martin Flodén, John Gathergood, Matteo Iacoviello, Ben Moll, Huw Pill, Christina Patterson, Ricardo Reis and Maria Olsson for comments. UK residential mortgage data are taken from the Product Sales Database (PSD) provided by the UK Financial Conduct Authority. The FCA Product Sales Data include regulated mortgage contracts only, and therefore exclude other regulated home finance products such as home purchase plans and home reversions, and unregulated products such as second charge lending and buy-to-let mortgages. This paper contains the views of the authors and not necessarily those of the Bank of England, the MPC, the FPC, or the PRC.

1 Introduction

When interest rates fall, do people spend more? The answer is fundamental to macroeconomics—for instance, determining the power of monetary policy over the economy. One factor that affects the transmission of rate cuts into consumption is household debt. Many households hold large amounts of debt, which exposes them to interest rate changes. For instance, in the United States and the United Kingdom, roughly a third of households have a mortgage; rate cuts lower the cost of servicing the debt, and encourage households to borrow more through their mortgage. A natural question is whether this "household debt channel"—the transmission of rate cuts into consumption via household debt—is large.

To measure the household debt channel, this paper combines millions of natural experiments with individual level data on mortgages and consumption. In the UK, households choose a mortgage deal that typically lasts for 2, 3, or 5 years. The interest rate is determined at the start of the deal, and fixed for the rest of the deal; floating rates are uncommon. When the deal expires, there are strong incentives to start a new deal. There is no cost of borrowing through the mortgage as a new deal starts, but large costs otherwise. Altogether, households face interest rate changes and refinancing opportunities at predetermined and staggered times.[1]

We pair these natural experiments with a dataset of individual and monthly level mortgage and consumption data for 2015-2024. We start with administrative data on mortgages for the universe of UK households, which we merge with consumption information from two personal finance apps. We observe individual-by-transaction level data from bank accounts linked to the app. We classify the transactions to arrive at measures of durable and non-durable spending. Using consumption data from apps has two challenges: selection into using the app is not random, and users might link only a subset of their accounts to the app. Fortunately, neither concern appears to be too great. App users are broadly representative of the population, and expenditure in the app is similar to official sources. There are 6.8 million natural experiments in the mortgage dataset, 3.7 million of which are repeat observations of the same household. We observe consumption information for 70 000 natural experiments.

To fix ideas about our empirical strategy, we study a simple incomplete markets model with mortgages. Households receive income, consume goods, hold mortgages, and hold a liquid asset subject to borrowing constraints. The mortgage rate depends on both aggregate rates and the household's loan-to-value ratio. We then define the household debt channel. When rates fall, households borrow more through their mortgage and face lower debt service

[1] Cloyne, Huber, Ilzetzki and Kleven (2019) were the first to uncover this natural experiment, with which they investigated the effect of house prices on borrowing—we instead study how interest rates affect consumption.

costs, meaning they can consume more. The household debt channel holds fixed the effect of rate cuts that do not operate via debt—such as increases in income that pass through to consumption, as stressed by Heterogeneous Agent New Keynesian (HANK) models (Kaplan, Moll and Violante, 2018; Auclert, 2019).[2] One reason why the debt channel may be large is the "household financial accelerator". When rates fall, house prices rise. With greater collateral, households' mortgage interest rates fall. Households can borrow more, and their debt service falls—meaning more cash-on-hand and consumption. The financial accelerator can affect both new borrowing, as Kiyotaki and Moore (1997) and Iacoviello (2005) emphasize, and payments to service existing debt (Bernanke, Gertler and Gilchrist, 1999).[3]

Our natural experiments identify the household debt channel. We regress the change in households' consumption, as their deal expires, on the change in aggregate rates between the start and the end of the deal, and include a time fixed effect. We also study the response of cash-on-hand—the sum of changes in borrowing and debt service. Our identification assumption is that the change in rates over the past deal is independent of idiosyncratic income shocks that happen as the deal expires. If so, then the natural experiment "differences out" movements in income—meaning we have identified the household debt channel. Our identification strategy exploits the particular feature of the UK mortgages: deal expiration is predetermined. In other countries such as the United States, refinancing is not predetermined and instead is a choice by households—which might well correlate with income shocks.

We use three sources of variation to estimate the household debt channel. First, in a given month, there is cross-sectional variation in the interest rate changes experienced by different households. Suppose that rates have been declining over the past five years. Then a household with a 5-year deal experiences a larger rate cut compared to a household with a 2-year deal expiring at the same time; we can isolate this variation with month fixed effects. Second, we can add deal length by year fixed effects, which compares two households with the same deal length, expiring at different times during the year. Third, since half of the 6 million natural experiments are repeat observations of the same household, we can add household fixed effects to the regression. The regression then compares the response of the same household as they receive different rate changes. Since our dataset contains millions of natural experiments, we can study each source of variation separately.

Our main result is that the household debt channel is large. In our baseline specification,

[2]As such, the household debt channel is closely related to what these papers term the "direct effect of interest rates on consumption" or the "interest rate exposure channel".

[3]A second reason why the household debt channel can be large is the "cashflow effect" emphasized by Di Maggio, Kermani, Keys, Piskorski, Ramcharan, Seru and Yao (2017) and Holm, Paul and Tischbirek (2021) among others—holding fixed house prices, lower rates reduce debt service and raise consumption. As we will discuss, there is abundant evidence that cashflow effects are important, but less evidence on how rates affect consumption via the household financial accelerator.

which combines all three sources of variation, a 1 pp rate cut causes consumption to rise by around 950 pounds, or 3.0% of consumption, in the 6 months after the deal expires.[4] Using each source of variation separately, the effect of a 1pp rate cut on the consumption of households with expiring mortgage deals ranges from around 850 to 1,800 pounds (2.6-7.6% of consumption). The response of cash-on-hand is 1,800 pounds (1.4%) in the baseline specification, with similar numbers using each source of variation alone. The marginal propensity to consume out of cash-on-hand ranges from 0.5 to 0.8 in the 6 months after deal expiry. Selection into the consumption data is unlikely to explain our results—the response of the cash-on-hand to rates is similar for households with and without the personal finance app. Our estimates of the household debt channel are sizable: for instance in the baseline specification, consumption increases by 0.6% of GDP after the 1 pp rate cut. The consumption response via the household debt channel is similar in size to the effect of aggregate monetary policy shocks on consumption in the UK (e.g. Cloyne, Ferreira and Surico, 2020).

Besides the stability of our results across different sources of variation, there are two pieces of evidence that support our identification assumption. First, there are no pre-trends, meaning that any confounding shock would have had to occur in precisely the same month that the deal expires. Second, after rate cuts, our measure of income falls at the same time as consumption rises. This pattern is consistent with standard wealth effects on labor supply and does not violate our identification assumption, whereas confounding shocks would likely increase income.

Our estimates inform an ongoing discussion about the channels through which monetary policy affects consumption in HANK models. Auclert (2019) studies a canonical HANK model with a financial accelerator, and describes the conditions under which the household debt channel matters for the overall transmission of monetary policy to consumption: the marginal propensity to consume must be relatively large for households who are exposed to rate cuts via their debt position. We provide sharp evidence for this condition.

In the second part of the paper we show that the household debt channel is large in part because of strong financial accelerator effects. We start with some preliminary evidence about how the two components of cash-on-hand, borrowing and debt service, respond to interest rate changes. After 6 months, borrowing accounts for around two-thirds of the response of cash-on-hand. We infer that the mechanism generating a large household debt channel must affect not only debt service but also borrowing—making the financial accelerator a promising candidate.

We then introduce a strategy to pin down how much the financial accelerator matters.

[4] All nominal pound values discussed in the paper are based in 2015. Rate cuts refers to changes in the mortgage rate not the Bank of England Base rate.

Given the millions of natural experiments in our data, we can separately estimate the household debt channel in each region of the UK. We then estimate how house prices respond to interest rates in each region. If the household financial accelerator is important, then borrowing should respond more to rate cuts in regions with sensitive house prices. For example, suppose that house prices respond strongly to interest rates in London, but do not respond to rate cuts in York. If the financial accelerator were important, then the response of borrowing to rates would be small in York and large in London. One might worry about omitted variables: perhaps the intertemporal elasticity of substitution is different in York versus London, confounding the role of house prices. Therefore we develop two instruments for the response of house prices to rates, based on the land supply elasticity and the duration of the housing stock (Guren, McKay, Nakamura and Steinsson, 2021; Giglio, Maggiori and Stroebel, 2015; Backer-Peral, Hazell and Mian, 2024).

We find that the household financial accelerator is important for how rates affect borrowing and consumption. For households in regions where house prices do not respond to interest rates, borrowing responds by on half as much as for households in the average region in our baseline specification. These estimates imply that the household financial accelerator accounts for half of the response of borrowing to interest rates. Across various specifications the household financial accelerator remains important, accounting for around 35%-65% of the overall borrowing response. The financial accelerator appears to be similarly important for how consumption responds to interest rate changes. However, the financial accelerator accounts for little of how rates affect debt service. Therefore variants of the financial accelerator model that stress changes in borrowing instead of debt service, such as Kiyotaki and Moore (1997), Iacoviello (2005) and Greenwald (2018), seem particularly relevant.

Related literature. There is already a thorough literature estimating how interest rates affect consumption at the micro level. To estimate causal effects, papers use a variety of identification schemes, such as aggregate monetary policy shocks (e.g. Flodén, Kilström, Sigurdsson and Vestman, 2020; Holm et al., 2021), event studies around refinancing (e.g. Beraja, Fuster, Hurst and Vavra, 2019; Berger, Milbradt, Tourre and Vavra, 2021), and regional differences in the exposure to rate cuts (e.g. Eichenbaum, Rebelo and Wong, 2022). Di Maggio et al. (2017) exploit that around the Great Recession, certain households experienced sharp falls in debt service costs as their adjustable-rate mortgages reset. The literature uses a variety of proxies for consumption, such as imputations from tax records, car purchases, and local economic activity. One consensus finding is that the "cashflow effect" of rates on consumption is important: holding fixed house prices, rate cuts raise consumption by lowering debt service. However, cashflow effects seem to be too small to explain all of the transmission of rate cuts to consumption—raising the question of what other mechanisms

matter (Cloyne, Ferreira and Surico, 2020).

We build on the previous work in three ways. First, we observe monthly-frequency transaction-level information on consumption, for both durables and non-durables. As such, we can estimate useful statistics such as the marginal propensity to consume, which connect our estimates to quantitative models. Second, our setting contains millions of natural experiments—offering powerful sources of variation for measuring how rates affect consumption. Third, we show that as well as the cashflow effects that past work emphasizes, the household financial accelerator also matters—favoring models such as Iacoviello (2005) and Greenwald (2018) that emphasize this channel.

A second literature estimates how interest rate cuts affect consumption in aggregate, using monetary policy shocks identified from the time series. Seminal papers include Romer and Romer (2004) and Gertler and Karadi (2015). One particularly relevant paper is Cloyne, Ferreira and Surico (2020), which shows that in the US and UK, the aggregate response of consumption to rate cuts concentrates among mortgage holders. We instead estimate how rates affect consumption through a particular channel—arising from household debt—while using cross-sectional variation. Our estimates of the household debt channel are similarly sized to past estimates of the overall effect of rate cuts on consumption. As such, household debt seems to be a crucial part of the monetary transmission mechanism.

2 Data

We now describe the mortgage and consumption datasets.

2.1 Mortgage Dataset

We measure mortgages from the Product Sales Database (PSD), an administrative dataset collected by the United Kingdom's Financial Conduct Authority. This dataset contains snapshots every 6 months of the universe of outstanding mortgages in the United Kingdom, from the start of 2015 until the start of 2024. At any point in time, there are roughly 8 million mortgages.

The PSD measures comprehensive data about mortgages, at monthly frequency. In particular, the PSD measures the initial loan size, the mortgage term, the initial loan interest rate, whether the interest rate is fixed or floating, whether the mortgage is refinanced, and how much has been paid back before refinancing. Finally, we measure how much equity is extracted at refinancing—that is, how much the household has borrowed using their mortgage; as well as the mortgage payments, i.e. the debt service cost of the mortgage. As we will

discuss later in more detail, in the UK mortgages typically have a discounted "deal" period of 2, 3, or 5 years. During the deal, interest rates are low. Afterwards, the mortgage interest rate sharply increases, but households can pay a relatively low fee to start a new deal. The PSD measures the length of the deal period and the fee paid at the start of the deal. We also observe the full date of birth for borrowers and the full postcode of their property, allowing us to track the same mortgage over time.[5] The PSD records the house price if the mortgage is for a new purchase.

In some exercises, we study a proxy for the house price associated with the mortgage. We use granular data on regional house prices from the UK House Price Index. This series captures changes in the value of residential properties bought both with cash and with a mortgage, and is calculated for 360 Local Authority Districts.

2.2 Consumption Data

We use two consumption datasets, from ClearScore and Money Dashboard. Both ClearScore and Money Dashboard are personal finance apps. The apps have a common owner and use essentially the same technology, but serve different customer bases. Both consumption datasets provide daily spending information and we create monthly aggregates to link to the monthly frequency of the mortgage dataset.

The apps collect information about consumers' finances across their various bank accounts. In each case, consumers who wish to simplify their finances can add their bank accounts to the app. The app then scrapes historical information from the consumers' bank accounts in order to form a comprehensive measure of their finances.

Money Dashboard is available from 2013 until mid 2021. Our sample of ClearScore data is available from July 2018 to March 2024. Money Dashboard contains around 130 thousand regular users, while ClearScore contains a little over 1.5 million users.

We observe each transaction by every individual in a bank account that they have linked to the app. The app measures what card type is associated with the transaction (i.e. credit or debit) and tags the transaction with an expenditure category based on the text of the transaction. These tags are detailed, at the level for instance of "gym equipment" or "take-away food". We then aggregate these spending categories into measures of durable, non-durable and services expenditure that follow the standard definition of the Bureau of Economic Analysis, as well as a comprehensive, broad measure of spending which we will refer to as "consumption".[6] We also observe various forms of income paid into accounts associated with the app. The most common is salary or wages, but other forms of income such as investment

[5] In the UK, postcodes typically contain around 15 properties.

[6] Appendix C.2 provides a detailed description of how we create the spending categories.

or rental income are also tracked. We also observe information on credit card debt, personal loan repayment and holdings of liquid assets in checking accounts.

To merge the finance app information to the mortgage dataset, we exploit that in the apps we can observe the sector-level (4-digit) postcode, year and month of birth, mortgage lender, and the monthly mortgage payment amount. In the mortgage dataset, we observe the full (6-digit) postcode, date of birth, mortgage lender, and the monthly mortgage payment amount. We merge on these variables, which selects unique individuals in the PSD in over 99% of cases. ClearScore merges 162 thousand households to the PSD, whereas Money Dashboard merges 37 thousand households.[7]

We then validate the consumption data. A first concern is that the apps contain only a subset of consumers' spending, because they do not link all of their accounts to the app. Fortunately, this concern does not appear to be too important. Appendix Figure B.1 plots mean monthly expenditure for mortgagors in Money Dashboard and ClearScore, compared to mean annual expenditure among mortgage holders in the Office for National Statistics' Living Costs and Food Survey. Both ClearScore and Money Dashboard have similar expenditure to the official data—suggesting that the apps do not underestimate spending.[8]

A second concern is that households with consumption data are not random. We can construct weights for the consumption sample, in order to correct for potential selection bias and to make the sample representative of the full population based on observables. Our sample weights are based on quintiles of age, income, loan-to-value, mortgage term and home value, as well as the broad region (NUTS1) of the property. We describe how we construct these weights in Appendix E. Appendix figure B.2 shows the distribution of these variables in the mortgage sample, and the consumption sample, with and without weighting. Households in the consumption sample are somewhat younger, own a less expensive house, have a higher loan to value ratio, and are less likely to live in London than the mortgage population as a whole. However, the results show that once our sample weights are included, the households in the consumption sample closely match the full mortgage population. We will use these weights for all of our results concerning consumption.[9]

[7] Appendix D provides a detailed description of the merge between the mortgage and consumption data.

[8] As we will discuss, if households fail to link all financial accounts to the app, then the response of consumption to rate cuts will be biased towards zero—making our estimates a lower bound.

[9] Later, we will show that the response of borrowing to rate cuts is similar for households in the consumption sample, compared to the rest of the population—suggesting that the consumption sample is not unobservably different in ways that affect our estimates.

3 Natural Experiments: Mortgage Deals in the UK

This section introduces the 6 million natural experiments arising from the expiry of mortgage deals in the UK.

3.1 Mortgage Deals in the UK

In the United Kingdom, mortgages have a particular structure, which leads to interest rate changes and opportunities to borrow using the mortgage, at predetermined and staggered times. Approximately 95% of mortgagors choose fixed interest rate mortgages. These mortgages have a term that depends on the borrower's income and preferences, and may be as long as 35 years. However mortgages typically offer a deal length of 2, 3, or 5 years. During the deal, the interest rate is fixed at a relatively low "teaser" rate, determined by macroeconomic conditions at the start of the deal. When the mortgage deal expires, households have a strong incentive to start a new deal. If households remain on the same deal, they must pay a significantly higher penalty rate, with an increase of at least 200 basis points. However they can also pay a fee of 1-2 thousand pounds to start a new deal. Households are typically reminded by banks that their deal will expire, around 3 months beforehand. Households do not face additional costs from extracting equity—that is, borrowing or saving using their mortgage—when their new deal starts. However extracting equity during a deal incurs large costs. Combining these features, households receive an interest rate change, as well as an opportunity to borrow using the mortgage, at predetermined and staggered moments.

There is a key difference compared to how refinancing works in other countries such as the United States. In the US, 30-year fixed rate mortgages are common, with an option to refinance that the borrower can exercise at any time. Therefore refinancing depends on aggregate shocks, such as interest rate changes, and idiosyncratic shocks such as changes in income. By contrast, in the UK the opportunity to refinance, and experience an interest rate change, is predetermined at the start of the deal. In the coming Section 4, we will see how the predetermined deal length is crucial for identification.

3.2 Summary Statistics on Deal Expiration

We will organize our empirical strategy around expiring deals. That is, we will study how households' consumption, borrowing, and debt service on the mortgage adjust as their deal expires—and then relate these adjustments to interest rate changes. Table 1 reports summary statistics on consumption and borrowing for our main samples. In total, we observe 6.8 million deal expiration events in the mortgages dataset, and 69,000 deal expirations in the

merged mortgage and consumption data.

In the first two columns of the table, we report summary statistics about the sample with mortgage information. In the middle two columns we report summary statistics about the sample with both mortgage and consumption information. In the final two columns we report summary statistics for the sample with consumption and mortgage information, after reweighting the sample in order to target the distribution of mortgage holders. We reweight based on quintiles of age, income, loan-to-value, mortgage term, and home value, as well as broad region (NUTS1) fixed effects. Comparing the first and last two columns, we can see that after reweighting, most observables are similar in the consumption and mortgage samples, including variables such as mortgage payment and interest rates that the reweighting does not explicitly target. In what follows, we will present results both for the full sample with mortgage information ("mortgage sample"), and the sample with consumption information ("consumption sample"). Where possible we will show that results from the two samples agree. Throughout, we restrict to households who do not move house in the 12 months after deal expiry (we cannot merge mortgage and consumption data for people who move house).

3.3 Household Behavior around Deal Expiry

We now study the behavior of households as their deals expire. Most households switch to a new deal as their old deal expires—unsurprisingly, given their financial incentives. Households also adjust consumption and borrowing.

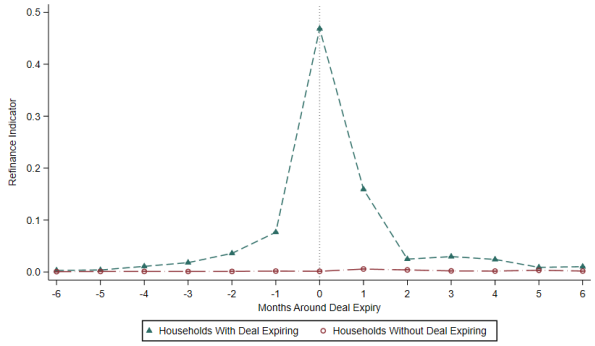
Figure 1 studies households with expiring deals. In the figure, t=0 is the point at which the deal expires in event time, with negative numbers indicating the preceding months and positive numbers the subsequent months relative to the expiry month. Then, in each month relative to the expiry month, we take the mean of a household-level outcome. We also plot outcomes for a comparison group of households. The comparison group is households whose deals do not expire during the 12-month window of the event study.

As Panel (a) of the figure shows, households with expiring deals start to switch deals around 3 months before their deal expires. Banks tend to notify households that their deal will expire around this time. The propensity to switch deals peaks in the month that the deal expires, with a monthly propensity to switch deals of 0.5, and remains elevated for several months afterwards. Overall, the cumulative probability to switch to a new deal around the refinancing trigger is around 90%.[10] Panel (b) shows that borrowing adjusts as the deal expires. There is a large adjustment, peaking at over 800 pounds in the month that the

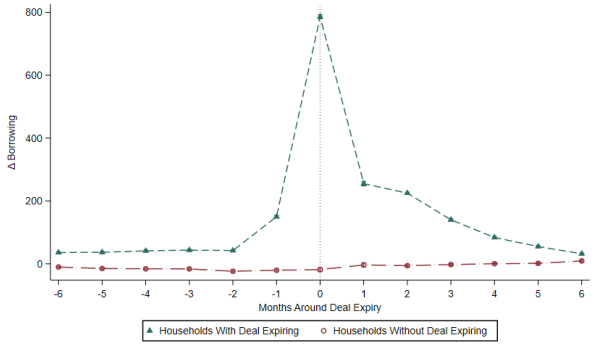
[10]The remaining households do not switch deals either because they are inattentive, because their mortgage has negative equity, or because their mortgage will finish soon.

Figure 1: Household Behavior around Deal Expiry

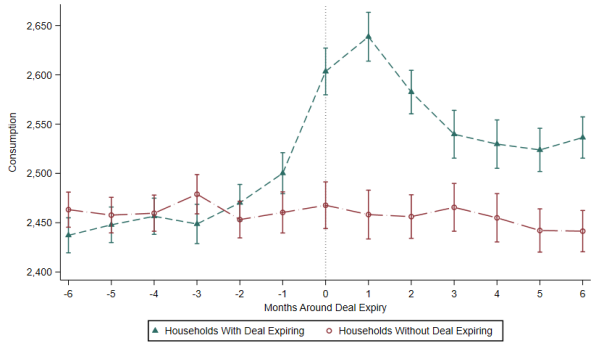
(a) Refinancing



(b) Borrowing



(c) Consumption



Notes: each panel plots household behavior as deals expire. In each month we report the mean of a household-level variable for households with expiring deals, and a comparison group comprises households whose deals do not expire within the 12-month event window. In panel (a), the variable is an indicator equal to one if the household refinances in that month and zero otherwise. In Panel (b) the variable is the mean change in household borrowing (in 2015 pounds). In panel (c) the variable is mean monthly consumption (in 2015 pounds). The bars represent 90% intervals. The sample is the mortgage population in panels (a) and (b), and the reweighted consumption sample in panel (c).

deal expires. Panel (c) shows how consumption changes, based on the smaller consumption sample; consumption increases substantially.

On average, consumption and borrowing rise as the deal expires, because during most of our sample, interest rates were falling. What is more relevant is that regardless of the direction, households are much more likely to adjust their behavior when deals expire. In the coming section, we will relate these adjustments to interest rate changes experienced as the deal expires.

4 The Household Debt Channel: A Simple Framework

This section develops a simple, empirically-oriented framework in order to define the household debt channel of interest rate changes, and the sense in which our natural experiments identify this channel. We also discuss how this channel can be large due to the presence of

a "household financial accelerator".

4.1 A Simple Model with Mortgages and Collateral

We study a standard model with mortgages and collateral, following for instance Laibson, Maxted and Moll (forthcoming) and Bernanke et al. (1999). We focus on the household's problem. The household is endowed with an old mortgage deal, that started at time $-n$ and ends at time 0. At time 0, the household starts a new mortgage deal, which ends at time N. During the course of the deal, the household chooses consumption c_t and labor supply n_t given utility $u(c_t, n_t)$, discounted with factor β; as well as holdings of liquid assets a_t subject to a borrowing constraint $\underline{a}$. At the start of the new deal, at time 0, the household can adjust their mortgage debt d_t. At all other times, mortgage debt declines by a fixed amount ξ, which is payment towards the principal of the mortgage. The household takes as given their wage w_t and the rate of return on the liquid asset r_t^a.

The mortgage rate is increasing with the loan-to-value ratio at the start of the deal, leading to a role for collateral. Specifically, the mortgage rate r_t^d is the sum of an aggregate rate r_t, as well as a household-specific spread μ_t. The spread is an increasing function of the loan to value ratio (LTV), i.e. $\mu_t = \mu\left(d_t / p(r_t)h\right)$. In this expression, the denominator depends on the stock of housing h, which we assume to be fixed, as well as the house price $p(r_t)$. House prices, in turn, can depend on aggregate rates r_t, which the household takes as given. The mortgage rate only adjusts at time 0, and is fixed before and afterwards.[11]

Overall, the household maximizes the discounted present value of their utility from time $-n$ onwards, solving

$$\max_{\{c_t, a_t, n_t, d_t\}_{t=-n}^{N}} E_{-n}\left[\sum_{t=-n}^{N} \beta^t u\left(c_t, n_t\right) + \beta^{N+1} v\left(a_N, d_N\right)\right]$$

by choosing consumption, labor, liquid assets and mortgage debt, while taking into account their continuation value $v\left(a_N, d_N\right)$ at the start of the subsequent deal after time N. Starting at time $-n$, households take expectations over the wage, which is random; all other exogenous variables are deterministic. They respect a budget constraint which allocates funds across consumption, liquid assets, and the mortgage, given labor income and the past rate of return on liquid assets and mortgage debt:

$$c_t + a_t - d_t = w_t n_t + \left(1 + r_{t-1}^a\right) a_{t-1} - \left(1 + r_{t-1}^d\right) d_{t-1},$$

[11] Following Laibson et al. (forthcoming), we abstract from moving house, meaning we ignore wealth effects from higher house prices.

and also respect the borrowing constraint on liquid assets $a_t \geq \underline{a}$; the requirement at times $t \neq 0$ mortgage debt satisfies $d_t - d_{t-1} = \xi$ and does not adjust; as well as the process for the mortgage interest rate, which adjusts only at time 0, and depends on both the aggregate rate and a spread, namely:

$$r_t^d = \begin{cases} r_0 + \mu_0 & t \geq 0 \\ r_{-n} + \mu_{-n} & t < 0 \end{cases} \qquad \mu_t = \mu\left(\frac{d_t}{p\left(r_t\right) h}\right).$$

Here, we have modeled a "soft" borrowing constraint, in which households can always borrow more, but only at a higher interest rate (e.g. Bernanke et al., 1999). This specification nests the case of a "hard" borrowing constraint (e.g. Kiyotaki and Moore, 1997; Iacoviello, 2005). With hard constraints, the slope of the mortgage spread with respect to loan-to-value, $\mu'(\cdot)$, becomes infinitely large at certain values, and the household can borrow only a fixed fraction of their home value.[12]

4.2 The Household Debt Channel and the Financial Accelerator

We can now define the main object of interest in this paper: the household debt channel. To define the channel, consider perturbations at time 0 to the aggregate interest rate r_0 starting from r_{-n}, as well as perturbations to the wage w_0 and the rate of return on the liquid asset r_0^a. The first order response of consumption at time 1 is

$$\Delta c_1 = \text{MPC}\left[\frac{\partial \text{cash-on-hand}_1}{\partial r_0}\right](r_0 - r_{-n}) + \mathcal{W}\Delta w + \mathcal{A}\Delta r^a \tag{1}$$

In this equation, Δw and Δr^a are the perturbations to wages and the return on liquid assets at time 1, which affect consumption with coefficients $\mathcal{W}$ and $\mathcal{A}$ that are defined in the Appendix. $r_0 - r_{-n}$ is the perturbation to the interest rate realized at time 0. $\partial \text{cash-on-hand}_1/\partial r_0$ is the response of cash-on-hand to the interest rate change, which can be decomposed into the response of borrowing and the response of debt service. That is, the response of cash-on-hand to changes in the aggregate rate satisfies

$$\frac{\partial \text{cash-on-hand}_1}{\partial r_0} = \underbrace{\frac{\partial d_1}{\partial r_0}}_{\substack{\text{borrowing} \\ \text{response}}} + \overbrace{\left(1 + r_0^d\right)\frac{\partial d_0}{\partial r_0} + \left(1 + \frac{\partial \mu_0}{\partial r_0}\right) d_0}^{\text{debt service response}} \tag{2}$$

[12] In the UK, spreads discontinuously increase at certain LTV thresholds, and are flat or moderately upward sloping otherwise (Best, Cloyne, Ilzetzki and Kleven, 2020), representing a functional form for $\mu(\cdot)$ that mixes hard and soft constraints.

where the first term on the right hand side is the response of mortgage borrowing $\partial d_1/\partial r_0$, the response of the mortgage policy function at time 1 to interest rate changes at time 0. The second term is the change in debt service—including servicing higher borrowing due to the rate cut $(1 + r_0^d)\partial d_0/\partial r_0$; and the changing cost of servicing the previous debt d_0, given how the rate cut affects mortgage spreads $\partial \mu_0/\partial r_0$. MPC $\equiv \Delta c_1/(\Delta c_1 + \Delta a_1)$ is the marginal propensity to consume out of the change in cash-on-hand, i.e. the relative response of consumption versus liquid asset accumulation to the interest rate cut.

The first term on the right-hand side of equation (1) is the household debt channel. After a rate cut, households borrow more and have lower debt service. As a result they have more cash-on-hand—which they use for either consumption or liquid asset accumulation, with the relative response of consumption determined by the MPC out of cash-on-hand from rate cuts. This MPC is different from the familiar MPC out of transitory income shocks (e.g. Johnson, Parker and Souleles, 2006). Here, the change in cash-on-hand is persistent and associated with a change in relative prices; both forces will tend to raise the MPC in our setting.[13]

Separate from the household debt channel, rate cuts affect consumption by raising income in general equilibrium. This force is represented by the $\mathcal{W}\Delta w$ term in equation (1). This channel has been emphasized by the Heterogeneous Agent New Keynesian (HANK) literature (e.g. Auclert, 2019). As such, the household debt channel is closely related to what Kaplan et al. (2018) term the "direct effect" of interest rates on consumption, whereas the $\mathcal{W}\Delta w$ term in equation (1) represents the "indirect effect" operating via general equilibrium movements in income. Interest rate changes outside the mortgage market can also matter for consumption, as summarized by the $\mathcal{A}\Delta r^a$ term.

In this setting, rate changes are anticipated from the start of the previous deal, at time $-n$. Nevertheless, the consumption response to rate cuts can still be positive around time 0, if borrowing constraints on the liquid asset bind before time 0.

What factors affect the size of the household debt channel? One important force is what we term the "household financial accelerator", which affects both components of cash-on-hand. When the interest rate r_t falls, house prices $p(r_t)$ rise. As a result, holding fixed mortgage debt d_t, the loan-to-value ratio $d_t/p(r_t)h$ falls. Since the household's mortgage interest rate depends on the loan to value, this interest rate falls as well. In response the household borrows more. Moreover the cost of servicing existing mortgage debt falls. As a result, cash-on-hand rises—meaning more resources available for consumption, as Bernanke

[13]The response of cash-on-hand to rate changes depends on the income and substitution effects of the rate change, how rate cuts affect the slackness of the borrowing constraint on the liquid asset, and the response of house prices to rates (see the Appendix).

et al. (1999) emphasizes. If the household is at a "hard" borrowing constraint, their loan to value ratio is fixed and their mortgage spread does not change. Therefore their borrowing will increase in proportion to the value of their home, as Kiyotaki and Moore (1997) and Iacoviello (2005) emphasize. In particular, we show in the Appendix that for the case of either soft or hard borrowing constraints, we can write the response of mortgage borrowing to rate cuts as

$$\frac{\partial d_1}{\partial r_0} = \beta_0 + \beta_{FA} \frac{d \log p\left(r_0\right)}{dr_0}. \tag{3}$$

Here, β_{FA} is a coefficient that represents how borrowing depends on the semi-elasticity of house prices to rate cuts, $d \log p\left(r_0\right) / dr_0$, which captures the household financial accelerator. Similarly, debt service depends on how interest rate cuts affect house prices, as the mortgage spread term $\partial \mu_0 / \partial r_0$ term in the debt service term of equation (2) indicates.

There are reasons why the household debt channel could be large other than the financial accelerator. One prominent reason is "cashflow effects"—holding fixed house prices, rate cuts still lower debt service. Equation (2) shows that debt service falls after rate cuts even if mortgage spreads μ_0 are constant and the household financial accelerator is not operational. If so, cash-on-hand–and possibly consumption–rises. There is abundant evidence that cashflow effects are important (Di Maggio et al., 2017, Holm et al., 2021); though they are unlikely to account for all of the transmission of rate cuts to consumption (Cloyne et al., 2020). However evidence on how rate cuts affect consumption via the household financial accelerator—the focus of this paper—is more scarce.

4.3 Natural Experiments and the Household Debt Channel

We now show how to use our natural experiment—the expiry of mortgage deals—in order to identify the household debt channel. We will then estimate the household debt channel in the coming Section 5, and in Section 6 we link to the household financial accelerator.

Equation (1), which defines the household debt channel, also shows the challenge to identifying it. There may be other shocks to consumption that do not operate through the household debt channel. These shocks could represent how rate cuts raise consumption through other channels, such as the general equilibrium movements in income stressed by the HANK literature. Or, there could be shocks to consumption that are unrelated to interest rate cuts, such as idiosyncratic income shocks. Either way, income movements confound estimates of the household debt channel.

We can use our natural experiment to estimate the household debt channel via a regression. Suppose that there is a population of households whose consumption behavior satisfies equation (1), with deals expiring across multiple time periods. Then the consumption of

household i with a deal expiring at time t satisfies

$$\Delta c_{it} = \text{MPC} \left[\frac{\partial \text{cash-on-hand}}{\partial r} \right] (r_t - r_{t-n_i}) + \alpha_t + \varepsilon_{it}, \tag{4}$$

where α_t is a time fixed effect, and ε_{it} includes idiosyncratic shocks to income—the difference between the income change of household i and the cross-sectional mean at time t—as well as other idiosyncratic factors affecting consumption.[14]

One can estimate equation (4) by regression. The outcome variable is the consumption change of the household as their deal expires, which one regresses on a time fixed effect, as well as the change in aggregate interest rates $(r_t - r_{t-n_i})$ between the current period and the start of the previous deal. The coefficient on interest rates is the product of the response of cash-on-hand to interest rate changes, and the MPC out of cash-on-hand. Alternatively, to recover the MPC, one can follow an instrumental variables procedure: regressing the consumption change on a time fixed effect and the change in cash-on-hand around deal expiry, while instrumenting for the change in cash-on-hand with $(r_t - r_{t-n_i})$. Finally, one can estimate the first stage associated with this regression—regressing the change in cash-on-hand on the rate change and time fixed effects.

The identification assumption is that the change in aggregate rates since the previous deal is orthogonal to idiosyncratic shocks to income and other factors that affect the household as their deal expires. With this assumption, the natural experiment "differences out" general equilibrium movements in income. This assumption is tenable because the time of refinancing is predetermined; we will probe the validity of the assumption as we discuss our results. In countries like the United States this assumption is unlikely to hold—refinancing is a choice of the household, which likely correlates with idiosyncratic shocks.

Could one find simpler alternatives to our identification strategy? For instance one could attempt to isolate the household debt channel by controlling for observed income in the regression. However, if households can vary labor supply, as in our model, then observed income contains not only shocks but also the endogenous response to rate cuts. We will present evidence that labor supply varies in this way. As such, income is a "bad control", which biases the response of consumption to rate cuts if it is included in the regression.

5 Estimates of the Household Debt Channel

This section presents our estimates of the household debt channel. We first describe details of how we implement our regression and the variation that it uses to estimate the household debt

[14] Equation (4) also removes time subscripts where there is no ambiguity.

channel, before we present results. The household debt channel is large—a 1 pp interest rate cut raises borrowers' consumption by 3%, which implies an aggregate decline in consumption worth 0.6% of GDP.

5.1 Implementing the Main Regression

We first adapt the regression equation (4) to our setting. We estimate a series of regressions

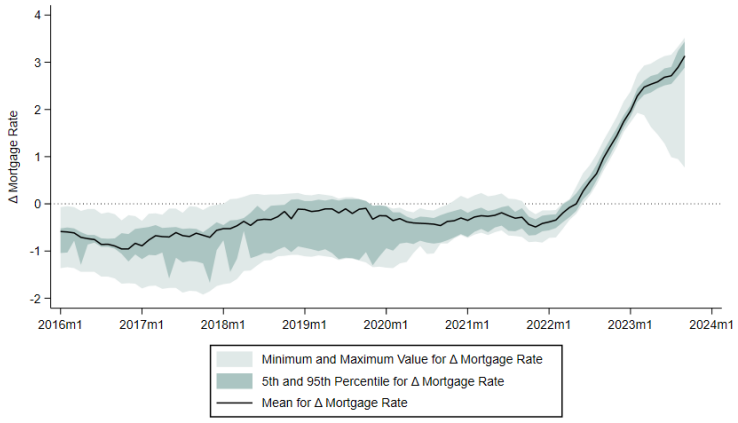
$$\sum_{j=-4}^{h} [Y_{i,t+j} - Y_{i,t_0}] = \alpha_t^h + \beta^h (r_t - r_{t-n_i}) + \text{fixed effects}_{it}^h + \varepsilon_{ith} \tag{5}$$

for each monthly horizon h. The primary outcome variables Y_{it} that we study are monthly consumption, and the monthly change in cash-on-hand (i.e the sum of the monthly payment towards debt service and the monthly change in mortgage borrowing). We take the difference between the outcome $Y_{i,t+j}$ and a base observation Y_{i,t_0}, which is the mean of Y_{it} between four and six months before refinancing. Finally, we sum over the h periods after the deal expires, in order to measure the cumulative response of consumption or cash-on-hand. The regression is estimated at monthly frequency. α_{th} is a year fixed effect, and fixed effects$_{it}^h$ represents other controls, always including fixed effects for age quintiles. With negative horizons h, we study the response of consumption and cash-on-hand prior to deal expiry, in order to account for anticipation effects. Standard errors are clustered by the month of deal expiry. Our measure of the aggregate rate r_t is the average interest rate across all new mortgage deals, which purges individual variation, provided by the Bank of England.[15]

There are three sources of variation that enter the baseline regression; we can estimate each source individually by including further fixed effects. First, there is purely cross-sectional variation. At a given point in time, households have different deal lengths, meaning they are exposed to different interest rate changes as their deals expire. For instance, suppose that interest rates have fallen over the past five years, and two households' deals expire. The first household has a five-year deal, and the second has a two-year deal. At the point of expiry, the first household receives a larger rate cut. We can isolate this source of variation only by adding month fixed effects to regression (5). Figure 2 shows this source of variation. The black line plots the mean change in the interest rate on new mortgage deals between 2015 and 2024. Mortgage interest rates were declining before 2022, and rising afterwards. As a result our estimates include both interest rate cuts and hikes. The figure also shows purely cross-sectional variation in interest rate changes. The light green shaded area illustrates the minimum and maximum interest rate change at a given point in time for expiring mortgages.

[15] We also discuss robustness exercises using different measures of the aggregate rate.

Figure 2: Interest Rate Variation from Natural Experiments



Notes: the figure shows time-series and cross-sectional variation in changes in mortgage rates on new deals. For each month from 2015-2024, we compute $\Delta r_{i,t} = r_t - r_{t-n_i}$ for households whose deals expire in month t. The black line plots the cross-sectional mean of $\Delta r_{i,t}$. The dark and light shades span the 5th and 95th percentiles, and the cross-sectional minimum and maximum of $\Delta r_{i,t}$, across expiring mortgages in month t. r_t is the average interest rate on new mortgage deals from the Bank of England. Units are percentage points; monthly frequency.

It also plots the 5th and 95th percentiles of these interest rate changes.

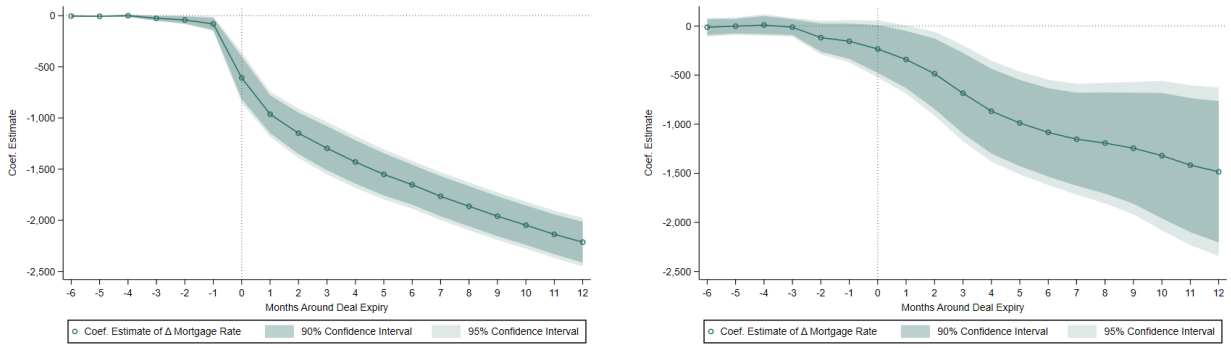
The second source of variation compares households of the same deal length, whose deal expires at different times during the year—for instance comparing a household with a 5-year deal expiring in January, to a household with a 5 year deal expiring in August of the same year. This specification controls for differences in the consumption behavior of households with long or short deals. We isolate this variation by adding deal length by year fixed effects to the regression. The third source of variation exploits that of the 6 million natural experiments, 3.5 million are repeat observations of the same household. Therefore, we can add household fixed effects to the baseline regression, exploiting how the same household responds to different-sized rate cuts over time.

What is clear from Figure 2, and the three sources of variation that we exploit, is the value of 6 million natural experiments. While one can think of omitted variables that might affect any single source of variation, it is more difficult to think of omitted variables affecting all three sources. If all sources of variation lead to similar estimates, one might have more confidence in the identification strategy.

Our identification strategy contrasts with a common approach in macroeconomics. This approach uses time-series variation to estimate how interest rate cuts, from monetary policy, affect consumption (e.g. Romer and Romer, 2004). Our cross-sectional approach offers potentially more convincing identification, albeit of only one channel through which rate cuts affect the economy.

Figure 3: Impulse Responses to Interest Rate Changes

(a) Cash-on-hand **(b)** Consumption



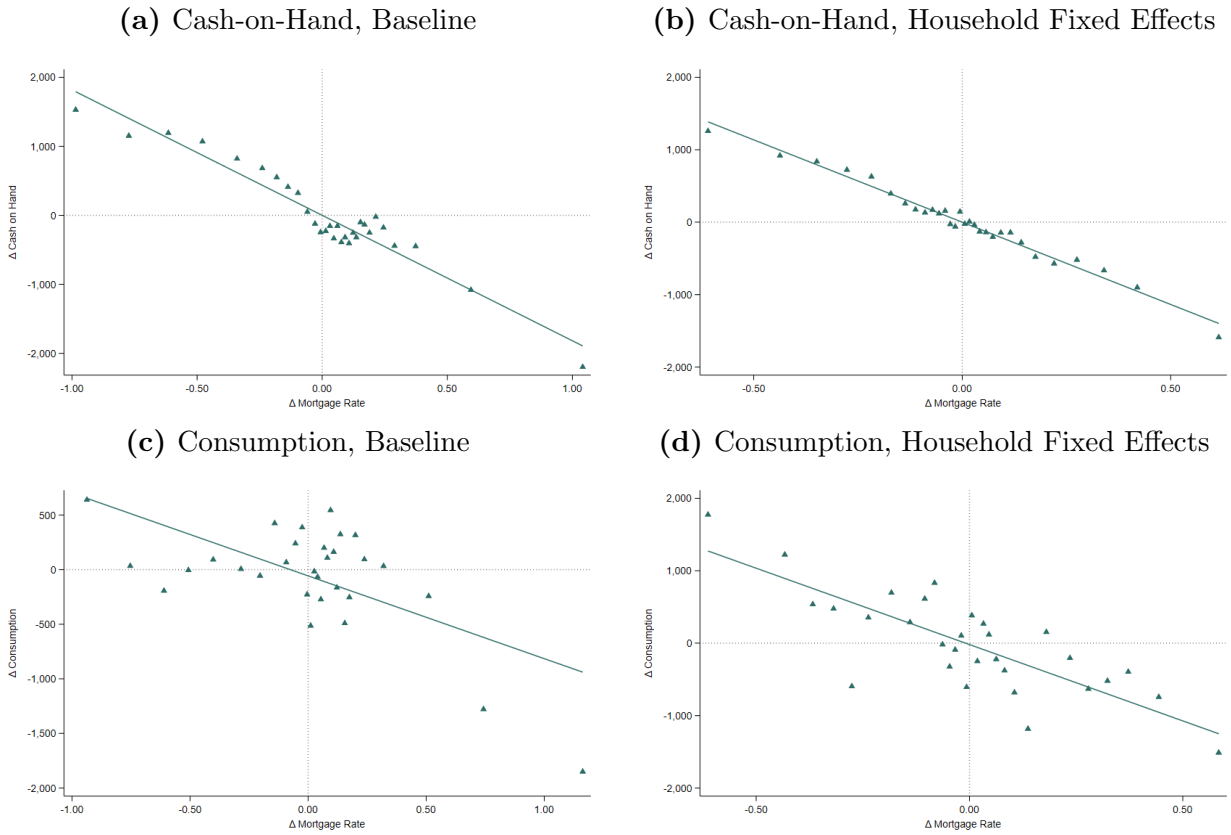
Notes: each panel plots impulse responses from the baseline monthly regression. For horizons $h \in [-6, 12]$ relative to deal expiry ($t = 0$), the figures plot the cumulative change in the outcome (pounds, 2015 base) per 1 percentage-point change in the interest rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. Shaded bands denote 90% and 95% confidence intervals. Panel (a): cash-on-hand for the universe of mortgagors. Panel (b): reweighted consumption for households with app-based spending

5.2 Estimates of the Household Debt Channel

First, we discuss the results for cash-on-hand. Figure 3, panel (a) presents the impulse response of cash-on-hand corresponding to our baseline regression. Each point plots the cumulative response of cash-on-hand at the relevant horizon, relative to deal expiry at time 0; starting from 6 months before deal expiry to 12 months afterwards. The shaded region around the coefficients represents 90% and 95% confidence intervals. The figure shows that in the month that the deal expires, a 1 pp rate cut causes cash-on-hand to increase by around 600 pounds (based in 2015, equivalent to roughly 800 US dollars). The effect grows, reaching over 2000 pounds after one year. There appears to be some anticipation—cash-on-hand starts to rise 3 months before the deal expires, at exactly the same time that refinancing starts according to Figure 1. Before then, there are no pre-trends.

We present the estimates underlying this figure in Table 2, column (1). This column contains the cumulative impulse response of cash-on-hand from 3 months before deal expiry to 6 months after. The estimated effect is an 1,818-pound increase after a 1pp rate cut. In columns (2)-(4) we study the different sources of variation associated with the baseline estimate. In Column (2) we introduce month fixed effects, meaning we compare households with deals of different length expiring in the same month. In Column (3) we add year-by-deal length fixed effects, which studies households of the same deal length expiring at different times during the year. Column (4) uses household fixed effects, on the reduced sample of households with repeat observations. In all cases, the estimates are similar to or greater than the baseline. Column (5) repeats the baseline, but with cumulated changes in log cash-on-hand as the outcome variable, now size-weighting each observation by mortgage principal

Figure 4: Binscatters of the Response of Cash-on-Hand and Consumption to Rate Changes



Notes: each panel plots binned averages of the cumulative change in the outcome (pounds, 2015 base) at the 6-month horizon relative to deal expiry against the change in the interest rate. Both variables are residualized on year and age fixed effects and so are centered around zero; panels (b) and (d) additionally residualize on household fixed effects. Observations are grouped into 30 equal-count bins of the interest-rate change; markers plot within-bin means. Axes units: x in percentage points, y in pounds. Panel (a): cash-on-hand, baseline. Panel (b): cash-on-hand with household fixed effects (repeat borrowers). Panel (c): reweighted consumption for households with app-based spending. Panel (d): reweighted consumption with household fixed effects.

before the deal expires. A 1 pp rate cut raises cash on hand by 1.4% over the event window. Columns (6)-(8) present similar estimates using each source of variation separately.

We also plot the variation associated with the estimates in Figure 4. We residualize both the 6-month response of cash-on-hand and the rate change against the baseline controls, i.e. year and age fixed effects. We then bin both variables into 50 groups, and plot them. Inspecting panel (a) of the figure, one can see a robust linear relationship between rate changes and cash-on-hand changes. One can see a similar pattern in panel (b), which additionally controls for household fixed effects.

We now present results on the impulse response of consumption, using an identical estimation framework. Figure 3, panel (b) contains the impulse response of consumption corresponding to the baseline regression. Each point plots the cumulative response of consumption at a horizon relative to the expiry of the deal at time 0; starting from 6 months

before deal expiry to 12 months afterwards. The figure shows that in the month that the deal expires, a 1 pp rate cut causes consumption to increase by around 300 pounds. The effect grows, reaching around 1500 pounds after one year. Again, there is anticipation starting 3 months before the deal expires, just as refinancing starts. However there are no pre-trends beforehand.

Table 3, column (1) presents the baseline estimates for consumption. This column contains the cumulative impulse response of consumption after 6 months, which is a 934 pound increase after a 1pp rate cut. In columns (2)-(4) we isolate the different sources of variation, again with month fixed effects in column (2), year-by-deal length fixed effects in column (3), and household fixed effects on a smaller sample in column (4). Again, all estimates are similar or greater than the baseline. Column (5) repeats the baseline, but with cumulated changes in log consumption as the outcome variable, again size-weighting each observation by mortgage principal before the deal expires.[16] A 1 pp rate cut raises consumption by 3.0% over the 6 months after deal expiry, relative to the previous 6 months. Columns (6)-(8) present similar estimates using each individual source of variation. Figure 4 contains the variation associated with these estimates, for the baseline (panel c) and the household fixed effects specification (panel d). In both cases there is a clear relationship between rate cuts and consumption increases—albeit noisier than for cash-on-hand, given the smaller sample size.[17]

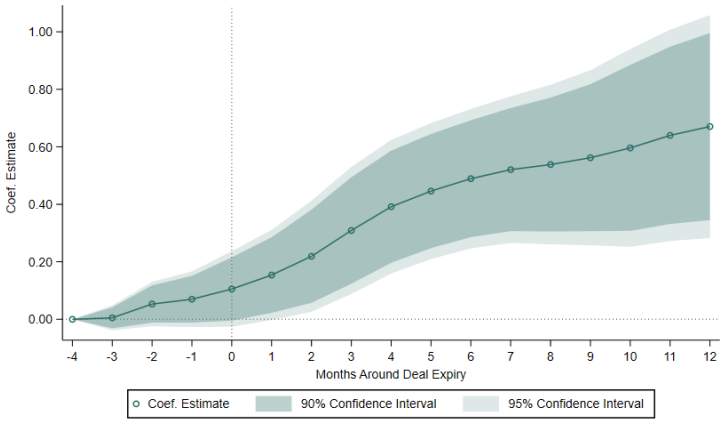
We can combine our estimates of the cash-on-hand and consumption response to estimate the marginal propensity to consume out of the cash-on-hand increase induced by rate cuts. The MPC is the ratio of the response of consumption to the response of cash-on-hand. Equivalently, one can regress the consumption response on the cash-on-hand response, while instrumenting for cash-on-hand with rate changes.

Figure 5 presents estimates of the MPC. We regress the change in consumption from 3 months before until h months after deal expiry, on the change in cash-on-hand from the 3 months before until the 12 months after deal expiry. The change in cash on hand is instrumented with the household's rate change as their deal expires. The figure plots the response of coefficient for every horizon h from 3 months before until 12 months after deal expiry. As such, we trace out the horizon-specific response of consumption to the cumulative change in cash-on-hand due to rate cuts. The MPC gradually increases, slightly exceeding

[16] Columns 1-4 weight the regressions using the sampling weights. Columns 5-8 weight using the product of the sampling weight and mortgage principal.

[17] We test for asymmetry and state dependence, by estimating the consumption response before and after 2022, i.e. during periods of falling and then rising interest rates (see Figure 2). We do not detect strong asymmetries or state dependence—the household debt channel is similar before and after 2022 (Appendix Table A.1).

Figure 5: Marginal Propensity to Consume out of Cash-on-Hand



Notes: the figure plots horizon-specific IV estimates of the marginal propensity to consume (MPC). For each horizon $h \in [-3, 12]$ months relative to deal expiry ($t = 0$), we regress the cumulative change in consumption from months -3 through h on the cumulative change in cash-on-hand from months -3 through $+12$, instrumenting the cash-on-hand change with the household-specific rate change at expiry. The y-axis reports the MPC coefficient (pounds per pound; unitless); the x-axis is event time h in months. Point estimates are shown with 90% and 95% confidence intervals.

0.6 after 12 months.

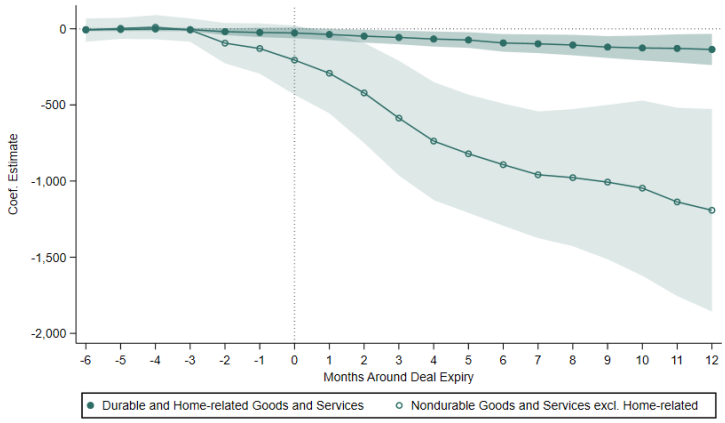
Table 4 presents the corresponding estimates of the MPC over 6 months. Column (1) is the baseline estimate. Over 6 months, households spend 52 pence out of every pound increase in cash-on-hand due to rate cuts. The estimates are similar or larger when we isolate the different sources of variation. The cumulative marginal propensity to consume that we have estimated may seem large—compared, for instance, to estimates of the marginal propensity to consume out of transitory income shocks (e.g. Johnson et al., 2006). However there are several differences which raise the MPC in our setting: our estimates include both durables and non-durable consumption, the shock is relatively persistent, and also involves changes in relative prices.[18]

Next, we exploit an unusual feature of our dataset—we can observe not only consumption, but also how households accumulate liquid assets by for instance holding funds in checking accounts, or paying off credit cards and personal loans. Columns (5)-(8) of Table 4 repeat the previous specifications, but study as an outcome the marginal propensity to either consume or save into liquid assets (what we call the marginal propensity to expend). The estimates are close to one with month or household fixed effects, meaning that households use the cash-on-hand changes from their mortgage either for consumption or for liquid assets.

We also exploit another advantage of our dataset, that we can disaggregate the con-

[18] In order to maximize power, we estimate the MPC via a two sample two stage least squares procedure (e.g. Angrist and Krueger, 1992; Inoue and Solon, 2010). We estimate the first stage relating cash-on-hand to rate cuts on the full population of mortgage borrowers, and the reduced form relating consumption to rate cuts on the subsample for which consumption is available. This procedure is valid if both populations respond similarly to rate cuts; we provide evidence for this proposition below.

Figure 6: Impulse Response of Consumption Components



Notes: the figure plots cumulative impulse responses for durables and non-durables consumption from the baseline monthly regression. For horizons $h \in [-6, 12]$ relative to deal expiry ($t = 0$), the series show the cumulative change in the outcome (pounds, 2015 base) per 1 percentage-point change in the interest rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. Shaded bands denote 90% and 95% confidence intervals. Sample: reweighted consumption for households with app-based spending; durables include home-related renovation expenses.

sumption response by durable versus non-durable consumption.[19] We separately estimate the response of the two consumption categories to the rate cut, and report the result in the impulse response of Figure 6. Evidently, the consumption response is concentrated in non-durables.

We next directly explore the link to monetary policy and consider how households respond to changes in mortgage rates that were caused by monetary policy. To identify monetary policy shocks we use the monthly series of Braun et al. (2025), which considers high frequency monetary policy surprises for the UK around Bank of England policy decisions. To cleanly isolate the impact of the monetary policy shocks, and absorb variation in aggregate mortgage rates at the time of deal expiry, we include month fixed effects. We thus compare two households with deals expiring in the same month, with the variation between them driven by the length of their deals, and the additional monetary policy shocks one of them experienced in the past. Table 5 presents the results. Column 1 shows a regression of the change in cash on hand on aggregate mortgage rates. Column 2 instruments the change in aggregate mortgage rates with the cumulative sum of monetary policy shocks experienced during the deal. The results are similar, the instrument is strong, and the first stage has the expected sign (Appendix table A.2). To tighten the identification further, in columns 3 and 4 we compare households with deals expiring in the same month, and whose deals started within 2-3 years of each other.[20] The results are again similar when the mortgage rates

[19] We include home-related renovation expenses in the durables category.

[20] Specifically we created coarse deal length fixed effects by forming two groups: (i) 1 and 2 year deals; (ii) 3 and 5 year deals. We then include coarse deal length x month fixed effects.

are instrumented. Columns 5-8 show similar findings for consumption: overall we find that households respond similarly when the change in mortgage rates was caused by monetary policy.[21]

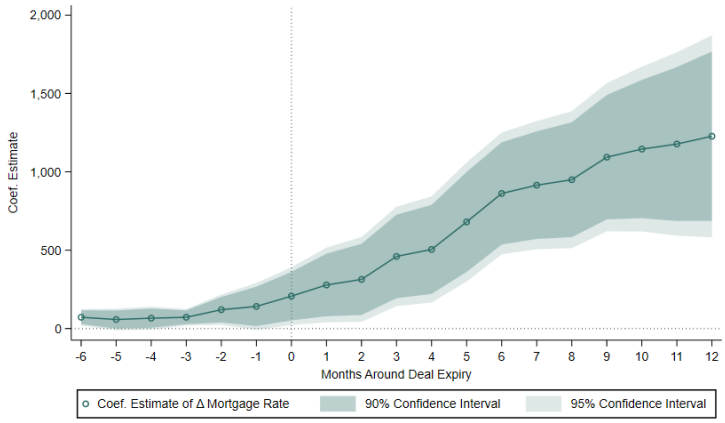
We now discuss the plausibility of our identification assumption. Already we have mentioned two pieces of evidence in its favor. First, there are no pre-trends in either cash-on-hand or consumption. Neither variable starts to rise until the months in which the household starts to refinance the deal, which favors our identification assumption. A confounding shock cannot have happened roughly around the same time as the deal expire—rather, the confounder must occur in precisely the same months. Secondly, estimates are similar across different sources of variation. One can think of confounders that might affect one source of variation. For instance, households of different deal lengths will receive different rate cuts, and may also have different consumption profiles. However this particular concern is eliminated in the specification with deal-by-year fixed effects. In general, shocks that confound all three sources of variation are harder to imagine.

As another assessment of the identification assumption, we study the response of income. We estimate our main regression (5) with monthly income as the outcome variable, as inferred from the spending apps. Figure 7 presents the results. After a rate cut, income falls: after 6 months, a 1 pp rate cut leads to a cumulative income fall of around 1200 pounds. Focusing on income from salary and wages only, these fall by around 800 pounds (see Appendix Figure B.3). This pattern is consistent with standard wealth effects on labor supply. As rates fall and cash-on-hand rises, households choose not only to spend more, but also to work less. If there were confounding income shocks around deal expiry that raise consumption, one would expect the opposite pattern. Income and consumption would rise at the same time—whereas we estimate that income falls when consumption rises.

Finally, we discuss some robustness. One important concern is selection into the consumption dataset. While all results with consumption are reweighted on observables to match the population, one worries about selection on unobservables. Perhaps, for instance, people using the app are more attentive to their finances, meaning their cash on hand responds more to rate cuts. Fortunately, we can partially test this concern by studying how cash-on-hand responds to rate cuts for the consumption sample, versus the rest of the mortgage population. As it turns out, the consumption sample behaves similarly to the rest

[21] Away from monetary policy, we consider a separate instrument for changes in aggregate mortgage rates, which uses long-term shocks to interest rates. These are constructed using unexpected changes in rates over the life of a mortgage deal. To proxy expected future rates on mortgages, we use yields on 2-year government bonds. Our constructed shock is the difference between the 2-year spot yield in the month the mortgage deal ends, less what the 2-year spot yield was expected to be in this month when the mortgage started. Appendix table A.3 shows similar estimates using this alternative instrument, with Appendix table A.4 showing the first stage.

Figure 7: Impulse Response of Income



Notes: the figure plots impulse responses of monthly income from the baseline monthly regression. For horizons $h \in [-6, 12]$ relative to deal expiry ($t = 0$), the series show the cumulative change in income (pounds, 2015 base) per 1 percentage-point change in the interest rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. Shaded bands denote 90% and 95% confidence intervals. Sample: mortgagor households with income inferred from linked spending apps.

(Appendix Table A.5).[22] A second concern is that we only observe a subset of consumption in the financial apps, because consumers only link a subset of their accounts. This concern would likely bias estimates of the MPC downwards, since some of the consumption response would be missed. If this concern is important, then the true MPC is likely to be even higher. Regardless, we saw in Table 4 that we observe the sum of consumption and liquid assets rise by roughly the same amount as cash on hand, meaning we can track where most of the cash-on-hand ends up being allocated.[23]

We now argue that our estimates of the household debt channel are large. Recall that 6 months after a 1pp rate cut, annual consumption increases by 3% relative. The increase is significant. Given the share of mortgagors in the population, and the share of consumption in GDP, GDP increases by 0.6% after 6 months. Moreover the household debt channel is a similar size to consensus estimates of the effect of rate cuts on consumption, in aggregate. For instance, using a long time series and identified monetary shocks, Cloyne et al. (2020) find that the response of consumption to rate cuts concentrates in mortgage holders. Moreover the overall consumption response of mortgage holders is of a similar size to the household debt channel that we have estimated.

As such, the household debt channel appears to be an important part of monetary policy transmission. However, several caveats are in order. First, the effect may be large in part

[22] An exception is the response of cash on hand with month fixed effects, which have notably higher estimates in the consumption sample. This motivates our choice of year fixed effects for our baseline estimates.

[23] We also show that our results are robust to dropping the Pandemic (Appendix Table A.6), using different measures of the aggregate rate (Appendix Table A.7), and using consumption data without reweighting (Appendix Table A.8).

because the shock is persistent, being expected to last for the whole of the next deal for each household. Second, the aggregate effect of rate cuts via the household debt channel will be delayed. Only a fraction of households have deals expiring at any point in time, meaning time must pass before most households experience a given rate cut. Third, there may be other forces that offset the household debt channel. For instance, savers will experience falling income after rate cuts, which offsets the increase in consumption from debtors.[24]

6 The Household Financial Accelerator

So far we have found that the household debt channel of interest rate changes is large—but why? This section shows that one important reason for the large response of consumption and borrowing is the "household financial accelerator" mechanism outlined in Section 4.

6.1 Response of Borrowing vs. Debt Service

As a preliminary finding that suggests the household financial accelerator is important, we study how the two components of cash-on-hand respond to rate cuts. Borrowing responds more than rate cuts, as Figure 8 shows. In the figure, the dark line is the cumulative impulse response of borrowing to a 1 pp interest rise, from 6 months before the deal expires to 12 months afterwards. The light line is the negative of the cumulative impulse response of mortgage payments. After 12 months, the cumulative response of borrowing is slightly larger than the response of monthly payments. At shorter horizons, borrowing is even more important for the response of cash-on-hand.[25]

This finding hints that the household financial accelerator could be important. As we have discussed, the financial accelerator leads households to not only make lower debt service payments but also borrow more after rate cuts. Other prominent mechanisms such as "cashflow effects" primarily operate through changes in debt service and not borrowing.

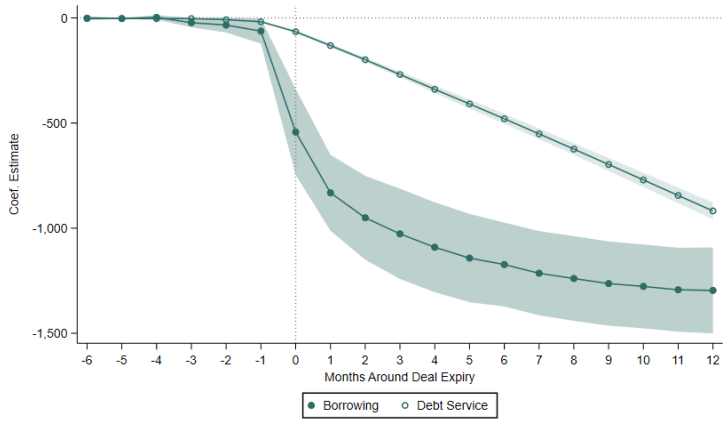
6.2 Identifying the Household Financial Accelerator

We now develop a method to identify the strength of the financial accelerator, with a particular focus on household borrowing. The main idea is to estimate how house prices respond to rate cuts—and to ask whether households with more sensitive house prices also have more responsive borrowing, as the financial accelerator predicts.

[24] If savers have low MPCs, as seems likely, then this offsetting channel will be small (Auclert, 2019).

[25] Appendix Table A.9 reports regressions underlying these results. One reason why payments respond, beyond the purely mechanical effect, is that households respond to rate cuts by lowering the term of the mortgage, as Appendix Table A.10 shows.

Figure 8: Impulse Response of Borrowing and Debt Service



Notes: the figure plots cumulative impulse responses from the baseline monthly regression. For horizons $h \in [-6, 12]$ relative to deal expiry ($t = 0$), the series show the cumulative change in the outcome (pounds, 2015 base) per 1 percentage-point change in the interest rate, expressed relative to the pre-period mean over months $t \in [-6, -4]$. The dark line is the cumulative response of mortgage borrowing. The light line is the negative of the cumulative response debt service, so both series are signed consistently with contributions to cash-on-hand. Shaded bands denote 90% and 95% confidence intervals. Sample: universe of mortgagors.

To start, recall that in our model, the response of borrowing to rate cuts admits the form

$$\frac{\partial d_1}{\partial r_0} = \beta_0 + \beta_{FA}\eta \qquad \eta \equiv \frac{d \log p\left(r_0\right)}{dr_0}.$$

If η is high, so that house prices are sensitive to rates, then households borrow more after rate cuts, with the strength of this effect mediated by the coefficient β_{FA}. If there are "soft" borrowing constraints, the reason is that higher house prices lower mortgage spreads and therefore the price of borrowing. If there are "hard" borrowing constraints, households' mortgage debt is a fixed fraction of the value of their home—-meaning as house prices increase after rate cuts, so too does debt. β_0 captures other reasons why households borrow more after rate cuts, such as intertemporal substitution in response to changes in the aggregate rate (holding fixed house prices).

We can embed this equation into our baseline regression equation (5) in order to estimate the strength of the household financial accelerator. Our main approach uses regional variation. In particular, we can estimate the modified regression equation

$$\Delta_{t,t+h}\text{borrowing}_{it} = \alpha_{j(i)t} + \beta_0\left(r_t - r_{t-n_i}\right) + \beta_{FA} \times \eta_{j(i)}\left(r_t - r_{t-n_i}\right) + \text{fixed effects}_{it} + \varepsilon_{it} \quad (6)$$

where the left-hand side is the cumulative response of borrowing over h months after deal expiry and we have dropped h superscripts from the regression coefficients to remove clutter.

$j(i)$ is the region (Local Authority District, LAD) to which household i belongs.[26] $\alpha_{j(i)t}$ is a region-by-year fixed effect. $\eta_{j(i)}$ is the response of house prices to rate cuts in region j, which we interact with the rate change upon refinancing $(r_t - r_{t-n_i})$. We estimate $\eta_{j(i)}$ in an auxiliary regression of the logarithm of regional house prices on the level of aggregate mortgage rates, with separate estimates for each region. With this variation, one can estimate both β_0, how rate cuts affect borrowing holding fixed house prices; and β_{FA}, how rate cuts affect borrowing via the household financial accelerator. With estimates of both parameters in hand, one can estimate how much of the overall response of borrowing to rate cuts is due to the household financial accelerator. We estimate similar regressions for consumption and debt service payments.

In effect, our approach estimates the household debt channel separately in every region. We then investigate whether the regions with sensitive house prices are also those in which the household debt channel is large. Our regression exploits the following variation. Imagine two regions in the UK, London and York. Suppose that after a period of falling aggregate interest rates, house prices rise in London; whereas house prices do not change in York. If the household financial accelerator is important, then borrowing should respond a lot in London, and little in York.

As we discussed in our prior estimates, some of the variation that identifies the regression comes from differences in deal length. Imagine that in both London and York, there is a first household with a 2-year deal and a second with a 5-year deal. In London, the household with the 2-year deal experiences a small increase in house prices, relative to the household with a 5-year deal. If the financial accelerator is important, then the London household with the 5-year deal should increase their borrowing by more than the household with the 2-year deal, as both households' deals expire. In York, however, neither household experiences rising house prices—meaning the increase in borrowing should be similar and small for both if the financial accelerator matters. As we have discussed, there are two other sources of variation that we have used to estimate the household debt channel. These, too, will vary between London and York if the financial accelerator matters.

One might worry about confounders. Regions that have more sensitive house prices might also behave differently for other reasons that affect the pass-through of rate cuts into borrowing. We stress that our regression includes region-by-time fixed effects. Therefore income shocks, including income shocks that vary by region, cannot bias the result. Nevertheless, there could be more subtle confounders. For instance the elasticity of intertemporal substitution might be different in London compared to York. If so, then borrowing will be different after rate cuts in the two regions, even if the household financial accelerator does

[26] A LAD is roughly equivalent to a county in the United States.

not matter.

We therefore isolate two plausibly exogenous sources of variation in households' sensitivity to rate cuts. First, we use an estimate of the land supply elasticity for each region in the United Kingdom, by adapting the method of Guren, McKay, Nakamura and Steinsson (2021) to the United Kingdom. In regions with inelastic land supply, house prices should respond more to rate cuts. Second, we use a measure of the duration of the home against which each mortgage is collateralized. As Giglio et al. (2015) and Backer-Peral et al. (2024) discuss, many homes in the UK are "freeholds", i.e. long duration assets that are held in perpetuity by the owners. Other properties are "leaseholds", i.e. shorter duration assets that are leased by the owners for many decades, but not forever. By definition, the price of the shorter-duration leasehold properties should respond less to rate cuts: we verify that this indeed holds. Helpfully, there are both leaseholds and freeholds in most regions of the UK, meaning we can isolate within-region variation in the strength of the financial accelerator. We discuss more details about how we isolate each of the two sources of variation as we introduce the results about them.

6.3 Estimates of the Household Financial Accelerator

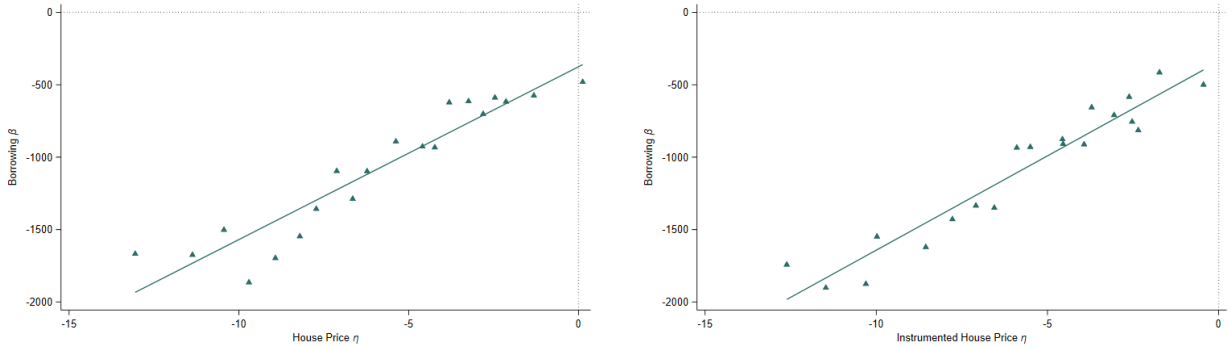
We now present estimates of how much the household financial accelerator matters. We start by implementing a version of our methods non-parametrically, in Figure 9, panel (a). First, we estimate the household debt channel separately for 20 equal-sized bins based on the size of the $\eta_{j(i)}$ estimates, by estimating equation (6) at the 6-month horizon. That is, we replace the $\eta_{j(i)}$ terms with bin indicators for η. Second, for every region we estimate the response of house prices to rate cuts. Finally, we present a scatter plot for the estimates of the 20 bins for the x and y variables.

There is evidence that the household financial accelerator matters for borrowing. In the figure, regions for which house prices respond strongly to rates are on the left side of the x-axis (a more negative value of η means that house prices are more sensitive to rates). These regions are also at the bottom of the y-axis, i.e. borrowing responds more strongly to rates. Regions whose house prices do not respond to rate cuts—such as York, in our example—are on the far right of the x-axis. For these households, the borrowing response to rate cuts is close to zero.

Table 6 presents regression evidence that is consistent with the scatter plot. Column (1) is the baseline specification with year x region fixed effects. The outcome variable is the response of borrowing from 3 months before to 6 months after the deal expires. Our measure of the house price response $\eta_{j(i)}$ is a series of three dummy variables; for the bottom 10%,

Figure 9: Response of Borrowing to Rates vs. Response of House Prices to Rates

(a) OLS Specification **(b)** Land Supply Elasticity Instrument



Notes: each panel plots cross-sectional relationships between the local response of borrowing to interest rates and the local response of house prices to interest rates at the 6-month horizon. The x-axis groups the region $\eta_{j(i)}$ estimates—i.e. the local house-price response per 1 percentage-point change in the interest rate—into 20 equal-sized bins based on the size of the $\eta_{j(i)}$ estimates, and the mean value is plotted for each bin. The y-axis is the estimate of the cumulative change in borrowing (pounds, 2015 base) per 1 pp change in the interest rate, obtained from Eq. (6) with $h = 6$ months, where we estimate β_{FA} for the bin indicators of η. Panel (a): non-parametric estimates for 20 bins. Panel (b): non-parametric estimates for 20 bins of fitted η values, where we instrument η with the land-supply-elasticity instrument and create 20 equal-sized bins based on the size of the predicted η outcomes. Sample: mortgagor households.

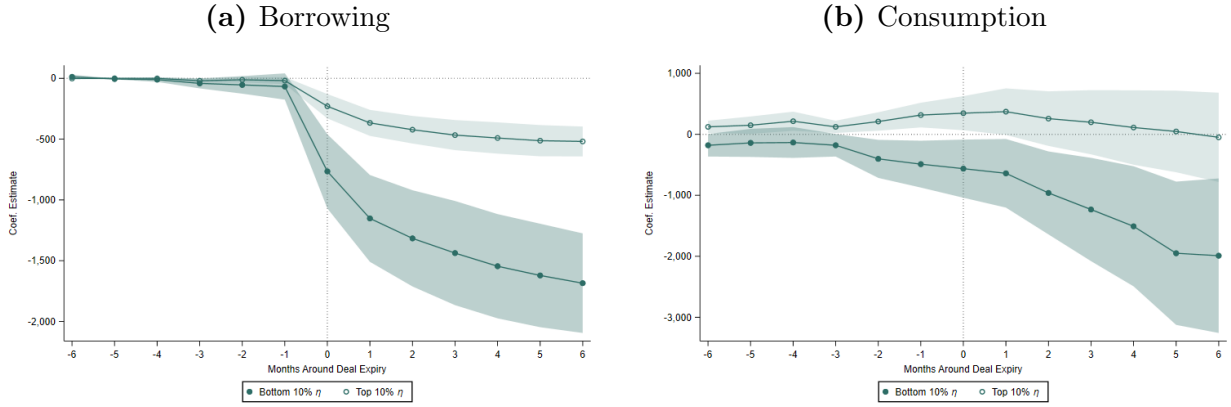
middle 80% and top 10% of the raw values of $\eta_{j(i)}$.[27] The household financial accelerator appears to be important: after a 1 pp rate cut, borrowing increases by roughly 1700 pounds after 6 months for households in the regions with the most sensitive house prices (the bottom 10% of the η distribution: note that η is negative); whereas borrowing increases by only around 500 pounds in the regions with the least sensitive house prices (the top 10% of the η distribution). Figure 10, panel (a) provides an illustration of these differences by showing an impulse response version of the results in column (1). The results are similar in columns (2)-(4), as we add either month, year-by-deal-length-by-region or household fixed effects. In column (5), we use the percent increase in borrowing as the outcome variable. After a 1 pp rate cut, principal rises by 0.98% for households with the most sensitive house prices; whereas principal rises by only 0.56% for the least sensitive households. Columns (6)-(8) show similar results as we add in the additional fixed effects.[28]

How important is the financial accelerator for the overall effect of rate cuts on borrowing? We present a simple calculation that suggests the financial accelerator accounts for around half of the borrowing response to rate cuts. As one can see from Figure 9, households with house prices in the least responsive 10% are those with prices that barely respond, if at

[27] We discretize instead of interacting rates with the value of $\eta_{j(i)}$, because the binscatter of Figure 9, panel (a) suggests that the relationship between house prices responses and borrowing responses is approximately piecewise linear around these thresholds.

[28] Appendix Table A.11 presents similar analysis for the response of cash-on-hand to rate cuts. The response of cash-on-hand is less related to house prices because of the behavior of debt service, as we will discuss.

Figure 10: Impulse Responses by House Price Sensitivity



Notes: the figure plots cumulative impulse responses from a modified version of the baseline monthly regression, where we allow the interest rate response to vary depending on a set of three dummy variables that are defined by the region $\eta_{j(i)}$ estimates—i.e. the local house-price response per 1 percentage-point change in the interest rate. The three dummy variables are for the bottom 10%, middle 80 and top 10% of the $\eta_{j(i)}$ estimates. For horizons $h \in [-6, 6]$ relative to deal expiry ($t = 0$), the series show the cumulative change in the outcome (pounds, 2015 base) per 1 percentage-point change in the interest rate, expressed relative to the pre-period mean over months $t \in [-6, -4]$. For each outcome, we show the estimates for the response to rates corresponding to the bottom 10% and the top 10% of the $\eta_{j(i)}$ estimates. Panel (a): is the cumulative response of mortgage borrowing for the sample of mortgagor households. Panel (b): is the cumulative response of consumption for the sample of reweighted consumption households with app-based spending. Shaded bands denote 90% confidence intervals.

all, to rate cuts (row 3 of Table 6). As such, the estimate of the borrowing response for these households approximates β_0 from regression equation (6), that is, how rate cuts affect borrowing holding fixed house prices. However row 2 of Table 6 measures how rate cuts affect borrowing for the average household. As such, row 2 measures $\beta_0 + \beta_{FA}\overline{\eta}$, where $\overline{\eta}$ is the average house price responsiveness. $\beta_0 + \beta_{FA}\overline{\eta}$ measures the power of the household debt channel on average, exactly as we estimated in Section 5. Therefore using information from rows 1 and 2 we can calculate $\beta_{FA}\overline{\eta}/(\beta_0 + \beta_{FA}\overline{\eta})$, which is the share of the household debt channel that is due to the household financial accelerator. In our baseline specification of column (1), by this calculation the financial accelerator accounts for half of the household debt channel. The estimates range from 35%-68% across the other specifications.

As we have discussed, one potential concern with these estimates is confounding effects—perhaps regions with sensitive house prices differ in other ways that might affect the household debt channel, such as their elasticity of intertemporal substitution. In response, we use two sources of plausible exogenous variation the response of house prices to rates. First, we use the land supply elasticity instrument.

Figure 9, panel (b) presents the scatter plot associated with this variation. As before, we estimate the household debt channel for the 20 bins of η. However, the bins are now determined by the fitted values of $\eta_{j(i)}$, where the fitted values come from an auxiliary cross-sectional regression of η on the land supply elasticity instrument. We again present a scatter plot for the estimates of the 20 bins for the x and y variables. Once more, the household

financial accelerator seems to matter for borrowing. In the figure, regions for which house prices respond strongly to rates are on the left side of the x-axis, and these regions are also at the bottom of the y-axis, i.e. borrowing responds more strongly to rates.

Table 7 presents the results associated with this variation. We categorize households as being in the top 10%, middle 80% or bottom 10% of the fitted value of the response of house prices to rate cuts in each region. Columns (1)-(8) repeat our previous specifications that test for the household financial accelerator, while each row measures how borrowing responses differ by land supply elasticity. As with the baseline estimates, borrowing responds much more in regions with inelastic land supply—whose house prices respond more to rate cuts.

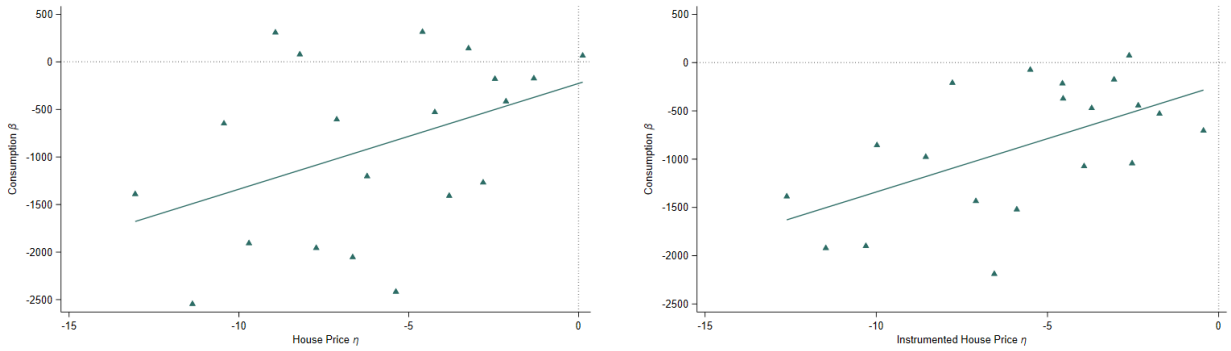
Our second source of variation is the duration of the house stock. Appealingly, duration gives us within-region variation in how house prices respond to changes in rates. To implement this in a simple way, we compare the responses of properties owned *freehold*, which are owned in perpetuity, with properties owned *leasehold*, which are only owned for a finite period of time. Table 8 presents the results. Columns (1)-(4) repeat our previous specifications that test for the household financial accelerator (with the addition of further interacted fixed effects for the duration of the property). Crucially, in these specifications, we are also able to control for the average regional response to rates over the whole sample, absorbing any heterogeneous regional responses that are due to house prices or other factors. Consistently, we find that principal contracts by around 750 pounds more for properties owned freehold. Going further, we can control not just average regional response to rates, but region-specific trends that occurred over a given mortgage deal. Column 5 shows similar estimates when we do this and compare the response of two households in the same local region whose mortgage deals start and finish in exactly the same months.

Throughout, the response of borrowing is greater for the longer duration properties: these are the properties with the greater increase in price in response to rate cuts.[29] To compare the magnitude of this effect with that derived from specifications using the aggregate regional sensitivity $\eta_{j(i)}$, we estimate how borrowing responds to house prices for alternative instruments, with results presented in Appendix table A.12.[30] Column 1 instruments the change in house prices using the interaction of the change in mortgage rates with $\eta_{j(i)}$, finding that a 1% increase in house prices causes a 229 pound increase in borrowing at deal expiry. Columns 2-4 instead instrument with the interaction of the change in rates with a dummy indicating whether the property is freehold or leasehold. To absorb region-specific responses linked to house price sensitivity, column 2 controls for the interaction of the change in mortgage rates with $\eta_{j(i)}$. This also provides a fair comparison to the results in column

[29] Appendix table A.13 provides direct evidence of this house price response for a range of specifications.
[30] First stage results are presented in Appendix table A.13.

Figure 11: Response of Consumption to Rates vs. Response of House Prices to Rates

(a) OLS Specification **(b)** Land Supply Elasticity Instrument



Notes: each panel plots cross-sectional relationships between the local response of consumption to interest rates and the local response of house prices to interest rates at the 6-month horizon. The x-axis groups the region $\eta_{j(i)}$ estimates—i.e. the local house-price response per 1 percentage-point change in the interest rate—into 20 equal-sized bins based on the size of the $\eta_{j(i)}$ estimates, and the mean value is plotted for each bin. The y-axis is the estimate of the cumulative change in borrowing (pounds, 2015 base) per 1 pp change in the interest rate, obtained from Eq. (6) with $h = 6$ months, where we estimate β_{FA} for the bin indicators of η. Panel (a): non-parametric estimates for 20 bins. Panel (b): non-parametric estimates for 20 bins of fitted η values, where we instrument η with the land-supply-elasticity instrument and create 20 equal-sized bins based on the size of the predicted η outcomes. Sample: reweighted consumption households with app-based spending.

1, & we find a very similar estimated effect with this specification. Columns 3-4 tighten the fixed effect specification, fully absorbing $\eta_{j(i)} * (r_t - r_{t-n_i})$, finding similar or larger effects. Overall, the estimated impact of the household financial accelerator on borrowing is at least as strong when exploiting this within-region variation.

The household financial accelerator matters not only for borrowing, but also for consumption. To see this point, we repeat the previous analysis of regression equation (6), with consumption as an outcome variable. We start with non-parametric evidence from a scatter plot, in Figure 11, panel (a). As with borrowing, regions for which house prices respond strongly to rates are on the left-hand side of the x-axis. These regions are also at the bottom of the y-axis. Regions whose house prices respond strongly to rate cuts also have a larger consumption response.

Table 9 contains the regression evidence underlying the scatter plot. In the baseline specification of column (1), the outcome is the response of consumption, from 3 months before until 6 months after deal expiry. Again, we discretize the house price response into three categories, corresponding to the bottom 10%, middle 80%, and top 10%. The household financial accelerator seems to matter: for households in the most responsive regions (top row), a 1 pp rate cut raises consumption by almost 2000 pounds; whereas for households in the least responsive regions (bottom row), consumption responds relatively little (although the estimates are imprecise). Figure 10, panel (b) visualizes these differences by presenting an impulse response version of the results in column (1). The results are similar, although occasionally imprecisely estimated, as we add in various combinations of fixed effects (columns

2-4), and repeat the analysis with consumption growth as the outcome variable (columns 5-8).

The household financial accelerator accounts for a meaningful share of the overall response of consumption to rate cuts. We extend our previous analysis, by comparing the effect of rate cuts on consumption for households with the least responsive house prices, to the average household. Our estimates with consumption are somewhat imprecise. Nevertheless, applying this calculation to each column of Table 9, we find that around 50% of the response of consumption to rate cuts is due to the financial accelerator.

Finally, we study the response of debt service—which does not appear to be greatly affected by the household financial accelerator. In particular, we repeat the analysis of regression equation (6), with a new outcome variable, the cumulative response of debt service payments. Table 10 presents the results. In columns (1)-(4) we present the level response of debt service payments. In columns (5)-(8) we present the log response of debt service payments. We can repeat the simple calculation to measure how much of the response of debt service is accounted for by the financial accelerator, which across the specifications ranges from 10% to 35%.[31] As such, the financial accelerator appears to matter less for debt service than for consumption and borrowing.[32]

We close this section by discussing the implications for models. As we have discussed, models of the household financial accelerator typically take two forms. First, there might be "soft constraints", as in Bernanke et al. (1999), which affect both borrowing and debt service. After rate increases, the higher house prices lead to lower loan-to-value ratios, that reduce mortgage spreads. As a result the household borrows more, and faces lower debt service costs. Second, there might be "hard constraints", as in Kiyotaki and Moore (1997) and Iacoviello (2005). Households can borrow at most a fixed fraction of their home

[31] Appendix table A.14 provides analysis of the impact of the financial accelerator on individual mortgage interest rates, which in turn affect the response of debt service. Columns 1-4 look at the level response of interest rates, which theoretically depend on the interaction of the log change in house prices and the initial LTV ratio. The estimates in column 1 imply that when the initial LTV ratio is 70%, the interest rate response does not depend on the regional house price sensitivity, consistent with the relatively flat interest rate schedule shown at this LTV ratio for the UK by Best et al. (2020). By contrast, the estimates imply that when the initial LTV ratio is 90%, the individual rate increases by 2 basis points more between the most and least responsive house price regions (although this difference is not statistically significant). Columns 5-8 show the response of the log change in interest rates, which should approximately vary with the log change in house prices. Column 5 shows that, following a 100bps increase in aggregate mortgage rates, individual rates increase by around 11% more in the regions with the most responsive house prices (39%) compared to the response in regions with the least responsive houses prices (34%), with this difference statistically significant.

[32] There are two nuances. First, after rate cuts, borrowing increases at the same time that house prices fall. Therefore the loan-to-value ratio could respond by little even if house prices respond by a lot, meaning that mortgage spreads change by little. Second, borrowing increases due to rate cuts, which all else equal raises debt service.

value. When home value rises, borrowing increases in proportion—while mortgage spreads are unchanged. Our evidence appears to favor the latter model. Movements in debt service appear to be relatively unimportant, whereas movements in borrowing are critical.

7 Conclusion

This paper asks whether the "household debt channel"—the transmission of rate cuts into consumption via household debt—is large.

To measure the household debt channel, we combine millions of natural experiments with individual level data on mortgages and consumption. The natural experiments arise from the staggered refinancing of mortgage deals, and the individual level data comes from a merge of administrative mortgage data and spending from financial apps.

In order to identify the household debt channel we regress the change in households' consumption, as their deal expires, on the change in aggregate mortgage rates between the start and the end of their deal. We study the response of cash-on-hand—the sum of changes in borrowing and debt service—in a similar way. Our design uses multiple sources of variation, including within the same household over time.

Our main result is that the household debt channel is large. After a 1pp rate cut, consumption increases by 3% for households with mortgage deals expiring, which aggregates to 0.6% of GDP. The consumption response via the household debt channel is of a similar size to the effect of aggregate monetary policy shocks on consumption in the UK (e.g. Cloyne, Ferreira and Surico, 2020).

Finally, we investigate why the household debt channel is so large, with a focus on the "household financial accelerator". Consistent with this mechanism, borrowing accounts for around two thirds of the response of cash-on-hand after rate cuts. We then introduce a strategy to pin down how much the financial accelerator matters. We use plausibly exogenous variation in how households' home values respond to rate cuts, based on the regional land supply elasticity and the duration of the housing stock. The financial accelerator accounts for roughly 50% of the borrowing and consumption response, though it matters less for debt service payments.

Tables

Table 1: Summary Statistics for Mortgage Population and Consumption Sample

	Mortgage Population		Unweighted Consumption Sample		Weighted Consumption Sample	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Mortgage Variables						
Mortgage Debt (£)	147,656.79	115,319.32	145,442.45	100,006.55	147,399.82	109,778.88
Home Value (£)	295,419.69	234,737.30	250,849.38	173,210.54	285,130.25	203,292.39
Mortgage Payment (£)	716.95	478.35	672.18	420.47	721.17	469.78
Mortgage Term (months)	252.48	99.36	285.18	90.91	256.21	97.40
Interest Rate (%)	2.34	0.82	2.40	0.84	2.32	0.81
Loan-to-value Ratio (%)	53.80	21.20	60.52	18.78	54.51	20.87
Loan-to-income Ratio	2.89	1.12	3.00	1.00	2.88	1.08
Debt-service Ratio (%)	17.56	6.44	16.96	5.48	17.55	6.22
Income (£)	58,424.47	48,968.15	55,496.83	36,535.74	58,472.90	40,969.75
Age	42.50	9.83	39.64	8.72	42.21	9.55
London Dummy Variable (%)	12.55	33.13	8.14	27.34	12.82	33.44
Consumption Variables						
Total Consumption (£)			2,457.25	2,907.43	2,567.47	3,041.63
Durable and Home-related (£)			133.07	587.06	139.55	640.52
Nondurables and Services excl. Home-related (£)			2,175.18	2,617.94	2,269.26	2,710.75
Observations	6,796,488	6,796,488	68,734	68,734	68,734	68,734

Notes: The table shows summary statistics for our datasets. In the first two columns we report summary statistics for observations with mortgage information, i.e. for the universe of mortgages from June 2015 to March 2024. In the middle two columns we report summary statistics for the observations with both mortgage and consumption information, i.e. for households that merge between the mortgage dataset (PSD) and the consumption dataset (ClearScore or MoneyDashboard). The final two columns present the consumption sample, but reweighted based on quintiles of age, income, loan-to-value, mortgage term and home value, as well as broad region (NUTS1) information, in order to match the mortgage sample. In the rows we present summary statistics on information, for either mortgages or consumption; including both the mean and standard deviation of variables. The observations are at the level of an expiring deal, i.e. we measure the characteristics of a household as its deal expires. The units are 2015 pounds.

Table 2: Effect of Mortgage Rates on Cash on Hand

	Δ Cash on Hand				Δ Log Cash on Hand			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1818.206***	-1816.882***	-1773.086***	-2272.775***	-1.440***	-1.827***	-1.278***	-1.869***
	(104.530)	(114.602)	(139.562)	(213.965)	(0.077)	(0.105)	(0.097)	(0.208)
Observations	6796488	6796488	6796488	3681291	6796379	6796379	6796379	3681212
Adjusted R^2	0.011	0.012	0.011	0.073	0.013	0.013	0.013	-0.024
Year FE	✓			✓	✓			✓
Month FE		✓				✓		
Deal Length x Year FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on the change in cash on hand, in the event window around the end of the mortgage deal. Results are shown for the full mortgage sample. In this table the change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal, with year and household age fixed effects. Column 2 includes month instead of year fixed effects. Column 3 includes deal length x year and household age fixed effects. Column 4 includes year and household fixed effects. Columns 5-8 repeat this with the change in log cash on hand as the outcome variable. This is defined as the ratio of the change in cash on hand, over the event window, to the mortgage balance 7 months prior to deal expiry. This ratio is multiplied by 100 to give a percentage interpretation. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table 3: Effect of Mortgage Rates on Consumption

	Δ Consumption				Δ Log Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-937.772***	-1364.088***	-834.740**	-1814.168***	-2.994***	-2.554**	-3.018**	-7.534***
	(277.406)	(388.150)	(320.379)	(584.804)	(1.028)	(0.998)	(1.217)	(2.073)
Observations	68734	68734	68734	17614	68734	68734	68734	17614
Adjusted R^2	0.003	0.011	0.003	0.040	0.006	0.025	0.006	0.022
Year FE	✓			✓	✓			✓
Month FE		✓				✓		
Deal Length x Year FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on the cumulative change in consumption, in the event window around the end of the mortgage deal. Results are shown for the matched mortgage-consumption sample. Column 1 shows the regression of the change in consumption on the change in aggregate mortgage rates over the deal, with year and household age fixed effects. Column 2 includes month instead of year fixed effects. Column 3 includes deal length x year and household age fixed effects. Column 4 includes year and household fixed effects. Columns 5-8 repeat this with the change in log consumption as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 1-4 are weighted by sampling weights. The regressions in columns 5-8 are weighted by the product of the sampling weight and the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table 4: Marginal Propensity to Consume and Expend Out of Cash on Hand

	Marginal Propensity to Consume				Marginal Propensity to Expend			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	0.516***	0.751***	0.471**	0.798***	0.549***	1.120***	0.414**	0.974***
	(0.153)	(0.214)	(0.181)	(0.257)	(0.162)	(0.393)	(0.180)	(0.324)
Second-Stage Observations	68734	68734	68734	17614	68734	68734	68734	17614
Adjusted R^2	0.003	0.011	0.003	0.040	0.002	0.004	0.002	0.054
First-Stage F-Stat	302.555	251.343	161.407	112.830	302.555	251.343	161.407	112.830
Year FE	✓			✓	✓			✓
Month FE		✓				✓		
Deal Length x Year FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the cumulative marginal propensity to consume, and expend, out of the change in cash on hand, in the event window around the end of the mortgage deal. In this table the change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. The results use the cumulative change in consumption or expenditure during the event window from 3 months before, to 6 months after the month of deal expiry. The regressions are estimated via Two-Sample Two-Stage Least Squares. The first stage regresses the change in cash in hand on the change in aggregate mortgage rates, in the full mortgage sample. The second stage regresses the average monthly change in consumption or expenditure on the fitted change in cash on hand from the first stage. The second stage regression is run in the matched mortgage-consumption sample. Columns 1-4 show results for the marginal propensity to consume. Column 1 includes year and household age fixed effects. Column 2 includes month and household age fixed effects. Column 3 includes deal length x year and household age fixed effects. Column 4 includes year and household fixed effects. Columns 5-8 show results for the marginal propensity to expend. Trigger month clustered standard errors are shown in parentheses.

Table 5: Effect of Mortgage Rates on Cash on Hand and Consumption: High Frequency IV

	Δ Cash on Hand				Δ Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1816.882***	-2320.410***	-1659.478***	-2073.929***	-1364.088***	-1313.654**	-1711.390**	-1339.903
	(114.602)	(193.350)	(171.015)	(310.854)	(388.150)	(607.004)	(738.344)	(1837.462)
Observations	6796488	6796488	6748651	6748651	68734	68734	68513	68513
Adjusted R^2	0.012	0.000	0.012	0.000	0.011	0.000	0.011	0.000
First-Stage F-Stat		290.056		189.108		194.095		67.984
Month FE	✓	✓			✓	✓		
Coarse Deal Length x Month FE			✓	✓			✓	✓
Age FE	✓	✓	✓	✓	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of high-frequency-identified monetary policy shocks on the change in cash on hand, and consumption, in the event window around the end of the mortgage deal. Columns 1-4 show results for the change in cash on hand, for the full mortgage sample. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal, with month and household age fixed effects. Column 2 instruments the change in aggregate mortgage rates over the deal with the sum of high-frequency-identified monetary policy shocks over the deal. Columns 3-4 instead include month x coarse deal length fixed effects. The coarse deal length fixed effects group together 1 and 2 year deals, and separately, 3 and 5 year deals. The regressions in columns 3-4 restrict to mortgage deals that begin from 2013 to 2021. Columns 5-8 repeat this exercise with the cumulative change in consumption as the outcome variable. These consumption regressions are run on the matched mortgage-consumption sample and are weighted by sampling weights. Trigger month clustered standard errors are shown in parentheses.

Table 6: Effect of Financial Accelerator on Mortgage Principal

	Δ Principal				Δ Log Principal			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1671.157***	-1671.796***	-1802.373***	-2617.762***	-0.983***	-0.717***	-1.162***	-1.930***
	(243.905)	(197.227)	(290.834)	(379.171)	(0.170)	(0.115)	(0.179)	(0.257)
ΔMortgage Rate x 𝕀(Middle 80% η)	-1092.502***	-1083.945***	-1090.033***	-1550.222***	-0.888***	-0.623***	-0.986***	-1.588***
	(112.113)	(109.571)	(141.686)	(221.915)	(0.120)	(0.085)	(0.133)	(0.196)
ΔMortgage Rate x 𝕀(Top 10% η)	-531.311***	-523.892***	-498.313***	-498.096***	-0.558***	-0.294***	-0.638***	-0.831***
	(71.626)	(102.632)	(99.643)	(147.401)	(0.110)	(0.091)	(0.125)	(0.165)
Observations	6796438	6796438	6796433	3681274	6793800	6793800	6793795	3679515
Adjusted R^2	0.012	0.012	0.013	0.075	0.014	0.014	0.015	0.061
Year x LAD FE	✓	✓		✓	✓	✓		✓
Month FE		✓				✓		
Deal Length x Year x LAD FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on mortgage principal, in the event window around the end of the mortgage deal. Each regression interacts the change in aggregate mortgage rates over the deal with dummies for three distinct groups: (i) households in the 10% of regions with the most sensitive house prices to changes in aggregate rates; (ii) households in the middle 80% of regions by sensitivity, and (iii) households in the 10% of regions with the least sensitive house prices to changes in aggregate rates. In columns 1-4 the dependent variable is the change in principal in the event window around the end of the deal. Column 1 includes year x Local Authority District (LAD), and household age fixed effects. Column 2 further adds a month fixed effect. Column 3 includes year x deal length x LAD and household age fixed effects. Column 4 includes year x LAD and household fixed effects. Columns 5-8 repeat this with the change in log principal as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table 7: Effect of Financial Accelerator on Mortgage Principal: Land Supply-Instrumented η

	Δ Principal				Δ Log Principal			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1818.784***	-1816.920***	-2063.571***	-2833.819***	-1.010***	-0.746***	-1.274***	-1.985***
	(251.844)	(211.010)	(276.406)	(362.873)	(0.189)	(0.134)	(0.184)	(0.251)
ΔMortgage Rate x 𝕀(Middle 80% η)	-1069.238***	-1059.618***	-1051.641***	-1508.013***	-0.889***	-0.624***	-0.967***	-1.572***
	(112.934)	(108.280)	(146.379)	(221.450)	(0.118)	(0.084)	(0.135)	(0.194)
ΔMortgage Rate x 𝕀(Top 10% η)	-543.330***	-536.971***	-522.667***	-575.003***	-0.467***	-0.202**	-0.591***	-0.858***
	(68.443)	(112.055)	(85.024)	(168.753)	(0.112)	(0.092)	(0.116)	(0.193)
Observations	6796438	6796438	6796433	3681274	6793800	6793800	6793795	3679515
Adjusted R^2	0.012	0.012	0.013	0.075	0.014	0.014	0.015	0.061
Year x LAD FE	✓	✓		✓	✓	✓		✓
Month FE		✓				✓		
Deal Length x Year x LAD FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on mortgage principal, in the event window around the end of the mortgage deal. Throughout, the house price η for each region is instrumented, building on the instrumental variable approach of Guren et al. (2021). Each regression interacts the change in aggregate mortgage rates over the deal with dummies for three distinct groups: (i) households in the 10% of regions with the most sensitive house prices to changes in aggregate rates; (ii) households in the middle 80% of regions by sensitivity, and (iii) households in the 10% of regions with the least sensitive house prices to changes in aggregate rates. In columns 1-4 the dependent variable is the change in principal in the event window around the end of the deal. Column 1 includes year x Local Authority District (LAD), and household age fixed effects. Column 2 further adds a month fixed effect. Column 3 includes year x deal length x LAD and household age fixed effects. Column 4 includes year x LAD and household fixed effects. Columns 5-8 repeat this with the change in log principal as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table 8: Effect of Financial Accelerator on Mortgage Principal: Duration Variation

	Δ Principal					Δ Log Principal				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ΔMortgage Rate x 𝕀(F.hold)	-724.115***	-720.710***	-764.521***	-828.606***	-804.211***	-0.508***	-0.519***	-0.557***	-0.431***	-1.071***
	(82.117)	(85.093)	(83.593)	(148.673)	(221.484)	(0.081)	(0.081)	(0.088)	(0.142)	(0.194)
Observations	4525675	4525675	4525664	2465329	4296642	4524225	4524225	4524214	2464290	4295200
Adjusted R^2	0.013	0.014	0.014	0.080	0.028	0.014	0.015	0.016	0.063	0.045
Year x LAD FE	✓	✓		✓		✓	✓		✓	
Month FE		✓					✓			
Deal Length x Year x LAD FE			✓					✓		
Year x 𝕀(F.hold) FE	✓	✓	✓	✓		✓	✓	✓	✓	
Lad x ΔMortgage Rate FE	✓	✓	✓	✓		✓	✓	✓	✓	
Start M. x Trig M x LAD FE					✓					✓
Month x 𝕀(F.hold) FE					✓					✓
Household FE				✓					✓	
Age FE	✓	✓	✓		✓	✓	✓	✓		✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on mortgage principal, in the event window around the end of the mortgage deal. Each regression interacts the change in aggregate mortgage rates over the deal with a dummy for whether a property is owned freehold or not. In columns 1-5 the dependent variable is the change in principal in the event window around the end of the deal. Column 1 includes the following fixed effects: (i) year x Local Authority District (LAD); (ii) household age; (iii) LAD x ΔMortgage Rate ; and (iv) year x 𝕀(Freehold), where 𝕀(Freehold) is a dummy for whether the property is owned freehold or not. Column 2 further includes a month fixed effect. Column 3 replace year x LAD fixed effects with year x deal length x LAD fixed effects. Column 4 replaces the household age fixed effect of column 1 with a household fixed effect. Column 5 includes the following fixed effects: (i) deal start month x deal expiry month x LAD; (ii) month x 𝕀(Freehold); and (iii) household age. Columns 6-10 repeat this with the change in log principal as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 6-10 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table 9: Effect of Financial Accelerator on Consumption

	Δ Consumption				Δ Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1946.229**	-2118.370**	-1526.866	-2800.440	-5.269***	-3.981**	-4.619**	-18.109***
	(761.825)	(859.724)	(921.325)	(2317.496)	(1.771)	(1.708)	(2.115)	(5.019)
ΔMortgage Rate x 𝕀(Middle 80% η)	-845.887***	-1012.382***	-864.150**	-1864.016***	-2.248**	-0.909	-2.677**	-6.285***
	(295.731)	(368.754)	(349.411)	(613.659)	(1.067)	(1.000)	(1.251)	(2.252)
ΔMortgage Rate x 𝕀(Top 10% η)	-64.910	-234.389	-399.141	-1056.824	-1.003	0.320	-1.446	-4.744
	(421.604)	(473.896)	(490.039)	(1137.919)	(1.901)	(1.934)	(2.089)	(3.430)
Observations	68541	68541	67078	16942	68541	68541	67078	16942
Adjusted R^2	0.021	0.029	0.030	0.114	0.036	0.055	0.054	0.118
Year x LAD FE	✓	✓		✓	✓	✓		✓
Month FE		✓				✓		
Deal Length x Year x LAD FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on cumulative consumption, in the event window around the end of the mortgage deal. Results are shown for the matched mortgage-consumption sample. Each regression interacts the change in aggregate mortgage rates over the deal with dummies for three distinct groups: (i) households in the 10% of regions with the most sensitive house prices to changes in aggregate rates; (ii) households in the middle 80% of regions by sensitivity, and (iii) households in the 10% of regions with the least sensitive house prices to changes in aggregate rates. In columns 1-4 the dependent variable is the cumulative change in consumption in the event window around the end of the deal. Column 1 includes year x Local Authority District (LAD), and household age fixed effects. Column 2 further adds a month fixed effect. Column 3 includes year x deal length x LAD and household age fixed effects. Column 4 includes year x LAD and household fixed effects. Columns 5-8 repeat this with the change in log consumption as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 1-4 are weighted by sampling weights. The regressions in columns 5-8 are weighted by the product of the sampling weight and the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table 10: Effect of Financial Accelerator on Monthly Payments

	−ΔPayments				−ΔLog Payments			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1016.289***	-954.502***	-996.284***	-951.326***	-13.070***	-13.411***	-12.323***	-13.056***
	(51.047)	(30.831)	(65.283)	(66.221)	(0.889)	(0.434)	(1.094)	(1.019)
ΔMortgage Rate x 𝕀(Middle 80% η)	-711.906***	-654.273***	-707.583***	-695.442***	-12.441***	-12.796***	-11.918***	-12.500***
	(37.818)	(25.978)	(48.950)	(44.523)	(0.800)	(0.427)	(1.020)	(0.907)
ΔMortgage Rate x 𝕀(Top 10% η)	-479.957***	-425.146***	-485.001***	-515.548***	-10.636***	-11.070***	-10.566***	-11.196***
	(24.551)	(26.132)	(32.609)	(28.511)	(0.580)	(0.450)	(0.770)	(0.781)
Observations	6796438	6796438	6796433	3681274	6756388	6756388	6756382	3649175
Adjusted R^2	0.128	0.132	0.133	0.088	0.167	0.173	0.171	0.102
Year x LAD FE	✓	✓		✓	✓	✓		✓
Month FE		✓				✓		
Deal Length x Year x LAD FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on cumulative monthly payments, in the event window around the end of the mortgage deal. Results are shown for the full mortgage sample. Each regression interacts the change in aggregate mortgage rates over the deal with dummies for three distinct groups: (i) households in the 10% of regions with the most sensitive house prices to changes in aggregate rates; (ii) households in the middle 80% of regions by sensitivity, and (iii) households in the 10% of regions with the least sensitive house prices to changes in aggregate rates. In columns 1-4 the dependent variable is −10 times the change in monthly payments, in the event window around the end of the mortgage deal.. Column 1 includes year x Local Authority District (LAD), and household age fixed effects. Column 2 further adds a month fixed effect. Column 3 includes year x deal length x LAD and household age fixed effects. Column 4 includes year x LAD and household fixed effects. Columns 5-8 repeat this with the negative of the change in log monthly payments as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

References

Angrist, Joshua D and Alan B Krueger, "The effect of age at school entry on educational attainment: an application of instrumental variables with moments from two samples," *Journal of the American statistical Association*, 1992, *87* (418), 328–336.

Auclert, Adrien, "Monetary policy and the redistribution channel," *American Economic Review*, 2019, *109* (6), 2333–2367.

Backer-Peral, Veronica, Jonathon Hazell, and Atif Mian, "Dynamics of the Long Term Housing Yield: Evidence from Natural Experiments," *Working Paper*, 2024.

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*, 2021, *111* (9), 2829–2878.

Bernanke, Ben S., Mark Gertler, and Simon Gilchrist, "The financial accelerator in a quantitative business cycle framework," in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1 of *Handbook of Macroeconomics*, Elsevier, 1999, chapter 21, pp. 1341–1393.

Best, Michael Carlos, James S Cloyne, Ethan Ilzetzki, and Henrik J Kleven, "Estimating the elasticity of intertemporal substitution using mortgage notches," *The Review of Economic Studies*, 2020, *87* (2), 656–690.

Braun, Robin, Silvia Miranda-Agrippino, and Tuli Saha, "Measuring monetary policy in the UK: The UK monetary policy event-study database," *Journal of Monetary Economics*, 2025, *149*, 103645.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico, "Monetary policy when households have debt: new evidence on the transmission mechanism," *The Review of Economic Studies*, 2020, *87* (1), 102–129.

_ , Kilian Huber, Ethan Ilzetzki, and Henrik Kleven, "The effect of house prices on household borrowing: A new approach," *American Economic Review*, 2019, *109* (6), 2104–2136.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong, "State-dependent effects of monetary policy: The refinancing channel," *American Economic Review*, 2022, *112* (3), 721–761.

Flodén, Martin, Matilda Kilström, Jósef Sigurdsson, and Roine Vestman, "Household Debt and Monetary Policy: Revealing the Cash-Flow Channel," *The Economic Journal*, 12 2020, *131* (636), 1742–1771.

Gertler, Mark and Peter Karadi, "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 44–76.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel, "Very Long-Run Discount Rates," *The Quarterly Journal of Economics*, February 2015, *130* (1), 1–53.

Greenwald, Daniel, "The mortgage credit channel of macroeconomic transmission," 2018.

Guren, Adam M, Alisdair McKay, Emi Nakamura, and Jón Steinsson, "Housing wealth effects: The long view," *The Review of Economic Studies*, 2021, *88* (2), 669–707.

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek, "The transmission of monetary policy under the microscope," *Journal of Political Economy*, 2021, *129* (10), 2861–2904.

Iacoviello, Matteo, "House prices, borrowing constraints, and monetary policy in the business cycle," *American economic review*, 2005, *95* (3), 739–764.

Inoue, Atsushi and Gary Solon, "Two-sample instrumental variables estimators," *The Review of Economics and Statistics*, 2010, *92* (3), 557–561.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

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

Kiyotaki, Nobuhiro and John Moore, "Credit Cycles," *Journal of Political Economy*, April 1997, *105* (2), 211–248.

Laibson, D, P Maxted, and B Moll, "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy," *Quarterly Journal of Economics*, forthcoming.

Maggio, Marco Di, 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), 3550–3588.

Romer, Christina D and David H Romer, "A new measure of monetary shocks: Derivation and implications," *American economic review*, 2004, *94* (4), 1055–1084.

Valliant, Richard and Jill A. Dever, *Survey Weights: A Step-by-step Guide to Calculation*, College Station, TX: Stata Press, 2018.

A Appendix Tables

Table A.1: Effect of Mortgage Rates Before and After Hiking Cycle

	Δ Cash on Hand	Δ Principal	−ΔPayments	Δ Consumption
	(1)	(2)	(3)	(4)
ΔMortgage Rate x 𝕀(Pre 2022)	-1777.717***	-1163.160***	-614.557***	-1186.017**
	(126.394)	(129.394)	(27.569)	(497.440)
ΔMortgage Rate x 𝕀(2022-)	-1837.742***	-1069.643***	-768.099***	-876.227***
	(138.629)	(169.807)	(60.159)	(330.239)
Observations	6796488	6796488	6796488	68734
Adjusted R^2	0.011	0.007	0.111	0.003
Year FE	✓	✓	✓	✓
Age FE	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on several outcome variables, before and after the interest rate hiking cycle. Results for columns 1-3 are shown for the full mortgage sample. In this table the change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal interacted with separate dummies for mortgage deals that expire before 2022, and those that expire from 2022 onward. This regression includes year and household age fixed effects. Columns 2-3 show the response of the sub-components of the change in cash on hand. In column 2 the dependent variable is the change in principal, in the event window around the end of the mortgage deal. In column 3 the dependent variable is −10 times the change in monthly payments, in the event window around the end of the mortgage deal. In column 4 the dependent variable is the cumulative change in consumption, in the event window around the end of the mortgage deal. This regression is run on the matched mortgage-consumption sample and the regression is weighted using the sampling weights. Trigger month clustered standard errors are shown in parentheses.

Table A.2: Effect of High Frequency Identified Monetary Policy Shocks on Mortgage Rates

	Δ Mortgage Rate			
	(1)	(2)	(3)	(4)
∑HF Monetary Policy Shocks	1.970***	1.319***	1.935***	1.148***
	(0.116)	(0.096)	(0.139)	(0.139)
Observations	6796488	6748651	68734	68513
Adjusted R^2	0.981	0.992	0.984	0.993
F-Stat	290.056	189.108	194.095	67.984
Month FE	✓		✓	
Coarse Deal Length x Month FE		✓		✓
Age FE	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of high-frequency-identified monetary policy shocks on the change in aggregate mortgage rates. Column 1 shows a regression of the change in aggregate mortgage rates on the sum of high frequency identified monetary policy shocks over the deal, including month and household age fixed effects. Column 2 instead includes month x coarse deal length fixed effects. The coarse deal length fixed effects group together 1 and 2 year deals, and separately, 3 and 5 year deals. This regression restricts to mortgage deals that begin from 2013 to 2021. Columns 3-4 repeat this for the consumption sample. These consumption regressions are run on the matched mortgage-consumption sample and are weighted by sampling weights. Trigger month clustered standard errors are shown in parentheses.

Table A.3: Effect of Mortgage Rates on Cash on Hand and Consumption: Long Term Rate Shocks IV

	Δ Cash on Hand				Δ Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1816.882***	-2236.023***	-1659.478***	-1792.931***	-1364.088***	-904.250*	-1711.390**	-1081.137
	(114.602)	(168.305)	(171.015)	(309.979)	(388.150)	(502.574)	(738.344)	(1231.941)
Observations	6796488	6733819	6748651	6689897	68734	67921	68513	67737
Adjusted R^2	0.012	0.000	0.012	0.000	0.011	0.000	0.011	0.000
First-Stage F-Stat		358.947		169.414		225.341		89.899
Month FE	✓	✓			✓	✓		
Coarse Deal Length x Month FE			✓	✓			✓	✓
Age FE	✓	✓	✓	✓	✓	✓	✓	✓

$^{\star}p < 0.10, {}^{\star\star}p < 0.05, {}^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of long-term interest rates shocks on the change in cash on hand, and consumption, in the event window around the end of the mortgage deal. Columns 1-4 show results for the change in cash on hand, for the full mortgage sample. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal, with month and household age fixed effects. Column 2 instruments the change in aggregate mortgage rates over the deal with long-term interest rate shocks, which are the difference between government bond yields at the time of refinance, less what they were expected to be (from the yield curve) at the start of the mortgage deal. Columns 3-4 instead include month x coarse deal length fixed effects. The coarse deal length fixed effects group together 1 and 2 year deals, and separately, 3 and 5 year deals. The regressions in columns 3-4 restrict to mortgage deals that begin from 2013 to 2021. Columns 5-8 repeat this exercise with the cumulative change in consumption as the outcome variable. These consumption regressions are run on the matched mortgage-consumption sample and are weighted by sampling weights. Trigger month clustered standard errors are shown in parentheses.

Table A.4: Effect of Long Term Interest Rate Shocks on Mortgage Rates

	Δ Mortgage Rate			
	(1)	(2)	(3)	(4)
Long Term Interest Rate Shocks	0.265***	0.197***	0.244***	0.177***
	(0.014)	(0.015)	(0.016)	(0.019)
Observations	6733819	6689897	67921	67737
Adjusted R^2	0.987	0.994	0.989	0.994
F-Stat	358.947	169.414	225.341	89.899
Month FE	✓		✓	
Coarse Deal Length x Month FE		✓		✓
Age FE	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of long-term interest rates shocks on the change in aggregate mortgage rates. Column 1 shows a regression of the change in aggregate mortgage rates on long-term interest rate shocks, which are the difference between government bond yields at the time of refinance, less what they were expected to be (from the yield curve) at the start of the mortgage deal. This regression includes month and household age fixed effects. Column 2 instead includes month x coarse deal length fixed effects. The coarse deal length fixed effects group together 1 and 2 year deals, and separately, 3 and 5 year deals. This regression restricts to mortgage deals that begin from 2013 to 2021. Columns 3-4 repeat this for the consumption sample. These consumption regressions are run on the matched mortgage-consumption sample and are weighted by sampling weights. Trigger month clustered standard errors are shown in parentheses.

Table A.5: Effect of Mortgage Rates on Cash on Hand: Consumption Sample

	Δ Cash on Hand				Δ Log Cash on Hand			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1918.781***	-3618.830***	-1574.334***	-1580.156**	-1.449***	-3.286***	-1.095***	-1.118*
	(252.587)	(768.936)	(257.800)	(770.884)	(0.194)	(0.575)	(0.172)	(0.623)
Observations	68734	68734	68734	17614	68734	68734	68734	17614
Adjusted R^2	0.009	0.011	0.010	0.112	0.005	0.008	0.008	-0.039
Year FE	✓			✓	✓			✓
Month FE		✓				✓		
Deal Length x Year FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on the change in cash on hand, in the event window around the end of the mortgage deal. Results are shown for the matched mortgage-consumption sample. In this table the change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal, with year and household age fixed effects. Column 2 includes month instead of year fixed effects. Column 3 includes deal length x year and household age fixed effects. Column 4 includes year and household fixed effects. Columns 5-8 repeat this with the change in log cash on hand as the outcome variable. This is defined as the ratio of the change in cash on hand, over the event window, to the mortgage balance 7 months prior to deal expiry. This ratio is multiplied by 100 to give a percentage interpretation. The regressions in columns 1-4 are weighted by sampling weights. The regressions in columns 5-8 are weighted by the product of the sampling weight and the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table A.6: Effect of Mortgage Rates: Omitting Pandemic

	Full Sample				Omit Pandemic (2020-2021)			
	Δ Cash on Hand	Δ Principal	−ΔPayments	Δ Consumption	Δ Cash on Hand	Δ Principal	−ΔPayments	Δ Consumption
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1818.206***	-1100.080***	-718.125***	-937.772***	-1749.650***	-1033.995***	-715.655***	-953.554***
	(104.530)	(119.383)	(38.187)	(277.406)	(115.184)	(131.344)	(40.593)	(290.665)
Observations	6796488	6796488	6796488	68734	4876119	4876119	4876119	39804
Adjusted R^2	0.011	0.007	0.111	0.003	0.013	0.008	0.124	0.003
Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Age FE	✓	✓	✓	✓	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on several outcome variables, on the full sample, and when we omit the Pandemic. Results for columns 1-3 are shown for the full mortgage sample. In this table the change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal. This regression includes year and household age fixed effects. Columns 2-3 show the response of the sub-components of the change in cash on hand. In column 2 the dependent variable is the change in principal, in the event window around the end of the mortgage deal. In column 3 the dependent variable is −10 times the change in monthly payments, in the event window around the end of the mortgage deal. In column 4 the dependent variable is the cumulative change in consumption, in the event window around the end of the mortgage deal. This regression is run on the matched mortgage-consumption sample and the regression is weighted using the sampling weights. Columns 5-8 repeat these 4 regressions omitting deals that expired in 2020 or 2021. Trigger month clustered standard errors are shown in parentheses.

Table A.7: Effect of Alternative Aggregate Mortgage Rates

	Δ Cash on Hand				
	OLS		IV: Agg. Rate Change	IV: Init. Deal Rate Change	IV: Init. LTV Rate Change
	(1)	(2)	(3)	(4)	(5)
ΔMortgage Rate	-1818.206***				
	(104.530)				
ΔIndividual Mortgage Rate		-1825.962***	-1664.483***	-1645.488***	-1494.208***
		(58.729)	(94.195)	(92.655)	(101.511)
Observations	6796488	6796488	6796488	6796488	6740237
Adjusted R^2	0.011	0.020	0.010	0.010	0.009
K-Papp F-Stat			444.975	238.655	280.483
Year FE	✓	✓	✓	✓	✓
Age FE	✓	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of alternative measures of changing aggregate mortgage rates on the change in cash on hand, in the event window around the end of the mortgage deal. Results are shown for the full mortgage sample. In this table the change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. Column 1 shows the regression of the change in cash on hand on the change in the aggregate rates on all new mortgages over the deal. Column 2 regresses the change in cash on hand on the change the household's mortgage rate around deal expiry. Column 3 regresses the change in cash on hand on the change in the household's mortgage rate, where the change in the household's rate is instrumented with the change in aggregate rates on all new mortgages over the deal. Column 4 repeats this, but instead uses as the instrument the change in aggregate rates on all mortgages with the same deal length as the household initially had. In column 5 the instrument is the change in aggregate rates on all mortgages within the bin of the household's initial Loan-to-Value ratio. All regressions include year and household age fixed effects. Trigger month clustered standard errors are shown in parentheses.

Table A.8: Effect of Mortgage Rates on Consumption: No Sampling Weights

	Δ Consumption				Δ Log Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-757.472***	-903.493***	-740.698**	-2105.807***	-2.670***	-1.914**	-2.816**	-7.971***
	(258.461)	(326.234)	(293.906)	(471.436)	(1.012)	(0.914)	(1.166)	(1.812)
Observations	68734	68734	68734	17614	68734	68734	68734	17614
Adjusted R^2	0.002	0.010	0.002	0.027	0.005	0.023	0.005	0.027
Year FE	✓			✓	✓			✓
Month FE		✓				✓		
Deal Length x Year FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on the cumulative change in consumption, in the event window around the end of the mortgage deal. Results are shown for the matched mortgage-consumption sample. The results in this table do not include the sampling weights. Column 1 shows the regression of the change in cash on hand on the change in aggregate mortgage rates over the deal, with year and household age fixed effects. Column 2 includes month instead of year fixed effects. Column 3 includes deal length x year and household age fixed effects. Column 4 includes year and household fixed effects. Columns 5-8 repeat this with the change in log consumption as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table A.9: Effect of Mortgage Rates on Cash on Hand Components

	Δ Principal	−ΔPayments	Δ Principal	−ΔPayments	Δ Principal	−ΔPayments	Δ Principal	−ΔPayments
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1100.080***	-718.125***	-1151.064***	-665.818***	-1065.620***	-707.466***	-1566.959***	-705.816***
	(119.383)	(38.187)	(112.038)	(27.020)	(157.025)	(49.488)	(227.213)	(45.409)
Second-Stage Observations	6796488	6796488	6796488	6796488	6796488	6796488	3681291	3681291
Adjusted R^2	0.007	0.111	0.008	0.115	0.007	0.114	0.071	0.068
Year FE	✓	✓					✓	✓
Month FE			✓	✓				
Deal Length x Year FE					✓	✓		
Household FE							✓	✓
Age FE	✓	✓	✓	✓	✓	✓		

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on the components of the change in cash on hand, in the event window around the end of the mortgage deal. Results are shown for the full mortgage sample and are decomposed into the change in principal and cumulative change in monthly payments. In column 1 the dependent variable is the change in principal, in the event window around the end of the mortgage deal. In column 2 the dependent variable is −10 times the change in monthly payments, in the event window around the end of the mortgage deal. Columns 1-2 include year and household age fixed effects. Columns 3-4 repeat this with month and household age fixed effects. Columns 5-6 include deal length x year and household age fixed effects. Columns 7-8 include year and household fixed effects. Trigger month clustered standard errors are shown in parentheses.

Table A.10: Effect of Mortgage Rates on Mortgage Term

	Δ Mortgage Term			
	(1)	(2)	(3)	(4)
ΔMortgage Rate	2.148***	3.068***	1.635***	1.796***
	(0.209)	(0.167)	(0.284)	(0.289)
Observations	6793674	6793674	6793674	3680041
Adjusted R^2	0.007	0.008	0.008	-0.013
Year FE	✓			✓
Month FE		✓		
Deal Length x Year FE			✓	
Household FE				✓
Age FE	✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of changing aggregate mortgage rates on the remaining months of the mortgage term, in the event window around the end of the mortgage deal. Results are shown for the full mortgage sample. Column 1 shows the regression of the change in mortgage term on the change in aggregate mortgage rates over the deal, with year and household age fixed effects. Column 2 includes month instead of year fixed effects. Column 3 includes deal length x year and household age fixed effects. Column 4 includes year and household fixed effects. Trigger month clustered standard errors are shown in parentheses.

Table A.11: Effect of Financial Accelerator on Cash on Hand

	Δ Cash on Hand				Δ Log Cash on Hand			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-2687.446***	-2626.298***	-2798.657***	-3569.089***	-1.585***	-1.972***	-1.458***	-2.182***
	(212.995)	(191.527)	(250.426)	(346.282)	(0.104)	(0.123)	(0.126)	(0.248)
ΔMortgage Rate x 𝕀(Middle 80% η)	-1804.408***	-1738.218***	-1797.617***	-2245.663***	-1.460***	-1.828***	-1.305***	-1.874***
	(97.124)	(109.225)	(122.965)	(210.002)	(0.072)	(0.102)	(0.088)	(0.207)
ΔMortgage Rate x 𝕀(Top 10% η)	-1011.268***	-949.038***	-983.314***	-1013.645***	-1.099***	-1.464***	-1.027***	-1.177***
	(66.668)	(97.822)	(90.825)	(138.944)	(0.066)	(0.103)	(0.089)	(0.169)
Observations	6796438	6796438	6796433	3681274	6796329	6796329	6796324	3681195
Adjusted R^2	0.017	0.018	0.018	0.079	0.016	0.017	0.018	-0.021
Year x LAD FE	✓	✓		✓	✓	✓		✓
Month FE		✓				✓		
Deal Length x Year x LAD FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on the change in cash on hand, in the event window around the end of the mortgage deal. Each regression interacts the change in aggregate mortgage rates over the deal with dummies for three distinct groups: (i) households in the 10% of regions with the most sensitive house prices to changes in aggregate rates; (ii) households in the middle 80% of regions by sensitivity, and (iii) households in the 10% of regions with the least sensitive house prices to changes in aggregate rates. In columns 1-4 the dependent variable is the change in cash on hand in the event window around the end of the deal. The change in cash on hand is defined as the change in mortgage principal less 10 times the change in monthly payments. Column 1 includes year x Local Authority District (LAD), and household age fixed effects. Column 2 further adds a month fixed effect. Column 3 includes year x deal length x LAD and household age fixed effects. Column 4 includes year x LAD and household fixed effects. Columns 5-8 repeat this with the change in log cash on hand as the outcome variable. This is defined as the ratio of the change in cash on hand, over the event window, to the mortgage balance 7 months prior to deal expiry. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table A.12: Effect of Financial Accelerator on Mortgage Principal: Alternative Instruments

	Δ Principal				Δ Log Principal			
	IV: η	IV: Duration			IV: η	IV: Duration		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
House Price Growth (%)	229.230**	249.289***	491.649***	234.081***	0.099**	0.165***	0.375***	0.344***
	(87.832)	(59.821)	(147.827)	(84.820)	(0.046)	(0.032)	(0.079)	(0.096)
ΔMortgage Rate	264.655	238.453			-0.293*	-0.144		
	(293.544)	(212.219)			(0.168)	(0.178)		
Observations	6796438	4495365	4495365	4270715	6793800	4493917	4493917	4269275
Adjusted R^2	-0.002	-0.002	-0.008	-0.000	-0.001	-0.002	-0.009	-0.001
K-Paap F-Stat	12.151	13.188	14.641	55.301	11.532	11.619	12.469	41.730
Year x LAD FE	✓	✓	✓		✓	✓	✓	
Year x 𝕀(F.hold) FE		✓	✓			✓	✓	
ΔMort Rate x η Control		✓				✓		
Lad x ΔMort Rate FE			✓				✓	
Start M. x Trig M x LAD FE				✓				✓
Month x 𝕀(F.hold) FE				✓				✓
Age FE	✓	✓	✓	✓	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on mortgage principal for alternative instruments. In columns 1-4 the dependent variable is the change in principal in the event window around the end of the deal. This is regressed on the change in the aggregate mortgage rate over the deal and 100 times the log change in the real Local Authority District (LAD) house price index over the deal. The change in the house price index is instrumented by ΔMortgage Rate x η, where η measures the average sensitivity of local house prices to interest rates. This regression includes year x LAD and household age fixed effects. Columns 2-4 repeat this regression, but instead use as an instrument ΔMort x 𝕀(Freehold), where 𝕀(Freehold) is a dummy for whether the property is owned freehold or not. In these columns, we use separate real house price indices for freehold and leasehold properties within each LAD. Column 2 includes the following fixed effects: (i) year x LAD; (ii) household age; and (iii) year x 𝕀(Freehold). It also includes ΔMortgage Rate x η as a control. Column 3 replaces the ΔMortgage Rate x η control with LAD x ΔMortgage Rate fixed effects. Column 4 includes the following fixed effects: (i) deal start month x deal expiry month x LAD; (ii) month x 𝕀(Freehold); and (iii) household age. Columns 5-8 repeat this with the change in log principal as the outcome variable. This variable is multiplied by 100 to give a percentage interpretation. The regressions in columns 6-10 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

Table A.13: Effect of Financial Accelerator Instruments on House Prices

	House Price Growth (%)			
	LAD House Price Index	LAD x FH/LH House Price Index		
	(1)	(2)	(3)	(4)
ΔMortgage Rate	-2.815***	-1.170*		
	(0.444)	(0.593)		
ΔMortgage Rate x η	0.519***	0.484***		
	(0.149)	(0.142)		
ΔMortgage Rate x $\mathbb{I}$(F.hold)		-1.734***	-1.476***	-3.430***
		(0.477)	(0.386)	(0.461)
Observations	6796438	4495365	4495365	4270715
Adjusted R^2	0.717	0.766	0.768	0.971
F-Stat	12.151	13.188	14.641	55.301
Year x LAD FE	✓	✓	✓	
Year x $\mathbb{I}$(F.hold) FE		✓	✓	
Lad x ΔMort Rate FE			✓	
Start M. x Trig M x LAD FE				✓
Month x $\mathbb{I}$(F.hold) FE				✓
Age FE	✓	✓	✓	✓

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of alternative financial accelerator instruments on house prices. In column 1 the dependent variable is 100 times the log change in the real Local Authority District (LAD) house price index over the deal. This is regressed on ΔMortgage Rate and ΔMortgage Rate x η, where η measures the average sensitivity of local house prices to interest rates. The regression includes year x LAD and household age fixed effects. In columns 2-4 the dependent variable includes separate real house price indices for freehold and leasehold properties within each LAD. Column 2 adds to the regression ΔMort x $\mathbb{I}$(Freehold), where $\mathbb{I}$(Freehold) is a dummy for whether the property is owned freehold or not. This regression includes the following fixed effects: (i) year x LAD; (ii) household age; and (iii) year x $\mathbb{I}$(Freehold). Column 3 further includes LAD x ΔMortgage Rate fixed effects. Column 4 includes the following fixed effects: (i) deal start month x deal expiry month x LAD; (ii) month x $\mathbb{I}$(Freehold); and (iii) household age. Trigger month clustered standard errors are shown in parentheses.

Table A.14: Effect of Financial Accelerator on Individual Mortgage Interest Rates

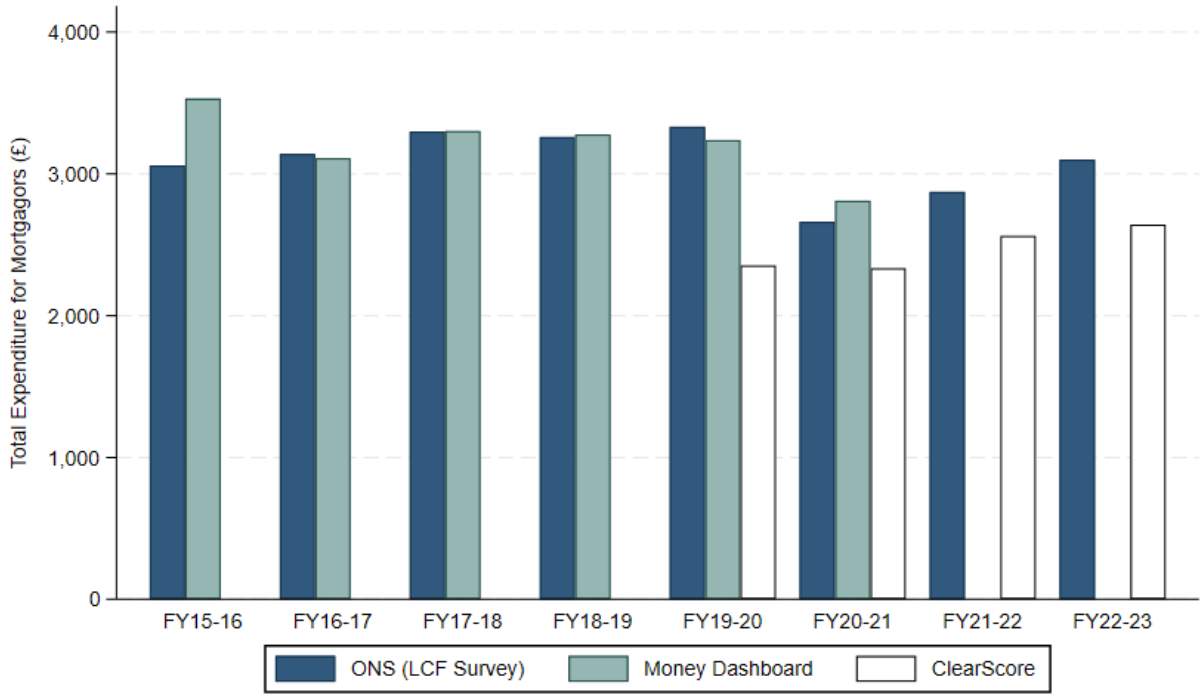
	Δ Individual Mortgage Rate				Δ Log Individual Mortgage Rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMort. Rate x 𝕀(Bottom 10% η)	1.246***	1.011***	1.249***	1.220***	39.193***	35.098***	39.511***	41.487***
	(0.068)	(0.098)	(0.065)	(0.070)	(2.901)	(4.589)	(2.771)	(2.844)
ΔMort. Rate x 𝕀(Middle 80% η)	1.283***	1.045***	1.293***	1.273***	37.905***	33.893***	38.690***	40.770***
	(0.072)	(0.096)	(0.066)	(0.073)	(2.941)	(4.590)	(2.810)	(2.706)
ΔMort. Rate x 𝕀(Top 10% η)	1.329***	1.093***	1.349***	1.394***	35.414***	31.575***	36.901***	39.954***
	(0.083)	(0.097)	(0.077)	(0.088)	(3.206)	(4.618)	(3.060)	(3.157)
ΔMort. Rate x 𝕀(Bottom 10% η) x LTV	-0.309***	-0.293***	-0.309***	-0.246***				
	(0.050)	(0.042)	(0.051)	(0.050)				
ΔMort. Rate x 𝕀(Middle 80% η) x LTV	-0.370***	-0.349***	-0.366***	-0.312***				
	(0.047)	(0.039)	(0.047)	(0.051)				
ΔMort. Rate x 𝕀(Top 10% η) x LTV	-0.425***	-0.402***	-0.420***	-0.400***				
	(0.050)	(0.041)	(0.048)	(0.056)				
Observations	6740211	6740211	6740206	3647778	6795567	6795567	6795562	3680597
Adjusted R^2	0.600	0.606	0.604	0.658	0.563	0.573	0.569	0.590
Year x LAD FE	✓	✓		✓	✓	✓		✓
Month FE		✓				✓		
Deal Length x Year x LAD FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	
Start Year	✓	✓	✓	✓	✓	✓	✓	✓
Start Year x LTV FE	✓	✓	✓	✓				
η x LTV FE	✓	✓	✓	✓				

$^{\star}p < 0.10, ^{\star\star}p < 0.05, ^{\star\star\star}p < 0.01$.

Notes: The table shows the impact of the financial accelerator on the change in individual mortgage interest rates. Each regression interacts the change in aggregate mortgage rates over the deal with dummies for three distinct groups: (i) households in the 10% of regions with the most sensitive house prices to changes in aggregate rates; (ii) households in the middle 80% of regions by sensitivity, and (iii) households in the 10% of regions with the least sensitive house prices to changes in aggregate rates. In columns 1-4 the dependent variable is the change in individual mortgage rates in the event window around the end of the deal. Columns 1-4 also include the triple interaction of the change in aggregate rates, the dummies for regional house price sensitivity, and the LTV ratio on the mortgage at the start of the deal. Columns 1-4 all include the following fixed effects: (i) start year; (ii) start year x LTV ratio; and (iii) regional house price sensitivity dummies x LTV ratio. In addition to this, column 1 includes year x Local Authority District (LAD) and household age fixed effects; column 2 includes year x LAD, month, and household age fixed effects; column 3 includes year x deal length x LAD and household age fixed effects; and column 4 includes year x LAD and household fixed effects. In columns 5-8 the dependent variable is the 100 times the log change in interest rates, and the triple interaction with the LTV ratio, and related fixed effects, are not included. The regressions in columns 5-8 are weighted by the mortgage balance 7 months prior to deal expiry. Trigger month clustered standard errors are shown in parentheses.

B Appendix Figures

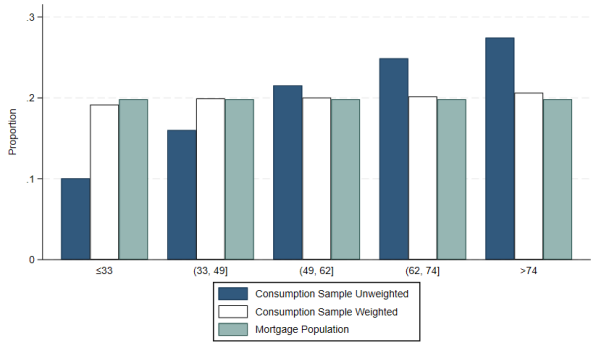
Figure B.1: Monthly Average Spending in Apps vs. Official Data



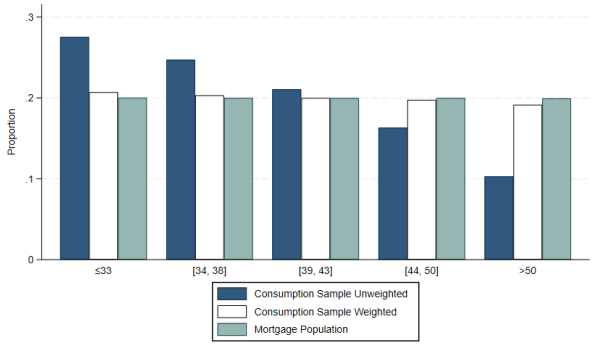
Notes: the figure shows the monthly average household total expenditure for mortgagors for the Office of National Statistics (Living Costs and Food Survey), Money Dashboard and ClearScore from 2015 to 2023. We use our total consumption measure to calculate the average values for the Money Dashboard and ClearScore apps. The units are pounds based in 2015.

Figure B.2: Distribution of Consumption Sample Versus Full Mortgage Sample

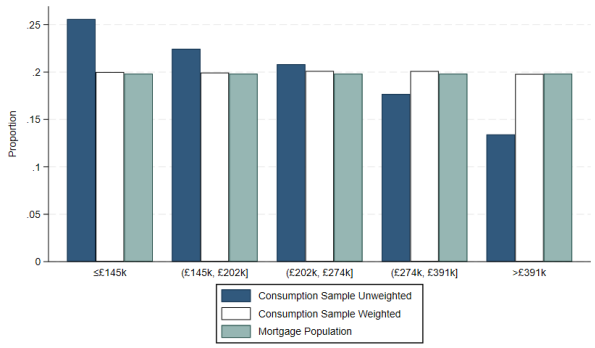
(a) Loan to Value



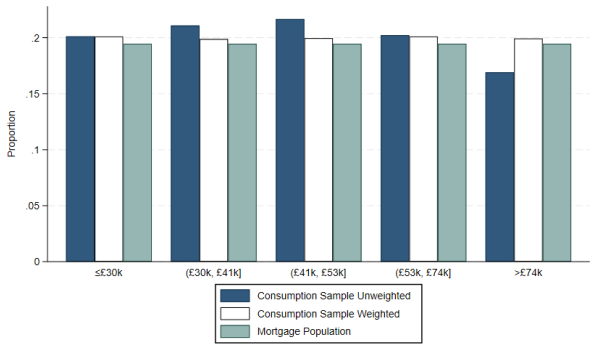
(b) Borrower Age



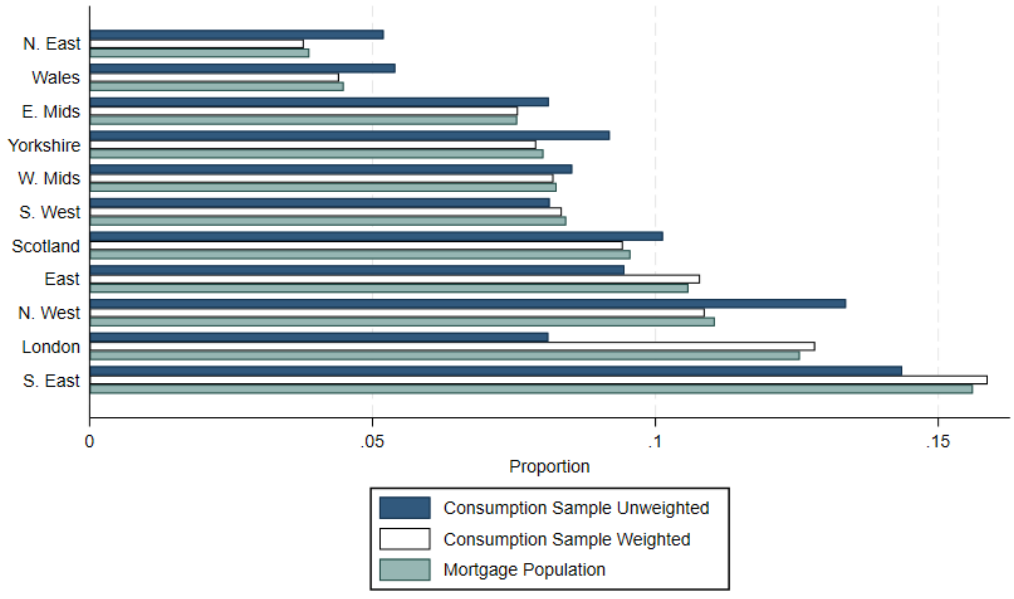
(c) Home Value



(d) Household Income

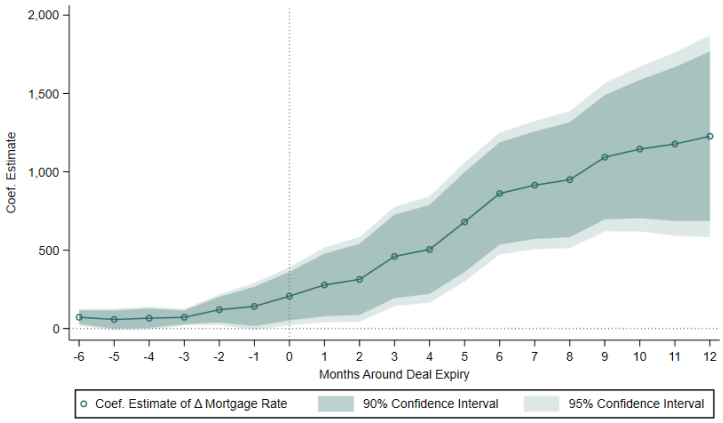


(e) Broad Region



Notes: the figure compares distributions of loan-to-value (a), borrower age (b), home value in 2015 pounds (c), household income in 2015 pounds (d) and region (e) in the mortgage population, the subset with observed app-based consumption, and the same subset after applying weights.

Figure B.3: Impulse Response of Salary and Wages



Notes: the figure plots impulse responses of monthly salary and wages from the baseline monthly regression. For horizons $h \in [-6, 12]$ relative to deal expiry ($t = 0$), the series show the cumulative change in salary and/or wages (pounds, 2015 base) per 1 percentage-point change in the interest rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. Shaded bands denote 90% and 95% confidence intervals. Sample: mortgagor households with income inferred from linked spending apps.

C Transaction-level Expenditure Data

C.1 Background About The Apps

Money Dashboard was a free personal budgeting application in the UK. The app aimed to help users understand how they spend money and to budget. It did this by collating information from the transactions of their current accounts, credit cards, investment and savings accounts, and providing users with an overview of their spending behavior. The app launched in 2010. In 2021, it was purchased by ClearScore and wound down in 2023. ClearScore is a similar free application in the UK that provides free advice about users credit scores. The app launched in 2018 and is still active today.

We use transaction-level data from both apps. For each transaction, there is information about:

- transaction category or "tag", e.g. mortgage payment, wages, online shopping, restaurant, personal loan repayments, credit card finance, cash withdrawal
- whether the transaction is a credit or debit
- transaction value
- account used to make the transaction, e.g. current account or credit card, including details of the financial intermediary associated with the account
- where relevant, details of the merchant attached to the transaction, e.g. Tesco, B&Q, British Gas, Amazon

C.2 Transaction-level Data Treatment

Expenditure Categories

There are around 300 distinct "tags" to categorize transactions in the Money Dashboard and ClearScore apps. We use this tag information, and augment the Bureau of Economic Analysis (BEA) approach[33], to create four expenditure categories:

1. *Non-durable consumption excluding home-related expenditures* includes spending on take-away food, groceries, alcohol, tobacco, clothing, footwear, and non-durable personal goods such as medication and toiletries, among others.

[33]See Table 5 here: https://www.bea.gov/resources/methodologies/nipa-handbook/pdf/chapter-05.pdf

2. *Services excluding home-related expenditures* includes spending on restaurants, hotels, travel fares, personal services, leisure services, education services, medical services, motoring expenditures, among others.

3. *Durable goods excluding home-related expenditures* includes spending on motor vehicles, other vehicles such as motorbikes and caravans, durable personal goods such as computers, durable leisure goods such as gym equipment, bikes and musical instruments, among others.

4. *Home-related expenditures* includes spending on furniture, soft furnishings, home lighting, garden-related expense, bedroom textiles, kitchenware, electric and home appliances, television set purchases, tradesmen fees, among others.

We create a "consumption" measure as the sum of the above four categories, as well as any cash expenditures. All expenditure categories and consumption exclude mortgage and rental-related costs.

Machine Learning for Unclassified Transactions

Around 10 percent of transactions have a tag of "No tag", which means they are not classified. These "No tag" transactions contain information such as the transaction value, a transaction description, whether the transaction is credit or debit, among other information, but nonetheless were not automatically classified by the Money Dashboard or ClearScore. We therefore use machine learning to predict the classification of the "No tag" transactions using the available information.

Our final prediction is a stacked model that combines: (i) a text prediction for the transaction description using a text classification model), (ii) a numeric predictions for the transaction amount using naive bayesian model, and (ii) categorical predictions for whether the transaction is a debit/credit.

The "No tag" items we classify with the stacked model are then included in the expenditure categories and other variables, based on their new assigned tag.

Other Variables

We also extract a number of other useful variables from the transactions including: Mortgage payments, salary and wages, other household income such as bonuses, benefits, investments, pension, tax refunds, insurance payments, savings, credit card finance and repayments, loan finance and repayments.

Sample Restrictions

We exclude: duplicate transactions, business accounts, users that are less than 16 years or more than 100 years. We also identify users that have periods for which they have no data or substantially missing data. This can occur if a user's online banking details change, leaving the app unable to collect the information. Observations and/or users are identified and excluded as those with mean total consumption in the top and bottom percentile, as well as the standard deviation of mean total consumption in the top percentile, and consumption observations than drop below 25 percent of the relevant user's mean consumption. We interpolate missing data when there is no more than two missing values.

D Merge Between Mortgage and Consumption Data

D.1 Treatment of PSD Mortgage Data

We use the combined PSD mortgage data for the merge to the MoneyDashBoard and ClearScore datasets. The PSD 007 dataset records, in June and December of each year from 2015 onwards, the stock of outstanding mortgages at that date. PSD 007 records the contractual monthly payment required from each mortgage (in a whole number of pounds), along with the lender, date of birth of the borrower, and full postcode of the property that secures the mortgage. There are around 1.8 million UK postcodes, and a typical one comprises 15 addresses.

The user location data received from MoneyDashBoard and ClearScore is anonymised from their full address, with their location provided at the sector level. The postcode sector omits the last two digits of the postcode and covers the postcode district plus the first character of the inward code. There are 12,463 postcode sectors in the UK.[34] Data on the date of birth of users is also partially obscured: in ClearScore we see the month and year of birth; in MoneyDashBoard we only see the year of birth.

We record the full dates of birth for one or two individuals for each mortgage in the PSD. For each of these dates of birth we form a dataset for matching to ClearScore based on: (i) the month and year of birth; (ii) the postcode sector of the property; (iii) the contractual monthly mortgage payment in pounds; (iv) the month of the mortgage snapshot; and (v) the mortgage lender (with some consolidation to capture which parts of banking groups report to the PSD and which parts of the group the mortgage payments are sent to). These variables identify a unique mortgage in the PSD data for at least 99.97% of observations for each of the two dates of birth. We remove mortgages which are duplicates in these matching variables. We follow the same procedure for matching to MoneyDashBoard, except, for (i), we only match on the year of birth. In this case the matching variables identify unique mortgages for at least 99.63% of observations for each of the two dates of birth.

D.2 Treatment of ClearScore Data

Transaction Floor To identify potential mortgage transactions in the ClearScore data we restrict the sample to debits from current accounts (the UK equivalent to checking accounts). Moreover, to keep the file size manageable, we restrict the sample to debits of at least 100 pounds. This limit includes the vast majority of monthly mortgage payments: over the matching period from 2018m12 to 2023m12, 98.7% of outstanding mortgages in our

[34] See https://www.ons.gov.uk/methodology/geography/ukgeographies/postalgeography.

sample have a contractual monthly payment of at least 100 pounds. We also only include transactions in June and December each year; the months for which we observe the stock of outstanding mortgages.[35] This initial sample has around 86 million transactions.

Lender Transaction Codes Current account providers give each transaction a code to indicate the transaction type. These codes differ by current account provider so we clean and harmonise the transaction codes used into 6 broad categories:

- *Direct Debit*: these are recurring transactions where the recipient can alter the amount taken, as required. These transactions are typically used for payment of recurring bills, and can be set to occur on a specified date, for example, the same day each month.
- *Standing Order*: these are recurring transactions where the sender chooses the amount and frequency. These can be used for recurring payments of fixed amounts, such as rental payments and mortgages. They can also be used to transfer a set amount of money to other individuals at a given frequency.
- *Funds Transfer*: this category covers the one-off transfer of funds between accounts.
- *Card Payments*: this category covers the use of the debit card linked to the current account. It covers both the in-person use of the card in shops, and the use of the card to make online purchases.
- *Cash*: this covers withdrawals of cash from the current account.
- *Other*: this covers a small number of observations not in the above categories, such as current account charges and interest charged on an overdrawn account.

The PSD mortgage data records the contractual monthly mortgage payment that is due. Thus, in ClearScore, we isolate transactions which are for recurring payments: that is, those that are made by direct debit or standing order (we also retain the small number of transactions missing a transaction code).[36] This leaves a dataset with just under 23 million observations.

Transactions to Banks/Building Societies We next use information in the *Transaction Description* to identify transactions that are made to banks and building societies (a mutual organisation that typically specialises in savings accounts and mortgage lending).

[35] This is without loss as UK mortgages are repaid monthly.

[36] Other categories, such as funds transfers, can include one-off mortgage payments, but these are more likely to be used for *ad hoc* over-payments.

The *Transaction Description* is a text string of up to around 160 characters which contains a description of what the payment is for, with personal information, such as specific account numbers redacted (e.g. replacing "direct debit to bank x, account 12345678" with "direct debit to bank x, account xxxx5678"). We assign transactions to different lenders by searching the *Transaction Description* for regular expressions, accounting for the different ways banks/building society names can be written. For a few lenders, we also assign certain forms of transaction description to that lender, even though the lender name is not explicitly present in the string. In these cases there are many transactions of this form, and in *ex-post* merge tests, almost all of these transactions match to the given lender, with no significant mass for any other lenders.

Transaction Categories Using the *Transaction Description,* ClearScore generates a variable *Purpose_Tag* which classifies the transaction into different categories. Using this variable, and our extracted lender names, we keep observations for 4 groups of transactions:

1. Transactions Tagged as Mortgages
    - This group predominantly consists of those tagged as "Mortgage payment", but also includes a smaller number of observations tagged as "Mortgage or Rent".
    - The observations in this category typically feature the word "mortgage" or the abbreviation "mtg" in the transaction description.
    - We keep all observations in these two categories where the *Transaction Description* can be mapped to a bank/building society, using the process described above. This latter restriction results in the loss of around 5% of observations in this category, but is used to ensure merge quality, with a matching lender in the PSD and ClearScore datasets.
2. Residual Tag Categories
    - A number of transactions are allocated to what are, effectively, residual transaction categories. These are "General Retail" and "No Tag".
    - From this broad set we retain the transactions which can be mapped to a bank/building society.
    - This includes transactions with a *Transaction Description* of the form "Direct Debit to X", where X is the name of a bank/building society. In these cases it is not clear what the direct debit is for, whether a mortgage repayment, or for example, a personal loan repayment.

3. Financial Transaction Categories

    - This group consists of transactions tagged as "Transfers" or "Personal Loans".
    - The "Transfers" category in this sample primarily consists of payments to the accounts of other individuals or recurring payments into savings accounts. However, it also contains a number of transactions to lenders that are not clearly of this form. In the "Transfers" category we retain payments that can be mapped to a bank/building society and also remove those which have regular expressions in the *Transaction Description* that can be clearly identified as transfers (e.g. containing the expressions "isa" or "saver").
    - The "Personal Loan" category primarily consists of unsecured personal loans, but also includes a smaller number of transactions to lenders which are categorised *en masse* as personal loans, although the lenders offer mortgage products.
    - As with transfers, we keep observations in this category which can be mapped to a bank/building society, and exclude those with regular expressions which clearly identify them as person loans (e.g. of the form "Bank X loans").

4. Mis-tagged Transactions to Banks/Building Societies in Other Categories

    - We also keep a small number of observations to specific lenders which are mis-tagged to other consumption categories. This can occur, for example, where a company offers both banking services and has a consumer retail presence.

After these filters are applied, we have a dataset with around 2.5 million observations.

Merging Variables As discussed above, the PSD mortgage dataset records the contractual monthly mortgage payment in a whole number of pounds. By contrast, ClearScore records the full transaction value, in pounds and pence. There is no clear guidance in the UK mortgage regulations on whether the reported contractual monthly mortgage payment should be rounded to the nearest pound, rounded down, or rounded up. Thus, to allow for both possibilities, and for different practices across lenders, we create two datasets: one in which transaction amounts are rounded down to the nearest pound (including those for a whole number of pounds), and one where they are rounded up.

In both datasets the matching variables are: (i) the month and year of birth; (ii) the postcode sector of the user's address; (iii) the rounded transaction amount in pounds; (iv) the month of the transaction; and (v) the lender identified in the transaction (with the same group consolidation applied to the lenders in both datasets). These variables identify unique

user observations in both ClearScore datasets in over 99% of cases. We remove observations which are duplicates in these matching variables.

D.3 Treatment of MoneyDashBoard Data

Transaction Floor As with ClearScore, we limit the sample to debits of at least 100 pounds from current accounts. Over the matching period from 2015m6 to 2021m6, 98.3% of outstanding mortgages in our sample have a contractual monthly payment of at least 100 pounds. We also only include transactions in June and December each year. This initial sample has around 10 million transactions.

Lender Transaction Codes Lender transaction codes are not reported as a separate variable in MoneyDashBoard. However, they can often be extracted from the *Transaction Description.* We group transactions into the same 6 categories as with ClearScore. In all, we are able to extract lender transaction codes for 59% of the data. We retain transactions which are for direct debit, standing order, or where the lender transaction code could not be extracted. This leaves a dataset with just over 6 million observations.

Transactions to Banks/Building Societies We follow a very similar procedure to that used for the ClearScore data and use information in the *Transaction Description* to identify transactions that are made to banks and building societies.

Transaction Categories Using the *Transaction Description,* MoneyDashboard generates a variable *Purpose_Tag* which classifies the transaction into different categories. Using this variable, and our extracted lender names, we keep observations for the same 4 groups of transactions as with the ClearScore data, using a very similar method. Compared to the ClearScore data, there is no "General Retail" tag in MoneyDashboard. Moreover, in MoneyDashboard, we identify mis-tagged transactions to banks/building societies in a few additional tagged categories. After these filters are applied, we have a dataset with around 600 thousand observations.

Merging Variables As with ClearScore, we create two matching datasets: one in which transaction amounts are rounded down to the nearest pound (including those for a whole number of pounds), and one where they are rounded up. For MoneyDashBoard the matching variables are: (i) the year of birth; (ii) the postcode sector of the user's address; (iii) the rounded transaction amount in pounds; (iv) the month of the transaction; and (v) the lender identified in the transaction (with the same group consolidation applied to the lenders in

both datasets). These variables identify unique user observations in both MoneyDashBoard datasets in around 97.5% of cases. We remove observations which are duplicates in these matching variables.

D.4 PSD-ClearScore/MoneyDashBoard Merge

Merge As discussed above, we (i) create a separate PSD matching dataset for up to two borrowers attached to each mortgage; and (ii) create separate ClearScore/MoneyDashBoard matching datasets where transaction amounts are rounded down or up to the nearest pound. This results in 4 separate merges for each of ClearScore/MoneyDashBoard which are then appended together.

Treating Multiple ClearScore/MoneyDashBoard Matches to PSD We first treat cases where multiple distinct users of the given app have been matched to the same mortgage in the PSD. This affects a little under 5% of matches to ClearScore and 1% of matches to MoneyDashBoard. The majority of these cases appear to be due to joint bank accounts. For example, where two mortgagors pay the mortgage from a joint bank account and have both shared this account information with the app. The joint accounts are identified by groups of users sharing the same: (i) postcode sector; (ii) current account provider; (iii) transaction date; and (iv) transaction amount. Within these groups we keep the app user matched to the PSD mortgage the greatest number of times, dropping the match for the other user.

After this treatment for joint accounts, we treat the remaining cases with multiple matches, which are assumed to be due to some false positives. Within this group we keep the app user matched to the PSD mortgage data the greatest number of times.

Additional Merge Treatment Finally, we perform two more refinements to the merged sample. First, we treat cases of app users matching to multiple different PSD households over time (identified using their date of birth and full postcode). We judge that this is more likely due to false positives than people moving house and retain the match with the PSD household that appears most frequently. This results in 1.3% of observations being dropped for the ClearScore merge and around 2% being dropped for the MoneyDashBoard merge. After this treatment, we are left with a unique mapping between app users and PSD mortgages, for each dataset.

Finally, we treat observations where bank codes cannot be identified for the transactions. This affects less than 1% of observations for the ClearScore merge and 44% of observations for the MoneyDashBoard merge. Where the bank code is missing, to ensure that we are picking

up repeat transactions, we restrict to merges where the app user and the PSD household are matched at least twice.

Merge Numbers and Quality: ClearScore Overall, after these restrictions, we match 824,999 ClearScore transactions to PSD mortgages. Of the transactions with an identified bank code, 99.7% are direct debits, with the rest standing order. Of the merged transactions, 98% are for users who are matched to the PSD data at least twice, with these matches at least 6 months apart, by construction. Moreover, 99.5% of the matched transactions are for a non-whole number amount, consistent with monthly mortgage payments which are the outcome of the calculations from an amortization schedule, and are thus typically not whole numbers. Overall, 64% of matches are tagged by ClearScore as "Mortgage Payment"; 25% are tagged as "General Retail"; and 9% are tagged as "No Tag". The remaining 2% of matches are distributed among the other categories.

Overall, we match 161,532 ClearScore users to a mortgage in the PSD. Of this, we match 148,642 ClearScore users (92%) to mortgages in the PSD dataset where the start date of the mortgage is known.

Merge Numbers and Quality: MoneyDashBoard Overall, after these restrictions, we match 180,155 MoneyDashBoard transactions to PSD mortgages. Of the transactions with an identified bank code, 99.9% are direct debits, with the rest standing order. Of the merged transactions, 97.5% are for users who are matched to the PSD data at least twice, with these matches at least 6 months apart, by construction. Moreover, 99.2% of the matched transactions are for a non-whole number amount. Overall, 42% of matches are tagged by MoneyDashBoard as "Mortgage Payment"; 26% are tagged as "No Tag"; 20% are tagged as "Personal Loan"; and 8% are tagged as "Mortgage or Rent". The remaining 3.4% of matches are distributed among the other categories.

Overall, we match 36,616 MoneyDashBoard users to a mortgage in the PSD. Of this, we match 30,949 MoneyDashBoard users (85%) to mortgages in the PSD dataset where the start date of the mortgage is known.

In total, across both datasets, we match 194,572 mortgagors to consumption app data (reflecting a small number of mortgagors matched to both ClearScore and MoneyDashBoard).

E Consumption Sample Weighting

Our consumption sample is non-random given that households self-select to join the apps. Our mortgage sample covers the population of refinancing activity and we have overlapping outcome variables present in both the consumption and mortgage samples. We therefore estimate pseudo-inclusion weights to correct for potential selection bias and coverage errors present in the consumption sample. Our pseudo-inclusion weights are estimated using the quasi-randmization weighting approach outlined in Valliant and Dever (2018).

We estimate the pseudo-inclusion probability of being in the consumption sample via a logistic regression model. The outcome variables is a dummy variable equal to 1 if the observations are included in the consumption sample and 0 for the remaining observations in the mortgage population. The control variables comprise the "pre" variables for age, income, LTV, mortgage term and home value (each divided into quintiles), as well as broad region fixed effects.

The pseudo-inclusion weights are the inverse of the estimated probability of being in the consumption sample. These weights correct for both selection bias and coverage error. The selection bias is well-understood but coverage error also occurs because not everyone in the mortgage population will have access to a smart phone and/or internet in order to select into the consumption sample in the first place. These under-covered groups tend to be older and are less likely to live in urban areas.

Table 1 compares summary statistics for the mortgage population, the unweighted consumption sample and the consumption sample weighted using our pseudo-inclusion weights. The statistics adjusted with the pseudo-inclusion weights show that our approach performs very well in adjusting the mean values in the consumption sample to match those in the mortgage population. Figure B.2 further demonstrates that our weights appropriately adjust the distribution (for loan-to-value, borrower age, home value, and income) in the consumption sample to match the mortgage population.