How Do Interest Rates Affect Consumption? Household Debt and the Role of Asset Prices*

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

February 23, 2026

Abstract

This paper estimates how rate cuts increase consumption, via debt and asset prices. Using administrative UK data on mortgages and consumption, we exploit the expiry of fixed-rate mortgages to construct six million household-level natural experiments. A 1 pp reduction in mortgage rates raises consumption by 3% in the following 6 months. Using plausibly exogenous variation in how house prices respond to rate cuts, we show that consumption increases mostly because households borrow against higher house prices; lower debt service after rate cuts matters less. These results suggest that in large part, monetary policy affects consumption through asset prices and borrowing.

*We thank Amil Mumssen, Verónica Bäcker-Peral and Alexander Prior for excellent research assistance; and Matias Bayas-Erazo, Gabriel Chodorow-Reich, Matthias Doepke, Thomas Drechsel, Martin Flodén, John Gathergood, Martin Holm, Kilian Huber, Matteo Iacoviello, Ethan Ilzetzki, Amir Kermani, Ben Moll, Huw Pill, Christina Patterson, Ricardo Reis and Maria Olsson for comments. UK residential mortgage data are taken from the Product Sales Data (PSD) provided by the UK Financial Conduct Authority (FCA). The PSD includes regulated mortgage contracts only, and therefore excludes 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 Monetary Policy Committee, the Financial Policy Committee, or the Prudential Regulation Committee. This paper was prepared under the European System of Central Banks ChaMP research network (Challenges for the Transmission of Monetary Policy in a Changing World).

1 Introduction

When interest rates fall, do people spend more? And if so, why? The answer is fundamental to macroeconomics—determining, for instance, the power of monetary policy. This paper estimates how rate cuts increase consumption through their effect on debt and asset prices. Roughly 40% of households in the United States and United Kingdom have a mortgage. Rate cuts encourage households to increase spending either by lowering debt service (the "mortgage cash flow effect"), or by increasing borrowing capacity as house prices go up. We collectively refer to this pass-through from interest rates to spending, via debt service and borrowing against asset prices, as the "household debt channel".

Our contribution is twofold. First, we estimate the household debt channel. We exploit the staggered expiration of fixed-rate mortgage deals in the United Kingdom, in order to construct six million household-level natural experiments. We find that the household debt channel is large: after a 1 pp decline in mortgage rates, mortgagors increase consumption by 3% over the next 6 months. Second, we study plausibly exogenous variation in how house prices respond to rate cuts. We show that consumption rises mostly because households borrow more against the house price increase after rate cuts; lower debt service is less important. Overall, monetary policy affects consumption in large part through asset prices, as in a large and evolving set of theoretical work (e.g. Iacoviello, 2005; Greenwald, 2018; Caballero and Simsek, 2024).[1]

Measuring the household debt channel is not easy because rate cuts affect consumption in many ways. For instance according to Heterogeneous Agent New Keynesian (HANK) models, rate cuts increase income, leading households to consume more (Kaplan, Moll and Violante, 2018; Auclert, 2019). To isolate the household debt channel, this paper studies mortgage deals in the UK, during 2015-2024. 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. When the deal expires, there are strong incentives to start a new one. There is no cost of borrowing as a deal starts, but large costs otherwise. Altogether, households face rate changes and borrowing opportunities at predetermined and staggered times—creating natural experiments that identify the household debt channel.[2]

We pair these natural experiments with individual, monthly level mortgage and consumption data. We merge administrative data on mortgages, for the universe of UK households, with consumption information from personal finance apps. App users are broadly represen-

[1] We often refer to the effect of "rate cuts", but most of our analysis is symmetric, and does not differentiate between cuts versus increases.

[2] Cloyne, Huber, Ilzetzki and Kleven (2019) were the first to exploit this natural experiment, by investigating the effect of house prices on borrowing. We instead study how interest rates affect consumption.

tative of the UK, and expenditure in the app is similar to official sources. There are 6.8 million expiring deals in the mortgage dataset, 3.7 million of which are repeat observations of the same household, with consumption information for 70 000 deals. Having so many expiring deals is useful. We want not only to estimate the household debt channel, but to find out what makes it large. Each expiring deal provides a single estimate. But since there are millions of deals in total, we can detect rich heterogeneity—which will reveal the role of asset prices.

We show how to identify the household debt channel using a simple model with mortgages. Households earn income and consume; and borrow via a mortgage, with an interest rate that depends on the loan-to-value ratio. There is a household debt channel: when rates fall, consumption rises because of lower debt service costs, and more borrowing through mortgages. That is, the first component of the household debt channel is the "mortgage cashflow effect" emphasized by Di Maggio et al. (2017) and Holm, Paul and Tischbirek (2021), which corresponds to lower debt service. However a second component, changes in borrowing, also affects the household debt channel. Other effects of rate cuts that do not operate via debt are held fixed, including the movements in income stressed by HANK models. Through the household debt channel, monetary policy affects consumption with "long and variable lags" (Friedman, 1961). Mortgages reset infrequently, meaning rate cuts pass through to aggregates with delay. In the model, monetary policy affects consumption in large part via asset prices. If rate cuts raise house prices, then households consume more by borrowing against their home (Iacoviello, 2005; Greenwald, 2018), meaning the household debt channel is large.

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, as well as a time fixed effect. Since households have different deal lengths, they experience different rate changes as they refinance, which affects their consumption. 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. This assumption is tenable because with UK mortgages, deal expiration is predetermined. In other countries such as the United States, refinancing is a choice of the household, which might correlate with income shocks.

Our first result is that the household debt channel is large. In our baseline specification a 1 pp fall in mortgage rates causes consumption to rise by around 950 pounds in the 6

months after the deal expires, an increase of around 3.0%.[3] The response of cash-on-hand is 1,800 pounds (1.4%) in the baseline specification. The marginal propensity to consume (MPC) out of the higher cash-on-hand due to the rate cut is around 0.5 in the 6 months after deal expiry. Our estimates of the household debt channel are large: if all deals expired at once, consumption would increase by around 0.7% of GDP after the 1 pp rate cut.

Our baseline specification combines three sources of variation; each individually leads to similar results. The first source is cross-sectional variation in the rate changes experienced by households in a given month. 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 isolate this variation with month fixed effects. Second, we add deal length by year fixed effects, which compares two households with the same deal length, expiring at different times during the year. Third, half of the 6 million natural experiments are repeat observations of the same household. We can add household fixed effects, which compares the response of the same household as they receive different rate changes. The results are similar with each kind of variation—limiting the chance that omitted variables explain our results.

Besides their stability, there are two more pieces of evidence that support our results. First, there are no pre-trends, meaning that any confounder would have had to occur in precisely the same month that the deal expires. Second, income falls as consumption rises after rate cuts. This pattern suggests wealth effects on labor supply and does not violate our identification assumption, whereas confounders would likely increase income.[4]

The household debt channel has significant aggregate implications. Its size is similar to time series estimates of how monetary policy affects consumption in the UK (e.g. Cloyne, Ferreira and Surico, 2020). Moreover due to the household debt channel, rate cuts affect consumption with the "long and variable lags" that appear in the time series. For this exercise, we combine our estimates of how rate cuts affect households as their deal expires, with estimates of how often deals expire. The lags of monetary policy are long, with the cumulative impact of rate cuts taking around 50 months to play out. The lags are also variable. The length of the typical mortgage deal rose over our sample. As a result, the response of consumption became more back-loaded, with the response of consumption on impact halving between 2016 and 2023.[5]

[3] All nominal pound values discussed in the paper are based in 2015.

[4] We carry out other robustness tests. Selection into the consumption data is unlikely to explain our results, since the response of the cash-on-hand to rates is similar for households with and without the finance app. The results also hold while instrumenting for mortgage rate changes with high frequency monetary shocks (Braun, Miranda-Agrippino and Saha, 2025).

[5] This finding relates to work showing that when mortgages have shorter duration, due to being floating-rate instead of fixed-rate, the response of consumption is more front-loaded (e.g. Garriga et al., 2017; World

In the rest of the paper, we show that the household debt channel is large because of how interest rates affect asset prices. Our second main result studies which of the two components of cash-on-hand, either borrowing or debt service, responds by more to rate cuts. We find that borrowing responds more than debt service at all horizons, accounting for about two thirds of the total response of cash-on-hand after 6 months. Therefore the mechanism generating a large household debt channel primarily works through borrowing—making the role of asset prices potentially important.

We develop a strategy to pin down how asset prices affect the household debt channel. Identifying the role of asset prices is not easy: one must separate house price movements from the many other forces that affect borrowing after rate cuts. At this point, the advantage of our setting—that there are millions of expiring deals—becomes especially important. Given this granularity, we can estimate the household debt channel separately for disaggregated groups of households, who differ in how interest rates affect their house price.

We estimate the household debt channel at the level of detailed regions throughout the UK.[6] We then estimate how house prices respond to interest rates in each region. If asset prices are important, then the household debt channel should be large in regions where house prices rise strongly after rate cuts. For example, suppose that house prices respond strongly to interest rates in London, but do not respond to rate cuts in Blackpool. If the asset prices were important, then the household debt channel would be small in Blackpool and large in London. We also develop two instruments for the response of house prices to rates, based on across-region variation in the housing supply elasticity (Guren, McKay, Nakamura and Steinsson, 2021), and within-region across-household variation in the duration of the housing stock (Giglio, Maggiori and Stroebel, 2015; Backer-Peral, Hazell and Mian, 2024).

Our third main result is that asset prices determine how rates affect borrowing and consumption. In regions where house prices do not respond at all to rates, the borrowing and consumption response are both around half of their size in the average region. In this average region, however, house prices rise by around 6% for every 1 pp rate cut, leading to a large borrowing response. We also study the marginal propensity to consume out of the higher housing wealth due to the rate cuts, estimating a value of 0.04. Previous work that estimates how house prices affect consumption—without connecting to monetary policy—tends to estimate similar MPCs out of housing wealth (e.g. Guren et al., 2021). Finally, we study heterogeneity. Consistent with the importance of asset prices and collateral, highly levered households are particularly sensitive to rate cuts.

Related literature. There is already a thorough literature estimating how interest

Economic Outlook, Chapter 2, 2024).

[6]Here, a region is equivalent to a county in the United States.

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. Wong, 2019; Flodén, Kilström, Sigurdsson and Vestman, 2020; Holm et al., 2021; Cumming and Hubert, 2023), 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 "mortgage cashflow effect" of rates on consumption is important: holding fixed house prices, rate cuts raise consumption by lowering debt service.[7] 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 two ways. First, we observe monthly-frequency transaction-level information on consumption, for both durables and non-durables. As such, we can estimate useful statistics that connect our estimates to quantitative models. Second, we show that movements in borrowing, due to asset prices, are a crucial part of why consumption responds to rates—alongside the cashflow effects that past work emphasizes.[8]

A second literature estimates how interest rate cuts affect consumption in aggregate, using monetary policy shocks identified from the time series (e.g. Romer and Romer, 2004; Gertler and Karadi, 2015). One 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 the household debt of mortgagors—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, and generate "long and variable lags" that are present in the time series. As such, household debt seems to be a crucial part of the monetary transmission mechanism.

[7] One dissent to this finding is Elias, Gillitzer, Kaplan, La Cava and Prasad (2025), who find that mortgage cashflow effects are small in Australia, which features adjustable-rate mortgages with redraw facilities that make mortgages liquid.

[8] Our paper also relates to recent work examining the effect of monetary policy during recent rate hikes of 2022-23. For instance, Bosshardt, Di Maggio, Kakhbod and Kermani (2024) show that debt-to-income ratios were a key factor mediating the effect of rate hikes in the United States. Caspi, Eshel and Segev (2024) study the effect of rising rates on consumption in Israel, using quasi-random exposure to adjustable rate mortgages. Bracke, Everitt, Fazio and Varadi (2024) use similar data to our paper and examine consumption and debt changes for mortgagors who refinanced in the UK, during the 2021-2023 policy hiking cycle.

A third literature emphasizes that monetary policy affects the economy in large part by raising asset prices. For instance, Caballero and Simsek (2024) review models in which monetary policy affects asset prices, which pass through to consumption with lags; a seminal paper is Iacoviello (2005). We provide empirical support for these ideas. Our paper focuses on household consumption, complementing a large literature linking monetary policy and asset prices to firm investment (e.g. Ottonello and Winberry, 2020; Bahaj, Foulis, Pinter and Surico, 2022).

Finally our paper relates to two previous studies of the UK housing market, Cloyne, Huber, Ilzetzki and Kleven (2019) and Best, Cloyne, Ilzetzki and Kleven (2020). Best et al. (2020) exploit discontinuities in the relationship between loan-to-value and mortgage interest rates to develop a bunching estimator. They show that holding fixed house prices, the response of borrowing to rate changes is small (i.e. a small intertemporal elasticity of substitution, IES). Our paper shows that the response of borrowing to rate changes is large when house prices vary. The two findings are consistent: even if the IES is small, borrowing responses can be large due to how rate cuts affect house prices. Cloyne et al. (2019) show that idiosyncratic movements in house prices lead to large changes in household borrowing, exploiting as we do the staggered expiry of mortgage deals. More broadly, there is evidence that house price increases spur home equity borrowing and raise consumption (e.g. Mian and Sufi, 2011; Mian, Rao and Sufi, 2013). This finding points to the importance of asset prices for consumption, which our paper links to monetary policy.

2 Data

We now describe our mortgage, housing, and consumption datasets.

Mortgage data. We measure mortgages from the Product Sales Data (PSD), an administrative dataset collected by the United Kingdom's Financial Conduct Authority. Our main dataset contains snapshots every 6 months of the universe of outstanding mortgages in the United Kingdom, from June 2015 until June 2024. At any point in time, there are roughly 8 million mortgages outstanding in the UK. We observe mortgage-related variables including the outstanding balance, the current interest rate, whether this rate is fixed or floating, the monthly payment associated with the mortgage, and the date the current deal ends. We also observe the mortgagor's full date of birth and the full property postcode, allowing us to track the same mortgage over time.[9] With this dataset we construct a monthly panel

[9] In the UK, postcodes typically contain around 15 properties, meaning the postcode and date of birth identify a property accurately.

tracking the evolution of each mortgage.[10]

Housing data. We use information 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 (equivalent to a county in the United States).

In the UK, the duration of properties varies. Some properties are owned in perpetuity ("freeholds"); whereas other properties are long duration leases, lasting decades or centuries ("leaseholds"). Therefore freehold properties have longer duration than leaseholds. We use data from the England and Wales Land Registry (specifically, the Price Paid and Registered Lease datasets) to measure the duration of each property. As described in Appendix C.1.2, we then merge property duration with our panel of mortgages.[11]

Consumption data. We use two consumption datasets, from ClearScore and Money Dashboard.[12] Both ClearScore and Money Dashboard are personal finance apps. The apps have a common owner and use 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 regular 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. We will refer to this measure as "consumption", with slight abuse of terminology because consumption of housing services is not included (we classify home-related renovation expenses as durables). We also observe various forms of income paid into accounts associated with the app. The most common is

[10] We discuss how we assemble the mortgage dataset in thorough detail, in Appendix C.1.

[11] More details on the housing data are given in Appendix C.2.

[12] Specifically, the dataset we use from ClearScore is the ExactOne dataset.

salary or wages, but other forms of income such as investment 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.[13]

To merge the finance app information to the mortgage dataset, we exploit that in the apps we observe the sector-level postcode (full postcode less the last two digits), year of birth, month of birth (ClearScore only), mortgage lender, and the monthly mortgage payment amount. In the mortgage dataset, we observe the full postcode, date of birth, mortgage lender, and the monthly mortgage payment amount. We merge on these variables, which select unique mortgage accounts in the PSD in over 99% of cases. ClearScore merges 162 thousand households to the PSD, whereas Money Dashboard merges 37 thousand households. Therefore the dataset is merged at the level of the mortgage holder—similar to, but distinct from, the level of the household.[14]

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 A.1 plots mean monthly expenditure for mortgagors in Money Dashboard and ClearScore, compared to mean annual expenditure among mortgage holding households 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.[15]

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 C.5. Appendix Figure A.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

[13] Appendix C.3 contains a detailed description of the consumption dataset.

[14] Appendix C.4 provides a detailed description of the merge between the mortgage and consumption data. After merging the datasets we drop the property postcode and date of birth variables and carry out all our analysis using an anonymised id key. Appendix Section C.4.4 explains how we deal with household-level bank accounts.

[15] 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.

will use these weights for all of our results concerning consumption.[16]

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.

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. In recent years, approximately 95% of new mortgage deals are 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, as shown in Appendix Figure A.3. 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 a typical increase of 200 basis points or more. 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.[17] Combining these features, households receive an interest rate change, as well as an opportunity to borrow using the mortgage, at staggered moments. In effect, these moments are pre-determined. In principle, households could borrow more during the deal length if the gains were large enough. However in practice, we shall see that borrowing more before the deal expires is rare.[18]

There is a 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

[16] Later, we will show that the response of cash-on-hand 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.

[17] A typical Early Repayment Charge is 3% of the outstanding mortgage balance.

[18] With certain mortgage deals, households can repay up to 10% of their mortgage per year, outside when the deal expires. This feature may introduce asymmetries, which we test for in our empirical strategy.

how the predetermined deal length is crucial for identification.

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. Appendix Table B.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 last two columns we report summary statistics about the sample with both mortgage and consumption information. In the middle 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 middle 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 event window around deal expiry, since we cannot easily track households in the mortgage data when they move properties.[19]

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.

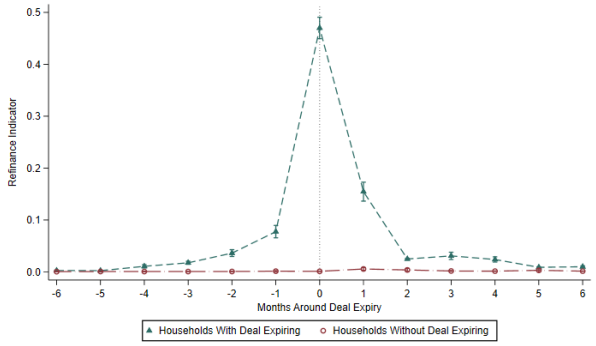
Figure 1 studies households with expiring deals. In the figure, t=0 is the point at which the deal expires in event time, with negative numbers indicating the preceding months and positive numbers the subsequent months relative to the expiry month. Then, in each month relative to the expiry month, we 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

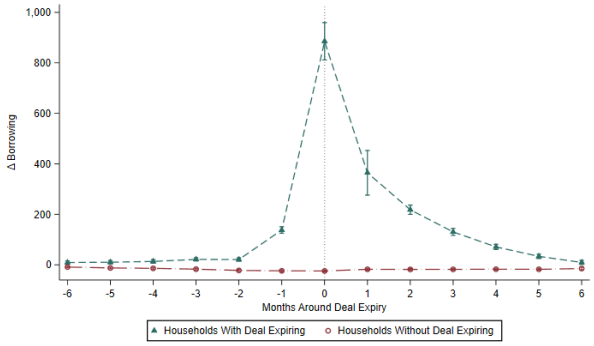
[19]In Appendix C.6 we show that home sales around deal expiry are relatively rare and mortgage rates have little effect on them.

Figure 1: Household Behavior around Deal Expiry

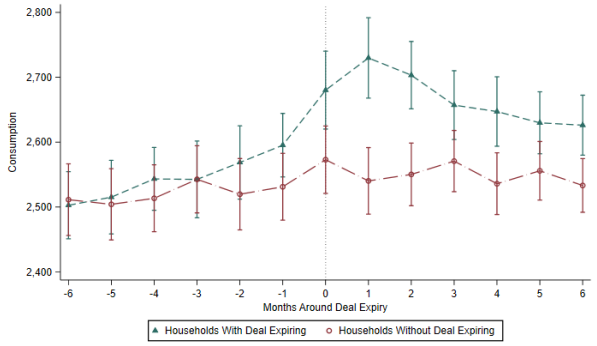
(a) Refinancing



(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 comprising of 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% confidence intervals for deal expiry month clustered standard errors. The sample is the mortgage population in panels (a) and (b), and the reweighted consumption sample in panel (c). All variables are winsorized, at the horizon level, separately for each group, at the 0.1% and 99.9% level.

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%.[20] 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 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 direc-

[20]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. Our strategy to identify the household debt channel, to be discussed in Section 4, does not depend on these issues.

tion, households are much more likely to adjust their behavior when deals expire. Moreover, households are very unlikely to adjust borrowing when their deal has not expired—suggesting that the costs to doing so are prohibitive. That is, the expiry of the deal is a pre-determined opportunity to borrow more. 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 to define the household debt channel of interest rate changes, and how our natural experiments identify this channel. We also show that the channel can be large when monetary policy affects asset prices: rate cuts increase house prices, letting households borrow more against their home. Finally, we link to the "long and variable lags" of monetary policy from Friedman (1961).

4.1 A Simple Model of Mortgages

We study a standard model with mortgages in which households can borrow against their home (e.g. Laibson, Maxted and Moll, forthcoming). 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 given utility $u(c_t) = c_t^{(1-1/\sigma)}/(1-1/\sigma)$, 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 labor income y_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. 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. $\mu_t = \mu\left(\frac{d_t}{p(r_t)h}\right)$. 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.[21]

Overall, the household maximizes the discounted present value of their utility from time

[21] We abstract from moving house, meaning we ignore wealth effects from higher house prices. Empirically, we will show that the effect of rate cuts on moving homes is relatively small.

$-N$ onwards, solving

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

by choosing consumption, 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 labor income, 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:

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

Households 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 otherwise 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:

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

Here, we have modeled a "soft" collateral constraint, in which households can always borrow more, but only at a higher interest rate (e.g. Bernanke, Gertler and Gilchrist, 1999). This specification nests the case of a "hard" collateral 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 a maximum loan-to-value ratio, and the household can borrow only a fixed fraction of their home value.[22]

4.2 The Household Debt Channel

We can now define 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 labor income y_0 and the rate of return on the liquid asset r_0^a. The response of consumption

[22] In the UK, spreads discontinuously increase at certain LTV thresholds, and are flat or moderately upward sloping otherwise, representing a functional form for $\mu(\cdot)$ that mixes hard and soft constraints (see Figure 3 of Best et al., 2020).

at time 1 is

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

In this equation, Δy and Δr^a are the perturbations to labor income and the return on liquid assets at time 1, which affect consumption with coefficients $\mathcal{Y}$ and $\mathcal{A}$. These coefficients, as well as all derivations, are in Appendix Section D.1. $r_0 - r_{-N}$ 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. The response of cash-on-hand is the sum of how borrowing and debt service changes. That is, the response of cash-on-hand to changes in the aggregate rate satisfies

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

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 $\partial \mu_0 / \partial r_0$. $\text{MPC}^{\text{cash-on-hand}} \equiv \Delta c_1 / \left(\Delta c_1 + \Delta a_1 - (1 + r_0^a)\, \Delta a_0\right)$ 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 associates with a change in interest rates.

We define the household debt channel in terms of reduced form objects, such as the MPC out of cash-on-hand or the response of cash-on-hand to rate cuts. These objects summarize the relationship between borrowing, debt service, consumption and income after rate cuts, without imposing strong assumptions. As we shall see these objects can be estimated using our natural experiment. Solving for these objects as a closed form function of primitives is not generally possible. However, Appendix Section D.3 does solve for all the objects in equation (1) in a simplified version of the model.

Separate from the household debt channel, rate cuts affect consumption by raising income in general equilibrium. This force is represented by the $\mathcal{Y} \Delta y$ 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 $\mathcal{Y}\Delta y$ 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. There may also be asymmetries, if households are unable to borrow to take advantage of future rate cuts, but can save against future rate increases. In our empirical analysis, we will test for such asymmetries.

4.3 Borrowing and the Effect of Monetary Policy on Asset Prices

The household debt channel is large when borrowing responds strongly to rate cuts. In turn, borrowing can respond to rate cuts for two reasons. First, there is a mechanism that operates through asset prices: rate cuts raise house prices, which raises consumption because households borrow against their home. This mechanism is easiest to see in the special case of our model with binding "hard" collateral constraints. Formally, in this special case, the function mapping loan-to-value to mortgage rates is flat beneath some maximum loan-to-value ratio, and vertical at the maximum loan-to-value. In the special case, we assume the collateral constraint is binding, so households borrow the maximum loan-to-value. In this case, the response of borrowing to rate cuts is

$$\frac{\partial d_1}{\partial r_0} = \text{maximum loan-to-value} \times h\, p'(r), \tag{3}$$

that is, the product of the maximum loan-to-value with how home values respond to rate cuts. Although we stress this effect less, with "soft" collateral constraints higher house prices also lower debt service after rate cuts, because lower loan-to-value reduces mortgage spreads (see equation 2).[23]

A second reason why households might borrow more after rate cuts is the standard intertemporal substitution mechanism: if borrowing via the mortgage has become cheaper, households bring consumption forward. To see this mechanism, consider a second special

[23]Since we abstract from the buying and selling of homes, we ignore "housing wealth effects", whereby higher house prices raise lifetime wealth of homeowners and increase their consumption. Section 6.4 finds that collateral appears to matter more than housing wealth effects, and Appendix Section C.6 shows that the effect of rate cuts on home moving is small.

case. Suppose that collateral constraints never bind, and households can borrow only with the mortgage. Formally, the mortgage rate spread μ would always equal zero, and the constraint on the liquid asset $\underline{a}$ would be sufficiently low. In that case, the response of borrowing is

$$\frac{\partial d_1}{\partial r_0} \approx -\sigma c N, \tag{4}$$

where c is steady state consumption at time 0. This expression, which we derive in Appendix Section D.2, shows that a rate cut raises consumption according to the intertemporal elasticity of substitution, and the length of time N for which the rate cut applies.[24]

Away from these two special cases, the response of borrowing to rate cuts depends on both intertemporal substitution and the response of asset prices to rate cuts; the precise form of this response cannot be characterized in closed form. Borrowing can respond to rate cuts even if intertemporal substitution is small (Best et al., 2020). Regardless of intertemporal substitution, borrowing responds if monetary policy affects asset prices.

There are reasons why the household debt channel could be large other than the effect of rates on asset prices. One prominent reason is "mortgage 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 effect of asset prices on rate cuts 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 and borrowing via the effect of rates on asset prices—the focus of this paper—is more scarce.

4.4 Long and Variable Lags from the Household Debt Channel

In the model, the effect of rate cuts on aggregate consumption will be delayed—the "long and variable lags" of monetary policy of Friedman (1961). The reason is that deals expire only infrequently. Therefore a rate cut passes through to the consumption of a household only when the household's mortgage deal expires, which may be some time after the rate cut takes place.

To see this point formally, suppose now that there is a population of otherwise-identical mortgagors whose behavior satisfies equation (1), but with deals expiring in various time periods. Then on impact, the response of aggregate consumption, amongst mortgagors, to

[24] This result requires some minor additional simplifying assumptions that we discuss in the derivation.

the rate cut at time 0 is

$$\frac{\partial \log c_0^{\text{agg}}}{\partial r_0} = \text{fraction of deals expiring at time 0} \times \frac{\partial \log c_0}{\partial r_0}. \tag{5}$$

That is, the impact response of aggregate consumption to rate cuts is smaller than what it would have been, if all households had been exposed to the rate cut. The reason is, of course, that only the subset of households had expiring deals at time 0; the other households will be exposed to the rate cut later on. Therefore the impulse response of aggregate consumption to rate cuts is more back-loaded than the impulse response for an individual with an expiring deal—the long and variable lags of Friedman (1961).

Put differently, the effect of rate cuts on consumption is "pent up". An interest rate cut does not affect consumption, via the household debt channel, at the moment the rate cut is realized. Rather, it is when deal expiration happens that the household can take advantage of the rate cut.

4.5 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, while in Section 6 we will link to the importance of asset prices.

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. According to equation 1, the consumption of household i with a deal expiring at time t satisfies

$$\Delta c_{it} = \text{MPC}^{\text{cash-on-hand}} \left[\frac{\partial \text{cash-on-hand}}{\partial r}\right] (r_t - r_{t-N_i}) + \alpha_t + \varepsilon_{it}, \tag{6}$$

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.[25]

One can estimate equation (6) by regression. The outcome variable is the consumption

[25] Equation (6) also removes time subscripts where there is no ambiguity.

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, 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.[26]

5 Estimates of the Household Debt Channel

This section presents our first main result: 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 channel. The household debt channel is large—a 1 pp mortgage interest rate cut raises borrowers' consumption by 3% over 6 months, which implies an aggregate increase in consumption worth 0.7% of GDP.

[26] Appendix Section D.3 presents an extension of our model with endogenous labor supply, which emphasizes this point.

5.1 Implementing the Main Regression

We first adapt the regression equation (6) to our setting. We estimate a series of regressions

$$\sum_{j=-6}^{h} [Y_{i,t+j} - Y_{i,t_0}] = \alpha_t^h + \beta_h (r_t - r_{t-N_i}) + \text{fixed effects}_{it}^h + \varepsilon_{ith} \tag{7}$$

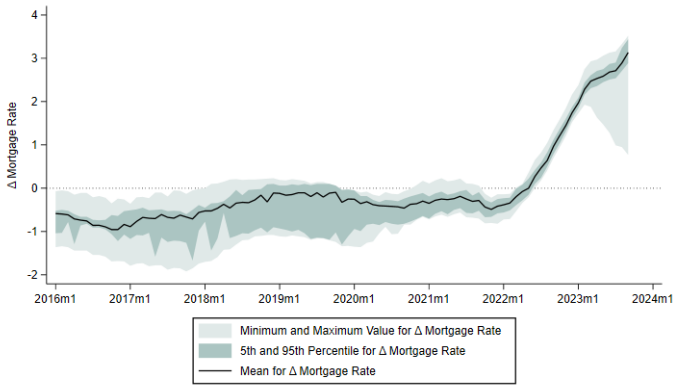
for each monthly horizon h. The primary outcome variables Y_{it} that we study are monthly consumption, and the monthly change in cash-on-hand (i.e the monthly change in mortgage borrowing minus the change in debt service costs). We take the difference between the outcome $Y_{i,t+j}$ and a base observation 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 from three months before the deal expires onward, in order to measure the cumulative response of consumption or cash-on-hand.[27] The regression is estimated at monthly frequency. α_t^h 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, provided by the Bank of England. This interest rate purges individual variation, which may be endogenous to household choices.[28]

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 (7). Figure 2 shows this source of variation. The black line plots the mean change in the interest rate on new mortgage deals between 2016 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. 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

[27] Appendix C.1.6 provides further details on the construction of the cumulative change in cash on hand.
[28] 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 2016-2024, 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.

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—or households with deals of slightly different length expiring at the same time of the year, for instance comparing an expiring 22-month deal with an expiring 26-month deal.[29] 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.

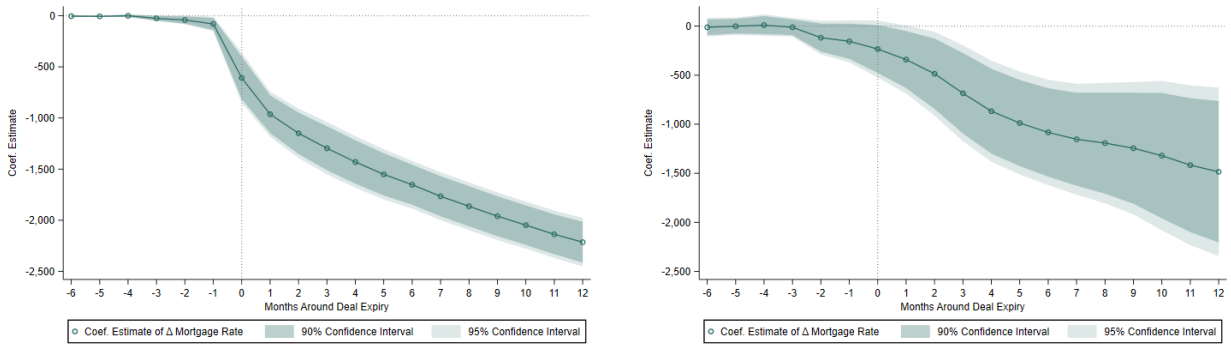
What is clear from Figure 2, and the three sources of variation that we exploit, is the value of 6 million natural experiments. While one can think of omitted variables that might affect any single source of variation, it is more difficult to think of omitted variables affecting all three sources. If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

[29] This latter comparison is possible as, whilst deal lengths cluster around a whole number of years, there is a distribution around these points, as shown in Appendix Figure A.3. This in turn occurs as lenders often post deals that expire at a specific calendar date, so households' deal lengths depend on when their deal started.

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 for deal expiry month clustered standard errors. Both regressions include year and age fixed effects. Panel (a): cash-on-hand for the universe of mortgagors. Panel (b): reweighted consumption for households with app-based spending. Outcomes are winsorized, horizon by horizon, at the 0.1% and 99.9% level.

5.2 Estimates of the Household Debt Channel

We now present our first main result: the household debt channel is large. We first 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 estimates of the change in cash-on-hand over the event window in Table 1, Panel A, column (1). This column contains the cumulative response of cash-on-hand from 3 months before deal expiry to 6 months after.[30] The estimated effect is an 1,818-pound increase after a 1 pp 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

[30]Formally, for the regression over the event window, we define the change in cash on hand as the change in mortgage principal over the event window less 10 times the estimated change in monthly payments. Appendix C.1.6 provides further details.

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 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.

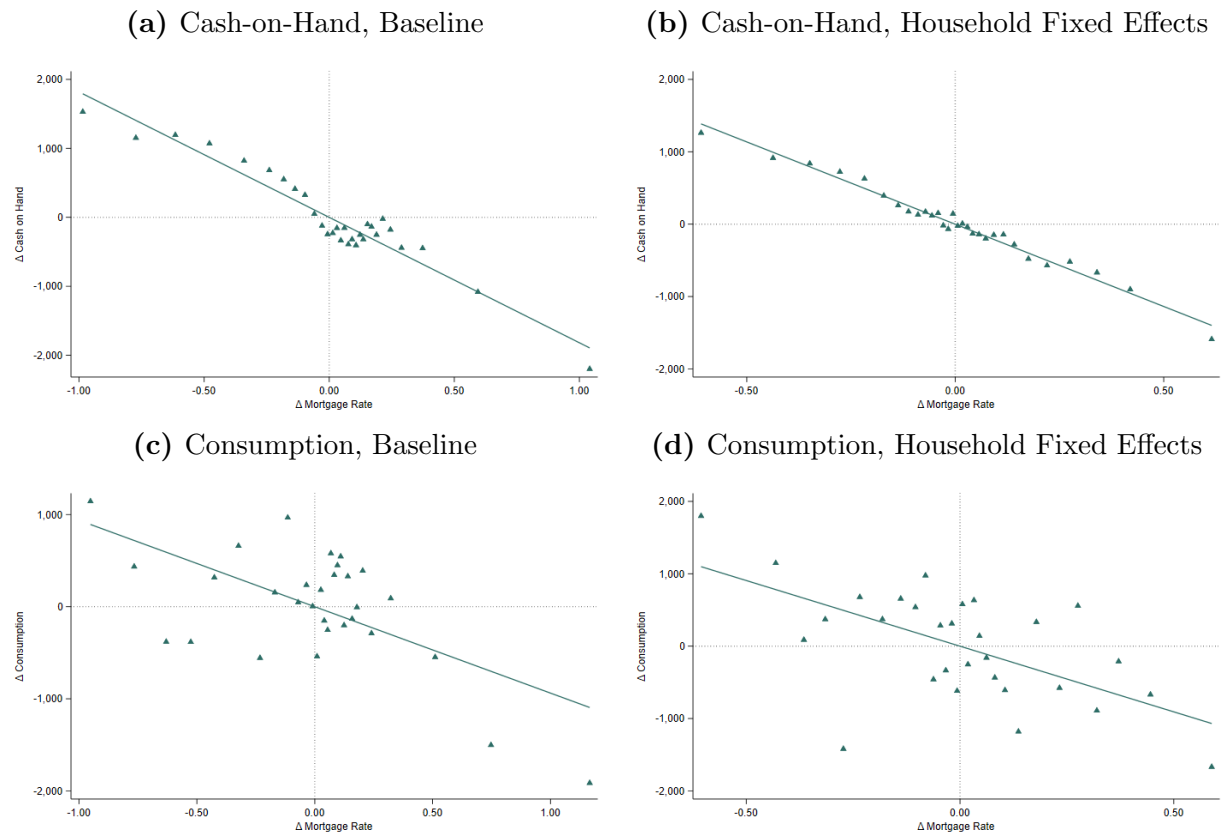
We also plot the variation associated with the estimates in Figure 4. We residualize both the 6-month response of cash-on-hand and the rate change against the baseline controls, i.e. year and age fixed effects. We then bin both variables into 30 groups, and plot them. Inspecting panel (a) of the figure, one can see a robust linear relationship between rate changes and cash-on-hand changes. That is, there do not appear to be strong asymmetries or size dependence. 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 1, Panel B, column (1) presents the baseline estimates for consumption over the event window. This column contains the cumulative response of consumption until 6 months after deal expiry, which is a 938 pound increase after a 1 pp rate cut.[31] 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 the log change in average consumption as the outcome variable, again size-weighting each observation by mortgage principal before the deal expires. A 1 pp rate cut raises average consumption by 3.0% in the window around deal expiry. 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

[31] Appendix Section C.4.5 provides further details on the regression specification for consumption.

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



Notes: each panel plots binned averages of the change in the outcome (pounds, 2015 base) in the window around 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) residualize on household fixed effects instead of age 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. Outcomes are winsorized, at the 0.1% and 99.9% level.

and consumption increases—albeit noisier than for cash-on-hand, given the smaller sample size.[32]

These estimates of the household debt channel are large. Consider the 1pp mortgage rate cut that increases consumption 3% by 6 months after deal expiry. Multiplying this effect by the share of consumption contributed by households with a mortgage, which is 39%, and the consumption share of GDP, which is 59%, this effect is approximately an increase in consumption worth 0.7% of GDP over 6 months.[33] This illustrative calculation assumes all

[32]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 (Appendix Table B.2).

[33]The source is the Office for National Statistics for 2024 Q1. The mortgagor share of household expenditure is taken from Table A32 (total household expenditure by tenure type, family spending workbook 2).

mortgagors face the change in mortgage rates immediately. In Section 5.3 we explore the dynamic response of consumption due to the staggered expiry of mortgage deals.

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 out of cash-on-hand. We regress the cumulative change in consumption until h months after deal expiry, on the cumulative change in cash-on-hand until 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 for every horizon h from 6 months before until 12 months after deal expiry. As such, we trace out the horizon-specific response of consumption, h periods after deal expiry; to the cumulative change in cash-on-hand due to rate cuts, in the 12 months after the deal expiry. The MPC gradually increases, slightly exceeding 0.7 after 12 months.[34] 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 durable and non-durable consumption, the shock is relatively persistent; and also involves changes in relative prices.

We exploit an advantage of our dataset, that we can disaggregate the consumption response by durable versus non-durable consumption. We separately estimate the response of the two consumption categories to the rate cut, and report the result in the impulse response of Appendix Figure A.4. The consumption response is concentrated in non-durables.[35]

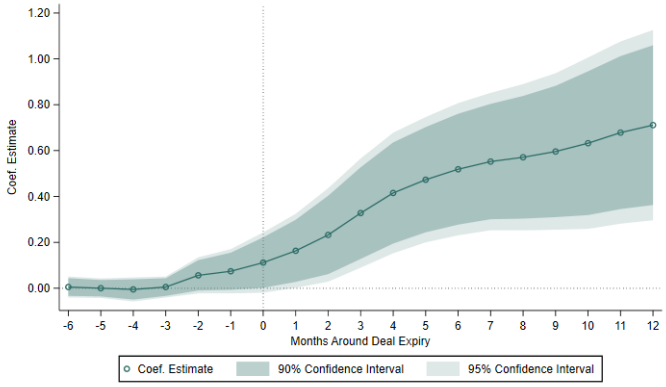
We next directly explore the link to monetary policy and consider how households respond to changes in mortgage rates that were caused by monetary policy. To identify monetary policy shocks we use the monthly series of Braun et al. (2025), which considers high frequency monetary policy surprises for the UK around Bank of England policy decisions. To cleanly

The consumption share of GDP is taken from Worksheet C1 on GDP by category of expenditure (series ABJR: Household Final Consumption Expenditure).

[34] Table 1, Panel C, presents the corresponding estimates of the MPC over 6 months. The estimates with household fixed effects are noisy, but Appendix Table B.3 uses a two sample two stage least squares procedure, with similar but more precise results (Angrist and Krueger, 1992). We estimate the first stage on the full population of mortgage borrowers, and the reduced form on the subsample for with consumption. This procedure is valid if both populations respond similarly to rate cuts, for which we provide evidence.

[35] Durables respond less than non-durables in levels, because the quantity of non-durable spending is much larger (£2 247 per month for non-durables and £131 for durables). However, in percent terms, durable spending is more elastic to rate cuts (after a 1 pp rate cut, durable spending increases by 8.6% over the next 12 months, non-durables increase by 4.4%).

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). The sample is the reweighted merged consumption-mortgage dataset. For each horizon $h \in [-6, 12]$ months relative to deal expiry ($t = 0$), we regress the cumulative change in consumption from months -6 through h on the cumulative change in cash-on-hand from months -3 through $+12$, instrumenting the cash-on-hand change with change in aggregate rates over the expiring deal. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. The y-axis reports the MPC coefficient; the x-axis is event time h in months. The regressions include year and age fixed effects. Point estimates are shown with 90% and 95% confidence intervals for deal expiry month clustered standard errors. Both the outcome and the regressor are winsorized, horizon by horizon, at the 0.1% and 99.9% level.

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 2 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 B.4). 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. The results are again similar when the mortgage rates 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.

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 (7) with monthly wage income as the outcome variable, as inferred from the spending apps. Appendix Figure A.5 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 800 pounds (see Appendix Figure A.6 for a version including non-labor income). 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.[36]

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 (Appendix Table B.5).[37] 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.[38]

5.3 Aggregate Impact and Long and Variable Lags

We now show that the household debt channel seems to matter for how consumption responds to rate cuts in aggregate. The magnitude of the household debt channel is similar to time series estimates of how rate cuts affect consumption. The household debt channel also leads to "long and variable lags", similar to the time series; we show that the extent of the lags

[36] The relative response of income and consumption are what the standard model predicts, provided that the intertemporal elasticity of substitution and the Frisch elasticity of labor supply are similar (see Appendix Section D.3).

[37] 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.

[38] We also show that our results are robust to dropping the Pandemic (Appendix Table B.6), using different measures of the aggregate rate (Appendix Table B.7), and using consumption data without reweighting (Appendix Table B.8). Finally, we study the effect of rate cuts around deal expiry on the propensity to move house, finding a small effect (Appendix Section C.6).

depends on both the duration of mortgage deals, and the persistence of the initial interest rate shock.

We previously estimated the response of consumption conditional on a deal expiring, in Figure 3. A 1 pp rate cut raises consumption of mortgage holders by a cumulative 1500 pounds after 12 months. We can compare this number to consensus estimates of how rate cuts affect consumption, in aggregate. For instance, Cloyne et al. (2020) estimate the aggregate response of consumption, for mortgage holders in the UK, to interest rate shocks. Shocks are identified as in Romer and Romer (2004). That paper estimates that after 4 years, a 1 pp rate cut raises consumption of mortgage holders by a cumulative 1500 pounds—similar to the contribution of the household debt channel alone.[39] Cloyne et al. (2020) also show that most of the response of consumption to rate cuts is due to mortgage holders, as opposed to renters or outright owners. We conclude that the household debt channel accounts for a large share of the aggregate response of consumption to rate cuts.[40]

The household debt channel also generates "long and variable lags" in the response of consumption to rate cuts. To see this point, we carry out a simple calibration informed by the model and our estimates. Generalizing equation (5) from the model, the response of aggregate consumption among mortgage holders, j periods after a rate cut at time 0 that lasts for at least j periods, is

$$\frac{d \log c_j^{\text{agg}}}{dr} = \text{frac. of deals expiring at } j \times \frac{\partial \log c_0}{\partial r} + ... + \text{frac. of deals expiring at } 0 \times \frac{\partial \log c_j}{\partial r}. \tag{8}$$

That is, j periods after a rate cut, the aggregate response of consumption is a weighted average of the individual-level responses of consumption to the rate cut. The individual response occurs conditional on a deal expiring, and the average weights across the fraction of households with expiring deals at each horizon. Therefore the response of aggregate consumption to rate cuts is delayed, because individuals alter their consumption only when their deal expires.

We now show that the lags of monetary policy are long. For this calculation, we use the impulse response function from Figure 3, which maps to estimates of $\partial \log c_j / \partial r$ in equation (8). We also use the sample average of the fraction of deals expiring at each horizon after the shock.[41] Figure 6, panel (a), shows the effect of a unit rate cut that lasts for 60 months,

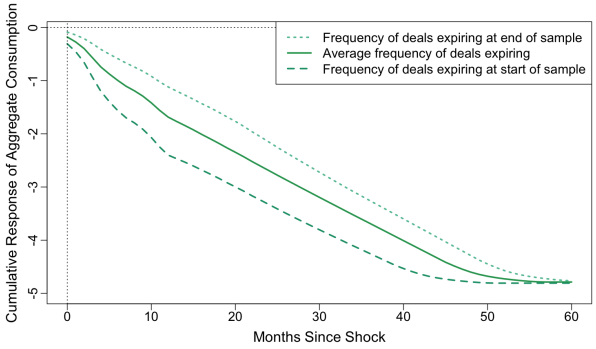
[39] Table 1 of Cloyne et al. (2020) shows that a 0.25 pp rate cut raises consumption by 600 US dollars (measured in 2007), which we convert into UK pounds (measured in 2015).

[40] Our estimate is not wholly comparable to Cloyne et al. (2020), for two reasons. First, our estimate is the 12-month response of consumption, compared to the 48-month response of consumption in the latter paper. Second, the persistence of the interest rate shock may be different in the two papers.

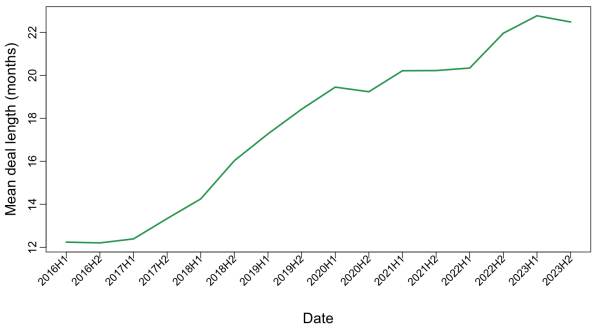
[41] We assume the impulse response from Figure (3) is constant after 12 months; in practice, the confidence intervals widen at this point. We treat floating rate mortgages as deals that expire at time 0, receiving an

Figure 6: Aggregate Impulse Response of Consumption

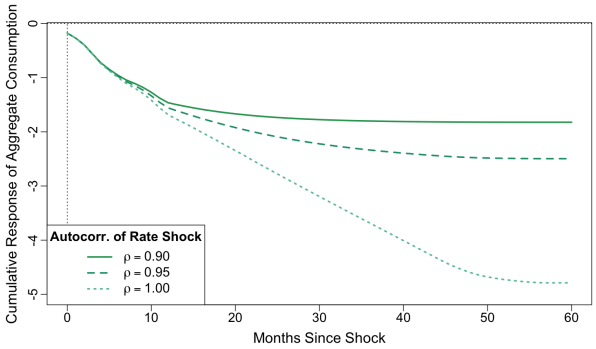
(a) Impulse Response, Varying Lags



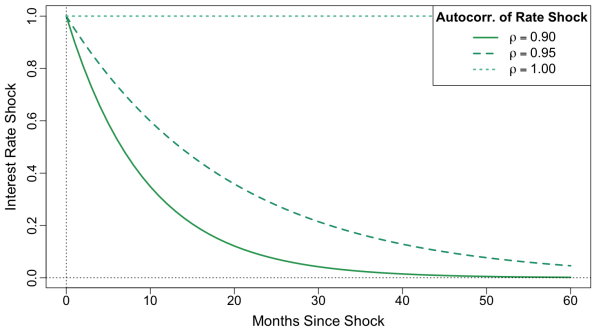
(b) Average Deal Length



(c) IRF, Varying Shock Persistence



(d) Persistence of Rate Hike



Notes: Panel (a) plots the implied aggregate impulse response of consumption, to a 1 pp rate cut that lasts for 60 months, combining the household-level impulse response in Figure 3 with the distribution of deal expiry. The units are the cumulative change in consumption, relative to average annual consumption, in percent. Separate lines show results using the average expiry distribution (black) and distributions at the start (2016) and end (2023) of the sample. Panel (b) shows the mean deal length over time and the implied fraction of mortgages resetting in each month, highlighting how refinancing frequency changes across the sample. Panel (c) shows how the IRF varies with interest rate hikes of differing autocorrelation ρ, these hikes are shown in panel (d).

after cumulating the impulse response of equation (8) and scaling by annual consumption. Clearly, the response of aggregate consumption to the rate cut is delayed. The aggregate response of consumption to rate cuts approaches its peak after around 50 months. Only then does the aggregate response equal the peak household level response, which occurs after 12 months. For context, Cloyne et al. (2020) estimate that the impulse response of aggregate consumption, without cumulating, peaks after around 50 months. Therefore there is a similar, but greater degree of persistence in the consumption response from the time series.

Friedman (1961) argued that the lags of monetary policy were not only long but also variable. We can calculate how variable these lags are due to mortgage deals. One could contemplate a world in which mortgage deals are either long or short, and use equation (8)

interest rate shock in this period only.

and our estimates to calculate the lags of monetary policy. We carry out these calculations for some relevant cases. Figure 6, panel (b), shows that deal lengths have been rising in the UK.[42] We re-calculate the aggregate response of consumption to rate cuts using the frequency of deals expiring either at the start of our sample, in 2016, or at the end of the sample, in 2023. Figure 6, panel (a) reports these results as well. Clearly, the delay in the consumption response is much greater in the second period. For instance, after 12 months, the aggregate impulse response is almost twice as large in 2016 as in 2023.

We can also use our estimates to guide monetary policymaking. The previous estimates study the response of consumption to a permanent interest rate shock. In reality, shocks to interest rates from monetary policy are persistent, but not permanent; the persistence is a key choice for the policymaker. However we can combine interest rate shocks of varying persistence, together with our estimates and a suitably modified version of equation (8). In doing so, we trace out how consumption responds, via the household debt channel, to various kinds of monetary shocks.[43]

Figure 6, panel (c) reports the result, for an autocorrelation of the monthly interest rate shock between 0.9 and 1.[44] One lesson that emerges from the figure is that the persistence of the shock affects its cumulative impact, with more persistent shocks having a larger cumulative impact. Transitory shocks have dissipated before many households' deals have expired, meaning a smaller share of aggregate consumption adjusts to the shock.

As such, the household debt channel appears to be an important part of monetary policy transmission. One caveat is that our calculation of the aggregate impulse response is particular to the UK. Nevertheless, we can speculate about the household debt channel in other countries. As equation (8) makes clear, the fraction of households that refinance in each period will shape the aggregate impulse response. In other mortgage markets such as the United States, refinancing is a choice by households, which is affected by rate cuts. These considerations affect the shape of the impulse response of aggregate consumption, and may introduce asymmetries.[45] However the long-run cumulative impulse response is not affected by when households refinance, when interest rate cuts are permanent. After all, all households must refinance eventually, meaning the long run effect of rate cuts on consumption will be less affected by the precise timing of when refinancing happens. Therefore the cumulative

[42] Rajan, Rodriguez-Tous and Salgado-Moreno (2025) analyze the reasons why the deal length varied. Histograms of new deal lengths at different points in time are shown in Appendix Figure A.3.

[43] Formally, with persistent interest rate shocks of autocorrelation coefficient ρ, the response of consumption is $\frac{d \log c_j^{\text{agg}}}{dr} = \sum_{k=0}^{j} \left[\rho^{(j-k)} \frac{\text{fraction refinancing at } (j-k)}{c} \frac{\partial c_k}{\partial r_0} \right]$.

[44] For comparison, Groth and Wheeler (2008) estimate a monthly autocorrelation of monetary shocks, inferred from a Taylor Rule, that is somewhat larger than 0.9.

[45] In particular, rate cuts tend to affect the likelihood of refinancing more than rate hikes in the US, since refinancing is optional for households (Berger et al., 2021).

impulse response of aggregate consumption may be more similar in the US and the UK, at least if interest rate shocks are persistent enough. A second caveat is that our exercise takes place in partial equilibrium. In general equilibrium, 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.[46]

6 Borrowing and the Role of Asset Prices

So far we have found that the household debt channel is large. This section shows that one important reason for the large response of consumption and borrowing is how monetary policy affects asset prices: rate cuts lead house prices to rise, and households can borrow more against their home in order to consume.

6.1 Response of Borrowing vs. Debt Service

As a first finding suggesting that asset prices are important, we study how the two components of the household debt channel respond to rate cuts. Borrowing responds more than rate cuts, as Figure 7 shows. In the figure, the dark line is the 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 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, being approximately two thirds after 6 months.[47]

This finding hints that asset prices could be important. As we have discussed, rising asset prices allow households to borrow more after rate cuts. Other prominent mechanisms such as "mortgage cashflow effects" primarily operate through changes in debt service and not borrowing.

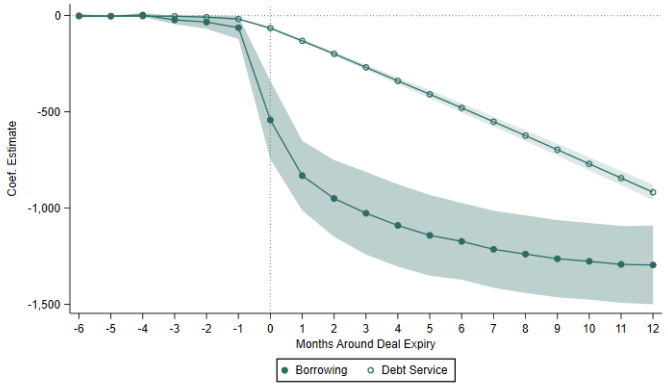
6.2 Identifying the Role of Asset Prices

We now develop a method to identify the role of asset prices, 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.

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

[47] Appendix Table B.9 reports regressions for this exercise over the event window around deal expiry. One reason why debt service responds, beyond the purely mechanical effect, is that households respond to rate cuts by lowering the term of the mortgage, as Appendix Table B.10 shows.

Figure 7: 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 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. Both regressions include year and age fixed effects. Shaded bands denote 90% confidence intervals for deal expiry month clustered standard errors. Sample: universe of mortgagors. We winsorize the outcome variable, horizon by horizon, at the 0.1% and 99.9% level.

To recap our model, borrowing responds more to rate cuts if house prices rise more. 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. There are also other reasons why households borrow more after rate cuts that do not depend on asset prices, such as intertemporal substitution.

Separating out the role of asset prices is not straightforward. One must disentangle house price movements from the other forces, such as intertemporal substitution, which affect borrowing after rate cuts. At this point the advantage of our setting—the millions of expiring deals—becomes important. Given this granularity, we can estimate the household debt channel for many disaggregated groups of households, whose house prices are affected differently by interest rates.

We modify the baseline regression equation (7) by using regional variation. In particular, we can estimate the modified regression equation

$$\Delta_{t,t+h}\text{borrowing}_{it} = \alpha_{j(i)t} + \beta_0 \left(r_t - r_{t-N_i}\right) + \sum_{q \geq 1} \beta_{FA,q} \times I\left[\eta_{j(i)} \in \text{quantile}_q\right] \left(r_t - r_{t-N_i}\right) + \text{FE}_{it} + \varepsilon_{it} \tag{9}$$

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, equivalent to a United States county) to which

household i belongs. $\alpha_{j(i)t}$ 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.[48] FE_{it} is a series of other fixed effects that we add to the regression. 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 asset prices. With estimates of both parameters in hand, one can estimate how much of the overall response of borrowing to rate cuts is due to asset prices. 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 Blackpool. Suppose that after a period of falling aggregate interest rates, house prices rise in London; whereas house prices do not change in Blackpool. If asset prices matter for the transmission of monetary policy, then borrowing should respond a lot in London, and little in Blackpool.

What variation identifies our regression? As we have discussed, some of the variation that identifies the regression comes from differences in deal length. Imagine that in both London and Blackpool, 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 asset prices are 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 Blackpool, however, neither household experiences rising house prices—meaning the increase in borrowing should be similar and small for both if asset prices matter. 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 Blackpool if asset prices matter.

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 Blackpool. If so, then borrowing will be different after rate cuts in the two regions, even if asset prices do not matter.

[48] Further details of this estimation, and well as robustness to alternative methods of estimating the regional $\eta's$, are given in Appendix C.7.

We therefore isolate two plausibly exogenous sources of variation in households' sensitivity to rate cuts. First, we use an estimate of the housing supply elasticity for each region in the United Kingdom, by adapting the method of Guren, McKay, Nakamura and Steinsson (2021) to the United Kingdom as described in Appendix C.7.2. In regions with inelastic housing 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 how asset prices affect mortgagors. 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 Role of Asset Prices

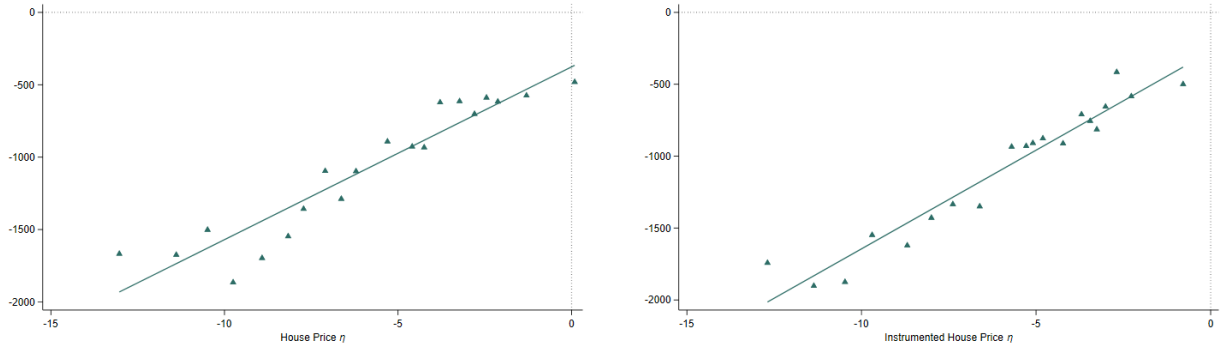
We now present estimates of how much asset prices matter for the pass-through of monetary policy into consumption. We start by implementing a version of our method non-parametrically, in Figure 8, panel (a). First, we estimate the household debt channel separately for 20 equal-sized bins based on the size of the $\eta_{j(i)}$ estimates, by estimating equation (9) over the event window until 6 months after deal expiry. 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 asset prices matter for how rate cuts affect 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 Blackpool, 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. Regions in the middle of the graph, whose house prices respond to rate cuts by the average amount, have a fairly strong response of borrowing. For the average region, a 1 pp rate cut increases house prices by around 6%, consistent with time series estimates of the effect of rates on house prices in the UK (Albuquerque, Lazarowicz and Lenney, 2025).

Table 3 presents regression evidence that is consistent with the scatter plot. Column (1)

Figure 8: Response of Borrowing to Rates vs. Response of House Prices to Rates

(a) OLS Specification **(b)** Housing 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 change in borrowing (pounds, 2015 base) in the event window until 6 months after deal expiry per 1 pp change in the interest rate, where we estimate β_{FA} from Eq. (9) 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. Both regressions include year x Local Authority District and age fixed effects. Sample: mortgagor households. We winsorize the outcome variable at the 0.1% and 99.9% level.

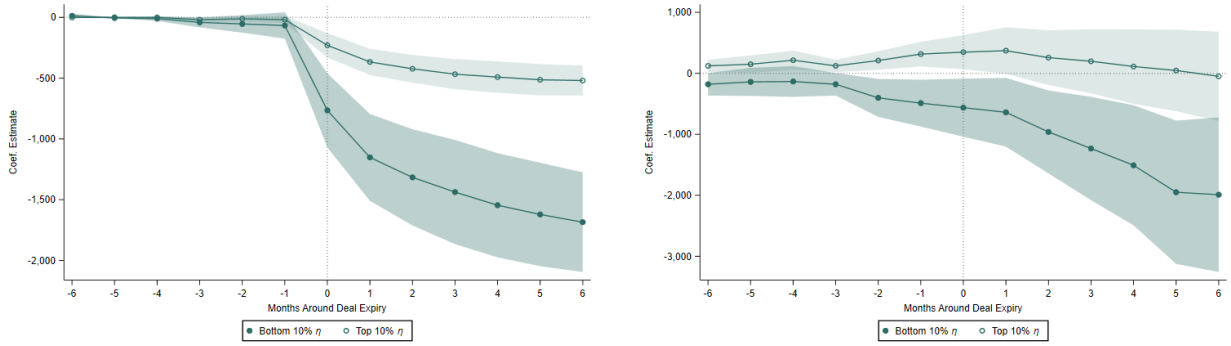
is the baseline specification with year x region fixed effects. The outcome variable is the response of borrowing from 3 months before to 6 months after the deal expires. Our measure of the house price response $\eta_{j(i)}$ is a series of three dummy variables; for the bottom 10%, middle 80% and top 10% of the raw values of $\eta_{j(i)}$.[49] Asset prices appear 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 9, 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.

How important are asset prices for the overall effect of rate cuts on borrowing? We present a simple calculation that suggests asset price movements accounts for around half of the borrowing response to rate cuts. As one can see from Figure 8, households with house

[49] We discretize instead of interacting rates with the value of $\eta_{j(i)}$, because the binscatter of Figure 8, panel (a) suggests that the relationship between house prices responses and borrowing responses is approximately piecewise linear around these thresholds.

Figure 9: Impulse Responses by House Price Sensitivity

(a) Borrowing **(b)** Consumption



Notes: the figure plots 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 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 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. Both regressions include year x Local Authority District and age fixed effects. Shaded bands denote 90% confidence intervals for deal expiry month clustered standard errors. We winsorize the outcome variable, horizon by horizon, at the 0.1% and 99.9% level.

prices in the least responsive 10% are those with prices that barely respond to rate cuts (row 3 of Table 3). As such, the estimate of the borrowing response for these households approximates β_0 from regression equation (9), that is, how rate cuts affect borrowing holding fixed house prices. However row 2 of Table 3 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 2 and 3 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 movements in asset prices. In our baseline specification of column (1), by this calculation asset price movements 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 housing supply elasticity instrument.

Figure 8, 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 housing supply elasticity instrument. We again present a scatter plot for the estimates of the 20 bins for the x and y variables. Once more, movements in rate cuts due to asset prices 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.

Appendix Table B.11 presents the results associated with this variation. We again 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 importance of asset prices, while each row measures how borrowing responses differ by housing supply elasticity. As with the baseline estimates, borrowing responds much more in regions with inelastic housing 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. Therefore we can absorb any purely regional variation in the household debt channel, including regional differences in intertemporal substitution. 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. As such, we replace the indicators for house price sensitivity, $\eta_{j(i)}$ in equation (9), with an indicator for whether the property is a freehold. We also add further controls to the regression: year by freehold fixed effects, to account for time varying shocks to freeholds; and changes in the mortgage rate interacted with Local Authority fixed effects, to control for differing regional sensitivity to rate cuts. These controls absorb any heterogeneous regional responses that do not vary across time.[50]

Table 4 presents the results. Columns (1)-(4) repeat our previous specifications that test for the role of asset prices (with the addition of the new fixed effects). Consistently, we find that borrowing contracts by around 750 pounds more for properties owned freehold, after a 1 pp rate hike. Column (5) introduces a saturated specification, controlling for the interaction of fixed effects at the start and end of the deal, interacted with Local Authority Fixed Effects. This regression compares the response of two households in the same local region whose mortgage deals start and finish in exactly the same months. The estimates are quantitatively similar to Column (1)-(4).[51]

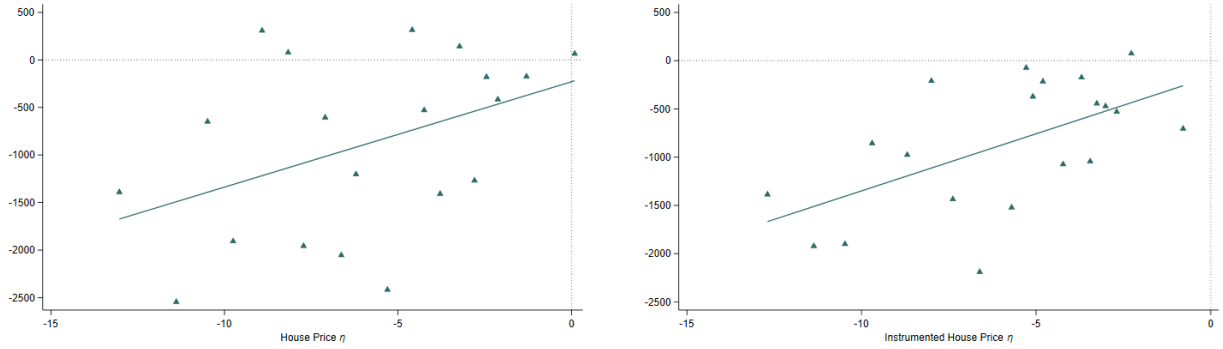
The response of asset prices to rate cuts matters not only for borrowing, but also for

[50] Appendix Table B.12 provides evidence that house prices respond more to rate cuts for longer duration properties.

[51] The role of asset prices is equally strong with both instruments for how house prices respond to rates, as well as the least squares variation. Appendix Table B.13 estimates the marginal propensity to borrow out of the house price changes due to rate cuts, and finds similar estimates for each source of variation.

Figure 10: Response of Consumption to Rates vs. Response of House Prices to Rates

(a) OLS Specification **(b)** Housing 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 change in consumption (pounds, 2015 base) over the event window until 6 months after deal expiry per 1 pp change in the interest rate, where we estimate β_{FA} from Eq. (9) for the bin indicators of η. The change in consumption is the cumulative change in consumption until 6 months after deal expiry. 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. Both regressions include year x Local Authority District and age fixed effects. Sample: reweighted consumption households with app-based spending. We winsorize the outcome variable, at the 0.1% and 99.9% level.

consumption. To see this point, we repeat the previous analysis of regression equation (9), with consumption over the event window as the outcome variable. We start with non-parametric evidence from a scatter plot, in Figure 10, 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 5 contains the regression evidence underlying the scatter plot. In the baseline specification of column (1), the outcome is the cumulative response of consumption in the event window 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%. Asset prices seem 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 9, 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 response of asset prices accounts for a meaningful share of the overall response of consumption to rate cuts. Using the same decomposition as we applied to borrowing, we find that around 50% of the response of consumption to rate cuts is due to the role of asset

prices. We also present similar analysis using the land supply elasticity instrument (Figure 10, panel (b), and Appendix Table B.14).[52,53]

We can summarize the importance of how rates affect consumption via asset prices, with an estimate of the marginal propensities to consume and borrow out of housing wealth (MPCH and MPBH, respectively). A large literature emphasizes how increases in house prices allow households to borrow and consume more—without linking the increase in house prices to monetary policy. This literature typically estimates a MPCH of around 5% (e.g. Campbell and Cocco, 2007; Guren et al., 2021). Our estimate is similarly large.

We estimate the MPC and MPB out of housing wealth as follows. Using the same notation as regression equation (9), we estimate a regression for consumption

$$\Delta_{t,t+h}\text{consumption}_{it} = \alpha_{j(i)t} + \text{MPCH} \times \text{house price}_{t,t-N_i} + \text{controls}_{it} + \varepsilon_{it}, \qquad (10)$$

where $\text{house price}_{t,t-N_i}$ is the change in the house price experienced over the past mortgage deal; and MPCH is the coefficient of interest in the regression. We estimate a similar regression with borrowing as the outcome variable. In both cases, we instrument for the house price change with $(r_t - r_{t-N_i}) \times \eta_{j(i)}$, the interaction of the rate change experienced by the household with the regional sensitivity to rate cuts. We use $(r_t - r_{t-N_i})$ and the initial home value as controls in this regression, and always include Local Authority District x year fixed effects. Table 6, column (1), reports the results, and finds an MPCH of 0.04, with similar estimates for borrowing, though the results for consumption are slightly imprecise.

6.4 Heterogeneity and the Importance of Collateral

So far, we have stressed that higher house prices, due to rate cuts, increase consumption—because households borrow against their home. Households could borrow more for two reasons: either because the higher house price loosens their collateral constraint (as in our model), or because higher house prices raise lifetime wealth. We now present evidence consistent with the importance of collateral.

One of the key determinants of the response is household leverage—that is, households' loan to value ratio. This heterogeneity suggests collateral is important, since collateral mat-

[52]We also studied how rate cuts affect consumption via asset prices using the duration instrument, though the results are imprecise.

[53]We also study the response of debt service—which does not appear to be as greatly affected by asset prices. Appendix Table B.15 presents the results for monthly payments, showing that areas with more responsive house prices see a greater absolute adjustment in monthly payments, though this effect is not as large as for borrowing. Appendix Table B.16 shows that this monthly payment effect is driven by the impact of house prices on individual mortgage rates. Columns 1-4 look at the level response, showing that for high LTV borrowers, individual mortgage rates are more responsive in the more house price responsive regions.

ters more for high loan-to-value households (see equation (3) from our model). In particular, we re-estimate equation (10) for various subgroups of households with: above or below 70% loan-to-value, measured 7 months before their deal expires; above or below the age of 35; and above or below median household income. Table 6, panels (a) and (b), report these results for consumption and for borrowing. The MPBH and MPCH are both much larger for indebted households, though the results are noisy for consumption (columns 2-3). Heterogeneity by income appears to be less important (columns 4-5), while being young is also an important determinant of MPCs out of housing wealth (columns 6-7). Amongst the young, high loan-to-value households again have higher MPCs and MPBs (columns 8-9).[54]

6.5 Active versus Passive Mechanisms

Overall, we have argued that monetary policy affects consumption by easing collateral constraints. However there is a separate but related mechanism that may also matter. If households have higher debt, then one component of the response of debt service to rate cuts—the product of initial debt and the change in rates—will "passively" respond more. The passive mechanism could be strong for the same households for whom collateral matters. Rates have been falling over our sample, meaning the households whose collateral responds to rate cuts will have higher debt.

We develop a strategy to differentiate the importance of collateral, the "active" mechanism; from the "passive" effect of higher debt. In particular, we separate the change in cash-on-hand around deal expiry into two terms. The first term, the passive change, is the product of the mortgage balance before refinancing, and the change in mortgage rates around refinancing.[55] This term captures exactly the passive mechanism that we have discussed. The second, residual component of cash-on-hand is due to "active" changes arising from collateral and related factors.

We then estimate the marginal propensity to consume, separately out of the active and passive components. That is, we estimate a regression

$$\Delta_{t,t+h}\text{consumption}_{it} = \alpha_{j(i)t} + \text{MPC}^{\text{active}}\ \Delta_{t,t+h}\text{active coh}_{it} + \text{MPC}^{\text{passive}}\ \Delta_{t,t+h}\text{passive coh}_{it} + \varepsilon_{it},$$

where "coh" stands for cash-on-hand, and we estimate the equation at a horizon of 6 months.

[54] As further evidence of the importance of collateral, Appendix Section C.8 carries out an analysis similar to Cloyne et al. (2019), concerning mortgage "notches", i.e. discrete jumps at certain loan-to-value thresholds. Consistent with the importance of collateral, households' borrowing responses are stronger when they cross notches.

[55] This is calculated using the change in aggregate mortgage rates over the deal, abstracting from endogenous household adjustments to the mortgage balance or amortization schedule at deal expiry.

In order to generate independent variation in both the active and passive components, we instrument for the active component with the interaction of the response of house prices to rates, $\eta_{j(i)}$, with the rate cut experienced by the household; without instrumenting for the passive component. This regression allows us to rule out the importance of the passive mechanism. In particular, if the MPC out of the passive component were small relative to the MPC out of the active component, then the passive mechanism is not relevant.

In practice, the passive mechanism indeed appears to be small, as Table 7 shows. Columns (1) and (2) show that both the active and passive components of cash-on-hand respond strongly to rate cuts. Column (3) replicates the baseline finding, from Figure 5, that the marginal propensity to consume out of cash-on-hand is around 0.5. Column (4), the main point of the table, separately estimates the MPC out of the active and passive components. The MPC out of the active component is large, around 0.7; while the MPC out of the passive component is near to zero. This result suggests that in practice, the passive mechanism is not as important as collateral in explaining the pass-through of rate cuts into consumption.

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.

Our first result is that the household debt channel is large. After a 1 pp reduction in mortgage rates, consumption increases by 3% for households with mortgage deals expiring, which aggregates to 0.7% 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.

Second, we show that the household debt channel is large because of the effect of rate cuts on asset prices. Borrowing accounts for around two thirds of the response of cash-on-hand after rate cuts. We link the borrowing response to asset prices, using plausibly exogenous variation in how households' home values respond to rate cuts, based on the regional housing supply elasticity and the duration of the housing stock. The response of asset prices to rate cuts accounts for over 50% of the borrowing and consumption response. Overall, monetary policy easing affects consumption in large part by raising asset prices, and encouraging household borrowing against the greater housing wealth.

Tables

Table 1: Effect of Rates on Cash on Hand and Consumption

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Panel A					
	Δ Cash on Hand				Δ Log Cash on Hand			
ΔMortgage Rate	-1818.2***	-1816.9***	-1773.1***	-2272.8***	-1.44***	-1.83***	-1.28***	-1.87***
	(104.5)	(114.6)	(139.6)	(214.0)	(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
			Panel B					
	Δ Consumption				Δ Log Consumption			
ΔMortgage Rate	-937.8***	-1364.1***	-834.7**	-1814.2***	-2.99***	-2.55**	-3.02**	-7.53***
	(277.4)	(388.2)	(320.4)	(584.8)	(1.03)	(0.998)	(1.22)	(2.07)
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
			Panel C					
	$\text{MPC}^{\text{cash-on-hand}}$							
Marginal Propensity to Consume	0.489***	0.377***	0.530**	1.148				
out of Cash-on-hand	(0.139)	(0.127)	(0.171)	(0.709)				
Observations	68734	68734	68734	17614				
Adjusted R^2	-0.336	-0.186	-0.404	-5.188				
K-Paap F-Stat	57.71	22.15	37.29	4.20				
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: Panel A reports the effect of changes in aggregate mortgage rates over the deal on the change in cash-on-hand over the event window around mortgage deal expiry, using the full mortgage sample. Δ Cash on Hand is defined as the change in mortgage principal minus 10 times the change in monthly payments. Δ Log Cash on Hand is defined as $\frac{\Delta\text{Cash on Hand}}{\text{Mortgage balance 7 months prior to deal expiry}} \times 100$; columns (5)–(8) are weighted by the mortgage balance 7 months prior to deal expiry. Panel B reports the effect of changes in aggregate mortgage rates over the deal on the change in consumption over the event window around mortgage deal expiry, using the matched mortgage–consumption sample. Δ Consumption is the cumulative change in consumption around deal expiry. Δ Log Consumption is the growth rate in average consumption around deal expiry, multiplied by 100 for a percentage interpretation. In Panel B, columns (1)–(4) are weighted by sampling weights; columns (5)–(8) are weighted by the product of the sampling weight and the mortgage balance 7 months prior to deal expiry. Panel C reports the marginal propensity to consume out of cash-on-hand estimated by 2SLS on the matched mortgage–consumption sample: the first stage regresses the change in cash-on-hand on the change in aggregate mortgage rates, and the second stage regresses the change in consumption on the fitted change in cash-on-hand. Fixed effects vary by column as indicated. Outcome variables in Panels A and B, and both outcome and regressor in Panel C, are winsorized at the 0.1% and 99.9% levels. Trigger-month clustered standard errors are in parentheses.

Table 2: Effect of Mortgage Rates—High Frequency Instrument for Rates

	Δ Cash on Hand				Δ Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1816.9***	-2320.4***	-1659.5***	-2073.9***	-1364.1***	-1313.7**	-1711.4**	-1339.9
	(114.6)	(193.4)	(171.0)	(310.9)	(388.2)	(607.0)	(738.3)	(1837.5)
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
K-Paap F-Stat		290.06		189.11		194.10		67.98
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. 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 change in consumption is the cumulative change in consumption around deal expiry. 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 change in consumption as the outcome variable. These consumption regressions are run on the matched mortgage-consumption sample and are weighted by sampling weights. Outcome variables are winsorized at the 0.1% and 99.9% levels. Trigger month clustered standard errors are shown in parentheses.

Table 3: Effect of Rates on Borrowing via Asset Prices

	Δ Principal				Δ Log Principal			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1671.2***	-1671.8***	-1802.4***	-2617.8***	-0.983***	-0.717***	-1.16***	-1.93***
	(243.9)	(197.2)	(290.8)	(379.2)	(0.170)	(0.115)	(0.179)	(0.257)
ΔMortgage Rate x 𝕀(Middle 80% η)	-1092.5***	-1083.9***	-1090.0***	-1550.2***	-0.888***	-0.623***	-0.986***	-1.59***
	(112.1)	(109.6)	(141.7)	(221.9)	(0.120)	(0.085)	(0.133)	(0.196)
ΔMortgage Rate x 𝕀(Top 10% η)	-531.3***	-523.9***	-498.3***	-498.1***	-0.558***	-0.294***	-0.638***	-0.831***
	(71.6)	(102.6)	(99.6)	(147.4)	(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 rate cuts on mortgage principal via asset prices, 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table 4: Effect of Rates on Borrowing via Asset Prices: Duration Variation

	Δ Principal					Δ Log Principal				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ΔMortgage Rate x 𝕀(F.hold)	-724.1***	-720.7***	-764.5***	-828.6***	-804.2***	-0.508***	-0.519***	-0.557***	-0.431***	-1.07***
	(82.1)	(85.1)	(83.6)	(148.7)	(221.5)	(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 rates on mortgage principal via asset prices, 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table 5: Effect of Rates on Consumption via Asset Prices

	Δ Consumption				Δ Log Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1946.2**	-2118.4**	-1526.9	-2800.4	-5.27***	-3.98**	-4.62**	-18.1***
	(761.8)	(859.7)	(921.3)	(2317.5)	(1.77)	(1.71)	(2.12)	(5.02)
ΔMortgage Rate x 𝕀(Middle 80% η)	-845.9***	-1012.4***	-864.2**	-1864.0***	-2.25**	-0.909	-2.68**	-6.29***
	(295.7)	(368.8)	(349.4)	(613.7)	(1.07)	(1.00)	(1.25)	(2.25)
ΔMortgage Rate x 𝕀(Top 10% η)	-64.9	-234.4	-399.1	-1056.8	-1.00	0.32	-1.45	-4.74
	(421.6)	(473.9)	(490.0)	(1137.9)	(1.90)	(1.93)	(2.09)	(3.43)
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	✓	✓	✓		✓	✓	✓	

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

Notes: The table shows the impact of rates on consumption via asset prices, 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 growth rate in average consumption around deal expiry 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table 6: Marginal Propensity to Borrow and Consume out of Housing Wealth

	All	$LTV \leq 70\%$	$LTV > 70\%$	Low Income	High Income	≤35 Years	> 35 Years	≤35 Years, $LTV \leq 70\%$	≤35 Years, $LTV > 70\%$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				Panel (a): Δ Principal					
ΔHome Value	0.042***	0.039***	0.075***	0.018**	0.051***	0.057***	0.040***	0.044**	0.066***
	(0.013)	(0.014)	(0.023)	(0.008)	(0.016)	(0.017)	(0.013)	(0.018)	(0.023)
Observations	6740340	4953735	1786405	3307122	3306675	1951329	4789011	981333	969945
Adjusted R^2	0.000	0.000	0.002	-0.000	-0.001	0.001	0.000	-0.002	0.002
K-Paap F Stat.	21.61	22.14	39.95	19.46	18.47	24.71	21.32	25.85	45.31
				Panel (b): Δ Consumption					
ΔHome Value	0.044	0.031	0.174	0.028	0.039	0.101**	0.024	0.061	0.343***
	(0.027)	(0.028)	(0.139)	(0.040)	(0.040)	(0.050)	(0.032)	(0.053)	(0.125)
Observations	68541	43841	24308	36606	31554	26292	41894	12356	13361
Adjusted R^2	-0.037	-0.054	-0.108	-0.059	-0.076	-0.099	-0.055	-0.170	-0.231
K-Paap F Stat.	17.56	17.01	24.09	18.13	14.73	19.53	15.68	16.62	27.59
Year x LAD FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Age FE	✓	✓	✓	✓	✓	✓	✓	✓	✓

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

Notes: The table reports IV estimates of borrowing and consumption responses to housing-wealth changes induced by aggregate mortgage-rate movements. In each column, the £ change in real home value over the expiring deal is instrumented with the interaction of the change in aggregate mortgage rates over the deal and regional house-price sensitivity, η. In all panels, the initial home value and the change in the mortgage rate are also included as regressors. Panel (a) uses the full mortgage sample and the dependent variable is the change in mortgage principal (Δ Principal) over the event window around deal expiry. Panel (b) uses the matched mortgage-consumption sample and the dependent variable is the cumulative change in consumption (Δ Consumption) over the event window around deal expiry. Regressions in Panel (b) are weighted by sampling weights. Columns (2)–(3) split by whether the LTV ratio (measured 7 months prior to deal expiry) is ≤ 70% or > 70%. Columns (4)–(5) split the sample by whether the mortgagor's most recently observed income is below/above the median in the year it is observed. Columns (6)–(7) split by whether the average mortgagor age at the property (measured 7 months prior to deal expiry) is ≤ 35 or > 35. Columns (8)–(9) apply the LTV split within the ≤ 35 group. All regressions include year × Local Authority District (LAD) and household age fixed effects. Outcome variables, the change in home value, and the initial home value, are winsorized at the 0.1% and 99.9% levels. Trigger-month clustered standard errors are in parentheses. The Kleibergen–Paap rk Wald F statistic is reported.

Table 7: Marginal Propensity to Consume Out of Active Cash on Hand

	Δ Active CoH	Δ Passive CoH	Δ Consumption	
	(1)	(2)	(3)	(4)
ΔMortgage Rate	-1259.1***	-658.9***		
	(238.0)	(18.6)		
ΔCash on Hand			0.454***	
			(0.133)	
ΔActive Cash on Hand				0.679**
				(0.280)
ΔPassive Cash on Hand				0.102
				(0.518)
Observations	68539	68539	68539	68539
Adjusted R^2	0.027	0.687	-0.332	-0.756
K-Paap F Stat.			65.79	14.18
Year x LAD FE	✓	✓	✓	✓
Age FE	✓	✓	✓	✓

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

Notes: The table shows the marginal propensity to consume out of 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. The *passive* change in cash-on-hand is defined as $-10 \times$ the passive change in monthly payments at refinance. The passive change in payments are the change in payments that occur due to the change in aggregate mortgage rates, abstracting from endogenous household adjustments at the time of refinance. The *active* change in cash-on-hand is defined as the change in cash-on-hand less the passive change in cash-on-hand. The change in consumption is the cumulative change in consumption around deal expiry. Column 1 shows the regression of the change in active cash-on-hand on the change in aggregate mortgage rates. In column 2 the outcome variable is instead the change in passive cash-on-hand. Column 3 shows the IV regression of the change in consumption on the change in cash-on-hand, where the change in cash-on-hand is instrumented with the change in aggregate mortgage rates. Column 4 shows the IV regression of the change in consumption on both active and passive cash-on-hand. The change in active cash-on-hand is instrumented by the interaction of the change in aggregate mortgage rates and the regional house price sensitivity, η. All regressions include year x Local Authority District (LAD), and household age fixed effects. All regressions are weighted by sampling weights. Outcome variables are winsorized at the 0.1% and 99.9% levels. Trigger month clustered standard errors are shown in parentheses.

References

Albuquerque, Daniel, Thomas Lazarowicz, and Jamie Lenney, "Monetary Transmission in a HANK Model with Housing and Rental Sectors," 2025.

Andaloussi, Mehdi Benatiya, Nina Biljanovska, Alessia De Stefani, and Rui Mano, "Feeling the Pinch? Tracing the Effects of Monetary Policy through Housing Markets," in "World Economic Outlook—Steady but Slow: Resilience amid Divergence," Washington, DC: International Monetary Fund, April 2024, chapter 2, pp. 43–64. Chapter 2 of the April 2024 WEO.

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.

Bahaj, Saleem, Angus Foulis, Gabor Pinter, and Paolo Surico, "Employment and the residential collateral channel of monetary policy," *Journal of Monetary Economics*, 2022, *131*, 26–44.

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.

Bosshardt, Joshua, Marco Di Maggio, Ali Kakhbod, and Amir Kermani, "The credit supply channel of monetary policy tightening and its distributional impacts," *Journal of Financial Economics*, 2024, *160*, 103914.

Bracke, Philippe and Silvana Tenreyro, "History dependence in the housing market," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 420–443.

_ , Matthew Everitt, Martina Fazio, and Alexandra Varadi, "When refinancing meets monetary tightening: heterogeneous impacts on spending and debt via mortgage modifications," *Bank of England Staff Working Paper*, 2024, (1,105).

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.

Caballero, Ricardo J and Alp Simsek, "Central banks, stock markets, and the real economy," *Annual Review of Financial Economics*, 2024, *16*.

Campbell, John Y. and João F. Cocco, "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics*, 2007, *54* (3), 591–621.

Caspi, Itamar, Nadav Eshel, and Nimrod Segev, "The Mortgage Cash-Flow Channel: How Rising Interest Rates Impact Household Consumption," *arXiv preprint arXiv:2410.02445*, 2024.

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.

Cumming, Fergus and Paul Hubert, "The Distribution of Households' Indebtedness and the Transmission of Monetary Policy," *Review of Economics and Statistics*, 2023, *105* (5), 1304–1313.

Davis, Steven J., John Haltiwanger, and Scott Schuh, *Job Creation and Destruction*, MIT Press, 1996.

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

Elias, Matthew, Christian Gillitzer, Greg Kaplan, Gianni La Cava, and Nalini V Prasad, "The Mortgage Debt Channel of Monetary Policy when Mortgages are Liquid," Technical Report, National Bureau of Economic Research 2025.

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.

Friedman, Milton, "The Lag in Effect of Monetary Policy," *Journal of Political Economy*, 1961, *69* (5), 447–466.

Garriga, Carlos, Finn E Kydland, and Roman Sustek, "Mortgages and monetary policy," *The Review of Financial Studies*, 2017, *30* (10), 3337–3375.

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," *MIT Sloan Research Paper*, 2018.

Groth, Charlotta and Tracy Wheeler, "The behaviour of the MPC: Gradualism, inaction and individual voting patterns," Technical Report, External MPC Unit Discussion Paper 2008.

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.

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.

Mian, Atif and Amir Sufi, "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis," *American Economic Review*, 2011, *101* (5), 2132–2156.

_ , Kamalesh Rao, and Amir Sufi, "Household Balance Sheets, Consumption, and the Economic Slump," *The Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

Ottonello, Pablo and Thomas Winberry, "Financial heterogeneity and the investment channel of monetary policy," *Econometrica*, 2020, *88* (6), 2473–2502.

Rajan, Aniruddha, Francesc Rodriguez-Tous, and Mauricio Salgado-Moreno, "Monetary Policy and Mortgage Fixation Lengths," Bank of England Working Paper, Bank of England 2025.

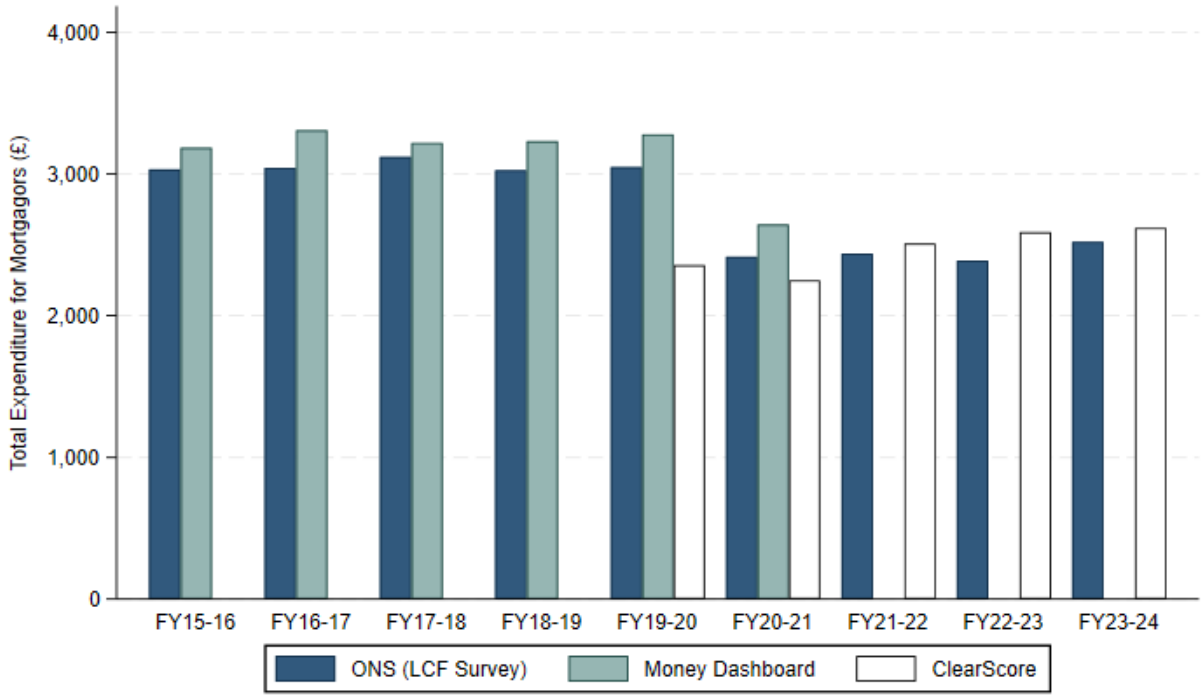
Romer, Christina D and David H Romer, "A new measure of monetary shocks: Derivation and implications," *American economic review*, 2004, *94* (4), 1055–1084.

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

Wong, Arlene, "Refinancing and the transmission of monetary policy to consumption," *Unpublished manuscript*, 2019, *20*.

A Appendix Figures

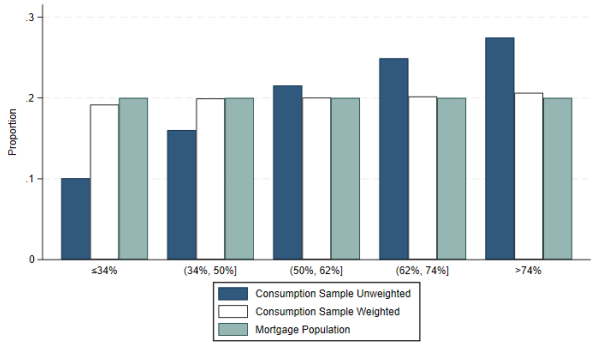
Figure A.1: Monthly Average Spending in Apps vs. Official Data



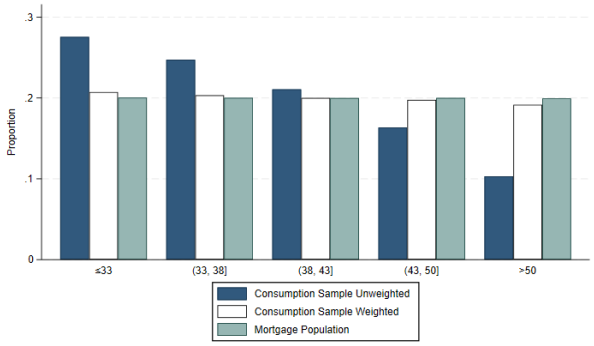
Notes: the figure shows the monthly average household total expenditure for mortgagors for the Office of National Statistics (Living Costs and Food Survey), Money Dashboard and ClearScore from 2015 to 2024. We use our total reweighted consumption measure to calculate the average values for the Money Dashboard and ClearScore apps. The sample is the same as Figure 1 (c) in the main text and the outcome is an average across the 13 month windows around deal expiry for both: (i) treated households whose deal expires within the window; and (ii) control households whose deal does not expire in the window. We winsorize the outcome variable, by calendar month and treatment status, at the 0.1% and 99.9% level. The units are pounds based in 2015. The Office of National Statistics measure of total expenditure includes mortgage interest payments that are excluded from the Money Dashboard and ClearScore measures. Total expenditure consumption averages are calculated within UK financial years, running from April to March.

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

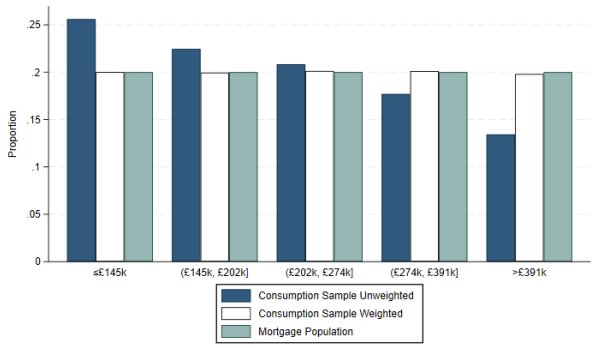
(a) Loan to Value



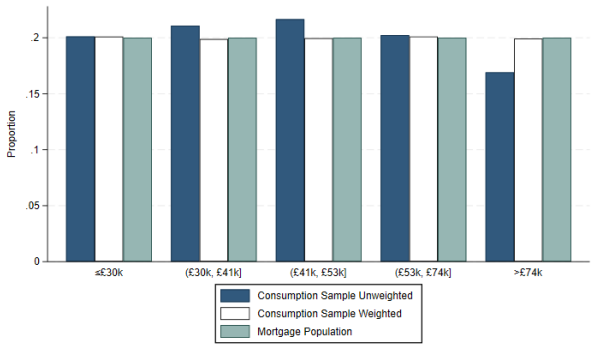
(b) Borrower Age



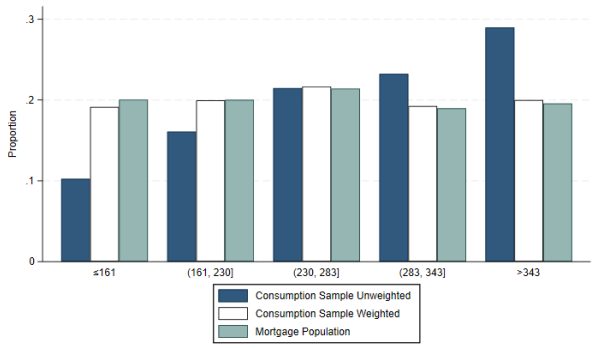
(c) Home Value



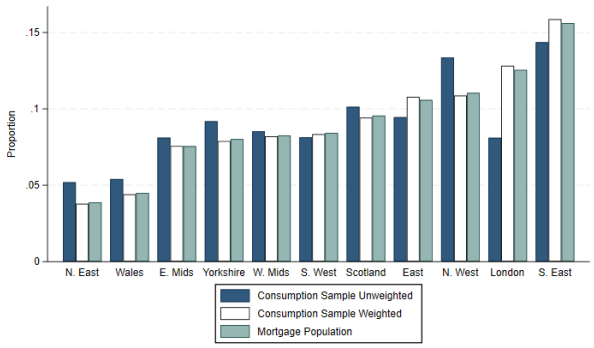
(d) Household Income



(e) Mortgage Term



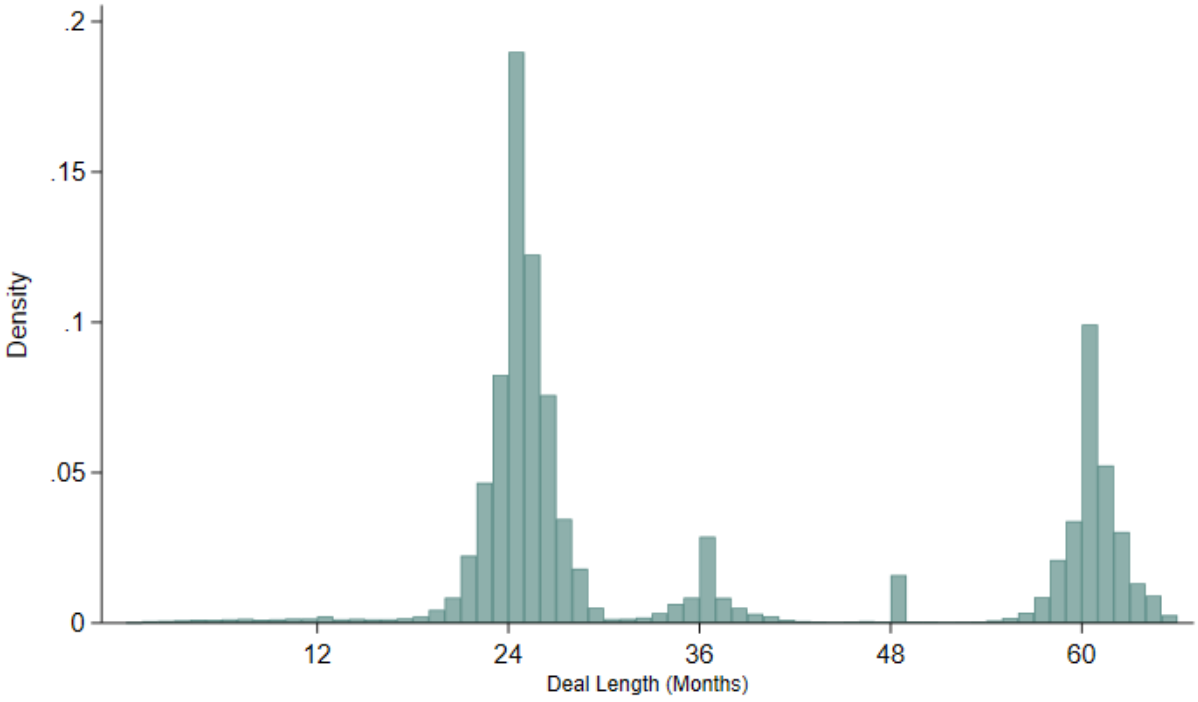
(f) Broad Region



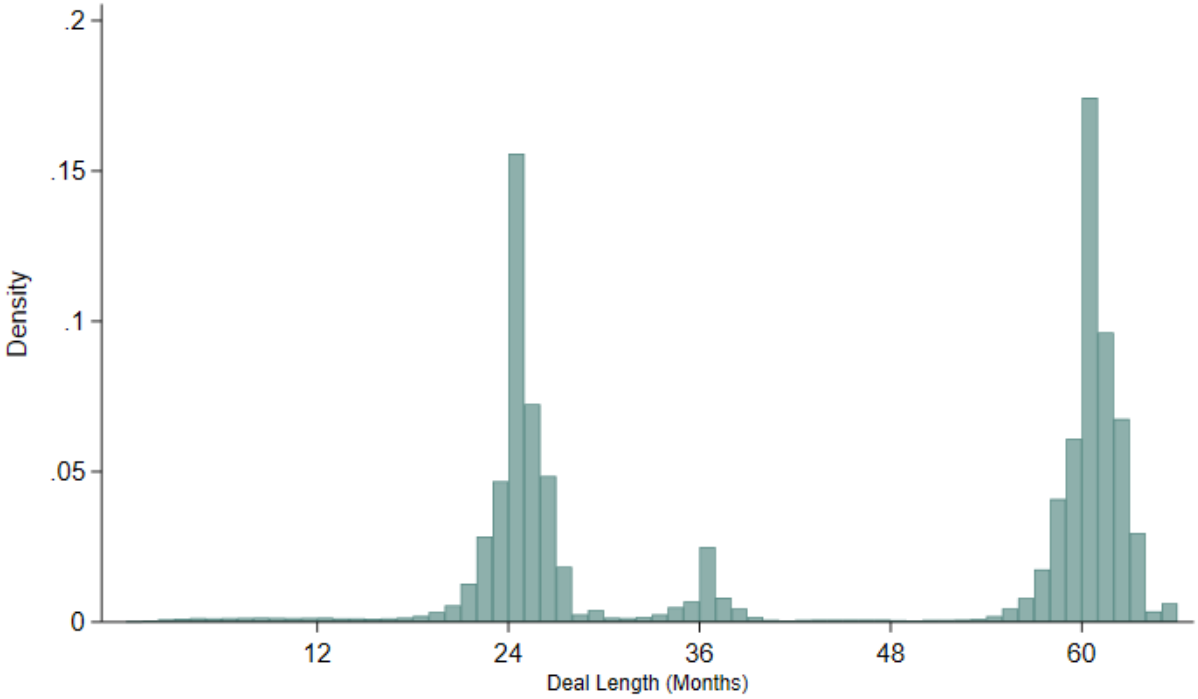
Notes: the figure compares distributions of loan-to-value (a), borrower age (b), home value in 2015 pounds (c), last observed household income in 2015 pounds (d), remaining mortgage term in months (e), and region (f) in the mortgage population, the subset with observed app-based consumption, and the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

Figure A.3: Histograms of New Fixed Rate Mortgage Deal Lengths

(a) 2015 H1

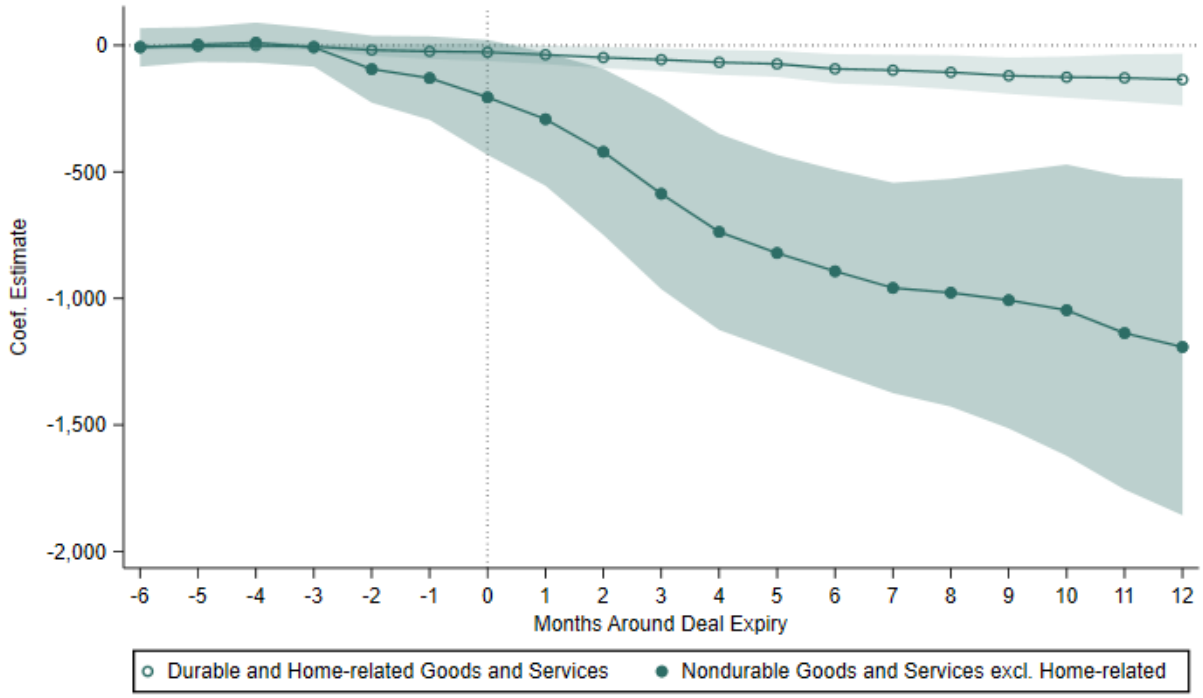


(b) 2019 H1



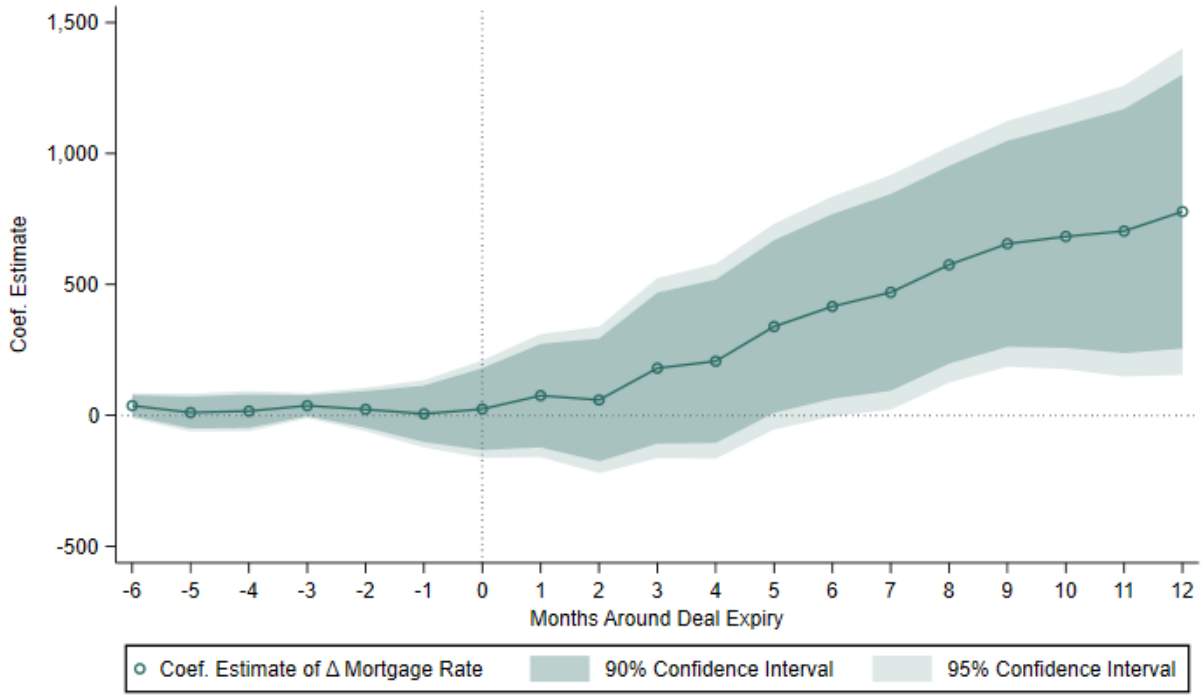
Notes: the figure shows histograms of new fixed rate mortgage deal lengths, in months, from the PSD 001 mortgage dataset. Panel a) includes mortgages originated in the first half of 2015. Panel b) includes mortgages originated in the first half of 2019. Both figures show mortgage deals of up to 66 months in length. In both time periods over 97% of new fixed rates mortgages in the PSD 001 dataset have a deal length of at most 66 months.

Figure A.4: Impulse Response of Consumption Components



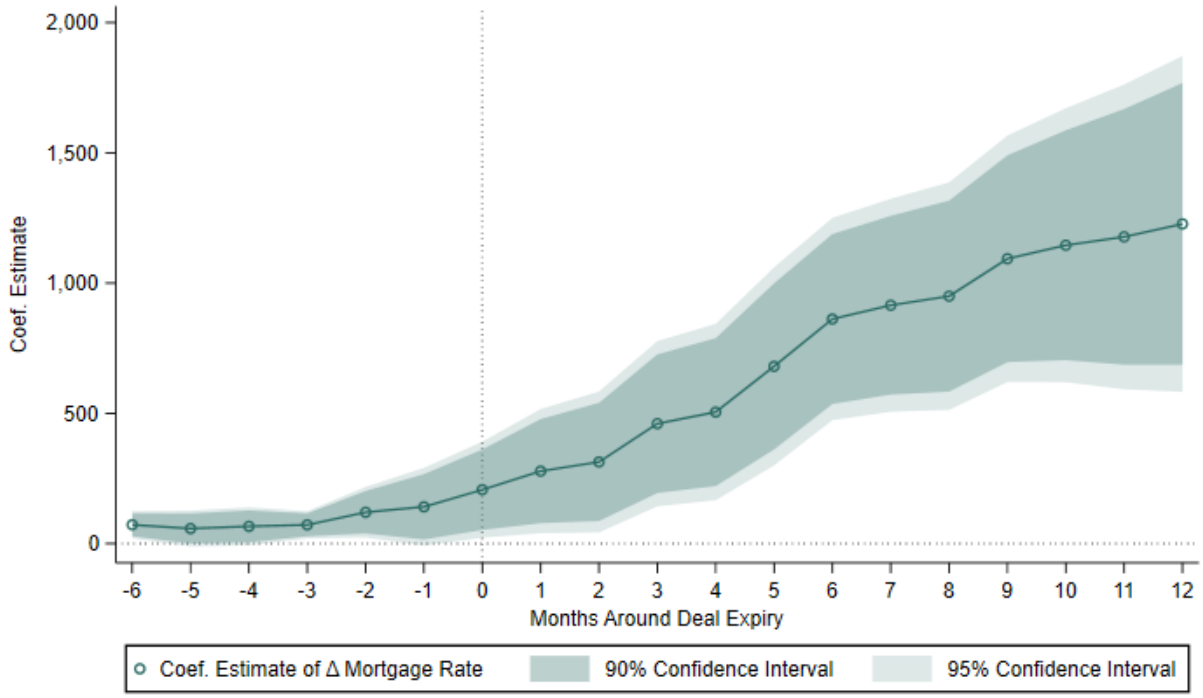
Notes: the figure plots cumulative impulse responses for durables and non-durables consumption from the baseline monthly regression. For horizons $h \in [-6, 12]$ relative to deal expiry ($t = 0$), the series show the cumulative change in the outcome (pounds, 2015 base) per 1 percentage-point change in the interest rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. Both regressions include year and age fixed effects. Shaded bands denote 90% confidence intervals for deal expiry month clustered standard errors. Sample: reweighted consumption for households with app-based spending; durables include home-related renovation expenses. Outcomes are winsorized, horizon by horizon, at the 0.1% and 99.9% level.

Figure A.5: 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 $t \in [-6, -4]$. The regression includes year and age fixed effects. Shaded bands denote 90% and 95% confidence intervals for deal expiry month clustered standard errors. We winsorize the outcome variable, at the relevant horizon, at the 0.1% and 99.9% level. Sample: reweighted mortgagor households with income inferred from linked spending apps.

Figure A.6: 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 $t \in [-6, -4]$. The regression includes year and age fixed effects. Shaded bands denote 90% and 95% confidence intervals for deal expiry month clustered standard errors. Sample: reweighted mortgagor households with income inferred from linked spending apps. Outcomes are winsorized, horizon by horizon, at the 0.1% and 99.9% level.

B Appendix Tables

Table B.1: Summary Statistics: Mortgage Population and Consumption Sample

	Mortgage Population		Weighted Consumption Sample		Unweighted Consumption Sample	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Mortgage Variables						
Mortgage Debt (£)	147,656.79	115,319.32	147,399.82	109,778.88	145,793.02	99,840.27
Home Value (£)	295,419.69	234,737.30	285,130.25	203,292.39	250,511.65	172,272.91
Mortgage Payment (£)	716.95	478.35	721.17	469.78	673.15	420.23
Mortgage Term (months)	252.48	99.36	256.21	97.40	286.04	90.71
Interest Rate (%)	2.34	0.82	2.32	0.81	2.40	0.84
Loan-to-Value Ratio (%)	53.80	21.20	54.51	20.87	60.64	18.74
Loan-to-Income Ratio	2.68	1.08	2.66	1.03	2.78	0.96
Income (£)	58,090.97	47,960.09	56,802.44	39,768.08	53,752.36	35,038.82
Age (Years)	41.53	9.60	41.07	9.18	38.49	8.31
London Dummy (%)	12.55	33.13	12.82	33.44	8.11	27.29
Consumption Variables						
Total Consumption (£)			2,543.39	2,553.44	2,433.75	2,378.23
Durable and Home-related (£)			131.43	401.26	126.37	377.65
Nondurables and Services excl. Home-related (£)			2,247.97	2,249.01	2,154.01	2,097.25
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 with expiring deals from January 2016 to September 2023. In the last 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 middle 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. Regarding consumption variables, total consumption is the sum of durable and home-related spending, non-durable and services excluding home-related, and cash expenditures. The observations are at the level of an expiring deal: we measure the characteristics of a household 7 months before its deal expires, except for income, which is the most recent income observed prior to this date. Outcome variables are winsorized at the 0.1% and 99.9% levels. The units are 2015 pounds.

Table B.2: 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)	-1777.7***	-1163.2***	-614.6***	-1186.0**
	(126.4)	(129.4)	(27.6)	(497.4)
ΔMortgage Rate x 𝕀(2022-)	-1837.7***	-1069.6***	-768.1***	-876.2***
	(138.6)	(169.8)	(60.2)	(330.2)
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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.3: Marginal Propensity to Consume Out of Cash on Hand: Two Sample

	Marginal Propensity to Consume			
	(1)	(2)	(3)	(4)
ΔCash on Hand	0.516***	0.751***	0.471	0.798*
	(0.185)	(0.212)	(0.306)	(0.455)
Observations	68734	68734	68734	17614
Second Stage Adjusted R^2	0.003	0.011	0.003	0.040
First Stage F-Stat	302.56	251.34	161.41	112.83
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 marginal propensity to consume 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 change in consumption is the cumulative change in consumption in the event window around 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 change in consumption on the fitted change in cash-on-hand from the first stage. The second stage regression is run in the matched reweighted mortgage-consumption sample. 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. Standard errors, shown in parentheses, are bootstrapped, with 1,000 replications, sampling by trigger month clusters. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.4: Effect of High Frequency Monetary Policy Shocks on Mortgage Rates

	Δ Mortgage Rate			
	(1)	(2)	(3)	(4)
∑HF Monetary Policy Shocks	1.97***	1.32***	1.94***	1.15***
	(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	290.06	189.11	194.10	67.98
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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.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	-1918.8***	-3618.8***	-1574.3***	-1580.2**	-1.45***	-3.29***	-1.10***	-1.12*
	(252.6)	(768.9)	(257.8)	(770.9)	(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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.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	-1818.2***	-1100.1***	-718.1***	-937.8***	-1749.7***	-1034.0***	-715.7***	-953.6***
	(104.5)	(119.4)	(38.2)	(277.4)	(115.2)	(131.3)	(40.6)	(290.7)
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. The change in consumption is the cumulative change in consumption in the event window around deal expiry. 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 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.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	-1818.2***				
	(104.5)				
ΔIndividual Mortgage Rate		-1826.0***	-1664.5***	-1645.5***	-1494.2***
		(58.7)	(94.2)	(92.7)	(101.5)
Observations	6796488	6796488	6796488	6796488	6740237
Adjusted R^2	0.011	0.020	0.010	0.010	0.009
K-Paap F-Stat			444.98	238.66	280.48
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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.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	-757.5***	-903.5***	-740.7**	-2105.8***	-2.67***	-1.91**	-2.82**	-7.97***
	(258.5)	(326.2)	(293.9)	(471.4)	(1.01)	(0.914)	(1.17)	(1.81)
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	✓	✓	✓		✓	✓	✓	

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

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. The results in this table do not include the sampling weights. 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 growth rate in average consumption around deal expiry 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. Outcomes are winsorized at the 0.1% and 99.9% level.

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

	Δ Borrowing	−ΔDebt Service	Δ Borrowing	−ΔDebt Service	Δ Borrowing	−ΔDebt Service	Δ Borrowing	−ΔDebt Service
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-1100.1***	-718.1***	-1151.1***	-665.8***	-1065.6***	-707.5***	-1567.0***	-705.8***
	(119.4)	(38.2)	(112.0)	(27.0)	(157.0)	(49.5)	(227.2)	(45.4)
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 borrowing and cumulative change in monthly debt service. In column 1 the dependent variable is the change in borrowing, in the event window around the end of the mortgage deal. In column 2 the dependent variable is −10 times the change in monthly debt service, 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. Outcomes are winsorized at the 0.1% and 99.9% level.

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

	Δ Mortgage Term			
	(1)	(2)	(3)	(4)
ΔMortgage Rate	2.15***	3.07***	1.64***	1.80***
	(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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.11: Borrowing and Asset Prices—Housing Supply Instrument

	Δ Principal				Δ Log Principal			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1818.8***	-1816.9***	-2063.6***	-2833.8***	-1.01***	-0.746***	-1.27***	-1.99***
	(251.8)	(211.0)	(276.4)	(362.9)	(0.189)	(0.134)	(0.184)	(0.251)
ΔMortgage Rate x 𝕀(Middle 80% η)	-1069.2***	-1059.6***	-1051.6***	-1508.0***	-0.889***	-0.624***	-0.967***	-1.57***
	(112.9)	(108.3)	(146.4)	(221.5)	(0.118)	(0.084)	(0.135)	(0.194)
ΔMortgage Rate x 𝕀(Top 10% η)	-543.3***	-537.0***	-522.7***	-575.0***	-0.467***	-0.202**	-0.591***	-0.858***
	(68.4)	(112.1)	(85.0)	(168.8)	(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 rates on mortgage principal via asset prices, 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.12: Effect of Asset Price Instruments on House Prices

	House Price Growth (%)			
	LAD House Price Index	LAD x FH/LH House Price Index		
	(1)	(2)	(3)	(4)
ΔMortgage Rate	-2.89***	-1.21*		
	(0.447)	(0.582)		
ΔMortgage Rate x η	0.505***	0.477***		
	(0.146)	(0.140)		
ΔMortgage Rate x $\mathbb{I}$(F.hold)		-1.73***	-1.46***	-3.37***
		(0.474)	(0.380)	(0.445)
Observations	6796438	4495365	4495365	4270715
Adjusted R^2	0.719	0.768	0.770	0.972
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 asset price 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 ΔMortage Rate x $\mathbb{I}$(Freehold) to the regression of column 1, 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.13: Marginal Propensity to Borrow out of Housing Wealth: Alt. IVs

	Δ Principal				
	IV: η	IV: Instr. η	IV: Duration		
	(1)	(2)	(3)	(4)	(5)
ΔHome Value	0.042***	0.046***	0.047***	0.065***	0.039***
	(0.013)	(0.014)	(0.009)	(0.013)	(0.011)
ΔMortgage Rate	-314.2***	-248.3**	-401.5***		
	(92.7)	(98.0)	(105.8)		
Initial Home Value	-0.002**	-0.002**	-0.002**	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	6740340	6740340	4495207	4495207	4270557
Adjusted R^2	0.000	-0.000	0.001	-0.001	0.001
K-Paap F-Stat	21.61	22.60	24.46	38.82	350.52
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 marginal propensity to borrow out of housing wealth for alternative instruments. 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 home value over the deal, the change in the aggregate mortgage rate over the deal, and the home value at the start of the deal. In column 1 the change in home value 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 Local Authority District (LAD) and household age fixed effects. Column 2 instead uses the interaction of the change in aggregate mortgage rates with the land-supply elasticity instrumented η. Columns 3-5 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. The change in home value here is calculated using the initial home value and separate real house price indices for freehold and leasehold properties within each LAD. Column 3 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 4 replaces the ΔMortgage Rate x η control with LAD x ΔMortgage Rate fixed effects. Column 5 includes the following fixed effects: (i) deal start month x deal expiry month x LAD; (ii) expiry month x 𝕀(Freehold); and (iii) household age. Trigger month clustered standard errors are shown in parentheses. Outcomes, the change in home value, and the initial home value, are winsorized at the 0.1% and 99.9% level.

Table B.14: Consumption and Asset Prices—Housing Supply Instrument

	Δ Consumption				Δ Log Consumption			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1652.6**	-1797.9**	-901.0	-149.4	-4.01*	-2.78	-3.34	-9.19*
	(779.7)	(879.9)	(955.0)	(2040.5)	(2.12)	(1.98)	(2.89)	(5.20)
ΔMortgage Rate x 𝕀(Middle 80% η)	-852.2***	-1023.1***	-902.9**	-2208.3***	-2.51**	-1.15	-2.93**	-8.45***
	(301.6)	(359.9)	(345.8)	(626.2)	(1.02)	(0.992)	(1.15)	(2.34)
ΔMortgage Rate x 𝕀(Top 10% η)	-293.8	-460.4	-635.5	-834.0	-0.364	0.910	-0.999	-0.916
	(510.4)	(592.5)	(656.4)	(1235.8)	(2.10)	(2.01)	(2.48)	(3.54)
Observations	68541	68541	67078	16942	68541	68541	67078	16942
Adjusted R^2	0.021	0.029	0.030	0.115	0.036	0.054	0.054	0.117
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 rates on consumption via asset prices, in the event window around the end of the mortgage deal. Results are shown for the matched mortgage-consumption sample. 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 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 growth rate in average consumption around deal expiry 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.15: Effect of Rates on Monthly Payments via Asset Prices

	−ΔPayments				−ΔLog Payments			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate x 𝕀(Bottom 10% η)	-1016.3***	-954.5***	-996.3***	-951.3***	-13.1***	-13.4***	-12.3***	-13.1***
	(51.0)	(30.8)	(65.3)	(66.2)	(0.889)	(0.434)	(1.09)	(1.02)
ΔMortgage Rate x 𝕀(Middle 80% η)	-711.9***	-654.3***	-707.6***	-695.4***	-12.4***	-12.8***	-11.9***	-12.5***
	(37.8)	(26.0)	(49.0)	(44.5)	(0.800)	(0.427)	(1.02)	(0.907)
ΔMortgage Rate x 𝕀(Top 10% η)	-480.0***	-425.1***	-485.0***	-515.5***	-10.6***	-11.1***	-10.6***	-11.2***
	(24.6)	(26.1)	(32.6)	(28.5)	(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 rates on cumulative monthly payments via asset prices, 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. Outcomes are winsorized at the 0.1% and 99.9% level.

Table B.16: Effect of Aggregate Rates on Individual Rates, via Asset Prices

	Δ Individual Mortgage Rate				Δ Log Individual Mortgage Rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMort. Rate x 𝕀(Bottom 10% η)	1.25***	1.01***	1.25***	1.22***	39.2***	35.1***	39.5***	41.5***
	(0.068)	(0.098)	(0.065)	(0.070)	(2.90)	(4.59)	(2.77)	(2.84)
ΔMort. Rate x 𝕀(Middle 80% η)	1.28***	1.05***	1.29***	1.27***	37.9***	33.9***	38.7***	40.8***
	(0.072)	(0.096)	(0.066)	(0.073)	(2.94)	(4.59)	(2.81)	(2.71)
ΔMort. Rate x 𝕀(Top 10% η)	1.33***	1.09***	1.35***	1.39***	35.4***	31.6***	36.9***	40.0***
	(0.083)	(0.097)	(0.077)	(0.088)	(3.21)	(4.62)	(3.06)	(3.16)
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 aggregate rates on the change in individual mortgage interest rates, via asset prices. 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. Outcomes are winsorized at the 0.1% and 99.9% level.

C Data Appendix

This section contains the data appendix. Section C.1 discusses the treatment of mortgage data. Section C.2 discusses the treatment of housing data. Section C.3 contains the treatment of the transaction level consumption data. Section C.4 contains the merge between the mortgage and housing data. Section C.5 explains how we reweight the consumption sample to make it representative of the population of mortgagors. Section C.6 discusses our treatment of home movers. Section C.7 explains how we estimate the response of house prices to rate cuts. Section C.8 contains our analysis of interest rate cuts and collateral constraints.

C.1 Treatment of Mortgage Data

This section describes the mortgage data used and how we treat it.

C.1.1 Product Sales Data

The Product Sales Data (PSD) is administrative mortgage data collected in the UK by the Financial Conduct Authority (FCA), and by the Financial Services Authority (FSA), prior to the creation of the FCA. The PSD mortgage data is split into two datasets, which we discuss in turn.

PSD 001: Sales Data PSD 001 covers data on the **flow** of new UK mortgages from April 2005 onward. The data provided under PSD 001 can be broken down into three broad regimes, which we discuss in turn.

PSD 001 Regime 1: 2005Q2-2014Q4 In this period PSD 001 includes data on all mortgages for home purchase, *external* remortgages (where the borrower switches mortgage provider), and *internal* remortgages (refinance with the same lender) where there is a change in loan terms, such as an increase in principal. For each of these observations, the PSD reports a number of detailed variables such as the value of the mortgage, the term of the mortgage, and financial ratios such as Loan-to-Value (LTV) and Loan-to-income (LTI) ratios. PSD 001 also records the interest rate of the mortgage and the date the incentive interest rate ends. However, reporting of these two variables was not mandatory in this period, with interest not reported for 35% of observations in Great Britain, and the end of the incentive interest rate not reported for 66% of observations in Great Britain.[56]

[56]The PSD data also covers mortgage activity in Northern Ireland, but we drop this from our analysis and focus on England, Wales, and Scotland.

Crucially, in this period, PSD 001 does not report data on *internal* remortgages (refinance with the same lender) where the loan terms are not modified. This includes cases where the principal either stays the same, or is reduced. Thus, using this dataset, we are only able to observe repeated mortgages for households which switch lenders or choose to change loan terms, such as increasing principal.

In this period the date of birth and income were only reported for the main borrower in the household.

PSD 001 Regime 2: 2015Q1-2021Q2 A number of changes were made to the PSD 001 data in this period. First, it became mandatory for lenders to report information on mortgage interest rates and the date incentive interest rates end. Second, lenders were required to report data on both the main and secondary borrowers in a household, including income and the date of birth for both. Third, lenders now had to report detailed data on the affordability of the mortgage for the borrowers, such as their committed monthly expenditures. In this period it remained the case that lenders were not required to record internal remortgages without a change in loan terms.

PSD 001 Regime 3: 2021Q2- The main change in the third regime is that lenders now had to record internal refinances without a change in loan terms. These make up a very large share of mortgage activity. For example, in 2023, PSD 001 data for Great Britain was broken down as follows:

- 60% were internal remortgages with no modification of loan terms.
- 13% were remortgages of the type identified previously in PSD 001 (external remortgages or internal remortgages with a modification of loan terms).
- 12% for first-time buyers.
- 11% were for home-movers.
- 4% were spread across other categories.

PSD 001: Selection From PSD 001 we exclude lifetime mortgages and shared appreciation mortgages. We also exclude second-charge and buy-to-let mortgages (with these latter two only reported in the dataset from 2021Q2 onwards).

PSD 001: Forming Household Identification Key PSD 001 reports the full postcode of the property which secures the mortgages. There are around 1.8 million UK postcodes, and a typical one comprises of 15 addresses. We use these postcodes along with the full dates of birth of borrowers to identify the same individual over time and across the two PSD datasets. As discussed, whilst lenders must report the date of birth of both the main and secondary borrower from 2015 onward, prior to this, PSD 001 only reports the date of birth for the main borrower. This could prevent us from identifying the same household over time if the selected main borrower changes in different PSD 001 observations prior to 2015. For example, suppose there is a property in postcode AB1 2CD, and the mortgage on the property has two borrowers, with dates of birth d_1 and d_2 where, without loss of generality, $d_1 < d_2$. Then we may observe the following sequence of mortgage activity in PSD 001:

Date	Postcode	Main Borr. DOB	Second Borr. DOB	DOB 1	DOB 2
25th January 2011	AB1 2CD	d_1		d_1	
25th January 2013	AB1 2CD	d_2		d_2	
25th January 2015	AB1 2CD	d_1	d_2	d_1	d_2
25th January 2017	AB1 2CD	d_2	d_1	d_1	d_2

If we identify the borrower based just on the postcode and date of birth of the main borrower, we would correctly identify that the January 2011 and January 2015 observations were mortgaging activity by the same household. We would also correctly link the remortgaging activity in 2013 and 2017. However, we would not link these two groups of observations together. To overcome this issue, in cases like this where we see multiple borrower dates of birth post 2015, we create a consolidated *household* identification key, which groups the household together if the postcode and either of the dates of birth match. To avoid the order of the dates of birth affecting the grouping, where there are multiple reported dates of birth, as in the third and fourth observations, we sort the dates of birth so that they appear in the same order in each row (as shown in the final two columns).

PSD 007: Performance Data PSD 007 covers data on the **stock** of all existing UK mortgages from June 2015 onward. PSD 007 data is reported twice per year, and provides data on all outstanding mortgages at the end of June and December each year, regardless of whether the mortgage was refinanced or not. PSD 007 reports detailed information on the mortgage, including the current outstanding balance, the current expected monthly payment, the current interest rate, the remaining term, and the date the current incentive interest rate ends (if applicable). Reporting of these variables is mandatory in PSD 007. PSD 007 also reports detailed information on the performance of the mortgage, for example, whether it is in arrears. PSD 007 captures information on the date of birth of the main borrower, though

it does not record changes to their income or the home value through time. We omit the small number of second charge mortgages reported in PSD 007 (these are reported from 2017 onward).

C.1.2 Merging Housing Data with Mortgage Data

Before building the borrower panel, we merge the combined Land Registry housing transaction data (described in Section C.2) with the mortgage data. This allows us to bring information on the housing *tenure* into the mortgage dataset, specifically, recording whether the property is owned *freehold* or *leasehold*.

We merge the Land Registry housing transaction data with the PSD 001 data on the flow of new mortgages using a similar procedure to Bracke and Tenreyro (2021). To carry out the merge we restrict the PSD sample to England and Wales, since only these areas are available in the Land Registry. We also exclude all remortgages from PSD 001, focusing on mortgages for property purchase. Both datasets record the full property postcode and the property value. The Land Registry records the date of the transaction, whilst PSD 001 records the date of mortgage origination.

We match the datasets using these pieces of information. Since the date of the mortgage completion is often different from the recorded transaction date of the house, we allow merges where the dates in each dataset are up to 31 days apart. We also include a small number of merges where the property values do not line up exactly across the two datasets. Overall, we merge 78% of PSD mortgage originations for property purchase in England and Wales to transactions in the Land Registry.

C.1.3 Building Household Panel

Our goal is to observe how households respond to the expiration of their mortgage deal, and in particular, how this was influenced by the change in aggregate mortgage rates over their deal. We wish to observe this behavior not just for external remortgages or internal remortgages with a change in loan terms, but also for internal mortgages without a change in loan terms. More generally, we also wish to observe the response of households *even if they do not remortgage*. Given this, our sample starts when PSD 007 begins, in June 2015. We use repeated household observations from PSD 007 to build a household panel over time. We augment this with data on the household from PSD 001, which provides additional important information, such as the house value and income of the borrowers, at the time these were observed in PSD 001.

To build the household panel, we begin with repeated observations from PSD 007, running

from June 2015 to June 2024. We drop mortgages that were originated prior to the beginning of PSD 001, in April 2005. We then merge in the generated household identification key, discussed above. We drop observations that are not matched to this identification key: we are unable to match such observations to any mortgage activity in PSD 001, so cannot, for example identify their home value or income at any point in time. We then merge PSD 001 information into this PSD 007 panel, based on the household identification key and the date the account opened. This then introduces key PSD 001 values, such as home value, to the panel for the date the account was opened.

We remove observations in the PSD 007 panel which are duplicates in the household identification key and the snapshot date. Almost all of these duplicates occur due to households having a closed mortgage account listed alongside their current mortgage, for example, if they switch their mortgage from one lender to another. In these cases we retain the information that the account closed in the subsequent snapshot, and then remove the observations for the closed account.

We then append the PSD 001 data to produce a combined panel of PSD 001 and PSD 007 data. We estimate the house value at all dates throughout this panel using the property value from the PSD 001 dates, and the change in the regional house price index for the Local Authority District the property is located in.

C.1.4 Identifying Remortgages

We use the borrower panel to identify the dates at which remortgages occur. Recall that, prior to 2021Q2, internal remortgages without a change in loan terms are not observable in PSD 001. To overcome this challenge we use successive PSD 007 snapshots as our basis for identifying all remortgages across our full sample period from 2015-2024. We split the identification of remortgages into two groups: (A) remortgages within the same account; and (B) a change in the mortgage account.

Case A: Remortgages Within The Same Account We first identify remortgages that occur within the same mortgage account. We identify the mortgage account based on the household identifier and the mortgage account open date. This type of remortgage corresponds to internal remortgages without a change in loan terms. Whilst PSD 007 does not explicitly record the date of this type of remortgage, we are able to infer such cases from a change in the date the incentivised mortgage rate ends between successive observations, six months apart, on the same mortgage account in the PSD 007 stock data. This covers cases (i) where the household was not on an incentivised mortgage rate at the previous PSD 007 snapshot (for example, if they did not remortgage immediately when their prior deal

expired); and (ii) cases where they were on an incentivised rate previously, and the new incentivised rate ends at a later date. This method is effective as it is mandatory to report the date incentivised mortgage rates end within PSD 007, and almost all remortgages in our sample period have an incentivised rate for a period. Our final regression sample will condition on mortgages with fixed rate deals whose incentivised rate is expiring.

For remortgages of this type, we know that the refinancing date occurs within the six-month window between successive observations in PSD 007. For these cases, we estimate the refinance date as follows:

1. In around two-thirds of cases, the expiry date of the new deal is in the same half of the year as the period between the two observed PSD 007 snapshots. For example, suppose we observe in the December 2015 snapshot that a household's new deal ends on 19th November 2017. Then, given deal lengths cluster around a whole number of years, a reasonable candidate refinance date is the day and month the new deal ends, and the year of the current snapshot. In this example, the candidate date is 19th November 2015. In this case, we take the candidate date as the refinance date, unless this date is in the 3 months (specifically 90 days) after the old deal expired. When the candidate date is in the 3 month window, we take the refinance date as the date the old deal expired (i.e. the household refinances exactly on time). This implies a new deal length of a whole number of years plus up to 3 months: such deal lengths are still common (Appendix Figure A.3). Outside of this 90 day window, we infer that the refinance occurred at the candidate date, for example, refinancing more than 90 days late.

2. In the second most common case the candidate date above does not fall within the snapshot window, but the old deal ended before the current snapshot. Modifying the example above, the household's old deal expired on 19th November 2015, and in December 2015, we observe that the new deal expires on 19th January 2018. Here the candidate date (19th January 2016) comes after the refinance has occurred, so cannot be the refinance date. In this case, we estimate that the refinance occurred at the expiration of the old deal, here 19th November 2015. This implies a deal length of a whole number of years plus a few months, which is common. In this example, the new deal length is inferred to be 2 years and 2 months.

3. In the least common case the candidate date also falls outwith the snapshot window, but the old deal had either already expired, or ends after the current snapshot. For example, suppose the original deal ended on 9th January 2016, and we observe in December 2015 that the new deal expires on 19th January 2018. Here, neither the old refinance date or the candidate date can be the refinance date as they both fall outside

the snapshot window. In this case, we modify the candidate date, based on the date the new deal ends, pushing it back by the minimum number of months required, up to a maximum of 6, to produce a refinance date within the snapshot window. We use the minimum number of months as this produces a new deal length closest to a whole number of years, which is more likely. In the example here, the raw candidate date is 19th January 2016, and the estimated refinance date is then 19th December 2015 (pulling it back 1 month), implying a deal length of 2 years and 1 month.

From 2021Q2 onward we observe internal remortgages in PSD 001 and can use the dates of these as a test of the accuracy of our method for estimating refinance dates. These tests suggest our method works well, with the median difference between the actual refinance date and our estimated refinance date being 0 days. From 2021Q2 onward we update our estimated refinance date with the exact date from the internal remortgage, where this is observed.

Case B: Change in Mortgage Account Our second method identifies remortgages which occur when the household changes their mortgage account. This can occur either where the household refinances with a new lender, or changes the mortgage terms with their existing lender, such as increasing the principal. In this case, we observe one mortgage account finishing for a household and a new one opening. In this case, the refinance date is the date the new account was opened, and this variable is observed in the PSD 007 dataset.

C.1.5 Monthly Household Panel

Mortgage Start and End Months The next step is to create a monthly panel of household mortgage information. To do this, we first identify the start and end months of mortgage deals. The end month is simply taken from the date the incentivised rate ends. The start month is more complex to construct. PSD 007 lists the date the mortgage account was first opened, but this need not be the start date of the most recent deal, if that deal began with an internal remortgage. We thus identify deal start dates as either (i) any observation in PSD 001, which, by construction, is for a new mortgage account; or (ii) a refinance identified between successive snapshots of PSD 007, as described in Section C.1.4. We restrict our mortgage sample to the mortgages with known deal start and end dates, as these are both required for our key empirical strategy.

Deal Length Using these start and end months, we can calculate the deal length. The length of fixed rate mortgage deals in the UK clusters around 2,3 and 5 years, with a smaller

number of observations at 1 and 7 years. We assign individual deal lengths in months into these broad deal length buckets as follows:

- 1 year deals: deals fixed for up to 17 months
- 2 year deals: deals fixed for 18 to 30 months
- 3 year deals: deals fixed for 31 to 53 months
- 5 year deals: deals fixed for 54 to 66 months
- 7 year deals: deal fixed for at least 67 months

Creating Monthly Panel With the deal start dates determined, and all relevant information extracted, we then drop the PSD 001 observations and focus on the PSD 007 snapshot data. We then expand the dataset to a monthly panel, creating new observations for the months between snapshot dates. Many variables are straightforward to fill in for the created dates, either because they're fixed within the mortgage account (such as property postcode), or because they're fixed within the mortgage deal (such as the initial length of the deal). More care must be taken for mortgage variables than can change over time. Here we consider three cases:

1. Case 1: a refinance occurs between two PSD 007 snapshots. For example, a refinance occurs in November 2015, between the June 2015 and December 2015 snapshot observations. In this case, the earlier snapshot contains details on the initial mortgage deal, and the later snapshot contains details on the new deal. Accordingly, we fill the mortgage variables such as interest rate and monthly payment forward from the prior snapshot until the month before the refinance, and fill the mortgage variables backwards from the subsequent snapshot to the refinance month. Using these values for the interest rate and monthly payment, and the balance observed at the snapshot dates, we then calculate the evolution of the mortgage balance between the snapshot dates.

2. Case 2: a deal expires between two PSD 007 snapshots, but there is no refinance. In this case we proceed in the analogous manner to Case 1, using the preceding snapshot to fill data prior to the deal expiration, and the terms from the subsequent snapshot (which here would be reversion to the lender's Standard Variable Rate) back to the date of deal expiry. It can of course occur that there is a deal expiry, followed by a refinance in a subsequent month, between successive PSD 007 snapshots. Here we can't map the terms of the deal in the period between the deal expiry and the refinance.

In this event, we map mortgage variables backwards and forwards either side of the refinance date (so this falls within Case 1), and the loss of accuracy should be small.

3. Case 3: there is no deal expiry or refinance between successive snapshots. Here, for variables such as interest rate and contractual monthly payment, we simply fill variables forward from the earlier snapshot. This is particularly appropriate in our setting, which will ultimately restrict the sample to the expiration of fixed rate mortgage deals. Mortgage principal requires more care, as it can still adjust from the amortization path outside of refinance, for example if households make small over-payments. If we simply filled the mortgage balance forward from the earlier snapshot, based on the amortization path, then any jumps in mortgage principal relative to this path would show up at snapshot dates. To avoid this undesirable feature, we instead smooth the path of principal between snapshots in this case, to ensure that any over/under payment of principal relative to the amortization path is constant each month.

Monthly Changes in Variables Using this monthly panel, we then create the monthly change in mortgage variables. For most variables, such as the contractual monthly payment, or interest rate, this is straightforward and is simply the value in the current month less the value in the prior month. For mortgage principal, we first calculate in a given month what the principal would have been if it evolved in line with the balance in the prior month, and the interest rate and contractual monthly payment. The monthly change in principal is then calculated as the principal this month relative to this contractual evolution of principal from the prior month. We thus calculate the change in principal relative to the contractual amortization path.

C.1.6 Event Study Windows

We next create windows around the event of deal expiry. We alternatively use windows which cover: (i) 6 months before until 6 months after deal expiry: these are used for our event window regressions; and (ii) 6 months before until 12 months after deal expiry, which are used for the impulse responses. We discuss these in turn.

Regressions: Event Windows from 6 Months Before to 6 Months After Trigger Here we create event windows which last 13 months, from 6 months before, to 6 months after the month a deal expires. We restrict the sample in this case to expiring deals where we observe the relevant mortgage variables throughout the whole event window. As the earliest snapshot of PSD 007 we observe is June 2015, the earliest deal expiry month for our event

windows is January 2016. We have consumption data until March 2024, so with a 6 month post sample, the latest trigger date we can observe is September 2023. To align with this, the 13 month event windows for the mortgage data have trigger months from January 2016 to September 2023.

Our final regression sample restricts to observations on expiring fixed rate deals with full event windows. We also drop a very small number of observations with long fixed rate deals (those in the "7 year" group), and the handful of observations missing data on the mortgage balance.

As described in Section C.1.5, for the relevant mortgage variables $V_{h,t}$ for household h at time t, we have calculated the monthly changes, $\Delta V_{h,t}$. To capture the impact of deal expiry on these variables, we compare the cumulative sum of $\Delta V_{h,t}$ before and after deal expiry. Based on the event study evidence presented in the main text, some households start to refinance 3 months before deal expiry. We thus compare the cumulative change in $V_{h,t}$ before and after this cut-off. Formally, working in event time e, relative to deal expiry at event time 0, we calculate our mortgage outcome variables as the cumulative post/pre difference $\widehat{V}_{h,e}$:

$$\widehat{V}_{h,e} := \underbrace{\sum_{i=-3}^{i=6} \Delta V_{h,i}}_{\text{Post}} - \underbrace{\sum_{i=-6}^{i=-4} \Delta V_{h,i}}_{\text{Pre}} \tag{1}$$

Two particularly important variables are the change in principal and the change in the monthly payment. To cumulate the change in the monthly payment over the refinancing window, we simply multiply the change in the monthly payment by 10, reflecting the length of the 'post' period. The change in cash on hand over the event window is then defined as the change in principal over the event window, less 10 times the change in the monthly payment.

Impulse Responses: Event Windows from 6 Months Before to 12 Months After Trigger Most of our impulse responses focus on 19 month event windows, which run until 12 months after the end of the mortgage deal. These event windows are created in the same way as the 13 month windows, with the additional restriction that the data must be observed for the full 19 months of the window. Given the availability of the consumption data, for this dataset we restrict to event windows with trigger months running from January 2016 to March 2023. The sample restrictions are the same for the 13 month window, except imposed throughout the full 19 month event window.

For the impulse responses, we calculate the change in monthly payments, relative to the average monthly payment over the pre-period (i.e. 6, 5 and 4 months before deal expiry),

and cumulate this from 3 months before deal expiry onward. The change in principal at a given horizon h is calculated as the change in principal, relative to its amortization path, between that date and the average principal over the pre-period. The cumulative change in cash on hand at horizon h is then defined as the change in principal at h less the cumulative change in monthly payments up to h.

C.2 Treatment of Housing Data

This section describes the housing data used in the paper.

C.2.1 Land Registry Data

We use two datasets from the Land Registry, which we discuss in turn.

Land Registry Price Paid Dataset The Land Registry Price Paid dataset records all residential property transactions in England and Wales since 1995. The dataset records, among other variables:

- the property address.
- the property type.
- property tenure: i.e. whether the property is owned *freehold* or *leasehold*.
- and the dates and prices of all transactions in this period.

Land Registry Registered Lease Dataset The Land Registry Registered Lease data records registered leasehold property in England and Wales. The dataset records, among other variables:

- the property address.
- the property tenure.
- the date the lease began.
- the length of the lease.

Combined Land Registry Data We merge properties identified as *leasehold* in the Price Paid data to the corresponding lease in the Registered Lease data. We merge the datasets based on the property address, and where multiple leases are associated with a given address (reflecting, for example, leasehold extensions), we select the lease that most closely precedes a given transaction in the Price Paid data. From the Price Paid data we retain information on all *freehold* properties, and information on *leasehold* properties which were successfully matched to their lease in the Registered Lease data.

C.2.2 Constructing Freehold and Leasehold House Price Series

We use the combined data series to estimate, for each Local Authority District, separate repeat sales house price indices for freehold and leasehold properties. We exclude observations for leasehold properties that have a lease extension between a pair of transactions. We drop data for the Isles of Scilly and the City of London (the *Square Mile* financial district in London) which have too few residential housing transactions. We also drop a further 7 local authority districts which have missing data for the leasehold house price series in more than 10% of months. We interpolate any missing data in the remaining series and smooth the series with a moving average process with 3 months of leads and 3 months of lags.

C.3 Transaction-level Expenditure Data

This section describes the transaction level expenditure data from finance apps.

C.3.1 Data Background

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 application in the UK that provides free advice about users' credit scores, as well as financial insights based on their transaction data. The app launched in 2018 and is still active today. The period with ClearScore data benefited from *Open Banking*, which launched in the UK in 2018 and required financial providers to make it easy for customers to share their financial data with third parties (such as financial apps) using standardised APIs. This enables a user to easily share years of their historical banking transactions data with third parties.

C.3.2 Data Description

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

- transaction description: a string of text describing the transaction as it would appear in a bank statement (with personal data such as account numbers redacted).
- 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.3.3 Transaction-level Data Treatment

Expenditure Categories

The Money Dashbord and ClearScore apps use the transaction description to classify transactions into around 300 distinct to categories known as "tags". We use this tag information, and augment the Bureau of Economic Analysis (BEA) approach[57], 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.
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.

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

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", meaning they were not automatically classified by Money Dashboard or ClearScore. That said, these "No tag" transactions still contain extensive useful information——such as the transaction value, a transaction description, whether the transaction is a credit or debit, among other fields. We use these fields to reclassify the "No tag" transactions.

We manually reclassify more than half of these "No tag" transactions using information contained in fields for the transaction description, whether the transaction is a credit or debit, and the merchant details. For example, debit transaction descriptions containing the text "tesco" tend to refer to expenditure on groceries at Tesco, a major food retailer in the UK. Similarly, debit transaction descriptions containing the text "rental" tend to refer to rental payments.

For the remaining "No tag" transactions, we use machine learning to predict their classification using the information available across other fields. Our final prediction is produced by a meta-model that combines outputs from several base learners including: (i) a text-classification model using a count vectorizer and TF-IDF transformer applied to the transaction description field; (ii) a numeric model using the transaction amount field as an input to a Naive Bayes classifier; (iii) a categorical model that maps the merchant identifier field to modal categories; and (iv) a similar categorical mapping based on combinations of the lender identified and lender transaction code fields. All of these models are trained on an over-sampled sample set.

The predictions from models (i) to (iv), along with the field for whether the transaction is a credit or debit, are taken as inputs into our meta-model. The meta-model is a regularized linear model estimated using stochastic gradient descent. It achieves an out-of-sample

accuracy of approximately 85%.

All "No tag" transactions reclassified—either manually or by machine learning—are then incorporated into the expenditure categories and related variables using their newly 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. Our "income" measure combines salary, wages and other household income.

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.

C.4 Merge Between Mortgage and Consumption Data

This section describes our merge between mortgage and consumption data.

C.4.1 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 anonymized 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.[58] 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.

C.4.2 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 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.[59] 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.

[58] See https://www.ons.gov.uk/methodology/geography/ukgeographies/postalgeography.
[59] This is without loss as UK mortgages are repaid monthly.

- *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).[60] 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). 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.

[60] 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.

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 personal 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.

C.4.3 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.

C.4.4 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 824,998 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. 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 161,531 ClearScore users to a mortgage in the PSD. Of this, we match 148,641 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 180,157 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. 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 36,617 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 196,538 mortgagors to consumption app data (reflecting a small number of mortgagors matched to both ClearScore and MoneyDashBoard).

C.4.5 Consumption Event Windows

Using the matched mortgage-consumption data, we next create event windows which cover: (i) 6 months before until 6 months after deal expiry: these are used for our event window regressions; and (ii) 6 months before until 12 months after deal expiry, which are used for the impulse responses.

Regressions: Consumption Event Windows from 6 Months Before to 6 Months After Trigger We begin with the 13 month mortgage samples, described in Online Appendix C.1.6, and merge in the monthly consumption data for the individuals that match to the mortgage data. In addition to the sample restrictions in place for the 13 month mortgage windows, for the consumption windows we further require that our broad measure of consumption is reported in each of the 13 months of the window.

With these restrictions in place, we observe just over 70 thousand consumption event windows around deal expiry. This is smaller than the roughly 195 thousand individuals we match from the financial apps to the mortgage data due mainly to: (i) the start date not being observed for every mortgage; (ii) not every matched mortgage having a deal expiry in the period consumption data is observed; and (iii) consumption data not being observed in a full window around deal expiry. In our final dataset, where we also include consumption

sample weights (Online Appendix C.5), we observe 68,734 consumption event windows, with a few households not reporting sufficient information (particularly on income) to assign the correct weight.

For the regressions, we calculate the change in average monthly consumption as the average consumption in the post-period (from 3 months before, to 6 months after deal expiry), less average consumption in the pre-period (from 6 to 4 months before deal expiry). The cumulative change in consumption over the event window is calculated as 10 times the change in these averages (note that this quantity is identical to the cumulative change in consumption relative to the pre-period, summing from 3 months before deal expiry onward). For the log change in average consumption, we calculate the growth rate of post/pre average consumption, using the Davis et al. (1996) measure.

Impulse Responses: Consumption Event Windows from 6 Months Before to 12 Months After Trigger The 19 month windows begin with the 19 month windows for the mortgage data and further restrict to our broad measure of consumption being reported in each of the 19 months of the window. The final sample has just under 53 thousand 19 month event windows for the matched mortgage-consumption dataset. The cumulative changes in the relevant variables are calculated relative to the average of the value across the pre-period (6,5,4 months before deal expiry).

C.5 Consumption Sample Weighting

This section describes how we reweight the sample of consumption data. Our matched mortgage-consumption sample is non-random given that households self-select into joining the apps. Our mortgage sample covers the full population of refinancing activity and thus we have overlapping outcome variables present in both the full mortgage sample and the matched mortgage-consumption sample. 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-randomization weighting approach outlined in Valliant and Dever (2018).

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

and measured 7 months before deal expiry; as well as broad region (NUTS1) fixed effects.[61]

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 B.1 in presents summary statistics for the mortgage population, the unweighted matched mortgage-consumption sample, and the matched mortgage-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 A.2 further demonstrates that our weights appropriately adjust the distribution (for loan-to-value, borrower age, home value, income, mortgage term, and region) in the consumption sample to match the mortgage population.

C.6 Home Movers

By construction, our analysis in this paper is restricted to households who do not move during the event windows around deal expiry. This is because, in the mortgage data, we identify households from their date of birth and property postcode: it is not straightforward to track households as their property changes (date of birth alone is not a good household identifier). Whilst we cannot easily track where mortgagors move to, we can observe when they sell their house, so we instead focus on mortgagors' selling propensity. This section explores what fraction of households sell their house during the event window, and what impact mortgage rates have on this selling propensity.

C.6.1 Measuring Home Sales

We measure home sales using the Land Registry Price Paid Dataset (Online Appendix C.2.1), which tracks all residential property transactions in England and Wales since 1995. In this section we restrict our sample to mortgagors whose property purchase in the Land Registry can be matched to the mortgage used for this purchase (Online Appendix C.1.2 provides details on the matching procedure). This transaction gives the date they bought their property: the next transaction date for the property is the date they sold their house.

[61] An exception is income, which is measured at the most recent income observed prior to this date.

Crucially, as this information comes from outwith the mortgage dataset, we can observe this sale date (or its absence) for households even if we no longer observe their mortgage.

We include all observations which have mortgage information 6 months before deal expiry. We then consider full event windows until 6 months post deal expiry, regardless of whether we observe mortgage data for the rest of the event window (as in our main sample) or not (for example, if they move).[62] We then create a dummy variable indicating whether the mortgagor sold their house in a given month throughout the event window, or not.

C.6.2 Results

Home Sales Around Deal Expiry Figure C.7 shows the unconditional response of home selling around deal expiry for treated households (green), whose mortgages expire at time 0, and for control households (red), whose deals start at the same as the treated group, but do not expire during the window.[63] The outcome variable is a dummy variable recording whether the house was sold in the given month or not (note that this variable returns to 0 in the months after the house is sold). Figure C.7 shows that there is little difference between treated and control households until the month the deal expires.[64] There is then a jump in the selling probability for the treated, peaking at around a 0.6% monthly probability of selling in the month after deal expiry. This jump for the treated suggests some households wait until their mortgage deal expires before selling, which may be more attractive financially.[65]

Overall, treated households are slightly more likely to sell during the event window, with a cumulative selling propensity that's 1.4pp higher. However, the overall magnitudes are small, with only 4.8% of treated households moving during this time.

Mortgage Rates and Home Sales We next consider whether the selling propensity during the event window is affected by mortgages rates. We follow the same regression spec-

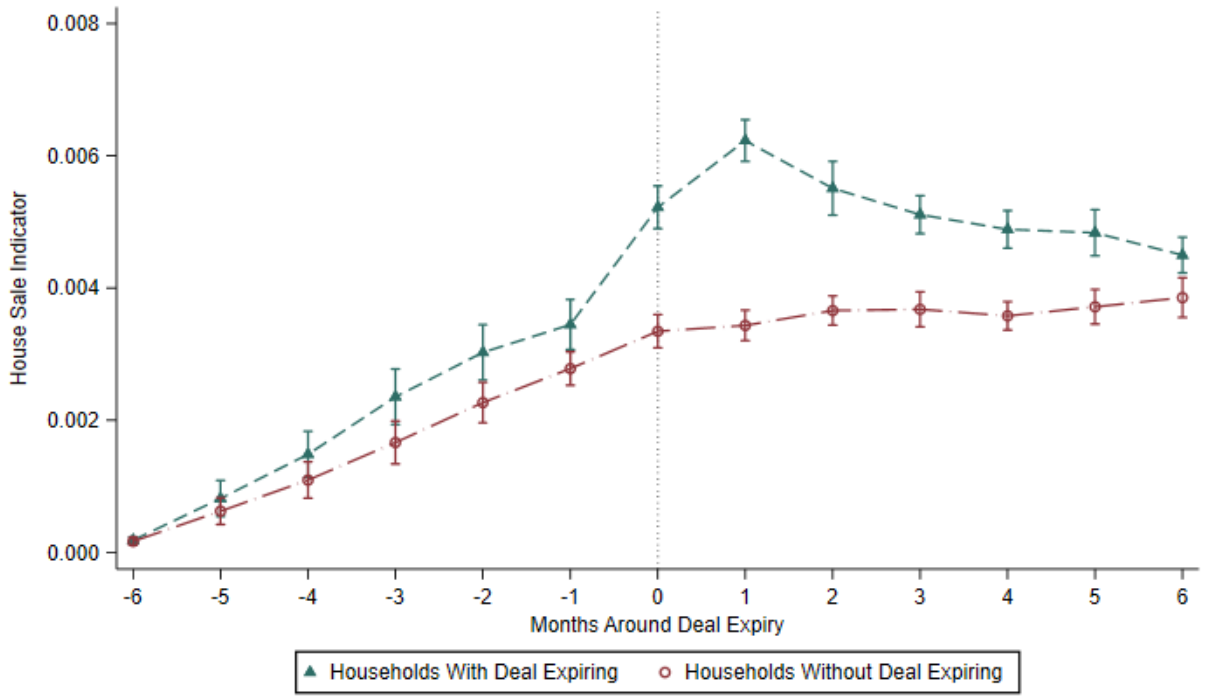
[62] Another reason for sample truncation would be the household paying off their mortgage in full during the event window and becoming an outright owner.

[63] These control groups are constructed in the same way as for the event studies in the main text.

[64] The upward trend for both treated and control households in the first 6 months of the event window likely reflects the construction of the dataset. To be in the sample, households must be observed in the mortgage dataset at the start of month −6. However, as the underlying mortgage snapshot data is observed every 6 months, many of these households would have been observed in the mortgage dataset several months after the start of the event window, and so cannot have moved until after this date.

[65] In the UK, many, but not all, mortgages can be ported. If a mortgage cannot be ported, the mortgage must be repaid in full when the house is sold, which will typically incur the early repayment charge if this happens before the deal expires. However, even where permitted, porting may not be straightforward if moving to a property of different value, resulting in a higher LTV (if moving to lower value property), or the need for additional funds or a larger mortgage (if moving to a higher value property). Material adjustments to the mortgage size during the deal could trigger the early repayment charge, which could make it more financially attractive to wait until the existing deal expires before moving.

Figure C.7: Home Selling Around Deal Expiry



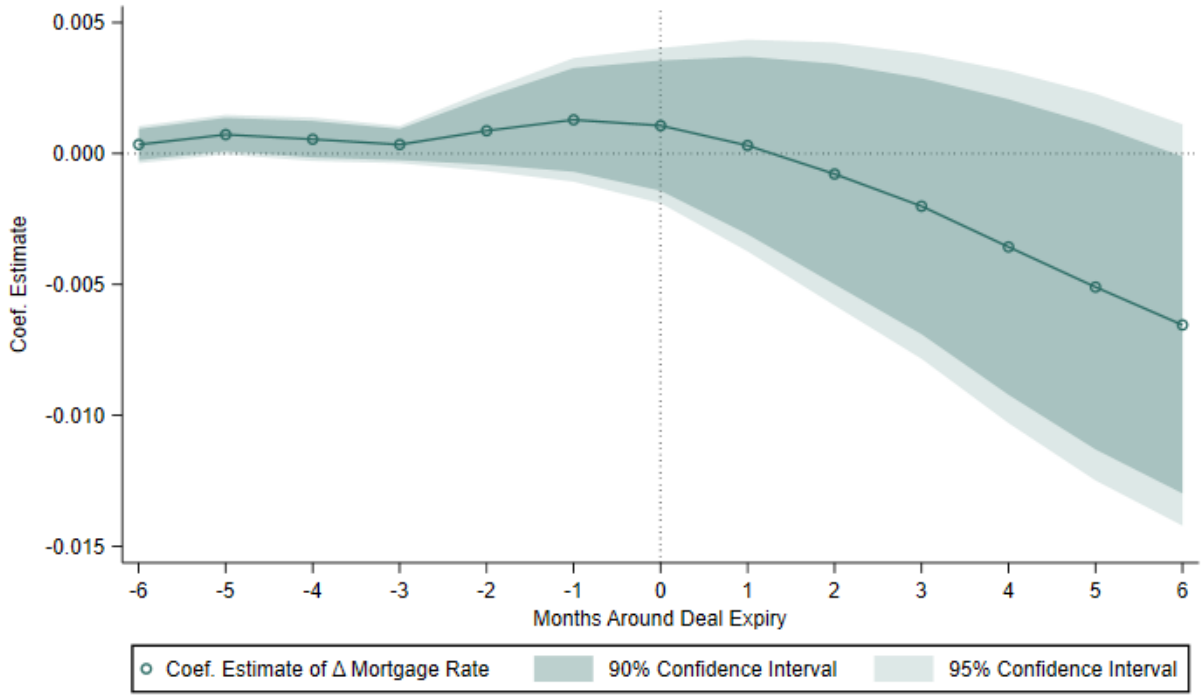
Notes: the figure plots the home selling propensity around deal expiry. In each month we report the mean of a home selling dummy for households with expiring deals, and a comparison group comprising of households whose deals do not expire within the 12-month event window. The bars represent 90% confidence intervals for deal expiry month clustered standard errors.

ification as in the main text, with the outcome variable the home selling dummy. Figure C.8 shows the cumulative impulse response of home selling in response to a 100bps increase in aggregate mortgage rates over the length of the expiring deal, including our baseline year and household age fixed effects. Overall, following deal expiry, the cumulative probability of moving decreases following an increase in rates. After 6 months the cumulative moving probability is down by around 0.5pp for a 100bps increase in mortgage rates, and is marginally significant at the 10% level, thought the estimates are quite imprecise throughout. Columns 1-4 of Table C.17 show the corresponding regression results, with cumulative outcomes after 6 months, for the standard alternative sources of variation explored in the main text. Overall, there is a rather mixed picture, with estimates of different signs across the different specifications.

The probability of a house sale may be affected both by the length of the ownership, and the time since the last deal: the household was presumably not planning on moving straight away when they bought or refinanced the house, so the selling propensity is likely to increase over time. Figure C.9 shows the impulse response where we add fixed effects to control for the deal length and the length of time the current home has been owner for. Overall, the shape is quite similar, though the error bands are quite a bit wider & the results insignificant. Columns 5-8 of Table C.17 show the corresponding regression results. The estimates are smaller in magnitude and mostly insignificant.

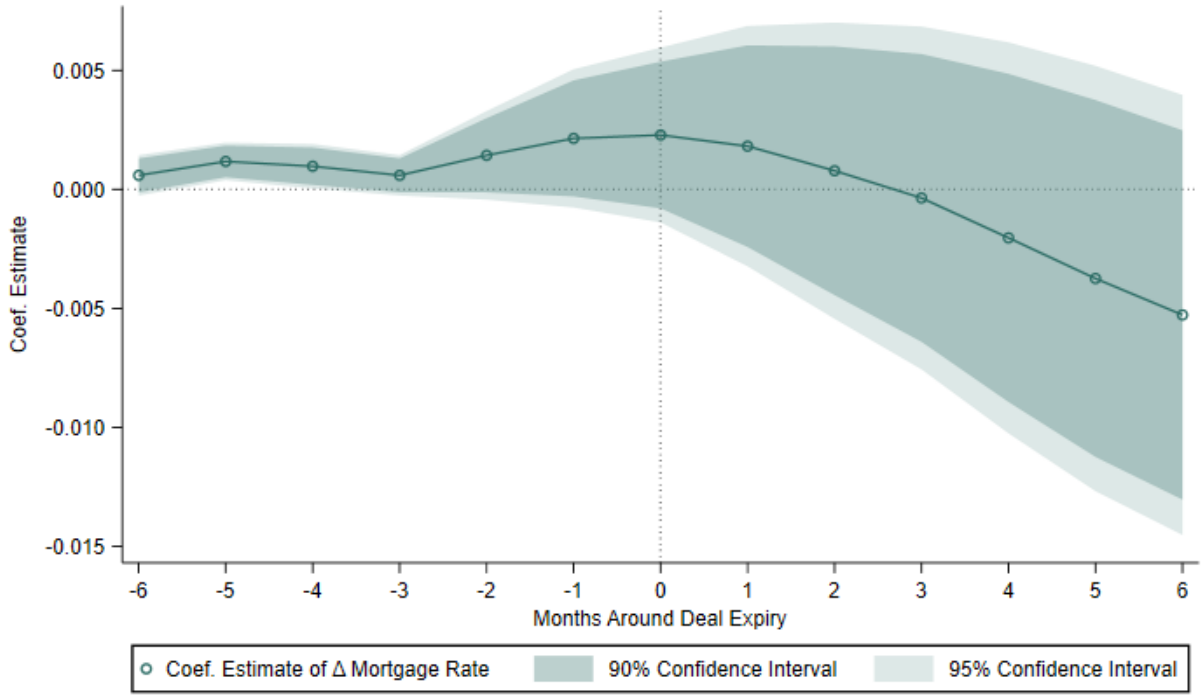
Overall, we find little robust evidence of an impact of mortgage rates on the moving probability. Moreover, the estimated impact is low relative to the total number of households present at the start of the event window.

Figure C.8: Impulse Response of Home Selling



Notes: the figure plots the impulse response for the home selling propensity. For horizons h$\in[-6, 6]$ relative to deal expiry ($t = 0$), the figure plots the cumulative change in the outcome per 1 percentage-point change in the mortgage rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. The regressions includes year and age fixed effects. Shaded bands denote 90% and 95% confidence intervals. Standard errors are clustered by the month of deal expiry.

Figure C.9: Impulse Response of Home Selling: Deal, Ownership Length Fixed Effects



Notes: the figure plots the impulse response for the home selling propensity, with additional fixed effects. For horizons h$\in[-6, 6]$ relative to deal expiry ($t = 0$), the figure plots the cumulative change in the outcome per 1 percentage-point change in the mortgage rate. Outcomes are expressed relative to the pre-period mean over months $t \in [-6, -4]$. The regression includes year and age fixed effects, and additionally fixed effects for the length of the expiring deal, and the length of time the house has been owned. Shaded bands denote 90% and 95% confidence intervals. Standard errors are clustered by the month of deal expiry.

Table C.17: Effect of Mortgage Rates on Probability of Home Selling

	Home Sale Dummy							
	Basic FE				Adding Ownership Length, Deal Length FE			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔMortgage Rate	-0.004*	-0.009***	-0.001	0.007*	-0.000	-0.003**	0.001	0.002
	(0.002)	(0.001)	(0.003)	(0.004)	(0.003)	(0.001)	(0.003)	(0.003)
Observations	5073807	5073807	5073807	2824832	5073807	5073807	5073807	2824832
Adjusted R^2	0.005	0.007	0.005	0.053	0.008	0.010	0.009	0.060
Year FE	✓			✓	✓			✓
Month FE		✓				✓		
Deal Length x Year FE			✓				✓	
Household FE				✓				✓
Age FE	✓	✓	✓		✓	✓	✓	
Deal Length FE					✓	✓		✓
Ownership Length 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 home selling propensity, in the event window around the end of the mortgage deal. The home selling propensity is a dummy variable which takes value 1 if the home is sold in a given month, and 0 otherwise. Column 1 shows the regression of the change in the home selling propensity 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 further addition throughout of fixed effects for the deal length and the length of ownership of the current property. Trigger month clustered standard errors are shown in parentheses.

C.7 Constructing House Price Sensitivity Measures

This section explains how we construct our measures of the sensitivity of house prices to rate cuts.

C.7.1 Estimating House Price $\eta's$

Baseline Estimation of η: Log-Level To estimate regional $\eta's$, we use aggregate data from the UK House Price Index for 349 Local Authority Districts in England, Wales, and Scotland. The sample runs from 2011-2023 to match the period of our mortgage deals (noting that the earliest deal expiry in our sample is January 2016, and the longest deals considered are 5 years). The baseline equation used to estimate the $\eta's$ is:

$$log\left(HP_{t,d}\right) = \alpha_d + \sum_{i \in D} \eta_i \mathbb{I}\left(i = d\right) \times rate_t + \epsilon_{d,t} \tag{2}$$

where $log\left(HP_{t,d}\right)$ is 100 times the log of the real house price index in region d in month t; α_d is a region fixed effect; and $rate_t$ is the average rate on all new mortgages completed in month t, as constructed by the Bank of England. In this regression, η_d measures the semi-elasticity of real house prices in region d with respect to aggregate mortgage rates.

We next consider alternative specifications for estimating $\eta's$.

Estimation of η: Log-Level, With Time Series Controls The first alternative approach considered retains the log-level approach, but adds in aggregate time series controls. This could include, for example: (i) a linear time trend; (ii) a linear and quadratic time trend; (iii) a time fixed effect. In the last case, we fully absorb the aggregate time series, and we estimate *relative* η's for 348 out of 349 regions, with the $\eta's$ relative to an omitted region.

It turns out that any of these methods will produce η's that are perfectly correlated with those calculated by equation (2). Specifically, controlling for aggregate time trends shifts all the $\eta's$ by a fixed constant (which depends on the time trends included). Therefore the level of η might change across specifications, but the dispersion will not. The dispersion is what identifies our estimate of how rates affect households via asset prices.

This result is shown in the following proposition.

Proposition C.1. *Let $\hat{\eta}_d$ be the estimated coefficients from equation 2. Consider adding to equation 2 a control $f(t)$ which is a function of time only. Let $\hat{\lambda}_d$ be the estimated semi-elasticities from this regression. Then $\exists c \in \mathbb{R} : \forall d\ \hat{\lambda}_d = \hat{\eta}_d + c$. In words: the regional coefficient estimates all shift by the same constant.*

Proof. Adding the time control, we have the following regression:

$$log\left(HP_{t,d}\right) = \alpha_d + \sum_{i \in D} \lambda_i \mathbb{I}\left(i = d\right) \times rate_t + \gamma' f(t) + \epsilon_{d,t} \tag{3}$$

This is then equivalent to running the regression

$$log\left(HP_{t,d}\right) - \gamma' f(t) = \alpha_d + \sum_{i \in D} \lambda_i \mathbb{I}\left(i = d\right) \times rate_t + \epsilon_{d,t} \tag{4}$$

As all parameters on the RHS vary by region, this regression can be run region by region.

Thus, the estimate for region d can be obtained from running the following regression for region d :

$$log\left(HP_{t,d}\right) - \gamma' f(t) = \alpha_d + \lambda_d rate_t + \epsilon_{d,t} \tag{5}$$

Note, crucially, that $\gamma' f(t)$ is the same for each region: $f(t)$ only varies in the time series, and the vector γ' doesn't vary by region.

Now, from equation 5, the estimate of λ_d is given by

$$\hat{\lambda}_d = \frac{Cov\left(rate_t, log\left(HP_{t,d}\right) - \gamma' f(t)\right)}{Var\left(rate_t\right)} = \frac{Cov\left(rate_t, log\left(HP_{t,d}\right)\right)}{Var\left(rate_t\right)} - \frac{Cov\left(rate_t, \gamma' f(t)\right)}{Var\left(rate_t\right)}$$

Now, from equation 2, given that all parameters vary at the regional level, we have

$$\hat{\eta}_d = \frac{Cov\left(rate_t, log\left(HP_{t,d}\right)\right)}{Var\left(rate_t\right)}$$

Thus, we have that

$$\hat{\lambda}_d = \hat{\eta}_d - \frac{Cov\left(rate_t, \gamma' f(t)\right)}{Var\left(rate_t\right)}$$

Let $c := -\frac{Cov(rate_t, \gamma' f(t))}{Var(rate_t)}$ and note that c is independent of the region d. This completes the proof. □

We note that in the special case where $f(t)$ is a 1x1 column vector, the coefficient c is equal to $-\gamma$ times the coefficient from a univariate regression of $f(t)$ on $rate_t$. The proposition shows that absorbing aggregate trends of arbitrary form doesn't affect the dispersion of the η estimates.

Estimation of η: Log Difference on Level Difference A second alternative approach is to estimate the log difference in prices on the log difference in rates.

$$log\left(HP_{t,d}\right) - log\left(HP_{t-L,d}\right) = \alpha_d + \sum_{i \in D} \eta_i \mathbb{I}\left(i = d\right) \times \left(rate_t - rate_{t-L}\right) + \epsilon_{d,t} \tag{6}$$

Given the prevalence of two-year fixed rate mortgage deals in the UK, for this calculation we regress the two-year change in real house prices on the two-year change in rates. In aggregate, this series is 91% correlated with the baseline η measure.

Alternative η Estimates: Results Table C.18 shows results for principal and consumption when the regional house price $\eta's$ are calculated using the alternative methods discussed above. In each case, following the main text, we split the observations into 3 groups by η, capturing, respectively, 10, 80, and 10 percent of the distribution. Column 1 shows the baseline results from the paper, where η is calculated following equation 2. Column 2 shows a regression with alternative $\eta's$ which have been estimated following equation 3, specifically a log-level specification with a month fixed effect. As shown in the proposition above, these $\eta's$ are perfectly correlated with the baseline η's, and so have exactly the same buckets and regression results. In Column 3, the $\eta's$ have been calculated using equation 6, for the 2 year change in log house prices on the 2 year change in mortgage rates. The estimates are very similar to the baseline. Columns 4-6 show results for consumption, with a very similar message.

Table C.18: Effect of Rates on Mortgage Principal & Consumption via Asset Prices: Alternative η Estimates

	Δ Principal			Δ Consumption		
	Log Level η	Log Level η: Time FE	Diff η: 2 Year	Log Level η	Log Level η: Time FE	Diff η: 2 Year
	(1)	(2)	(3)	(4)	(5)	(6)
ΔMortgage Rate x $\mathbb{I}$(Bottom 10% η)	-1671.2***	-1671.2***	-1807.7***	-1946.2**	-1946.2**	-1769.7**
	(243.9)	(243.9)	(285.7)	(761.8)	(761.8)	(837.1)
ΔMortgage Rate x $\mathbb{I}$(Middle 80% η)	-1092.5***	-1092.5***	-1056.7***	-845.9***	-845.9***	-862.6***
	(112.1)	(112.1)	(108.3)	(295.7)	(295.7)	(278.7)
ΔMortgage Rate x $\mathbb{I}$(Top 10% η)	-531.3***	-531.3***	-691.8***	-64.9	-64.9	-201.2
	(71.6)	(71.6)	(75.1)	(421.6)	(421.6)	(427.8)
Observations	6796438	6796438	6796438	68541	68541	68541
Adjusted R^2	0.012	0.012	0.012	0.021	0.021	0.021
Year x LAD 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 rates on the change in mortgage principal and consumption, via asset prices, for alternative measures of house price sensitivity η. 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 1-3 the dependent variable is the change in mortgage principal in the event window around the end of the deal, shown for the full mortgage sample. Column 1 shows results with the baseline η measure; in column 2 the regional $\eta's$ are calculated in the log-level specification, including a month fixed effect. In column 3, the regional $\eta's$ are calculated from the two-year difference of real house prices on the 2 year difference in average mortgage rates. Each regression includes year x Local Authority District fixed effects and household age fixed effects. Columns 4-6 repeat this with the cumulative change in consumption around deal expiry as the outcome variable. This is run on the matched mortgage-consumption dataset and the observations are weighted by the sample selection weights. Trigger month clustered standard errors are shown in parentheses. Outcomes are winsorized at the 0.1% and 99.9% level.

C.7.2 Housing Supply Elasticity Instrument

To instrument our district-level η_d, we use a similar approach to that outlined by Guren, McKay, Nakamura and Steinsson (2021). We first estimate district-level proxies of housing supply elasticities for the United Kingdom, as proposed by Guren et al. (2021) for the United States. These proxies reflect the systematic historical sensitivity of district-level house prices to broader regional housing cycles. We estimate these historical sensitivity estimates as:

$$\Delta log\left(HP_{drt}\right) = \psi_d + \phi_d \Delta log\left(HP_{rt}\right) + v_{drt} \tag{7}$$

where $\Delta log\left(HP_{drt}\right)$ is the annual log change in the real house price index of the granular local authority district d contained within the broader (NUTS1) region r in month t, $\Delta log\left(HP_{rt}\right)$ is the annual log change in the real house price index of the broad region r in month t, and ϕ_d is a district-specific coefficient. We estimate equation 7 over the period from 2011-23. We then construct our instrument for η_d as $z_{dr} = \hat{\phi}_d \eta_r$, where η_r is the house price responsiveness to interest rates at the broad region r level.[66]

C.8 Interest Rate Notches and Collateral Constraints

We investigate the role of collateral in driving the household financial accelerator by following a similar approach to Cloyne et al. (2019). In the UK, mortgage interest rates increase discretely at specific LTV thresholds (typically at LTV ratios of 50 percent, 60 percent, 70 percent, 75 percent, 80 percent, 85 percent and 90 percent). Cloyne et al. (2019) argue these thresholds represent soft collateral constraints and use these interest notches to explore whether the collateral channel is driving a household borrowing response to house price growth. If collateral constraints matter, they argue that the borrowing response should differ depending on whether a homeowner's LTV crosses an interest rate notch threshold, or not.

We estimate the following regression:

$$\Delta \text{borrowing}_{ikt} = \alpha_{j(i)t} + \beta_k \Delta log\left(HP_{j(i)t}\right) \times N_k + \text{fixed effects}_{ikt} + v_{ikt} \tag{8}$$

where the left-hand side is the cumulative response of borrowing from 3 months before deal expiry to 6 months after. $\triangle log\left(HP_{idt}\right)$ is 100 times the log change in the real house price index over the deal for local authority district j to which household i belongs. Similar to Cloyne et al. (2019), N_k is a variable with categories whereby real house price growth

[66] In Guren et al. (2021), they create a time-varying instrument and interact the time-invariant historical sensitivity estimate ϕ_d with today's change in regional house prices.

moves homeowners' LTV either: below their current interest rate notch ("relaxed"); no move to their interest rate notch ("unchanged"); or above their current interest rate notch ("reinforced"). We consider changes to interest rate notches by comparing two different LTV ratios: one based on the pre-refinance mortgage balance and the home value at the start of the last deal; the other is based on the pre-refinance mortgage balance and the updated home value at the time of refinancing.

Table C.19 presents the results. The borrowing response is the largest when the collateral constraints are relaxed in column 1 (baseline fixed effects) and column 2 (where we add household fixed effects). In columns 3 and 4, we instrument the real house price growth and notch move interaction term with: the interaction of the notch change variable, the change in aggregate mortgage rates over the deal, and the regional house price η. The results are similar although the magnitude of the borrowing response is larger. Columns 5-8 repeat this exercise with the log change in borrowing as the dependent variable. Here we consistently see that the borrowing response is smallest when the collateral constraint is reinforced, but it does not uniformly increase as the collateral constraint moves from unchanged to relaxed.

Table C.19: Heterogeneity in Borrowing Response to Real House Price Changes by Notches Moved

	Δ Principal				Δ Log Principal			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ΔLog HPI x Relaxed Notch	68.5***	168.1***	152.1***	454.3***	0.022**	0.073***	0.060*	0.198***
	(7.64)	(11.7)	(40.0)	(55.2)	(0.009)	(0.009)	(0.035)	(0.040)
ΔLog HPI x Unchanged Notch	20.6**	79.7***	217.3***	325.9***	0.014	0.129***	0.232***	0.450***
	(8.65)	(12.1)	(56.2)	(57.2)	(0.010)	(0.014)	(0.050)	(0.046)
ΔLog HPI x Reinforced Notch	11.5	23.7	26.2	6.6	-0.007	0.051***	-0.116	0.042
	(11.4)	(22.0)	(111.3)	(310.2)	(0.010)	(0.018)	(0.072)	(0.142)
Observations	6740139	3647742	6740139	3647742	6737541	3646012	6737541	3646012
Adjusted R^2	0.012	0.077	-0.002	-0.790	0.014	0.061	-0.004	-0.795
K-Paap F-Stat			17.16	9.58			36.86	16.26
Year x LAD FE	✓	✓	✓	✓	✓	✓	✓	✓
Household FE		✓		✓		✓		✓
Age FE	✓		✓		✓		✓	
Notch Change FE	✓	✓	✓	✓	✓	✓	✓	✓

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

Notes: The table shows how the response of borrowing to real house price growth varies as households cross LTV notches. Interest rate notches occur at LTV thresholds of 50 percent, 60 percent, 70 percent, 75 percent, 80 percent, 85 percent, and 90 percent. We define the collateral constraint as being relaxed (reinforced) when house price growth moves a mortgagor at least one LTV notch down (up). In columns 1-4 the dependent variable is the change in principal in the event window around the end of the deal. Column 1 regresses this change in principal on the interaction of the notch change categorical variable with real house price growth over the expiring deal. Column 1 includes year x Local Authority District (LAD), household age, and notch change fixed effects. Column 2 replaces household age fixed effects with household fixed effects. Columns 3 and 4 instrument the interaction of the notch change variable and real house price growth with the interaction of the notch change variable, the change in aggregate mortgage rates over the deal, and the regional house price η. Columns 5-8 repeat columns 1-4 with the log of borrowing 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. Outcomes are winsorized at the 0.1% and 99.9% level.

D Model Appendix

This section contains the Model Appendix. Section D.1 contains our derivation of the equations characterizing the household debt channel. Section D.2 explains how rate cuts affect borrowing in our model. Section D.3 solves for the household debt channel in closed form, for a special case, and also discusses labor supply effects.

D.1 Deriving the Household Debt Channel

This section derives equations (1) and (2) from the main text, which summarize the household debt channel. To do so formally, we need some additional notation. Define the exogenous state variables of the household by the vector $s_t = (y_t, r_t^a, r_t)$, with a corresponding history $s^t = (s_{-N}, ..., s_t)$.

The old mortgage deal started in period $-N$, meaning that $d\left(s^{-N}\right)$ is pre-determined. The household therefore solves

$$\max_{\{c(s^t),a(s^t)\}_{t=-N,s^t}^{N},\{d(s^0)\}_{s^0}} \sum_{t=-N}^{N}\sum_{s^t}\beta^{t+N}u\left(c\left(s^t\right)\right)\pi\left(s^t\right)+\sum_{s^N}\beta^{2N+1}v\left(a\left(s^N\right),d\left(s^N\right)\right)\pi\left(s^N\right)$$

subject to

$$c\left(s^t\right)+a\left(s^t\right)-d\left(s^t\right)=y\left(s^t\right)+\left(1+r^a\left(s^t\right)\right)a\left(s^{t-1}\right)-\left(1+r^d\left(s^t\right)\right)d\left(s^{t-1}\right)$$

and

$$d\left(s^t\right)-d\left(s^{t-1}\right)=-\xi \quad t\neq 0$$

and

$$a\left(s^t\right)\geq \underline{a}$$

with $d\left(s^{-N}\right)$ and $a\left(s^{-N}\right)$ pre-determined.

The mortgage rate solves

$$r^d\left(s^t\right)=\begin{cases} r\left(s^{-N}\right)+\mu\left(s^{-N}\right) & t<0 \\ r\left(s^0\right)+\mu\left(s^0\right) & t\geq 0 \end{cases}$$

$$\mu\left(s^t\right)=\mu\left(\frac{d\left(s^t\right)}{p\left(r_t\right)h}\right),\mu'(\cdot)>0.$$

We will study a particular history $\tilde{s}^1$, and the continuation path leading to this history $\tilde{s}^0$.

For brevity notation-wise, we will write

$$c\left(\tilde{s}^1\right) = c_1 \quad a\left(\tilde{s}^1\right) = a_1 \quad d\left(\tilde{s}^1\right) = d_1 \quad y\left(\tilde{s}^1\right) = y_1 \quad c\left(\tilde{s}^0\right) = c_0,$$

and so on.

We will then allow $y_1 = y = \bar{y} + dy$, $r_0 = r = \bar{r} + dr$, and $r_0^a = r^a = \bar{r}^a + dr^a$ to vary around steady state values $\bar{y}$, $\bar{r}$, and $\bar{r}^a$, while holding all other values of $y\left(s^t\right)$, $r\left(s^t\right)$ and $r^a\left(s^t\right)$ fixed. Then we can write policy functions

$$c_1 = c_1\left(y, r, r^a\right) \quad a_1 = a_1\left(y, r, r^a\right) \quad d_1 = d_1\left(y, r, r^a\right) \quad c_0 = c_0\left(y, r, r^a\right),$$

and so on, where we are suppressing dependence on other values of $y\left(s^t\right)$, $r\left(s^t\right)$ and $r^a\left(s^t\right)$ since they are not varying. For notational simplicity we will also write

$$\mu_0\left(y, r, r^a\right) = \mu\left(\frac{d_0}{p\left(r_0\right)h}\right)$$

Substituting the policy functions into the time 1 and state $\tilde{s}^1$ budget constraint yields

$$c_1 + a_1 - d_1 = y_1 + \left(1 + r_0^a\right)a_0 - \left(1 + r_0^d\right)d_0.$$

Then we have

$$\begin{aligned}
dc_1 + da_1 - dd_1 &= dy + d\left[\left(1 + r_0^a\right)a_0\right] - d\left[\left(1 + r_0^d\right)d_0\right] \\
\implies dc_1 + da_1 - dd_1 &= dy + \left(1 + r_0^a\right)da_0 + a_0 dr^a - \left(1 + r_0^d\right)dd_0 - d_0 dr_0^d \\
\implies dc_1 + da_1 - \left(1 + r_0^a\right)da_0 &= dd_1 + dy + a_0 dr^a - \left(1 + r_0^d\right)dd_0 - d_0 dr_0^d \qquad (9)
\end{aligned}$$

Substituting in the definitions of policy functions, the right hand side becomes

$$\begin{aligned}
&dd_1 + dy + a_0 dr^a - \left(1 + r_0^d\right)dd_0 - d_0 dr_0^d \\
&= dd_1\left(y, r, r^a\right) + dy + a_0 dr^a - \left(1 + r_0^d\right)dd_0\left(y, r, r^a\right) - d_0 d\left[r_0 + \mu\left(y, r, r^a\right)\right] \\
&= \left[\frac{\partial d_1}{\partial y}dy + \frac{\partial d_1}{\partial r}dr + \frac{\partial d_1}{\partial r^a}dr^a\right] + dy + a_0 dr^a \\
&\quad - \left(1 + r_0^d\right)\left[\frac{\partial d_0}{\partial y}dy + \frac{\partial d_0}{\partial r}dr + \frac{\partial d_0}{\partial r^a}dr^a\right] - d_0\left[dr + \frac{\partial \mu}{\partial y}dy + \frac{\partial \mu}{\partial r}dr + \frac{\partial \mu}{\partial r^a}dr^a\right] \\
&= \left[\frac{\partial d_1}{\partial y} + 1 - \left(1 + r_0^d\right)\frac{\partial d_0}{\partial y} - d_0\frac{\partial \mu}{\partial y}\right]dy + \left[\frac{\partial d_1}{\partial r} - \left(1 + r_0^d\right)\frac{\partial d_0}{\partial r} - d_0\left(1 + \frac{\partial \mu}{\partial r}\right)\right]dr
\end{aligned}$$

$$+\left[\frac{\partial d_1}{\partial r^a}+a_0-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r^a}-d_0\frac{\partial \mu}{\partial r^a}\right]dr^a$$

$$=\tilde{\mathcal{Y}}_0 dy+\tilde{\mathcal{A}}_0 dr^a+\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]dr \tag{10}$$

where

$$\tilde{\mathcal{Y}}_0=\left[\frac{\partial d_1}{\partial y}+1-\left(1+r_0^d\right)\frac{\partial d_0}{\partial y}-d_0\frac{\partial \mu}{\partial y}\right]$$

$$\tilde{\mathcal{A}}_0=\left[\frac{\partial d_1}{\partial r^a}+a_0-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r^a}-d_0\frac{\partial \mu}{\partial r^a}\right].$$

The preceding expressions (9) and (10) imply

$$dc_1+da_1-\left(1+r_0^a\right)da_0=\tilde{\mathcal{Y}}_0 dy+\tilde{\mathcal{A}}_0 dr^a+\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]dr$$

$$\implies dc_1\times\frac{dc_1+da_1-\left(1+r_0^a\right)da_0}{dc_1}=\tilde{\mathcal{Y}}_0 dy+\tilde{\mathcal{A}}_0 dr^a+\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]dr$$

$$\begin{aligned}\implies dc_1&=\frac{dc_1}{dc_1+da_1-\left(1+r_0^a\right)da_0}\left[\tilde{\mathcal{Y}}_0 dy+\tilde{\mathcal{A}}_0 dr^a+\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]dr\right]\\&=\text{MPC}_0\left[\tilde{\mathcal{Y}}_0 dy+\tilde{\mathcal{A}}_0 dr^a+\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]dr\right]\\&=\text{MPC}_0\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]dr+\mathcal{Y}_0 dy+\mathcal{A}_0 dr^a\\\text{MPC}_0&\equiv\frac{dc_1}{dc_1+da_1-\left(1+r_0^a\right)da_0}\mathcal{Y}_0\equiv\text{MPC}_0\tilde{\mathcal{Y}}_0,\mathcal{A}_0\equiv\text{MPC}_0\tilde{\mathcal{A}}_0.\end{aligned}$$

We can evaluate MPC_0, $\mathcal{Y}_0$ and $\mathcal{A}_0$ at their steady state values, with $dr=0$, $dy=0$ and $dr^a=0$, we will denote these objects by $\text{MPC}^{\text{cash-on-hand}}$, $\mathcal{Y}$ and $\mathcal{A}$. With these simplifications, the final expression reads that to a first order

$$\Delta c_1\approx\text{MPC}^{\text{cash-on-hand}}\left[\frac{\partial d_1}{\partial r}-\left(1+r_0^d\right)\frac{\partial d_0}{\partial r}-d_0\left(1+\frac{\partial \mu}{\partial r}\right)\right]\Delta r+\mathcal{Y}\Delta y+\mathcal{A}\Delta r^a$$

for changes $\Delta r,\Delta y$ and Δr^a. This final equation derives equations (1) and (2).

D.2 The Effect of Rate Cuts on Borrowing via Asset Prices

This subsection derives the effects of rate cuts on borrowing, equations (3) and (4) from the main text.

Hard collateral constraint. Suppose the spread function $\mu(\cdot)$ is flat up to a maximum loan-to-value ratio θ, and vertical at θ. If the constraint binds at the start of the deal, then mortgage borrowing at time 1 is pinned down by the collateral value,

$$d_1 = \theta p\left(r_0\right) h.$$

Taking the total derivative with respect to r_0 and evaluating around the initial steady state r_{-N} yields

$$\frac{\partial d_1}{\partial r_0} = \theta h\, p'\left(r_{-N}\right),$$

which is equation (3) in the main text.

Case with intertemporal substitution. We now make the following simplifying assumptions:

- Households can only borrow with the mortgage asset, i.e. $\underline{a} \to \infty$
- Labor income is constant, i.e. $y_t = y$
- The rate cut is unanticipated until time 0, with the interest rate at its steady state value r before time 0 and from time N onwards
- The repayment is interest only, i.e. $\xi = 0$
- The household is patient, so $\beta\left(1+r\right) = 1$
- Interest rates are small at steady state, so $r \approx 0$

With these assumptions, the budget constraint is

$$c_t - d_t = y - \left(1 + r^d_{t-1}\right) d_{t-1}.$$

Consumption satisfies:

- At time 0:
$$c_0 = d_0 + y - \left(1+r\right) d_{-1}. \tag{11}$$
- At time $1 \leq t < N$:
$$\begin{aligned} c_{\text{mid}} - d_{t-1} &= y - \left(1+r_0\right) d_{t-1} \\ \implies c_{\text{mid}} &= y - r_0 d_0 \end{aligned} \tag{12}$$

- At time $t = N$:

$$c_N + (1 + r_0) d_0 = y + d_N \implies c_N = y + d_N - (1 + r_0) d_0$$

- From time $t \geq N$, consumption is optimally constant, since the economy has returned to steady state, which implies

$$c_N = y - r d_N.$$

Equating the two expressions for c_N, we have

$$\begin{aligned} y + d_N - (1 + r_0) d_0 &= y - r d_N \\ \implies d_N &= \frac{1 + r_0}{1 + r} d_0 \\ \implies c_N &= y - r \left(\frac{1 + r_0}{1 + r} d_0 \right). \end{aligned} \tag{13}$$

Therefore we have written consumption at all times as a function of d_0, so that the agent solves

$$\max_{d_0} \sum_{t=0}^{\infty} \beta^t u(c_t) = \max_{d_0} \left[u(d_0 + y - (1 + r) d_{-1}) + \sum_{t=1}^{N-1} \beta^t u(y - r_0 d_0) + \sum_{t=N}^{\infty} \beta^t u \left(y - r \left(\frac{1 + r_0}{1 + r} d_0 \right) \right) \right],$$

which has a first order condition

$$u'(c_0) = r_0 \left(\sum_{t=1}^{N-1} \beta^t \right) u'(c_{\text{mid}}) + \left(r \frac{1 + r_0}{1 + r} \right) \left(\sum_{t=N}^{\infty} \beta^t \right) u'(c_N),$$

which simplifies with to

$$\begin{aligned} u'(c_0) &= r_0 (N - 1) u'(c_{\text{mid}}) + (1 + r_0) \left(\frac{r}{1 + r} \right) \left(\frac{\beta^N}{1 - \beta} \right) u'(c_N) \\ &= r_0 (N - 1) u'(c_{\text{mid}}) + (1 + r_0) (1 + r)^{-N} u'(c_N) \\ &\approx r_0 (N - 1) u'(c_{\text{mid}}) + (1 + r_0) u'(c_N) \end{aligned} \tag{14}$$

where the second line uses $\beta (1 + r) = 1$, and the third line uses $r \approx 0$. Finally, we solve for dc_0/dr_0.

From equation (11) we have

$$\frac{dc_0}{dr_0} = \frac{dd_0}{dr_0}$$

and from equation (12) we have

$$\frac{dc_{\text{mid}}}{dr_0} = -\left(d_0 + r_0 \frac{dd_0}{dr_0}\right)$$

$$\implies \frac{dc_{\text{mid}}}{dr_0} = -\left(d_0 + r_0 \frac{dc_0}{dr_0}\right).$$

Finally, from equation (13) we have

$$\frac{dc_N}{dr_0} = -\frac{rd_0}{1+r} - \frac{r\left(1+r_0\right)}{1+r}\frac{dc_0}{dr_0} \approx 0,$$

where the final approximation uses that r is small. Differentiating equation (14) implies

$$u''\left(c_0\right)\frac{dc_0}{dr_0} = \left(N-1\right)u'\left(c_{\text{mid}}\right) - r_0\left(N-1\right)u''\left(c_{\text{mid}}\right)\left(d_0 + r_0\frac{dc_0}{dr_0}\right) + u'\left(c_N\right)$$

which is approximately

$$u''\left(c\right)\frac{dc_0}{dr_0} \approx \left(N-1\right)u'\left(c\right) + u'\left(c\right)$$

where we have used the approximations that consumption is in a neighborhood of the steady state, and r_0 is small. Rearranging yields

$$\frac{dc_0}{dr_0} = -cN\sigma$$

where σ is the elasticity of intertemporal substitution. As a result, we also have

$$\frac{dd_1}{dr_0} = -cN\sigma.$$

D.3 Deriving the Household Debt Channel in Closed Form

This section extends the baseline model, for two purposes. First, we derive closed form expressions for the household debt channel, equation (1) from the main text, in a simplified version of the model that is analytically tractable. We first present the model, then the expression for the household debt channel, and finally the derivations. Second, we also use this model to discuss the effects of endogenous labor supply.

D.3.1 Model Setup

There are two types of households $i \in \{P, I\}$, Patient and Impatient, with fraction π patient. Discount factors satisfy $\beta_I < \beta_P$.

Households' mortgage deals expire every N periods. The main focus is the refinancing at time 0. Type i chooses consumption c_{it}, labour ℓ_{it}, and mortgage debt d_{it} to solve

$$\max_{\{c_{it},\ell_{it},d_{it}\}_{t=-N}^{\infty}} \sum_{t=-N}^{\infty} \beta_i^t \big[u(c_{it}) - v(\ell_{it})\big],$$

where $u(c) = \frac{c^{1-1/\sigma}}{1-1/\sigma}$ and $v(\ell) = \frac{\ell^{1+1/\varphi}}{1+1/\varphi}$, subject to:

- **Budget constraint** ($d_{it} > 0$ is debt/liability):

$$c_{it} - d_{it} = w_t \ell_{it} - (1 + r_t^d) d_{i,t-1}.$$

- **Adjustment**: debt can be re-set only at times $t = kN$, $k \geq -1$. Otherwise principal does not change:

$$d_{it} = d_{i,t-1} \quad (t \neq kN).$$

- **Initial condition**: $d_{i,-N-1} = 0$.

- **Deal rate**: $r_{kN+j}^d = r_{kN}$ for $j = 0, \ldots, N-1$.

- **Collateral constraint** at refinancing: with house size h and house-price function $p(r)$,

$$d_{ikN} \leq \theta h p(r_{kN}).$$

At $t = -N$ the economy is in steady state with $w_t = w$ and $r_t = r$. Agents expect this to persist. Between $t = -N$ and $t = 0$ there are unexpected shocks to interest rates

$$r_0 = r + \Delta r,$$

as well a shock to wages at time 0

$$w_0 = w + \Delta w,$$

after which the economy reverts to the steady state.

The patient type satisfies $\beta_P(1 + r) = 1$. The impatient type is sufficiently impatient that the collateral constraint binds at $t = -N$ and at $t = 0$.

D.3.2 Household Debt Channel in Closed Form

For small r, the first-order responses to a small rate surprise $\Delta r \equiv r_0 - r_{-N}$ are:

Impatient households.

$$\Delta d_{I,0} \approx \theta h p'(r_{-N}) \Delta r, \tag{15}$$

$$\Delta c_{I,0} \approx \frac{\ell_I(1+\varphi)}{1+\varphi/\sigma} \Delta w_0 + \frac{\theta\, h\, p'(r_{-N}) - d_I}{1+\varphi/\sigma} \Delta r, \tag{16}$$

$$\Delta(w_0 \ell_{I,0}) \approx \ell_I(1+\varphi) \Delta w_0 - \frac{\varphi}{\sigma}\, \Delta c_{I,0}. \tag{17}$$

i.e. the consumption response depends on the labor income response and the borrowing response. Consumption responds by less than borrowing since some of the higher borrowing is "consumed" by more leisure. The labor income response is increasing in wages but decreasing in consumption (wealth effects on labor supply).

Patient households.

$$\Delta d_{P,0} \approx -(1+\varphi)\, \ell_P\, \Delta w - (N-1)\, c_P(\sigma+\varphi)\, \Delta r, \tag{18}$$

$$\Delta c_{P,0} \approx -(N-1)\sigma\, c_P\, \Delta r, \tag{19}$$

$$\Delta(w_0 \ell_{P,0}) \approx \ell_P(1+\varphi)\, \Delta w_0 - \frac{\varphi}{\sigma}\, \Delta c_{P,0}, \tag{20}$$

Regarding consumption, since we have permanent income consumers and r is small, there is only the intertemporal substitution effect, which is amplified since the rate cut lasts for N periods. The cashflows and income movements do not matter since the MPC is zero for these households.

Average consumption.

$$\begin{aligned}
\mathbb{E}[\Delta c_{i0}] = \Big[&- \pi(N-1)\sigma c_P + (1-\pi)\tfrac{\theta\, h\, p'(r_{-N})}{1+\varphi/\sigma} \Big] \Delta r \quad \to \quad \text{borrowing} \\
&- (1-\pi)\frac{d_I}{1+\varphi/\sigma} \Delta r \quad \to \quad \text{cashflows} \\
&+ (1-\pi)\frac{\ell_I(1+\varphi)}{1+\varphi/\sigma}\, \Delta w_0 \quad \to \quad \text{income}
\end{aligned}$$

naturally the borrowing response can be decomposed into the response of intertemporal substitution and asset prices.[67]

[67] In the consumption response, the intertemporal elasticity of substitution σ is multiplied by (N-1) and not N, as in the main text. The reason for the difference is the timing of the budget constraint in the model extension, versus the main text.

Relation between labor income and consumption. Equations (17) and (20) show that, holding fixed wages, the relationship between consumption changes and labor income changes equals φ/σ, i.e. the ratio of the Frisch elasticity of labor supply to the elasticity of intertemporal substitution. Therefore after an interest rate shock, and holding fixed wages, the response of labor income and consumption should be the same if these two parameters are the same. These equations also make it clear that changes in labor income are correlated with changes in consumption due to endogenous movements in labor supply. Hence controlling for changes in labor income, in regressions of consumption changes on rate cuts, is a "bad control".

D.3.3 Derivations

We now show how to derive the household debt channel in closed form for this model. The first order borrowing response of the constrained agent is

$$\Delta d_{I,0} = \theta h \big[p(r_0) - p(r_{-N})\big] \approx \theta h \, p'(r_{-N}) \, (r_0 - r_{-N}). \tag{21}$$

Regarding their response of consumption and labor, the static FOC $w_0 u'(c_{I,0}) = v'(\ell_{I,0})$ implies $\ell_{I,0} = w_0^{\varphi} c_{I,0}^{-\varphi/\sigma}$. Linearising around (c_I, ℓ_I, w):

$$\Delta c_{I,0} - \Delta d_{I,0} \approx w \, \Delta \ell_{I,0} + \ell_I \Delta w_0 - d_I (r_0 - r_{-N}), \tag{22}$$

$$\frac{\Delta \ell_{I,0}}{\ell_I} \approx \varphi \frac{\Delta w_0}{w} - \frac{\varphi}{\sigma} \frac{\Delta c_{I,0}}{c_I}. \tag{23}$$

Substituting the second into the first and using $c_I \simeq w \ell_I$ (small r) yields

$$\Delta c_{I,0} \approx \frac{\ell_I (1+\varphi)}{1+\varphi/\sigma} \, \Delta w_0 + \frac{\theta h \, p'(r_{-N}) - d_I}{1+\varphi/\sigma} \, (r_0 - r_{-N}).$$

Then, their labor income response is

$$\Delta(w_0 \ell_{I,0}) \approx \ell_I (1+\varphi) \, \Delta w_0 - \frac{\varphi}{\sigma} \, \Delta c_{I,0}.$$

We now discuss the consumption plan patient agent. With $\beta_P (1+r) = 1$ and $t \geq N$, $c_{P,t} = c_{P,N}$ and $\ell_{P,t} = \ell_{P,N}$ satisfy

$$c_{P,N} + r d_{P,N-1} = w^{1+\varphi} c_{P,N}^{-\varphi/\sigma}.$$

Using $w_t u'(c_{P,t}) = v'(\ell_{P,t})$ gives $w_t \ell_{P,t} = w_t^{1+\varphi} c_{P,t}^{-\varphi/\sigma}$ and thus:

- $t = 0$: $c_{P,0} - d_{P,0} = w_0^{1+\varphi} c_{P,0}^{-\varphi/\sigma} - (1+r) d_{P,-1}$.
- $0 < t < N$: $c_{P,mid} - w^{1+\varphi} c_{P,mid}^{-\varphi/\sigma} = -\, r_0 d_{P,0}$.
- $t = N$: $c_{P,N} + r d_{P,0} = w^{1+\varphi} c_{P,N}^{-\varphi/\sigma}$.

Let $S = \sum_{t=1}^{N-1} \beta^t$. With the above sign convention,

$$u'(c_{P,0}) = (r_0 S)\, u'(c_{P,mid}) + (1+r)\beta^N\, u'(c_{P,N}).$$

Define $F(d_{P,0}, w_0, r_0) = u'(c_{P,0}) - (r_0 S) u'(c_{P,mid}) - (1+r)\beta^N u'(c_{P,N}) = 0$. Let $Y(c; w) = w^{1+\varphi} c^{-\varphi/\sigma}$. Using implicit-function derivatives (evaluated at steady state):

$$\frac{\partial c_{P,0}}{\partial d_{P,0}} = \frac{+1}{1 - Y'(c_P)}, \quad \frac{\partial c_{P,0}}{\partial w_0} = \frac{\partial Y/\partial w}{1 - Y'(c_P)}, \tag{24}$$

$$\frac{\partial c_{P,mid}}{\partial d_{P,0}} = \frac{-\, r_0}{1 - Y'(c_P)}, \quad \frac{\partial c_{P,mid}}{\partial r_0} = \frac{-\, d_P}{1 - Y'(c_P)}, \quad \frac{\partial c_{P,N}}{\partial d_{P,0}} = \frac{-\, r}{1 - Y'(c_P)}. \tag{25}$$

Hence

$$\frac{\partial F}{\partial d_{P,0}} = \frac{u''(c_P)}{1 - Y'(c_P)} \Big[1 + r_0^2 S + r(1+r)\beta^N \Big], \tag{26}$$

$$\frac{\partial F}{\partial w_0} = \frac{u''(c_P)}{1 - Y'(c_P)} \frac{\partial Y}{\partial w}, \tag{27}$$

$$\frac{\partial F}{\partial r_0} = -S\, u'(c_P) + \frac{u''(c_P)\, d_P\, r_0 S}{1 - Y'(c_P)}. \tag{28}$$

Therefore

$$\Delta d_{P,0} \approx -\frac{1}{\partial F/\partial d_{P,0}} \left[\left(\frac{\partial F}{\partial w_0} \right) \Delta w + \left(\frac{\partial F}{\partial r_0} \right) (r_0 - r_{-N}) \right].$$

Let $\Omega' = 1 + r_0^2 S + r(1+r)\beta^N$. Using $\frac{u'}{u''} = -\sigma c$,

$$\Delta d_{P,0} \approx \frac{1}{\Omega'} \left[-\left(\frac{\partial Y}{\partial w} \right) \Delta w + \big(-\, d_P\, r_0 S - \sigma c_P S \big(1 - Y'(c_P)\big) \big) (r_0 - r_{-N}) \right].$$

With $\beta(1+r) = 1$ and $S = \frac{1-\beta^{N-1}}{r}$, $\lim_{r\to 0} \Omega' = 1$ and $\lim_{r\to 0} S = N - 1$. Since $d_P r S \to 0$,

$$\Delta d_{P,0} \approx -\,(1+\varphi) w^{\varphi} c_P^{-\varphi/\sigma}\, \Delta w - (N-1)\sigma c_P \left(1 + w^{1+\varphi} \frac{\varphi}{\sigma} c_P^{-\varphi/\sigma - 1} \right) (r_0 - r_{-N}).$$

Using $\ell_P = w^{\varphi} c_P^{-\varphi/\sigma}$ and $c_P \simeq w \ell_P$ (small r) gives the intuitive form

$$\Delta d_{P,0} \approx -\,(1+\varphi) \ell_P\, \Delta w - (N-1)(\sigma c_P + w \ell_P \varphi)\, (r_0 - r_{-N}).$$

Combining the linearised $t = 0$ budget and FOC,

$$\Delta c_{P,0}\left(1 + \tfrac{w\ell_P}{c_P}\tfrac{\varphi}{\sigma}\right) \approx \ell_P(1+\varphi)\Delta w_0 + \Delta d_{P,0},$$

and substituting $\Delta d_{P,0}$ yields cancellation of the wage terms and

$$\Delta c_{P,0} \approx -(N-1)\sigma c_P\,(r_0 - r_{-N}).$$