



Refinancing cross-subsidies in the mortgage market

Jack Fisher^{a,1}, Alessandro Gavazza^{b,c,1}, Lu Liu^{d,1}, Tarun Ramadorai^{e,c,*,1}, Jagdish Tripathy^{f,1}

^a Harvard Business School, United States

^b London School of Economics, United Kingdom

^c CEPR, United Kingdom

^d The Wharton School, University of Pennsylvania, United States

^e Imperial College London, United Kingdom

^f Bank of England, United Kingdom

ARTICLE INFO

Dataset link: [Refinancing Cross-Subsidies Repository \(Original data\)](#)

JEL classification:

G21

G50

N20

R21

R31

L51

Keywords:

Mortgages

Refinancing

Cross-subsidies

Wealth inequality

Household inaction and inertia

Household finance

ABSTRACT

In household finance markets, inactive households can implicitly cross-subsidize active households who promptly respond to financial incentives. We assess the magnitude and distribution of cross-subsidies in the mortgage market. To do so, we build a structural model of household mortgage refinancing and estimate it on rich administrative data covering the stock of outstanding mortgages in the UK. We estimate sizeable cross-subsidies that flow from relatively poorer households and those located in less-wealthy areas towards richer households and those located in wealthier areas. Our work highlights how the design of household finance markets can contribute to wealth inequality.

1. Introduction

In retail financial markets, households often face complex contracts that require prompt action in response to changing financial incentives. Households who do not swiftly respond to these incentives can unwittingly provide revenues to financial firms. Conversely, such products can be beneficial to more sophisticated customers who are quicker to take appropriate action. This can result in regressive cross-subsidies in financial markets that flow from less sophisticated customers, who are often poorer and less educated, to those who are more sophisticated, wealthy, or educated. In this way, the design of household finance products can be a powerful contributor to wealth inequality.

Our paper provides a new approach to quantify such household finance cross-subsidies and to identify how they are distributed across the population. We apply the method in the setting of residential mortgage refinancing. Mortgages are the largest household liability (Campbell, 2006; Badarinza et al., 2016; Goetzmann et al., 2021), but despite their importance in their budgets, many households do not appropriately manage this debt. A crucial determinant of sound mortgage management is timely refinancing in response to financial incentives, and evidence has built up that lower-income and less-educated households fall short on this dimension (Agarwal et al., 2016; Keys et al., 2016; Andersen et al., 2020; Byrne et al., 2023).

* Corresponding author.

E-mail addresses: jfisher@hbs.edu (J. Fisher), a.gavazza@lse.ac.uk (A. Gavazza), lliu1@wharton.upenn.edu (L. Liu), t.ramadorai@imperial.ac.uk (T. Ramadorai), jagdish.tripathy@bankofengland.co.uk (J. Tripathy).

¹ Toni Whited was the editor for this article. We thank the editor, an anonymous referee, and Steffen Andersen, Elliot Anenberg, Nuno Clara, Mikhail Golosov, Gregor Matvos, Pascal Noel, Otto Toivanen, and numerous seminar and workshop participants for useful comments and discussions. Alessandro Gavazza gratefully acknowledges financial support from the European Research Council (under Consolidator Grant Agreement No. 771004). Any views expressed in this article are not meant to represent the views of the Bank of England, members of its Financial or Monetary Policy Committee, or the Financial Conduct Authority. This work was supported in part through the NYU IT High Performance Computing resources, services, and staff expertise. All errors are our own.

<https://doi.org/10.1016/j.jfineco.2024.103876>

Received 12 July 2023; Received in revised form 27 May 2024; Accepted 27 May 2024

Available online 4 June 2024

0304-405X/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

We assess the magnitude of these cross-subsidies in the UK, which is an environment uniquely suited to this purpose. Crucially, as in many other countries, the dominant UK mortgage form is a “dual-rate” mortgage contract.² At origination, borrowers fix their mortgage rate for a short initial period at a low “discounted rate”. To fully take advantage of such contracts, it is imperative to swiftly refinance when the initial fixation period ends, to avoid being rolled on to a significantly more expensive “reset rate”.³ Households who fail to do so pay these higher reset rates, and these higher payments benefit those who quickly refinance back into lower discounted rates.

Our strategy to uncover cross-subsidies in this market addresses a range of empirical and conceptual challenges. We set out to measure the difference between the total costs and benefits of a mortgage in the current dual-rate regime versus that in a single-rate counterfactual. This requires accounting for how different mortgage contracts affect decisions to participate in homeownership, mortgage payments, and any associated tangible or psychological costs incurred for contract management. To do so, we pursue a structural approach – rather than a reduced-form approach such as a counterfactual comparison based solely on observational data – for at least three reasons. First, the decision of whether or not to refinance is driven by both financial incentives (i.e., it is less beneficial to refinance mortgages with low outstanding balances or short remaining time to maturity), as well as tangible and psychological costs attached to refinancing. A reduced-form approach cannot separate the contribution of these different drivers to observed refinancing behavior, which can lead to very different conclusions about both the level and distribution of welfare under alternative mortgage contracts. Second, loan demand and optimal mortgage size may also change as households re-optimize when faced with alternative mortgage contracts. For richer predictions of such responses, we need to model household valuations for housing alongside mortgage costs and benefits. Third, unlike reduced-form approaches, structural models are well suited to evaluating non-marginal counterfactuals, such as a change in mortgage contract structure. Practically, we can evaluate households’ potentially non-linear responses to changes in incentives in the single-rate counterfactual across a range of plausible alternative interest rates.

The first step of our approach is therefore to build and estimate the structural parameters of a model of household refinancing by matching a broad set of moments from high-quality administrative data on the entire stock of mortgages in the UK. To address the conceptual challenges outlined above, we model households as heterogeneous along two dimensions. The first dimension is households’ valuation for owned housing (we model renting as an outside option), which allows us to match the loan size distribution in the stock of outstanding mortgages. The second dimension is the costs that households face at the point of refinancing. To imbue the exercise with greater realism, we model these costs as comprised of both a persistent component, as well as temporary fluctuations or shocks around this persistent component. Households in the model optimize subject to these costs, and thus this setup allows us to capture household inertia and inattention to beneficial refinancing opportunities through a realized high refinancing cost shock in any given period. Moreover, at the point of deciding whether to buy a house (and, if they do so decide, their mortgage size), households have only noisy information about their future costs of promptly refinancing, thereby capturing potentially imperfect “self-knowledge”.

Households face a dual-rate structure in the model, and all households are initially on the discounted rate. When the initial fixation period ends, households choose to refinance when the benefits of refinancing, driven mainly by the difference between the discounted and

the reset rates and loan size, outweigh the costs of refinancing whose structure is described above. Larger loans are therefore more likely to pay discounted rates: in the cross-section, these loans correspond to households with a greater valuation for housing; in the time-series, these loans correspond to households who recently originated their mortgages.

The model makes it easy to aggregate loans, thus generating intuitive expressions for aggregate mortgage loan balances on the discounted rate and on the reset rate. This allows us to estimate the model’s parameters to match a rich set of moments in granular and comprehensive data from the UK’s Financial Conduct Authority (FCA) which track the stock of all outstanding UK mortgage loans issued by all regulated financial institutions in the country at a semi-annual frequency between 2015H1 and 2017H2.⁴ The granular nature of the data means that we observe household-level mortgage refinancing behavior; and the comprehensive coverage of the entire mortgage stock allows us to compute cross-subsidies.

The second step in our approach is to then use the estimates of the structural parameters of our model to compare household outcomes in the status quo with outcomes under a counterfactual scenario. In this counterfactual, borrowers face a simpler contract that pays a single rate until mortgage maturity, with no refinancing required. We conduct these counterfactual comparisons for households across the income distribution and regions of the UK. This procedure allows us to make the invisible cross-subsidies in the current system visible, and to assess their distribution across the population. The approach can be more widely applied to uncover and analyze cross-subsidies in other market settings.

Our analysis mainly focuses on the stock in 2015H1, when the total stock of household mortgage debt in our sample equals £470 billion. The majority of this stock (65 percent) pays the discounted rate, but the remaining 35 percent pays the reset rate, with a weighted average rate spread of 52 basis points (bps). Hence, there is an appreciable spread between reset rates and discounted rates, and many households pay these different rates.

We estimate the model parameters assuming that the market is in steady state and match the data well. Our estimates imply that average refinancing costs equal £4042 among mortgage borrowers, with a standard deviation equal to £15,102; we later discuss how these money-metric estimates relate to the broader literature on mortgage refinancing.

In the counterfactual single-rate economy, households adjust both their individual loan sizes (intensive margin) as well as their participation in the housing and mortgage markets (extensive margin), in response to the different paths of mortgage rates and to the elimination of refinancing costs. Assuming that the interest rate in the single-rate economy equals the weighted average rate in our sample, aggregate mortgage debt increases by 3.55 percent relative to that in the baseline dual-rate economy.⁵ High-refinancing-cost households mostly account for the increase in aggregate debt; as they no longer pay either the punitive reset rate or refinancing costs, they are more likely to enter the market, raising the total number of mortgages in the counterfactual economy. However, the mean initial loan balance falls in the counterfactual economy by 2.85 percent of the baseline average loan size, because the composition of borrowers changes: marginal households who enter the mortgage market in the single-rate economy have smaller loan sizes than inframarginal households whose participation does not change.

The richer versions of the model, in which we allow all parameters to vary separately for 12 income groups as well as for different

² Such dual-rate contracts are also ubiquitous in other retail markets, including in credit cards, and cellphone and electricity plans (Armstrong and Vickers, 2012).

³ This feature of the UK mortgage market has prompted prominent calls for reform which highlight the implicit cross-subsidy (Miles, 2004).

⁴ In what follows, we denote the first and second observations in each year of our sample by H1 and H2 respectively to denote “half-years”.

⁵ As we later discuss, the interest rate elasticity of mortgage loans in our setup is comparable to other estimates in the literature.

geographical regions in the UK, feature considerable differences in refinancing propensities across income groups and regions. These differences presage significant variation across both the income distribution and regions in cross-subsidies that are paid and received. Most notably, we find clear evidence that higher-income households (and households in the richer South-West of the UK) would pay higher rates under the single-rate structure, and lower-income households (and households in the relatively poorer North-East and North-West of the UK) would pay lower mortgage rates under the counterfactual single-rate scenario than they do in the dual-rate economy.

The counterfactual single-rate economy also displays striking differences across groups in their endogenous adjustments to mortgage take-up and mortgage sizes. Average mortgage debt shrinks for higher-income groups and wealthier regions in response to the counterfactual single rate, since they no longer have access to the discounted rate. In contrast, the counterfactual single-interest rate economy induces lower-income households to enter the mortgage market because they expect to pay lower rates and incur no refinancing costs. This is evident in increases in the home-ownership rate, mainly driven by low-income households. This “democratization” of mortgage take-up under the counterfactual is another important indicator of the regressive effect of cross-subsidies in the dual-rate economy.

Finally, on aggregate, consumer surplus increases by 4.14 percent in the single-rate economy relative to the dual-rate economy, with larger surplus increases among low-income households, who benefit from lower interest rates and increased participation, and smaller increases among high-income households, who nonetheless benefit from the elimination of high refinancing costs. This makes our findings quite nuanced: Low-income households are penalized in the status quo dual-rate economy largely because they have smaller loan balances and thus lower incentives to refinance, whereas high-income households refinance more frequently due to greater loan balances, although they actually incur higher refinancing costs from doing so. On net, our counterfactual suggests that the removal of refinancing costs more than compensates high-income households for the increase in paid rates and associated reductions in loan balances. Overall, this leads to a positive increase in consumer surplus across all household groups.

The remainder of this section discusses related literature. Section 2 describes the administrative data and the UK institutional setting. Section 3 lays out the model and describes the computation of cross-subsidies. Section 4 discusses parameter estimation and model fit. Section 5 discusses counterfactual analysis and the flow of cross-subsidies across income groups and regions. Section 6 concludes. The Appendices include additional analyses.

1.1. Related literature

Our paper contributes to several strands of the literature. First, our work complements many empirical papers that document switching costs, inertia, and inattention in insurance and household finance markets, such as health insurance (e.g., [Handel, 2013](#)), car insurance (e.g., [Honka, 2014](#)), retirement plans (e.g., [Luco, 2019](#); [Illanes, 2016](#)), credit cards (e.g., [Ausubel, 1991](#); [Stango and Zinman, 2016](#); [Nelson, 2022](#)), pension contributions (e.g., [Choi et al., 2002](#)), and portfolio rebalancing (e.g., [Brunnermeier and Nagel, 2008](#)), among others.⁶ However, none of these papers focuses on documenting regressive cross-subsidies, though this possibility has been raised in theory (e.g., [Gabaix and Laibson, 2006](#); [Armstrong and Vickers, 2012](#)).

The papers that document inaction and frictions in mortgage refinancing (e.g., [Agarwal et al., 2016](#); [Keys et al., 2016](#); [Scharfstein and](#)

[Sunderam, 2016](#); [DeFusco and Mondragon, 2020](#); [Byrne et al., 2023](#); [Berger et al., 2023](#)) are more directly related to our work. We advance this literature, developing a novel framework that relies on counterfactual analysis to quantify the magnitudes of cross-subsidies across households arising from refinancing. This approach differs from [Andersen et al. \(2020\)](#), who model a fixed refinancing cost (“state-dependent inaction”), but with intervals of “time-dependent inaction” where refinancing is not possible, using a periodic “Calvo” shock to borrowers, and [Berger et al. \(2021\)](#), who adopt a similar approach in their analysis of US refinancing behavior. These approaches imply that the costs of refinancing are always higher than the benefits during periods of time-dependent inaction, but do not quantify these costs. In contrast, our model features a household-specific fixed refinancing cost with a time-varying shock; thus our estimation recovers the full distribution of the costs of inaction across households and over time. Apart from the differences in setting, this different modeling approach explains why the average refinancing costs that we estimate are higher than those in [Andersen et al. \(2020\)](#) and [Berger et al. \(2021\)](#).⁷

Complementary work by [Zhang \(2022\)](#) and [Berger et al. \(2023\)](#) studies the distributional effects of refinancing frictions in the US. Using a lifecycle model, [Zhang \(2022\)](#) highlights that US borrowers that do not pay upfront “points” (i.e., closing costs) to reduce mortgage rates are worst affected under the status quo. Intuitively, such borrowers pay higher interest rates for longer, thereby contributing more to lender revenues. [Berger et al. \(2023\)](#) also study the US market, and endogenize mortgage rates under simplifying assumptions on the behavior of mortgage borrowers and investors. They consider counterfactuals that include alternative contracts as well as policies that reduce borrowers’ frictions. In contrast with these papers, our paper features a more detailed model of household inaction alongside a more stripped-down supply side. Moreover, we document the redistributive consequences of cross-subsidies across income groups and regions. Finally, our focus on the UK market reduces confounds arising from unobservable moving propensities, and enables a simple calculation of refinancing inaction based on the dual-rate structure. The UK setting bears similarities with many mortgage markets outside of the US, making our work also broadly applicable to such markets.

Second, our paper is connected to a growing body of work on the design of mortgage markets around the world ([Campbell, 2013](#); [Piskorski and Seru, 2018](#)). For example, several mortgage markets also feature fixed rates for a shorter interval than the maturity of the mortgage. [Allen and Li \(2020\)](#) study borrower refinancing and lender pricing in the Canadian mortgage market; similarly, [Thiel \(2021\)](#) studies a ban on price discrimination between new and existing customers in the Dutch mortgage market. We focus on implicit cross-subsidies across borrowers in the cross-section, whereas [Allen and Li \(2020\)](#) and [Thiel \(2021\)](#) focus on intertemporal price discrimination within borrowers.⁸

Finally, our structural model provides a money-metric assessment of cross-subsidies in an important household finance market, and shows that these cross-subsidies are regressive. This showcases how the design of the financial system can contribute to inequality, connecting our work to the growing literature on wealth inequality ([Alvaredo et al., 2017](#); [Benhabib and Bisin, 2018](#); [Fagereng et al., 2020](#); [Hubmer et al., 2020](#)) and, more specifically, to that on inequality in financial

⁷ [Andersen et al. \(2020\)](#) estimate an average total psychological plus fixed refinancing cost of £1,852 in the Danish mortgage market. [Berger et al. \(2021\)](#) estimate an average refinancing cost of \$1,934 in the US mortgage market. These are lower than our estimate of the average cost across both refinancing and non-refinancing borrowers, which equals £4042.

⁸ Our paper also contributes to the growing literature on UK mortgage markets, including [Benetton \(2021\)](#), [Robles-Garcia \(2022\)](#), [Cloyne et al. \(2019\)](#), [Best et al. \(2020\)](#), [Belgibayeva et al. \(2020\)](#), [Benetton et al. \(2023\)](#), [Liu \(2022\)](#). Most of these studies focus on the flow of newly originated mortgages, whereas we focus on the stock of mortgages.

⁶ [Farrell and Klemperer \(2007\)](#) present a survey of the literature on switching costs, with a theoretical focus; [Heidhues and Köszegi \(2018\)](#) survey the literature on behavioral industrial organization; and [Gavazza and Lizzeri \(2021\)](#) the literature on markets with frictions.

wealth (Campbell et al., 2019; Greenwald et al., 2021). We also show that cross-subsidies vary across the UK, demonstrating that regional redistribution can occur directly as a result of differential efficiency in the use of financial products, a novel channel for such redistribution in mortgage markets (Hurst et al., 2016; Beraja et al., 2019).

2. Institutional setting and data moments

This section describes the UK dual-rate mortgage rate environment, and introduces the data and targeted data moments that we match when estimating the structural model.

2.1. UK mortgage market: Institutional features

Several features of the UK mortgage market make it ideal for our analysis. First, this market features posted prices at the national level, with no variation across regions, as Cloyne et al. (2019), Benetton (2021), Robles-Garcia (2022), and Benetton et al. (2023) document, among others. Borrower-specific pricing, common in US mortgage markets, is virtually non-existent in the UK.

Second (and crucial for our purposes), the vast majority of UK mortgages are issued with discounted interest rates that are fixed for a set time period, usually between one and five years (the modal fixation period is two years), depending on the contract chosen by the borrower. During the discounted period, households typically incur substantial prepayment penalties (between 3–5 percent of the loan balance), which means that households typically refinance after the end of the fixed period (Cloyne et al., 2019; Belgibayeva et al., 2020). At the end of the discounted period, the mortgage rate automatically rolls over into a higher reset rate known as the “standard variable rate”, unless borrowers refinance the mortgage into another discounted rate (for a detailed treatment of the characteristics of the UK mortgage market see Miles, 2004).⁹

This “dual-rate” structure is a feature of many mortgage systems, including Canada, Australia, India, Ireland, Germany, and Spain, meaning that our study is more broadly applicable around the world.¹⁰ We do not study the origins of this rate structure, which likely reflects mortgage lenders’ funding structures and price-discrimination strategies between active and inactive borrowers (Ellison, 2005; Gabaix and Laibson, 2006; DellaVigna and Malmendier, 2006), and we focus instead on its implications for borrowers’ refinancing. That said, in Appendix D, we follow and extend the analysis of Cloyne et al. (2019), who perform a thorough comparison between borrowers who pay the discounted and the reset rates, suggesting that the dual-rate structure does not seem designed for lenders to screen borrowers based on their default risk.¹¹

This dual-rate contract structure provides strong incentives for households to refinance at the expiration of the fixation period. UK households are free to take advantage of these incentives to refinance, as there are no further credit checks when households refinance with their existing lender, and any upfront fees can be rolled into the loan balance, meaning that liquidity constraints do not inhibit refinancing (Best et al., 2020). In Appendix F, we also rule out the possibility

that borrowers rationally stay on the reset rate to exploit the real option of timing their refinancing to coincide with interest rate declines.

Third, UK mortgages are portable, meaning that households can retain their existing mortgage contract when they move, subject to the new collateral being verified.¹² This feature stands in contrast with the US, where the lack of portability means that moving probabilities are a more important driver of prepayment/refinancing and contract choice (Stanton, 1995; Stanton and Wallace, 1998; Zhang, 2022).

2.2. Dataset description

Our primary data source is the FCA, which comprehensively tracks the stock of outstanding mortgage loans issued by all regulated financial institutions in the UK. The specific FCA dataset that we use is the Product Sales Database 007 (henceforth PSD007), which reports information about the stock of mortgage loans between June 2015 (henceforth 2015H1), and December 2017 (2017H2) at a semi-annual frequency.¹³

At each reporting date, PSD007 records the original loan amount, outstanding balance, original loan term, remaining term to maturity, current interest rate, current monthly payment, and performance status (i.e., whether the loan is in arrears and if so, for how long this has been the case) for each outstanding mortgage. The database also includes information on the property location at the most granular level in the UK (6-digit postcode), and borrower characteristics such as date of birth and the opening date for the bank account associated with the mortgage. Table A.1 in Appendix 2 provides more detailed descriptions of the main variables from the PSD007 dataset used in this paper.

The PSD007 dataset does not include information on borrower incomes. We therefore merge borrowers in the stock data with comprehensive loan-level data on borrower characteristics, including their income, shared with lenders at the time of loan origination. We also measure the current loan-to-value (LTV) ratio on each outstanding loan following a common approach in the literature, dividing the outstanding loan balance by the scaled house price at mortgage origination, using Local Authority district-level house price indices. Appendix A.2 provides details of the procedure used to merge borrower and house characteristics at loan origination to our stock data.

We further complement the PSD007 dataset with data on UK homeownership rates sourced from the Office for National Statistics (ONS) dataset *Dwelling stock by tenure*. These homeownership data allow us to measure households’ extensive margin decision of whether to buy a house and take a mortgage, or rent.

Using rich data on the stock of mortgages offers several advantages over using the flow of originations. Notably, the stock allows us to accurately capture refinancing behavior across all mortgage maturities, including mortgages originated in the past. Moreover, the structural parameters of a model estimated using the stock of mortgages rather than the flow depend less on changes in refinancing behavior or refinancing waves over short periods of time. Finally, using the mortgage stock facilitates computing average mortgage rates and aggregate lender revenues, which proves useful in our counterfactual analyses.

⁹ There is a third type of interest rate known as a tracker rate, paid on around 15 percent of all mortgages outstanding, which is a floating rate linked to the Bank of England base rate. We exclude such mortgages from our analysis because such mortgages are subject to rate fluctuations, and there are rarely transitions from the rest and discounted rate category into this category. Appendix A.1 reports further details.

¹⁰ Badarinza et al. (2018) provide information on mortgage interest-rate fixation periods across a broad set of countries and show that many large economies have similar average mortgage-rate fixation periods to the UK.

¹¹ This pricing structure with a discount for new or active customers is common in many other retail markets, including electricity, telecoms, and magazines, in which default concerns play a negligible role.

¹² Among other countries, Australia, Canada, and Germany share this feature (Lea, 2010).

¹³ Regulated financial institutions in the UK are legally required to report these details within 30 working days following the end of each calendar half-year. The group of regulated financial institutions in the UK includes deposit-taking institutions (including building societies), as well as some non-bank financial institutions. Our sample focuses on the owner–occupier segment of the mortgage borrowing population, and excludes “buy-to-let” mortgages which are issued to landlords on rental properties.

2.3. Sample selection: Borrowers ineligible to refinance

One potential challenge to our empirical analysis and our cross-subsidy calculations is to distinguish between households who can refinance, but do not do so promptly, from households who are constrained and unable to take advantage of refinancing opportunities. To address this potential confounding effect, we filter our data to remove borrowers who are potentially ineligible for refinancing—i.e., borrowers who are “involuntarily” on the reset rate, but who would potentially like to switch if they were allowed to do so.

To identify these ineligible borrowers, we follow studies by the FCA (Financial Conduct Authority, 2019, 2021) and a 2018 industry agreement that unified and codified refinancing eligibility criteria across 65 UK lenders, with a market share of around 95 percent. Passing these eligibility criteria means that a mortgage borrower can refinance into a new contract with their lender, without any affordability assessment, meaning no additional credit or income checks.¹⁴ The criteria are that the borrowers are first-charge owner-occupiers that are existing borrowers of an active lender, up to date with their payments, with a minimum remaining term of 2 years, and a minimum outstanding balance of £10,000 (Financial Conduct Authority, 2019). We broaden out these eligibility criteria to filter out borrowers potentially ineligible for refinancing, under the assumption that the 2018 agreement ratified pre-existing practice that was prevalent in the 2015H1 stock.

Table A.6 in Appendix A shows the exact proportion of loans that are potentially ineligible for refinancing using these criteria as well as broader definitions of ineligibility. Borrowers who have LTVs greater than 95 percent comprise approximately 2 percent of the sample. These borrowers may find it difficult to refinance, even though they are strictly eligible under the industry agreement if they fulfill all the other criteria. Borrowers with remaining loan balances smaller than £30,000 constitute approximately 6 percent of the total sample. And about 5 percent of loans are non-performing (in arrears, or under forbearance or possession orders). Applying these filters together removes around 14 percent of the mortgage stock in 2015H1–2017H2.

We note here that we estimated our model on both unfiltered and filtered samples. Filtering does not materially affect our main qualitative results on the regressive nature of cross-subsidies, for two main reasons. First, in the filtered sample the share of mortgage debt paying the reset rate is still quite large, and lower-income borrowers are more likely to pay the reset rate than higher-income borrowers. Second, the largest fraction of excluded borrowers are those with small loan balances, for whom refinancing benefits are small, because the refinancing benefit is proportional to the loan balance. Appendix A.4 provides more information on these filtered borrowers.¹⁵

2.4. Mortgage stock and data moments

Our analysis focuses on the 2015H1 mortgage stock, which comprises 3.59 million mortgages of borrowers eligible to refinance and

for whom we have estimates of current income; 65 percent of these mortgages pay discounted rates in this 2015H1 sample. In addition to using cross-sectional moments from 2015H1, we are also able to track mortgages across the reporting snapshots. We discuss below the key features from the cross-sectional and panel dimensions of our data.¹⁶

Table 1 shows summary statistics for variables in the filtered 2015H1 sample, most of which serve as moments that we target when we estimate our model. On average, the mean outstanding loan balances on the discounted and reset rate equal £140,647 and £112,692, respectively. The mean loan balance at origination across all loans is £142,333 (the difference is attributable mainly to amortization over time). This aggregates to a total stock of outstanding mortgage debt of £470 billion.

Households pay a weighted average discounted rate of 320 bps, and a weighted average reset rate of 372 bps, implying a difference of 52 bps, which reflects the refinancing incentive.¹⁷ Mortgages on the discounted rate have an average remaining term to maturity of 20.6 years, whereas mortgages on the reset rate have a remaining term of 16.8 years, consistent with the intuition that older loans with shorter remaining terms typically have smaller loan balances, and so provide less of an incentive to refinance back into the discounted rate. The average remaining discounted period equals 2.1 years.

Table 1 also reveals considerable cross-sectional variation in these variables, most notably in the outstanding loan balance and the remaining mortgage term. When the outstanding loan balance and/or the remaining term are low, borrowers should be less likely to refinance given the lower financial incentive from any interest rate reduction associated with doing so. Furthermore, there is considerable overlap between the distributions of these variables across rate types, which suggests that substantial variation in underlying refinancing cost may be needed to justify different behaviors observed even conditional on a given loan balance and loan term. The difference in average loan balances could be driven by both cross-sectional variation in the initial loan balance, as well as time-series variation in the remaining term of the loan, which we reflect in the model. Table 1 also shows that there is substantial variation in the age of borrowers, which lines up with variation in the remaining term of the loan, so we do not match this dimension separately in the model.

Two additional statistics not reported in Table 1 constitute important targets for our model. First, the ONS dwelling data report that 63 percent of households are homeowners in 2015. Second, rate types are highly persistent.

The panel dimension of our data allows us to track mortgages over time and highlight both within- and across-borrower variation in refinancing patterns. Figure B.2 in Appendix B shows the transition probabilities between discounted and reset rates for our sample. Households on a reset or discounted rate are much more likely to stay on the same rate type over the next 24 months than to switch—i.e., 76.2 percent of 2015H1 borrowers pay the same rate type in

¹⁴ The UK is somewhere between the US and Denmark in this respect. In the US, refinancing triggers a credit check (Keys et al., 2016), whereas in Denmark, even delinquent borrowers are able to refinance as long as there is no cash out (Andersen et al., 2020). The UK system does not trigger a credit check at the point of refinancing as long as the borrower satisfies the eligibility criteria.

¹⁵ A 2018 FCA report of the mortgage market (Financial Conduct Authority, 2019) studied 2 million reset rate mortgages using the same data that we employ and concluded that only approximately 30,000 of these mortgages were unable to switch despite being up to date with payments. The report finds that two-thirds of these mortgages were associated with an inactive, failed lender (e.g., Northern Rock, famously subject to a run during the financial crisis); and the remainder were either interest-only mortgages that were subject to changes in lending standards following the financial crisis, or in negative home equity. We expect that our filters catch many of these mortgages.

¹⁶ The main statistics of the mortgage stock are quite stable between 2015H1 and 2017H2, consistent with the idea that short-run changes have small effects on the stock of long-term debt contracts. Appendix B describes the evolution of the mortgage stock between 2015H1 and 2017H2, which exhibits two main patterns: (1) the fraction of mortgage debt paying the reset rate decreases by 2017H2, and (2) the spread between the average reset rate and the discounted rate increases over the same period. While the first pattern should decrease the magnitude of cross-subsidies across borrowers, the second one should increase them, with a small net effect.

¹⁷ Our dataset includes mortgages by two large lenders who offered to cap reset rates at 250 bps for mortgages issued up to and during the 2007–09 financial crisis. Excluding these lenders (around 900k observations) pushes up the average rate for reset rate mortgages substantially (with no change in the average rate for discounted mortgages), increasing the spread to 110 bps. We have kept mortgages by these two large lenders in our sample to provide conservative cross-subsidy estimates.

Table 1
Summary statistics for the mortgage stock in 2015H1.

	Mean	SD	Median
Current Loan Balance, Discounted Rate (£)	140,647	105,062	114,953
Current Loan Balance, Reset Rate (£)	112,692	79,684	93,916
Interest Rate, Discounted (%)	3.20	0.95	3.14
Interest Rate, Reset (%)	3.72	0.98	3.99
Original Loan Balance (£)	142,333	100,661	118,399
Original Term (Years)	23.32	7.07	25.00
Remaining Term, Discounted Rate (Years)	20.57	7.73	20.92
Remaining Term, Reset Rate (Years)	16.84	6.95	16.50
Remaining Discounted Period (Years)	2.11	1.52	1.83
Borrower Age (Years)	41.97	10.02	41.00
Observations	3,590,228		

Notes: The table above shows summary statistics of mortgages from the stock data reported in 2015H1. The sample includes mortgages in two categories, namely, those paying discounted interest rates, and those paying the reset rate ("Standard Variable Rate"). The total sample comprises around 3.59 million mortgages, of which 65 percent are discounted rate mortgages at this point in time. Table A.1 in Appendix A contains a description of the underlying variables.

Table 2
Summary statistics for the mortgage stock in 2015H1, by income quantiles.

Quantiles	Income (£)	Homeowners (%)	Balance (£)	Discounted (%)
0–10	24,604	0.50	61,522	0.65
10–20	30,504	0.61	76,141	0.64
20–30	35,631	0.64	86,894	0.64
30–40	40,701	0.68	96,752	0.64
40–50	46,198	0.72	107,326	0.64
50–60	52,687	0.75	118,920	0.64
60–70	60,974	0.80	132,885	0.65
70–80	73,154	0.82	152,206	0.66
80–85	82,349	0.84	173,235	0.66
85–90	96,616	0.86	195,018	0.67
90–95	126,236	0.91	232,354	0.68
95–100	214,486	0.96	351,530	0.69

Notes: The table above shows summary statistics of mortgages from the stock data reported in 2015H1, split by income quantiles of borrowers. Table A.1 in Appendix A contains a description of the underlying variables.

2017H1. However, some borrowers do switch over time. Switches from the lower discounted to the higher reset rate may reflect inattention and inertia, or more generally, refinancing costs, while refinancing benefits decline as the loans amortize. Against this backdrop, switches from the reset to the discounted rate suggest that these costs vary over time within households. In Appendix Figure B.3, the propensity to switch from reset to discounted rates changes non-monotonically with income, suggesting that refinancing costs are higher for higher-income borrowers. Taken together, these switching patterns suggest that a combination of household-specific fixed refinancing costs and time-varying stochastic shocks may capture the high persistence of rate types, as well as the occasional switches across rate types over time.

Table 2 shows summary statistics across quantiles of the income distribution of borrowers; this is the main dimension along which we later evaluate cross-subsidies. The third column of the table shows that the homeownership rate rises monotonically with the level of income—it equals 50 percent in the lowest-income group and attains 96 percent in the highest-income group. The remaining columns refer to borrowers. Their loan balance increases with their income, as expected. More importantly, the share of mortgages on the discounted rate (fifth column) also tends to increase with borrower income. These patterns document that lower-income borrowers are less likely to refinance than higher-income borrowers, hinting at the likely direction of cross-subsidies. Table C.1 in Appendix C provides a similar table across UK regions, confirming that borrowers in higher-income regions are more likely to pay discounted rates than those in lower-income regions.

In Appendix D, we look for empirical evidence on the two main reasons that could explain why higher-income households are more likely to refinance and pay the discounted rate than lower-income

households, namely: (1) higher-income borrowers tend to have higher loan balances, which increases the benefit of refinancing, and (2) higher-income borrowers could have lower unobserved refinancing costs.

More specifically, Table D.1 provides a descriptive analysis of the observables associated with prompt refinancing behavior. As expected, higher-income borrowers with larger loan balances are more likely to refinance promptly. However, emphasizing the complexity of this issue, Table D.2 shows that in the stock of mortgages, for loans with comparable sizes, borrowers paying the reset rates tend to have larger incomes. This last observation hints at the possibility that higher-income borrowers pay higher costs to refinance. Taken together, these stylized facts further motivate the need to carefully model variation in both the benefits and costs of refinancing. This is to ensure that we correctly estimate cross-subsidies resulting from refinancing behavior across income groups.

Overall, the summary statistics reported in Tables 1 and 2 document that the UK mortgage market comprises a mix of borrowers paying discounted rates and reset rates. While mortgages on discounted rates constitute the main share, a large fraction of the outstanding mortgage stock (35 percent of all loans) pays the reset rate, at an average spread of 52 bps over the discounted rate. In the next Section, we develop a model that we map to these data features in our structural estimation. We use the model to quantitatively assess the benefits and costs of refinancing, as well as the magnitude of the cross-subsidies that the dual-rate structure embeds and how they are distributed across borrowers.

3. Model

We model a mortgage market in which a measure M of households enters in each period. When they enter the market, households choose whether to buy a house with a mortgage or rent a property. If a household i chooses to buy, they pay a one-time origination cost k_i^0 and obtain per-period flow utility from their house equal to $v_i h_i^\alpha - m(l_i, r, T)$, where v_i is household i 's per-period valuation for housing, h_i is the size of the house that the household i chooses, and $0 < \alpha < 1$ is a parameter governing the utility from housing. $m(l_i, r, s)$ is the per-period mortgage payment of a household with a mortgage with current loan balance l_i , interest rate r , and remaining term s , which follows from the amortization of the loan:

$$m(l_i, r, s) = l_i \frac{r(1+r)^s}{(1+r)^s - 1}. \quad (1)$$

Renting a property yields per-period utility $\bar{u}$, which we assume is common to all households and fixed over time. All households discount the future at the common rate β .

Mortgages are long-term contracts for T periods that pay a discounted rate r_d for an initial time interval T_d , and subsequently pay

a reset rate $R > r_d$ following this interval, unless the household refinances back into the discounted rate. To simplify and facilitate evaluating counterfactuals, we take both rates as given constant values. We also assume that T/T_d is a (positive) integer and we normalize by the length of this initial fixation period, treating it as a single time unit, i.e., we assume $T_d = 1$ and $T = 15$, and all rates are computed over the period T_d . Moreover, we assume that households do not change their loan balance (i.e., we rule out “cash-out refinancing”), and rule out maturity extensions (i.e., households in the model do not change the maturity of their loan at the point of refinancing). Households receive the loan amount at time $t = 0$, but make the first repayment at $t = 1$, which is also the first refinancing period. Hence, the loan balance of a mortgage with interest rate r evolves over time as follows:

$$l_{i,t+1}(r, l_{i,t}) = l_{i,t}(1 + r) - m(l_{i,t}, r, s). \quad (2)$$

Mortgages are fully repaid after T periods. Thus, each household makes T payments over the life of the loan, the same as the duration of the mortgage contract.

At time $t = 0$, if they choose to buy a house, households choose the size of their mortgage loan $l_{i,0}$ to finance their house h_i , where $\omega_i = h_i/l_{i,0}$ denotes the inverse of the loan-to-value at origination. In each subsequent period, households can refinance their mortgage at the discounted rate r_d ; to do so, they have to pay refinancing costs equal to $k_{i,t} = k_i \varepsilon_{i,t}$, where k_i is a persistent component of the refinancing cost for household i and $\varepsilon_{i,t}$ is a transitory component. We assume that $\varepsilon_{i,t}$ is a non-negative random variable, independent and identically distributed across households and over time, with mean equal to one, with cumulative distribution function $F(\varepsilon_{i,t})$ and density $f(\varepsilon_{i,t})$. Hence, each household's average refinancing costs equal their persistent component of refinancing costs, i.e., $E(k_{i,t}) = k_i$.

Households are heterogeneous in their per-period valuation for housing v_i (to capture the heterogeneity of initially chosen loan sizes seen in the data) and in their persistent component k_i of the cost of refinancing (to capture the household heterogeneity in refinancing for a given loan balance). We assume that, at the time of originating a mortgage, households perfectly know their valuation for owned housing v_i , but only receive a signal of their persistent component k_i of refinancing costs and thus of their average refinancing costs over time. Specifically, we assume that k_i is correlated with the origination cost k_i^o according to $k_i = k_i^o \varepsilon_{i,0}$, where $\varepsilon_{i,0}$ is a non-negative random variable that is realized after the origination of the mortgage and before the first refinancing opportunity. Thus, the precision of the signal negatively depends on the variance of $\varepsilon_{i,0}$. We assume that $\varepsilon_{i,0}$ is independent and identically distributed across households, with mean equal to one, with cumulative distribution function $F_0(\varepsilon_{i,0})$ and density $f_0(\varepsilon_{i,0})$.

Valuations and origination costs are distributed according to the cumulative joint distribution function $G_o(v_i, k_i^o)$ with density $g_o(v_i, k_i^o)$. Hence, the joint density of valuations and persistent refinancing costs equals $g(v_i, k_i) = \int_0^{+\infty} g_o(v_i, k_i/\varepsilon_{i,0}) f_0(\varepsilon_{i,0}) \frac{1}{\varepsilon_{i,0}} d\varepsilon_{i,0}$.

Intuitively, in the model, households learn about their persistent ongoing mortgage refinancing costs from the costs/hassle that they experience during the process of mortgage origination. The variance of $\varepsilon_{i,0}$ determines how informative is the origination process about ongoing mortgage refinancing costs. If this variance is zero, the initial process of mortgage origination perfectly informs households about the future persistent cost of refinancing. Alternatively, if this variance is high, households learn little about the future process of refinancing from their experience during origination, since k_i is likely quite different from k_i^o .

We now solve the model to determine two household choices: (1) whether or not to refinance at each opportunity; and (2) the optimal size of the initial loan $l_{i,0}^*(v_i, k_i^o)$.

3.1. Optimal refinancing

Households refinance when their refinancing costs are below a threshold that depends on their loan size. Hence, households with larger loans are more likely to refinance. Similarly, because the loan is amortizing, each household's incentives to refinance decline over time as the outstanding balance decreases; notably, some households (almost) always refinance because they have a low value of the persistent component k_i of the cost of refinancing.

We solve for the optimal refinancing path by backward induction. Consider period T , which is the last refinancing period, and households with a beginning-of-period (i.e., before making a payment) loan balance l_i (we suppress the subscript t for simplicity). Such households refinance if their refinancing cost $k_{i,T}$ is below the benefit of refinancing $k_i^*(T)$:

$$\begin{aligned} k_i^*(T) &= m(l_i, R, 1) - m(l_i, r_d, 1) \\ &= l_i(R - r_d). \end{aligned}$$

The benefit of refinancing depends on the difference between the interest rates $R - r_d$, as well as on the loan balance l_i .

We can define the expected (i.e., prior to the realization of the transitory component $\varepsilon_{i,t}$) value function $V_T(k_i, l_i)$ of a household with persistent cost k_i as the expected payment:

$$\begin{aligned} V_T(k_i, l_i) &= \int_0^{+\infty} \min(m(l_i, r_d, 1) + k_i \varepsilon_{i,T}, m(l_i, R, 1)) dF(\varepsilon_{i,T}) \\ &= \int_0^{k_i^*(T)/k_i} (m(l_i, r_d, 1) + k_i \varepsilon_{i,T}) dF(\varepsilon_{i,T}) \\ &\quad + \int_{k_i^*(T)/k_i}^{+\infty} m(l_i, R, 1) dF(\varepsilon_{i,T}). \end{aligned} \quad (3)$$

where $k_i^*(T)/k_i$ is the cutoff point in the distribution of the transitory component $\varepsilon_{i,t}$ that determines household refinancing.¹⁸

Similarly, in the previous period $T - 1$, households' expected value function equals the discounted sum of expected future payments:

$$\begin{aligned} V_{T-1}(k_i, l_i) &= \int_0^{+\infty} \min(m(l_i, r_d, 2) + k_i \varepsilon_{i,T-1} + \beta V_T(k_i, l_i(1 + r_d) - m(l_i, r_d, 2)), \dots \\ &\quad m(l_i, R, 2) + \beta V_T(k_i, l_i(1 + R) - m(l_i, R, 2))) dF(\varepsilon_{i,T-1}) \\ &= \int_0^{k_i^*(T-1)/k_i} (m(l_i, r_d, 2) + k_i \varepsilon_{i,T-1} \\ &\quad + \beta V_T(k_i, l_i(1 + r_d) - m(l_i, r_d, 2))) dF(\varepsilon_{i,T-1}) + \\ &\quad \int_{k_i^*(T-1)/k_i}^{+\infty} (m(l_i, R, 2) + \beta V_T(k_i, l_i(1 + R) - m(l_i, R, 2))) dF(\varepsilon_{i,T-1}), \end{aligned}$$

where

$$\begin{aligned} k_i^*(T-1) &= m(l_i, R, 2) + \beta V_T(k_i, l_i(1 + R) - m(l_i, R, 2)) + \\ &\quad - m(l_i, r_d, 2) - \beta V_T(k_i, l_i(1 + r_d) - m(l_i, r_d, 2)) \end{aligned}$$

defines the monetary benefits of refinancing, such that households with $k_{i,t} \leq k_i^*(T-1)$ refinance, and households with $k_{i,t} > k_i^*(T-1)$ do not.

In a generic period t , the expected value function equals:

$$\begin{aligned} V_t(k_i, l_i) &= \int_0^{+\infty} \min(m(l_i, r_d, T-t+1) + k_i \varepsilon_{i,t} \\ &\quad + \beta V_{t+1}(k_i, l_i(1 + r_d) - m(l_i, r_d, T-t+1)), \dots \\ &\quad m(l_i, R, T-t+1) + \beta V_{t+1}(k_i, l_i(1 + R) - m(l_i, R, T-t+1))) dF(\varepsilon_{i,t}), \end{aligned}$$

and the benefits $k_i^*(t)$ determine the cutoff point in the cost distribution that characterizes household refinancing decisions.

Therefore, we can describe the optimal refinancing policy as follows:

$$r(l_i, k_{i,t}) = \begin{cases} r_d & \text{if } k_{i,t} \leq k_i^*(t) \\ R & \text{otherwise.} \end{cases} \quad (4)$$

¹⁸ When $\varepsilon_{i,T}$ equals $k_i^*(T)/k_i$, then $k_i \cdot \varepsilon_{i,T}$ equals $k_i^*(T)$ and the cost is exactly equal to the benefit of refinancing.

Hence, households with a lower persistent component k_i are more likely to refinance and pay the discounted rate r_d than households with a higher k_i . Moreover, the refinancing behavior of each household varies over time depending on the realization of the transitory shock $\varepsilon_{i,t}$. Because the transitory shock $\varepsilon_{i,t}$ is multiplicative, its realization has a smaller effect on the refinancing activity of borrowers with low k_i , and a larger effect on that of borrowers with high k_i .

3.2. Optimal loan size

Households choose the loan size that maximizes their value function at origination, given their valuation for housing v_i and origination cost k_i^o . The value at origination equals:

$$W_0(v_i, k_i^o) = \max_{l_{i,0}} \sum_{t=0}^{+\infty} \beta^t v_i(\omega_i l_{i,0})^\alpha - k_i^o - \beta \int_0^{+\infty} V_1(k_i^o \varepsilon_{i,0}, l_{i,0}) dF_0(\varepsilon_{i,0}), \quad (5)$$

where the loan-to-value at origination equals $l_{i,0}/h_i = 1/\omega_i$ and k_i^o is the mortgage origination cost described above. Households do not know the exact value of their future refinancing cost and thus they form their expectations based on the available signal, which is their origination cost.

The optimal loan size $l_{i,0}^*(v_i, k_i^o)$ satisfies the first-order condition

$$\frac{\alpha \omega_i v_i (\omega_i l_{i,0}^*)^{\alpha-1}}{1-\beta} - \beta \frac{\partial}{\partial l_{i,0}} \int_0^{+\infty} V_1(k_i^o \varepsilon_{i,0}, l_{i,0}^*) dF_0(\varepsilon_{i,0}) = 0. \quad (6)$$

Hence, the optimal loan size depends directly on the household valuation for housing v_i , and indirectly on the origination costs k_i^o , because it is correlated with the expected future mortgage payments through the optimal refinancing policy $r(l_{i,t-1}, k_{i,t})$ described above. The refinancing policy in (4) highlights that refinancing costs determine the extent to which households make mortgage payments at the higher reset rate rather than at the lower discounted rate. This is because obtaining the cheaper discounted rate in a greater number of periods requires incurring the refinancing cost $k_{i,t}$ across a greater expected number of refinancing opportunities.

Given the optimal loan size, we define $v_i^*(k_i^o)$ as the valuation for housing of a household that is indifferent between buying a house and getting a mortgage or renting a property:

$$W_0(v_i^*, k_i^o) = \frac{\bar{u}}{1-\beta}, \quad (7)$$

where $\bar{u}$ is a per-period utility of the outside rental option. This extensive-margin condition determines whether or not households enter the housing market rather than rent: households with a high valuation v_i and a low cost k_i^o enter the housing and mortgage market.

The precision of information that households have about their future refinancing costs plays into both optimal loan size (the intensive margin described in Eq. (6)) and whether or not households enter the housing market in the first place (the extensive margin described in Eq. (7)). On the intensive margin, a higher k_i generates an incentive to scale back the size of the initial loan, and on the extensive margin, a higher k_i may be a deterrent to entering the mortgage market in the first place. Conditional on the other parameters including their housing valuation v_i , the extent to which this effect operates depends on the variance of $\varepsilon_{i,0}$. If this variance is small, households choose an initial loan size that is strongly correlated with the origination costs k_i^o and thus with the persistent component k_i of refinancing costs. If the variance of $\varepsilon_{i,0}$ is larger, households have less precise information at origination to evaluate their future mortgage costs. Hence, their initial loan size will be weakly correlated with the cost k_i . The variance of the transitory component $\varepsilon_{i,t}$ of refinancing costs similarly affects households' optimal initial loan size, because a larger variance of $\varepsilon_{i,t}$ makes it more difficult for households to predict their refinancing activity, and thus the rates of their future mortgage payments.

Eq. (7) shows that origination costs k_i^o also capture any household constraints to becoming homeowners. Once again, the precision of

households' information at origination, captured by the variance of $\varepsilon_{i,0}$, critically affects this adjustment. This condition will play an important role in our counterfactual analysis as it determines how initial homeownership and mortgage take-up change. We return to these issues in greater detail when evaluating counterfactuals.

3.3. Aggregation: Mortgage stocks in steady-state

We calculate the total stock of mortgages that pay the discounted rate and the reset rate, assuming that the economy is in steady state.

It is useful in this calculation to recursively define the endogenous cumulative distribution function $H_t(\cdot)$ and its associated density $h_t(\cdot)$ of loan balances t periods after origination, given the evolution of the loan balances in (2), and the refinancing policy described in (4). This distribution evolves as follows:

$$H_0(z) = \iint_{\{(v_i, k_i^o) : v_i \geq v_i^*(k_i^o) \cap l_{i,0}^*(v_i, k_i^o) \leq z\}} g_0(v_i, k_i^o) dv_i dk_i^o,$$

$$H_t(z) = \int_{\{l_{i,t-1} : l_{i,t}(r(l_{i,t-1}, k_{i,t})) \leq z\}} h_{t-1}(l_{i,t-1}) dl_{i,t-1}.$$

We next define three groups (0, 1, 2) of mortgages. Group 0 comprises the mortgages of households who took a mortgage of initial size $l_{i,0}^*(v_i, k_i^o)$ and are on their initial discount period. Their aggregate number $N_0(r_d)$ and aggregate balance $Q_0(r_d)$ of mortgages equal:

$$N_0(r_d) = M \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g(v_i, k_i^o) dv_i dk_i^o, \quad (8)$$

$$Q_0(r_d) = N_0(r_d) \int_0^{+\infty} z h_0(z) dz = M \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} l_{i,0}^*(v_i, k_i^o) g_0(v_i, k_i^o) dv_i dk_i^o. \quad (9)$$

To gain an intuition for Eq. (8), recall that a mass M of households enters the market in each time period. (Discounted) mortgage take-up among these households is determined by whether or not they satisfy the extensive margin condition $v_i \geq v_i^*(k_i^o)$, with the outer integral integrating across the k_i^o distribution. Eq. (9) follows by weighting these mortgages by their initial loan sizes.

The second group comprises the mortgages of all households who refinanced and pay the discounted rate. In each period $t \in \{1, \dots, T-1\}$, the number $N_{1,t}(r_d)$ of mortgages in this group equals:

$$N_{1,t}(r_d) = N_0(r_d) \int_{\{l_{i,t} : r(l_{i,t}, k_{i,t}) = r_d\}} h_t(l_{i,t}) dl_{i,t} \quad (10)$$

Eq. (10) combines all borrowers who have refinancing costs $k_{i,t}$ lower than the benefits $k_i^*(t+1)$, and thus have policy functions $r(l_{i,t}, k_{i,t}) = r_d$. Thus, the aggregate number $N_1(r_d)$ of mortgages of this group equals:

$$N_1(r_d) = \sum_{t=1}^{T-1} N_{1,t}(r_d). \quad (11)$$

The aggregate balance Q_1 of this group is the sum of the balances $Q_{1,t}(r_d)$ of the different cohorts who pay the discounted rate r_d : $Q_1(r_d) = \sum_{t=1}^{T-1} Q_{1,t}(r_d)$, where $Q_{1,t}(r_d)$ evolves as follows:

$$Q_{1,t}(r_d) = N_0(r_d) \int_{\{l_{i,t} : r(l_{i,t}, k_{i,t}) = r_d\}} l_{i,t} h_t(l_{i,t}) dl_{i,t}.$$

The third group comprises the mortgages of all households who did not refinance, and pay the reset rate. In each period $t \in \{1, \dots, T-1\}$, the number $N_{2,t}(R)$ of mortgages in this group equals:

$$N_{2,t}(R) = N_0(r_d) \int_{\{l_{i,t} : r(l_{i,t}, k_{i,t}) = R\}} h_t(l_{i,t}) dl_{i,t}, \quad (12)$$

which is the set of borrowers who have refinancing costs above the benefits $k_i^*(t+1)$, and thus have policy functions $r(l_{i,t}, k_{i,t}) = R$. Thus, the aggregate number of households who pay the reset rate equals

$$N_2(R) = \sum_{t=1}^{T-1} N_{2,t}(R). \quad (13)$$

The aggregate balance $Q_2(R)$ of this group is the sum of the balances of the different cohorts who pay the reset rate R : $Q_2(R) = \sum_{t=2}^T Q_{2,t}(R)$, where $Q_{2,t}(R)$ evolves as follows:

$$Q_{2,t}(R) = N_0(r_d) \int_{\{l_{i,t}: r(l_{i,t}, k_{i,t})=R\}} l_{i,t} h_t(l_{i,t}) dl_{i,t}.$$

The above expressions can be directly mapped to the empirically observed stock of mortgages in each category, under the assumption that the market is in steady state.

3.4. Cross-subsidy

To calculate the cross-subsidy across different households, we consider a benchmark case in which all mortgages have a constant interest rate r_c for their entire duration. In Section 5, we consider several values of this constant interest rate.

Under the constant interest rate r_c , households do not need to refinance and their mortgage payments are constant over time. Hence, their optimal loan size $l_{i,0}^*(v_i, k_i^o)$ maximizes the value function at origination (5) evaluated at $\epsilon_{i,t} = 0$ for all $t > 0$, with a constant payment stream $m(l_{i,0}, r_c, T)$. The expression for optimal loan size simplifies to:

$$\begin{aligned} l_{i,0}^*(v_i, k_i^o) &= \frac{1}{\omega_i} \left(\frac{1-\beta}{\alpha \omega_i v_i} \left(\sum_{t=1}^T \beta^t \frac{\partial m(l_{i,0}, r_c, T)}{\partial l_{i,0}} \right) \right)^{\frac{1}{\alpha-1}} \\ &= \frac{1}{\omega_i} \left(\frac{\beta(1-\beta^T)}{\alpha \omega_i v_i} \frac{r_c(1+r_c)^T}{(1+r_c)^T - 1} \right)^{\frac{1}{\alpha-1}}. \end{aligned} \quad (14)$$

The aggregate number $N(r_c)$ and aggregate balance $Q(r_c)$ of mortgages then equal:

$$\begin{aligned} N(r_c) &= MT \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g_o(v_i, k_i^o) dv_i dk_i^o, \\ Q(r_c) &= M \sum_{t=1}^T \gamma_{r_c}(t-1) \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} l_{i,0}^*(v_i, k_i^o) g_o(v_i, k_i^o) dv_i dk_i^o, \end{aligned}$$

where we define

$$\gamma_{r_c}(t-1) = \frac{l_{i,t}(r_c, l_{i,0})}{l_{i,0}} = \frac{(1+r_c)^T - (1+r_c)^t}{(1+r_c)^T - 1},$$

as the beginning-of-period- t share of the initial loan still to be repaid, and $v_i^*(k_i^o)$ is the valuation of a household that is indifferent between buying a house and getting a mortgage, or renting a property in this constant rate scenario. Thus, households still face the origination cost k_i^o that, as we recount above, includes additional household constraints to homeownership, but no subsequent refinancing costs, i.e., $k_{i,t} = 0$ for $t > 0$.

Based on the estimated parameters, the observed discounted rate r_d and reset rate R , and this counterfactual constant rate r_c , we can calculate the differences in mortgage market outcomes between the current and counterfactual scenarios for each household (v_i, k_i^o) . These outcomes include differences in loan sizes and mortgage payments between current and counterfactual scenarios. They also include a measure of the lifetime cross-subsidy paid or received by the household, measured as the household-level reduction or increase (when comparing current and counterfactual scenarios) in the interest rate. These household-level calculations can be aggregated up at the group level using the baseline model, or indeed, using an extended version of the model in which we estimate group-specific parameters. We describe this extended model next.

3.5. Multiple groups

The richness of our data allows us to calculate subsidies across different groups based on observable demographic characteristics. We focus on two specific household groupings: The first groups households by income deciles, and the second by UK regions.

Understanding variation in the extent of cross-subsidies paid or received along the income distribution helps us to understand how the design of the financial system contributes to the inequality of financial wealth, to the extent that wealth and income are correlated. We also look at the extent of regional variation in mortgage cross-subsidies given the importance of regional re-distribution through the mortgage market.

We extend the model to accommodate and interpret such heterogeneity. We index different groups based on observable characteristics by $j = 1, \dots, J$. Let M_j and $G_{oj}(v_i, k_i^o)$ be the measure and the cumulative distribution function of household housing preferences v_i and origination costs k_i^o in group j , respectively. Following the analysis of previous subsections, we can define the variables $N_{0,j}(r_d), Q_{0,j}(r_d), \dots, Q_{2,j}(R)$ for each group j , and proceed with our counterfactual comparisons as before using this extended model.

We next turn to acquire quantitative estimates of the model's parameters and an assessment of the model-implied cross-subsidy by mapping the model to the data.

4. Quantitative analysis

The model does not admit an analytic solution for all endogenous outcomes. As a result, we choose the parameters that best match moments of the data with the corresponding moments computed from the numerical solution of the model in steady state. We then study the quantitative implications of the model evaluated at the estimated parameters.

4.1. Estimation

We fix a subset of parameters, often reading them directly from the data, and we estimate the remaining parameters of the model to best match key moments of the mortgage data.

Specifically, we set the unit of time in the model to be $T_d = 2$ years, which is the modal initial fixation period in the UK mortgage market over the sample period; we then set the mortgage maturity at $T = 15$ periods, to give us the modal mortgage origination maturity of 30 years. We set the discount rate at $\beta = 0.95^2 = 0.9025$ to correspond to our assumption on the unit of time.

We read the annual interest rates on discounted and reset rate mortgages directly from the underlying data, using value-weighted averages of the corresponding rates in the 2015H1 sample, and compound them to correspond to $T_d = 2$ years. Annual average discounted and reset rates equal 320 bps and 372 bps in our sample, meaning $r_d = 650$ bps and $R = 759$ bps over two years.

We set the loan-to-value ratio at origination common across households at 80 percent, close to the modal value in our data, so $\omega = 1.25$.

We read market size M from the data, as follows. The total number of mortgages in the model equals:

$$N_0(r_d) + N_1(r_d) + N_2(R) = MT \int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g_o(v_i, k_i^o) dv_i dk_i^o. \quad (15)$$

Hence, we compute the market size M by dividing the total number of mortgages $N_0(r_d) + N_1(r_d) + N_2(R)$ by their maturity T and by the share of households who own a property $\int_{-\infty}^{+\infty} \int_{v_i^*(k_i^o)}^{+\infty} g_o(v_i, k_i^o) dv_i dk_i^o$.

We estimate all remaining parameters by applying some assumptions about distributions. We assume that households' valuation v_i follows a lognormal distribution, i.e., $\log(v_i)$ follows a normal distribution with mean μ_v and standard deviation σ_v . We assume that the origination cost k_i^o follows a mixture distribution of two lognormal distributions, which allows for a bimodal distribution of origination costs. We model the persistent component of refinancing costs as $k_i = k_i^o \epsilon_{i,0}$, which is correlated with the origination cost. As a result, the distribution of refinancing costs could be bimodal as well, with some households with low refinancing costs and others with high refinancing

costs, and heterogeneity within these two household groups. With probability η , $\log(k_i^o)$ follows a normal distribution with mean μ_{k1} and standard deviation σ_{k1} ; with probability $1 - \eta$, $\log(k_i^o)$ follows a normal distribution with mean μ_{k2} and standard deviation σ_{k2} .

Without loss of generality, we denote type-1 as households whose average cost is lower than the average cost of type-2 households—i.e., $\exp\left(\mu_{k1} + \frac{\sigma_{k1}^2}{2}\right) \leq \exp\left(\mu_{k2} + \frac{\sigma_{k2}^2}{2}\right)$. We do not restrict the variances of these distributions, and thus the type-2 distribution does not necessarily first-order stochastically dominate the type-1 distribution. Clearly, the assumption of lognormality implies that some type-1 households have higher costs than some type-2 households. We set $\eta = 0.5$, and the correlation between v_i and k_i^o to zero, because the empirical moments that we employ in the estimation do not allow us to separately identify these parameters, as we explain in more detail below.

We further assume that $\varepsilon_{i,t}$ follows a lognormal distribution with parameters μ_ε and σ_ε . We set the mean of $\varepsilon_{i,t}$ equal to one, meaning that $E(k_{i,t}) = k_i$, and hence $\mu_\varepsilon = -\sigma_\varepsilon^2/2$.

We assume that $\varepsilon_{i,0}$ is governed by the same distribution as $\varepsilon_{i,t}$. In other words, if shocks to household attention/distraction (i.e., temporary refinancing costs shocks) are drawn from a high-variance distribution, this also means that households learn less from their mortgage origination costs about their persistent refinancing costs and vice versa. Hence, households obtain a noisy signal of the persistent component of their refinancing costs (and thus of their average refinancing costs) at the point of mortgage origination. Therefore, their loan size and participation decisions will exhibit a weaker correlation with their refinancing costs than in the case in which households know their refinancing costs.

Finally, our estimation recovers the parameter α of the utility function and the level of the outside option $\bar{u}$.

We search for the vector of 9 parameters $\psi = (\mu_v, \sigma_v, \mu_{k1}, \sigma_{k1}, \mu_{k2}, \sigma_{k2}, \sigma_\varepsilon, \alpha, \bar{u})$ that minimizes the distance between selected moments in the data and the corresponding moments of the model. More specifically, for each combination of these unknown parameters, we solve the model shown in Section 3 to find households' optimal policies, characterized by their choice between buying a house with a mortgage or renting a property, and, if they choose to participate in the mortgage market, their mortgage loans at origination $l_{i,0}^*(v_i, k_i^o)$ and their optimal sequence of refinancing. Based on these household policies, we compute the following aggregate moments:

1. the average loan balance for mortgages on the discounted rate;
2. the standard deviation of the loan balance of mortgages on the discounted rate;
3. the average loan balance for mortgages on the reset rate;
4. the standard deviation of the loan balance of mortgages on the reset rate;
5. the average remaining maturity of mortgages on the discounted rate;
6. the standard deviation of the remaining maturity of mortgages on the discounted rate;
7. the average remaining maturity of mortgages on the reset rate;
8. the standard deviation of the remaining maturity of mortgages on the reset rate;
- 9–14. the shares of mortgages on the discounted rate for the following partition of the loan balance distribution: [0–5] percentile, (5–25] percentile, (25–50] percentile, (50–75] percentile, (75–95] percentile, and (95–100] percentile;
15. the share of mortgages on the reset rate in 2015H1 that paid the discounted rate in 2017H1;
16. the share of homeowners, i.e., the fraction of households that enter the housing market and choose to purchase a house and take on a mortgage.

Section 2 outlines several filters that we apply to the data. One of these filters is that the outstanding mortgage balance exceeds £30,000, and for consistency, we apply the same filter when computing simulated moments 1 to 15 in the model.

The minimum-distance estimator chooses the parameters that minimize the criterion function:

$$(\mathbf{m}(\psi) - \mathbf{m}_S)' \Omega (\mathbf{m}(\psi) - \mathbf{m}_S),$$

where $\mathbf{m}(\psi)$ is the vector of moments computed from the model at the parameter vector ψ , $\mathbf{m}_S$ is the vector of corresponding sample moments, and Ω is the inverse of the covariance matrix of our sample moments, which is the optimal weighting matrix (Hansen, 1982) that we compute by stacking influence functions (Erickson and Whited, 2002).

We estimate three versions of the model. In the baseline version, we pool together all mortgages in our data and assume that all households can be characterized by a single distribution $G_o(v_i, k_i^o)$, as well as common σ_ε , α , and $\bar{u}$ parameters. This entails estimating 9 parameters using the 16 moments listed above.

We also pursue the estimation in versions with richer borrower heterogeneity. The first one estimates the model separately for different income groups, and the second one estimates the model separately for different geographic areas of the UK. In each case, we set group-specific market sizes M_j and to estimate group-specific parameters of the distributions $G_{oj}(v_i, k_i^o)$, as well as σ_{ε_j} , α_j , and $\bar{u}_j$ for each group j (denoting either income groups or geographical areas). This gives us additional flexibility to capture heterogeneity across groups in preferences, costs, and thus refinancing activities. Of course, when we estimate these parameters, we do so using an expanded set of group-specific moments in each case.

We consider 12 income groups based on the following percentiles of the distribution of reported incomes in the PSD: 0–10, 10–20, 20–30, 30–40, 40–50, 50–60, 60–70, 70–80, 80–85, 85–90, 90–95, and 95–100. We also consider 12 broad regions and devolved administrations of the UK, namely North-East, North-West, Yorkshire and the Humber, East Midlands, West Midlands, East of England, Greater London, South East, South West, Wales, Scotland, and Northern Ireland.¹⁹ Hence, in each case, we estimate a total of 108 parameters (9 parameters for each of the 12 groups) using a total of 192 moments (16 moments listed above for each of the 12 groups).

4.2. Sources of identification

The model is highly nonlinear, so (almost) all parameters affect all outcomes. That said, the identification of certain parameters does rely more heavily on particular data moments.

More specifically, moments characterizing the distributions of loan sizes on the discounted and the reset rate, those characterizing the distributions of remaining maturities in each mortgage category, and the shares of mortgages in the two categories together identify the parameter α , and the parameters of the distributions of household preferences v_i and the persistent component of costs k_i . Notably, households' initial loan amounts – and, thus over time, their loan balances – depend on their housing preferences v_i , as well as their expected refinancing costs which on average equal k_i . Moreover, for every mortgage, the parameter α affects the sensitivity of the initial loan size to expected mortgage payments, and thus to interest rates, as Eqs. (6) and (14) show.

If the cost $k_{i,t}$ was prohibitively high for all borrowers, almost all mortgages would be on the reset rate, and conversely, if $k_{i,t}$ was extremely low for all borrowers, all mortgages would be on the discounted rate. Hence, the shares of mortgages on the reset rate are

¹⁹ These are the 12 NUTS-1 regions of the UK, where NUTS stands for Nomenclature of Territorial Units for Statistics.

informative about the parameters of the distribution of the refinancing cost $k_{i,t}$ and its components.

Given a value of $k_{i,t}$, borrowers have stronger financial incentives to refinance if they have a large loan balance, meaning that the share of mortgages on the discounted rate should be increasing in the loan balance. The rate of change of the share of mortgages on each rate as loan size changes is informative about the heterogeneity in $k_{i,t}$. The increase is fast if the heterogeneity across households is small, whereas it is slow if the heterogeneity is large. Our assumption that k_i follows a mixture distribution allows us to flexibly capture different rates of increase in the share of mortgages on the discounted rate at different percentiles of the loan balance distribution. This means that the change in the share in the two categories of mortgages at different levels of the loan balance contributes to the identification of the refinancing cost heterogeneity parameters σ_{k1} and σ_{k2} of the mixing distribution. Because we allow the support of the two k_i distributions to overlap, it is difficult to separately identify the mixing probability η ; therefore, we set it to $\eta = 0.5$. Moreover, identifying any correlation between v_i and k_i would require rich within-borrower moments; our mapping to the stock means that all our moments are cross-sectional (except for the share of mortgages on the reset rate that later pays the discounted rate, moment 15 above), so we set this correlation to zero by assumption.

The share of mortgages that transition from paying the reset rate to paying the discounted rate is informative about the within-borrower heterogeneity in refinancing costs, and thus identifies the parameter σ_ϵ . If refinancing costs were fixed over time for each borrower, because loan balances decline over time, borrowers' optimal refinancing policy would be deterministic: it would be characterized by a borrower-specific cutoff date $T_{max}(v_i, k_i)$, such that a (v_i, k_i) -borrower always refinances before $T_{max}(v_i, k_i)$ and never does after $T_{max}(v_i, k_i)$. Transitions from the reset rate to the discounted rate violate this deterministic refinancing policy, and thus identify the transitory component of refinancing costs governed by σ_ϵ .

Moreover, our data does not allow us to identify households' beliefs at origination about their future refinancing costs, captured by the variance of $\epsilon_{i,0}$. Hence, we set it to equal σ_ϵ .²⁰

Finally, the share of owners versus renters identifies the level of outside option utility $\bar{u}$.

4.3. Parameters, model fit, and refinancing behavior

Table 3 reports the parameters of the model for the three cases of the estimated model: aggregate, income group-specific, and geography-specific. The top of the table reports the fixed parameters, which are common across cases and across groups.

The main body of the table reports the estimated parameters. Column (1) reports the parameter estimates for the baseline versions that use UK-wide moments, and their asymptotic standard errors in parentheses. Columns (2) and (3) report the estimates for the case that uses separate moments for each income group and for each region and devolved administration, respectively. In columns (2) and (3), we report the weighted averages of the parameters across groups, as well as the weighted standard deviations of the parameters across groups (in parentheses), where the weights are the estimated market sizes M_j .

The bottom of Table 3 reports the calibrated market size M computed using Eq. (15); in columns (2)–(3), they correspond to the unweighted averages and standard deviations of M_j across groups. Note that several parameters are not easily comparable across columns. For example, the outside options $\bar{u}$ differ across groups in columns (2) and (3), and affect the estimated parameters of the valuation distribution. Other parameters, such as α , are more easily comparable across columns.

²⁰ Of course, we could set the value of $\epsilon_{i,0}$ to alternative values than the one that we choose.

4.3.1. UK-wide

The estimated parameters in column (1) imply that households' valuation v_i has a median equal to 1.001, a mean equal to 1.013 and a standard deviation equal to 0.160 in the full population of borrowers (homeowners) and non-borrowers (renters). In the model, households with the lowest valuations are less likely to participate in the mortgage market, choosing instead to rent a property. This means that, among borrowers, valuations are higher, with median v_i equaling 1.069, mean 1.083, and standard deviation 0.135.

The estimate of the parameter $\alpha = 0.786$ implies modest concavity in household utility from housing. This value implies that a household with average v_i enjoys a utility flow of $v_i h^\alpha$ equal to £10,232 over a two-year period from a house worth £125,000. This translates into an annual yield of 4.012%, which is slightly lower than the UK average rental yield, but broadly in line with average rental yields in London in this period (see, for example, Savills, 2015).

Among homeowners/borrowers, the median origination cost k_i^o equals £1294, its mean equals £4321, and its standard deviation equals £7190. However, households with the highest origination are less likely to participate in the mortgage market and choose to rent a property. Because our moments do not report any information on households who do not borrow (except for their share in the population), we obtain the distribution of origination costs k_i^o (as well as that of preferences v_i) in the full population by extrapolating those of borrowers out of sample. This leads us to estimate origination costs and refinancing costs across all households, including those that do not borrow, with a median that equals £3837, a mean of £13,135, and a standard deviation of £243,229 in the full population. It is worth noting that in the counterfactual exercises, we retain origination costs k_i^o , but remove refinancing costs $k_{i,t}$. This results in a relatively small effect on our calculations of the large k_i^o values estimated for non-participants in the baseline dual-rate economy.²¹

Interestingly, while we assume that preferences v_i and costs k_i^o are uncorrelated in the population of households, they are correlated among borrowers because of households' endogenous selection into the mortgage market—borrowers with high k_i^o enter the market only if their v_i is sufficiently high. The correlation coefficient among borrowers equals 0.352, suggesting that the effect of selection is appreciable.

In turn, the estimates in column (1) of Table 3 imply that median refinancing cost of borrowers $k_{i,t}$ equals £456, its mean equals £4042, and its standard deviation equals £15,102. Moreover, The estimate of $\sigma_\epsilon = 0.925$ means that the standard deviation of $\epsilon_{i,t}$ equals 1.164, which implies that the within-household variation in refinancing costs is non-trivial: The ratio $\frac{St.Dev.(k_i)}{St.Dev.(k_{i,t})}$ equals 0.65 in the population and 0.64 among borrowers—that is, the persistent household component k_i (cross-household variation) accounts for a larger share of the standard deviation of the refinancing costs $k_{i,t}$ than the transitory component $\epsilon_{i,t}$ (within-household variation). Below, we provide more statistics on borrower refinancing costs $k_{i,t}$ and compare them to the benefits of refinancing.

The value of the per-period outside option utility $\bar{u}$ equals £1062, which implies an annual net utility from renting equal to $\frac{\bar{u}}{1+\beta^{1/2}} = £544$. Households with a net utility value (over and above all mortgage payments and refinancing costs) greater than this level from purchasing a house enter the mortgage market.

Table 4 presents a comparison between the empirical moments and the moments calculated from the model at the estimated parameters reported in column (1) of Table 3. Overall, the model seems to fit the data well, successfully matching two crucial features of the data that underscore refinancing incentives across households and over time: on average, mortgages on the discounted rate have higher balances and

²¹ We could reduce the large standard deviation by setting an upper bound to k_i^o equal to the maximum forgone refinancing benefits in the data.

Table 3
Parameters.

r	650	R	759	T	15
β	0.9025	ω	1.2500	η	0.5000
	UK-wide (1)	Income Groups (2)		Regions (3)	
μ_v	0.0008 (0.0000)	0.0007 (0.0002)		0.0009 (0.0006)	
σ_v	0.1566 (0.0002)	0.2193 (0.0119)		0.3427 (0.0268)	
μ_{k1}	5.7454 (0.0076)	5.1239 (0.7421)		5.1073 (0.5263)	
σ_{k1}	2.6460 (0.0304)	1.7335 (0.2466)		1.5336 (0.9701)	
μ_{k2}	9.1875 (0.0178)	9.3004 (0.1201)		9.1938 (0.0643)	
σ_{k2}	0.9868 (0.0002)	0.9547 (0.0794)		0.9754 (0.0969)	
σ_e	0.9253 (0.0010)	0.8995 (0.0470)		0.8104 (0.0690)	
α	0.7856 (0.0000)	0.7896 (0.0076)		0.7879 (0.0026)	
$\bar{u}$	1062 (4.0946)	1,564 (590.5706)		1556 (430.2231)	
M	379,145	28,164 (11,890)		31,894 (15,298)	

Notes: This table reports the estimated parameters. In column (1), the numbers in parentheses refer to asymptotic standard errors of the parameter estimates. In columns (2) and (3), the numbers in parentheses refer to standard deviations of the parameter estimates across income and region groups, respectively.

Table 4
Model fit.

	Data	Model	t-statistic
Mean Loan Balance, Discounted Rate	140,647	141,240	−8.63
Standard Deviation Loan Balance, Discounted Rate	105,062	107,233	−11.49
Mean Loan Balance, Reset Rate	112,692	111,688	14.12
Standard Deviation Loan Balance, Reset Rate	79,684	78,120	6.65
Mean Remaining Years, Discounted Rate	20.57	18.82	28.96
Standard Deviation Remaining Years, Discounted Rate	7.73	7.85	−3.30
Mean Remaining Years, Reset Rate	16.84	15.51	17.95
Standard Deviation Remaining Years, Reset Rate	6.95	7.39	−8.67
Share of Mortgages on Discounted Rate, 0–5 Percentile	52.72	52.88	−1.39
Share of Mortgages on Discounted Rate, 5–25 Percentile	56.36	57.62	−21.53
Share of Mortgages on Discounted Rate, 25–50 Percentile	61.48	59.95	29.80
Share of Mortgages on Discounted Rate, 50–75 Percentile	67.76	63.97	76.82
Share of Mortgages on Discounted Rate, 75–95 Percentile	73.77	71.88	36.55
Share of Mortgages on Discounted Rate, 95–100 Percentile	81.19	81.93	−8.02
Transition from Reset Rate to Discounted Rate	16.52	18.01	−45.10
Share of Owners	63.13	63.12	0.97

Notes: This table reports the values of the empirical moments and of the moments calculated at the estimated parameters reported in column (1) of [Table 3](#).

are closer to issuance (have greater remaining maturity) than those on the reset rate. However, the absolute value of most t -statistics exceed standard critical values, because our sample size of 3.59 million mortgages is extremely large and thus the standard errors of the data moments are tiny.

Refinancing: Benefits and costs. The parameters of column (1) in [Table 3](#) have some interesting implications for borrowers' refinancing behavior.

The left panel of [Fig. 1](#) displays the distribution of the net benefits of refinancing $k_i^*(t) - k_{i,t}$ of all borrowers in the model (i.e., including also those with a simulated loan balance of less than £30,000 that we excluded from the computation of the moments used in the estimation). The heterogeneity of net benefits is striking: borrowers have positive net benefits in 59 percent of refinancing opportunities, which corresponds to the share of all mortgages on the discounted rate.²² The median estimated net benefit is positive (it equals £264), but the

average net benefit is negative (it equals −£2827), driven by the long left tail (the standard deviation equals £15,067). Some borrowers have extremely low measured net benefits, reflecting the fact that the model requires high costs to rationalize the non-refinancing behavior of a small group of borrowers with high loan balances and long maturities that would otherwise be expected to refinance.

The distribution of gross benefits of borrowers who refinance has a median of £1003, an average of £1329, and a standard deviation of £1099; their costs have a median of £73, an average of £266, and a standard deviation of £472. The corresponding distribution of gross benefits of borrowers who do not refinance has a median of £836, an average of £1049, and a standard deviation of £854; the costs of these non-refinancing borrowers have a median of £3704, an average of £9554, and a standard deviation of £22,572. The comparison of these statistics between borrowers who refinance and borrowers who do not

²² Because the mortgages excluded from the estimation have small balances and thus small gross refinancing benefits, the share of mortgages on the

discounted rate in the full distribution, equal to 59 percent, is lower than that in the truncated distribution, equal to 64 percent.




Fig. 1. Distribution of net benefits of refinancing.

Notes: The left panel displays the histogram of the net benefits of refinancing. The right panel displays the histograms of the number of periods in which borrowers pay the discounted rate.




Fig. 2. Transition probabilities by loan balance.

Notes: This figure displays the shares of mortgages that pay reset rates in 2015H1 and discounted rates in 2017H1 across loan balance deciles in the data (red squares) and in the model (blue dots). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

shows that the difference in their respective costs is larger than that in their benefits. Hence, heterogeneity in refinancing costs $k_{i,t}$ is the main driver in the model of the heterogeneity in refinancing behavior observed across borrowers.

The heterogeneity of refinancing behavior is also apparent in the right panel of Fig. 1, which displays the distribution of the number of periods on the discounted rate across individuals. No borrower always pays the reset rate because all of them receive the discounted rate at origination, in period $t = 0$. Thereafter, 14 percent of borrowers never refinance, many borrowers refinance occasionally, and 33 percent of borrowers always refinance. This heterogeneous distribution obtains because borrowers with low values of their persistent component k_i of refinancing costs (almost) always refinance, whereas borrowers with high values of k_i refinance only when they receive a temporary shock $\varepsilon_{i,t}$ that is low enough.

Fig. 2 provides further insights into refinancing behavior. It displays the shares of mortgages that transition from paying reset rates in 2015H1 to paying discounted rates in 2017H1 across loan balance deciles in the data (red squares) and in the model (blue dots). Our estimation targets (and successfully matches) the aggregate transition rate in the sample, but the rates across loan balance deciles are untargeted moments. The figure shows that our model reproduces the feature that the probability of switching from the reset rate to the discounted rate increases with the size of the loan balance. This is useful validation that

our model is well able to capture important variation across income groups in refinancing incentives.

Whereas our primary focus is to model refinancing behavior, we note that the heterogeneity of borrower refinancing propagates into substantial heterogeneity in borrowers' elasticities of loan demand with respect to discounted and reset rates. We estimate that the mean borrower elasticity with respect to the discounted rate r_d equals -1.624 and its standard deviation equals 0.551 . These magnitudes are close to those reported by other papers that use UK mortgage origination data: Benetton et al. (2023) estimate an average elasticity of the mortgage loan demand to the initial rate equal to 1.23 percent among first-time buyers and home movers in 2010–2014, and Taburet (2023) estimates a semi-elasticity to the initial interest rate of 0.52 with an average interest rate of 250 bps in 2018 data, corresponding to an average elasticity of loan demand equal to 1.3 percent.

Borrowers with a lower k_i are more elastic to the discounted rate and less elastic to the reset rate than borrowers with a higher k_i because they are more likely to refinance regularly and thus pay the discounted rate. The elasticities to the discounted rate of the lowest- and highest- k_i borrowers equal -2.461 and -0.300 , respectively. The mean borrower elasticity with respect to the reset rate R equals -0.402 and its standard deviation equals 0.629 . Interestingly, in our setting some borrowers display a positive elasticity with respect to the reset rate, because if the reset rate increases (while keeping the discounted rate fixed), the




Fig. 3. Refinancing costs across groups. Notes: This figure displays the average refinancing cost $k_{i,t}$ of households with a mortgage in each income group.

benefits of refinancing increase, and thus some borrowers are more likely to refinance. This leads to a lower average interest rate, meaning that such borrowers will increase their initial loan size in response to a higher reset rate.

4.3.2. Multiple groups

Columns (2) and (3) in Table 3 report parameters for the model estimated on the 12 different income groups and on the 12 UK regions and devolved administrations, respectively. As discussed above, some parameters are not easily comparable between columns (1)–(3), though many are similar in magnitudes. The parameters in column (2) for the different income groups exhibit some differences from those in column (3) because the heterogeneity across and within income groups differs from the heterogeneity across and within regional groups, which in turn affects the average and the standard deviations of some parameters.

The parameters that exhibit the most interesting heterogeneity across groups are those that determine the distribution of refinancing costs $k_{i,t}$. Fig. 3 displays borrowers' average refinancing cost across income groups: Perhaps surprisingly, but consistent with findings in Andersen et al. (2020), they increase with income, from £1896 among the lowest-income borrowers to £6515 among the highest-income borrowers. Our analysis of devolved administrations also displays a similar pattern in that more deprived regions tend to have lower refinancing costs than less deprived regions.

This finding is pertinent to our understanding of household refinancing behavior because it implies that the relatively prompt refinancing of higher-income groups, shown in Table 2, follows from their greater financial incentives due to larger loan sizes, and not because they have lower refinancing costs.²³ These estimated refinancing costs have consequences for our study of counterfactuals below. In particular, replacing the dual rate structure of mortgages with a constant interest rate allows borrowers to save these recurring refinancing costs, which needs to be taken into account as a benefit even for those borrowers who end up paying higher average interest rates.

Moreover, the outside option $\bar{u}$ also displays significant heterogeneity in the population. This parameter is a key input into the “extensive margin” decision of households, i.e., whether or not they enter the mortgage market. The heterogeneity in this parameter across groups means that there are different sensitivities across groups of this extensive margin decision to changes in interest rates. This factor contributes

to differences between the sensitivity of household participation decisions to interest rates in the multiple-group model and that in the baseline model.

Although we do not report measures of goodness-of-fit across groups, we note that the model fits the group-specific moments well. This is perhaps not surprising given that Table 4 shows that the UK-wide model fits the aggregate data well; the same model might therefore be expected to fit as well or better at a lower level of aggregation.

5. Counterfactual analyses: Constant interest rate

We compare the outcomes for households in our estimated models under the dual rate structure, with a counterfactual in which all households pay a constant interest rate and have no need to refinance. We present these comparisons for a constant interest rate equal to the weighted average of the discounted and the reset rates, i.e.,

$$r_c = \frac{r_d(Q_0(r_d) + Q_1(r_d)) + RQ_2(R)}{Q_0(r_d) + Q_1(r_d) + Q_2(R)}. \quad (16)$$

We calculate this weighted average using the aggregate balances in the data and obtain $r_c = 683$ bps. Appendix G reports counterfactual results for alternative values of this common interest rate.

We note here that our model focuses on cross-household differences in borrowers' inaction—i.e., the demand side of the mortgage market. Clearly, changes in the profile of interest rates affect lender profits and revenues as well, and their supply-side responses could constitute an important ingredient for further analysis.²⁴ Furthermore, households make optimal loan size decisions for a given interest rate based on their own valuation of housing and total costs of originating mortgages. We do not separately consider financial origination fees; as we show in Appendix E, such fees and costs vary only very slightly with loan size, and more importantly, and similar to prior literature (e.g., Andersen et al., 2020), are an order of magnitude smaller than the total costs (including all financial as well as psychological/behavioral costs) estimated using the model.

Table 5 reports the results of the counterfactual mortgage market outcomes for the different estimated models (in different columns) as ratios of their respective baseline values (i.e., in the dual-rate economy). We first describe the changes for the UK-wide case, and then for the multiple-group cases.

²³ In Figure H.1 in the Appendix, we also show the variance of the refinancing cost shock across groups, which is relatively stable across the income distribution, meaning it is not the principal driver of cross-income-group differences in refinancing behavior.

²⁴ Among others, Gurun et al. (2016), Guiso et al. (2022), Benetton et al. (2023), Allen and Li (2020), and Thiel (2021) study supply-side incentives in mortgage markets.

Table 5
Market outcomes with constant interest rates.

	UK-wide	Income groups	Regions
Constant Interest Rate=683 bps			
Number of Mortgages	1.06	1.09	1.07
Mean Initial Loan Amount	0.97	0.97	0.97
Standard Deviation Initial Loan Amount	0.96	0.93	0.94
Mean Loan Balance	0.97	0.97	0.97
Standard Deviation Loan Balance	0.96	0.94	0.95
Consumer Surplus	1.04	1.03	1.03

Notes: This table reports the statistics on the mortgage market in counterfactual markets with constant interest rates, as ratios of those of the estimated market with dual interest rates. The statistics are calculated using a constant interest rate equal to the average interest rate equal to (16).




Fig. 4. Change in loan size at origination.

Notes: The left panel displays the distribution of the changes in loan sizes at origination between the counterfactual economy with constant interest rates and the baseline economy with discounted and reset rates. The right panel displays the average change in loan sizes for households with different values of their origination costs (in bins of £2500). All statistics displayed are computed including only households who either participate in the mortgage market in the baseline dual-rate economy, or in the counterfactual single-rate economy, or in both.

UK-wide case. Table 5 reports that the change in the profile of interest rates to a single-rate structure yields two main aggregate adjustments in opposite directions: the number of mortgages increases, but the average loan size decreases.

More precisely, the first row of Table 5 reports that the number of mortgages increases by 6.44 percent relative to the number of mortgages in the baseline economy. The reason for this increase is that there are many households with valuation v_i and with moderate or high costs k_i^o just below the entry threshold $v_i^*(k_i^o)$ in the baseline economy who switch from renting a property in the dual-rate economy to taking a mortgage to buy a house in the counterfactual single-rate economy. These households would rarely refinance, and thus expect to pay an average rate close to the reset rate in the baseline dual-rate economy, which raises the costs of taking on a mortgage. These households, therefore, choose to rent in the dual-rate world, but since they pay a lower rate in the counterfactual single-rate economy, they choose to buy a house by taking on a mortgage in the counterfactual. The mass of these households, on net, is greater than the mass of households with low k_i^o who pay an average rate close to the discounted rate in the baseline economy but pay a higher rate in the single-rate economy. Such low- k_i^o households switch from owning with a mortgage in the dual-rate economy to renting a property in the counterfactual single-rate economy, but their exit is more than offset by new entrants in the single-rate mortgage market.

The second row shows that the average initial loan size decreases by 2.85 percent of the average loan size in the baseline case, corresponding to a mortgage size reduction of £5465. The main reason for this decline is the change in the composition of borrowers: marginal households who enter the mortgage market in the single-rate economy have smaller

loan sizes than inframarginal households whose participation does not change.

More generally, the change in the average loan size combines borrowers who increase their mortgage amounts with borrowers who decrease them. The left panel of Fig. 4 shows the full distribution of the changes in mortgage amounts of those households who participate in the mortgage market in the baseline dual-rate economy, in the counterfactual single-rate economy, or both. The heterogeneity of the changes in mortgage amounts at origination is apparent, with decreases in mortgage amounts more concentrated than increases.

The right panel of Fig. 4 helps to rationalize the asymmetric adjustment in loan sizes. It displays how the average change in mortgage size varies with the origination cost k_i^o , which is correlated with refinancing costs. Borrowers with the lowest k_i^o likely pay an interest rate close to 650 bps in the estimated mortgage market, because they almost always refinance, but they pay 683 bps in the counterfactual market with a constant interest rate. This higher rate induces them to reduce their loan sizes. In contrast, borrowers with the highest k_i^o likely pay an interest rate close to 759 bps in the baseline market, because they never refinance, but pay 683 bps in the counterfactual market. As a result, these borrowers increase their loan sizes. The increases in loan sizes are more dispersed than the decreases, because there is a bigger difference between the rates in the dual- and single-rate worlds paid by households with high k_i^o than that between the interest rates in the economies paid by those with low k_i^o .

The third row in Table 5 reports that the standard deviation of initial loan sizes declines quite substantially, by 4.34 percent of the standard deviation of the initial loan size (corresponding to £5440) in the estimated baseline model. The reason is that one dimension of

household heterogeneity, namely k_i , contributes to the determination of the loan size in the baseline model with refinancing. However, this dimension of heterogeneity becomes irrelevant when interest rates are constant. More specifically, the previous arguments suggest – and Fig. 4 shows – that borrowers with larger loans in the baseline economy decrease their loan sizes in the counterfactual, whereas borrowers with smaller loans in the baseline economy increase their loan sizes in the counterfactual with constant interest rates and no refinancing. A common interest rate thus pushes loan sizes to be more homogeneous.

The decline in initial loan size and the increase in the number of mortgages together combine to increase aggregate mortgage debt by 3.55 percent relative to the model with dual rates. Cross-subsidies are eliminated in the counterfactual, and one consequence of this change is that the mortgage market increases in size, although the effect is tempered by the opposing effects on the extensive and intensive margins.

The fourth and fifth rows of Table 5 report that the patterns in the initial loan size distribution described above transfer to the aggregate loan balance distribution (i.e., including different cohorts of mortgages), with one additional subtle effect. In the baseline economy, on average, borrowers who originate large loans pay lower rates than borrowers with small loans. Hence, as loans amortize over time, the loan balances of borrowers with large loans tend to decline at a faster rate than the loan balances of borrowers with small loans, which compresses the distribution of loan balances over time. This force is absent in the counterfactual single-rate economy as all borrowers pay the same rate. Hence, the standard deviation of loan balances (normalized by that observed in the baseline economy) reported in the last row is slightly larger than the standard deviation of initial loan balances (also normalized with respect to the baseline economy) reported in the third row.

Finally, the last row summarizes all the changes in a single money-metric ex-ante measure of consumer surplus, calculated for each household as $\max(W_0(v_i, k_i^0), \frac{\bar{u}}{1-\beta})$. Consumer surplus increases by 4.14 percent in the single-rate economy relative to the dual-rate economy.

Income groups. The cases with multiple groups allow us to explain some of the observable heterogeneity in refinancing rates across income groups and geographies of the UK with heterogeneity in preferences v_i and costs k_i . These richer cases help us to evaluate whether and how the shift to a single mortgage rate structure leads to different outcomes for households in these groups.

Column (2) of Table 5 reports aggregated counterfactual estimates when the model is estimated using moments for different income groups. When we compare these aggregate statistics with those of the UK-wide model in column (1), the differences appear small. The main difference is that column (2) exhibits a slightly larger adjustment in the extensive margin (i.e., the number of mortgages) than column (1). We now analyze how the results differ across income groups.

Fig. 5 plots selected changes to mortgage market outcomes for each income group. The top-left panel shows that interest rates (in bps) are lower in the counterfactual economy for income groups up to roughly the 80th percentile of the income distribution in the sample, and are higher for the very highest income groups. This pattern is consistent with the regressive nature of the cross-subsidies in the dual-rate economy. The highest income group pays higher interest rates in the single-rate economy than the average rates they pay in the dual-rate economy. This is primarily because high-income households have larger loans, which gives them greater incentives to refinance promptly in the dual-rate economy.

The top-right panel shows that these changes in interest rates translate into an aggregate increase in the share of households with mortgages. Critically, the percentage-point increase is larger for lower-income groups, and minimal for the highest-income groups. Table 2 reports that the homeownership rate among low-income households equals 50 percent and rises steeply with income, and our model suggests

that the design of the mortgage market may be a contributing factor to these patterns. In the single-rate market, there is a greater entry of these low-income households into the housing and mortgage markets.

The bottom-left panel plots the average percent differences between initial loan sizes in the single-rate economy and those in the dual-rate economy. While there are also important changes within groups, the across-group comparison highlights that higher-income groups adjust their average initial loan size downward more than lower-income groups. The adjustment in the average initial loan size of the highest-income group is sizable, because many of these borrowers – i.e., a larger fraction than among lower-income groups – almost always refinance and thus suffer a substantial increase in the interest rate that they pay, from $r_d = 650$ to $r_c = 683$ bps.

Finally, the bottom-right panel summarizes all changes into changes in consumer surplus, confirming that all income groups would benefit in the single-rate economy relative to the dual-rate economy, because they either pay lower interest rates or they save the refinancing costs $k_{i,t}$. Lower-income groups attain larger percentage-change increases in consumer surplus than higher-income groups. This makes our findings quite nuanced: Low-income households are penalized in the status quo dual-rate economy largely because they have smaller loan balances and thus lower incentives to refinance, whereas high-income households refinance more frequently due to greater loan balances, although they actually incur higher refinancing costs from doing so. On net, our counterfactual suggests that the removal of refinancing costs more than compensates high-income households for the increase in paid rates and associated reductions in loan balances. Overall, this leads to a positive increase in consumer surplus across all household groups.

Overall, these panels suggest that the richer model with greater household heterogeneity across the income distribution implies that higher-income households pay lower rates and enjoy greater benefits than lower-income households in the current dual-rate structure. These patterns are consistent with the idea that the dual-rate structure fosters regressive cross-subsidies. The panels of Fig. 5 suggest that different income groups would respond to a single-rate structure with different types and levels of adjustments on both the intensive and extensive margins. In particular, raising participation in the mortgage market is the main adjustment for lower-income groups, whereas lowering initial loan sizes is the main adjustment for higher-income groups.

Regions. Column (3) of Table 5 reports aggregated counterfactual estimates when we estimate the model with parameters and moments for different UK regions. Once again, as with the model which incorporates greater heterogeneity across income groups, the aggregate statistics reported in column (3) are remarkably similar to those of the UK-wide estimation reported in column (1).

Fig. 6 presents maps that display some of the changes to mortgage market outcomes across different UK regions. In each panel, darker colors indicate larger (positive) changes in the counterfactual market with constant interest rates when compared with the baseline dual-rate economy with discounted and reset rates.

The top-left map displays the change in average interest rates paid on mortgages, reported in bps. Households in the more prosperous regions of Greater London, the South East of England, and the East of England experience the largest increases in mortgage rates, whereas households in relatively less well-off regions and devolved administrations such as Northern Ireland, Wales, and the North East of England would experience the largest decreases in rates when moving to a single rate. These regional patterns are consistent with the dual-rate structure featuring regressive cross-subsidies across UK regions.

In the counterfactual scenario, as seen earlier in the case of income groups, households endogenously adjust their mortgage market participation as well as their mortgage amounts. The top-right plot shows the change in the number of mortgages, which broadly increases the most in regions and devolved administrations that experience the largest decrease in mortgage rates, such as Scotland and Northern Ireland. In




Fig. 5. Changes in market outcomes by income groups.

Notes: The top-left panel displays the change in interest rates (in bps); the top-right panel displays the percentage-point change in the share of households with mortgages; the bottom-left panel displays the percentage change in the average initial loan size; and the bottom-right panel displays the percentage change in average consumer surplus for each income group in the counterfactual case with a constant interest rate equal to $r_c = 683$ bps relative to the baseline case.

contrast, southern regions of the UK experience smaller adjustments to mortgage market participation.

The bottom-left map displays the changes in the average initial loan size rate across regions. The differences in these averages when moving to the single-rate world mask larger within-region changes. That said, once again, southern regions' average initial loan sizes do shrink more than less well-off regions' average initial loan sizes in the counterfactual single-rate world. Overall, the top-right and bottom-left maps confirm the pattern that the change in the profile of interest rates affects mostly the extensive margin in lower-income regions, and mostly the intensive margin in higher-income regions.

Finally, the bottom-right map shows that all regions would enjoy higher consumer surplus in the single-rate economy than in the dual-rate economy, because they either pay lower interest rates or they save the refinancing costs $k_{i,t}$, with lower-income regions experiencing higher percentage-change increases in consumer surplus than higher-income regions.

6. Conclusion

We develop a model of mortgage refinancing and estimate it using rich data from the UK mortgage market. Our model matches broad features of the data, and the parameters reveal considerable heterogeneity in mortgage refinancing benefits and costs across households, echoing findings in prior literature. Our estimates imply that the current dual-rate structure penalizes households with small loan balances, such as low-income households; hence, they are more likely to pay the high reset rate than high-income households.

We use the estimated parameters to uncover and quantify regressive cross-subsidies in this market by conducting a counterfactual comparison with an alternative mortgage contract that features a constant interest rate and no need for refinancing. Using 2015 data, we set

annual interest rates in our main counterfactual single-rate scenario to lie 16 bps above the average discounted rate and 36 bps below the average reset rate that borrowers are rolled on to at the expiration of the discounted rate fixation period. These are material changes given the importance of mortgages to household budgets.

The counterfactual scenario features different adjustments by low- and high-income groups. Low-income households enter the mortgage market in greater numbers in response to the lower interest rates that they would pay and the elimination of refinancing costs. In contrast, high-income households mainly take on smaller loans in the single-rate counterfactual economy in response to their inability to take advantage of the discounted rate.

Our results suggest that simplifying the design of mortgage refinancing and eliminating the costs associated with refinancing can induce households (most notably those with lower incomes) to participate more extensively in the mortgage market, even if they do not have perfect foresight about their ongoing refinancing costs at origination. These findings highlight an important and novel dimension of inequality that would be invisible without our approach: Changes in mortgage rates make entry and participation in the housing market more equal across the income distribution; they also tend to push loan sizes to be more uniform across high- and low-income households.

Our work has both methodological and economic contributions beyond the specific context that we study. We believe that our approach to estimating financial cross-subsidies by comparing the current and counterfactual market structures is a useful way to provide a money-metric assessment of the impacts of heterogeneity in household inaction. This has potentially wider implications for the field of household finance, where such heterogeneity is widely prevalent in many markets including credit and insurance. Our findings on the regressive nature of these cross-subsidies highlight that other household finance settings where high-income households benefit more due to




Fig. 6. Regional changes.

their larger stakes and their greater propensity to take action may also contribute to inequality. In a broader sense, our results on the distribution of financial cross-subsidies in this important market show that

studying household finance is helpful for the agenda of identifying the sources and consequences of wealth inequality, a continuing concern for society.

CRediT authorship contribution statement

Jack Fisher: Writing – review & editing, Formal analysis. **Alessandro Gavazza:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Lu Liu:** Writing – review & editing, Writing – original draft, Investigation, Formal analysis, Data curation. **Tarun Ramadorai:** Writing – review & editing, Writing – original draft, Validation, Supervision, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jagdish Tripathy:** Writing – review & editing, Writing – original draft, Project administration, Investigation, Formal analysis, Data curation.

Declaration of competing interest

I declare that I have no relevant or material financial interests that relate to the research described in the paper, and I hold no paid or unpaid positions as officer, director, or board member of relevant organizations.

This paper did not involve collection of data on human subjects, so the project was considered exempt from IRB.

Data availability

Refinancing Cross-Subsidies Repository (Original data) (Mendeley Data)

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jfineco.2024.103876>.

References

- Agarwal, S., Rosen, R.J., Yao, V., 2016. Why do borrowers make mortgage refinancing mistakes? *Manage. Sci.* 62 (12), 3494–3509.
- Allen, J., Li, S., 2020. Dynamic competition in negotiated price markets. *Staff Working Paper 2020–22*, Bank of Canada.
- Alvaredo, F., Chancel, L., Piketty, T., Saez, E., Zucman, G., 2017. Global inequality dynamics: New findings from WID.world. *Amer. Econ. Rev.* 107 (5), 404–409.
- Andersen, S., Campbell, J.Y., Nielsen, K.M., Ramadorai, T., 2020. Sources of inaction in household finance: Evidence from the Danish Mortgage market. *Amer. Econ. Rev.* 110 (10), 3184–3230.
- Armstrong, M., Vickers, J., 2012. Consumer protection and contingent charges. *J. Econ. Lit.* 50 (2), 477–493.
- Ausubel, L.M., 1991. The failure of competition in the credit card market. *Amer. Econ. Rev.* 81 (1), 50–81.
- Badarizna, C., Campbell, J.Y., Ramadorai, T., 2016. International comparative household finance. *Annu. Rev. Econ.* 8, 111–144.
- Badarizna, C., Campbell, J.Y., Ramadorai, T., 2018. What calls to ARMs? International evidence on interest rates and the choice of adjustable-rate mortgages. *Manage. Sci.* 64 (5), 2275–2288.
- Belgibayeva, A., Bono, T., Bracke, P., Cocco, J., Majer, T., 2020. When discounted rates end: The cost of a taking action in the mortgage market. *Occasional Paper 54*, Financial Conduct Authority.
- Benetton, M., 2021. Leverage regulation and market structure: An empirical model of the UK mortgage market. *J. Finance* 76 (6), 2297–3053.
- Benetton, M., Gavazza, A., Surico, P., 2023. Mortgage pricing and monetary policy. Discussion Paper 16456, Centre for Economic Policy Research.
- Benhabib, J., Bisin, A., 2018. Skewed wealth distributions: Theory and empirics. *J. Econ. Lit.* 56 (4), 1261–1291.
- Beraja, M., Fuster, A., Hurst, E., Vavra, J., 2019. Regional heterogeneity and the refinancing channel of monetary policy. *Q. J. Econ.* 134 (1), 109–183.
- Berger, D., Milbradt, K., Tourre, F., Vavra, J., 2021. Mortgage prepayment and path-dependent effects of monetary policy. *Amer. Econ. Rev.* 111 (9), 2829–2878.
- Berger, D., Milbradt, K., Tourre, F., Vavra, J., 2023. Refinancing frictions, mortgage pricing and redistribution. Working Paper, Northwestern University.
- Best, M.C., Cloyne, J.S., Ilzetzki, E., Kleven, H.J., 2020. Estimating the elasticity of intertemporal substitution using mortgage notches. *Rev. Econ. Stud.* 87 (2), 656–690.
- Brunnermeier, M.K., Nagel, S., 2008. Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals. *Amer. Econ. Rev.* 98 (3), 713–736.
- Byrne, S., Devine, K., King, M., McCarthy, Y., Palmer, C., 2023. The last mile of monetary policy: Inattention, reminders, and the refinancing channel. Working Paper Series 31043, National Bureau of Economic Research.
- Campbell, J.Y., 2006. Household finance. *J. Finance* 61 (4), 1553–1604.
- Campbell, J.Y., 2013. Mortgage market design. *Rev. Finance* 17 (1), 1–33.
- Campbell, J.Y., Ramadorai, T., Ranish, B., 2019. Do the rich get richer in the stock market? Evidence from India. *Am. Econ. Rev. Insights* 1 (2), 225–240.
- Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A., 2002. Defined contribution pensions: Plan rules, participant choices, and the path of least resistance. *Tax Policy Econ.* 16, 67–113.
- Cloyne, J., Huber, K., Ilzetzki, E., Kleven, H., 2019. The effect of house prices on household borrowing: A new approach. *Amer. Econ. Rev.* 109 (6), 2104–2136.
- DeFusco, A.A., Mondragon, J., 2020. No Job, No Money, No Refi: Frictions to refinancing in a recession. *J. Finance* 75 (5), 2327–2376.
- DellaVigna, S., Malmendier, U., 2006. Paying not to go to the gym. *Amer. Econ. Rev.* 96 (3), 694–719.
- Ellison, G., 2005. A model of add-on pricing. *Q. J. Econ.* 120 (2), 585–637.
- Erickson, T., Whited, T.M., 2002. Two-step GMM estimation of the errors-in-variables model using high-order moments. *Econometric Theory* 18 (3), 776–799.
- Fagereng, A., Guiso, L., Malacrino, D., Pistaferri, L., 2020. Heterogeneity and persistence in returns to wealth. *Econometrica* 88 (1), 115–170.
- Farrell, J., Klemperer, P., 2007. Coordination and lock-in: Competition with switching costs and network effects. In: Armstrong, M., Porter, R. (Eds.), *Handbook of Industrial Organization*, vol. 3, Elsevier, pp. 1967–2072.
- Financial Conduct Authority, 2019. Mortgage Market Study: Final Report. Financial Conduct Authority, Market Study MS 16/2.3.
- Financial Conduct Authority, 2021. Mortgage Prisoners Review. Financial Conduct Authority, Policy Paper CP 576.
- Gabaix, X., Laibson, D., 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Q. J. Econ.* 121 (2), 505–540.
- Gavazza, A., Lizzeri, A., 2021. Frictions in product markets. In: Ho, K., Hortaçsu, A., Lizzeri, A. (Eds.), *Handbook of Industrial Organization*, vol. 4, Elsevier, pp. 433–484.
- Goetzmann, W.N., Spaenjers, C., Van Nieuwerburgh, S., 2021. Real and private-value assets. *Rev. Financ. Stud.* 34 (8), 3497–3526.
- Greenwald, D.L., Leombroni, M., Lustig, H., Van Nieuwerburgh, S., 2021. Financial and total wealth inequality with declining interest rates. Working Paper Series 28613, National Bureau of Economic Research.
- Guiso, L., Pozzi, A., Tsou, A., Gambacorta, L., Mistrulli, P.E., 2022. The cost of steering in financial markets: Evidence from the mortgage market. *J. Financ. Econ.* 143 (3), 1209–1226.
- Gurun, U.G., Matvos, G., Seru, A., 2016. Advertising expensive mortgages. *J. Finance* 71 (5), 2371–2416.
- Handel, B.R., 2013. Adverse selection and inertia in health insurance markets: When nudging hurts. *Amer. Econ. Rev.* 103 (7), 2643–2682.
- Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 1029–1054.
- Heidhues, P., Köszegi, B., 2018. Behavioral industrial organization. In: Bernheim, B.D., Vigna, S.D., Laibson, D. (Eds.), *Handbook of Behavioral Economics - Foundations and Applications 1*, vol. 1, Elsevier, pp. 517–612.
- Honka, E., 2014. Quantifying search and switching costs in the US auto insurance industry. *Rand J. Econ.* 45 (4), 847–884.
- Hubmer, J., Krusell, P., Smith Jr., A.A., 2020. Sources of US wealth inequality: Past, present, and future. In: *NBER Macroeconomics Annual 2020*, vol. 35, University of Chicago Press.
- Hurst, E., Keys, B.J., Seru, A., Vavra, J., 2016. Regional redistribution through the US mortgage market. *Amer. Econ. Rev.* 106 (10), 2982–3028.
- Illanes, G., 2016. Switching costs in pension plan choice. Working Paper, Northwestern University.
- Keys, B.J., Pope, D.G., Pope, J.C., 2016. Failure to refinance. *J. Financ. Econ.* 122 (3), 482–499.
- Lea, M., 2010. International Comparison of Mortgage Product Offerings. Technical Report, Research Institute for Housing America.
- Liu, L., 2022. The demand for long-term mortgage contracts and the role of collateral. Working Paper, Imperial College London.
- Luco, F., 2019. Switching costs and competition in retirement investment. *Am. Econ. J. Microecon.* 11 (2), 26–54.
- Miles, D., 2004. The UK mortgage market: Taking a longer-term view. Working Paper, Imperial College London.
- Nelson, S., 2022. Private information and price regulation in the US credit card market. Working Paper, University of Chicago.
- Piskorski, T., Seru, A., 2018. Mortgage market design: Lessons from the great recession. *Brook. Pap. Econ. Act.* 2018 (1), 429–513.
- Robles-García, C., 2022. Competition and incentives in mortgage markets: The role of brokers. Working Paper, Stanford University.
- Savills, 2015. Rents and returns. Available at https://www.savills.co.uk/research_articles/229130/193120-0.
- Scharfstein, D., Sunderam, A., 2016. Market power in mortgage lending and the transmission of monetary policy. Working Paper, Harvard University.

- Stango, V., Zinman, J., 2016. Borrowing high versus borrowing higher: Price dispersion and shopping behavior in the US credit card market. *Rev. Financ. Stud.* 29 (4), 979–1006.
- Stanton, R., 1995. Rational prepayment and the valuation of mortgage-backed securities. *Rev. Financ. Stud.* 8 (3), 677–708.
- Stanton, R., Wallace, N., 1998. Mortgage choice: What's the point? *Real Estate Econ.* 26 (2), 173–205.
- Taburet, A., 2023. Screening using a menu of contracts: A structural model for lending markets. Working Paper, Duke University.
- Thiel, J., 2021. Price discrimination, switching costs and welfare: Evidence from the dutch mortgage market. Working Paper, Copenhagen Business School.
- Zhang, D., 2022. Closing costs, refinancing, and inefficiencies in the mortgage market. Working Paper, Harvard University.