The cost of homeownership*

João B. Duarte,† Francisco F. Rodrigues† and Luís Teles Morais†

31st October 2025 - **work in progress**

Abstract

Conventional housing affordability metrics contradict the widespread perception of a secular rise in the cost of homeownership. The main index used by policymakers — based on mortgage-payment-to-income ratios — suggests affordability today is similar to 2000 and better than in the 1980s. Yet homeownership has declined among younger generations, who report feeling 'priced out'. We develop a microfounded measure of homeownership cost that resolves this disconnect. Within a standard income fluctuations model with explicit housing finance, we compute the welfare cost of becoming a homeowner as the consumption-equivalent loss relative to free housing, isolating the cost of ownership from shifts in rent-price ratios. Unlike conventional metrics, this captures the full intertemporal burden: saving for a downpayment, purchasing the house, and servicing a mortgage. Our measure shows that median first-time buyers in the US faced a 30% increase in costs since 2000, rising to 60% for low-earners, while the top income quintile saw no increase — consistent with conventional metrics, which reflect the experience of wealthier households. Tighter macroprudential policies contributed to the increase by raising downpayment requirements. The measure is data-light and theoretically sound, providing a practical tool to assess housing affordability over the income distribution and across time and countries.

*Preliminary draft. We are thankful to Nuno Pereira for excellent research assistance. For helpful conversations and suggestions we are grateful to Pedro Brinca, Daniel B. Greenwald, João B. Sousa, Nick Flamang, Marta Cota, as well as to participants in the Nova SBE PhD workshops for useful comments. All errors and omissions remain our own. We gratefully acknowledge the financial support of the Portuguese Science Foundation (FCT) through PhD grants no. SFRH/BD/140788/2018 and SFRH/BD/06616/2021 of the 'la Caixa' Foundation's Social Research Call 2023 under the project code LCF/PR/SR23/57000006.

†Nova School of Business and Economics

1 Introduction

The cost-of-living crisis has dominated policy and public discussions and motivated substantial research. Housing costs — its key component — are widely acknowledged to have risen in recent decades, with many contributing factors: rapid urbanization, supply constraints, demographic shifts, to name a few. Of particular concern is the asymmetric effect of higher housing costs across generations. There is a growing sense that the barrier to enter the housing market has increased for young households. The consequences of this are wide-reaching, affecting wealth distribution and the efficacy of monetary and fiscal policy transmission. Housing has been documented as the main vehicle of upward mobility on the wealth distribution, and many have recently linked the increased uncertainty and higher housing costs to lower MPC[1].

Figure 1: IMF's HAI for US and 40 OECD countries

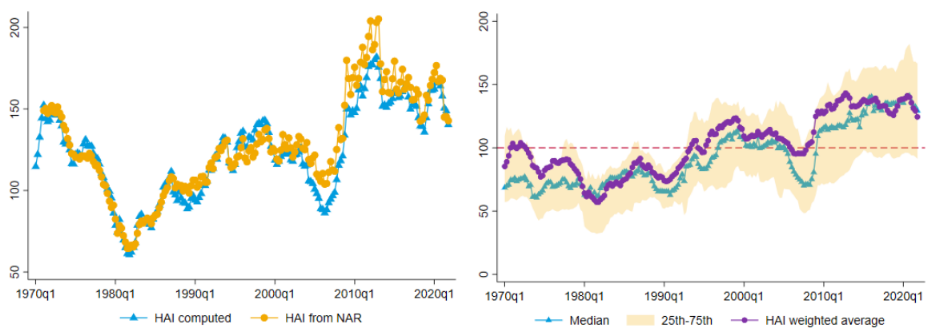

(a) HAI for US

(b) HAI for 40 OECD countries

Notes: figures 1 and 2 from Biljanovska et al. (2023)

Yet the available measures to gauge the evolution of housing costs seem disconnected from this perception. Figure 1, borrowed from Biljanovska et al. (2023) — an IMF working paper documenting a Housing Affordability Index (HAI) for several advanced economies — portrays housing in the early 2020s as more affordable than in the early 2000s, 1990s, and 1980s. Both for the US and the full sample, this broad picture contradicts the experience of many potential homebuyers who feel increasingly "priced out" of ownership. Paz-Pardo (2024) documents falling homeownership rates for younger generations, confirming this disconnect.

Following the great financial crisis (GFC), the HAI depicts an abrupt increase in affordability. This cannot be rationalized by falling house prices, which had started to recover by then. It is even more puzzling given tighter macroprudential regulation that lower loan-to-value (LTV) caps, thus imposing larger downpaymnets. This raises a fundamental question: do conventional metrics adequately capture the cost of transitioning to homeownership? Is it the case that stricter access to credit — more salient to lower income households — simply fade out on the aggregate? Or are

[1] Marginal propensities to consume were shown to be key statistics in shaping the outcomes of monetary policy Kaplan et al. (2018) and fiscal policy Auclert et al. (2024).

these metrics overlooking important aspects to access homeownership?

In this paper, we revisit how to measure housing costs and propose a novel and fully micro-founded framework to do so. We show the conventional metrics fail to accurately account for the evolution of the aggregate housing conditions, let alone conditions for specific segments of the income distribution. Our measure encompasses the housing cost of long-term renters, mortgage or outright-owners and — most importantly — that of households trying to get on the housing ladder: the 1st time homebuyers. Yet it is as data-frugal as conventional metrics, relying on immediately accessible public data and allowing for comparability across time and space (either across countries or regions).

When applied to US data since the early 2000s, our measure tells a different story to that of conventional metrics. Rather than depicting comparable housing costs between 2023 and 2000, *median* 1st time homebuyers faced a 30% rise in costs over this period. In fact, equivalence between today's costs and those 20 years ago is only observed for the top 20% of earners, while 1st time buyers in the lowest income quintiles have seen a rise by 60%. Even when calibrated to the each income quintile, conventional metrics fail to capture such wide variation of cost changes across the distribution. We discuss in detail how the macroprudential policies implemented in the aftermath of the GFC contributed to the deterioration of housing cost for 1st time homebuyers.

At the micro level, it is clear that housing affordability means a different thing for different households. Even before considering deeper heterogeneity aspects, housing costs are shaped by tenure status. Renters are directly exposed to the path of rents, while mortgage-homeowners are primarily exposed to interest rate shifts and changes in their foregone returns of housing wealth — following a *user cost* perspective. However, for 1st time homebuyers — those committed to buying and who rent while saving for a downpayment — both exposures apply. A commitment to buy a house in the near future — say within months — not only changes their budget constraint after purchase, when servicing the mortgage, but also brings about a higher liquidity requirement in each period before purchase, as they build up a savings stock to meet the downpayment, on top of the usual hedging against income uncertainty. While post-purchase implications on budget constraints are well understood and accounted for in conventional metrics, the burden of saving for the downpayment is overlooked by design. Since houses are durable goods typically subject to regulatory LTV caps, such pre-purchase effects are a relevant driving force of housing costs.

The price-to-income ratios is the most common metric in public discourse, both for time and regional comparisons. Its main merit is its interpretability and low data requirements. It is both great for headlines and easy to compute when median house prices and median incomes are available. On top of that, it can also be argued as a proxy for loan values relative to income and downpayment relative to income. However, holding fixed house prices, improving credit conditions in the form of lower rates leaves this indicator unchanged. Moreover, since house prices and policy rates — tracked closely by mortgage rates — tend to be negatively correlated, house prices typically rise when mortgage rates fall. Thus, from a mortgage repayment perspective, price-to-income ratios overstate the importance of house price movements given their impermeability to interest rate shifts.

Mortgage-based indicators, just as the HAI, improve on this. Still with low data requirements, these indices measure affordability through a version of income-to-mortgage-payment ratio. Typically, a fixed-rate and fixed-payment schedule is assumed for a 30-year mortgage contract[2], and given data on mortgage rates, house price, and LTV cap, the index can be produced at a quarter or annual frequency. To restore the interpretation of a cost measure, one could simply take the reciprocal of the index to consider the payment-to-income ratio. In turn, since the mortgage payment is a rescaling of the house price, that reciprocal is essentially a refinement of price-to-income ratios, where mortgage rates and LTV caps are accounted for.

Yet this captures exclusively the burden to repay the mortgage only after meeting any downpayment and payment-to-income (PTI) requirements and thus qualifying for the mortgage. It is, by design, silent on how costly it is to save for the downpayment amid rising rents, higher labor income risk and borrowing constraints. Note also that, even without a binding LTV cap, a PTI constraint might enforce a *de facto* binding downpayment requirement, for the only way to lower the payment-to-income ratio below the cap is to reduce the loan size with a larger downpayment — given a fixed permanent income.

To further illustrate the pitfall of ignoring the pre-purchase effort, consider a tightening of macroprudential regulations that lowers the LTV cap . For the same house price and mortgage rate, now the mortgage payment is lower since the loan size is also lower. Mortgage-based indices then register an increase in housing affordability — i.e, lower housing costs. However, the reduced loan size implies a larger downpayment. Hence, with no change in the income profile, households for whom the LTV cap binds are required to increase their savings. This extra downpayment-motivated savings commitment reduces the cash-on-hand households use to meet negative income shocks, effectively increasing their precautionary savings motive. The effect of this shifting of household resources away from consumption is therefore omitted in mortgage-based indices, but it turns out to be relevant enough to flip the direction of housing costs when properly accounted for.

Our framework avoids this weakness by relying on a self-contained structural approach that keeps data needs to a minimum. We take the standard income fluctuations problem (IFP), with uninsurable labor income risk and a borrowing limit, and include in budget constraints housing expenditures that depend on tenure-status — nesting renters and owners, with or without a mortgage. We keep tenure-status exogenous, with households making the usual consumption-savings choice. Exploiting the main strength of any model, we consider a *free-user* counterfactual in which the household in question simply *does not* incur any housing expenditure. We then use this counterfactual as benchmark to assess the welfare loss accruing from lower non-housing consumption that a given tenure-status entails. In particular, we evaluate the welfare loss of the housing expenditure profile for 1st time buyers: paying rent while accumulating for near-future downpayment and

[2] Mortgage payment is given by $P\theta r_m \left[1 - \frac{1}{(1+r_m)^{T_m}}\right]^{-1}$, where P is house price, θ max LTV, r_m the mortgage rate, and T_m is the mortgage maturity.

thereafter servicing the mortgage while also paying maintenance expenses. Thus, we produce a measure of the cost of becoming a homeowner that fully captures the impact on the intertemporal budget budget constraint — before the purchase, at the purchase and thereafter.

It is important to note that our framework aims to characterise the cost of a *tenure-status*, not to model tenure choice. In fact, it describes one component of that choice — the cost of one tenure alternative — while abstracting from the value households assign to any specific alternative. That is because we measure the welfare loss from the forgone non-housing consumption and abstract altogether from housing preferences. All that we require of housing preferences is separability with respect to non-housing consumption. Moreover, while we emphasize the application of this framework to 1st time buyers, it is perfectly suitable to measure the cost of renters and owners as well. Ad hoc measures applied specifically to renters and to owners are somewhat consensual, given the less complex implications that renting and owning have on the intertemporal budget constraint. Our framework then provides a unified way to measure housing costs across any tenure status.

Related literature Housing affordability has traditionally been assessed through three main lenses: the user-cost model, mortgage-based affordability indexes, and housing supply metrics.

The user-cost framework, pioneered by Poterba (1984), defines affordability in terms of the opportunity cost of housing capital. This measure includes mortgage interest, property taxes, maintenance, and the foregone return of house investment net of the expected capital gains on the house. While theoretically sound, this approach is limited in its practical application by data availability issues — particularly the need to pin down expectations on capital gains and foregone returns — and by its abstraction from credit constraints and household heterogeneity. It also does not rule out a negative cost if expected capital gains are substantial. Recent applications by Battistini and Gareis (2024) to Euro area data show the volatility of this measure in response to interest rate fluctuations, raising questions about its stability as a policy indicator. Although our model counterfactual captures the foregone return when buying a house, in our current application we shut down the capital gains margin to keep data requirements minimal and perform the most conservative horse-race with HAI index. However, such extension could be perfectly accommodated within our framework[3].

Perhaps the most widely used metrics are mortgage-based affordability indexes, such as the HAI, originally developed by the National Association of Realtors and recently calculated for a sample 40 countries by the IMF (Biljanovska et al., 2023). These indexes typically measure the ratio of median household income to the "qualifying income" necessary to obtain a mortgage for a median-priced home. Gan and Hill (2009) extend this logic to a wider distributional perspective on affordability, focusing on the entire distribution of household incomes rather than just the median. As mentioned above, by focusing on housing-related expenses after mortgage issuance, this approach fails to capture the burden to saving for a downpayment — a critical aspect of the cost for 1st time buyers.

[3] And while still avoiding negative costs to the extend the *free-user* counterfactual also earns those capital gains.

A third strand of literature focuses on supply-side metrics, such as the ratio of construction costs to home values as proposed by Glaeser et al. (2005) or the availability of "affordable" housing units relative to the number of households in specific income or age groups Bogdon and Can (1997). The supply perspective has been further developed in recent work by Favilukis et al. (2023), who analyse the role of supply constraints in driving housing affordability issues. While valuable for understanding market dynamics, these metrics tend to focus on broader housing affordability rather than the specific challenges of transitioning from renting to homeownership.

Lastly, also related to our work is the recent literature on incorporating housing costs into overall inflation measures. As discussed in Chodorow-Reich et al. (2025), this requires tracking how tenure choices change with underlying housing costs, which substantially raises data requirements given the need to discipline preferences for owning — in fact they conducted their on survey to elicit these preferences. Although the resulting metric is useful for time comparisons, it is unclear how it could be used for comparisons across space.

The paper is organized as follows. Section 2 details our proposed measure of homeownership costs. Section 3 describes how it can be parametrized with US data to describe the cost of 1st time buyers. Section 4 presents the results on the time-series of cost for a representative household and compares them to the HAI counterpart measure. Section 5 presents the results across the income distribution and section 6 concludes summarizing next steps.

2 Measuring the cost of homeownership

The decision to become a homeowner involves complex intertemporal trade-offs. Buying commits households to substantial future outlays: accumulating sufficient savings for a downpayment, meeting ongoing mortgage obligations (including interest payments and principal amortization), and covering recurring ownership-related costs such as property taxes, insurance, and maintenance. These commitments reduce households' liquidity and limit their flexibility to sustain consumption levels in response to negative future income shocks or other economic changes. In contrast, renting allows households greater flexibility by eliminating long-term financial commitments related to homeownership. This flexibility, however, comes at the cost of exposing households to rent volatility, potentially leaving them vulnerable to rising housing costs that reduce long-term consumption opportunities.

Homeownership also offers several potential benefits. For most households, it represents a unique opportunity to undertake a highly leveraged investment. By financing the purchase largely through debt, households can achieve substantial wealth accumulation over time. Historically, leveraged housing investments have yielded high returns due to house-price appreciation, with long-run returns comparable on a risk-adjusted basis to equities Jorda et al. (2019). Homeownership can also yield additional non-pecuniary benefits, such as stability, control over one's living environment, and social status. Finally, owning a home may grant households access to housing types or locations that are scarce or entirely unavailable in the rental market.

A key aspect affecting this trade-off is the relative cost of renting compared to that of transitioning to ownership. On the one hand, renting costs are fairly easy to measure, given that rents are observed prices that enter households' budget constraint in a trivial way each period. On the other hand, since house purchases are predominantly leveraged, observed house prices do not affect budget constraints in trivial way over time. Buyers incur in mortgage payments after the purchase and anticipate the downpayment before, with a savings intensity that depends on the initial wealth stock.

To rigorously quantify these underlying costs, we propose a structural, dynamic framework that explicitly measures the welfare value of foregone non-housing consumption when transitioning to homeownership. Specifically, we measure the welfare cost of homeownership as the difference in expected lifetime utility between the path of going from renter to owner, at a given point in time, and a counterfactual *free-user* path in which households have no housing related outlays. This approach leverages the strength of the model in producing a counterfactual with virtually no empirical counterpart.

Importantly, by relying on the assumption of separability of preferences between housing and non-durable consumption, c, the approach is able to fully abstract from housing preferences. This abstraction represents a methodological advantage, as housing preferences — such as tastes for location, size, or amenities — are notoriously difficult to measure empirically. Our measure, by construction, thus does not depend on these hard-to-observe parameters, enhancing its robustness and comparability.

Next we detail the model framework used to construct the cost measure.

2.1 Model

We take the standard income fluctuations problem — in which households face uninsurable idiosyncratic income risk and a borrowing constraint — and add the simplest housing cost flow that embeds a renter, outright-owner and mortgage-owner housing outlays. More specifically, households coming into the period with liquid assets b, housing wealth H, housing services consumption h and with a possible mortgage payment x; after observing their income z, make a consumption-savings choice (c, b'). The Bellman equation is:

$$
\begin{aligned}
V_t(z,\, b;\, h,\, H,\, x) = \max_{c,\, b'}\; & u(c) + \beta \mathbb{E}\Big[V_{t+1}(z',\, b';\, h,\, H,\, x)\Big] \\
\text{s.t.} & \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad (2.1) \\
c + b' = & \underbrace{z}_{\text{labor income}} + \underbrace{(1+r_t^b)b}_{\text{return on liq.wealth}} - \underbrace{q_t(h-H)}_{\text{rent}} - \underbrace{x}_{\text{mortgage payment}} - \underbrace{\delta p_t H}_{\text{maintenance}} \\
b' \geq & 0 \\
x' = & \pi_t^{-1} x
\end{aligned}
$$

The housing outlay is made up of (i) rent, deducted of rent income — where q_t is the rental rate per housing unit –; (ii) house maintenance cost for home owners — proportional to housing wealth $p_t H$ –; and (iii) nominally set mortgage payment, x, whose real value decays over time with gross inflation π. We will assume throughout a typical mortgage design of a fixed payment set at origination, given maturity, T_m, and a fixed mortgage interest rate, r^m — that equals the risk-free return rate of liquid savings plus a spread. The model frequency is monthly.

As specified in 2.1, there is no tenure choice, we explicitly model that as exogenous to the household decision. Yet the framework allows for heterogeneity in housing consumption, h, — which is likely linked to permanent income — and the corresponding outlays. Just as housing consumption, housing wealth, H, captures a composite of house-size-quality whose unit is priced at p_t. Clearly, households assign value to their level of housing consumption and also to ownership. However, we single those out in the value function V, which reflects the value assigned to expected lifetime non-durable consumption c alone.

This framework, then nests the value of different housing arrangements, such as renters, who hold no housing wealth, $V_t^{\text{renter}}(z,b,h) \equiv V_t(z,b;h,H=0,x=0)$, or owners $V_t^{\text{owner}}(z,b;h,x) = V_t(z,b;h,H=h,x)$[4]. The counterfactual taken as benchmark to measure housing costs has expected value of $V_t^{\text{free-user}}(z,b) \equiv V_t(z,b;h=0,H=0,x=0)$.

To capture the cost of first-time homebuyers — which henceforth is denoted as *buyer cost* —, we consider the expected value over the buyer path. At an exogenous specified time of purchase, t^*, on top of the typical consumption-savings choice, households choose the mortgage size used to

[4]Strictly speaking, for owners there is a another exogenous state which is time until the mortgage being fully repaid.

finance the purchase, subject to an LTV cap. Upon the purchase, the buyer problem is therefore:

$$
\begin{aligned}
V_{t^*}^{\text{buyer}}(z,\, b;\, h,\, H) = & \max_{c,\, b',\, m^*}\; u(c) + \beta \mathbb{E}\left[V_{t^*+1}^{\text{owner}}(z',\, b';\, H,\, x')\right] \\
& \text{s.t.} \\
c + b' = & \underbrace{z}_{\text{labor income}} + \underbrace{(1+r_t^b)b}_{\text{return on liq.wealth}} - \underbrace{q_t h}_{\text{rent}} + \underbrace{m^*}_{\text{loan issued}} - \underbrace{p_t H}_{\text{house value}} \\
b' \geq & 0 \\
m^* \leq & \theta_t p_t H \\
x' = & \frac{r_t^m (1+r_t^m)^{T_m}}{(1+r_t^m)^{T_m} - 1} m^*
\end{aligned}
\tag{2.2}
$$

while for $t < t^*$:

$$
\begin{aligned}
V_t^{\text{buyer}}(z,\, b;\, h,\, H) = & \max_{c,\, b'}\; u(c) + \beta \mathbb{E}\left[V_{t+1}^{\text{buyer}}(z',\, b';\, h,\, H)\right] \\
& \text{s.t.} \\
c + b' = & \underbrace{z}_{\text{labor income}} + \underbrace{(1+r_t^b)b}_{\text{return on liq.wealth}} - \underbrace{q_t h}_{\text{rent}} \\
b' \geq & 0
\end{aligned}
\tag{2.3}
$$

We thus define the *buyer cost* as

$$
V_0^{\text{free-user}} - V_0^{\text{buyer}}
$$

which captures the welfare loss of buying a house at t^*, with as many periods of anticipation, relative to the counterfactual of no housing expenditures throughout[5]. Crucially, since V_0^{buyer} includes expected $V_1^{\text{buyer}}, \ldots, V_{t^*}^{\text{buyer}}, \ldots V_{t^*+n}^{\text{onwer}} \ldots$, the measure captures, in tandem, the pre and post-purchase impact on the intertemporal budget constraint.

To simplify the analysis, in the results of section 4 and 5, we also abstract away form the role of the house as a financial asset. Homeownership is then strictly a vehicle to save on rents. Although such extension can be accommodated in this framework by accounting for some form of home-equity extraction, it would demand higher data requirements to discipline the expected path of house prices and therefore the expected path of home-equity, as one repays the mortgage and moving forward. This simplification allows this metric to be a potential substitute of conventional ones by keeping at par with data requirements.

Lastly, the way the buyer value is specified also accommodates transitions into homeownership which coincide with moves to housing units that differ significantly in location, quality, or size[6].

[5]One can see this cost measure as the welfare loss of a *purchase shock* at t^* that is anticipated at 0.

[6]Namely, in $V_{t^*}^{\text{buyer}}(z, b;\, h, H)$, the housing unit rented h does not need to be the same as the housing unit that is later bought H.

This is relevant given the empirical evidence on market segmentation: some house-types cannot be found in the rental market while others can only be found in the rental market.

3 Data and calibration

This section describes the data used to construct a yearly time series of the buyer cost. We applied our buyer cost measure to U.S. data between 2000 and 2023. This data also feeds our computation of the cost measure underlying the HAI.

3.1 Stylized facts

Figure 2 shows the path of median house prices, median asking rent, and median household income, all from US Census Bureau and in real terms. Notably, rents depict a steady increase throughout the period, while house prices followed more closely the business cycle, with a slump during the GFC, a soaring period during the pandemic years, and a recent downward trajectory at the end of the sample. Household median income took 15 years to sustainably surpass the 2000 level. It recorded a substantially lower average yearly growth rate than house prices and rents, within the sample.

These series feed the most back-of-the-envelop measure of housing costs. Given the sizeable gap in average growth rate of price and rents with respect to income, both the rent-to-income and the price-to-income ratios track closely the time series for rents and prices, respectively.

3.2 Income process

The income dynamics are standard. There is a constant permanent income component, z^p, and log income follows an AR(1) process:

$$lnz_t = (1-\rho)lnz^p + \rho lnz_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0,\sigma_z)$$

We use the time-path of median household income in figure 2 to pin down the permanent income component for each year of the buyer cost measure $V_0^{\text{free-user}} - V_0^{\text{buyer}}$. Throughout the sample, the persistent component is set to 0.91 and the standard deviation of the idiosyncratic component is set to 0.25, given its stability in the GRID repository Guvenen et al. (2022).

3.3 House size and Initial Wealth

Figure 2: Median of house prices, rent and income over time

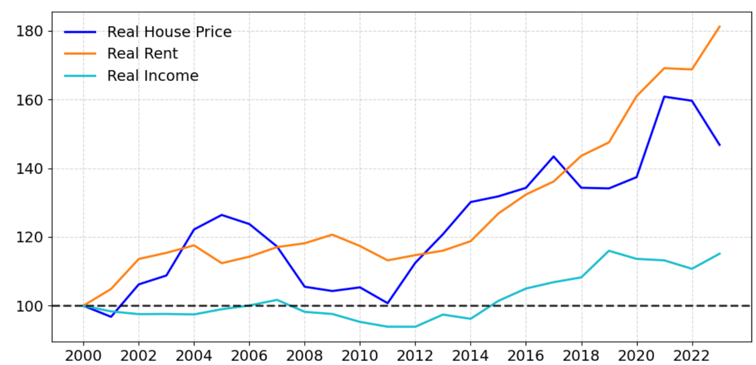

(a) Medians over time

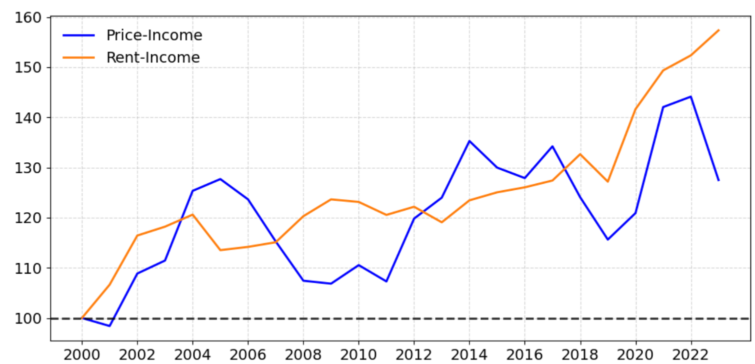

(b) Price and rent to income ratios

Notes: House price is the median sales price of houses sold. Rent is the median asking rent. Income is median household income. Data from *US Census Bureau*, in 2023 US dollars, with 2000 as base year.

We normalize house size[7] to unity, both the rented house before the purchase and the one bought. In effect, this translates into a buyer path where the household buys the house it inhabited as a renter. Any deviation from this assumption in the data is to some extent picked up by the rental rate, q_t and the unit price p_t. We again use the data from figure 2 to pin down the path of these two prices in the buyer cost time-series, given the normalization of house size h.

Regarding initial liquid wealth b, we define it as median holdings of deposits, money market funds, stocks and US government and municipal bonds. Figure 3 shows its time-series data from *Distributional Financial Accounts*. For each year, given the permanent income level z_t^p, we set $b_{0t}$to match the path of liquid-wealth-to-income ratio.

3.4 Mortgage Design

According to the mortgage contract, households pay a fix payment every month, starting one

[7]We interpret house size as general notion of housing services volume. House size and house quality are used interchangeably.

Figure 3: Median liquid wealth

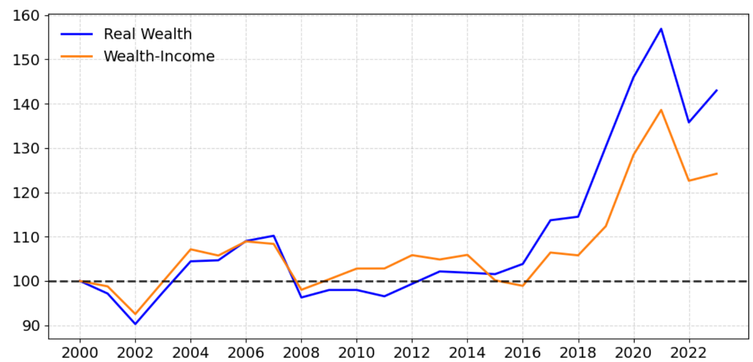

Notes: Liquid wealth includes deposits, money market funds, stocks and US government and municipal bonds. Data from *Distributional Financial Accounts*, in 2023 US dollars, with 2000 as base year.

month after the issuance of the loan, and until the end of the mortgage maturity T_m. In this fixed nominal payment $\tilde{x}$, there is an interest and amortization component which decreases and increases over time, respectively. The whole stream of nominal payments fully amortizes the initial mortgage balance, m^*, at T_m:

$$m^* = \sum_{t=t^*+1}^{T_m} \frac{\tilde{x}}{(1+r_m)^t}$$

We impose the median maturity in the data to set $T_m = 360$, that is, a 30 years mortgage, while the mortgage rate comes from the annual average from *Freddie Mac* data presented in figure 4.

Figure 4: Mortgage interest rate

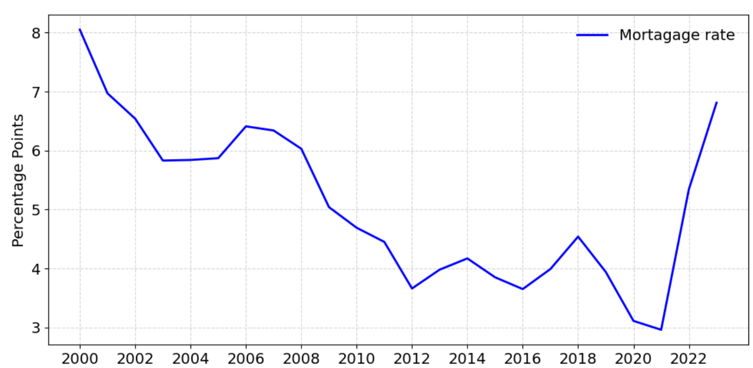

Notes: Average interest rate for 30-year fixed-rate mortgages. Data from *Freddie Mac*.

3.5 Calibration summary

The remaining parameters are set to either standard values or the sample average. However, the LTV cap, which for the US is not easy to pin down, is set to 97.5% up until 2010, and thereafter it goes down to 90%. This is to reflect the tighter macroprudential policies that led banks to be stricter in mortgage issuance after the GFC. We see this as a conservative reduction, for instance Paz-Pardo (2024) assumes LTV limits were lowered to 80% after 2010.

	Description	Value	Source
σ	EIS non-durable consumption	1	Standard
r^b	Risk-free rate (monthly)	0.14%	1.67% (T-Bill year average)
r^m	Mortgage rate		Freddie Mac
θ	LTV cap	$\{0.975, 0.9\}$	Standard
T_m	Mortgage maturity (months)	360	Standard
π	Gross inflation (monthly)	1.0016	2% Fed target
z^p	Permanent income		US Census Bureau
p	Unit house price		US Census Bureau
q	Rental rate		US Census Bureau
b_0	Initial liquid wealth		DFA
δ	House depreciation (monthly)	0.002	Berger et al. (2017)
β	Discount factor (monthly)	0.98	Standard
ρ_z	wrt log(z) process	0.91	Standard
σ_z	wrt log(z) process	0.25	Guvenen et al. (2022)

Table 1: Calibration

4 Headline Buyer Cost

This section presents results for our headline homeownership cost measure, which we refer to as Buyer Cost. The parametrization described in section 3 allows us to compute the Buyer Cost, $V_0^{\text{free-user}} - V_0^{\text{buyer}}$, for each single year. Specifically, this means the calibrated parameters such as prices (p, q, r^m), permanent income z^p, and initial wealth b_0, are fixed at their values in, say, year 2002 and with them in hand$V_0^{\text{free-user}} - V_0^{\text{buyer}}$is computed for that year. Throughout the sample, we set the purchase timing, t^*, to 6 months. Hence we are measuring the welfare cost of buying a house 6 months from now[8].

Figure 5 compares the time-series of Buyer Cost and of the cost underlying HAI — in every figure, we plot the reciprocal HAI^{-1} to preserve the interpretation of cost, a rise in HAI^{-1} reflects an increase in measured costs. Remarkably, the two series track each other closely up until the GFC, after which a wide gap ensues. This has to do with the lower LTV cap introduced by macroprudential policy at the time and the pernicious way in which this is captured by the HAI. The stricter LTV limit enforces a larger downpayment requirement to qualify for the mortgage and demands of households a higher savings effort to buy a house, for any given price. However, since indicators based on mortgage repayment, such as the HAI, are blind to the pre-purchase burden, the larger downpayment exclusively reflects a smaller mortgage loan and thus a lower monthly

[8]Reducing the time-frame for anticipating the purchase would exacerbate the savings effort to meet the downpayment, thus increasing costs.

Figure 5: Buyer Cost vs HAI

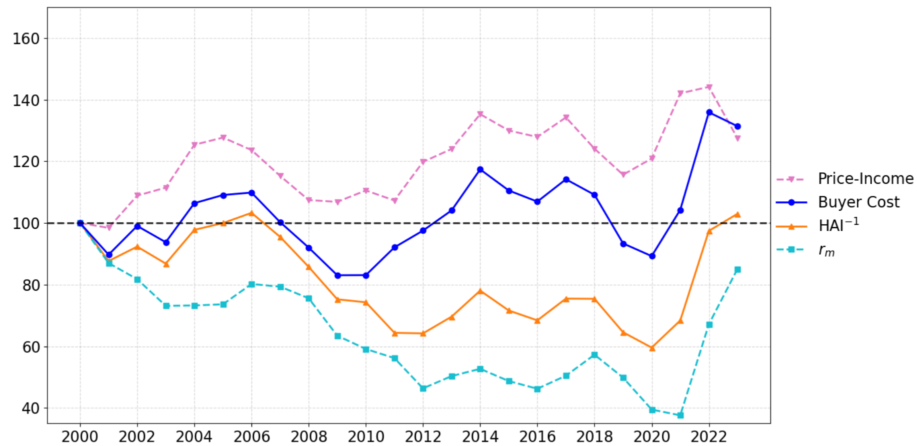

Notes: The solid blue and orange lines correspond to the time-series of Buyer Cost and of the HAI reciprocal. The latter is used to preserve the interpretation of cost: $HAI^{-1}$is virtually a mortgage payment to income ratio. A rise in HAI^{-1}reflects an increase in measured costs. Dashed lines are the time-series of data on price-to-income ratio and mortgage rate, used when computing both cost measures. Base year is 2000 for all.

instalment, which culminates mechanically into lower costs. On the other hand, our Buyer Cost measure factors in this lower payment-to-income effect along side the higher saving effort it implies. In turns out the latter was the dominant effect for the period, leading to an increase in recorded Buyer Cost of about 40%, between 2010 and 2014, whilst the HAI cost grew by no more than 10%.

The evolution Buyer Cost was closer to that of the price-to-income ratio rather than the path of mortgage rates, conversely to the HAI. Yet, convergence in the variation from the start to the end of sample — between the Buyer Cost and price-to-income — is only achieved given the concurrent lower house prices and higher mortgage rates in the last years. Absent the steep increase in mortgage rates between 2022 and 2023, the Buyer Cost would have decreased at par with price-to-income, preventing the observed convergence.

All in all, the main takeaway from figure 5 is that while HAI depicts similar housing costs between 2000 and 2023, our Buyer Cost measure records a more than 30% rise over the period. The major driver of this discrepancy is the larger burden of saving for a downpayment after a tightening of LTV limits, which exacerbated the effect of rising house prices in the second half of the sample.

For a deeper insight into our results, we can dissect each cost measure by singling out each of its time-varying components. For the HAI, those are the house price, income, LTV and mortgage rate variation. Figure 6 shows the time-series of HAI-cost that we would observe if only each of those variables had changed over time. Clearly, the HAI with only mortgage rate variation is the closest to the full-variation HAI, with the most significant detachment occurring between 2003 and 2008, when mortgage rates barely moved while house prices had a clear inverted u-shape. In contrast, figure 7 shows that the Buyer Cost attributes a lower weight to mortgage rate movements and is therefore more sensitive to house price changes. Of course, this measure of cost has other variables at play, its full decomposition is included in appendix A.

Figure 6: HAI decomposition

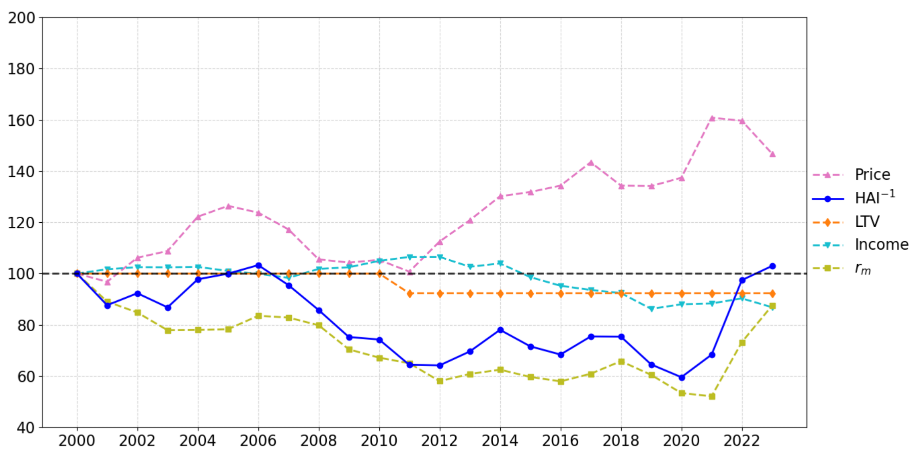

Notes: Solid blue line is the implied HAI cost measure, as in figure 5. Each dashed lines corresponds to a counterfactual HAI measure where only a single item is allowed to vary overtime, while other items are fixed at their initial sample values.

Figure 7: Buyer Cost decomposition

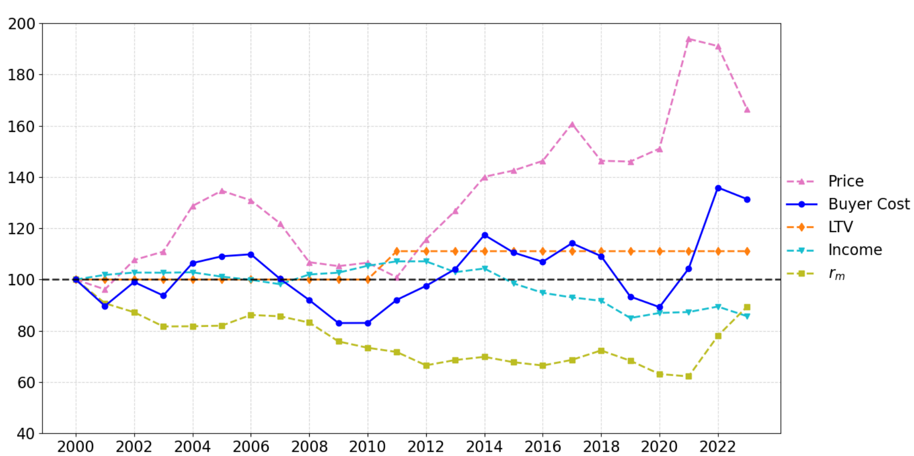

Notes: Solid blue line is the Buyer Cost measure, as in figure 5. Each dashed lines corresponds to a counterfactual Buyer Cost where only a single item is allowed to vary overtime, while other items are fixed at their initial sample values.

For a closer look on how the two indices weigh each component differently, figure 8 overlaps the two when time-variation comes from a single component — it basically overlaps the dashed lines in figure 6 and figure 7. In the top-left panel we can see how the Buyer Cost index translates rising house prices into a larger costs than what HAI does — whilst the effect that higher prices have on higher mortgage instalments is present in both indices, the effect of a larger downpayment is only present in the Buyer Cost. The LTV panel illustrates the key driver of the sizeable discrepancy between the full-variation indices: the lower LTV cap translates into more than 10% increase in costs according to our measure but into a 7.5% reduction according to the HAI. The interaction between larger downpayment, for any given price, and rising house prices is the key driver of rampant economic costs for becoming a homeowner. Finally, in the bottom row, we can see that whilst the Buyer Cost is slightly more responsive to income fluctuations — again, higher income not only lowers the burden of mortgage instalments but also makes it easier to save up when restricted by the LTV limit – and that mortgage rate fluctuations have a greater effect on the HAI index.

Figure 8: Discrepancies between Buyer Cost and HAI

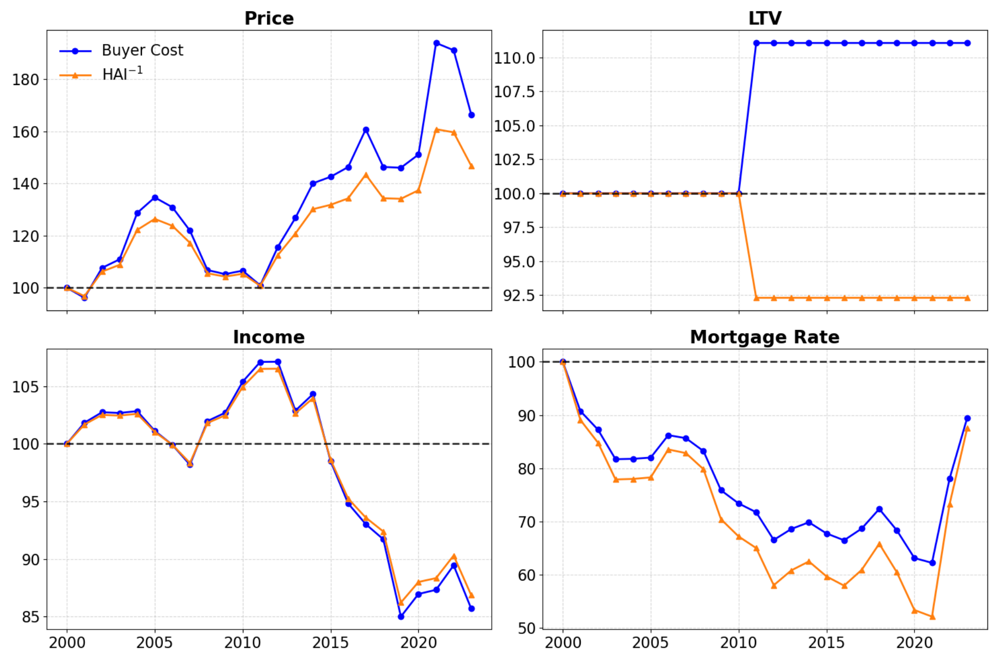

Notes: Each panel depicts the counterfactual Buyer Cost and HAI time-series where only the designated component varies overtime, all other components are fixed at their initial sample values.

5 Buyer Cost over income distribution

5.1 Income quintiles calibration

In this section we now turn to the question of how did costs evolved across the income distribution and of what do we loose out by just focusing on indicators akin to the HAI.

We proceed by calibrating a representative household for each income quintile. Across quintiles the main source of heterogeneity is permanent income, z_i^p, which we pin down for every year in the sample using the average household income within each quintile[9], taken from the *US Census Bureau - Historical Income Tables* and shown in figure 9.

In our simple exercise, house type is exogenous to households choice. This ensures our cost measure is insulated from preferences. However, in real life, different households select into dif-

[9] The top 1% are excluded from the fifth quintile.

ferent house types. So, to account for that, we take the rental rate per unit of housing, q, from our headline calibration at year 2000, and set the house type h_i for each quintile to match the observed quintile average of rent-to-income ratios, qh_i/z_i^p, of that year. Just as with the median household we then fix the house type per quintile constant throughout the sample. For rent-to-income data we use the *Survey of Consumer Finances (SCF)*. The result, in figure 10, is that higher quintiles occupy 'better houses'.

Regarding initial wealth b_{0i}, we proceed accordingly and set it conditional on permanent income, such that initial wealth-to-income ratios b_{0i}/z_i^p match the median of each quintile of liquid financial wealth, as it is defined in section 3.

All other parameters, including rental rates, unit house prices and mortgage rates and LTVs are taken from the previous calibration and thus vary over time but not across quintiles. Crucially, this assumes that rents and unit prices of all house types evolved in a same manner.

Figure 9: Average Income by Quintile

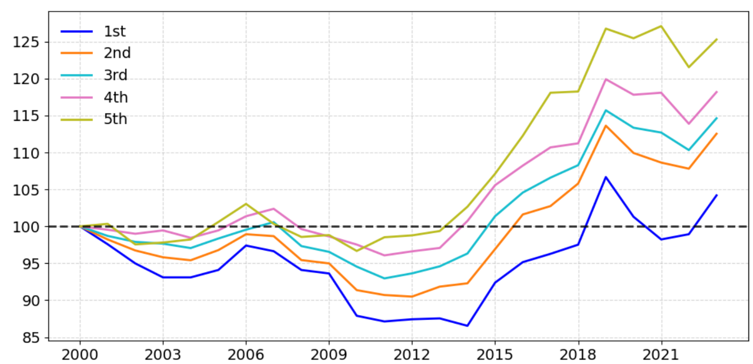

Notes: Average household income by quintile. Data from *US Census Bureau - Historical Income Tables*. In 2023 US dollars with 2000 as base year.

Figure 10: Calibrated house type h_i per income quintile

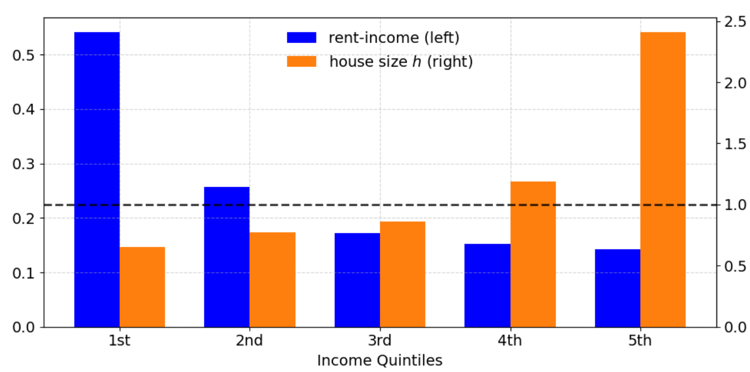

Notes: On left-hand-side axis, average rent-to-income ratio for each quintile in year 2000. Data from *SCF* (imputed for in-between waves). On right-hand-side axis, calibrated house type that matches rent-to-income data, given each quintile's income level and the rental rate q, calibrated in section 3.

Figure 11: Average liquid wealth to income ratio by quintile

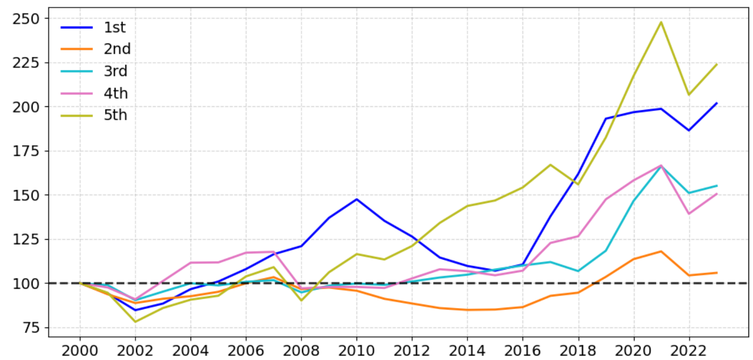

Notes: Liquid wealth includes deposits, money market funds, stocks and US government and municipal bonds. Data from *Distributional Financial Accounts*, in 2023 US dollars, with 2000 as base year.

5.2 Income quintiles Buyer Cost

Firstly, as shown in figure 12, the HAI depicts a small variation of time changes across income quintiles. Over the whole sample the first quintile saw costs increase by more than 10% while the top 20% of earners saw cost reduced by about 5%. The third and fourth quintile track closely the costs for the median household.

Figure 12: HAI across the income distribution

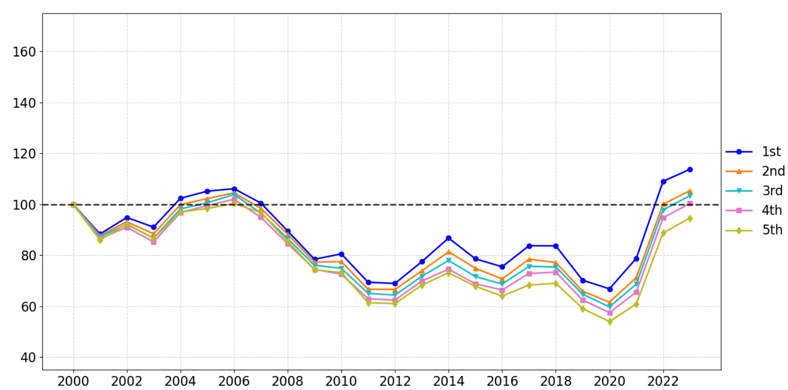

Notes: HAI reciprocal, HAI^{-1} for each income quintile.

As we describe the same quintiles and time period with the Buyer Cost measure, the picture substantially different. First, for those in the lowest quintile becoming homeowner is prohibitively expensive throughout the sample [10]. Second, relative to the headline Buyer Cost, the effect of the pre-purchase burden are exacerbated for the 2nd income quintile and much more muted for the

[10] This means that, with the current calibration, this household cannot sustain a savings profile to afford a downpayment within 6 months without ruling out the possibility of ending up with no cash-on-hand to consume when faced with a negative income shock. This highlights a limitation of our approach. However, we could extract information on the cost for these households by computing the lowest purchase timing that would be affordable. That would be an auxiliary statistic for such cases.

top quintile. This stems not only from the higher income growth for top earners but also from a much higher growth in liquid wealth holdings, specially from 2018 onwards, figure 11. In sum, over the whole sample, while the HAI and Buyer Cost converge for the top earners, but if the Buyer Cost is taken as a benchmark, the HAI significantly understates the rise in costs for the 2nd quintile, a gap of about 50 percentage points.

Figure 13: Buyer Cost across the income distribution

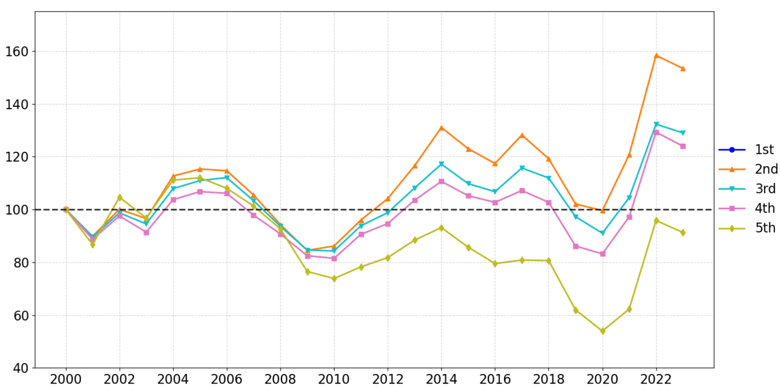

Notes: Buyer Cost for each income quintile. For the 1st quintile, buying the house in a 6-months horizon is prohibitively costly (infinite cost) through out the sample.

Figure 14 and figure 15 single out the 2nd and 5th quintile indices, respectively. These figures emphasize that while the two cost measures diverge on the rate of change for the lower quintiles, they converge for the top earners in the latter part of the sample period. Such pattern results from the significant increase of liquid wealth for top earners, which renders the saving effort to meet the downpayment relatively trivial, and consequently renders pre-purchase effects not so relevant.

Figure 14: Buyer cost vs HAI, 2nd quintile

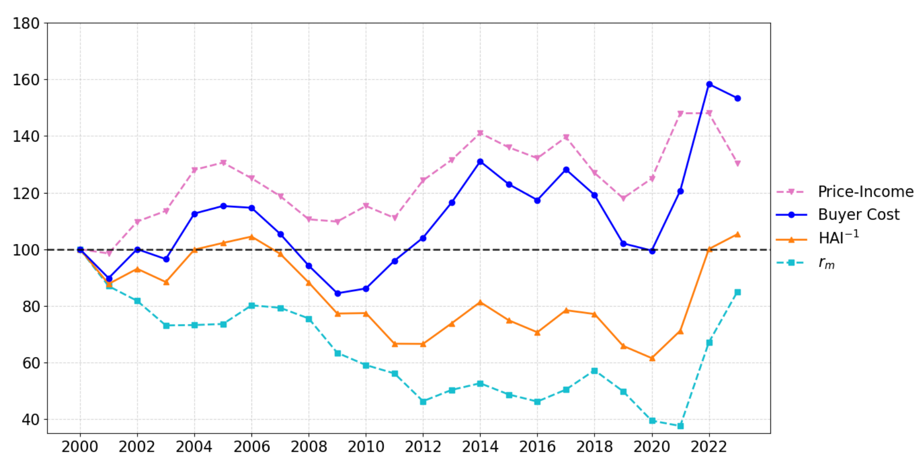

Notes: The solid blue and orange lines correspond to the time-series of Buyer Cost and of the HAI reciprocal. The latter is used to preserve the interpretation of cost: $HAI^{-1}$is virtually a mortgage payment to income ratio. A rise in HAI^{-1}reflects an increase in measured costs. Dashed lines are the time-series of data on price-to-income ratio and mortgage rate, used when computing both cost measures. Base year is 2000 for all.

Figure 15: Buyer cost vs HAI, 5th quintile

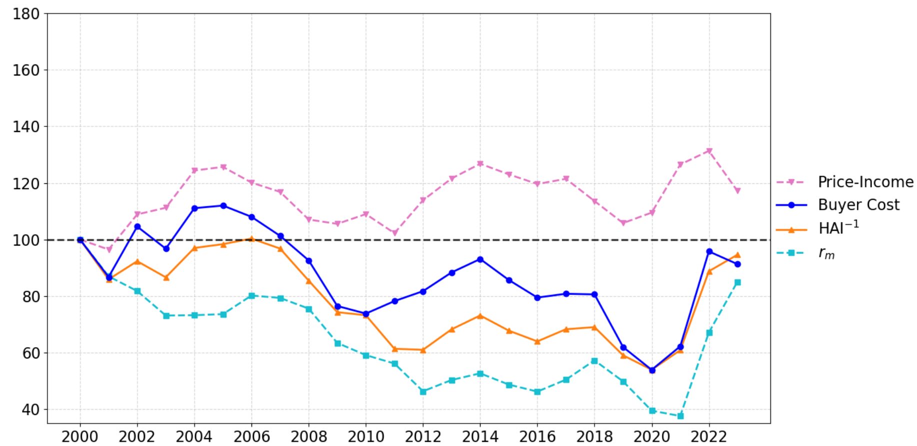

Notes: The solid blue and orange lines correspond to the time-series of Buyer Cost and of the HAI reciprocal. The latter is used to preserve the interpretation of cost: $HAI^{-1}$is virtually a mortgage payment to income ratio. A rise in HAI^{-1}reflects an increase in measured costs. Dashed lines are the time-series of data on price-to-income ratio and mortgage rate, used when computing both cost measures. Base year is 2000 for all.

6 Conclusion

This paper addresses the disconnect between conventional housing affordability metrics and the lived experience of potential homebuyers. While indices like the Housing Affordability Index suggest stable conditions, younger generations face falling homeownership rates and report feeling increasingly priced out. We show this discrepancy arises because existing measures overlook the pre-purchase burden of saving for a downpayment while paying rent and facing income uncertainty.

We develop a microfounded measure of homeownership cost that captures the full intertemporal welfare loss of becoming a homeowner. Applied to US data from 2000 to 2023, our measure reveals that median first-time homebuyers faced a 30% cost increase — not the near-zero change suggested by the HAI. The lowest income quintiles experienced a 60% increase, while only the top 20% saw costs slightly below 2000 levels. Tighter macroprudential policy after the Great Financial Crisis, which raised downpayment requirements, is a key driver—an effect entirely missed by mortgage-payment-based indices.

Our framework provides policymakers with a theoretically grounded yet empirically tractable tool for evaluating how economic conditions and regulations affect homeownership affordability across different household types, reconciling measured costs with perceived barriers to homeownership.

References

Auclert, A., M. Rognlie, and L. Straub (2024). The intertemporal keynesian cross. *Journal of Political Economy 132*(12), 4068–4121.

Battistini, N. and J. Gareis (2024). Housing investment and the user cost of housing in the euro area. *Economic Bulletin Boxes 3*.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2017, 10). House prices and consumer spending. *The Review of Economic Studies 85*(3), 1502–1542.

Biljanovska, N., C. Fu, and D. Igan (2023, 12). Housing affordability: A new dataset, wp/23/247, december 2023.

Bogdon, A. S. and A. Can (1997). Indicators of local housing affordability: Comparative and spatial approaches. *Real Estate Economics 25*(1), 43–80.

Chodorow-Reich, G., E. Glaeser, X. Jaravel, and A. Lipton (2025). Homeownership and the cost of living.

Favilukis, J., P. Mabille, and S. V. Nieuwerburgh (2023, 1). Affordable housing and city welfare. *Review of Economic Studies 90*, 293–330.

Gan, Q. and R. J. Hill (2009, 6). Measuring housing affordability: Looking beyond the median. *Journal of Housing Economics 18*, 115–125.

Glaeser, E. L., J. Gyourko, and R. E. Saks (2005). Why have housing prices gone up? *American Economic Review 95*(2), 329–333.

Guvenen, F., L. Pistaferri, and G. L. Violante (2022). The global repository of income dynamics. https://www.grid-database.org/. Accessed DD.MM.YYYY.

Jorda, O., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). The rate of return on everything, 1870-2015. *The Quarterly Journal of Economics*, 1225–1298.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Paz-Pardo, G. (2024). Homeownership and portfolio choice over the generations. *American Economic Journal: Macroeconomics 16*, 207–237.

Poterba, J. M. (1984). Tax subsidies to owner-occupied housing: an asset-market approach. *The quarterly journal of economics 99*(4), 729–752.

A Full decomposition of Buyer Cost measure

Figure 16: Decomposition of Buyer Cost, median household

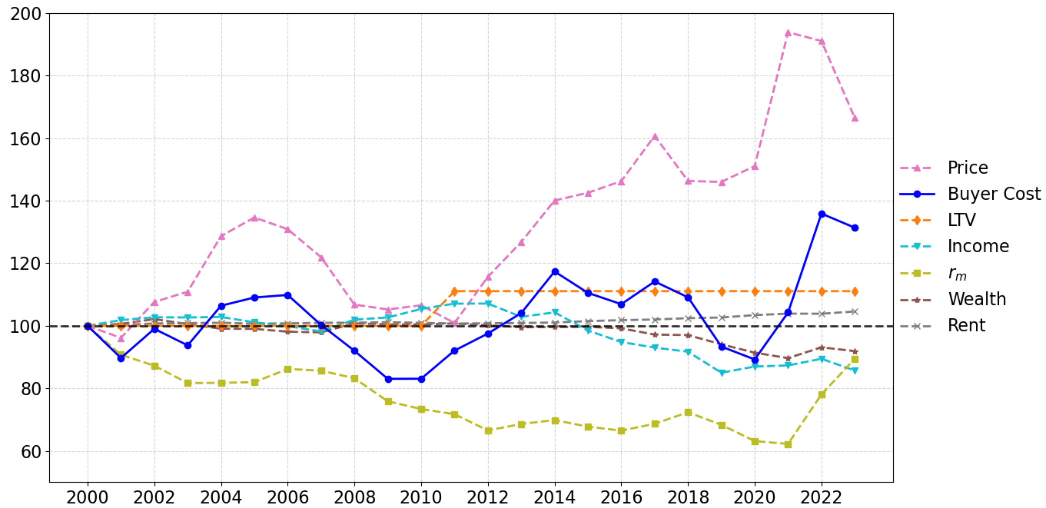

Notes: Solid blue line represents the time-series of Buyer Cost; dashed lines represent a counterfactual Buyer Cost where only a single component varies over time.

Figure 17: Decomposition of Buyer Cost, 2nd quintile

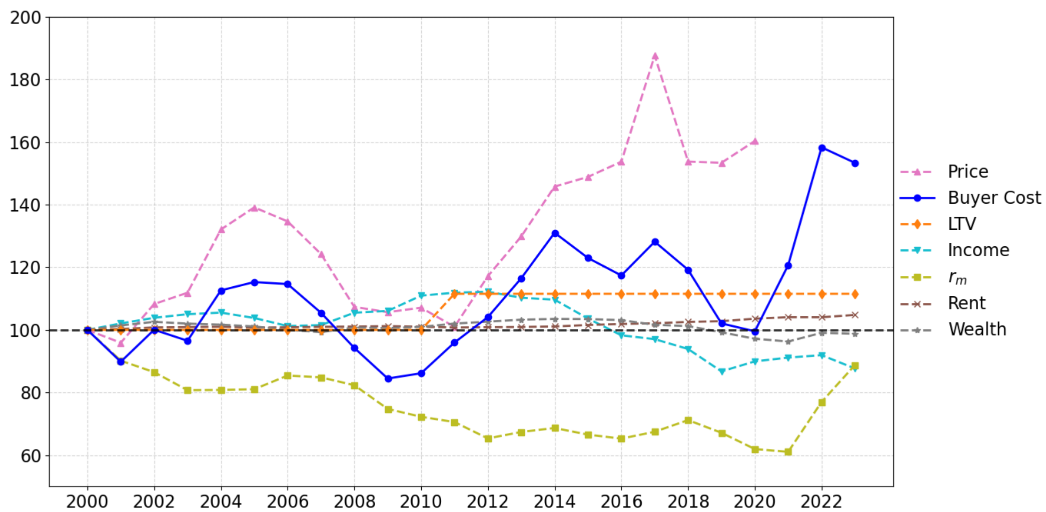

Notes: Solid blue line represents the time-series of Buyer Cost; dashed lines represent a counterfactual Buyer Cost where only a single component varies over time.

Figure 18: Decomposition of Buyer Cost, 3rd quintile

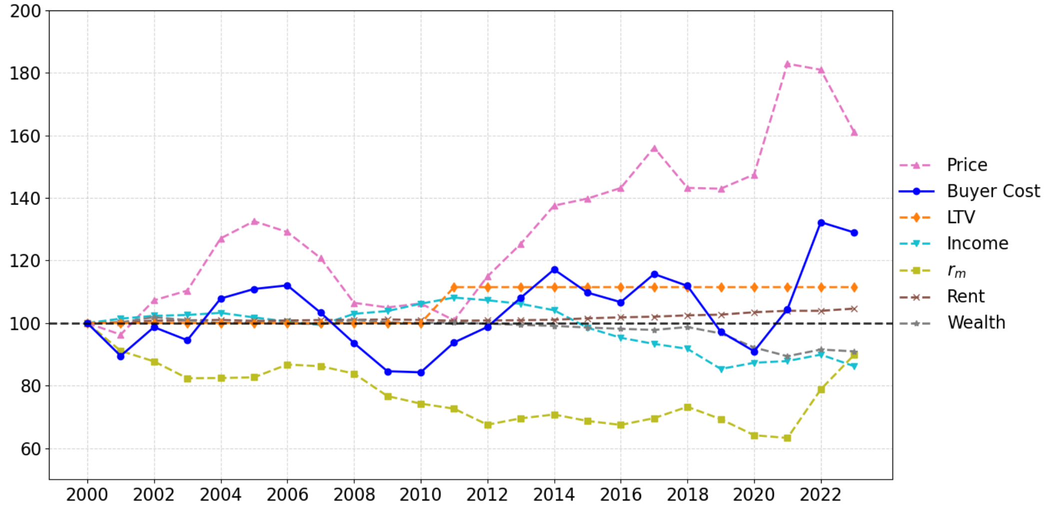

Notes: Solid blue line represents the time-series of Buyer Cost; dashed lines represent a counterfactual Buyer Cost where only a single component varies over time.

Figure 19: Decomposition of Buyer Cost, 4th quintile

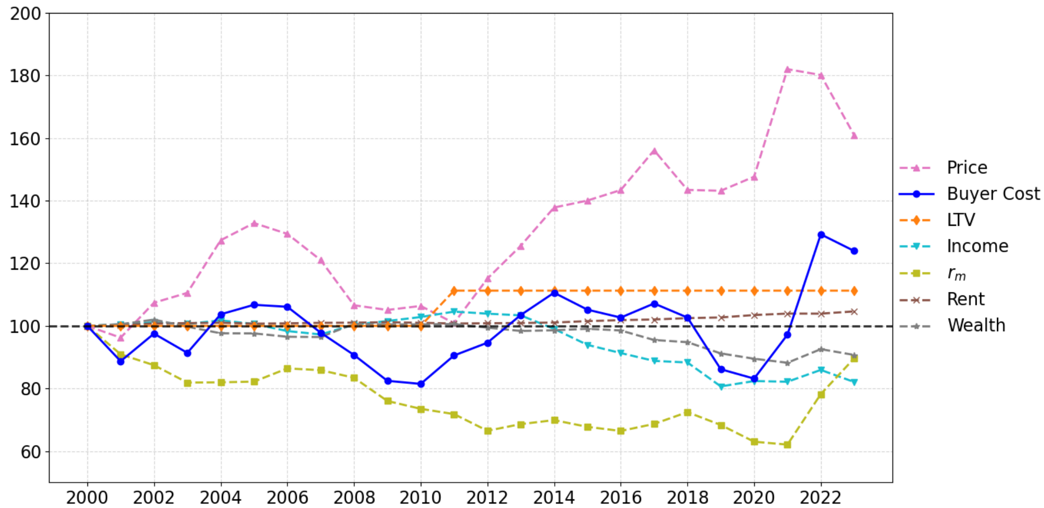

Notes: Solid blue line represents the time-series of Buyer Cost; dashed lines represent a counterfactual Buyer Cost where only a single component varies over time.

Figure 20: Decomposition of Buyer Cost, 5th quintile

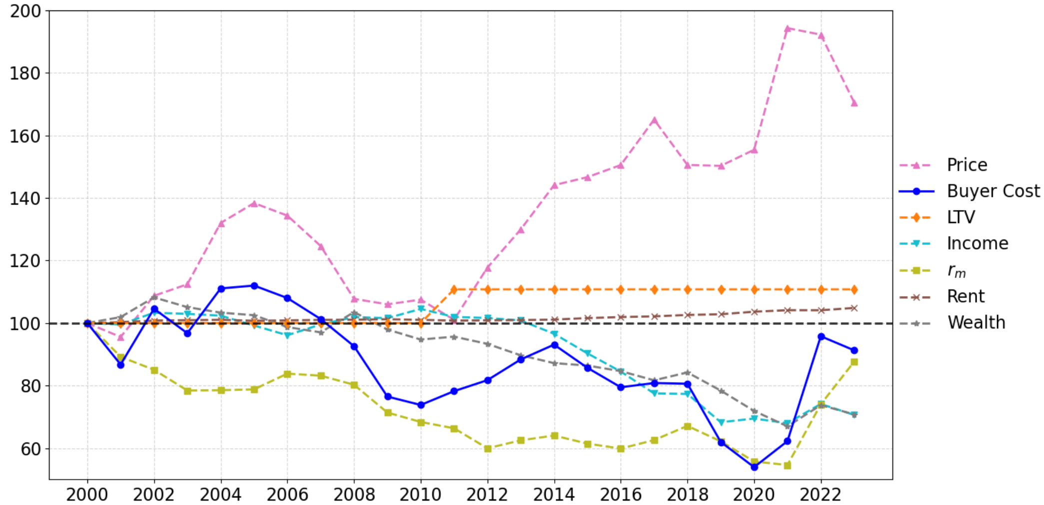

Notes: Solid blue line represents the time-series of Buyer Cost; dashed lines represent a counterfactual Buyer Cost where only a single component varies over time.