Social Sorting and the Refinancing Channel of Monetary Policy

Marta Cota*
Nova School of Business and Economics

Mauro Mastrogiacomo
De Nederlandsche Bank and Vrije Universiteit Amsterdam

Ante Šterc†
Banco de Portugal and Nova School of Business and Economics

This draft: May 1st, 2026
First draft: January 5th, 2026

Abstract

Note: This paper is pending data disclosure approval from the Dutch National Bank and Statistics Netherlands.

This paper measures and decomposes the contribution of social networks to the income gradient in mortgage refinancing pass-through. The gradient has two sources: direct effects of income on refinancing decisions, and indirect effects through social networks, which sort by income and therefore expose higher-income households to peers who refinance more. Using Dutch administrative registers that link mortgage contracts to whole-population family, workplace, and neighborhood networks, we provide the first nationally representative causal estimate of peer effects in mortgage refinancing, identified through the 2016 revision to the loan-to-income norm. A one-standard-deviation increase in peer housing-market activity raises active refinancing by 8 percentage points in family networks, 0.74 percentage points in neighbor networks, and 5.7 percentage points in neighbor networks, sizable estimate against the basline rates of refinancing at 8-10 percent. The response rises sharply with income, with a top-quintile coefficient roughly three times the bottom-quintile coefficient in family and workplace networks. We then build a heterogeneous-agent model of refinancing with a reduced-form network friction and an endogenous term premium, calibrated to the IV estimate. The model quantifies the aggregate amplification of the rate-cut response, the share of the income gradient in pass-through attributable to network sorting rather than to direct income effects, and the consumption-equivalent welfare effects by income quintile under counterfactuals with and without sorting. The framework also informs macroprudential design: loan-to-income limits alter the composition of peer networks and therefore the strength of the network channel through which monetary policy operates.

*Email: marta.cota@novasbe.pt, Universidade NOVA de Lisboa, Campus de Carcavelos, 2775-405 Carcavelos, Portugal

†Email: asterc@bportugal.pt, Banco de Portugal, Avenida Álvaro Pais 2, 1600-007 Lisbon, Portugal

Keywords: Mortgage Refinancing, Network Effects, Inequality, Monetary Policy Transmission

1 Introduction

Mortgage refinancing is the central margin through which interest-rate cuts reach household balance sheets. Refinancing rates fall substantially below frictionless benchmarks, and this inertia is concentrated among lower-income households, so monetary policy reaches the top of the income distribution more strongly than the bottom. The income gradient has two sources. First, direct effects of income on the refinancing decision: high-income households hold larger mortgages, maintain better liquid buffers, and pay closer attention to rate movements. Second, indirect effects through social networks: networks sort by income, so high-income households are connected to peers who themselves refinance more, amplifying their own response. This paper measures both pieces and decomposes their relative contributions. Using Dutch administrative registers that link mortgage contracts to whole-population family, workplace, and neighborhood networks, we provide the first nationally representative causal estimate of peer effects in mortgage refinancing. A one-standard-deviation increase in peer housing-market activity raises active refinancing by 8 percentage points in family networks and 5.7 in neighbor networks, against baseline rates of 9.4 percent. The response rises sharply with income, with a top-quintile coefficient roughly three times the bottom-quintile coefficient in family and workplace networks. We then build a heterogeneous-agent model of refinancing with a reduced-form network friction and an endogenous term premium, calibrated to the family-network IV estimate, and use it to quantify three outputs: the aggregate amplification of the refinancing response to a rate cut, the share of the income gradient in pass-through that is attributable to network sorting rather than to direct income effects on the refinancing decision, and the consumption-equivalent welfare effects of the rate cut by income quintile, with and without sorting.

The Dutch institutional setting is well suited to identify the causal effect of network exposure on refinancing. Mortgage contracts are highly standardized, and underwriting rules are centrally regulated by the government following advice from a national budget institute (NIBUD), so variation in refinancing behavior cannot be attributed to differential contractual incentives (Groot and Lejour, 2018; Thiel, 2019; Caloia, 2024a). Comprehensive administrative registers link individual mortgage contracts, household wealth records, and whole-population family, workplace, and neighborhood networks (van der Laan et al., 2022). Identification exploits the 2016 revision to the loan-to-income (LTI) norm, which mechanically shifted borrowing capacity for potential first-time homebuyers across income and interest-rate cells. The revision affects entry into homeownership but does not bind for refinancing of existing loans, so it shifts peer purchases without altering the refinancing incentives of existing borrowers. We instrument network-level peer purchase rates with the level of network borrowing capacity in 2015 and the policy-driven change in the LTI norm induced by the 2016 NIBUD revision, evaluated at peer

income and the 2015 interest rate.

The estimates document heterogeneity across both networks and households. The family channel is the largest, with a one-standard-deviation increase in peer purchase activity raising refinancing by 8.0 percentage points, an 86 percent increase relative to the family-network baseline rate. The neighbor channel is comparable in magnitude (5.7 percentage points, 61 percent of baseline) but less robust to identification choices, attributable to spatial correlation in network borrowing capacity that is absorbed once network composition controls and postcode clustering are imposed. The coworker channel is small in absolute terms and statistically fragile under clustered inference. Within networks, the response to peer exposure rises with household income in family and workplace networks, with the top-quintile coefficient roughly 3 times larger than the bottom-quintile coefficient. In neighbor networks the gradient is non-monotonic, with elevated responses at both income tails, consistent with homophily and with the location of marginal home purchases at intermediate incomes. This income gradient is the empirical anchor for the distributional channel developed in the model.

The estimates discipline the strength of a network friction but are silent on three questions only the structural model can answer. First, when the policy rate falls, how much does the network channel amplify the aggregate refinancing response, and how much does the endogenous response of mortgage pricing offset it? Second, how much of the income gradient in refinancing pass-through is attributable to network sorting as opposed to direct income effects on the refinancing decision? Third, if sorting were removed, how would the gap between top-quintile and bottom-quintile pass-through change? To answer these, we embed a reduced-form network friction in a discrete-time Bewley environment with fixed-rate mortgage contracts, automatic renewal, and an endogenous term premium pinned down by a competitive zero-profit lender. The structural network parameter θ is calibrated by matching the family-network 2SLS coefficient through an explicit moment condition, and the model is validated against the income gradient in refinancing as an untargeted moment.

We use the model to compute three outputs of the network channel. First, the aggregate amplification of the rate-cut response, identified as the difference between impulse responses with the network channel active and shut down, with the endogenous term premium operating in both. Second, the share of the income gradient in pass-through attributable to social sorting, identified by replacing household-specific network exposure with the population average. Third, the consumption-equivalent welfare effects of the rate cut by income quintile, both with and without sorting, to assess the distributional incidence of monetary policy that is socially transmitted versus driven by direct balance-sheet exposure.

The framework speaks to the design of macroprudential policy. Loan-to-income limits change the composition of peer networks by altering who enters homeownership, and therefore alter

the strength of the network channel through which monetary policy operates. The trade-off between containing housing risk and enabling broad-based pass-through is quantifiable in the model because θ is identified from a macroprudential policy change but does not depend on the source of variation in peer activity. The remainder of the paper proceeds as follows. Section 2 describes the data, the institutional setting, the identification strategy, and the empirical results. Section 4 develops the structural model and reports the counterfactuals and welfare implications. Section 5 outlines our way forward with the model counterfactuals.

1.1 Related Literature

This paper contributes to the social finance literature studying how information, beliefs, and behavior diffuse through social networks to shape financial decisions (Kuchler and Stroebel, 2021). Administrative data have expanded peer-effects research from education and health into saving, borrowing, and portfolio choice (De Giorgi and Pellizzari, 2013; Beshears et al., 2015; Knüpfer et al., 2017; Kalda, 2020), though a central identification challenge remains disentangling endogenous social effects from correlated unobservables (Manski, 1993). Within mortgage markets specifically, existing work documents localized spillovers in refinancing through workplace interactions and geographic proximity (McCartney and Shah, 2022, 2018), while Rehbein and Rother (2020) show that social connectedness shapes equilibrium mortgage outcomes. We depart from these studies in two ways: we focus on exposure to housing-market *participation*—rather than peers' own refinancing choices—as the source of spillovers, and we address the identification challenge using policy-induced variation in borrowing capacity as an instrument for network-level housing activity.

A large literature documents heterogeneity in refinancing behavior driven by search costs, limited attention, and informational barriers (Allen et al., 2014; Bhutta et al., 2020; Cota and Šterc, 2022; Agarwal et al., 2024). Financial literacy externalities play a direct role: less literate households benefit from proximity to more literate peers (Haliassos et al., 2019). At the same time, comprehensive administrative network data reveal that social interactions amplify economic shocks through sizable social multipliers, whether in consumption (Agarwal et al., 2021; De Giorgi et al., 2020) or product adoption (Bailey et al., 2022; Gilchrist and Sands, 2016; Bailey et al., 2018). These studies typically examine a single behavioral margin. Our contribution is to bring the whole-network perspective to mortgage refinancing and to trace information spillovers from housing-market participation through to aggregate refinancing dynamics and welfare.

Our structural model builds on the quantitative literature on mortgage refinancing in heterogeneous agent economies. Closest to our work, Berger et al. (2021) develop a structural prepayment model with idiosyncratic taste shocks and show that the distribution of outstanding mortgage rates generates path-dependent effects of monetary policy. We share their use of Gumbel

taste shocks to smooth the discrete refinancing decision but extend the framework with a network friction that generates an endogenous social multiplier (Glaeser et al., 2003). Greenwald (2018) models the interaction between mortgage credit constraints and monetary transmission in general equilibrium; our contribution is to introduce network-mediated frictions as a complementary channel. Beraja et al. (2019) combine regional empirical evidence on the refinancing channel with a structural model, an approach closely related to ours. On the mortgage choice side, Campbell and Cocco (2003) study optimal fixed-rate versus adjustable-rate choice under income risk; our 5-year versus 10-year term choice adapts similar trade-offs to the Dutch institutional environment with automatic contract renewal (Campbell et al., 2021).

Finally, our paper relates to the literature on household heterogeneity in monetary policy transmission through the refinancing channel. Existing work documents substantial variation in consumption responses to rate changes (Wong, 2021; Eichenbaum et al., 2022) and the role of unhedged interest rate exposure (Auclert, 2019). Our interpretation is consistent with models emphasizing inattention in household finance (Andersen et al., 2020; Berger et al., 2025). Closest in spirit, Zhou (2025) shows that social learning about cheaper lenders strengthens monetary policy transmission; we focus instead on housing-market participation as the source of social learning. While this literature concentrates on U.S. markets, European mortgage markets—more concentrated and characterized by automatic rate resets—remain less well understood. Evidence linking financial sophistication to mortgage choice extends to the Netherlands (Van Ooijen and van Rooij, 2016; Coen et al., 2023), and recent work shows that macroprudential policy changes increase housing turnover (Caloia et al., 2025). We exploit this setting to trace spillovers from first-time buyers to refinancing decisions of existing homeowners and quantify their aggregate consequences.

2 Data and Descriptives

Our empirical analysis combines two main Dutch administrative data sources and covers the period 2016 to 2017 to a set of policy parameters (maximum debt limits). Individual and household level registry data are merged with loan level mortgage records, allowing us to observe mortgage refinancing decisions, a unique feature of our data, together with detailed borrower characteristics and social network exposure. The loan level mortgage data cover approximately 80 percent of outstanding residential mortgages in the Netherlands during this period. The analysis sample is obtained after matching the mortgage records with household registry and network data and therefore reflects the set of mortgages for which all required information is jointly observed.

The resulting dataset allows us to study mortgage refinancing decisions at the individual

level while exploiting rich information on household characteristics, mortgage contract terms, and multiple dimensions of network exposure. The unit of observation throughout the analysis is the household reference person. More detailed loan-level information is used to recover the assigned mortgage interest rate, which allows us to precisely match the relevant debt limits. All individual and network characteristics are measured prior to 2017, housing purchase activity within networks is observed in 2016, and refinancing outcomes are measured in 2017.

2.1 Dutch Household Registry Data

Administrative household registry data from Statistics Netherlands (CBS) provide the core individual and household level characteristics used in the analysis. These data are organized at the individual level and linked to households using unique identifiers that allow consistent tracking across administrative sources. Several registries are used for this purpose. These include the municipal register, which provides information on sex, country of birth, and household position, as well as education records, tax files, and land registry data. The latter are essential to identify spells of cohabitation and residential moves. So, for each individual, we observe a rich set of socio-economic characteristics, including age, gender, income, education, immigration status, degree of urbanization of the area of residence, and detailed balance sheet information covering financial assets, housing wealth, and debt.

Household members are observed together with information on their position within the household. We define the household decision maker as the main income earner (breadwinner) and attribute mortgage refinancing decisions to this individual. This definition is consistent with the institutional setting in the Netherlands, where mortgage contracts are typically held jointly, but only the main earner's income fully counts toward mortgage qualification, while financial decisions are taken at the household level.

Network Data Administrative personal network data produced by Statistics Netherlands (CBS) are used to measure network exposure. These data are drawn from a whole population network constructed from linked administrative registers (van der Laan et al., 2022). The network records institutional and legally observed relationships that imply repeated exposure and a higher likelihood of interaction, conditional on relationship type, rather than direct measures of social contact.

The analysis focuses on three network layers: family, coworker, and neighborhood networks. Family networks are defined using legally recorded parent child and co residential partnership relations and include nuclear family members, partners, in laws, and step relatives. Coworker networks link individuals employed at the same firm during the year, accounting for multiple jobs and job changes. In large firms, coworker links are restricted to geographically closest hun-

dred colleagues to proxy for shared workplaces and more likely interaction. Neighborhood networks are based on residential address information and connect individuals to residents of the ten closest addresses, with randomization at the cutoff to ensure complete population coverage, including individuals living in rural areas.

For each household reference person, network measures are constructed by aggregating characteristics of connected individuals within each network layer. These measures capture exposure to network level events rather than direct interaction or communication. All network links and network characteristics are measured prior to refinancing activity, so network composition and exposure are predetermined with respect to the outcome of interest.

Loan-Level Mortgage Data We complement the household registry data with loan level mortgage data collected by De Nederlandsche Bank (DNB). The Loan Level Data (LLD) are based on regulatory reporting by mortgage lenders and cover approximately 80 percent of outstanding residential mortgages in the Netherlands over the period 2012 to 2018. The data are collected at the loan level and allow individual loans to be observed together with detailed contract characteristics (Mastrogiacomo et al., 2015). The information contained in the data is comparable to that collected under the ECB Loan-Level Data Initiative, through which the European DataWarehouse (EDW) gathers and publishes detailed loan-level information on securitized loans in Europe as part of the post-crisis transparency framework. The Dutch central bank (DNB) collected the same type of information up to 2018, and extended the coverage to include non-securitized loans.

For each loan, the LLD report information on the outstanding balance, interest rate, contract type, interest rate fixation period, and key leverage measures, including loan to value and loan to income ratios at origination of the entire mortgage contract. Fixed rate mortgages in the Netherlands are characterized by predetermined interest rate fixation periods, after which the interest rate is reset unless the borrower actively refinances, with fixation periods most commonly of 10 years. Borrower level information at origination, including income and employment status at origination, is also observed.

Importantly for our analysis, the LLD allow us to identify refinancing events as active changes in mortgage interest rates or contract terms that occur outside of scheduled interest rate resets. This enables us to distinguish voluntary refinancing decisions from mechanical interest rate adjustments implied by contract structure. More precisely, we identify refinancing as a change in the loan-level interest rate between two adjacent quarters in a year in which no automatic rate reset is scheduled. In addition, we verify that the reported loan purpose in that year is remortgaging only, excluding other purposes such as equity release.

Maximum Borrowing Capacity Data The final dataset contains the statutory borrowing limits implied by the NIBUD affordability tables. Figure 1 summarizes the level of, and year-to-year changes in, maximum loan-to-income (LTI) limits by income group.

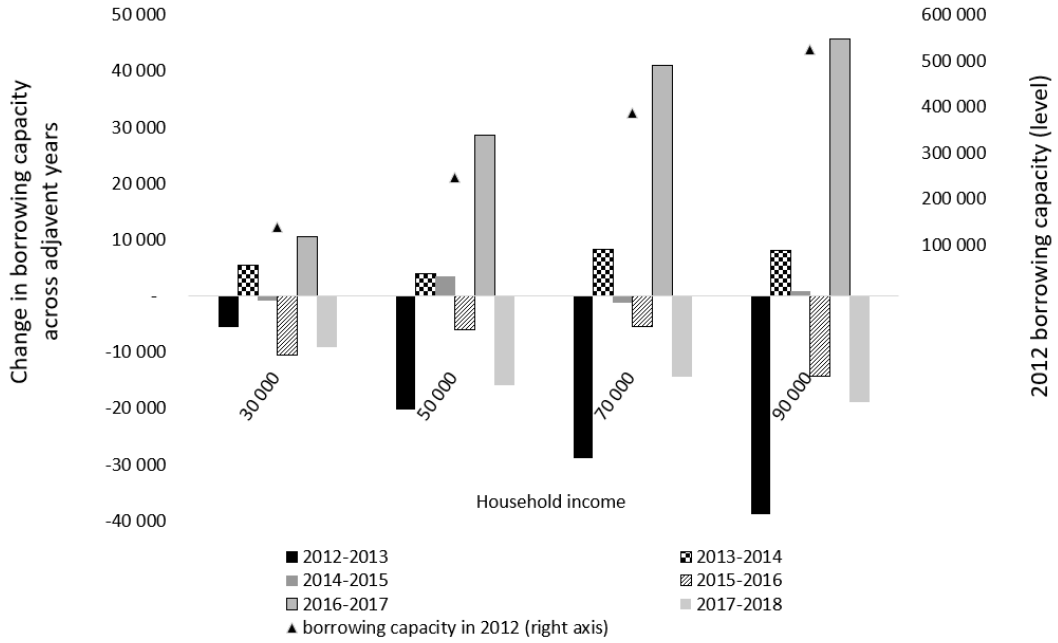


Figure 1: **Maximum borrowing capacity according to LTI limits, a summary.**

Two features are central. First, borrowing capacity increases more than proportionally with income, reflecting higher permissible debt-service ratios at higher income levels. Second, revisions to the affordability tables generate heterogeneous shifts across income groups. In particular, the 2016 reform increased borrowing capacity across much of the distribution, with larger absolute gains at middle and higher incomes. Because these changes are set centrally and apply mechanically to income–interest-rate cells, they generate exogenous, income-specific variation in mortgage eligibility. We exploit this 2016 policy-induced heterogeneity in the empirical analysis below.

2.2 Sample Construction and Timing

The analysis sample consists of households with an active mortgage in both 2016 and 2017 that can be observed in the loan level mortgage data and linked to the household registry. Mortgage records are merged to household level information using unique identifiers, and network measures are attached to the household reference person. We restrict the sample to borrowers who have held their mortgage for at least two years by 2016 in order to exclude refinancing activity

mechanically related to recent mortgage origination.

Individual and household characteristics, as well as network level aggregates, are measured in 2015. These variables are therefore predetermined with respect to peers housing purchase decisions in 2016. Mortgage contract characteristics used as controls, including interest rates, contract type, loan to value ratios, and remaining time until the next scheduled interest rate reset, are measured in 2016 prior to the refinancing outcome.

The outcome variable is an indicator equal to one if the borrower actively refinances their mortgage during 2017.[1] This timing structure ensures that peer housing purchases precede refinancing decisions and establishes a clear temporal ordering between peer exposure and individual refinancing behavior.

Final Sample Descriptives Table 1 summarizes borrower and mortgage characteristics across samples constructed using different network layers. Borrowers in the analysis sample are typically in their early forties, have stable employment and income, and hold mortgages with current loan to value ratios in the range of 70 to 80 percent. Leverage at origination is moderate, with loan to income ratios between 3.5 and 3.8 across samples. Mortgage contracts are similar across network definitions, with most borrowers holding fixed rate mortgages, being customers of one of the four largest banks, and facing their next scheduled interest rate reset several years ahead. Refinancing activity over the sample period ranges from 5 to 8 percent.

Table 2 reports descriptive statistics for network characteristics. Network peers are broadly similar across coworker, family, and neighborhood networks along dimensions such as household structure, education, and housing market exposure. Family networks tend to include slightly older and wealthier peers, while coworker and neighborhood networks include relatively younger peers with lower average income and wealth, reflecting differences in underlying social environments. Importantly for the analysis, the share of peers purchasing a home in 2016 is of comparable magnitude across network layers, indicating that variation in peer exposure is not mechanically driven by large differences in network composition.

2.3 Institutional Details

Mortgage lending in the Netherlands is subject to borrower based macroprudential regulation aimed at limiting household risk taking and safeguarding financial stability. Central to this framework are loan to value (LTV) and loan to income (LTI) limits, which are set at the na-

[1] We observe refinancing only when it occurs with the same lender. Cross-lender refinancing is not observed in our data. However, switching in the Dutch mortgage market is limited. Monetary switching costs, covering notary, appraisal, and insurance fees, amount to approximately €3,500 for an average household, in addition to non-monetary hassle costs (Thiel, 2019). Consistent with these frictions, switching rates at renewal are low, around 3 percent in recent years. We therefore capture the dominant refinancing margin in this institutional setting.

Variable	Coworkers	Family	Neighbors
Refinancing and demographics			
Actively refinances	8.5%	9.3%	9.4%
Age	45.1	48.1	48.4
Male	64%	66%	66%
Education and background			
Elementary education	6%	7%	7%
Secondary education	25%	24%	23%
Bachelor	20%	18%	18%
Master or higher	12%	10%	10%
Dutch born	86%	90%	86%
Household composition			
Single no children	19%	18%	19%
Couple no children	24%	30%	30%
Couple with children	50%	46%	45%
Single parent	7%	6%	6%
Share dual-earner couples	14%	14%	13%
Income, wealth, and leverage			
Gross income 2015 (EUR 1,000)	61.3	56.7	56.8
Household wealth excl. housing (EUR 1,000)	67.5	88.7	87.9
Original LTI	3.5	3.4	3.4
Current LTV (%)	71.1	67.6	67.4
Mortgage characteristics			
Mean interest rate	3.6%	3.5%	3.5%
Mortgage insurance	41%	34%	34%
Customer of top-4 bank	91%	92%	92%
First future reset year	2023.2	2022.9	2022.9
Location			
Highly urban	21%	18%	20%
Urban	33%	32%	32%
Somewhat urban	18%	18%	18%
Almost rural	20%	22%	21%
Rural	8%	9%	9%
Observations	794,293	1,015,757	1,136,017

Table 1: **Borrower and mortgage characteristics by network sample.** Descriptive statistics for borrower, household, and mortgage characteristics in the analysis sample. Columns correspond to samples constructed using different network layers (coworkers, family, neighbors). Borrower characteristics are measured at the household level; mortgage characteristics refer to the active mortgage. All variables are measured prior to refinancing outcomes.

tional level by the Dutch government and applied by mortgage lenders at origination. While the LTV limit restricts borrowing relative to collateral value, the LTI limit constrains borrowing capacity relative to household income and plays a key role in determining mortgage eligibility (Caloia, 2024a).

In the Netherlands, LTI limits are set by the government following advice from the National Institute for Family Finance Information (NIBUD). These limits are derived from affordability

Variable	Coworkers	Family	Neighbors
Household structure and demographics			
Mean age of peers	40.9	44.1	40.4
Share singles without children	16%	12%	14%
Share couples without children	24%	32%	27%
Share couples with children	52%	49%	51%
Share singles with children	7%	5%	7%
Share individuals in a couple	64%	67%	53%
Education and background			
Share with lower education	13%	20%	30%
Share with secondary education	29%	23%	19%
Share with bachelor degree	17%	13%	10%
Share with higher education	10%	7%	5%
Share Dutch born	81%	89%	81%
Economic characteristics			
Mean gross personal income 2015 (EUR 1,000)	48.7	39.8	38.5
Mean non-housing wealth (EUR 1,000)	99.1	123.7	123.9
Housing market exposure			
Share of peers purchasing in 2016	1.5%	1.3%	1.2%
Observations	794,293	1,015,757	1,136,017

Table 2: **Peer network characteristics by network type.** Descriptive statistics for network characteristics constructed from administrative personal network data. Network measures are computed at the level of the household reference person by aggregating characteristics of connected individuals within each network layer (coworkers, family, neighbors). All network characteristics are measured in 2015 or earlier and are predetermined with respect to mortgage refinancing outcomes observed in 2017.

standards that cap the debt-service-to-income (DSTI) ratio, based on detailed household budget calculations. These calculations account for living costs, taxation, and other mandatory expenditures.

The DSTI constraint is then translated into a maximum loan-to-income (LTI) ratio, which determines the largest mortgage a household may take out as a multiple of gross annual income. Since debt service depends on the mortgage interest rate, the applicable LTI varies mechanically with both income and prevailing interest rates, using an annuity formula that assumes full repayment over 30 years (Caloia, 2024a).

This procedure results in a set of tables that specify the maximum allowable LTI for each combination of income and market interest rate. Each table contains roughly 1,000 such combinations. Separate tables exist for different borrower groups, most importantly working-age households (our focus) and retirees, as well as for specific cases such as purchases of energy-efficient homes. These distinctions reflect differences in taxation and utility costs across household types and over time.

The LTI norms are recalibrated regularly following updates to the underlying NIBUD bud-

get tables, typically announced at the end of the year and implemented the following January. As a result, borrowing capacity can change over time even within the same household type, or even for households with unchanged income and interest rates. Around 2016, revisions to the affordability calculations led to increases in borrowing capacity for much of the income distribution, with smaller adjustments at lower income levels (Caloia, 2024a), as illustrated in Figure 1. These centrally determined revisions apply heterogeneously across households, depending on their characteristics.

Importantly, observed borrowing capacity is potentially endogenous, as it may also respond to household income choices, including labor supply adjustments. Changes in the interest rate may have both an endogenous component, related to saving decisions, since down payments and interest rates are inversely related, and an exogenous component driven by movements in market interest rates. Recent work exploiting the Dutch institutional setting explicitly decomposes changes in borrowing capacity into components driven by income changes, interest rate changes, and changes in the LTI norm itself (Caloia et al., 2025). This decomposition shows that policy-driven changes in the LTI norm represent exogenous and unexpected shifts in credit availability, unrelated to local housing-market conditions, whereas income-driven changes may partially reflect household behavioral responses.

A key institutional feature of the Dutch LTI framework is that it follows a comply or explain principle. While LTI limits serve as binding benchmarks at mortgage origination, lenders are permitted to exceed the cap in specific and well documented cases, including mortgage refinancing, energy saving investments, bridge loans, and situations involving credibly expected future income growth. As a result, the LTI framework primarily constrains home purchase decisions rather than refinancing behavior (Caloia, 2024a).

Observed borrowing capacity can be written as

$$BC_{i,t} = f(y_{i,t}, r_t, \theta_t)\,, \tag{1}$$

where $y_{i,t}$ denotes gross household income, r_t the prevailing mortgage interest rate, and θ_t the policy-determined LTI norm embedded in the NIBUD affordability tables. In the Dutch institutional framework, the LTI cap is derived from a debt-service-to-income constraint that is translated into a maximum loan amount using an annuity formula that depends mechanically on income and interest rates (Caloia, 2024b; Caloia et al., 2025). As a result, borrowing capacity varies with income, market interest rates, and revisions to the policy parameter θ_t.

Because $f(\cdot)$ is nonlinear, changes in borrowing capacity between $t-1$ and t may reflect simultaneous changes in income, interest rates, and the LTI norm. Following Caloia et al. (2025), we isolate the policy-driven component by constructing counterfactual borrowing capacity hold-

ing income and interest rates fixed at their pre-revision levels. Specifically, define

$$\Delta^{\theta} BC_{i,t} = f(y_{i,t-1}, r_{t-1}, \theta_t) - f(y_{i,t-1}, r_{t-1}, \theta_{t-1}) \,. \tag{2}$$

This term captures the mechanically implied change in maximum borrowing capacity arising solely from revisions to the LTI norm. By construction, $\Delta^{\theta} BC_{i,t}$ abstracts from endogenous adjustments in income and from movements in market interest rates. It therefore represents an exogenous shift in credit eligibility induced by national policy changes.

The remaining change in borrowing capacity,

$$f(y_{i,t}, r_t, \theta_t) - f(y_{i,t-1}, r_{t-1}, \theta_t) \,,$$

reflects income growth and interest rate movements. Income changes may partly reflect labor supply decisions, although labor supply elasticities of primary earners in the Netherlands are small (de Boer et al., 2022; Mastrogiacomo et al., 2017). Interest rate changes may reflect both macroeconomic conditions and borrower-level adjustments in leverage. However, our identification strategy relies exclusively on the policy-driven component $\Delta^{\theta} BC_{i,t}$, which is determined centrally and does not respond to individual borrowing or refinancing decisions.

Table 3 illustrates this decomposition with a stylized example, where borrowing capacity increased due to substantial income, interest rate and policy effects. It considers a household with an annual income of 64,000 euros that borrows in period $t-1$ at an interest rate of 3%. In the hypothetical table, this household is assigned a borrowing cap of 2.5, implying a maximum borrowing capacity of 160,000 euros. In period t, the household's income increases to 80,000 euros and, at an interest rate of 1%, its maximum borrowing capacity rises to 400,000 euros. Notice that while the cap has doubled, the borrowing capacity has more than doubled, because of the income effect. This is a total increase of 240,000 euros, but in the example $\Delta\theta_t$ is only equal to 32,000 euros.

Information versus interest rate exposure Revisions to the LTI norm do not alter prevailing interest rates, contractual reset schedules, or the mechanical gains from refinancing for existing mortgage holders. We condition on a rich set of borrower and mortgage contract characteristics, including the existing interest rate, remaining maturity, loan balance, reset schedule, and borrower observables, which together summarize the financial incentive to refinance. As a result, any effect of network housing market activity on refinancing behavior cannot be attributed to differential exposure to interest rate movements. Instead, identification operates through variation in housing market activity within individuals' observed networks, and the remaining differences in refinancing take up reflect informational or attentional frictions associated with network ex-

		$t-1$	
		r	
		1%	3%
Y	64 000	cap = 2 *128 000*	cap = 2.5 *160 000*
	80 000	cap = 3.5 *280 000*	cap = 3 *240 000*

		t	
		r	
		1%	3%
Y	64 000	cap = 4.5 *288 000*	cap = 3 *192 000*
	80 000	cap = 5 *400 000*	cap = 3.8 *304 000*

Notes: The increased borrowing capacity in t can be decomposed into four components. First, the baseline borrowing capacity in period $t-1$ ($BC_{i,t-1} = 160{,}000$ euros). Second, the change attributable to the policy parameter, $\Delta\theta_t$, holding income and interest rates constant (32,000 euros, corresponding to $192{,}000 - 160{,}000$). Third, the change due to the reduction in the interest rate, holding income constant (Δr_t = 96,000 euros, corresponding to $288{,}000 - 192{,}000$). Fourth, the change due to the increase in income, holding the interest rate constant ($\Delta y_{i,t}$ = 112,000 euros, corresponding to $400{,}000 - 288{,}000$).

Table 3: **Stylized borrowing caps in $t-1$ and t**

posure rather than refinancing incentives themselves.

3 Identification and Empirical Strategy

We examine whether exposure to housing market activity within an individual's social network affects the likelihood of mortgage refinancing. A simple comparison of borrowers with high and low network housing activity would not identify a causal effect. Network housing purchases may reflect shared economic environments, common exposure to local housing market conditions, or similarities in household characteristics such as financial sophistication or attitudes toward debt. These factors may simultaneously influence both peers' decisions to purchase homes and an individual's refinancing behavior. As a result, the observed correlation between network housing activity and refinancing decisions may not reflect a causal spillover effect. To address this, we use two instruments constructed at the network level. The first is the level of network borrowing capacity in 2015, BC_{t-1}, which combines peer income, the prevailing 2015 interest rate, and the 2015 LTI norm. The second is the change in the LTI norm itself, $\Delta\theta_t$, induced by the 2016 NIBUD revision and evaluated at peer income and the 2015 interest rate. The level identifies which networks contain peers near the homeownership margin. The change in θ activates them through a centrally determined rule unrelated to local housing market conditions or to the borrower's own refinancing incentives.

Sources of Variation in Borrowing Capacity Our identification strategy uses two instruments derived from the Dutch borrowing-capacity rule $BC_{i,t} = f(y_{i,t}, r_t, \theta_t)$ in equation (1). The first is the level of network borrowing capacity in 2015, $BC_{i,n,2015}$, computed from peer income and the 2015 LTI schedule. The second is the change in the LTI norm induced by the 2016 NIBUD

revision, $\Delta\theta_{i,n,2016}$, evaluated at peer income and the 2015 interest rate. The level identifies which networks contain peers near the homeownership margin; the change activates them through a centrally determined rule.

Both instruments are exogenous to individual refinancing decisions, for distinct reasons. The level is fixed at 2015 values, before the policy revision and before the 2017 refinancing outcome, and cannot respond to individual decisions. The change is set centrally through the NIBUD update and applies uniformly across households, generating variation in mortgage eligibility that differs across networks only through heterogeneity in income distributions. Sorting on income and observable network characteristics is absorbed by the controls described in Section 3.3.

LTI limits constrain new mortgage origination but do not bind for refinancing of existing mortgages without a new property purchase (Caloia, 2024a). The 2016 revision therefore affects entry into homeownership without altering prevailing mortgage rates, reset schedules, or the mechanical gains from refinancing for existing borrowers. Conditional on borrower contract characteristics, the primary channel through which the revision can influence existing-borrower refinancing is by shifting housing purchase activity in their networks. We measure peer purchases in 2016, immediately following the revision, and refinancing decisions in 2017, so network exposure precedes the refinancing outcome.

3.1 Descriptive Evidence

Figures 2 and 3 provide descriptive motivation for the empirical strategy.

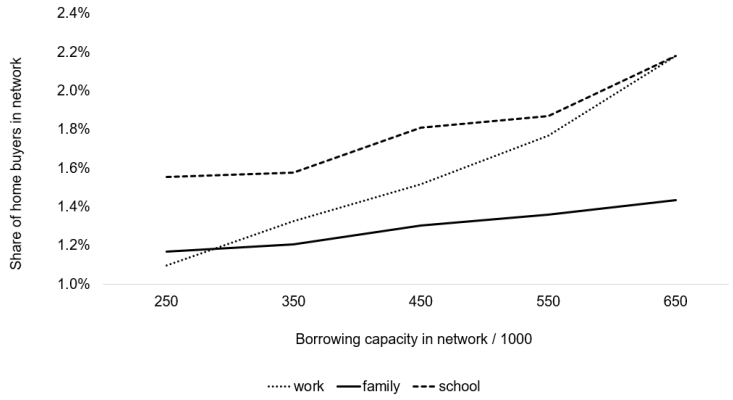


Figure 2: **Share of home buyers in network versus borrowing capacity.** Borrowing capacity is computed using statutory loan-to-income limits and income information measured prior to 2016.

Figure 2 plots the share of home buyers within each network against average borrowing capacity in that network. Borrowing capacity is computed using statutory LTI limits applied to income measured prior to 2016, holding interest rates fixed. In family and workplace networks,

higher borrowing capacity is associated with a higher share of peers purchasing a home, with a particularly steep gradient in workplace networks. In contrast, the relationship is flat in neighborhood networks. These patterns are consistent with the first-stage mechanism: policy-induced increases in borrowing capacity translate into higher entry into homeownership within certain network layers.

Figure 3 compares network housing activity for borrowers with recently originated mortgages to those with older contracts. Recent originators are defined as borrowers whose current mortgage was originated within the previous two years. Across all network definitions, recent originators are embedded in networks with substantially higher housing purchase rates. This correlation underscores the potential endogeneity of network exposure and highlights the need for an instrument that isolates exogenous variation in network-level housing market activity.

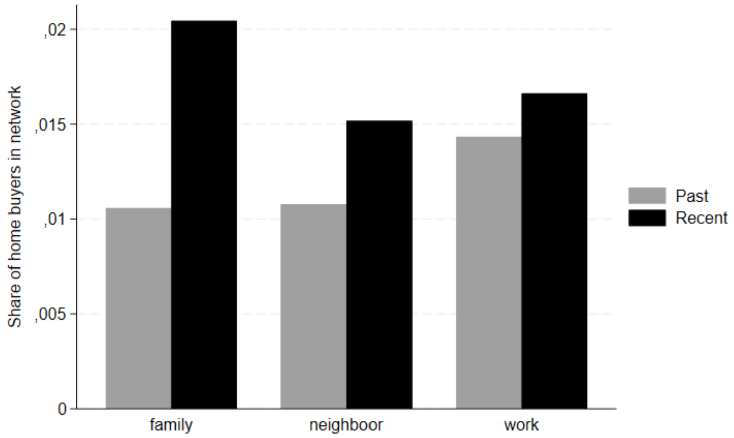


Figure 3: **Share of home buyers in network versus type of originator (past/recent).**

Instrumental Variables Specification Let i index borrowers and $n(i)$ denote the set of individuals connected to borrower i in network layer $n \in \{\text{family}, \text{coworker}, \text{neighborhood}\}$. The outcome variable is $Refinance_{i,2017}$, an indicator equal to one if borrower i actively refinances during 2017. The endogenous regressor is network housing purchase exposure, $PeerPurchase_{i,n,2016} = \frac{1}{|n(i)|}\sum_{j\in n(i)} Purchase_{j,2016}$, where $Purchase_{j,2016}$ equals one if peer j purchases a home in 2016.

We instrument $PeerPurchase_{i,n,2016}$ with two network-level objects. The first is the level of network borrowing capacity in 2015, $BC_{i,n,2015} = \frac{1}{|n(i)|}\sum_{j\in n(i)} BC_{j,2015}$, where $BC_{j,2015} = f(y_{j,2015}, r_{2015}, \theta_{2015})$ is defined in equation (1). The second is the network-level change in the LTI norm induced by the 2016 NIBUD revision, $\Delta\theta_{i,n,2016} = \frac{1}{|n(i)|}\sum_{j\in n(i)} \Delta\theta_{j,2016}$, where $\Delta\theta_{j,2016} = \theta(y_{j,2015}, r_{2015}, 2016) - \theta(y_{j,2015}, r_{2015}, 2015)$ measures the change in the LTI norm at peer j's pre-revision income and interest rate cell.

The first-stage equation is

$$PeerPurchase_{i,n,2016} = \alpha + \beta_1 \, BC_{i,n,2015} + \beta_2 \, \Delta\theta_{i,n,2016} + X'_{i,2016}\gamma + \varepsilon_i,$$

where $X_{i,2016}$ collects predetermined borrower demographics, income, education, household composition, migration background, liquid wealth, mortgage contract characteristics (loan-to-value ratio, loan-to-income ratio at origination, mortgage age, interest rate type, reset horizon, lender fixed effects, origination cohort indicators), and local urbanization and network composition controls. The second-stage equation is

$$Refinance_{i,2017} = \delta + \beta \, \widehat{PeerPurchase}_{i,n,2016} + X'_{i,2016}\pi + u_i,$$

where $\widehat{PeerPurchase}_{i,n,2016}$ is the fitted value from the first stage. Identification arises from cross-network differences in income composition, which generate heterogeneous exposure to the 2015 LTI schedule and to the 2016 revision.

3.2 Identification Assumption

Identification requires that $Cov(BC_{i,n,2015}, u_i) = 0$ and $Cov(\Delta\theta_{i,n,2016}, u_i) = 0$, so that both instruments affect refinancing only through peer housing purchases. Three institutional facts support this assumption. First, LTI limits constrain mortgage origination but do not bind for refinancing of existing loans without a new property purchase. Second, NIBUD affordability revisions are set at the national level and respond to changes in household budgets, not to local housing-market conditions, which addresses reverse causality. Third, the level instrument $BC_{i,n,2015}$ is fixed at pre-revision values and cannot respond to 2017 outcomes. The principal remaining concern is that peer 2015 borrowing capacity correlates with local housing-market conditions through spatial clustering of networks, and these conditions could enter the borrower's refinancing through channels other than peer purchases. We absorb this variation through controls for the borrower's own mortgage rate, reset schedule, leverage, and lender fixed effects, through controls for network composition (income, education, age, urbanization), and by clustering standard errors at the postcode level. Conditional on these, the Hansen J test does not reject joint exogeneity in the family and coworker networks, and yields $p = 0.07$ in the neighbors specification.

Family networks may include kin who share housing-market information through cohabitation; we show this is not the case by overlapping the family and neighbor network, and find a low share of the overlap. Peers may also directly transmit information about the LTI revision, but the revised norm does not change refinancing eligibility for existing borrowers, so this channel

is ruled out.

3.3 Overview of Empirical Estimates

Table 4 reports 2SLS estimates of the effect of peer housing purchases on active mortgage refinancing in 2017, separately for the three network layers. The columns add controls in blocks and probe inference and outcome definitions, with column 4 the baseline specification.

The first stage is strong across all three networks, with the Kleibergen–Paap F statistic well above conventional thresholds in every specification. The two instruments have opposite signs in the partial regression, $BC_{i,n,2015}$ positive and $\Delta\theta_{i,n,2016}$ negative, but their joint contribution to predicted peer purchases is positive in every specification, as reported in the row $BC_{t-1} + \Delta\theta_t > 0$.[2]

The estimated effects vary systematically across networks. In family networks, the IV coefficient is 2.16 percentage points, significant at the 1% level, and stable across the robustness specifications. In neighbor networks, the coefficient is 2.09 percentage points, also significant at 1% and stable. In coworker networks, the coefficient is 0.36 percentage points, smaller in magnitude and sensitive to the inclusion of mortgage controls, and it loses significance once standard errors are clustered at the postcode level. Work-based exposure therefore plays a more limited role in shaping refinancing decisions than family or neighborhood exposure.

Table 5 translates the IV coefficients into magnitudes. A one-percentage-point increase in the share of peer 2016 buyers raises the probability of active refinancing by 2.16 percentage points in family networks, 2.09 in neighbor networks, and 0.36 in coworker networks. Relative to the baseline refinancing rates of 9.4% in family and neighbor networks and 8.5% in coworker networks, this corresponds to a 23% increase in refinancing probability for family, 22% for neighbors, and 4% for coworkers. Scaled to a one-standard-deviation increase in peer share, the implied effects span 0.7 to 8.0 percentage points across networks, with family at the upper bound. The cross-network gradient is therefore both statistical and economic: the spillover from family and neighbor exposure is large in relative terms, while the workplace channel is quantitatively limited.

The Hansen J test sharpens this picture across the three networks. Family networks pass joint exogeneity comfortably in every specification, with the smallest p-value at 0.13. Coworker

[2]The opposite partial signs follow from the economic content of the two instruments. $BC_{i,n,2015}$ captures predominantly cross-sectional variation in peer income, since income is the largest driver of variation in borrowing capacity across peers. $\Delta\theta_{i,n,2016}$ isolates the policy-driven shift in eligibility from the 2016 NIBUD revision. The revision raised borrowing capacity disproportionately at higher income levels (Figure 1, Caloia, 2024a), so the two instruments are positively correlated in the data, and the partial coefficient on $\Delta\theta_{i,n,2016}$ measures residual policy variation across networks of comparable income composition. Caloia et al. (2025) document that the revision was contractionary in a subset of income–rate cells, consistent with the negative partial.

		(1)	(2)	(3)	(4)	(5)	(6)
Coworkers	First stage						
	BC_{t-1}	0.00758***	0.00740***	0.00660***	0.00814***	0.00814***	0.00814***
	$\Delta\theta_t$	-0.000331***	-0.000309***	-0.000304***	-0.000157***	-0.000157***	-0.000157***
	$BC_{t-1} + \Delta\theta_t > 0$	Yes	Yes	Yes	Yes	Yes	Yes
	Second stage						
	Share of home buyers in network	0.874***	0.859***	0.339*	0.3610	0.361**	0.3780
	N	794293	794293	794293	794293	794293	794293
	Hansen J stat p-value	0.42	0.44	0.09	0.12	0.09	0.29
	Kleibergen–Paap F	71.2	94.2	91.1	151.1		151
Family	First stage						
	BC_{t-1}	0.00191***	0.00186***	0.00199***	0.00149***	0.00149***	0.00149***
	$\Delta\theta_t$	-0.000150***	-0.000149***	-0.000149***	-0.000154***	-0.000154***	-0.000154***
	$BC_{t-1} + \Delta\theta_t > 0$	Yes	Yes	Yes	Yes	Yes	Yes
	Second stage						
	Share of home buyers in network	2.480***	2.456***	1.975***	2.156***	2.156***	2.770***
	N	1015757	1015757	1015757	1015757	1015757	1015757
	Hansen J stat p-value	0.77	0.82	0.20	0.28	0.20	0.13
	Kleibergen–Paap F	300.5	286.9	264.9	261.6		262
Neighbors	First stage						
	BC_{t-1}	0.000933***	0.000682***	0.00101***	0.00113***	0.00113***	0.00113***
	$\Delta\theta_t$	-0.000163***	-0.000146***	-0.000148***	-0.000121***	-0.000121***	-0.000121***
	$BC_{t-1} + \Delta\theta_t > 0$	Yes	Yes	Yes	Yes	Yes	Yes
	Second stage						
	Share of home buyers in network	2.397***	2.416***	2.160***	2.093***	2.093***	2.018**
	N	1136017	1136017	1136017	1136017	1136017	1136017
	Hansen J stat p-value	0.00	0.00	0.00	0.07	0.00	0.14
	Kleibergen–Paap F	79.0	155.2	156.3	169.5		169.5
Controls	Demographic controls & IR	Yes	Yes	Yes	Yes	Yes	Yes
	Household welfare indicators	No	Yes	Yes	Yes	Yes	Yes
	Mortgage characteristics	No	No	Yes	Yes	Yes	Yes
	Network controls	No	No	No	Yes	Yes	Yes
	Clustered standard errors	Yes	Yes	Yes	Yes	No	Yes
	Includes passive refinancers	No	No	No	No	No	Yes

Table 4: **Peer effects in mortgage refinancing: IV estimates by network type.** 2SLS estimates of the effect of peer housing purchases on the probability of active mortgage refinancing, separately for three network definitions. The dependent variable is $Refinance_{i,2017}$, an indicator for active refinancing during 2017. The endogenous regressor is $PeerPurchase_{i,n,2016}$, the share of peers in network n purchasing a home in 2016. The instruments are the level of network borrowing capacity in 2015, $BC_{i,n,2015}$, and the network-level change in the LTI norm induced by the 2016 NIBUD revision, $\Delta\theta_{i,n,2016}$. The row "$BC_{t-1}+\Delta\theta_t > 0$" reports a Wald test of the joint contribution of the two instruments to predicted peer purchases at network-average values. Column (4) is the baseline; columns (1) to (3) progressively add household welfare, mortgage, and network composition controls; column (5) is identical to (4) but does not cluster standard errors at the postcode level; column (6) is identical to (4) with passive refinancers added to the dependent variable. All specifications include demographic and interest-rate controls. Standard errors are clustered at the postcode level except in column (5). Hansen J is the over-identification test p-value; Kleibergen–Paap is the first-stage F statistic. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

networks pass at the 5% level but are borderline at 10% in columns 3 and 5 ($p = 0.09$); the small magnitude of the coworker estimate makes the test less informative for this network. Neigh-

		Per 1-pp increase in peer share		Per 1-SD increase in peer share		
	Baseline refi rate (%)	Effect on refi (pp)	Relative to baseline (%)	SD of peer share (pp)	Effect on refi (pp)	Relative to baseline (%)
Coworkers	8.46	0.36	4.3	2.06	0.74	8.8
Family	9.35	2.16	23	3.73	8.04	86
Neighbors	9.36	2.09	22	2.73	5.71	61

Table 5: **Implied magnitudes of network spillovers.** The first column reports the mean active refinancing rate in 2017. The next two columns report the percentage-point increase in refinancing probability implied by a one-percentage-point increase in the share of peer 2016 buyers, and the corresponding relative change against the baseline rate. The last three columns report the standard deviation of peer share in the network, the percentage-point increase in refinancing probability implied by a one-standard-deviation increase in peer share, and the corresponding relative change. Coefficients are taken from column 4 of Table 4.

bor networks reject in columns without the full set of network composition controls or without postcode clustering (columns 1–3 and 5, $p = 0.00$), but do not reject in the baseline column 4 ($p = 0.07$) or in column 6 ($p = 0.14$). The rejection pattern reflects two specific features of the neighbor network: peer 2015 borrowing capacity correlates with local housing-market conditions through shared geography, which the network composition controls absorb, and residual spatial dependence in the standard errors requires postcode clustering to size the test correctly. If correct, this interpretation predicts that the single-instrument estimate using only $\Delta\theta_{i,n,2016}$ is close to the joint estimate. We report single-instrument robustness for all three networks in Appendix Table 8.

The block-wise structure of Table 4 serves as the main robustness check. Columns 1 through 4 sequentially add household welfare indicators, mortgage characteristics, and network composition controls to a baseline of demographic and interest-rate controls. Column 5 drops postcode clustering, and column 6 adds passive refinancers to the dependent variable. The family and neighbor IV coefficients are stable in magnitude and statistical significance across all six specifications, indicating that the result is not driven by borrower selection on observable household or contract characteristics, by observable differences in network composition, by inference choices, or by the definition of refinancing. The coworker estimate is sensitive to mortgage controls and to clustering, consistent with the smaller magnitude of the workplace channel.

From data to model Our estimates have a specific local interpretation. The 2SLS coefficient is the average response among compliers, namely homeowners whose realized peer purchase rate is shifted by the 2016 LTI revision: individuals embedded in networks of peers near the homeownership margin, in income–rate cells affected by the policy. The structural model in Section 4 takes this estimate as the moment that disciplines the network parameter θ, calibrated

so that the same local moment computed on simulated data reproduces the 2SLS coefficient. Whether θ extrapolates to a population-average response is a separate question, addressed by the untargeted income-gradient fit reported in Table 7.

4 A Model of Mortgage Refinancing with Network Frictions

The empirical analysis identifies a causal effect of peer exposure to housing market activity on individual refinancing decisions. These estimates discipline the strength of a network friction but are silent on how such frictions operate in the aggregate when interest rates change. A natural question follows: does the micro effect scale up, or does it wash out once general equilibrium responses are accounted for? To answer this, we build a heterogeneous-agent model of mortgage refinancing in which network exposure shifts the likelihood that households act on refinancing incentives, and use it to quantify the aggregate refinancing response to monetary policy with and without the network channel.

4.1 Environment

Time is discrete and the period is one year. The economy consists of a continuum of infinitely lived households who may be either renters or homeowners, and a competitive banking sector that prices fixed-rate mortgage contracts. Households save in a risk-free bond that pays the interest rate r, determined in equilibrium by market clearing in the bond market. The banking sector funds mortgages at the same rate r and sets mortgage rates to break even; the term premium therefore captures all non-funding costs of mortgage provision (Section 4.2).

Renters choose consumption and savings each period, and decide whether to purchase a home by taking out a fixed-rate mortgage subject to a loan-to-income constraint. Homeowners face an additional choice: whether to refinance their existing mortgage at the prevailing market rate or continue with the current contract. This refinancing decision depends on the gap between the existing mortgage rate and current market rates, on the remaining term of the contract, and on the recent housing market activity of the household's peers. If a homeowner does not refinance and the contract reaches its maturity date, the rate resets automatically to the current market rate.

Households face idiosyncratic income risk, governed by a permanent-transitory process:

$$\begin{aligned} \epsilon_t &= z_t + \xi_t, \\ z_t &= \rho\, z_{t-1} + \eta_t, \end{aligned} \tag{3}$$

where $\xi_t \sim \mathcal{N}(0, \sigma_\xi^2)$ is a transitory shock and $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$ is a permanent shock. We discretize

the process using the Rouwenhorst method. Earnings are $y = w\epsilon$, where w is a fixed wage rate normalized to one.

Housing is indivisible: each homeowner holds a single property of fixed value h. Mortgages are interest-only, so the outstanding principal m remains constant throughout the contract. Fixed-rate contracts have term $\kappa \in \{5, 10\}$ years, consistent with the two dominant fixation periods in the Dutch mortgage market. The variable $g \in \{0, 1, \ldots, \kappa - 1\}$ tracks the remaining periods until automatic renewal. When $g = 0$ and the household does not refinance, the contract renews at the prevailing market rate r_κ and g resets to $\kappa - 1$. When $g > 0$ and the household does not refinance, the existing rate r_m continues and g decreases by one. Formally, the passive transition is:

$$(r_m, \kappa, g) \xrightarrow{\text{no refinance}} \begin{cases} (r_\kappa, \kappa, \kappa - 1), & \text{if } g = 0, \\ (r_m, \kappa, g - 1), & \text{if } g > 0. \end{cases} \tag{4}$$

The loan-to-income constraint applies only at mortgage origination:

$$m' \leq \overline{\text{LTI}} \cdot w\epsilon. \tag{5}$$

Crucially, this constraint does not bind at refinancing, where the household retains the existing principal m. This asymmetry mirrors the Dutch institutional setting and is central to the identification strategy in Section 2: the LTI reform affects potential buyers (who face the constraint at origination) but not existing homeowners (who refinance without re-qualifying).

We describe the lender's pricing problem first, since mortgage rates enter all household decisions, and then turn to renters and homeowners in sequence.

4.2 Lender and mortgage pricing

A representative, competitive, risk-neutral bank issues fixed-rate mortgages of term $\kappa \in \{5, 10\}$, each paying a fixed coupon r_κ per period that the loan survives. Borrowers may prepay in any period by refinancing. Under free entry, each contract yields zero expected profit. With a flat prepayment hazard across horizons and a unit bank discount factor, the zero-profit condition pins down the equilibrium mortgage rate (Appendix B.2):

$$r_\kappa^* = \frac{1}{\kappa} \sum_{s=0}^{\kappa-1} \mathbb{E}_t[r_{t+s}] + \frac{c_\kappa}{1 - p_\kappa}, \tag{6}$$

where p_κ is the aggregate prepayment rate for term-κ contracts, computed endogenously from the household distribution as the fraction of term-κ borrowers who refinance in a given period. The first term is the average expected funding cost over the contract horizon. The second term is

the endogenous term premium $\Theta_\kappa(p_\kappa) \equiv c_\kappa/(1-p_\kappa)$, which rises with prepayment risk: higher refinancing shortens the expected life of the mortgage and forces the lender to spread the per-period cost c_κ over fewer periods.

We linearise around the steady-state prepayment rate $\bar{p}_\kappa$:

$$r^*_\kappa \approx \frac{1}{\kappa}\sum_{s=0}^{\kappa-1} \mathbb{E}_t[r_{t+s}] + \hat{c}_\kappa + \bar{\kappa}_\kappa\, p_\kappa, \tag{7}$$

where $\bar{\kappa}_\kappa = c_\kappa/(1-\bar{p}_\kappa)^2$ governs the sensitivity of the term premium to prepayment risk. The per-period cost c_κ is calibrated to match the average observed mortgage-swap spread for each term; $\hat{c}_\kappa$ and $\bar{\kappa}_\kappa$ then follow from the structural restriction without additional free parameters.

4.3 Renters

At the beginning of each period, renters draw idiosyncratic taste shocks τ^{stay} and τ^{mort}, independently distributed as Gumbel$(0, \sigma^{own})$, and choose between continuing to rent and taking out a mortgage:

$$V^r(a,\epsilon,r) = \mathbb{E}_\tau \max\left\{V^{r,stay}(a,\epsilon,r) + \tau^{stay},\; V^{r,mort}(a,\epsilon,r) + \tau^{mort}\right\}. \tag{8}$$

The parameter σ^{own} governs the dispersion of unobserved preferences for homeownership. Taste shocks smooth the ownership entry decision and are independent of the network friction.

Households who continue to rent solve a standard consumption-savings problem:

$$V^{r,stay}(a,\epsilon,r) = \max_{c,\,a'} \left\{u(c) + \beta\, \mathbb{E}_{\epsilon',r'}\left[V^r(a',\epsilon',r')\right]\right\} \tag{9}$$

subject to

$$\begin{aligned} &c + a' \le w\epsilon + (1+r)a - \text{rent}(w\epsilon), \\ &c \ge 0, \quad a' \ge 0. \end{aligned}$$

Liquid savings a earn the risk-free rate r, determined in equilibrium by bond market clearing (Section 4.5).

Households who take out a mortgage choose the contract term $\kappa \in \{5, 10\}$ to maximize value:

$$V^{r,mort}(a,\epsilon,r) = \max\left\{V^{r,mort,5}(a,\epsilon,r),\; V^{r,mort,10}(a,\epsilon,r)\right\}. \tag{10}$$

Conditional on the chosen term κ, new homeowners solve

$$V^{r,mort,\kappa}(a,\epsilon,r) = \max_{c,\,a',\,m'} \left\{ u(c) + \beta\, \mathbb{E}_{\epsilon',r'}\left[V(\mathbf{x}')\right] \right\} \tag{11}$$

subject to

$$\begin{aligned}
&c + a' \leq w\epsilon + (1+r)a - \text{rent}(w\epsilon) - \mathcal{T}^h(m'),\\
&m' \leq \overline{\text{LTI}} \cdot w\epsilon, \quad m' \leq \overline{\text{LTV}} \cdot h,\\
&c \geq 0, \quad a' \geq 0,\\
&g' = \kappa - 1, \quad r'_m = r_\kappa, \quad \kappa' = \kappa,
\end{aligned}$$

where $\mathcal{T}^h(m')$ is the cost of mortgage origination (including the down payment), $\overline{\text{LTV}}$ is the loan-to-value cap, and $\mathbf{x}' = (a', m', \epsilon', r'_m, r', \kappa', g')$ collects next-period homeowner states. The first mortgage payment occurs next period; only origination costs are paid at entry.

4.4 Homeowners

Let $\mathbf{x} \equiv (a, m, \epsilon, r_m, r, \kappa, g)$ collect the individual state variables of a homeowner: liquid assets, mortgage balance, productivity, mortgage rate, the short rate, contract term, and remaining periods until renewal. At the beginning of each period, each homeowner draws idiosyncratic taste shocks τ^{refi} and τ^{norefi}, independently distributed as Gumbel$(0, \sigma^{refi})$, and chooses between voluntary refinancing and continuing with the current contract:

$$V(\mathbf{x};\, \iota) = \mathbb{E}_\tau \max \left\{ V^{norefi}(\mathbf{x}) + \tau^{norefi},\; V^{refi}(\mathbf{x}) + \tau^{refi} + v(\iota) \right\}. \tag{12}$$

The parameter σ^{refi} governs the dispersion of unobserved preferences over refinancing, with $\mu = 0$ by normalisation. Taste shocks smooth the refinancing decision and are independent of the network friction.

The term $v(\iota)$ captures network peer effects on the refinancing decision, where $\iota \equiv \Delta\mathbb{h}_{-1}$ denotes the lagged change in the aggregate homeownership rate. We specify $v(\iota)$ as linear:

$$v(\iota) = \theta\, \Delta\mathbb{h}_{-1}, \tag{13}$$

where θ governs the strength of social influence. When homeownership rises in the household's network ($\Delta\mathbb{h}_{-1} > 0$), the utility payoff from refinancing increases, capturing the idea that households become more inclined to act when peers are active in the housing market. The converse holds when $\Delta\mathbb{h}_{-1} < 0$. Since ι is an aggregate variable, each household takes it as given. The

parameter θ is disciplined by the empirical estimates through the moment condition derived in Section 4.6.[3]

The model uses a single aggregate network variable, common to all households. Cross-sectional heterogeneity in network responses arises endogenously from the interaction between this common signal and household-specific financial states. Households near the refinancing margin, where the financial gain from a rate reduction is moderate relative to the cost, are most responsive to $v(\iota)$. Since proximity to the margin varies systematically with income, mortgage balance, and contract terms, the model generates an income gradient in network effects without requiring household-specific network exposure. Section 4.7 assesses the model's ability to reproduce the empirical income gradient documented in Section 2.

In each period, homeowners face an exogenous reallocation shock: with probability $\lambda \in (0, 1)$ the household must sell the property and revert to renter status at the beginning of the next period. Upon reallocation, the household liquidates the property and pays off the outstanding mortgage, receiving net equity $h - m'$ as liquid assets. With probability $1 - \lambda$, the household continues as a homeowner. This shock generates realistic turnover in the housing market. All continuation values below incorporate it.

Non-refinancing households. If $g = 0$, the contract renews automatically at the prevailing market rate:

$$V^{norefi}(\mathbf{x})\big|_{g=0} = \max_{c,\, a'} \left\{ u(c) + \beta\, \mathbb{E}_{\epsilon', r'} \left[(1-\lambda)\, V(\mathbf{x}') + \lambda\, V^{r}(a' + h - m', \epsilon', r') \right] \right\} \tag{14}$$

subject to

$$\begin{aligned} & c + a' \leq w\epsilon + (1+r)a - r_m \cdot m, \\ & c \geq 0, \quad a' \geq 0, \\ & m' = m, \quad g' = \kappa - 1, \quad r'_m = r_\kappa, \quad \kappa' = \kappa. \end{aligned}$$

If $g > 0$, the contract continues under its existing terms:

$$V^{norefi}(\mathbf{x})\big|_{g>0} = \max_{c,\, a'} \left\{ u(c) + \beta\, \mathbb{E}_{\epsilon', r'} \left[(1-\lambda)\, V(\mathbf{x}') + \lambda\, V^{r}(a' + h - m', \epsilon', r') \right] \right\} \tag{15}$$

[3]The value V^{refi} does not depend on the current mortgage rate r_m: once the household refinances, the old rate is replaced and is therefore not a payoff-relevant state for that branch.

subject to

$$
\begin{aligned}
&c + a' \leq w\epsilon + (1+r)a - r_m \cdot m, \\
&c \geq 0, \quad a' \geq 0, \\
&m' = m, \quad g' = g - 1, \quad r'_m = r_m, \quad \kappa' = \kappa.
\end{aligned}
$$

In both cases, $\mathbf{x}' = (a', m', \epsilon', r'_m, r', \kappa', g')$ with transitions as specified in the constraints.

Refinancing households. Households who refinance lock in a new contract at the current market rate and may switch the contract term:

$$
V^{refi}(\mathbf{x}) = \max\left\{V^{refi,\kappa}(\mathbf{x}),\ V^{refi,\tilde{\kappa}}(\mathbf{x})\right\}, \tag{16}
$$

where $\tilde{\kappa} \neq \kappa$ is the alternative term. The current-period mortgage payment is based on the existing rate r_m; refinancing costs $\mathcal{T}(m)$ are paid this period, and the new rate takes effect next period.

Retaining the current term:

$$
V^{refi,\kappa}(\mathbf{x}) = \max_{c,\, a'} \left\{u(c) + \beta\, \mathbb{E}_{\epsilon', r'}\left[(1-\lambda)\, V(\mathbf{x}') + \lambda\, V^r(a' + h - m', \epsilon', r')\right]\right\} \tag{17}
$$

subject to

$$
\begin{aligned}
&c + a' \leq w\epsilon + (1+r)a - r_m \cdot m - \mathcal{T}(m), \\
&c \geq 0, \quad a' \geq 0, \\
&m' = m, \quad g' = \kappa - 1, \quad r'_m = r_\kappa, \quad \kappa' = \kappa.
\end{aligned}
$$

Switching to the alternative term $\tilde{\kappa}$:

$$
V^{refi,\tilde{\kappa}}(\mathbf{x}) = \max_{c,\, a'} \left\{u(c) + \beta\, \mathbb{E}_{\epsilon', r'}\left[(1-\lambda)\, V(\mathbf{x}') + \lambda\, V^r(a' + h - m', \epsilon', r')\right]\right\} \tag{18}
$$

subject to

$$
\begin{aligned}
&c + a' \leq w\epsilon + (1+r)a - r_m \cdot m - \mathcal{T}(m), \\
&c \geq 0, \quad a' \geq 0, \\
&m' = m, \quad g' = \tilde{\kappa} - 1, \quad r'_m = r_{\tilde{\kappa}}, \quad \kappa' = \tilde{\kappa}.
\end{aligned}
$$

Refinancing costs $\mathcal{T}(m)$ are proportional to the outstanding principal. Households cannot adjust

the mortgage balance at refinancing; only the rate and term change.

4.5 Equilibrium

Let μ_t denote the cross-sectional distribution of households over all individual states at date t. Define aggregate net liquid asset holdings as $A_t \equiv \int a \, d\mu_t$, where the integral is over both renters and homeowners. The economy has an exogenous net asset supply $\bar{B}$. The short rate r_t clears the asset market:

$$A_t(r_t) = \bar{B}. \tag{19}$$

A stationary equilibrium consists of a short rate r^*, mortgage rates (r_5^*, r_{10}^*), household value functions and policy functions, and a time-invariant distribution μ^* such that: (i) household value functions and policy functions solve the Bellman equations in Sections 4.3–4.4, given (r^* ,r_5^*, r_{10}^*, ι^*); (ii) the distribution μ^* is a fixed point of the law of motion induced by the policy functions, the income transition matrix, the contract renewal rules, and the reallocation shock; (iii) the asset market clears at r^*; (iv) mortgage rates satisfy the pricing equation (7); and (v) the network variable is consistent with the distribution: $\iota^* = \mathbb{h}(\mu^*) - \mathbb{h}(\mu_{-1}^*) = 0$, where $\mathbb{h}(\mu) \equiv \int \mathbf{1}\{\text{homeowner}\} \, d\mu$ is the homeownership rate.

The last condition deserves emphasis. In stationary equilibrium the homeownership rate is constant, so $\iota^* = 0$ and the network friction is absent: $v(0) = 0$. The friction activates only on the transition path, when aggregate shocks alter the flow into or out of homeownership and generate $\iota_t \neq 0$. This property has two implications. First, the steady state is identical to a standard Bewley model without network effects, which isolates the network channel cleanly in transition experiments. Second, the network friction creates a feedback loop along the transition path: a rate cut stimulates refinancing and raises homeownership, which increases ι_t, which further encourages refinancing through $v(\iota_t)$. At the same time, the rise in aggregate prepayment feeds back into mortgage pricing through the endogenous term premium $\Theta_\kappa(p_\kappa)$, partially offsetting the stimulus. The net effect on aggregate refinancing is the outcome of these two opposing forces, and its magnitude is a quantitative question that the model is designed to answer.

4.6 Calibration - TBD

Table 6 summarizes the parameter values. We group them into externally calibrated parameters, set to match Dutch data or standard values from the literature, and internally calibrated parameters, determined jointly within the model to match targeted moments.

Preferences and income. Households have CRRA preferences with $u(c) = c^{1-\gamma}/(1-\gamma)$. We set risk aversion $\gamma = 2$, a standard value in the household finance literature. The discount factor β

is calibrated internally to match the median wealth-to-income ratio from CBS data. The income process follows (3) and is discretized using the Rouwenhorst method. We set the persistence parameter ρ, the permanent shock standard deviation σ_η, and the transitory shock standard deviation σ_ξ to match estimates from Dutch administrative income data.

Housing and mortgages. House value is fixed at $h = 4\bar{y}$, where $\bar{y}$ is average income, consistent with the Dutch average house-price-to-income ratio. The loan-to-value cap is $\overline{\text{LTV}} = 0.90$ and the loan-to-income cap is $\overline{\text{LTI}} = 4$, both reflecting Dutch regulatory limits over the sample period. Rent is set to $\text{rent}(w\epsilon) = 0.25 \cdot w\epsilon$, consistent with CBS data on rental expenditure shares for tenants. The mortgage origination cost $\mathcal{T}^h(m')$ equals the required down payment: $(1 - \overline{\text{LTV}}) \cdot h$. Refinancing costs are set to $\mathcal{T}(m) = 0.02 \cdot m$, corresponding to approximately EUR 3,500 on a median Dutch mortgage of EUR 175,000, in line with the estimates in Thiel (2019).

Contract terms and term premia. Contract terms are $\kappa \in \{5, 10\}$ years, matching the two dominant fixation periods in the Dutch mortgage market. The per-period lender cost c_κ is calibrated to match the average observed spread between mortgage rates and the corresponding interest rate swap for each term over the sample period.[4] The steady-state term premium is $\hat{c}_\kappa = c_\kappa/(1 - \bar{p}_\kappa)$ and the prepayment sensitivity is $\bar{\kappa}_\kappa = c_\kappa/(1 - \bar{p}_\kappa)^2$, both pinned down by c_κ and the model-implied steady-state prepayment rate $\bar{p}_\kappa$.

Reallocation and market clearing. The reallocation probability λ is set to match the annual homeownership turnover rate in the Netherlands, approximately 3–4% based on CBS housing transaction data. The net asset supply $\bar{B}$ is determined residually by asset market clearing at the target steady-state short rate r^*, set equal to the average three-month Euribor over the sample period.[5]

Taste shocks. The scale parameter σ^{own} on the ownership margin is calibrated to match the aggregate homeownership rate in the Netherlands (approximately 57%). The scale parameter σ^{refi} on the refinancing margin is calibrated to match the baseline annual active refinancing rate in the data (approximately 9%). In steady state, $v(\iota^*) = 0$, so the refinancing rate depends on financial incentives and σ^{refi} alone, and the two parameters are identified separately.

Network effect. The parameter θ is the object that links the structural model to the empirical estimates. The Gumbel taste shock structure implies that the refinancing probability for a

[4] [TBD: Spread values from DNB mortgage rate data and swap curves.]
[5] [TBD: Average Euribor value over sample period.]

homeowner with state $\mathbf{x}$ and network exposure ι is

$$P^{refi}(\mathbf{x};\, \iota) = \frac{\exp\left[(V^{refi}(\mathbf{x}) + v(\iota))/\sigma^{refi}\right]}{\exp\left[(V^{refi}(\mathbf{x}) + v(\iota))/\sigma^{refi}\right] + \exp\left[V^{norefi}(\mathbf{x})/\sigma^{refi}\right]}. \tag{20}$$

The marginal effect of network exposure on the individual refinancing probability is $\partial P^{refi}/\partial \iota = (\theta/\sigma^{refi}) \cdot P^{refi}(\mathbf{x})(1 - P^{refi}(\mathbf{x}))$. Aggregating over the stationary distribution of homeowners:

$$\frac{\partial \bar{P}^{refi}}{\partial \iota} = \frac{\theta}{\sigma^{refi}} \cdot \mathbb{E}_{\mu}\big[P^{refi}(\mathbf{x})(1 - P^{refi}(\mathbf{x}))\big]. \tag{21}$$

We set θ so that the model-implied average marginal effect equals the 2SLS coefficient from the family network specification in Table **??** ($\hat{\theta}_{2SLS} \approx 1.59$). The moment condition equates the model's population-average marginal effect to the empirical LATE. To the extent that compliers (homeowners whose networks were activated by the LTI reform) are more responsive than the average homeowner, θ represents a lower bound on the population network effect. We assess sensitivity to alternative values of θ calibrated to the coworker and neighborhood estimates.

4.7 Model fit

Before turning to counterfactual experiments, we assess the model's ability to reproduce key features of the Dutch mortgage market that were not directly targeted in the calibration. Table 7 compares model-implied moments to their data counterparts across four dimensions: aggregate refinancing behavior, the income gradient in refinancing, homeownership patterns, and the mortgage leverage distribution.

The aggregate active refinancing rate is targeted and matched by construction through σ^{refi}. The relevant test is whether the model reproduces the cross-sectional patterns that were not targeted. First, we examine the refinancing rate by income quintile. The empirical analysis documents a positive income gradient: higher-income households refinance at higher rates. In the model, this pattern emerges endogenously because higher-income households hold larger mortgages (and therefore save more from a rate reduction), maintain better liquid asset buffers (reducing the effective cost of paying $\mathcal{T}(m)$), and are more likely to hold contracts near the renewal date. Second, we compare the distribution of mortgage rates across households, which reflects the history of refinancing and renewal decisions. Third, we assess whether the homeownership rate and the distribution of mortgage leverage (loan-to-income ratios at origination) are consistent with the data.

Table 6: Model Calibration

Parameter	Symbol	Value	Target / Source	Method
Preferences and income				
Risk aversion	γ	2	Standard	External
Discount factor	β	[TBD]	Wealth-to-income ratio, CBS	Internal
Income persistence	ρ	[TBD]	CBS income panel	External
Permanent shock SD	σ_η	[TBD]	CBS income panel	External
Transitory shock SD	σ_ξ	[TBD]	CBS income panel	External
Housing and mortgages				
House value	h	$4\bar{y}$	Avg. house/income ratio	External
LTV cap	$\overline{\text{LTV}}$	0.90	Dutch regulation	External
LTI cap	$\overline{\text{LTI}}$	4	NIBUD tables	External
Rent-to-income	rent	$0.25 \cdot w\epsilon$	CBS rental data	External
Origination cost	$\mathcal{T}^h$	$(1 - \overline{\text{LTV}}) \cdot h$	Down payment	External
Refinancing cost	$\mathcal{T}(m)$	$0.02 \cdot m$	**?**	External
Contract structure				
Term types	κ	$\{5, 10\}$	Dutch market	External
Lender cost, 5-yr	c_5	[TBD]	Mortgage-swap spread	External
Lender cost, 10-yr	c_{10}	[TBD]	Mortgage-swap spread	External
Taste shocks and network				
Ownership scale	σ^{own}	[TBD]	Homeownership rate ($\approx$57%)	Internal
Refinancing scale	σ^{refi}	[TBD]	Active refi rate ($\approx$9%)	Internal
Network effect	θ	[TBD]	2SLS, family ($\hat{\theta} \approx$1.59)	Internal
Reallocation and market clearing				
Reallocation prob.	λ	[TBD]	Homeownership turnover	External
Net asset supply	$\bar{B}$	[TBD]	Market clearing at r^*	Internal
Steady-state rate	r^*	[TBD]	Avg. 3-month Euribor	External

4.8 Counterfactual experiments - TBD

The model is designed to answer a specific question: how do network frictions shape the aggregate refinancing response to monetary policy? We study this through three experiments, each computed as a perfect-foresight transition following an unanticipated, persistent 100 basis point reduction in the short rate.

Experiment 1: Baseline with network effects ($\theta > 0$). The economy starts in the stationary equilibrium and is hit by the rate shock. The network friction is active: the initial round of refinancing raises the homeownership rate (through increased mortgage take-up by renters who observe cheaper rates), which generates $\iota_t > 0$, which in turn encourages further refinancing through $v(\iota_t)$. This feedback loop produces a hump-shaped refinancing response, with the peak arriving later than in the standard model because the network channel operates with a one-period lag. At the same time, the rise in aggregate prepayment feeds back into mortgage pricing

Table 7: Model Fit: Untargeted Moments

Moment	Data	Model
Refinancing by income quintile (annual rate, %)		
Q1 (bottom)	[TBD]	[TBD]
Q2	[TBD]	[TBD]
Q3	[TBD]	[TBD]
Q4	[TBD]	[TBD]
Q5 (top)	[TBD]	[TBD]
Homeownership and leverage		
Homeownership rate (%)	≈57	[TBD]
Median loan-to-income at origination	[TBD]	[TBD]
Share with $\kappa = 10$	[TBD]	[TBD]
Mortgage rate distribution		
Std. dev. of mortgage rates	[TBD]	[TBD]
Share with rate > current market	[TBD]	[TBD]

through the endogenous term premium, partially offsetting the stimulus. We track the impulse response of the aggregate refinancing rate, the homeownership rate, aggregate consumption, and the distribution of mortgage rates across households.

Experiment 2: No network effects ($\theta = 0$). The network friction is shut down. Households refinance based solely on financial incentives and idiosyncratic taste shocks. This isolates the standard monetary transmission channel and serves as the benchmark against which the network amplification is measured.

The difference between the two impulse responses identifies the network multiplier: the fraction of the total refinancing response attributable to peer effects. This is the central quantitative result of the paper. To illustrate: if the peak aggregate refinancing rate rises by 6 percentage points with $\theta > 0$ and by 4 percentage points with $\theta = 0$, the network multiplier is $(6-4)/6 \approx 33\%$, meaning one-third of the refinancing response is generated by social influence rather than direct financial incentives.

Experiment 3: Remove sorting. The empirical heterogeneity results show that network effects vary with income, consistent with homophily in social networks: high-income households are connected to other high-income households, who are closer to the refinancing margin and more responsive to rate changes. This sorting amplifies inequality in the refinancing response. To quantify this channel, we replace the household-specific distance to the refinancing margin with the population average, giving all households the same effective network exposure. The difference between this experiment and the baseline isolates the contribution of sorting to the aggregate response and to the cross-sectional dispersion in refinancing outcomes.

Welfare. We compute welfare effects in consumption-equivalent variation (CEV) terms: the permanent percentage increase in consumption that would make a household in the no-network economy ($\theta = 0$) indifferent to living in the baseline economy ($\theta > 0$) following the rate shock. The CEV is computed at the individual level and then aggregated, preserving the distributional dimension. We report the population-weighted average CEV as well as the CEV by income quintile, to assess whether network effects are welfare-improving on average and how the gains or losses are distributed across the income distribution.

5 Way Forward

This draft presents the empirical evidence and the structural framework that connects peer effects in mortgage refinancing to the aggregate and distributional incidence of monetary policy. The remaining steps complete the calibration of the structural model and produce the three quantitative outputs the framework is designed to deliver, with two natural extensions on the policy side and the empirical side.

Calibration. The structural model is partially calibrated. External parameters tied to the Dutch institutional setting (LTV cap, LTI cap, contract terms of 5 and 10 years, rent share of income, origination and refinancing costs) are pinned down from regulation and from prior empirical work. The remaining internal parameters are the discount factor (matching the wealth-to-income ratio in CBS data), the dispersion of preference shocks for ownership and refinancing (matching the homeownership rate of approximately 57 percent and the active refinancing rate of approximately 9 percent), the network parameter θ (matching the family-network 2SLS coefficient through an explicit moment condition), and the lender per-period costs c_5 and c_{10} (matching observed mortgage-swap spreads, with the term premium then pinned down without further free parameters). The income process and the reallocation shock are set from CBS estimates and registry-based homeownership turnover.

Three quantitative outputs. Once calibrated, the model delivers the three numbers the introduction promises. The aggregate amplification of the refinancing response to a 100 basis point rate cut is computed as the difference between impulse responses with the network channel active and shut down. The endogenous term premium operates in both cases, so the amplification is reported as the net effect after the pricing offset. The share of the income gradient in pass-through that is attributable to social sorting is computed by replacing household-specific network exposure with the population average and recomputing the cross-sectional refinancing response by income quintile; the narrowing of the top-bottom gap quantifies the contribution of

sorting to the distributional incidence of monetary policy. The consumption-equivalent welfare effects of the rate cut are computed at the household level and aggregated by income quintile, both with and without sorting. The contrast between top and bottom quintiles, and how it changes when sorting is removed, isolates the share of distributional welfare that is socially transmitted versus driven by direct balance-sheet exposure.

Macroprudential and monetary policy. A natural extension uses the calibrated framework to study the interaction between macroprudential and monetary policy. Loan-to-income limits affect the composition of peer networks by altering who enters homeownership at the margin, and therefore alter the strength of the network channel through which monetary policy operates. The model is well suited to this question because θ is identified from a macroprudential policy change but does not depend on the source of variation in peer activity. The counterfactual tightens the LTI norm and reports how the aggregate amplification and the income gradient in pass-through change, expressing the trade-off between containing housing risk and enabling broad-based pass-through in numbers comparable to the baseline counterfactuals.

Empirical extensions. Three robustness exercises support the structural claims and pre-empt the most predictable referee objections. We will report heterogeneity in the IV estimates by income quintile within each network, providing the untargeted moment against which the model fit on the income gradient is tested. We will report single-instrument robustness for the three networks, distinguishing the contribution of the level of network borrowing capacity from the policy-driven change in the LTI norm. We will report placebo network and placebo outcome tests, which address the cross-sectional design's sensitivity to local-economy shocks. Together with the income-quintile heterogeneity table, these exercises strengthen the empirical perimeter on which the structural results rely.

References

Agarwal, S., Grigsby, J., Hortaçsu, A., Matvos, G., Seru, A., and Yao, V. (2024). Searching for Approval. *Econometrica*, 92(4):1195–1231.

Agarwal, S., Qian, W., and Zou, X. (2021). Thy neighbor's misfortune: Peer effect on consumption. *American Economic Journal: Economic Policy*, 13(2):1–25.

Allen, J., Clark, R., and Houde, J.-F. (2014). Price Dispersion in Mortgage Markets. *The Journal of Industrial Economics*, 62(3):377–416.

Andersen, S., Campbell, J. Y., Nielsen, K. M., and Ramadorai, T. (2020). Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review*, 110(10):3184–3230.

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

Bailey, M., Cao, R., Kuchler, T., and Stroebel, J. (2018). The Economic Effects of Social Networks: Evidence from the Housing Market. *Journal of Political Economy*, 126(6):2224–2276.

Bailey, M., Johnston, D., Kuchler, T., Stroebel, J., and Wong, A. (2022). Peer effects in product adoption. *American Economic Journal: Applied Economics*, 14(3):488–526.

Beraja, M., Fuster, A., Hurst, E., and Vavra, J. (2019). Regional heterogeneity and the refinancing channel of monetary policy. *The Quarterly Journal of Economics*, 134(1):109–183.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2021). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review*, 111(9):2829–78.

Berger, D., Milbradt, K., Tourre, F., and Vavra, J. (2025). Optimal mortgage refinancing with inattention. *American Economic Review: Insights*, 7(4):497–515.

Beshears, J., Choi, J. J., Laibson, D., Madrian, B. C., and Milkman, K. L. (2015). The Effect of Providing Peer Information on Retirement Savings Decisions. *The Journal of Finance*, 70(3):1161–1201.

Bhutta, N., Fuster, A., and Hizmo, A. (2020). Paying Too Much? Price Dispersion in the U.S. Mortgage Market. Finance and Economics Discussion Series 2020-062, Board of Governors of the Federal Reserve System (U.S.).

Caloia, F., Mangan, M., and Mastrogiacomo, M. (2025). Macroprudential policy, household credit and house prices. Working paper, SSRN. Available at SSRN: https://ssrn.com/abstract=5559281.

Caloia, F. G. (2024a). Borrower-based measures, house prices and household debt. *Journal of International Money and Finance*, 143:103051.

Caloia, F. G. (2024b). Borrower-based measures, house prices and household debt. *Journal of International Money and Finance*, 143:103051.

Campbell, J. Y., Clara, N., and Cocco, J. F. (2021). Structuring mortgages for macroeconomic stability. *The Journal of Finance*, 76(5):2525–2576.

Campbell, J. Y. and Cocco, J. F. (2003). Household risk management and optimal mortgage choice. *The Quarterly Journal of Economics*, 118(4):1449–1494.

Coen, J., Kashyap, A. K., and Rostom, M. (2023). Price discrimination and mortgage choice. Working Paper 31652, National Bureau of Economic Research.

Cota, M. and Šterc, A. (2022). Financial Skills and Search in the Mortgage Market. CERGE-EI Working Papers 780, The Center for Economic Research and Graduate Education - Economics Institute.

de Boer, H.-W., Jongen, E. L., and Kabatek, J. (2022). The effectiveness of fiscal stimuli for working parents. *Labour Economics*, 76:102152.

De Giorgi, G., Frederiksen, A., and Pistaferri, L. (2020). Consumption Network Effects. *The Review of Economic Studies*, 87(1):130–163.

De Giorgi, G. and Pellizzari, M. (2013). Understanding social interactions: Evidence from the classroom. *The Economic Journal*, 124(579):917–953.

Eichenbaum, M., Rebelo, S., and Wong, A. (2022). State-dependent Effects of Monetary Policy: The Refinancing Channel. *American Economic Review*, 112(3):721–61.

Gilchrist, D. S. and Sands, E. G. (2016). Something to talk about: Social spillovers in movie consumption. *Journal of Political Economy*, 124(5):1339–1382.

Glaeser, E. L., Sacerdote, B. I., and Scheinkman, J. A. (2003). The social multiplier. *Journal of the European Economic Association*, 1(2-3):345–353.

Greenwald, D. (2018). The mortgage credit channel of macroeconomic transmission.

Groot, S. P. and Lejour, A. M. (2018). Financial incentives for mortgage prepayment behavior: Evidence from dutch micro data. *Journal of Housing Economics*, 41:237–250.

Haliassos, M., Jansson, T., and Karabulut, Y. (2019). Financial literacy externalities. *The Review of Financial Studies*, 33(2):950–989.

Kalda, A. (2020). Peer Financial Distress and Individual Leverage. *The Review of Financial Studies*, 33(7):3348–3390.

Knüpfer, S., Rantapusks, E., and Sarvimäki, M. (2017). Formative experiences and portfolio choice: Evidence from the finnish great depression. *The Journal of Finance*, 72(1):133–166.

Kuchler, T. and Stroebel, J. (2021). Social finance. *Annual Review of Financial Economics*, 13(Volume 13, 2021):37–55.

Manski, C. F. (1993). Identification of Endogenous Social Effects: The Reflection Problem. *The review of economic studies*, 60(3):531–542.

Mastrogiacomo, M., Bosch, N. M., Gielen, M. D. A. C., and Jongen, E. L. W. (2017). Heterogeneity in labour supply responses: Evidence from a major tax reform. *Oxford Bulletin of Economics and Statistics*, 79(5):769–796.

Mastrogiacomo, M., Van der Molen, R., et al. (2015). Dutch Mortgages in the DNB Loan Level Data. Technical report, Netherlands Central Bank, Research Department.

McCartney, W. B. and Shah, A. (2018). 'I'll Have What She's Having': Identifying Social Influence in Household Mortgage Decisions. *Consumer Finance Protection Bureau*.

McCartney, W. B. and Shah, A. M. (2022). Household mortgage refinancing decisions are neighbor influenced, especially along racial lines. *Journal of Urban Economics*, 128:103409.

Rehbein, O. and Rother, S. (2020). The role of social networks in bank lending. Technical report, ECONtribute Discussion Paper.

Thiel, J. (2019). Price discrimination, switching costs and welfare: Evidence from the dutch mortgage market.

van der Laan, J., de Jonge, E., Das, M., Te Riele, S., and Emery, T. (2022). A whole population network and its application for the social sciences. *European Sociological Review*, 39(1):145–160.

Van Ooijen, R. and van Rooij, M. C. (2016). Mortgage Risks, Debt Literacy and Financial Advice. *Journal of Banking & Finance*, 72:201–217.

Wong, A. (2021). Refinancing and the Transmission of Monetary Policy to Consumption. Working Paper 2021-57, Princeton University, Economics Department.

Zhou, X. (2025). Financial technology and the transmission of monetary policy: The role of social networks. *Journal of Political Economy Macroeconomics*, 3(1):122–164.

A Additional Data Estimates and Auxiliary Tables

		Joint IV	$\Delta\theta_t$ only	BC_{t-1} only
Coworkers	First stage			
	BC_{t-1}	[TBD]	—	[TBD]
	$\Delta\theta_t$	[TBD]	[TBD]	—
	Share of home buyers in network	[TBD]	[TBD]	[TBD]
	Hansen J stat p-value	[TBD]	—	—
	Kleibergen–Paap F	[TBD]	[TBD]	[TBD]
Family	First stage			
	BC_{t-1}	[TBD]	—	[TBD]
	$\Delta\theta_t$	[TBD]	[TBD]	—
	Share of home buyers in network	[TBD]	[TBD]	[TBD]
	Hansen J stat p-value	[TBD]	—	—
	Kleibergen–Paap F	[TBD]	[TBD]	[TBD]
Neighbors	First stage			
	BC_{t-1}	[TBD]	—	[TBD]
	$\Delta\theta_t$	[TBD]	[TBD]	—
	Share of home buyers in network	[TBD]	[TBD]	[TBD]
	Hansen J stat p-value	[TBD]	—	—
	Kleibergen–Paap F	[TBD]	[TBD]	[TBD]

Table 8: **Single-instrument robustness across networks.** Each panel reports 2SLS estimates for one network with three IV configurations: joint use of $BC_{i,n,2015}$ and $\Delta\theta_{i,n,2016}$, single instrument $\Delta\theta_{i,n,2016}$ only, and single instrument $BC_{i,n,2015}$ only. All specifications include the baseline controls of column 4 of Table 4. Standard errors are clustered at the postcode level. The Hansen J test is defined only for the joint column. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Single-instrument specifications across networks.

	Coworkers	Family	Neighbors
Active refinancing rate, 2017 (mean)	8.46%	9.35%	9.36%
Share of 2016 buyers in network (mean)	1.55%	1.27%	1.22%
Share of 2016 buyers in network (SD)	0.0206	0.0373	0.0273
N	794,293	1,015,757	1,136,017

Table 9: **Outcome and exposure descriptives.** Mean and standard deviation of the active refinancing rate in 2017 and the share of peer buyers in 2016, by network layer. Sample sizes match the IV regressions reported in Table 4.

Measures used for coefficient interpretations

B Model Extensions and Alternative Specifications

This appendix collects alternative specifications and extensions to the baseline model presented in the main text. Section B.1 describes alternative network effect mechanisms considered during model development. Section **??** extends the baseline mortgage pricing to allow for an endogenous term premium.

B.1 Alternative Network Effect Specifications

The baseline model introduces network peer effects through a utility cost $v(\iota) = \theta\, \Delta \mathbb{h}_{-1}$ added to the refinancing branch of the homeowner's Bellman equation. Below we document an alternative specification considered during model development.

Alternative A: Network effects via taste shock scale. In this specification, the network variable ι affects the *scale* of the idiosyncratic taste shocks for refinancing, rather than entering as a separate utility term. Starting from the homeowner's refinancing decision, let $x = (a, m, \epsilon, r_m, r, \kappa, g)$ denote the full individual state. Each household draws iid taste shocks τ^{refi} and τ^{norefi}, so the choice becomes

$$V(x) = \max\left\{V^{norefi}(x) + \tau^{norefi}, \quad V^{refi}(x) + \tau^{refi}\right\}.$$

By assuming $\tau^i \sim \text{Gumbel}(\mu^i, \sigma(\iota))$ with $\iota = \Delta \mathbb{h}_{-1}$, the expected value is

$$\mathbb{E}V = \sigma(\iota) \log\Big(\exp\big((V^{refi} + \mu^{refi})/\sigma(\iota)\big) + \exp\big((V^{norefi} + \mu^{norefi})/\sigma(\iota)\big)\Big) + \sigma(\iota)\gamma,$$

where $\mu^i \in (-\infty, \infty)$ is the location parameter, $\sigma(\iota) > 0$ is the scale parameter dependent on network exposure, and $\gamma \approx 0.57721$ is the Euler–Mascheroni constant. The corresponding refi-

nancing probability is

$$\mathbf{P}^{refi}(x) = \frac{\exp\big((V^{refi} + \mu^{refi})/\sigma(\iota)\big)}{\exp\big((V^{refi} + \mu^{refi})/\sigma(\iota)\big) + \exp\big((V^{norefi} + \mu^{norefi})/\sigma(\iota)\big)}.$$

Note: if $\sigma(\iota) \to \infty$, both probabilities converge to $1/2$ (random choice). If $\sigma(\iota) \to 0$, the choice becomes deterministic.

The network influence is specified as follows. The scale parameter depends on the network variable such that

$$\begin{aligned} \Delta\mathbb{h}_{-1} > 0 &\implies \sigma(\iota) \searrow, \\ \Delta\mathbb{h}_{-1} < 0 &\implies \sigma(\iota) \nearrow . \end{aligned}$$

When network activity is high ($\Delta\mathbb{h}_{-1} > 0$), the taste shock variance falls, making households more responsive to financial incentives. When network activity is low, the large variance of taste shocks dilutes financial incentives and generates inertia.

The baseline specification (utility cost $v(\iota) = \theta\,\Delta\mathbb{h}_{-1}$, combined with independent taste shocks for smoothing) is preferred because it has a direct mapping to the 2SLS coefficient: the parameter θ can be pinned down by matching the estimated average network exposure elasticity. The taste-shock-scale alternative requires additional parametric assumptions on the functional form of $\sigma(\iota)$.

B.2 Derivation of Mortgage Pricing

This appendix derives the mortgage pricing equation used in Section 4.2 from the zero-profit condition of a competitive, risk-neutral bank. The derivation proceeds in two stages. The first stage establishes the general pricing equation under the zero-profit condition. The second stage shows that the term premium is $c_\kappa/(1 - p_\kappa)$ under a flat prepayment approximation, and linearises this expression around the steady-state prepayment rate to obtain the specification $\Theta_\kappa(p_\kappa) = \hat{c}_\kappa + \bar{\kappa}_\kappa \cdot p_\kappa$ used in the main text.

B.2.1 Zero-Profit Condition

Consider a representative bank that funds at the short rate r_t and issues fixed-rate mortgages. For a new mortgage of term κ with principal normalised to one, the bank receives a fixed coupon r_κ each period the loan survives. Let $p_{\kappa,s}$ denote the endogenous prepayment probability at horizon s, determined by household refinancing decisions. Under risk neutrality and a bank

discount factor β_b, the expected present value of the mortgage profit is

$$\Pi_\kappa(r_\kappa; S_t) = \mathbb{E}_t\left[\sum_{s=0}^{\kappa-1} \beta_b^s \left((1-p_{\kappa,s})\, r_\kappa - (1-p_{\kappa,s})\, r_{t+s} - c_\kappa\right)\right], \tag{22}$$

where $c_\kappa > 0$ is a per-period cost specific to term κ (capturing origination, servicing, and regulatory costs) and S_t denotes the aggregate state. Under perfect competition and free entry, each contract yields zero expected profit:

$$\Pi_\kappa(r_\kappa^*; S_t) = 0, \qquad \forall\, \kappa. \tag{23}$$

Setting $\beta_b = 1$ (the bank is risk-neutral and funds at r_t, so its effective discount rate is r_t itself; under perfect foresight the distinction is immaterial) and solving for the equilibrium mortgage rate yields

$$r_\kappa^*(t) = \frac{\sum_{s=0}^{\kappa-1}(1-p_{\kappa,s})\, \mathbb{E}_t[r_{t+s}]}{\sum_{s=0}^{\kappa-1}(1-p_{\kappa,s})} + \frac{\kappa\, c_\kappa}{\sum_{s=0}^{\kappa-1}(1-p_{\kappa,s})}. \tag{24}$$

The first term is a survival-weighted average of expected funding costs: periods with higher prepayment risk receive less weight because fewer loans remain outstanding. The second term is a markup that covers per-period operating costs, scaled by the total expected survival mass of the loan.

B.2.2 Decomposition into Funding Cost and Term Premium

Define the total expected survival mass $S_\kappa \equiv \sum_{s=0}^{\kappa-1}(1-p_{\kappa,s})$ and the horizon-specific survival weights $\omega_s \equiv (1-p_{\kappa,s})/S_\kappa$, which satisfy $\sum_{s=0}^{\kappa-1}\omega_s = 1$. The first term of (24) can be decomposed as

$$\sum_{s=0}^{\kappa-1} \omega_s\, \mathbb{E}_t[r_{t+s}] = \frac{1}{\kappa}\sum_{s=0}^{\kappa-1} \mathbb{E}_t[r_{t+s}] + \underbrace{\sum_{s=0}^{\kappa-1}\left(\omega_s - \frac{1}{\kappa}\right)\mathbb{E}_t[r_{t+s}]}_{\equiv\ \rho_\kappa}, \tag{25}$$

where ρ_κ is the refinancing premium: the wedge between survival-weighted and uniform-weighted funding costs. This wedge captures the covariance between prepayment risk and the funding rate path. When prepayment is more likely in periods with low funding costs (the empirically relevant case, since rate cuts trigger refinancing), the survival-weighted average exceeds the uniform average, and $\rho_\kappa > 0$. Substituting back into (24):

$$r_\kappa^* = \frac{1}{\kappa}\sum_{s=0}^{\kappa-1} \mathbb{E}_t[r_{t+s}] + \underbrace{\rho_\kappa + \frac{\kappa\, c_\kappa}{S_\kappa}}_{\equiv\ \Theta_\kappa}. \tag{26}$$

The term premium Θ_κ therefore collects two components: the refinancing premium ρ_κ and the cost markup $\kappa c_\kappa / S_\kappa$.

B.3 Flat Prepayment Approximation

For tractability, assume the prepayment probability is approximately constant across horizons within a contract: $p_{\kappa,s} \approx p_\kappa$ for all s. This approximation is reasonable in our setting because contract terms are short ($\kappa \in \{5, 10\}$) and the MIT shock framework generates slowly moving aggregate states. Under this approximation, the survival weights become uniform ($\omega_s = 1/\kappa$ for all s), so the refinancing premium vanishes ($\rho_\kappa = 0$) and the survival mass simplifies to $S_\kappa = \kappa(1 - p_\kappa)$. The mortgage rate reduces to

$$r_\kappa^* = \frac{1}{\kappa} \sum_{s=0}^{\kappa-1} \mathbb{E}_t[r_{t+s}] + \frac{c_\kappa}{1 - p_\kappa}, \tag{27}$$

and the exact term premium under this approximation is

$$\Theta_\kappa(p_\kappa) = \frac{c_\kappa}{1 - p_\kappa}. \tag{28}$$

B.4 From Exact to Linear: The Endogenous Term Premium

In the stationary equilibrium, the aggregate prepayment rate is constant at $\bar{p}_\kappa$ and the term premium is simply $\Theta_\kappa = c_\kappa/(1-\bar{p}_\kappa)$. On the transition path, however, p_κ varies over time as interest rate changes and network effects alter refinancing incentives. To capture this variation while preserving computational tractability, we take a first-order Taylor expansion of (28) around $\bar{p}_\kappa$:

$$\frac{c_\kappa}{1 - p_\kappa} \approx \frac{c_\kappa}{1 - \bar{p}_\kappa} + \frac{c_\kappa}{(1 - \bar{p}_\kappa)^2} (p_\kappa - \bar{p}_\kappa). \tag{29}$$

Collecting terms:

$$\Theta_\kappa(p_\kappa) \approx \underbrace{\frac{c_\kappa(1 - 2\bar{p}_\kappa)}{(1 - \bar{p}_\kappa)^2}}_{\equiv\ \hat{c}_\kappa} + \underbrace{\frac{c_\kappa}{(1 - \bar{p}_\kappa)^2}}_{\equiv\ \bar{\kappa}_\kappa} \cdot p_\kappa, \tag{30}$$

which is the specification $\Theta_\kappa(p_\kappa) = \hat{c}_\kappa + \bar{\kappa}_\kappa \cdot p_\kappa$ used in Section 4.2. The slope parameter $\bar{\kappa}_\kappa = c_\kappa/(1-\bar{p}_\kappa)^2$ has a direct structural interpretation: it equals the per-period bank cost scaled by the square of the steady-state survival rate. Higher bank costs or higher baseline prepayment rates both increase the sensitivity of mortgage pricing to refinancing waves. The intercept $\hat{c}_\kappa$ and the slope $\bar{\kappa}_\kappa$ are pinned down by a single data target per contract term (the observed steady-state mortgage-swap spread), together with the structural restriction $\bar{\kappa}_\kappa = \hat{c}_\kappa/(1 - \bar{p}_\kappa)^2$, as discussed

in Section 4.6.

This linear specification captures the key economic channel: when a rate cut triggers a refinancing wave, the rise in p_κ feeds back into higher term premia, partially offsetting the pass-through of lower short rates to new mortgage contracts. The strength of this feedback is governed by $\bar{\kappa}_\kappa$, which is disciplined by the zero-profit micro-foundation rather than set as a free parameter.

Remark on interest rate variance. A more general specification would include the conditional variance of the short rate, $\text{Var}_t(r_{t+s})$, as an additional determinant of the term premium (higher rate volatility increases the option value of prepayment and hence the lender's exposure). However, in our baseline computation we use perfect-foresight MIT shocks, under which $\text{Var}_t(r_{t+s}) = 0$ by construction. The parameter $\bar{\kappa}$ therefore absorbs the unconditional variance of the short rate, which is constant. In an extension with aggregate uncertainty (e.g., Krusell–Smith), the variance term could be reintroduced explicitly.