Behavioral Lock-In: Housing Markets with Reference Dependent Agents

Cristian Badarinza, Tarun Ramadorai,

Juhana Siljander, Jagdish Tripathy*

August 16, 2026

Abstract

We embed optimizing agents with reference-dependent and loss-averse preferences into a dynamic equilibrium search and matching model of the housing market with rich heterogeneity and realistic financial constraints. We calibrate and evaluate the model using granular administrative data from the U.K. housing market. Behavioral frictions act as a nominal rigidity, a cause of hysteresis in response to shocks, and a source of distortion for ongoing property taxes. At the aggregate level, a simple statistic, the prevalence of "paper losses" in the stock of properties, captures variation in prices and volumes across regions, and predicts local responses to policy interventions.

*Badarinza: Frankfurt School of Finance & Management, Email: c.badarinza@fs.de. Ramadorai (corresponding author): London School of Economics, Imperial College London and CEPR, Email: t.ramadorai@lse.ac.uk. Siljander: Bank of Finland, Email: juhana.siljander@bof.fi. Tripathy: Bank of England, Email: jagdish.tripathy@bankofengland.co.uk. Badarinza and Tripathy were visiting researchers at Imperial College London during the production of this draft. We thank Ed Burgess, Tim Bannister, Natalie Jakomis, Rob Murphy, Leanne Pienaar, Amit Tiwari, and Oscar Walet at Rightmove for data access and numerous useful conversations, and John Campbell, Stefano Della Vigna, Daniel Greenwald, Ben Lockwood, Ulrike Malmendier, Rachel Ngai, seminar participants at the London School of Economics Macro Workshop, the Imperial-LSE-UCL-IFS-CEPR Micro-Macro Household Finance Seminar, the London Behavioural Finance Workshop, Princeton University, Emory University, INSEAD, Swansea University, the 5th Joint BoC-ECB-NY Fed Conference, EDHEC, the Stanford Institute for Theoretical Economics (Psychology in Economics and Financial Regulation), the Imperial-Bank of England Household Finance and Housing Conference, the 2024 Nordic Summer Symposium in Macroeconomics, the Yale SoM Finance Seminar, the 14th Helsinki Finance Summit, the 2024 BoF-CEPR conference, the CREI @30 Conference, the Berkeley Psychology and Economics seminar, the AEA-AFA joint luncheon, Oxford University, and the MIT Finance Seminar for comments. Any views expressed are solely those of the author(s) and so cannot and should not be taken to represent, or reported to represent, the views of the Bank of England, or members of the Monetary Policy Committee, Financial Policy Committee or Prudential Regulation Committee, or to state Bank of England policy.

1 Introduction

Evidence from high-stakes field settings, including job search, income tax filing, and financial investment shows that reference dependence and loss aversion systematically shape economic behavior (DellaVigna et al., 2017; Rees-Jones, 2018). Housing provides a particularly transparent setting, where international evidence shows that owners are visibly reluctant to sell below the nominal price that they originally paid (Genesove and Mayer, 2001; Andersen et al., 2022). This reluctance leaves a precise empirical signature in realized transaction prices, listing premia, sale probabilities, and time on the market as sellers move across the gain/loss boundary defined by the original purchase price. While these micro-level patterns are intriguing, does this behavioral mechanism materially affect aggregate house prices and quantities, and the transmission of monetary and housing-market fiscal policy?

We answer these questions by embedding reference-dependent, loss-averse homeowners in a dynamic equilibrium search-and-matching model with rich heterogeneity and realistic financial constraints. We calibrate the model on micro-level moments from granular administrative data covering housing transactions, property listings, buyer search, and the universe of mortgage contracts in England and Wales. The calibrated model is well able to capture a range of untargeted aggregate moments, and delivers useful predictions which we validate in the data.

The central idea is that individual homeowners' nominal anchoring behavior affects their pricing and trading decisions. This individual behavior, when aggregated, affects posted prices and equilibrium sale probabilities, which feed back into individual decision problems. This feedback makes the distribution of nominal gains and losses a payoff-relevant aggregate state variable, since current shocks alter the distribution of individual gain/loss positions and thus, the trading environment facing market participants. An important consequence is "behavioral lock-in": sellers facing the prospect of realizing a nominal loss defer trading even after large idiosyncratic mobility shocks, forgoing moves that would be valuable in a frictionless world. The same mechanism makes the response of prices and transaction volumes to common shocks and policy interventions depend on local transaction histories.

The quantitative significance of reference dependence arises for four related reasons. First, while 1–2% of sales transact exactly at the original purchase price, exact bunching is not the relevant measure of aggregate exposure. What matters is the density of homeowners near

the nominal anchor, since small movements in valuations can push many sellers across the gain/loss domain. Second, reference dependence operates inside a search and matching equilibrium (Guren, 2018; Genesove and Han, 2012), so a higher reservation price slows the match rate, and buyers' and sellers' positions relative to the reference point govern whether a purchase intention converts into a completed sale. Third, reference points are endogenous to past transaction histories, so the cross-sectional distribution of nominal gains and losses evolves slowly with the realized history of prices, generating hysteresis and persistent geographic dispersion even when fundamentals are stationary. Fourth, mortgage balances are set relative to the same nominal purchase price, so reference dependence and home-equity constraints (Stein, 1995) have complementary effects linked to the same historical anchor. Since housing activity is cyclical and a major contributor to the economy (Leamer, 2007), the amplification of transaction-volume responses generated by this micro-level nonlinearity also affects broader quantities.

Loan-level data covering the universe of U.K. mortgage contracts allow us to construct home-equity positions at the time of listing and to study sellers both with and without an outstanding mortgage. The data show that reference-dependent listing and transaction patterns remain present among sellers without an outstanding mortgage, while low home equity introduces an additional financial force among mortgaged sellers. The model incorporates realistic mortgage amortization and home-equity constraints, and throughout the quantitative analysis we compare the full model with a financial-only specification in which these constraints remain operative but reference dependence is switched off. This design allows us to quantify the separate and joint contributions of behavioral and financial frictions to individual listing and pricing decisions and to aggregate regional prices and volumes.

A direct implication of the mechanism is that a single summary statistic, the "paper loss share," the fraction of homeowners in a region whose property would sell today at a loss relative to its original purchase price, helps predict aggregate outcomes. Across regions of England and Wales, the loss share alone explains roughly 47% of the cross-sectional variation in local price–volume comovement, and it predicts that high-loss-share regions have substantially more volatile transaction volumes, and marginally less volatile prices, than low-loss-share regions.[1]

[1]Importantly, this regional variation is driven by the post-listing behavior of sellers and the matching of buyers to listings, rather than by the rate at which new listings arrive: the loss share affects *how* a market

Because each transaction establishes a new reference price, the loss share is itself the cumulative imprint of past shocks, and captures the endogenous link between short- and long-run housing-market dynamics.

Equipped with the calibrated model, we first show that in the presence of behavioral and financial frictions demand shocks are absorbed more by quantities and less by prices precisely in the segments of the market where many homeowners face nominal losses or low home equity. This helps account for the long-standing puzzle in housing economics that house prices and housing transaction volumes move together strongly and positively, both at business-cycle frequencies and across local housing markets within the same country.[2] The pattern is evident in recent data: across all regions of the U.K. between 2010 and 2024, a 6% decline in local house prices is associated with a 30% decline in transaction volumes; in the contemporaneous U.S. state-level panel, the same price decline is associated with a 20% volume decline. These facts are hard to reconcile with a frictionless model in which prices clear the market and changes in fundamentals leave quantities largely unaffected. Moreover, while mortgage market structures in the two countries are substantially different, the joint behavior of prices and volumes is not.[3]

The same mechanism also imparts hysteresis: because reference points reset when properties trade, a temporary shock leaves an imprint that fades slowly, making the market's response history-dependent, with a large effect when a shock reverses the prevailing trend, and a negligible one when it merely deepens it. In our first application, high-frequency monetary-policy surprises provide empirical support for this prediction: after recent appreciation, contractionary shocks reverse the prevailing trend, move homeowners back toward and potentially across their reference points, and trigger stronger volume contractions. Identical interventions therefore play out markedly differently across regions and over the cycle, making the transmission of monetary policy depend on local housing-market trends.

Our second application revisits the long-standing debate over the appropriate balance between housing transaction taxes (levied at the point of sale) and ongoing property taxes (levied

clears, not whether owners decide to put their houses on the market in the first place.

[2]The positive comovement of prices and volumes was emphasized in early work by Stein (1995) and Wheaton (1990), and has since been documented in many settings; Ngai and Tenreyro (2014) and Guren (2018) review the evidence and propose competing mechanisms.

[3]The U.S. market is dominated by long fixed-rate mortgages, which generate "mortgage lock-in" effects on homeowner mobility (Fonseca and Liu, 2023). The U.K. market is dominated by short-duration adjustable-rate mortgages with a modal fixation period of two years (Fisher et al., 2023), leaving little scope for analogous mortgage lock-in. The price–volume comovement is nonetheless of similar magnitude in the two markets.

on owners year after year), given that housing is a large and quantitatively important source of public revenue.[4] Stock-based taxes are conventionally viewed as more efficient because they are less distortionary (Mirrlees et al., 2012; Saez and Zucman, 2019), but such comparisons hinge on how aggregate quantities respond (Best and Kleven, 2018; Eerola et al., 2021; Han et al., 2023; Coven et al., 2025; Cho et al., 2024). In our model, behavioral and financial frictions amplify the distortions of the transaction tax and introduce a short-run inefficiency in the ongoing property tax that would be neutral in the frictionless benchmark. Because these distortions scale with the loss share, a uniform national tax can have substantially different effects across regions depending on the local stock of accumulated nominal gains and losses. Empirical evidence from bunching of transaction prices around notches in the U.K. stamp duty schedule supports the prediction that the transaction tax has lower price effects in high-loss-share regions.

Our work adds to the growing heterogeneous-agent literature on the role of behavioral frictions for macroeconomic and financial outcomes (Ahn et al., 2018; Kaplan and Weidner, 2014; Kaplan et al., 2018; Gabaix, 2020; Maxted, 2023) by showing how a distribution of historically determined reference points becomes an aggregate state variable, and to a parallel literature on the consequences of behavioral frictions for optimal policy (O'Donoghue and Rabin, 2006; Mullainathan et al., 2012; Chetty, 2015; DellaVigna et al., 2017; Taubinsky and Rees-Jones, 2018; Rees-Jones, 2018; Allcott et al., 2019; Farhi and Gabaix, 2020). Second, we add to the empirical literature on reference dependence in housing markets (Genesove and Mayer, 2001; Bracke and Tenreyro, 2021; Andersen et al., 2022) by quantifying how well-documented micro behavior aggregates to shape cyclical dynamics, hysteresis, and the effects of housing taxation. Third, we extend the search-and-matching tradition in housing economics (Wheaton, 1990; Krainer, 2001; Diaz and Jerez, 2013; Ngai and Tenreyro, 2014; Han and Strange, 2015; Guren, 2018; Allen et al., 2019; Anenberg and Bayer, 2020; Moen et al., 2021; Piazzesi et al., 2020; Anenberg and Ringo, 2022; Grindaker et al., 2021) by embedding reference dependence and realistic financial constraints into a dynamic equilibrium setup that simultaneously matches a rich set of micro and aggregate moments. Fourth, our empirical strategy employs a bunching estimator in the public-finance tradition (Kleven, 2016; Allen et al., 2017; Rees-Jones, 2018;

[4]Revenue from housing market taxation accounts for roughly 30% of local tax take in the U.S. (Urban-Brookings Tax Policy Center, 2024) and above 25% in the U.K. (Office for Budget Responsibility, 2024).

Anagol et al., 2022), applied here to the cross-sectional distribution of realized nominal gains.

Section 2 describes the data. Section 3 empirically links micro-level reference-dependent behavior with aggregate housing-market outcomes. Section 4 develops the model. Section 5 calibrates and validates the model using untargeted aggregate moments. Section 6 shows applications of the model. Section 7 concludes.

2 Data

Our analysis uses five main data sets: i) address-level information on the U.K.'s residential property stock from the U.K. Ordnance Survey; ii) the universe of residential property transactions in England and Wales from HM Land Registry; iii) hedonic characteristics and listings information for virtually all properties listed for sale, and detailed information about initiated contact requests between buyers and sellers from Rightmove.com; iv) additional hedonic characteristics on properties sourced from U.K. Energy Performance Certificates (EPCs); and v) loan-level data on the outstanding stock of all mortgage contracts in the U.K. from the Bank of England.

2.1 Residential property stock and transactions

The U.K. residential property stock comprises 27.3 million postal residential addresses, recorded in the Royal Mail Postcode Address File released by the Ordnance Survey. We combine these address data with the universe of residential property transactions in England and Wales, recorded by HM Land Registry.[5] These data are publicly available and report, for each transaction, the property address (including the postcode), the contract date, and the recorded sale price. Between January 1995 and December 2024, the registry records close to 30 million transactions, implying an annual turnover rate of around 3.6% of the housing stock. For comparison, we also obtain publicly available U.S. data on residential sales from Zillow, which we aggregate to a state-level monthly panel of median prices and transaction volumes covering January 2010 to December 2024.

[5]Although our primary transaction and listings data restrict our empirical sample to England and Wales, we use "U.K." throughout the paper for expositional simplicity.

2.2 Property listings and the hedonic model

We acquire listing information for individual properties from Rightmove.com, the largest online portal for house listings in the U.K. These data include listings for sale and searches for these properties from potential buyers, and track all listings on the platform: those resulting in a successful transaction as well as those that were withdrawn and archived. We merge these listings with the transactions registry. The merge rate is high: 81% of Land Registry transactions in 2024 match a preceding listing, consistent with the high market share of Rightmove.com and the tendency of sold properties in the U.K. to be first listed online.

For each property, the data contain information on the exact address, the date of the initial listing, the listing price, and a comprehensive set of hedonic characteristics which include the number of bedrooms and bathrooms, the floor space area, the building type (detached and terraced houses, vs. apartments), and indicator variables for the construction type, whether the property is a development property, retirement home, or classified as affordable housing. The data cover 22.4 million listings between January 2010 and December 2024.

To merge listings with final transactions, we restrict the geographical coverage to England and Wales and remove potentially misreported address units (those with either several transactions or several listings associated with a single address on the same day). We use property-level address information to implement two merges: (i) we match listings to a preceding transaction, which gives the initial purchase price for the current owner; and (ii) we match transactions with preceding listings, to estimate a hedonic pricing model using the characteristics in the Rightmove data. Appendix Table A.1 reports the size of each cleaning step and the resulting merged samples used throughout the paper.

For the hedonic valuation model, we exclude properties transacted within 30 days of the first listing, to avoid measurement error in hedonic characteristics driven by potentially spurious listings information. We also augment the hedonic characteristics with information from energy performance certificates (EPCs), which record the floor area for properties for which this information is missing in the Rightmove data. Using the final data set of realized transactions, we estimate an adjusted-for-characteristics price for each property, predicting the log of the

sale price $P_{i,m,t}$ of all sold properties i in month m of year t:

$$\ln\left(P_{i,m,t}\right) = \zeta_w + \xi_{l,t} + \psi_{r,m} + \beta'_{\mathbf{x}}\mathbf{X}_{\mathbf{i},\mathbf{m},\mathbf{t}} + \varepsilon_{i,m,t}, \tag{1}$$

where ζ_w are electoral ward fixed effects, $\xi_{l,t}$ are local authority district (LAD) $\times$ year fixed effects, $\psi_{r,m}$ are region $\times$ month fixed effects, and $\mathbf{X}_{\mathbf{i},\mathbf{m},\mathbf{t}}$ is a vector of potentially time-varying property characteristics, namely, a second order polynomial of floor area (m^2) augmented with a logarithmic term, number of bedrooms and bathrooms, and dummy variables for property type, whether the property is a development property, retirement home, or affordable housing property.[6] The baseline model has strong explanatory power with an R^2 of 0.87; Appendix Figure A.1 reports the in-sample fit.[7]

2.3 Search and matching

The Rightmove data also record online search activity for all Rightmove.com listings observed between January 2019 and December 2024. We distinguish between two types of engagements between an online user and a property listing.[8] The first, which we denote as a "search visit," is an instance of the user clicking the "detail view" hyperlink on a listing, which provides the full set of property characteristics (e.g., the floor area) and qualitative descriptions of features that the seller/landlord deems relevant. The second layer of engagement occurs when an online user enters a unique request for direct contact with the owner of the property. When we compute market tightness, i.e., the number of buyers divided by the number of listings, we use the average number of contact requests as our measure of the number of potential buyers interested in each property.

[6] For alternative geographic aggregations, we replace the estimated fixed effects accordingly in each specification, moving between roughly 7,200 Middle Layer Super Output Areas (MSOAs), 320 local authority districts, and 35 ITL2 regions as each empirical exercise requires.

[7] Not all variables are available for all properties. Most importantly, we frequently miss the number of bathrooms, which we impute by regressing it as a continuous variable on floor area with fixed effects on number of bedrooms, electoral ward, and property type, and dummy variables on whether the property is a retirement home, auctioned, a development, or affordable housing property. The hedonic model uses the imputed number of bathrooms throughout.

[8] Each "online user" in our setting corresponds to a unique IP address tagged with a fully anonymized (hashed) user identifier. We filter the data to remove online users with only one search visit over the entire sample period and those who search over more than 1,000 listings, as these are likely bots.

2.4 Mortgage loans

We acquire information about the outstanding stock of mortgages in the U.K. using the Financial Conduct Authority's PSD007 data set, provided by the Bank of England. These data cover the period between 2015 and 2024 and are available in half-yearly snapshots. To obtain information about the outstanding mortgage balance at the time of listing for eventually transacted listings, we identify individual properties using the postcode, the previous transacted value and the mortgage contract account opening date. The combination of these three variables is a near-unique identifier across our datasets. In cases where no exact match is found, the closest combination within 30 days or 5% of the reported dates and transacted prices, respectively, is considered a match. Appendix Table A.2 gives a breakdown of the progressive matches made using this algorithm.[9]

We refer to a transaction as belonging to the "non-mortgage" sample when it cannot be matched to any outstanding contract in the mortgage stock, so that the seller is unencumbered at the time of listing, and potential home equity is de facto $\widehat{H} = 1$. To estimate the average amortization schedule, we track mortgages across stock snapshots based on the postcode and borrower date of birth (the latter variable also allows us to uniquely identify a borrower-property combination), and rely on a sub-sample of mortgages originating in the first half of 2015. Appendix Table A.3 reports summary statistics. The final sample for this exercise comprises 70,000 observations with an average origination loan size of £159,692, an initial interest rate of 2.88%, and a monthly payment of £782. On average, mortgages originating in the first half of 2015 amortize 2.53% of the remaining loan balance every six months and are amortized by a cumulative 30.14 percentage points by the end of 2022. Appendix Section A.1 provides further details on our approach to tracking mortgages across different snapshots of the mortgage stock.

2.5 Nominal "paper losses" and home-equity

We compute nominal "paper" gains and losses in our sample of property listings as the percent difference between the estimated hedonic value of the property at the time of listing, and the

[9]For analysis requiring granular home-equity data, we use averaged values of variables (in particular realized premia) by binned values of home equity and potential gain positions of each listing to preserve privacy across the Rightmove listings and Bank of England mortgage datasets.

initial purchase price.[10] The "potential gain" distribution for the full set of listings suggests that approximately 25% of sellers would face nominal losses if they sold at the estimated hedonic value, varying between 13.6% and 40.0% across ITL2 regions. Appendix Figure A.2 shows how the regional "paper loss share" has evolved across U.K. locations and over time.

For properties financed with a mortgage, the seller's "potential home equity" position $\widehat{H}$ is calculated as the *share* of the hedonic value left unencumbered by the outstanding mortgage at the time of listing, i.e., the hedonic value minus the outstanding mortgage balance, divided by the hedonic value. In the U.K. mortgage market, interest costs start to increase for loan-to-value ratios above 70% (Liu, 2022). A seller with home equity below 30% therefore finds it more costly to take a loan to finance a property of similar value to the one they are selling, and may have to consider downsizing. In the data, we estimate that a small fraction (1.2%) of all outstanding mortgage contracts across all time periods are underwater, and 23.7% have a home equity level below 30%.

3 From micro behavior to aggregate outcomes

3.1 Reference dependence

Andersen et al. (2022) use administrative data from Denmark and identify the nominal purchase price of a property as a quantitatively relevant reference point for sellers of residential properties. They test for reference-dependent preferences using a joint set of *ex-ante* and *ex-post* empirical moments that capture the behavior of sellers throughout the transaction process. They rationalize the empirical patterns using a model in which preferences exhibit reference dependence and loss aversion, with a change in slope (a "kink") at the reference point.

In Figure 1, we verify that such patterns are also evident in granular U.K. housing transactions data.[11] Panel A shows the distributions of nominal gains on all housing transactions in Denmark and the U.K., computed as the percent gain of the sale transaction price relative

[10] The model treats as an aggregate state variable the loss share over the *stock* of homeowners, i.e., the fraction of all owners in a region whose property would deliver a paper loss at current valuations. Because hedonic characteristics are only systematically available for properties that have been listed for sale, we cannot construct this stock-level measure directly and instead proxy for it with the analogous share computed across current listings. Reassuringly, Andersen et al. (2022) show that the listing probability does not vary much with nominal paper losses, which limits the role of endogenous selection into, and persistence in, the listings pool.

[11] Similar patterns are also evident in U.S. data, see, e.g., Ramadorai (2026).

to the original purchase value for each property, expressed in percentage points to match the binning of the figure. To extract a measure of "excess mass" from the distribution of realized transaction prices, we compare these distributions with the counterfactual frequency of nominal gains that would occur if properties were sold exactly at their estimated hedonic values. The plots in the U.K. and Denmark are highly similar, with missing mass immediately below the reference point, and significant bunching at exactly zero nominal gains. A build-up of additional mass is observed as diffuse bunching for nominal gains in the low positive domain, which is consistent with reference dependence in the presence of optimization frictions (Anagol et al., 2022).

Panel B of Figure 1 turns to sellers' *ex ante* listing decisions, plotting the average listing premium (ℓ), the percent gap between the asking price and the hedonic value, against the potential gain. The schedule traces a pronounced "hockey stick": premia are high and steeply declining in the loss domain, where sellers facing nominal losses list well above fair value, and flatten in the gain domain. Comparing the mortgage and non-mortgage samples separates the financial channel from the behavioral one: the hockey stick is present even among non-mortgaged sellers, who face no home-equity constraint, but in the loss domain mortgaged sellers post visibly higher premia, consistent with financial constraints reinforcing the reference-dependence motive around the same nominal reference point. Low home equity introduces down-payment constraints for both seller strategy and final transaction outcomes (Stein, 1995), and we incorporate these forces in our theoretical and empirical analyses.

The seller's pricing and bargaining strategy is linked with *ex-post* transaction outcomes beyond the realized price. Plot a) of Figure 2 reports the average probability of sale within the listing window, plotted against the nominal potential gain associated with the listing. This combines the sellers' listing behavior documented in Figure 1 with a standard downward sloping demand curve: the 6-month probability of a transaction rises monotonically from approximately 28% for properties facing a 20% nominal loss to approximately 39% for properties with gains of 20% or more, with a visibly steeper slope in the loss domain and a plateau in the gain domain. Plot b) reports a closely related outcome, the elapsed time between the initial listing date and the final sale date, against the same measure. Properties with negative nominal potential gains take approximately 20% longer to sell than comparable properties in

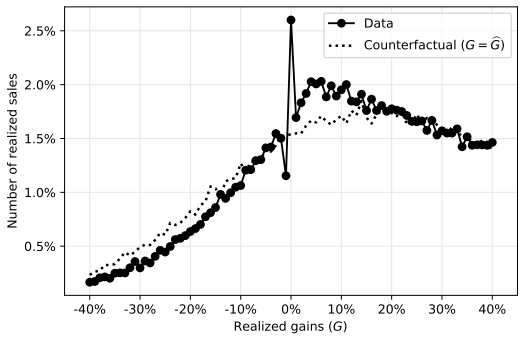
Figure 1
Reference dependence and financial constraints

Panel A compares the distribution of realized nominal gains (percent difference between the sale price and the initial purchase price) with the counterfactual distribution of potential gains for the same properties, using percentage-point bins closed to the left and open to the right. Plot a) is from Andersen et al. (2022); plot b) uses HM Land Registry transactions. Panel B reports the average listing premium (percent difference between the listing price and the hedonic valuation at listing) as a function of the potential gain ($\widehat{G}$), separately for sellers with and without an outstanding mortgage.

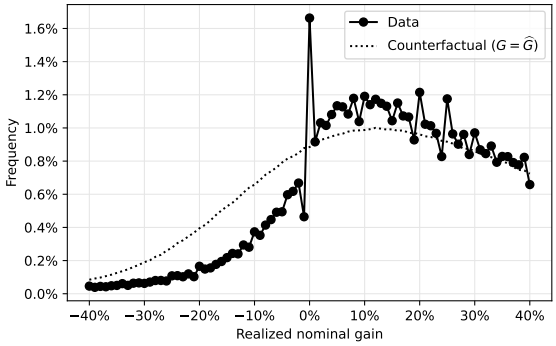
Panel A

Bunching and excess mass

a) Denmark

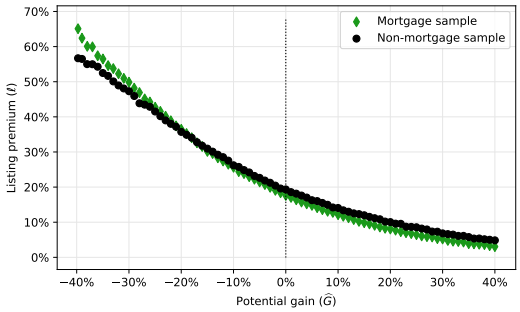


b) U.K.



Panel B

Ex ante listing prices



the gain domain. Together, the two plots indicate a substantial delay in market clearing when sellers face potential losses, consistent with the same asymmetric reluctance to transact at nominal losses that is visible in listing prices around the reference point.[12]

Panel B of Figure 2 formalizes this evidence, estimating a Cox proportional hazards model linking the probability of sale to the nominal "paper gain" on a given property, allowing for separate slopes in the loss ($\widehat{G} < 0$) and gain ($\widehat{G} \geq 0$) domains. Tracking listings in the transactions register also lets us treat properties that are eventually withdrawn without a sale

[12]Appendix Figures A.3 and A.4 show that the bunching and hockey-stick patterns documented for the aggregate U.K. sample also hold in each of the nine ITL1 regions of England.

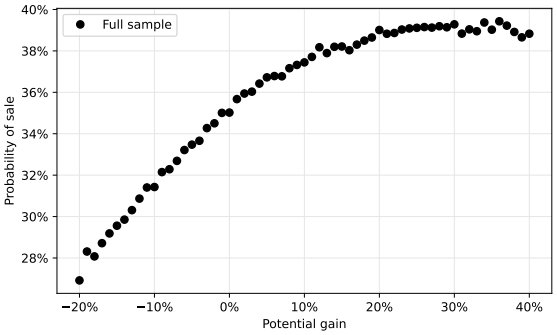
as censored observations. The estimated hazard implies that a seller with a 50% nominal return since purchase is close to twice as likely to sell within 6 months of listing as a seller facing a 50% nominal loss. Plot b) shows the same patterns in the non-mortgage sample, i.e., when we restrict the view to those listings for which the seller has full home-equity at the time of listing.

Figure 2
Time on the market and sale probability

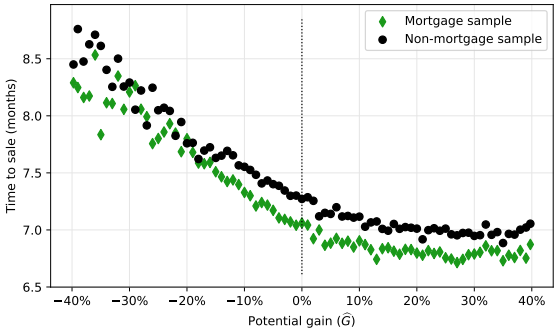
Panel A reports two non-parametric measures of transaction outcomes against the nominal potential gain ($\widehat{G}$) associated with the listing: plot a) reports the average probability of sale within the listing window, and plot b) reports the elapsed time between the initial listing date and the final sale date. Panel B reports estimated sale probabilities from a Cox proportional hazards model, controlling for holding period at the time of sale, and with separate coefficients in the nominal loss ($\widehat{G} < 0$) and the nominal gain ($\widehat{G} \geq 0$) domain. Plot a) reports estimates on the full sample; plot b) reports the same estimates on the restricted sample of properties for which the seller has no outstanding mortgage at the time of listing, i.e., full home equity.

Panel A
Transaction outcomes

a) Probability of sale

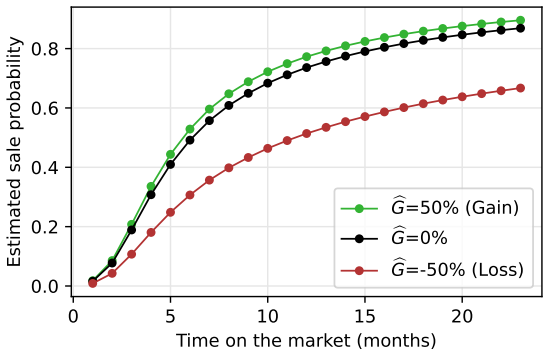


b) Time on the market

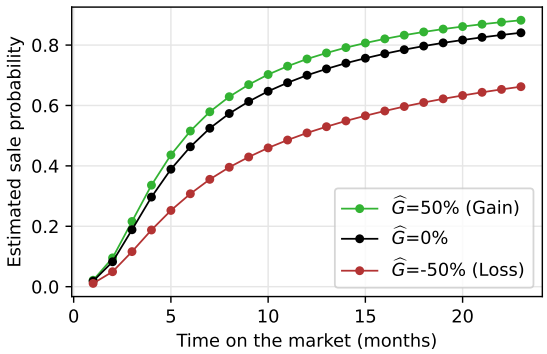


Panel B
Cox proportional sale hazard

a) Full sample



b) Non-mortgage sample



3.2 Aggregate outcomes

Does nominal anchoring manifest in aggregate prices and quantities in the housing market? We investigate the relationship between prices and transaction volumes at the ITL2 regional level in the U.K. (and at the equivalent state-level in U.S. data) at the monthly frequency between January 2010 and December 2024.

We first construct a price growth series using the location × month fixed effects from the hedonic pricing model in equation (1), and compute the corresponding number of transactions from HM Land Registry in each month and location. For the U.S., we use the median transaction price and the reported number of residential property sales in each state and month. We demean volumes and prices using their location-specific means to clean out persistent differentials between locations. Panel A of Figure 3 shows a robust positive price–volume comovement in both markets. This unconditional comovement at a location-month level is material: a 6% decrease in prices is associated with a 20% decrease in transaction volumes in the U.S., and 30% in the U.K.

We next calculate the share of sellers facing nominal losses in each month in each ITL2 region of the U.K. Panel B of Figure 3 conditions the price–volume comovement on this loss share, comparing region-months where the share of nominal losses lies in the top quartile of the full-sample distribution ("High" in the plot's legend), with those where the share of nominal losses lies in the bottom quartile ("Low"). We lag the classification of region-months by one month relative to the time at which we measure price growth and volumes to avoid any mechanical correlation. The figure shows that the positive relationship between prices and volumes is far more pronounced when the share of homeowners facing nominal losses is high, and much more muted when it is low. This effect is driven by the post-listing behavior of individual sellers, as there seems to be no strong link between the evolution of prices and the number of new listings that arrive on the market.

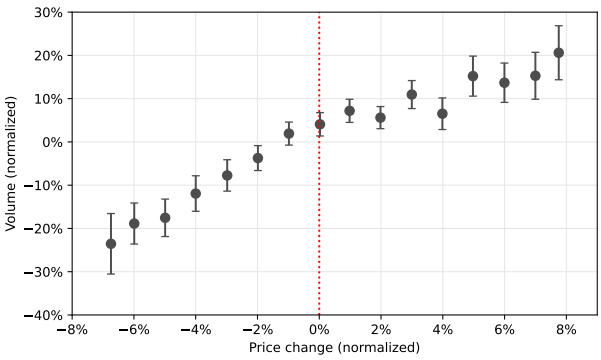
To acquire more granular evidence, we next analyze prices and volumes across regions. Plot a) of Appendix Figure A.5 documents the link between bunching in realized transaction prices and the loss share at the regional level. The excess mass rises steeply with the local loss share, from close to zero in the lowest-loss-share regions to roughly 50% in the highest, tying the ex-post pile-up of realized prices at original purchase values to the ex-ante share of sellers

Figure 3
Comovement of prices and volumes

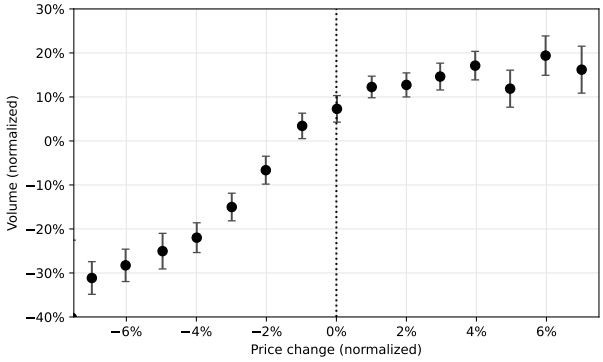
Panel A reports the correlation between year-over-year log changes in housing prices and log transaction volumes. Both variables are expressed as deviations from their location-specific means. We calculate prices and volumes at the level of ITL2 regions (U.K.) and the level of states (U.S.), respectively, measured monthly for the period between January 2010 and December 2024. On the horizontal axis, we report price growth rate bins, constructed by rounding to the second decimal. Panel B reports the correlation between year-over-year changes in housing prices and the volume of listings and realized transactions. We distinguish between periods and regions for which the share of sellers that face the possibility of a loss is below/above the respective bottom/top quartile ("Low/High loss share"). Error bars indicate 95% confidence intervals for bin-specific means.

Panel A
Aggregate effects

a) U.S. (Zillow)

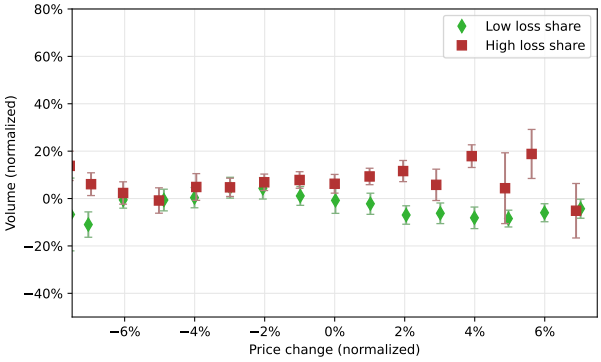


b) U.K. (Land Registry)

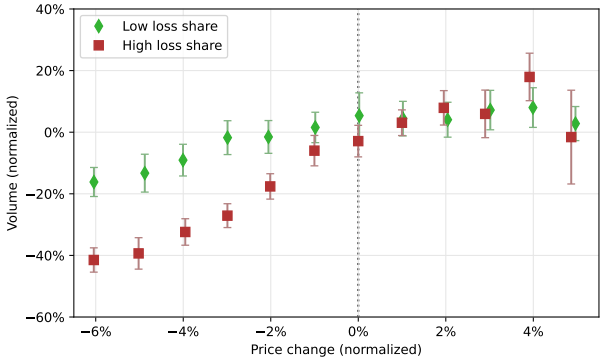


Panel B
The role of the loss share

a) Listing volumes



b) Transaction volumes



at risk of a nominal loss.

Plot a) of Figure 4 estimates the price–volume comovement separately for each location, using the following specification: $\ln V_{i,t} = \gamma_i + \beta_{i,P} \Delta \ln P_{i,t} + \varepsilon_{i,t}$, where $V_{i,t}$ is the number of transactions recorded in ITL2 region i and month t, and $P_{i,t}$ is the corresponding price level. We plot the estimated location-specific coefficients $\beta_{i,P}$ against the location-specific share of properties that face nominal losses. The plot shows a positive association, with a univariate $R^2 = 0.47$. Plot b) repeats the exercise on the sub-sample of properties for which the seller has no outstanding mortgage contract at the time of listing, showing that the relationship is of similar magnitude in this restricted sample. Plot b) of Appendix Figure A.5 documents a parallel positive relationship between the price–volume comovement and the share of sellers with low home equity at the time of listing.

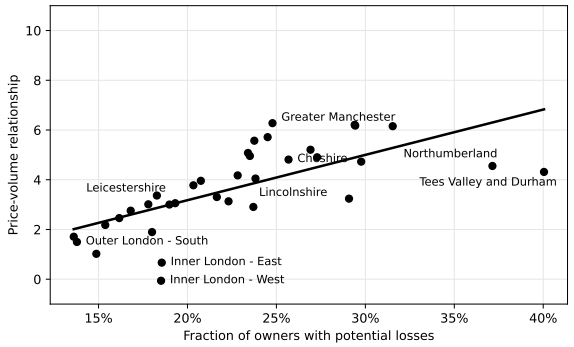
Plot d) turns to the intensive margin, i.e., whether or not properties sell once they are listed, reporting how strongly the conversion rate of listings into sales comoves with price changes across regions. Together with plot c), which reports the corresponding extensive-margin coefficients for listing volumes, this shows that across time and regions the conversion rate comoves far more strongly with price changes, and is responsible for the cross-regional variation in comovement between total housing volumes and price changes. Put differently, the price–volume comovement mainly comes from faster turnover of already listed properties when price growth increases; in contrast, listing volumes comove only weakly with price growth. This important feature of the data is one that we revisit when discussing the predictions of the model.

Finally, plots e) and f) report the time-series volatility (i.e. standard deviation) of volumes and prices across regions against the regional share of potential nominal losses. Where the loss share is high, the volatility of volumes is substantially higher, and the volatility of prices slightly lower. This finding aligns with broad intuition about the behavioral nominal rigidity: in areas where many agents are reluctant to realize nominal losses when housing values decline, volumes collapse and absorb shocks that would otherwise be manifest in price volatility, and vice versa when housing values rise. We label this phenomenon *nominal rigidity at the aggregate level*: because sellers condition their decisions on gains and losses relative to their original purchase prices, nominal variables become linked to real outcomes, similar to the way in which firm-level

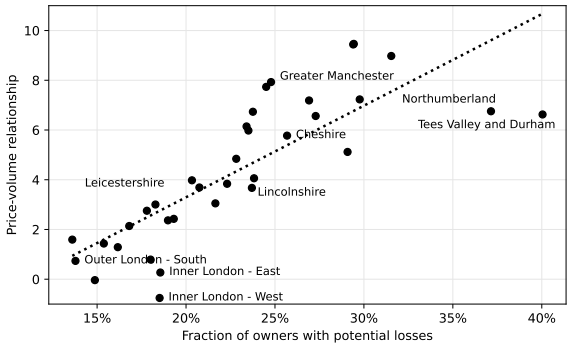
Figure 4
Cross-regional variation in the price–volume relationship

In all plots the horizontal axis indicates the average "loss share" for U.K. ITL2 regions between January 2010 and December 2024, i.e., the fraction of owners that list a property for sale for which the estimated hedonic value is below the initial purchase price. Plot a) reports coefficients from regressions of the realized monthly transaction volume on the year-on-year price change in each ITL2 region during the period, and b) shows the same coefficients estimated on the restricted sample of properties for which the seller has no outstanding mortgage at the time of listing. Plot c) reports coefficients from regressions of listing volumes on the year-on-year price change in each region, and plot d) the corresponding coefficients for the conversion rate of listings into sales. Plots e) and f) report region-specific volatilities of transaction volumes and transaction prices, respectively, measured as the standard deviation of monthly volumes and of monthly price growth rates during the full sample period.

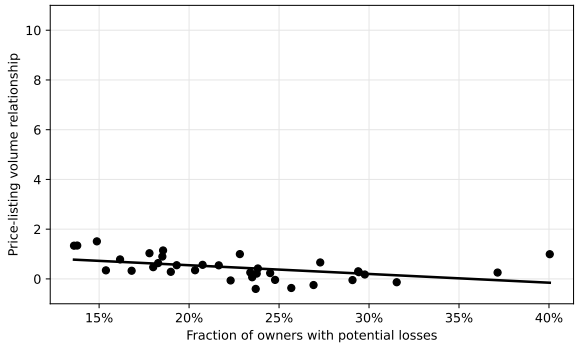
a) Price-volume relationship



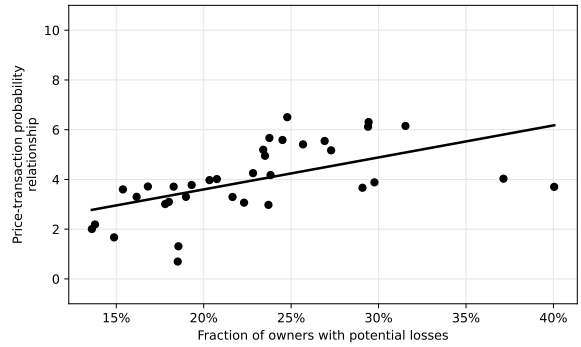
b) Non-mortgage sample



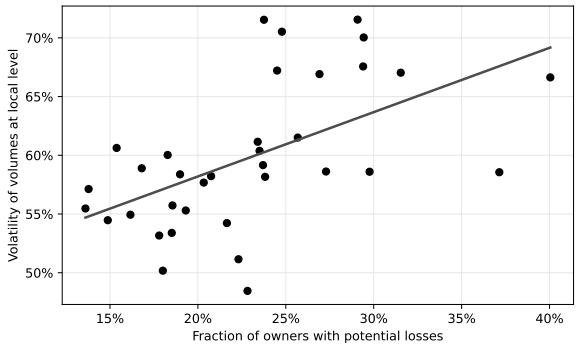
c) Listing volumes



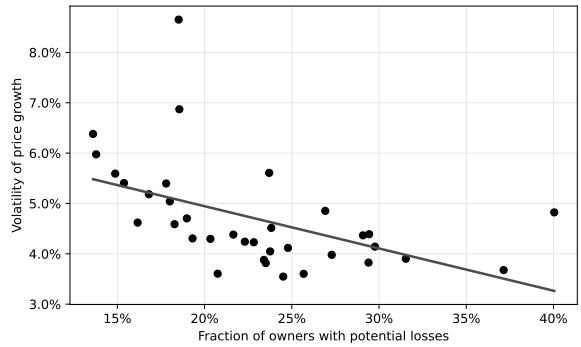
d) Probability of sale



e) Volatility of transaction volumes



f) Volatility of transaction prices



price and wage stickiness keeps quantities from reflecting fundamentals.

3.3 Listing premia and aggregate outcomes

To isolate the separate contribution of financial constraints to the price–volume relationship documented in the previous section, we regress average transaction outcomes, namely time on the market (measured in months) and the probability of sale, on average listing premia at the MSOA×year level, controlling for information on each seller's home-equity position. Average listing premia are constructed using hedonic prices estimated via a leave-one-out procedure: for each listing, we estimate a hedonic pricing model that excludes its own MSOA, and use the predicted price to compute the listing premium. These hedonic prices also form the basis for the nominal potential gain associated with each listing.

Owing to data availability, the analysis focuses on listings initiated between 2015 and 2024. For each MSOA j and year t, we construct average listing premia, potential gains, and home-equity positions. Appendix Table A.4 reports summary statistics for these MSOA×year aggregates. We stress that this exercise should be read as a descriptive decomposition rather than a causal instrumental-variables design: the excluded instruments are themselves economically meaningful state variables in our mechanism, not exogenous shifters of the listing premium alone. Conditional on the posted premium, they may also affect transaction outcomes through seller selection into listing, reservation and withdrawal behavior, and local liquidity, so the exclusion restriction need not hold strictly. We therefore interpret the 2SLS estimates as isolating the component of the price–volume relationship that comoves with these financial and reference-point states.

The two left-most columns of Table 1 report the first stage, in which we instrument the average listing premium with three measures of the local nominal position. The first is the average potential loss, for listings for which the hedonic value is below the original purchase price. The second is the average potential gain, defined symmetrically over listings in the gain domain. The third is average constrained equity, which records how far the typical seller falls short of the 30% home-equity threshold. The asymmetry documented at the micro level reappears clearly once the data are aggregated: the average potential loss is a far stronger predictor of the listing premium than the average potential gain, and its coefficient is essentially

unchanged when the constrained-equity measure is added, so the behavioral and the financial margin carry largely independent information.

The remaining columns turn to the second stage, relating average time on the market and the average sale indicator to the instrumented listing premium. For each of the two outcomes we report an OLS estimate alongside three instrumental-variables specifications, which differ only in how the constrained-equity measure enters: it is first left out of the instrument set, then added to it, and subsequently included as a control variable. The estimated effect is stable across these alternatives, and the first-stage F-statistics remain significantly large. The OLS estimates are strongly attenuated relative to their instrumented counterparts; for the sale indicator, the OLS coefficient is close to zero and statistically insignificant, consistent with measurement error in the leave-one-out premium and with offsetting local demand conditions. The constrained-equity measure itself enters with an economically meaningful magnitude: a larger average equity shortfall among constrained sellers is associated with materially longer times on the market. In the specification with all three instruments, a 10 percentage point higher listing premium lengthens the expected time on the market by about a month, and lowers the probability of sale by approximately 3 percentage points.

4 The Model

We develop a dynamic heterogeneous agent search and matching model of the housing market where homeowners are reference dependent and loss-averse around a pre-determined anchor, and face realistic financial constraints. Seller decisions are conditional on heterogeneous realizations of shocks to mobility, heterogeneous levels of reference prices, and heterogeneous outstanding mortgage amounts. Buyer decisions are conditional on heterogeneous realizations of housing matching quality shocks. The participation and pricing decisions of homeowners in the search and matching process determine aggregate housing market outcomes including prices and transaction volumes.

Table 1
Aggregate transaction outcomes

This table reports MSOA×year level regressions of aggregate transaction outcomes on average listing premia, for listings initiated between 2015 and 2024. The average listing premium is instrumented using (i) the average potential loss (defined as $-\hat{G}$ for listings with $\hat{G} < 0$), (ii) the average potential gain (for $\hat{G} \geq 0$), and (iii) average constrained home equity (defined as $0.3 - \hat{H}$ for listings with $\hat{H} < 0.3$). The hedonic price used to calculate the listing premium is estimated via a leave-one-out procedure, which excludes the MSOA in which a listing is located. All specifications include MSOA and listing-year fixed effects. Standard errors clustered at the more conservative ITL2 geographic level are reported in parentheses; *, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

	Avg. list. prem.		Time on market (months)				Sale indicator			
	OLS	OLS	OLS	2SLS ($\hat{G}$)	2SLS ($\hat{G}, \hat{H}$)	2SLS ($\hat{G}$)	OLS	2SLS ($\hat{G}$)	2SLS ($\hat{G}, \hat{H}$)	2SLS ($\hat{G}$)
Avg. list. prem			4.024***	8.819***	9.552***	8.389***	-0.013	-0.289***	-0.300***	-0.283***
			(0.386)	(1.509)	(1.538)	(1.474)	(0.008)	(0.044)	(0.044)	(0.044)
Avg. potential loss	0.257***	0.255***								
	(0.014)	(0.014)								
Avg. potential gain	-0.008***	-0.008***								
	(0.002)	(0.002)								
Avg. equity (Constrained)		0.016***				1.189***				-0.017***
		(0.004)				(0.164)				(0.004)
MSOA fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Listing year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of obs.	57,874	57,874	57,873	57,873	57,873	57,873	57,874	57,874	57,874	57,874
R^2	0.864	0.864	0.683	0.678	0.677	0.680	0.835	0.825	0.824	0.825
Adjusted R^2	0.848	0.849	0.647	0.642	0.641	0.644	0.816	0.805	0.804	0.806
F-test (1st stage)				1,622	1,093	1,600		1,622	1,093	1,600

4.1 Environment

Consider a discrete-time formulation of an economy with a housing stock N_H potentially available for sale, and a mass N_B of interested buyers. Current homeowners choose whether to list their houses for sale, and buyers randomly search over the outstanding stock of available listings. We denote the endogenous number of sellers in a generic period t by N_{St} and the resulting market tightness by $q_t = N_B/N_{St}$. The probability that a seller meets a buyer is modeled using a constant returns to scale matching function $\chi(q_t)$, which we specify as $\chi(q_t) \equiv \bar{\chi} \cdot q_t^{\upsilon}$, for some constants $\bar{\chi} > 0$ and $\upsilon \in [0, 1]$. The probability that a buyer meets a seller is then equal to $\chi(q_t)N_{St}/N_B = \chi(q_t)/q_t$, and is by assumption bounded between 0 and 1. Upon a successful transaction, the buyer becomes a homeowner and the seller begins their search for a new home as a buyer, which implies that the overall number of buyers N_B remains constant through time.[13]

[13]Thus a positive demand shock leads households to "move up the housing ladder", and individual decisions spill over to the broader housing market (Aiello et al., 2022).

4.2 Homeowners

Owning a home generates flow utility u. At the beginning of period t, each homeowner i draws a mobility shock $\theta_{it} \sim F_\theta(\cdot)$ with probability ρ. The mobility shock represents a newly available outside opportunity with a payoff that only gets realized if the house is successfully sold. Given the draw, the homeowner either lists the property for sale and pays a one-time cost $\phi \geq 0$ for doing so, or ignores the opportunity and takes another independently distributed draw during the next period (with probability ρ). If the homeowner chooses to list, they post a log price p_{it} as a take-it-or-leave-it offer.

With probability $\alpha_t(p_{it})$, the sale transaction completes successfully at the posted price p_{it} and the seller receives a one-time utility payoff $U(p_{it}, r_i, m_{it}) + \theta_{it}$, where r_i denotes the log reference price at which they originally purchased their property, and m_{it} their log mortgage balance; and they start searching for a new property, from the position of a new buyer. With probability $1 - \alpha_t(p_{it})$, the homeowner is unable to sell their property, and they draw another independent and identically distributed realization of the mobility shock in the next period.

We denote the homeowner's value function by $V_t^h(r_i, m_{it})$ and define it through the following Bellman equation:

$$\begin{aligned} V_t^h(r_i, m_{it}) = u + \mathbb{E}_{\rho,\theta} \max\{ & \max_{p_{it}} \underbrace{\alpha_t(p_{it})[U(\cdot) + \theta_{it} + \beta \mathbb{E}_t[V_{t+1}^b]]}_{\text{Listing success}} \\ & + \underbrace{(1 - \alpha_t(p_{it}))\beta \mathbb{E}_t[V_{t+1}^h(\cdot)]}_{\text{Listing failure}} - \phi, \underbrace{\mathbb{E}_t[\beta V_{t+1}^h(\cdot)]}_{\text{No listing}} \}. \end{aligned} \tag{2}$$

The operator $\mathbb{E}_{\rho,\theta}$ denotes the compound expectation over both the arrival of a mobility opportunity (which occurs with probability ρ) and the conditional draw of θ, so that with probability $1-\rho$ no opportunity arrives and the homeowner simply retains the no-listing continuation value $\beta \mathbb{E}_t[V_{t+1}^h(\cdot)]$; the listing option is therefore exercised with probability at most ρ each period, consistent with the listing share $\rho[1 - F_\theta(\theta_t^*)]$ used in the aggregation below.

4.2.1 Seller utility

We assume the seller utility function includes reference dependence and loss aversion, and set:

$$U(p_{it}, r_i, m_{it}) = p_{it} + \underbrace{W(p_{it}, r_i)}_{\substack{\text{Reference} \\ \text{dependence}}} - \underbrace{\mu(\gamma - (p_{it} - m_{it}))_+^2}_{\substack{\text{Downsizing} \\ \text{penalty}}}, \tag{3}$$

where:

$$W(p_{it}, r_i) = \begin{cases} \eta_+(p_{it} - r_i), & \text{if } p_{it} \geq r_i, \\ \eta_-(p_{it} - r_i), & \text{if } p_{it} < r_i. \end{cases} \tag{4}$$

The function $W(\cdot)$ describes reference-dependent and loss-averse preferences through separate coefficients on the gain side (η_+) and the loss side (η_-). We also include a convex downsizing penalty, motivated by Stein (1995) and Andersen et al. (2022). If the seller's home equity is lower than the exogenously set threshold γ, they expect to face additional costs to finance another house of similar size or quality after selling their home. Therefore, they dislike the prospect of selling at a price that reduces home equity beneath the threshold, because that would require either painful downsizing if they cannot acquire financing, or costly financing to bridge the gap.[14]

4.2.2 Homeowners' extensive margin decision

Conditional on drawing a θ shock and listing, the maximization problem for the posted price p produces an optimal price $p_t^*(r, m, \theta)$ for each pair of state variables r and m, and for each draw of θ. The extensive margin decision is governed by a threshold rule, where homeowners list their property for sale if and only if:

$$\alpha_t(p_t^*(r_i, m_{it}, \theta_{it}))\left[U(p_t^*(r_i, m_{it}, \theta_{it}), r_i, m_{it}) + \theta_{it} + \beta\mathbb{E}_t[V_{t+1}^b] - \beta\mathbb{E}_t[V_{t+1}^h(r_i, m_{i,t+1})]\right] \geq \phi.$$

[14] In the calibration and numerical implementation, the penalty argument $p_{it} - m_{it}$ and the threshold γ are expressed in log price-to-mortgage units, i.e., $p_{it} - m_{it} = \log(P/M)$, so that the value of $\gamma = 0.30$ that we use is a cutoff on $\log(P/M)$ corresponding to a loan-to-value ratio of about 74% and home equity of about 26%.

This yields a threshold rule for listing:

$$\theta_{it} \geq \frac{\phi}{\alpha_t(p^*_{it})} - \left[U(p^*_{it}, r_i, m_{it}) + \beta \mathbb{E}_t[V^b_{t+1}] - \beta \mathbb{E}_t[V^h_{t+1}(r_i, m_{i,t+1})]\right] \equiv \theta^*_{it}(r_i, m_{it}), \tag{5}$$

where we have denoted $p^*_{it} \equiv p^*_t(r_i, m_{it}, \theta^*_{it})$, and V^b_t is the buyer's value function at time t. Notice that the equality in (5) is a consistency condition for θ^*_{it}. Together with the optimal price setting condition from the seller's optimization problem, these constitute two equations in two unknowns p^*_{it} and θ^*_{it} for each (r_i, m_{it}) pair.

4.3 Buyers

Each period, buyers search for houses randomly. Upon meeting with a seller i, buyer j draws an idiosyncratic match quality shock $\varepsilon_{jt} \sim F_\varepsilon(\cdot)$, which leads them to accept the sales offer and to purchase the property, if and only if:

$$\beta \mathbb{E}_t[V^h_{t+1}(p_{it}, m_{j,t+1})] + \varepsilon_{jt} \geq p_{it} + \beta \mathbb{E}_t[V^b_{t+1}],$$

where p_{it} is the log price posted by seller i, at which the sale completes, and $m_{j,t+1}$ is the log mortgage amount taken to purchase the property. The transaction price becomes the buyer's reference price as they become a new homeowner upon sale completion.

The buyer's optimal choice is thus also governed by a threshold rule:

$$\varepsilon_{jt} \geq p_{it} + \beta \mathbb{E}_t[V^b_{t+1}] - \beta \mathbb{E}_t[V^h_{t+1}(p_{it}, m_{j,t+1})] \equiv \varepsilon^*_t(p_{it}). \tag{6}$$

Note that both match quality and mobility shocks are measured in units of log price. In particular, the model implicitly assumes log utility on transaction prices for buyers and sellers.

The buyer's Bellman equation can then be expressed as:

$$V^b_t = \underbrace{\beta \mathbb{E}_t[V^b_{t+1}]}_{\substack{\text{Value of remaining} \\ \text{as buyer}}} + \frac{\chi(q_t)}{q_t} \int_p \int_{\varepsilon^*_t(p)}^{\infty} \underbrace{\left[\varepsilon - \varepsilon^*_t(p)\right]}_{\text{Option value for high taste}} d\Phi_\varepsilon(\varepsilon)\, d\Omega_t(p), \tag{7}$$

where $d\Omega_t(\cdot)$ is the endogenous density of available posted prices.

After the purchase of a property at price R_i, the buyer draws a random mortgage bal-

ance $M_{i,t+1} \sim F_M([0, R_i])$ and enters the economy as a homeowner with state $(r_i, m_{i,t+1}) = (\ln R_i, \ln M_{i,t+1})$. We assume the mortgage loan-to-value ratio (relative to the original purchase price), denoted $\widetilde{m}_{i,t+1}$, reduces deterministically: $\widetilde{m}_{i,t+2} = \widetilde{m}_{i,t+1} - \delta_m$, for a constant δ_m. This implies that the mortgage amortization rate is independent of the original purchase price or issuance LTV. Therefore, the expectation $\mathbb{E}_t$ integrates over the distribution of the initial mortgage balance $m_{j,t+1}$ conditional on the purchase price, allowing the buyer's acceptance threshold $\varepsilon_t^*(p_{it})$ to be defined strictly as a function of price. Appendix Figure A.6 provides empirical evidence consistent with this assumption: outstanding balances decline as a roughly constant fraction of the original loan size, so the estimated amortization rate is independent of the issuance LTV.

4.4 Aggregation and model dynamics

We denote the (time-varying) density of homeowners by $f_t(\cdot, \cdot)$ over reference prices r and mortgages m_t. At each (r, m_t) pair, each existing homeowner draws a mobility shock θ_t with probability ρ, chooses whether to list the property for sale, and chooses the corresponding asking price $p_t^*(r, m_t, \theta_t)$. These optimal choices induce an endogenous price distribution, which we denote by $d\Omega_t(p) = \omega_t(p)\, dp$.

At each reference price r and mortgage balance m_t, the share of existing homeowners who decide to list their property for sale is $\rho[1 - F_\theta(\theta_t^*(r, m_t))]$; with probability $\alpha_t(p_t^*(r, m_t, \theta))$, sales go through. This is endogenously determined by the number of meetings and the probability of buyers accepting their offers, conditional on a meeting:

$$\alpha_t(p) = \underbrace{\chi(q_t)}_{\substack{\text{Probability that} \\ \text{seller meets a buyer}}} \times \underbrace{[1 - F_\varepsilon(\varepsilon_t^*(p))]}_{\substack{\text{Probability of acceptance} \\ \text{conditional on a meeting}}} . \tag{8}$$

On the other hand, after a successful transaction, the new homeowner draws a random log mortgage balance and enters the economy with reference price r and mortgage balance $m_{t+1} \sim F_m(\cdot \mid r)$. The mass of new homeowners with reference price r and mortgage balance

m_{t+1} is:[15]

$$\underbrace{N_{St}\omega_t(r)}_{\substack{\text{Number of sellers}\\ \text{who list at}\\ \text{price } r}} \times \underbrace{\alpha_t(r)}_{\substack{\text{Probability of}\\ \text{sale conditional}\\ \text{on listing}}} \times \underbrace{f_m(m_{t+1} \mid r)}_{\substack{\text{Density of new}\\ \text{homeowners with}\\ \text{mortgage } m_{t+1}, \text{ given } r}},$$

where we denote by $\omega(r), \alpha(r)$ the listing density and probability of transaction, respectively, for a posted price equal to r, should the buyer accept the offer.

Taken together, it follows that the law of motion for the homeowner distribution is given by:

$$f_{t+1}(r, m_{t+1}) = \text{Existing homeowners not selling} + \text{New homeowners}$$

$$= \underbrace{f_t(r, m_t)}_{\text{Existing homeowners}} \times \underbrace{\left[1 - \rho \int_{\theta_t^*(r,m_t)}^{\infty} \alpha_t(p_t^*(r, m_t, \theta))\, d\Phi_\theta(\theta)\right]}_{\text{Non-sellers and failed sellers}} \tag{9}$$

$$+ \underbrace{N_{St}\omega_t(r)\alpha_t(r) f_m(m_{t+1} \mid r)}_{\text{New homeowners}},$$

where $m_{t+1} = \ln(M_{t+1})$ and $M_{t+1} = M_t - R \cdot \delta_m$, i.e., the mortgage is reduced by a constant fraction of the purchase price R. Thus, the distribution updates by shifting the surviving mass of existing homeowners deterministically from each current balance m_t to its subsequent amortized state m_{t+1}.

Finally, we close the model with

$$N_{St} = \rho \int_{r,m} \int_{\theta^*(r,m)}^{\infty} f_\theta(\theta)\, d\theta\, dF_t(r, m), \tag{10}$$

which constitutes a fixed point problem, since the optimal threshold rule for listing depends on aggregate equilibrium outcomes.

4.5 Equilibrium

We define the equilibrium concept as follows.

[15] The posted price becomes the buyer's reference price $r = p$ upon the transaction, so the demand function and the price distribution can equivalently be indexed by the resulting reference price.

Definition. An equilibrium consists of:

- Seller and buyer value functions $V_t^h(r,m), V_t^b$ defined through the Bellman equations (2) and (7), respectively;
- Buyer policy function $\varepsilon_t^*(p)$ satisfying (6);
- Seller price setting policy $p_t^*(r,m,\theta)$, which solves the seller's price setting problem, and extensive margin policy function $\theta_t^*(r,m)$, which solves (5);
- Aggregate transaction probability $\alpha_t(p)$ satisfying (8);
- Distribution of reference prices and mortgage balances $F_t(r,m)$, whose law of motion satisfies (9);
- Distribution of posted prices $\Omega_t(p)$, which arises endogenously from the sellers' extensive and intensive margin decisions; and
- Market tightness q_t satisfying (10).

In our analyses, we concentrate on studying a stationary equilibrium with $F(\cdot,\cdot) \equiv F_t(\cdot,\cdot) = F_{t+1}(\cdot,\cdot), V^b \equiv V_t^b, V^h(\cdot,\cdot) \equiv V_t^h(\cdot,\cdot)$.

5 Numerical calibration

We solve the model numerically, using value function iteration to solve for a deterministic steady state. We define a time period in the model to correspond to 6 months in the data, and set the time discount factor equal to $\beta = 0.99$ (i.e., we assume an annualized real discount rate of 2%). For a convenient interpretation of aggregate quantities, we normalize the housing stock to $N_H = 1$ (which means that all volumes in the model are relative to the housing stock), and the average log price level to 1.

Table 2 summarizes how we discipline the model using a combination of externally sourced inputs and internally calibrated structural parameters. Panel A collects parameters anchored outside the model (e.g., the semi-annual discount factor, institutional and accounting quantities such as the rate of amortization, and matching parameters informed by listings and search data). We calibrate the distribution of mortgage draws $F_M([0,R])$ non-parametrically using Bank of England data on initial loan-to-value (LTV) ratios. The mortgage amortization rate

is set equal to $\delta_m = 1.8\%$, corresponding to observed amortization patterns in the data, for a six-month horizon.[16]

For the calibration of the aggregate matching function, we use observed listings and search behavior between 2019 and 2024. We set the elasticity to be equal to $\upsilon = 0.571$, following Badarinza et al. (2024) who use the same data to estimate the housing market matching function. We use the number of residential addresses in the Ordnance Survey address data as a proxy for the stock of residential properties. The average listing probability of 5.6% is an internal target in the calibration; combined with the semi-annual transaction probability conditional on listing of 32.5% in the data, it implies a semi-annual turnover rate of 1.8% of the housing stock. The average listing probability and the steady-state level of market tightness also imply a steady-state value of demand (i.e., the number of potential buyers relative to the housing stock) equal to $N_B = 0.181 = 0.056 \times 3.229$.

To calibrate the mortgage down-payment penalty parameters, we use the observed pattern of mortgage pricing shown in Appendix Figure A.7. For loan-to-value (LTV) ratios below 70% (i.e., for down payment amounts above 30%), the mortgage interest rate schedule is flat, suggesting that mortgage contract originators do not seek to be compensated for marginal changes in the debt amount below this threshold. However, as the LTV ratio rises above 70%, the mortgage rate starts to increase significantly. We therefore set $\gamma = 0.30$ to approximate the minimum level of home equity below which the current transaction value of the house, net of debt, is insufficient to cover the down payment on a subsequent property purchase of a similar size and quality.

We exploit the granularity of our data to pin down a number of key model parameters, building on the growing micro literature on the effects of reference dependence and loss aversion in the housing market. Our calibration targets three sets of micro-level moments that are well-known in the existing literature (Andersen et al., 2022; Guren, 2018; Genesove and Mayer, 2001) and reported in Figure 5: (i) sellers who face nominal losses substantially mark up their listing prices over the hedonic value of the property, consistent with preferences that are

[16] We iterate on a two-dimensional grid over reference prices and mortgage balances. The mortgage grid is an equally spaced set of loan-to-value ratios relative to the original purchase price, running from 42.8% to 95% LTV, with newly issued mortgages drawn according to calibrated weights and amortization implemented as a one-step advance down the grid. The highest balance receives no amortizing inflow and is populated only by newly issued mortgages, while the lowest balance is absorbing, retaining its own surviving mass in addition to the mass amortizing in from the step above.

reference-dependent and loss-averse; this results in a hockey-stick pattern of realized premia, i.e., differences between the final transaction price and the property's hedonic value; (ii) the analogous home-equity hockey stick: financially constrained sellers mark up their listing prices in an attempt to avoid costly downsizing; and (iii) the probability of a transaction going through increases in the potential gain.

Table 2
Numerical calibration

Panel A reports the values and sources for externally calibrated parameters. Panel B reports the calibrated structural parameters, obtained through a simulated method of moments approach, targeting realized premia and probabilities of sale by nominal "paper gains" and potential home-equity.

Panel A
External calibration

Parameter		Value	Target/Source
Time discount	β	0.990	Semi-annual calibration
Housing Stock	N_H	1.000	Normalization
Matching function elasticity	υ	0.571	Badarinza et al. (2024)
Market tightness	q	3.229	Badarinza et al. (2024)
Mortgage amortization	δ_m	0.018	Bank of England mortgage data
Down-payment constraint threshold	γ	0.300	Mortgage interest rate schedule

Panel B
Structural parameters

Parameter		Value	Parameter		Value
Flow utility	u	0.076	Mobility shock (shape)	ξ_θ	-0.460
Matching function constant	$\bar{\chi}$	0.357	Listing cost	ϕ	1.16%
Taste shock (st. dev.)	σ_ε	0.968	Gain-side reference dependence	η_+	0.083
Probability of θ draw	ρ	0.069	Loss-side reference dependence	η_-	0.934
Mobility shock (location)	μ_θ	1.209	Down-payment penalty	μ	14.42
Mobility shock (scale)	σ_θ	0.207			

Internal calibration We specify the distribution $F_\theta(\cdot)$ of seller shocks as generalized extreme value (GEV) and the buyer shock distribution $F_\varepsilon(\cdot)$ as normal with mean zero and standard deviation σ_ε. We use the standard GEV cumulative distribution function:

$$F_\theta(\theta) = \exp\left\{ -\left[1 + \xi_\theta \frac{\theta - \mu_\theta}{\sigma_\theta}\right]^{-1/\xi_\theta} \right\}, \qquad 1 + \xi_\theta \frac{\theta - \mu_\theta}{\sigma_\theta} > 0, \tag{11}$$

in which μ_θ is the location parameter, $\sigma_\theta > 0$ the scale parameter, and ξ_θ the shape parameter. These and other remaining structural parameters are determined by minimizing a quadratic loss function, where the model-implied moments $\mathbf{M_m}(\mathbf{x})$ are calculated numerically from the equilibrium solution of the model for any given set of parameters $\mathbf{x}$. The corresponding empirical target moments are denoted by $\mathbf{M_d}$:

$$\widehat{\mathbf{x}} = \arg\min_{\mathbf{x}} \left(\mathbf{M_m}(\mathbf{x}) - \mathbf{M_d}\right)' \mathbf{W} \left(\mathbf{M_m}(\mathbf{x}) - \mathbf{M_d}\right).$$

We target the following set of empirical moments: the normalized average log price level equal to 1, the realized price premium by potential gains (in the interval -20% to 40%) and potential home equity (in the interval 0% to 40%), the average per-period listing probability of 5.6%, and the probability of a successful sale by potential gains (in the interval -20% to 40%). The weighting matrix $\mathbf{W}$ is chosen to normalize all empirical target moments to 1. Panel B of Table 2 reports the calibrated parameters for the full version of the model, which allows for reference dependence, loss aversion and financial constraints, pinned down through the two behavioral parameters $\eta_+, \eta_- > 0$ (gain-side and loss-side reference dependence) and the down-payment penalty parameter $\mu > 0$.

Recall that the average log price level is normalized to one in steady state, so all model quantities should be interpreted in units of the average price. The calibrated parameters have economically meaningful magnitudes. The flow utility from homeownership $u = 0.076$ corresponds to a 7.6% per-period (half-yearly) flow benefit relative to non-ownership, and the listing cost $\phi = 1.16\%$ of the property price is a plausible approximation of monetary and non-monetary costs of maintaining a property for sale on the market. Homeowners receive mobility opportunities with probability $\rho = 0.069$ per half-year (about 14% per year). In units of the normalized average price, the loss-side coefficient $\eta_- = 0.934$ implies that a seller accepting a price 10% below their reference point incurs a utility cost of roughly 9% of the average house price, and the gain-side coefficient $\eta_+ = 0.083$ implies a utility gain of about 0.8% for a sale 10% above it. Finally, Appendix A.2 provides an interpretation of the calibrated value of μ in terms of the higher mortgage borrowing costs faced by a constrained seller.

Figure 5 reports the model fit along three dimensions: the realized premium as a function of potential gains, the realized premium as a function of home equity, and the probability

of a successful transaction within six months of listing as a function of potential gains. In each panel we plot the empirical pattern alongside the model-implied pattern under three calibrations: the full model with reference dependence and financial constraints, the frictionless benchmark ($\eta_+ = \eta_- = 0$, $\mu = 0$), and an intermediate "financial-only" specification with financial constraints but no reference dependence ($\eta_+ = \eta_- = 0$, $\mu > 0$).[17]

The comparison shows that both frictions are needed to fit the data. The frictionless benchmark generates an essentially flat listing-premium schedule along both the potential-gain and home-equity dimensions and a flat transaction probability across potential gains, missing the data patterns entirely. Adding financial constraints alone produces a marked hockey stick along the home-equity dimension, but leaves the gain-side schedule and the transaction-probability slope nearly flat: financial constraints rationalize the home-equity hockey stick but cannot account for the asymmetric responses around the nominal reference price. Only the full model with reference dependence simultaneously reproduces all three empirical patterns. We next describe the transmission mechanism of shocks and assess the model's ability to explain the motivating aggregate-level facts.

5.1 Impulse response analysis

This subsection traces how valuation, demand, and supply shocks translate into prices and transaction volumes in the calibrated model, and contrasts these dynamics with the frictionless benchmark.[18] Figure 6 reports impulse responses to three temporary shocks. The first is a -5% valuation shock to u, a change in the average utility flow from housing services that symmetrically affects buyers and sellers. The second is a -5% demand shock, an additive shock to the average value of ε, which could be mapped to changes to buyer composition, such as increased foreign demand (Badarinza and Ramadorai, 2025), or to mortgage or credit market conditions. The third is a +5% supply shock, an additive shock to the average value of θ, which could be mapped to shocks to mobility, family dynamics, or the attractiveness of distant

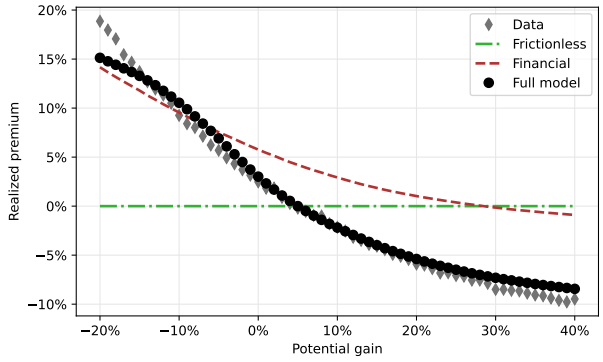
[17] We use the term "frictionless benchmark" to indicate the absence of reference dependence and the down-payment penalty, i.e., it does not imply a perfectly competitive allocation, as the classic search-and-matching trading frictions remain active in this baseline.

[18] Temporary shocks to monetary and fiscal policy can also be mapped to the set of shocks that we consider, so this analysis also hints at how the model can be used to evaluate the effects of policy changes on housing market prices and quantities.

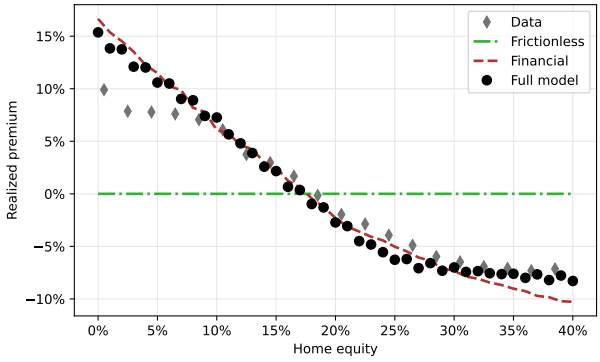
Figure 5
Model fit

The figure reports the realized premium (difference between the final transaction price and the hedonic value) as a function of potential gains (difference between hedonic value and purchase price) and potential home equity (difference between hedonic value and outstanding mortgage amount) in the data as well as under three model calibrations: the full model with reference dependence and financial constraints, the frictionless benchmark ($\eta_+ = \eta_- = 0$, $\mu = 0$), and an intermediate "financial-only" specification with financial constraints but no reference dependence ($\eta_+ = \eta_- = 0$, $\mu > 0$). In plot b), the empirical profile is constructed by averaging the realized premium over potential gains $\widehat{G} \in [0.1, 0.4]$ within each home-equity bin (and de-meaning the resulting profile over the plotted range), so that it isolates the home-equity gradient among sellers in the positive nominal-gain domain. Plot c) reports the probability of a successful transaction within six months of listing as a function of potential gains, in the data and under the three model specifications.

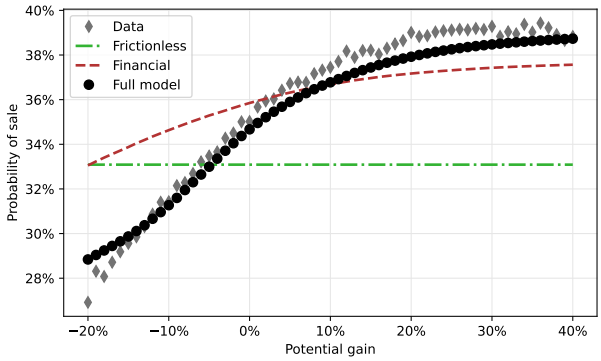
a) Realized prices by potential gains



b) Realized prices by home equity



c) Probability of sale



job opportunities. The shocks are unanticipated and decay at a rate of 0.5 each period, and we consider the perfect foresight responses of prices and transaction volumes, assuming that agents form rational expectations about the future evolution of the economy.

In the frictionless benchmark, the negative valuation shock leads to a decrease in prices of approximately 5% on impact. Transaction volumes, however, remain exactly unchanged. With $U(p, r, m) = p$, the seller's problem depends on the posted price only through the buyer's threshold $\varepsilon_t^*(p)$, so the shock is fully capitalized into the price level and leaves the surplus

from trade, and hence both the buyer's acceptance margin and the seller's listing margin, untouched. The demand shock generates a significant decline in prices, but also induces changes in transaction volumes, as it initially creates a willingness-to-pay-willingness-to-accept gap. Mobility-driven supply shocks lead to selling pressure, which generates a significant decrease in prices and an increase in transaction volumes, so that the price–volume correlation is negative.

In the full model with frictions, these responses differ markedly. For valuation and demand shocks, prices decrease less than in the frictionless case, while transaction volumes drop more, reflecting the role of reference dependence and financial frictions in dampening price adjustment and amplifying quantity responses. The seller's optimal pricing strategy is a non-linear schedule that depends on the location of the reference point. A higher (lower) accumulated potential gain after a price increase (decrease) makes the seller list for a lower (higher) price, in an attempt to attract more buyers and realize the gain (offset their losses with a higher realized premium). This mechanism generates an observed positive correlation between aggregate prices and realized transaction volumes, consistent with the results reported in Figure 3.

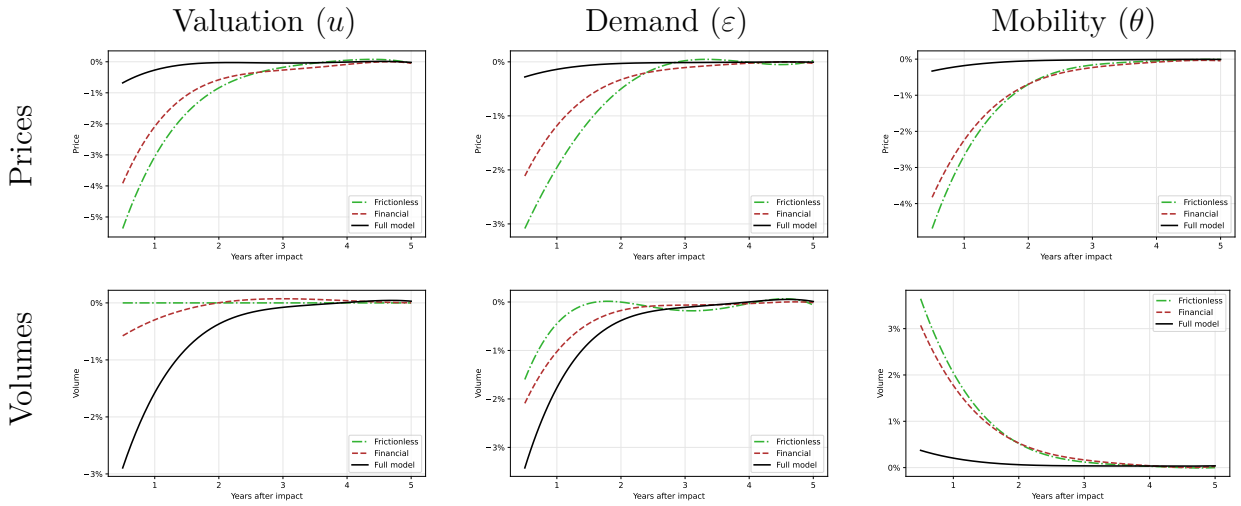
In response to mobility-driven supply shocks, the responses of both prices and transaction volumes are more muted than in the frictionless model; this is the "behavioral lock-in" mechanism. At the calibrated parameters, we calculate that to generate the same increase in transaction volumes in a frictionless economy, the increase in the mobility shock needs to be roughly ten times greater in the economy with behavioral and financial frictions. While the supply shock does lead to an increase in the number of sellers, the presence of reference dependence and financial frictions means that some of these sellers are "locked in" by nominal losses, a low level of home equity, or both.

5.2 Untargeted aggregate moments

Having traced how the model transmits shocks, we now assess its ability to reproduce a salient aggregate pattern that was *not* targeted in the calibration: the cross-regional variation in the price–volume comovement documented in section 3.2. Different regions have been subject to different histories of shocks, which leave their imprint on the local "paper loss" share, so we examine how the model-implied response to a current shock depends on the regional paper-loss share. We map regions to sub-populations in the model, calculating the implied price and

Figure 6
Impulse responses to valuation, demand, and mobility shocks

We start from the model steady-state and introduce an unanticipated change in the path of u, ε and θ, with perturbations equal to -5%, -5%, and +5%, respectively. All shocks decay at rate 0.5 each period. We consider how the agents in the model respond to this perfect foresight shock, assuming they form rational expectations about how the economy evolves after the initial impact. Each plot compares the full model (solid black), the specification with financial frictions only (dashed red), and the frictionless benchmark (dash-dotted green); responses are expressed as deviations from the steady state, in units of the normalized average price.



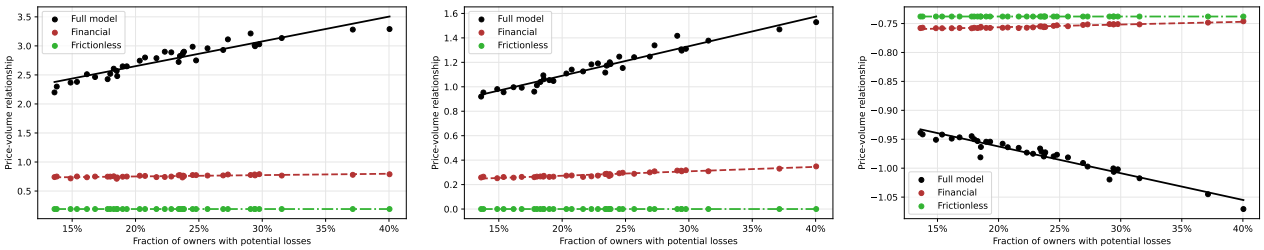
volume response to temporary shocks in each region by feeding the model with the regional joint distribution of reference prices and mortgage amounts from the data. Each region therefore differs along two dimensions, the local distribution of potential gains and losses relative to the reference point, and the local distribution of home equity, and both jointly determine the response of prices and volumes to a given shock.

Figure 7 reports that the model-implied price–volume comovement is stronger in regions with higher loss shares, capturing the idea that reference dependence and financial frictions amplify the price–volume comovement by flattening the supply curve and steepening the demand curve. The model-predicted range of price–volume comovement values runs from approximately 2.2 to 3.3 for the case of demand shocks; the level of the comovement is lower for valuation shocks, although its cross-sectional gradient is steeper in relative terms. In all three panels the frictionless benchmark is flat in the loss share, and for the valuation shock it is identically zero, since that shock leaves volumes unchanged at every loss share. The entire comovement in plot b), both its level and its slope, is therefore generated by the frictions. Plot c) shows that mobility shocks work in the opposite direction: they increase volumes while depressing prices, so

Figure 7
Untargeted price–volume comovement across regions

Each panel reports the price–volume comovement calculated as the ratio between the volume and price response to a temporary -5% demand shock (plot a), a -5% valuation shock (plot b), and a +5% mobility shock (plot c), each decaying at rate 0.5 per period. We calculate model-implied responses in each ITL2 region, using the observed joint distribution of reference prices and mortgage amounts in that region, i.e., regions differ both in the distribution of potential gains and losses and in the distribution of home equity.

a) Demand (ε) b) Valuation (u) c) Mobility (θ)



the implied comovement is negative, at around -0.94 in the lowest-loss-share sub-populations and -1.07 in the highest. The gradient again goes in the direction of amplification, and is once more entirely attributable to the frictions, since the frictionless benchmark is flat at about -0.74. The sign of the observed comovement is therefore informative about the combination of shocks hitting a region: only demand and valuation shocks generate the positive comovement observed in the data. While a precise quantitative match is not the goal of this exercise (the model is not calibrated to identify the time series of alternative types of shocks in the data), the model can account for a significant portion of the observed cross-regional variation.[19]

The model also reproduces the margin along which this comovement operates. In the data, the price–volume comovement is driven almost entirely by the intensive margin: the rate at which listed properties convert into sales comoves strongly with price changes, while listing volumes respond only weakly (plots c) and d) of Figure 4). Appendix Figure A.8 validates this pattern in the model: the model-implied comovement between prices and listing volumes is essentially zero for demand, valuation, and mobility shocks alike, and does not vary systematically with the regional loss share. Shocks thus move total transaction volumes

[19] A combination of shocks that varies across regions can also generate the steeper gradient observed in the data. Since demand and valuation shocks generate positive comovement and mobility shocks negative comovement, the comovement measured in a region depends on the history and magnitude of shocks that affect a particular location. If mobility shocks are more prevalent around London, where the loss share is low, and demand shocks in the Northern regions, where it is high, the implied cross-sectional slope is steeper than under any single shock.

through the probability that a given listing results in a sale, rather than through the entry of new listings, mirroring the empirical decomposition in section 3.2.

6 Applications

6.1 Monetary policy and the housing market

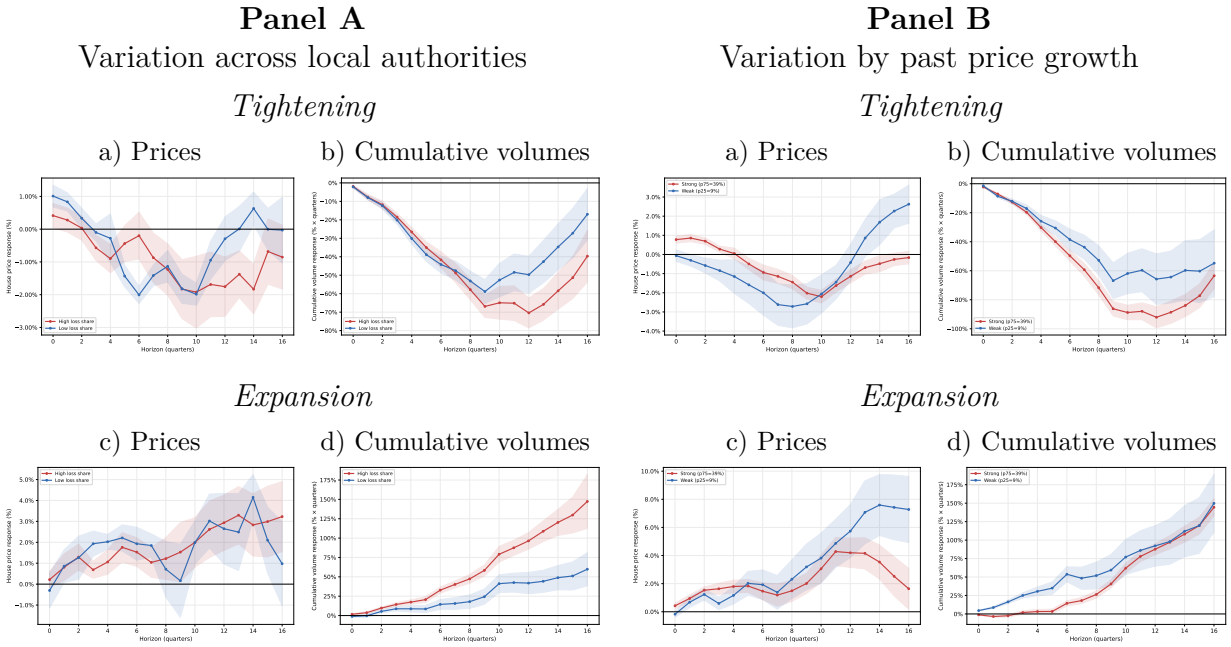
Shocks in our model also change the state variables that govern future outcomes, so their effects persist after the initial impulse has passed. Because housing transactions occur infrequently, reference points are updated only slowly: a temporary price movement shifts the distribution of nominal gains and losses across homeowners, and this shift reverses only as turnover gradually returns reference points to fundamentals. The market's response to a given shock is state-dependent, and is stronger when a larger share of homeowners are close to or below their reference points because of recent price movements. Identical policies can therefore have substantially different short-run effects depending on the phase of the housing cycle, which helps explain why temporary interventions can have persistent effects on housing activity.

We bring this prediction to the data using monetary policy surprises as a source of exogenous variation in aggregate conditions. We measure the monetary surprise using the interest rate surprise ("Target") factor of the U.K. Monetary Policy Event-Study Database (UKMPD; Braun et al., 2025), released by the Bank of England, which measures the surprise component of the policy rate from asset-price movements in a 30-minute window around scheduled announcements. Corresponding house prices and transaction volumes are constructed at the level of 280 English local authority districts (LADs) from the HM Land Registry Price Paid micro-data: for each district and quarter we form the transaction count and a mix-adjusted average price, using fixed national property-type weights. Our macroeconomic control variables are the quarterly real GDP and consumer price index (CPI) from the Office for National Statistics, and the Bank of England Official Bank Rate as the policy-stance indicator. For the loss-share split (Panel A of Figure 8), we rank districts within each year by the four-year moving average of their loss share, the district-level analogue of the ITL2 loss share behind Figure 4, computed from the listings data and lagged so that group membership is predetermined (years before 2011 use the earliest available ranking), and compare the top and bottom

decile of districts ("High loss share" and "Low loss share"). The estimation sample runs from 1997Q2 to 2024Q4.

Figure 8
Empirical asymmetry and hysteresis in the response to monetary-policy shocks

Panel A reports impulse responses of U.K. local-authority-district house prices (plots a and c) and cumulative transaction volumes (plots b and d) for the top ("High loss share") and bottom ("Low loss share") decile of the predetermined local loss share. Panel B splits districts instead by recent house-price growth over a four-year look-back. Within each panel, plots a) and b) report the response to a one-standard-deviation contractionary monetary policy surprise ("Tightening"), and plots c) and d) the response to an expansionary one ("Expansion"). Shaded bands indicate 90% confidence intervals based on standard errors clustered at the ITL2 level.



For each horizon $h = 0, \ldots, 16$ quarters, each outcome y (the log mix-adjusted house price, or the log quarterly transaction volume), and each district i, we estimate

$$y_{i,t+h} = \alpha_{i,h} + \beta_h \, z_t + \sum_{j=1}^{P} \left(\gamma_j \, z_{t-j} + \boldsymbol{\delta}_j' \, W_{i,t-j} \right) + \varepsilon_{i,t+h}, \qquad W_{i,t} = \left(R_t, \; y_{i,t}, \; gdp_t, \; cpi_t \right), \quad (12)$$

with $P = 4$ lags. The control vector W contains the Official Bank Rate R_t, the lagged dependent variable, and log real GDP and CPI; we also include four lags of the surprise component itself. The main specification includes district fixed effects $\alpha_{i,h}$ and controls for the district's mortgage share, the volume-weighted share of mortgage-financed transactions, interacted separately with the tightening and easing components of the surprise and centered

at its mean, so that differences between district groups are not driven by mortgage-share composition; the level of the (time-varying) loss share, and of the price-growth state in Panel B, is also controlled for. Standard errors are clustered by ITL2 region. Because tightenings and easings need not move the market symmetrically, in the asymmetry diagnostics of Figure 8 we split the surprise into its tightening ($z_t^+ \equiv \max\{z_t, 0\}$) and expansionary ($z_t^- \equiv \min\{z_t, 0\}$) components and replace $\beta_h z_t$ in equation (12) with $\beta_h^+ z_t^+ + \beta_h^- z_t^-$, reporting the separately estimated tightening response β_h^+ (top row) and expansionary response $-\beta_h^-$ (bottom row), each scaled to a one-standard-deviation surprise.

Figure 8 reports two diagnostics. Panel A splits districts into the top and bottom decile of the local loss share. In high-loss-share districts, the cumulative volume contraction following a tightening shock is larger and the house-price decline at medium horizons (one to two years after the shock) smaller than in low-loss-share districts, while the turnover rebound following an expansionary shock is markedly larger. The trough cumulative volume response to a tightening surprise is roughly -70 percent$\times$quarters in high-loss-share districts versus roughly -50 in low-loss-share districts, while after about a year and a half prices have fallen by about 2% in low-loss-share districts versus less than half a percent in high-loss-share ones. This is consistent with the asymmetric supply curve in the model: tightening shocks push more sellers into the loss domain and activate the lock-in margin, whereas expansionary shocks release sellers from it. Panel B documents state dependence: following a tightening shock, the house-price decline is markedly more persistent in periods of high recent local house-price growth (top quartile of four-year price growth); prices remain below their pre-shock level throughout the four-year horizon, whereas after weak recent growth (bottom quartile) they recover fully, and the cumulative volume contraction points the same way; expansionary shocks, in turn, produce their largest price response in periods of weak recent growth, where they reverse the prevailing depreciation. Strong recent growth leaves the economy in a gain region; a tightening shock reverses the appreciation and moves homeowners back toward and across their reference points, where loss aversion amplifies the response, while a shock that extends the prevailing trend leaves this margin largely inactive. To connect these patterns to the model's primitives, we simulate the structural model under a sequence of demand shocks and isolate the marginal effect of the final shock as a function of the preceding history, reported in Appendix Figure A.9.

6.2 Housing market taxation

We now consider the role of housing market fiscal policy, focusing on two distinct instruments. First, we analyze transaction taxes, which we assume (in keeping with the institutional reality in the U.K.) are paid by the buyer and therefore enter the model as a wedge τ^b on the buyer's acceptance condition: the threshold in equation (6) becomes $\varepsilon_{jt} \geq p_{it} + \tau^b + \beta\mathbb{E}_t[V^b_{t+1}] - \beta\mathbb{E}_t[V^h_{t+1}(p_{it}, m_{j,t+1})]$, where τ^b is the log of one plus the stamp duty rate. Second, we examine recurrent property taxes, which reduce the homeowner's flow utility in the Bellman equation (2): the flow u is replaced by $u - \tau^h$, where $\tau^h = \ln\left(\frac{1}{1-\text{property tax}}\right)$.

Policy impact The two tax instruments map directly onto the shocks studied earlier: the transaction tax τ^b enters the buyer's acceptance threshold and is therefore isomorphic to a shift in the buyer match quality ε, while the property tax τ^h shifts the homeowner's flow utility and corresponds to a shock to u. The impulse-response patterns documented above therefore carry over: relative to the frictionless benchmark, the full model with reference dependence and financial frictions delivers muted price responses and stronger volume responses to either tax instrument. The contrast is stronger for the property tax: in the frictionless benchmark its price elasticity is flat at -1.62 and its volume elasticity is identically zero, irrespective of the distribution of reference prices and mortgage balances, so the entire cross-sectional gradient reported in Figure 9 is generated by the frictions. Appendix Figure A.10 reports the same elasticities for the two restricted model versions: the financial margin on its own produces a gradient of the same sign but a considerably smaller cross-sectional spread, while the frictionless benchmark is flat in the loss share throughout.

Distortionary property taxes A key implication of this analysis is that short-run changes in the ongoing property tax are *distortionary* in our model. This stands in marked contrast to the neoclassical benchmark, in which ongoing property or land taxes are capitalized into prices but leave real activity unaffected (Mirrlees et al., 2012). In our environment, by contrast, a property tax compresses sellers' nominal gain/loss positions relative to their reference points, which moves listing decisions and depresses transaction volumes on impact. Over time, housing turnover causes both reference points and mortgage balances to converge to the post-tax price level: this restores homeowners' nominal gain/loss and home-equity positions, and returns

aggregate volumes to the initial level, recovering the no-volume-distortion benchmark in the long run.

Cross-regional variation in tax distortion Figure 9 plots model-implied responses of prices and volumes to a change in the levels of taxes, conditional on the share of "paper losses" in the stock of properties. To implement this analysis, we consider a temporary +5% shock to both the transaction tax rate and the property tax rate, which decays at rate 0.5 each period. We implement the analysis for different sub-populations of the stationary equilibrium, matching each sub-population with the empirical joint distribution of reference prices and mortgage amounts for each ITL2 region. The vertical axis shows the model-implied percentage change at impact; the horizontal axis shows the loss share of the particular sub-population for which the policy change is evaluated.

Panel A shows that prices respond less to either instrument in regions with higher loss shares, consistent with reference dependence flattening the supply curve and dampening the price response. The price response is uniformly stronger for the property tax than for the transaction tax. The transaction tax enters only the buyer's acceptance threshold, while the loss share captures seller-side motivations. The property tax, by contrast, affects buyers and sellers symmetrically, and its price impact depends on the full joint distribution of reference prices and mortgage balances in equilibrium. Quantitatively, the absolute price elasticity with respect to the transaction tax declines from roughly 0.22 in the sub-populations with the lowest loss shares to roughly 0.17 in those with the highest; the corresponding elasticity with respect to the property tax declines from roughly 0.72 to roughly 0.55.

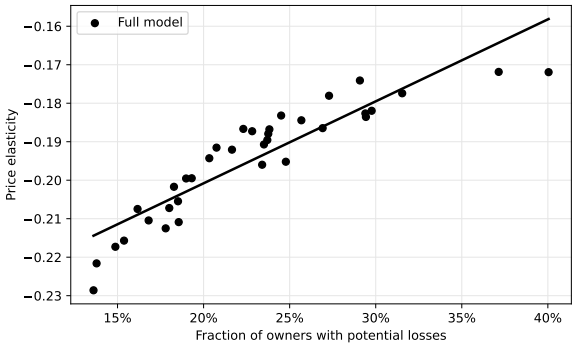
Panel B reports the corresponding pattern for volumes. The distortionary effect of taxes is smaller around London (low loss share) and larger in the Northern regions (high loss share); the absolute volume elasticity with respect to the transaction tax rises from about 0.51 in the lowest-loss-share sub-populations to about 0.57 in the highest, and the corresponding elasticity with respect to the property tax rises from about 0.68 to about 0.85. When sellers' nominal positions are concentrated in the loss domain, behavioral frictions bind more tightly: the volume response absorbs a larger share of the overall policy impact, while the price response is correspondingly compressed. Both panels reinforce the interpretation of the regional loss share as a parsimonious statistic for the cross-sectional incidence of housing market tax changes.

Figure 9
Model-implied tax elasticities: Cross-sectional variation

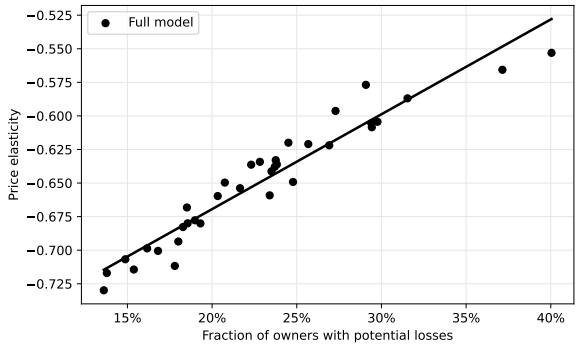
The figure reports the model-implied cross-sectional variation of the short-run response in prices and transaction volumes to changes in transaction taxes and ongoing property taxes. We consider a temporary +5% shock to both the transaction tax τ^b and the ongoing property tax τ^h, which decays at rate 0.5 each period, and calculate the first-period response elasticities of prices and total transaction volumes.

Panel A
Prices

a) Transaction tax

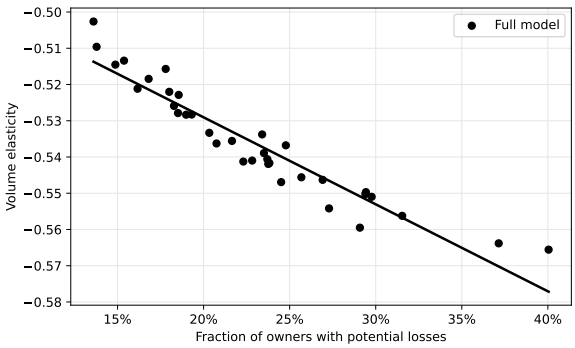


b) Property tax

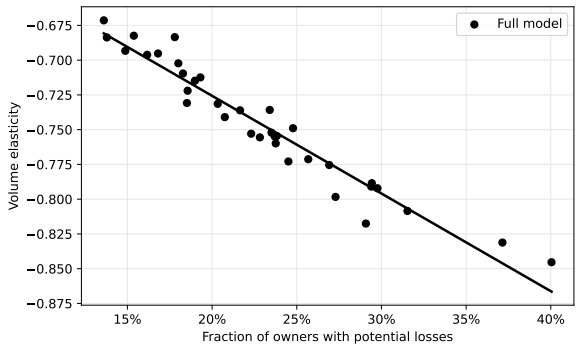


Panel B
Volumes

a) Transaction tax



b) Property tax



Evidence from U.K. transaction tax notches Before December 2014, the U.K. stamp duty schedule featured "notches": the tax rate jumped discretely at threshold prices such as £125,000, £250,000 and £500,000, and crossing a threshold raised the rate on the *entire* purchase value. Buyers therefore had a strong incentive to push transaction prices to just at or below a threshold, generating visible bunching in the realized price distribution. The December 2014 reform replaced each notch with a "kink", a continuous increase in the schedule, leaving the marginal fiscal cost of pricing just above a threshold only slightly higher and so substantially weakening that incentive. Panel A of Appendix Figure A.11 shows the two

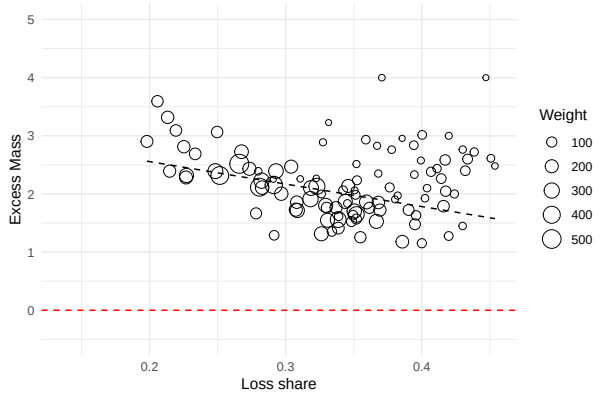
schedules, and Panel B the resulting fall in bunching at the £250,000 threshold.

Reference dependence on the seller side works against this bunching. A seller who would have to realize a nominal loss at the threshold price is less willing to accommodate the buyer's demand for a price reduction down to the notch, so bunching should be weaker precisely where many sellers sit in the loss domain; Appendix Section A.3 provides a graphical illustration of the underlying incentives. We test this by measuring excess bunching mass at the £250,000 threshold, which lies close to the median U.K. house price. Excess mass is the difference, within a £5,000 price band, between the number of realized transactions and the counterfactual number that would occur if every house traded at its hedonic price.

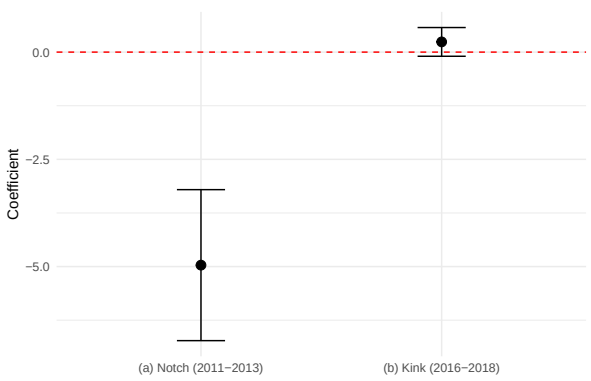
Figure 10
Cross-regional variation of bunching at tax thresholds

Plot a) reports the excess bunching mass for transaction prices in the interval (£245,000, £250,000] against the loss share at the level of ITL2 regions. The dark dotted line shows a linear fit between bunching mass and loss shares weighted by the number of expected transactions for the price band (the weights are shown by the relative size of the circles). Plot b) reports regression coefficients obtained by regressing excess bunching mass on loss shares at the ITL2 level during periods when the tax schedule had a notch (2011-2013) at the price £250,000 vs. a kink (2016-2018), using annual data.

a) Bunching of transaction prices



b) Pre- vs. post-reform



First, excess bunching mass at the £250,000 threshold is negatively related to regional loss shares. Plot a) of Figure 10 plots the annual excess mass at this threshold for each U.K. ITL2 region over 2011–2013 against the region's loss share. We lag the loss shares by six months to avoid mechanical correlation. The association is negative, both in the weighted fit line and in the regression coefficients reported in plot b): a higher proportion of potentially loss-making listings goes together with less excess bunching, consistent with sellers' reluctance to shave prices even marginally to accommodate the buyer's stamp duty saving. Second, this

association weakens once the notch is replaced by a kink. Plot b) of Figure 10 compares the regression of regional excess bunching mass on loss shares over the notch period (2011–2013) with the same regression over the kink period (2016–2018): the coefficient is significantly negative under the notch, and much smaller and statistically insignificant under the kink, as expected once the underlying incentive to bunch is largely removed.

7 Conclusions

We study how reference dependence and loss aversion aggregate in the residential housing market. Using granular data on property listings, transactions, buyer search, and outstanding mortgage balances, we calibrate a dynamic equilibrium search-and-matching model with rich heterogeneity and realistic behavioral and financial constraints. The model is disciplined by micro-level pricing and transaction moments and reproduces untargeted aggregate patterns in housing prices and volumes.

The central result is that the cross-sectional distribution of homeowners' nominal gain/loss positions, inherited from past transactions, is a payoff-relevant state variable. Its simple regional summary, the "paper loss share," predicts how shocks are divided between prices and transaction volumes: price–volume comovement is stronger where more homeowners face paper losses, and such regions exhibit substantially greater volume volatility and slightly lower price volatility. Because sellers near or below their reference points are reluctant to transact, nominal price movements affect real mobility and turnover. Because reference points reset only when properties trade, these effects are also persistent and history-dependent. Consistent with this mechanism, contractionary monetary-policy surprises generate larger volume declines after stronger recent appreciation, when the shock reverses the prevailing trend and moves homeowners back toward their reference points. Housing-tax incidence is similarly state-dependent: in high-loss-share regions, tax changes distort volumes more and prices less, and transaction-price bunching around U.K. stamp-duty notches is weaker. Taken together, these findings show that household-specific nominal reference points can aggregate into a persistent market-level rigidity. Housing-market adjustment and policy transmission consequently depend on the transaction histories embodied in the current distribution of gains and losses.

References

Ahn, SeHyoun, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf, "When inequality matters for macro and macro matters for inequality," *NBER macroeconomics annual*, 2018, *32* (1), 1–75. 5

Aiello, Darren, Jason D. Kotter, and Gregor Schubert, "When Money Moves In: The Consequences of Housing Wealth," Technical Report, Brigham Young University, Working Paper 2022. 20

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky, "Regressive sin taxes, with an application to the optimal soda tax," *The Quarterly Journal of Economics*, 2019, *134* (3), 1557–1626. 5

Allen, Eric J, Patricia M Dechow, Devin G Pope, and George Wu, "Reference-dependent preferences: Evidence from marathon runners," *Management Science*, 2017, *63* (6), 1657–1672. 5

Allen, Jason, Robert Clark, and Jean-François Houde, "Search frictions and market power in negotiated-price markets," *Journal of Political Economy*, 2019, *127* (4), 1550–1598. 5

Anagol, Santosh, Benjamin B Lockwood, Allan Davids, and Tarun Ramadorai, "Diffuse Bunching with Frictions: Theory and Estimation," Technical Report, Working Paper 2022. 6, 11

Andersen, Steffen, Cristian Badarinza, Lu Liu, Julie Marx, and Tarun Ramadorai, "Reference dependence in the housing market," *American Economic Review*, 2022, *112*, 3398–3440. 2, 5, 10, 12, 22, 27

Anenberg, Elliot and Daniel Ringo, "The propagation of demand shocks through housing markets," *American Economic Journal: Macroeconomics*, 2022, *14* (3), 481–507. 5

_ **and Patrick Bayer**, "Endogenous Sources Of Volatility In Housing Markets: The Joint Buyer–Seller Problem," *International Economic Review*, 2020, *61* (3), 1195–1228. 5

Badarinza, Cristian and Tarun Ramadorai, "International dimensions of housing markets," *Journal of Economic Perspectives*, 2025, *39* (3), 87–106. 30

_ **, Vimal Balasubramaniam, and Tarun Ramadorai**, "In Search of the Matching Function in the Housing Market," Technical Report, Working Paper 2024. 27

Best, Michael Carlos and Henrik Jacobsen Kleven, "Housing market responses to transaction taxes: Evidence from notches and stimulus in the UK," *The Review of Economic Studies*, 2018, *85* (1), 157–193. 5

Bracke, Philippe and Silvana Tenreyro, "History dependence in the housing market," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 420–443. 5

Braun, Robin, Silvia Miranda-Agrippino, and Shreyosi Saha, "Measuring monetary policy in the UK: The UK monetary policy event-study database," *Journal of Monetary Economics*, 2025, *149*. 35

Chetty, Raj, "Behavioral economics and public policy: A pragmatic perspective," *American Economic Review*, 2015, *105* (5), 1–33. 5

Cho, Yunho, Shuyun May Li, and Lawrence Uren, "Stamping out stamp duty: Housing mismatch and welfare," *Quantitative Economics*, 2024, *15* (2), 381–426. 5

Coven, Joshua, Sebastian Golder, Arpit Gupta, and Abdoulaye Ndiaye, "Property Taxes and Housing Allocation Under Financial Constraints," *American Economic Review*, 2025, p. forthcoming. 5

DellaVigna, Stefano, Attila Lindner, Balázs Reizer, and Johannes F Schmieder, "Reference-dependent job search: Evidence from Hungary," *The Quarterly Journal of Economics*, 2017, *132* (4), 1969–2018. 2, 5

Diaz, Antonia and Belén Jerez, "House prices, sales, and time on the market: A search-theoretic framework," *International Economic Review*, 2013, *54* (3), 837–872. 5

Eerola, Essi, Oskari Harjunen, Teemu Lyytikäinen, and Tuukka Saarimaa, "Revisiting the effects of housing transfer taxes," *Journal of Urban Economics*, 2021, *124*, 103367. 5

Farhi, Emmanuel and Xavier Gabaix, "Optimal taxation with behavioral agents," *American Economic Review*, 2020, *110* (1), 298–336. 5

Fisher, Jack, Alessandro Gavazza, Lu Liu, Tarun Ramadorai, and Jagdish Tripathy, "Refinancing cross-subsidies in the mortgage market," Technical Report, Working Paper 2023. 4

Fonseca, Julia and Lu Liu, "Mortgage Lock-In, Mobility, and Labor Reallocation," Technical Report, Working Paper 2023. 4

Gabaix, Xavier, "A behavioral New Keynesian model," *American Economic Review*, 2020, *110* (8), 2271–2327. 5

Genesove, David and Christopher Mayer, "Loss aversion and seller behavior: Evidence from the housing market," *The quarterly journal of economics*, 2001, *116* (4), 1233–1260. 2, 5, 27

_ **and Lu Han**, "Search and matching in the housing market," *Journal of Urban Economics*, 2012, *72* (1), 31–45. 3

Grindaker, Morten, Artashes Karapetyan, Espen Moen, and Plamen Nenov, "Transaction Sequence Decisions and Housing Market Volatility: Evidence from Norway," Technical Report 2021. 5

Guren, Adam M, "House price momentum and strategic complementarity," *Journal of Political Economy*, 2018, *126* (3), 1172–1218. 3, 4, 5, 27

Han, Lu and William C Strange, "The microstructure of housing markets: Search, bargaining, and brokerage," *Handbook of regional and urban economics*, 2015, *5*, 813–886. 5

_ **, Rachel Ngai, and Kevin Sheedy**, "To Own or To Rent? The Effects of Transaction Taxes on Housing Markets," 2023. 5

Kaplan, GreG and Justin Weidner, "The Wealthy Hand-to-Mouth," *Brookings Papers on Economic Activity*, 2014, pp. 77–138. 5

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante, "Monetary policy according to HANK," *American Economic Review*, 2018, *108* (3), 697–743. 5

Kleven, Henrik Jacobsen, "Bunching," *Annual Review of Economics*, 2016, *8*, 435–464. 5

Krainer, John, "A theory of liquidity in residential real estate markets," *Journal of Urban Economics*, 2001, *49* (1), 32–53. 5

Leamer, Edward E., "Housing IS the Business Cycle," NBER Working Paper 13428, National Bureau of Economic Research 2007. 3

Liu, Lu, "The demand for long-term mortgage contracts and the role of collateral," Technical Report, Working Paper 2022. 10, 11

Maxted, Peter, "A macro-finance model with sentiment," *Review of Economic Studies*, 2023, p. rdad023. 5

Mirrlees, James, Stuart Adam, Tim Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, and James Poterba, "The Mirrlees review: A proposal for systematic tax reform," *National Tax Journal*, 2012, *65* (3), 655–683. 5, 38

Moen, Espen R, Plamen T Nenov, and Florian Sniekers, "Buying first or selling first in housing markets," *Journal of the European Economic Association*, 2021, *19* (1), 38–81. 5

Mullainathan, Sendhil, Joshua Schwartzstein, and William J Congdon, "A reduced-form approach to behavioral public finance," *Annu. Rev. Econ.*, 2012, *4* (1), 511–540. 5

Ngai, L Rachel and Silvana Tenreyro, "Hot and cold seasons in the housing market," *American Economic Review*, 2014, *104* (12), 3991–4026. 4, 5

O'Donoghue, Ted and Matthew Rabin, "Optimal sin taxes," *Journal of Public Economics*, 2006, *90* (10-11), 1825–1849. 5

Office for Budget Responsibility, "Council Tax," 2024. 5

Piazzesi, Monika, Martin Schneider, and Johannes Stroebel, "Segmented housing search," *American Economic Review*, 2020, *110* (3), 720–59. 5

Ramadorai, Tarun, "Household Finance in Action," AEA-AFA Joint Luncheon Address, ASSA Annual Meeting, Philadelphia 2026. Slideshow, 3 January 2026. 10

Rees-Jones, Alex, "Quantifying loss-averse tax manipulation," *The Review of Economic Studies*, 2018, *85* (2), 1251–1278. 2, 5

Saez, Emmanuel and Gabriel Zucman, "Progressive wealth taxation," *Brookings Papers on Economic Activity*, 2019, *2019* (2), 437–533. 5

Stein, Jeremy C, "Prices and trading volume in the housing market: A model with down-payment effects," *The Quarterly Journal of Economics*, 1995, *110* (2), 379–406. 3, 4, 11, 22

Taubinsky, Dmitry and Alex Rees-Jones, "Attention variation and welfare: Theory and evidence from a tax salience experiment," *The Review of Economic Studies*, 2018, *85* (4), 2462–2496. 5

Urban-Brookings Tax Policy Center, "The Tax Policy Briefing Book," 2024. 5

Wheaton, William C, "Vacancy, search, and prices in a housing market matching model," *Journal of Political Economy*, 1990, *98* (6), 1270–1292. 4, 5

Behavioral Lock-In: Housing Markets with Reference Dependent Agents

Online Appendix

Cristian Badarinza, Tarun Ramadorai,

Juhana Siljander, Jagdish Tripathy

August 16, 2026

A.1 Construction of the mortgage data

To merge granular listings with the mortgage data, we use mortgage observations that report property values together with a unique combination of granular postcodes, property value, and transaction dates. This combination is able to uniquely identify observations across the datasets, though we also allow tolerances in the reported property value and transaction dates to account for potential misreporting.

To track home-equity positions for individual mortgages, we further restrict the sample to observations that report a unique combination of granular postcodes and borrower date of birth, dropping duplicate observations and cases where a borrower has multiple liens on the same property. We follow these mortgages across snapshots based on the unique postcode–date of birth combination.

We use the mortgage balances tracked over multiple snapshots to calculate an amortization schedule. Since mortgage borrowers in the U.K. may close their accounts when moving houses or refinancing, we restrict this sample to contracts with a non-zero balance in each half-yearly snapshot between 2015H2 and 2022H2.[1] Thus, we condition on the particular postcode–date

[1]Mortgage closures and residential moves limit our ability to track mortgages over long horizons. As a result, we are able to follow only around 70,000 of the approximately 400,000 mortgages originated in 2015H1 through to 2022H2.

of birth combination being present in all subsequent snapshots, with a non-zero balance on the associated mortgage. In results available upon request, we find that the amortization schedule is similar across origination dates and borrower types, and does not differ greatly for mortgages with different loan-to-value ratios at origination.

A.2 Interpretation of μ

The down-payment penalty μ in the seller's utility (3) admits a quantitative interpretation in terms of the higher borrowing costs faced by a constrained seller who, when financing a subsequent purchase, is pushed into a high loan-to-value (LTV) mortgage bracket. Consider a mortgage of size M and maturity T carrying an effective per-period interest rate i. The constant amortizing payment is $A(M,i) = M\,i/[1-(1+i)^{-T}]$. Discounting this stream at the borrower's own discount rate r, its present value is

$$NPV(M,i) = A(M,i)\,\frac{1-(1+r)^{-T}}{r} = M\,\frac{1-(1+r)^{-T}}{1-(1+i)^{-T}}\,\frac{i}{r}.$$

Normalizing the house value to $P = 1$ and setting $M = LTV$, the additional present-value cost of financing at a penalized rate i_1 rather than the unpenalized rate i_0 is

$$\kappa^{data} = LTV\left(\frac{1-(1+r)^{-T}}{1-(1+i_1)^{-T}}\,\frac{i_1}{r} - \frac{1-(1+r)^{-T}}{1-(1+i_0)^{-T}}\,\frac{i_0}{r}\right).$$

Using the empirical U.K. mortgage-rate schedule (Figure A.7), moving from a 70% to a 90% LTV raises the mortgage rate from $i_0 = 7\%$ to $i_1 = 11\%$. For a discount rate $r = 3.8\%$ (the average return on long-term mortgage bonds) and a maturity of $T = 25$ years, this implies a penalty κ^{data} of approximately 47% of the house value at the 90% LTV bracket. By comparison, our calibrated value $\mu = 14.42$ implies, through the quadratic penalty $\mu(\gamma - \log(P/M))_+^2$ with $\gamma = 0.30$ and the approximation $\log(P/M) \approx \widehat{H}$ for moderate equity levels, a penalty of $14.42 \times (0.30 - 0.10)^2 \approx 0.58$, i.e., about 58% of the house value for a seller with 10% home equity. The calibrated $\mu = 14.42$ is thus somewhat larger than the penalty implied by interest-rate differentials at a discount rate of $r = 3.8\%$, with the residual plausibly reflecting non-pecuniary and non-interest-rate costs of downsizing. Since a wide range of discount rates is supported by the literature, Figure A.12 plots how κ^{data} varies with r: the interest-rate-

implied penalty ranges from about 65% (at $r = 1\%$) to 23% (at $r = 12\%$) of the house value, so the calibrated $\mu = 14.42$ lies inside the empirically relevant range, close to the value implied by discount rates at the lower end of that range.

A.3 Theoretical intuition: Bunching at tax threshold levels

To quantify bunching at tax kinks and notches, let $m = p(1 + \tau)$ denote the total cash outflow of the buyer for a property of price p subject to a tax rate τ. Figure A.13 illustrates the resulting incentive: buyers who would otherwise transact at prices in the interval $[p^*, p^* + \Delta p]$ instead locate at the threshold p^*, producing excess mass at the kink.

The size of this bunching mass is governed by the curvature of the buyer's indifference curves, with a steeper indifference curve implying a smaller bunching region. The mechanism parallels the canonical labor-supply bunching problem, in which a worker adjusts hours of work, and hence taxable income, along a consumption-leisure trade-off. In our setting, the buyer's only margin of adjustment with respect to a given property and seller is the timing of the purchase: the buyer may transact immediately at a higher price, or delay in anticipation of transacting later at a lower price, should the seller fail to sell in the interim. The counterpart of a worker who reduces hours is in our case a buyer who credibly threatens to delay or renege on the transaction. The strength of this margin is determined by the seller's urgency, that is, by how readily the seller lowers the price as time passes. Reference dependence shapes this willingness to negotiate: a seller in the loss domain is more reluctant to reduce the price, which narrows the buyer's scope for delay and reduces the magnitude of bunching.

A.4 Additional tables

Table A.1
Transactions and listings data: sample description and cleaning

The table reports the data cleaning process for two different samples used in the paper. The left column begins with the Rightmove listings data, and for each listing, looks for the last transaction on the property in the Land Registry transactions data. The merged data is used to estimate the potential gain associated with each property. The right column begins with the Land Registry transaction data and looks for the listing associated with the transaction in the Rightmove data. The merged data is used to run the hedonic model, using the hedonic characteristics from the Rightmove data and the transaction prices from the transactions data. Exclusion categories overlap, so the sum of individual exclusions exceeds the total number of dropped observations.

Rightmove listings data		HM Land Registry Price Paid data	
All unique listings	**23,323,691**	**All transactions**	**29,953,770**
Listings before 2010	-955,528	Transactions before 2010	-16,199,932
Ambiguous addresses	-3,924,639	Ambiguous addresses	-2,334,876
Not in England or Wales	-1,078,869		
Final sample after filters	18,652,386		12,550,718
Development prop.	378,888		
Non-development prop.	17,469,450		
Listings/transactions match			
Unmatched to a preceding transaction	-6,720,224	Unmatched to a preceding listing	-3,841,805
Development prop.	-225,214		
Non-development prop.	-6,495,010		
	11,128,114		8,708,913
Development prop.	153,674		
Non-development prop.	10,974,440		
Hedonic sample			
Auction properties	-133,518		-147,736
Transaction within 30 days of first listing	-75,163		
Incomplete hedonic characteristics	-754,532		-567,798
Final listings data	**10,164,901**	**Hedonic model data**	**7,993,379**

Table A.2
Matching listings (2015-2024) with mortgages

The table reports the outcome of each step of the algorithm matching transacted and untransacted listings to mortgage records in the PSD007 snapshots. Listings are matched on the postcode (PCD), the previous transaction price, and the mortgage account opening date; where no exact match exists, the closest combination within 30 days of the reported dates or 5% of the transacted prices is considered a match.

	Transacted	**Untransacted**	**Total**
Total	5,519,145	1,505,024	7,024,169
Prev. Trans Year < 2005	1,299,685	306,960	1,606,645
Can Match	**4,219,460**	**1,198,064**	**5,417,524**
Unmatched	1,538,742	518,889	2,057,631
Matched but zero balance	574,421	52,994	627,415
Matched with non-zero balance			
Exact match	477,847	135,245	613,092
Exact [PCD + Price] + Closest Date (30 days)	1,344,711	390,748	1,735,459
Exact [PCD + Date] + Closest Price (5%)	58,134	17,960	76,094
Exact [PCD] + Closest Date (30 days) + Closest Price (5%)	225,605	82,228	307,833
Total	**2,106,297**	**626,181**	**2,732,478**
With mortgage (%)	50%	52%	50%
Matched	64%	57%	62%

Table A.3
Summary statistics on mortgages originated in 2015

The table shows summary statistics for mortgages originated in the first half of 2015 (2015H1) that have a non-missing balance in each snapshot until 2022H2. We observe the original loan and subsequent balance in half-yearly snapshots over 7 years for this sample, and use it for the amortization schedule shown in Figure A.6 used to fit the model. Incentivised ($\mathbb{D}$) is an indicator equal to one if the mortgage carries an incentivised introductory rate (e.g., an initial fixed-rate or discounted-rate period).

	mean	sd	p10	p25	p50	p75	p90
Original Loan (£)	159,692	129,195	55,635	84,382	129,960	197,542	291,925
Balance (£)	156,387	127,008	54,185	82,144	126,803	193,359	285,515
Interest Rate (%)	2.88	0.96	1.84	2.25	2.78	3.39	4.29
Original Term (months)	275	89	155	204	288	348	408
Monthly Payment (£)	782	575	305	443	650	964	1,392
Incentivised ($\mathbb{D}$)	0.96	0.20	1	1	1	1	1

Table A.4
Summary statistics for variables aggregated at an MSOA×Year level

The table reports summary statistics for the variables used in the analysis of aggregate transaction outcomes, averaged at the MSOA×year level: the time on market (in months), the sale indicator, the average listing premium, average potential losses and gains, and the average home-equity measures.

Variable	N	Mean	SD	P10	P50	P90
Time on Market (Months)	57874	10.194	3.848	6.037	9.237	15.664
Sale Indicator	57874	0.776	0.151	0.545	0.823	0.921
Av. Listing Premium	57874	0.122	0.143	-0.035	0.105	0.300
Avg. Pot. Loss	57874	0.128	0.069	0.043	0.123	0.219
Avg. Pot. Gain ($\hat{g} \geq 0$)	57874	0.534	0.367	0.257	0.456	0.787
Avg Equity ($\hat{h} < 0.3$)	57874	0.144	0.094	0.000	0.165	0.240
Avg. Equity (Constrained)	57874	0.156	0.094	0.060	0.135	0.300

A.5 Additional figures

Figure A.1
Actual price vs. hedonic price

The figure illustrates the in-sample fit of our hedonic model. The horizontal axis shows the realized log transaction price, and the vertical axis the predicted log price from the hedonic model. We report a bin-scatter plot based on 20 frequency bins.

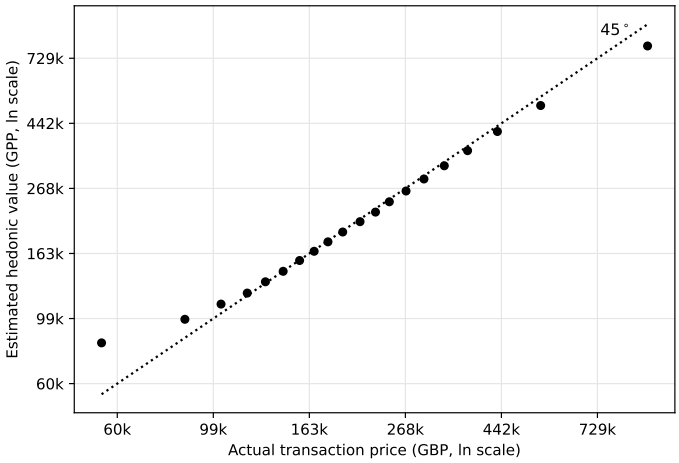


Figure A.2
Variation of "paper loss" share across locations and time

The figure reports, for each local authority district, the average share of nominal losses across listings outstanding during the month indicated below each map. Darker (blue) shades indicate lower loss shares and lighter (yellow) shades higher loss shares, on a color scale ranging from 0% to 50%.

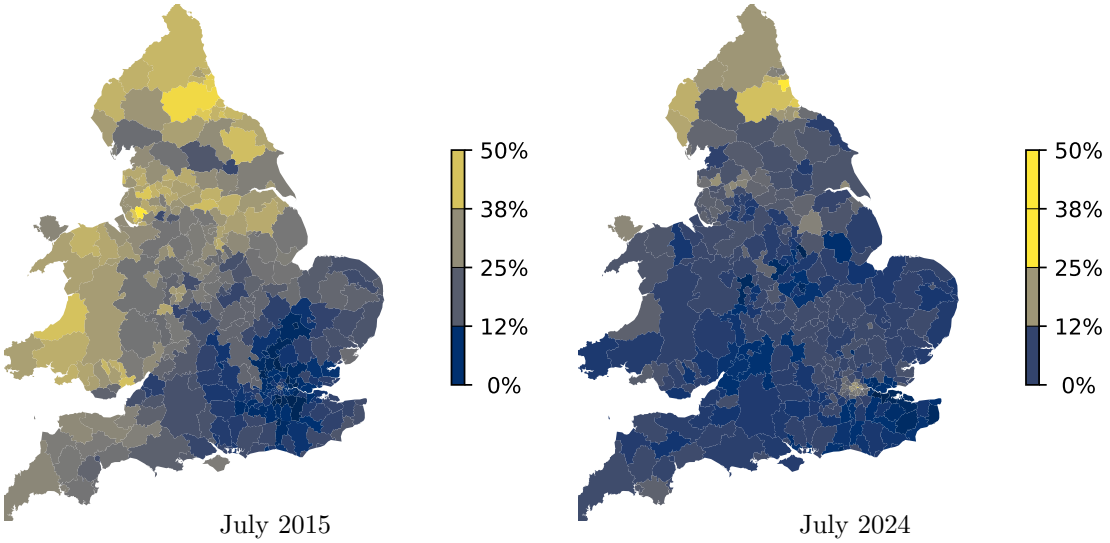


Figure A.3
Regional variation in bunching of transaction prices

The figure reports the distribution of the realized nominal gain, defined as the percent difference between the realized sale transaction price and the initial purchase price, computed for the nine ITL1 regions of England.

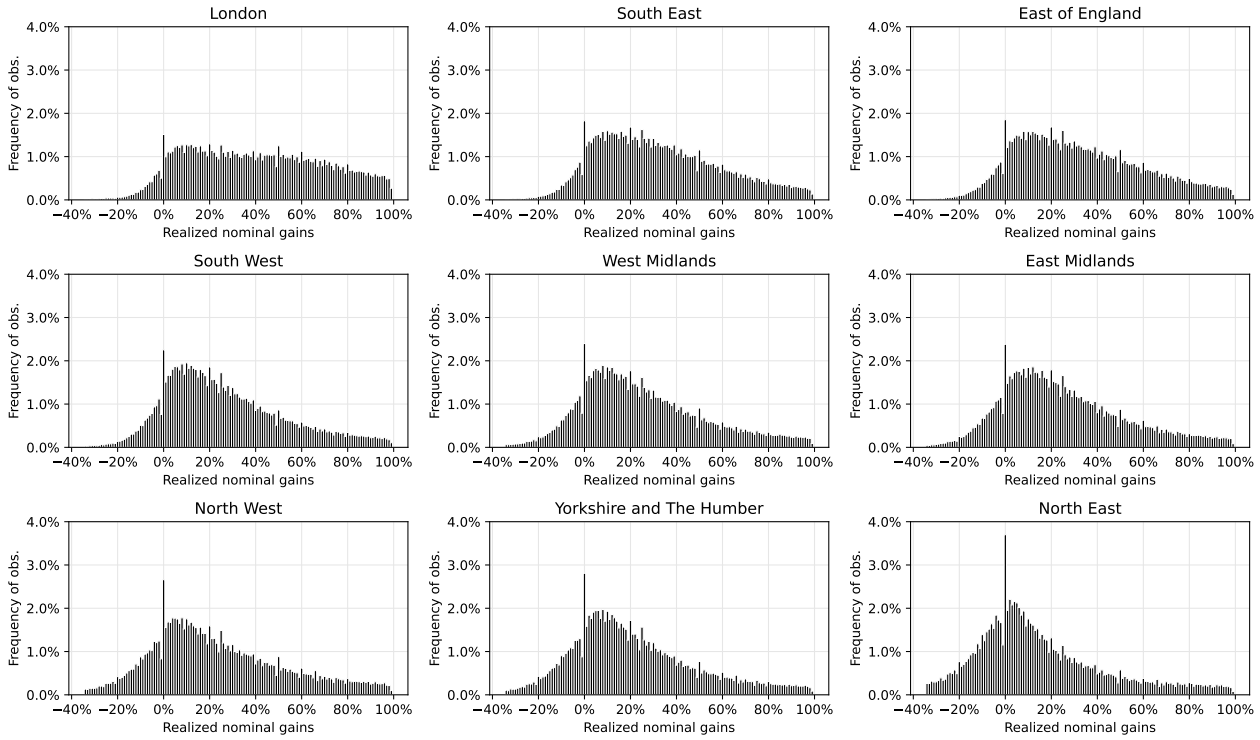


Figure A.4
"Hockey stick" pattern across regions

The figure reports average listing premia, i.e., the percent differences between the listing price and the hedonic property value at the time of listing, for different values of potential gains. We restrict the horizontal axis, for a more convenient graphical representation of effects around zero potential gains. The frequency distribution of potential gains is shown in gray shading. The numbers reported in the left upper corner of each plot indicate the share of homeowners in each ITL1 region for whom the nominal potential gain has a value below zero.

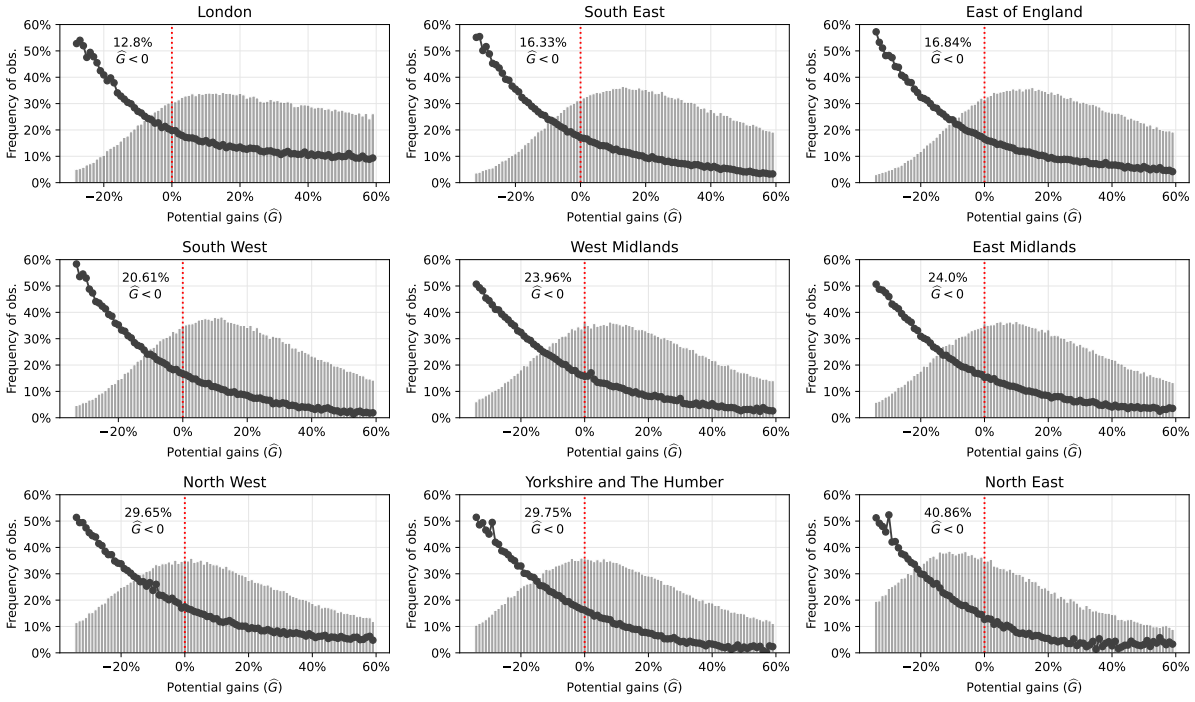
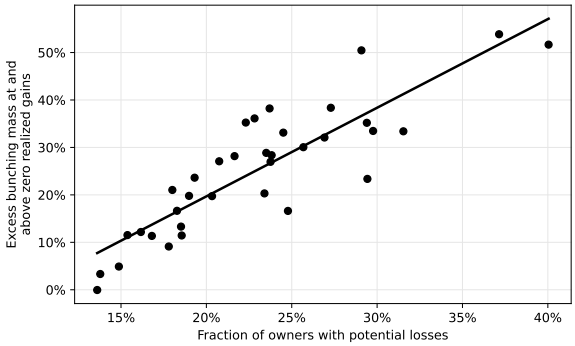


Figure A.5
Excess mass and home equity across regions

In plot a) the horizontal axis is the average "loss share" for each U.K. ITL2 region between January 2010 and December 2024 (the fraction of owners that list a property for sale for which the estimated hedonic value is below the initial purchase price), and in plot b) it is the fraction of listings with low home equity at the time of listing. Plot a) reports the region-specific excess mass at the reference point, measured as the bunching at zero nominal gains in the distribution of realized transaction prices relative to the counterfactual hedonic-value distribution. Plot b) reports the region-specific price–volume comovement coefficients of plot a) of Figure 4 against the share of listings with low home equity.

a) Excess mass at the reference point



b) Home equity

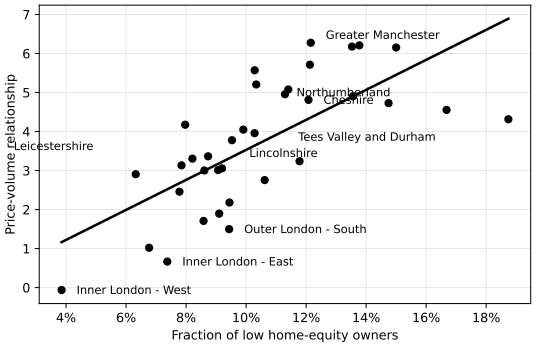


Figure A.6
Amortization and repayment of mortgage loans

The figure illustrates the evolution of outstanding mortgage balances in our data in half-yearly snapshots for a sub-sample of mortgages originated in the first half of 2015. The estimated coefficients reflect the average amortization rate at each half-yearly lag from origination, based on the following specification:

$$b_{it} = \sum_{\tau} \beta_{\tau} \cdot \mathbb{D}[\tau = t] + \varepsilon_{it},$$

where $b_{it} = l_{it}/l_{i,2015H1}$ is based on the outstanding balance for mortgage i in half-yearly snapshot t (l_{it}) and original loan size ($l_{i,2015H1}$), and $\mathbb{D}[\tau = t]$ are half-yearly dummies.

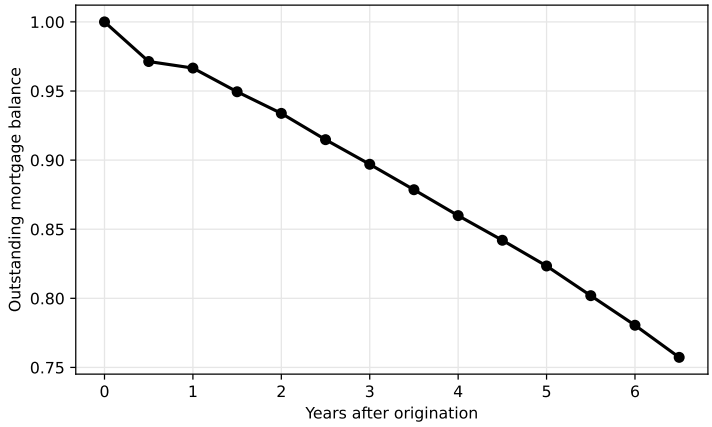


Figure A.7
Initial loan-to-value ratios and costs

The figure reports the distribution of loan-to-value ratios at origination and the mortgage credit interest rate spread paid on 5-year fixed-rate mortgages across loan-to-value bands estimated by Liu (2022) for mortgage contracts issued between 2013 and 2017.

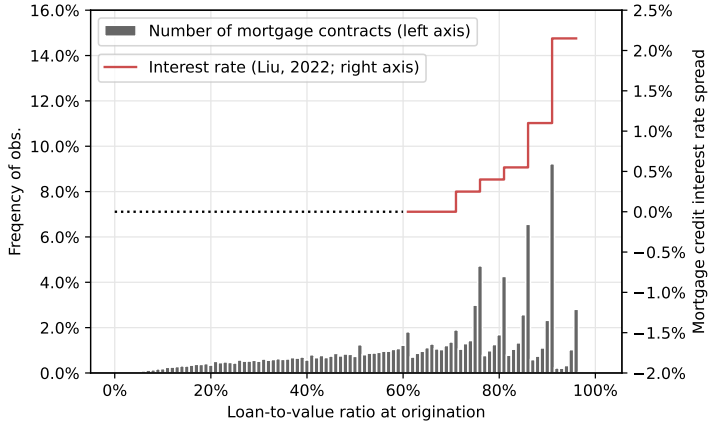


Figure A.8
Model-implied comovement of prices and listing volumes

Each panel reports the model-implied comovement between prices and listing volumes, calculated as the ratio between the listing-volume response and the price response to a temporary shock that decays at rate 0.5 each period (a -5% demand shock, a -5% valuation shock, and a +5% mobility shock, respectively). We calculate model-implied responses in each ITL2 region, using the observed joint distribution of reference prices and mortgage amounts in that region; the horizontal axis shows the loss share of each region. Each plot compares the full model, the specification with financial frictions only, and the frictionless benchmark.

a) Demand (ε) b) Valuation (u) c) Mobility (θ)

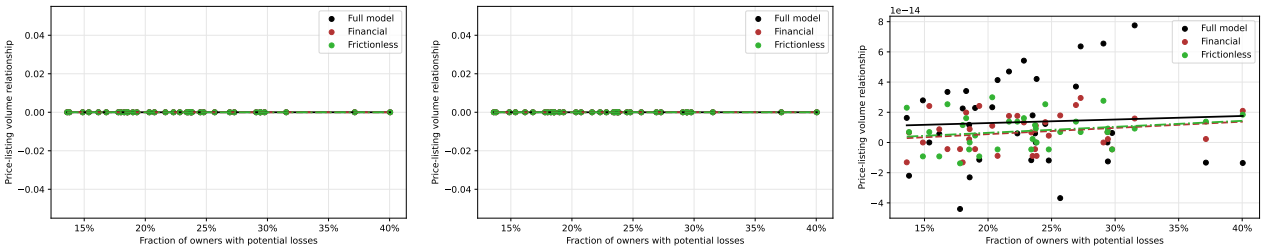


Figure A.9
Marginal effect of a terminal demand shock, conditional on recent price history

Each bar reports the change in transaction volume (plot a) and in the log price level (plot b) attributable to a demand shock that occurs in the ninth period in the calibrated model, after eight sequential shocks of ±1% that build an appreciating (+) or depreciating (−) recent path. Bars are green for a positive terminal shock and red for a negative one; adjacent pairs hold the terminal shock fixed and vary only the recent history.

a) Transaction volumes b) Prices

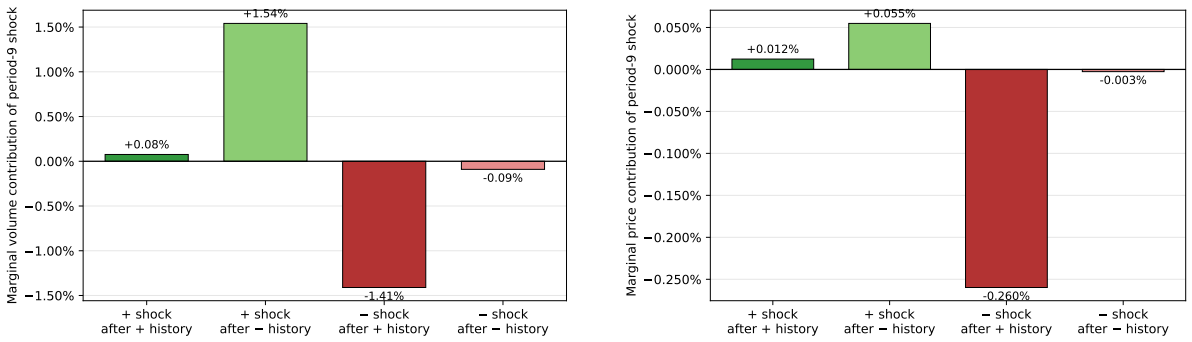
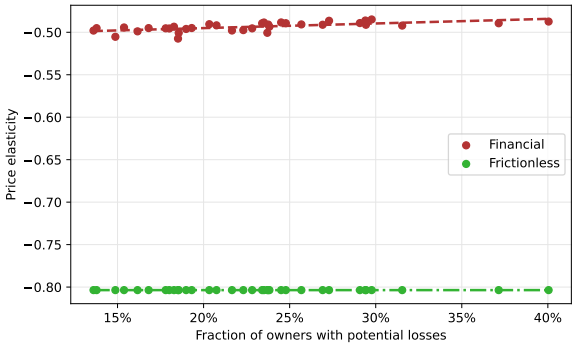


Figure A.10
Model-implied tax elasticities: Alternative model versions

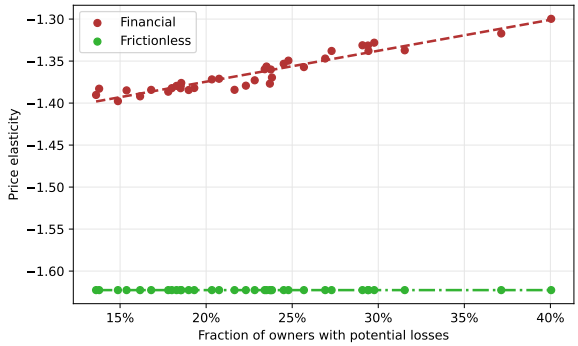
The figure replicates Figure 9 for the two restricted versions of the model: the specification with financial frictions only ($\eta_+ = \eta_- = 0$, $\mu > 0$) and the frictionless benchmark ($\eta_+ = \eta_- = 0$, $\mu = 0$). As in Figure 9, we consider a temporary +5% shock to both the transaction tax τ^b and the ongoing property tax τ^h, which decays at rate 0.5 each period, and calculate the first-period response elasticities of prices and total transaction volumes for sub-populations of the stationary equilibrium matched to the empirical joint distribution of reference prices and mortgage amounts in each ITL2 region.

Panel A
Prices

a) Transaction tax

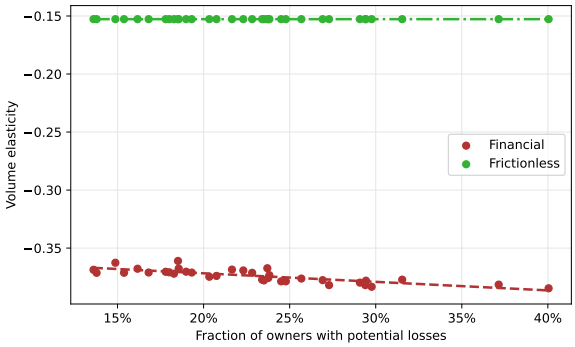


b) Property tax



Panel B
Volumes

a) Transaction tax



b) Property tax

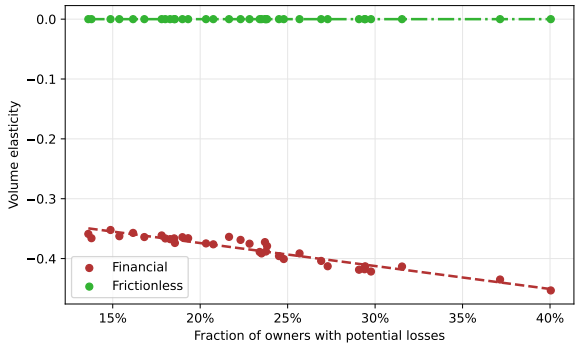
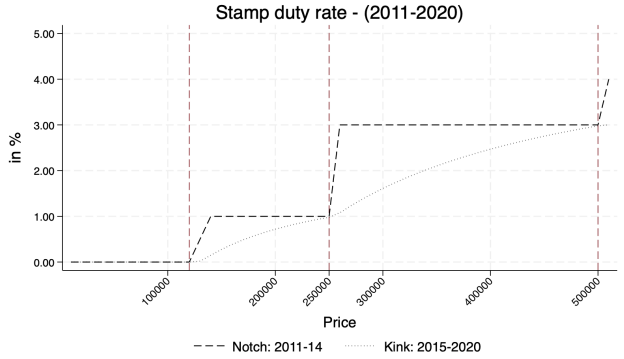


Figure A.11
Transaction tax schedule and bunching

Panel A plots the average stamp duty rate paid for different levels of housing transaction prices in the U.K. during 2011-2014 (notch schedule) and 2015-2020 (kink schedule). The average rate paid changes at threshold prices indicated by the dashed vertical red line. Panel B shows histograms of all transactions in 2013 (under notch tax schedule) and 2016 (under kink tax schedule) by the hedonic price and the transaction mass. The difference in the mass by transaction and hedonic price for any given price band gives the bunching mass for the price level.

Panel A

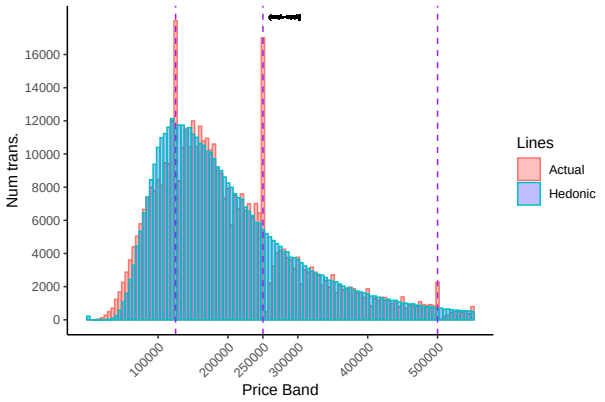
Transaction tax rates for different prices (2011-2020)



Panel B

Histograms of transaction and hedonic prices in 2013 (notch schedule) and 2016 (kink schedule)

(a) 2013 (Notch)



(b) 2016 (Kink)

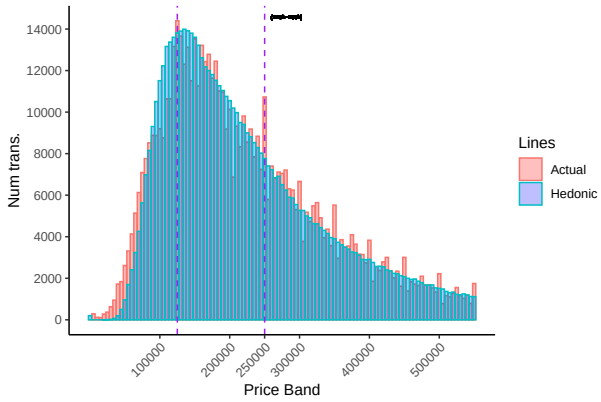


Figure A.12
Borrowing-cost penalty implied by the mortgage-rate schedule

The figure plots the present-value borrowing-cost penalty κ^{data}, expressed as a share of the house value, incurred by a seller who finances a subsequent purchase at a 90% LTV rather than a 70% LTV, as a function of the borrower's discount rate r. The penalty is computed from the empirical U.K. mortgage-rate schedule (Figure A.7), setting the unpenalized rate to $i_0 = 7\%$ at 70% LTV, the penalized rate to $i_1 = 11\%$ at 90% LTV, and the maturity to $T = 25$ years.

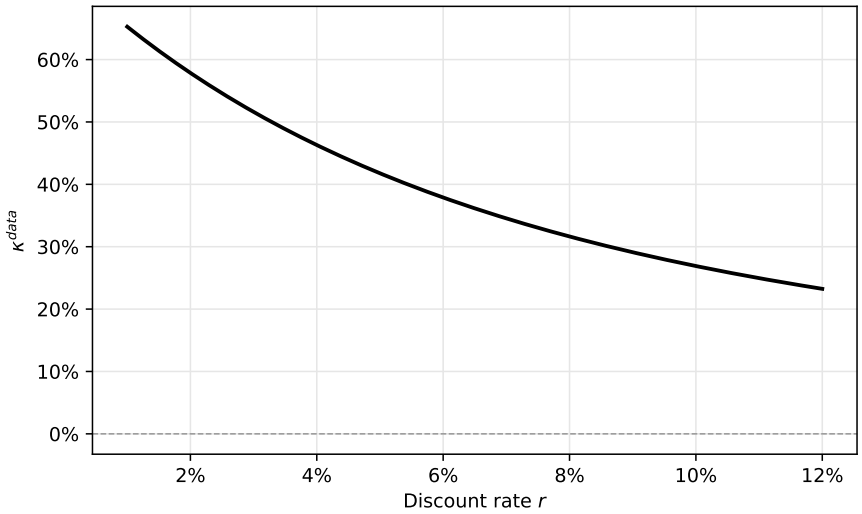


Figure A.13
Illustration of bunching at tax kink points

The figure illustrates the buyer's incentive to transact at the tax threshold price p^*. The solid blue line is the buyer's total cash outflow $-m = -p(1+\tau)$, which changes slope discretely at the threshold. Buyers whose indifference curves cross the outlay schedule in the interval $[p^*, p^* + \Delta p]$ locate at the threshold, producing excess mass. When indifference curves are steeper (right plot), reflecting a more limited scope for the buyer to delay the transaction, the bunching interval $\Delta p'$ is smaller.

