BANK OF CANADA
BANQUE DU CANADA

Staff working paper / Document de travail du personnel—2026-43

Last updated: November 25, 2026

Demand-Supply Decomposition of House-Price Inflation and the Impact of Monetary Policy

Edward Booth
McGill University
EdwardBooth47@hotmail.com

Yaz Terajima
Financial Stability Department
Bank of Canada
yterajima@bankofcanada.ca

Bank of Canada staff research is produced independently from the Bank's Governing Council and may support or challenge prevailing views. The views expressed in this paper are solely those of the authors and may differ from official Bank of Canada positions. No responsibility for them should be attributed to the Bank.

DOI: https://doi.org/10.34989/swp-2026-43 | ISSN 1701-9397

© 2026 Bank of Canada

Demand-Supply Decomposition of House-Price Inflation and the Impact of Monetary Policy*

Edward Booth†
McGill University

Yaz Terajima‡
Bank of Canada

October 29, 2026

Abstract

How does monetary policy affect house prices through housing demand and supply? We classify regional Canadian house-price inflation into demand- and supply-driven components using a sign-restriction classification based on the co-movement of prices and quantities, and estimate each component's response to high-frequency monetary-policy surprises around Bank of Canada announcements. In the resale market, the response is concentrated in the demand-classified component, while the supply-classified component shows little systematic response. Supply plays a materially larger role in new housing, particularly for the combined and structure components of the New Housing Price Index. For land, the balance depends on the construction-stage quantity used: the response is demand-dominated under starts or units under construction but often supply-dominated under completions. These results indicate that monetary-policy transmission differs sharply across resale and new-housing markets. Given the relative scale of resale activity, the results suggest that demand is likely to account for a substantial share of monetary-policy transmission to house-price inflation overall.

Keywords: House-price inflation; monetary policy transmission; demand–supply decomposition; resale housing; new housing construction; high-frequency identification

JEL Codes: E52, R31, R21, E31

*We thank Rodrigo Sekkel for sharing the monetary policy shock series, and Jing Yang, Oleksiy Kryvtsov, Nishaad Rao, Brian Peterson, and seminar participants at the Bank of Canada for helpful comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Canada.

†McGill University. Email: EdwardBooth47@hotmail.com.

‡Corresponding author. Bank of Canada. Email: yterajima@bankofcanada.ca.

1 Introduction

House prices and housing costs are highly sensitive to interest rates, and the housing market is widely viewed as a central channel through which monetary policy is transmitted to the broader economy (Iacoviello, 2005; Piazzesi and Schneider, 2016; Williams, 2015). Yet the observed response of house prices to a monetary policy shock is a reduced-form object that conflates two economically distinct forces. A tightening in policy can cool house-price inflation by discouraging buyers—reducing the inflow of demand into the market—or by altering the willingness of sellers to list existing homes and of developers to build new ones. These two channels have very different implications for quantities: a contraction driven by weaker demand reduces prices and transactions together, whereas one driven by tighter supply pushes prices and quantities in opposite directions. Knowing which force dominates is essential for understanding both how monetary policy affects housing and what its transmission implies for construction, transaction volumes, and ultimately the supply of housing.

This paper makes three contributions. First, it adapts the sign-restriction decomposition of Shapiro (2026) to housing markets, where quantities measure resale transactions or construction activity rather than product sales. Second, it applies this decomposition separately to resale housing and new housing, distinguishing the combined, structure, and land components of new house prices and three stages of the construction pipeline. Third, it estimates how each classified component responds to high-frequency Bank of Canada monetary policy surprises. Because the classification is constructed using only price and quantity data, the subsequent local projections ask whether the same monetary policy shock moves demand- and supply-driven house-price inflation differently.

A distinctive feature of housing is that the relevant quantity differs across market segments. In resale markets, the observed quantity is completed sales, and supply-classified movements are interpreted through the search model as reflecting changes in listings or seller participation; in new housing markets, quantities are construction activity and supply-classified movements are interpreted as reflecting developers' pipeline decisions. This dis-

tinction lets us compare transmission across the demand- and supply-classified margins in each segment.

Our main findings point to a pronounced asymmetry in how monetary policy transmits to house-price inflation. In the resale market, the response is concentrated in the demand-driven component: a contractionary monetary policy shock significantly reduces demand-driven house-price inflation, while the supply-driven component shows no statistically significant response. In the new housing market, by contrast, the supply-driven component responds more systematically, and for the price of new structures the supply-driven response is comparable in magnitude to the demand-driven one. For the land component of new housing prices, the balance between the two channels depends on which construction-stage quantity measures supply: land is demand-dominated when supply is measured by housing starts or units under construction, mirroring the resale market, but supply plays a larger and often dominant role once supply is measured by completions. These patterns are broadly robust across the alternative specifications we consider, although the estimated contribution of resale supply is more sensitive to specifications that allow for longer price–quantity dynamics. We document this fully in the robustness analysis. These segment-level findings suggest that demand is likely to play the larger role in overall housing-market transmission, given the relative scale of resale activity and the overwhelmingly demand-driven resale price response; we do not, however, construct a formal aggregate response combining the resale and new-housing segments.

Our results connect several strands of the housing and monetary economics literature. The classification method builds directly on Shapiro (2026), who decomposes U.S. personal consumption expenditure (PCE) inflation into supply- and demand-driven components. In a related decomposition that additionally distinguishes persistent from transitory demand and supply shocks, Sheremirov (2022) attributes a larger quantitative role to supply shocks in the post-pandemic inflation increase, while finding that demand shocks contain a relatively important persistent component. We adapt the same logic to housing, where the mapping

from economic concepts to observable prices and quantities requires reinterpreting "quantity" as transaction volumes in resale markets and construction activity in new housing markets: housing-search models link prices and transaction volumes through market tightness and match formation (Ngai and Tenreyro, 2014), while Anenberg and Ringo (2024) use listings and a search model to decompose short-run sales and price fluctuations into demand and for-sale-supply factors.

Lee (2024) uses a sign-restricted VAR to quantify the contributions of structural housing-supply, housing-demand, credit, and speculative-demand shocks to aggregate U.S. house-price fluctuations. We ask a different question: we classify regional price–quantity innovations into reduced-form demand- and supply-consistent components and estimate how each responds to an independently identified monetary-policy announcement surprise. Our interpretation of this mapping is grounded in the housing search-and-matching literature, in which prices, sales, and liquidity are jointly determined by the inflows of buyers and sellers (Wheaton, 1990; Krainer, 2001; Díaz and Jerez, 2013; Head et al., 2014; Ngai and Sheedy, 2020). This literature micro-founds the sign restriction we use and motivates our separate treatment of the resale and new housing segments.

The interpretation of demand shocks as operating through credit conditions and borrowing costs follows the macro-housing literature emphasizing collateral and credit constraints (Iacoviello, 2005; Justiniano et al., 2019; Kaplan et al., 2020; Greenwald and Guren, 2025). The finding that supply elasticity determines how demand shocks are divided between prices and quantities relates to work documenting how geography and regulation shape housing supply (Saiz, 2010; Gyourko et al., 2008; Glaeser and Gyourko, 2018).

Most closely related on the transmission side, Aastveit and Anundsen (2022) show that the response of house prices to monetary policy across U.S. metropolitan areas depends on local supply elasticities and on the sign of the shock; we differ in decomposing house-price inflation itself into demand- and supply-driven components and in asking how monetary policy acts on each component, rather than relating the aggregate price response to cross-

sectional supply conditions.

Our estimates can also be read alongside the literature on demand-side housing policies—mortgage stress tests and other macroprudential measures. Borrower-based macroprudential tools primarily operate through credit availability and housing demand rather than by directly expanding housing supply (Cerutti et al., 2017). The demand-driven responses we recover provide a monetary-policy benchmark against which the magnitudes of such interventions can be gauged. A related literature on mortgage lock-in emphasises a supply-side channel whereby higher rates reduce the mobility of existing owners holding low fixed-rate debt (Quigley, 1987; Ferreira et al., 2010; Fonseca and Liu, 2024). As we discuss in Section 5, this mechanism *reinforces* rather than contradicts our finding of a muted resale-supply response, since it implies that monetary tightening need not raise listing flow. Finally, our identification of monetary policy shocks relies on the high-frequency approach of Gürkaynak et al. (2005), using the Bank of Canada surprise series of Sekkel et al. (2025), while remaining attentive to the information-effect concerns raised by Bauer and Swanson (2023a).

The remainder of the paper is organised as follows. Section 2 sets out the conceptual framework and identification strategy for distinguishing demand- and supply-driven house-price inflation and for identifying monetary policy shocks. Section 3 constructs the decomposition, describes the data, and presents the resulting demand- and supply-driven inflation series. Section 4 specifies the local projections and reports the estimated effects of monetary policy on each component. Section 5 discusses the economic interpretation of the results, and Section 6 concludes.

2 Conceptual Framework and Identification Strategy

Our empirical strategy proceeds in two steps. First, we decompose house-price inflation into demand- and supply-driven components using a sign-restriction framework adapted from Shapiro (2026). Second, we estimate the dynamic effects of monetary policy shocks on each

component using local projections (Jordà, 2005).

2.1 Identification Strategy

Our empirical strategy rests on two distinct identification problems, which we deliberately keep separate. The first is a *classification* problem: how to split observed house-price inflation, at each point in time and in each region, into movements whose price–quantity co-movement is consistent with demand or supply forces. This step does not uniquely recover structural shocks; it assigns innovations to demand- or supply-driven categories under the maintained supply-and-demand geometry. The second is a *policy-identification* problem: how to identify plausibly predetermined policy-announcement variation and estimate the response of each classified component once it has been constructed. Conflating the two would make it impossible to know whether a given empirical result reflects the way demand and supply are measured or the way monetary policy shocks are identified. We address each problem with a separate, purpose-built identification scheme, and only combine them in the second-stage local projections.

Classifying demand- and supply-driven inflation. We classify observed price and quantity movements using the sign-restriction approach of Shapiro (2026), itself a modern implementation of a classical identification argument in supply-and-demand analysis: under an upward-sloping supply curve and a downward-sloping demand curve, a shock that shifts the demand curve moves price and quantity in the *same* direction (tracing out points along supply), while a shock that shifts the supply curve moves price and quantity in *opposite* directions (tracing out points along demand). This sign pattern is the only restriction imposed; no assumption is made about the source of the underlying shock, its persistence, or its relation to monetary policy, which is what allows the classification to be constructed independently of the second stage described below. In practice, the restriction is applied to the reduced-form residuals of rolling price and quantity regressions rather than to raw

price and quantity changes, so that the classification isolates innovations net of each series' own predictable dynamics (Section 3.1.1). Shapiro (2026) develops this methodology for U.S. PCE categories; we adapt it to housing markets, where the same logic applies once "quantity" is reinterpreted as transaction volumes in resale markets or construction activity in new housing markets.

Because the sign-only rule classifies even very small price and quantity innovations, we impose in the baseline a minimum-magnitude threshold, measured in standard-deviation units, to avoid assigning economically small co-movements to demand or supply; observations that do not meet this threshold are labelled ambiguous rather than forced into either category (Section 3.1.1). This magnitude requirement is our conservative baseline for the housing application. We also report in Appendix A.3 a sign-only, threshold-zero classification in the spirit of Shapiro (2026), who classifies observations by the sign of price and quantity innovations and treats precision filters as robustness checks.

What the classification captures. Because the sign restriction is applied to the *co-movement* of prices and quantities, it classifies any force that shifts the demand or supply curve, including forward-looking behaviour. Expectations are therefore not outside the framework. When lower mortgage rates stimulate purchases because households anticipate future capital gains or easier financing, buyer inflow and prices rise together and the episode is classified as demand-driven; when developers accelerate or postpone building in anticipation of future demand or future rate changes, the induced price–quantity co-movement is classified as supply-driven. Appendix B grounds this interpretation in a search-and-matching framework, where buyers enter based on the forward-looking value of owning and sellers or developers enter based on the option value of waiting. What the decomposition does *not* do is identify the structural *source* of a given shift: a demand-classified movement driven by a contemporaneous rate change and one driven by a revision in expected future prices are both labelled demand-driven. The classification thus measures the *incidence* of shocks on the

demand or supply side of the market rather than their structural origin—the reduced-form view we maintain throughout.

Identifying monetary policy shocks. For the second stage, we require a measure of policy-announcement surprises that is plausibly insulated from contemporaneous reverse causality and macroeconomic news released outside the event window. We use the high-frequency identification approach of Kuttner (2001) and Gürkaynak et al. (2005), which measures monetary policy surprises from the change in interest rate futures and other short-term instruments within a short window around scheduled policy announcements. Because the event window is narrow, changes in these instruments within it are commonly interpreted as being driven primarily by the policy announcement rather than by other macroeconomic news released outside the window. This supports interpreting the estimates as responses to the measured announcement surprise, subject to the central-bank information-effect qualification below.

Specifically, we use the one-quarter-ahead Bank of Canada policy-rate-announcement surprise supplied directly by Sekkel et al. (2025). In their original sample, this series is BAX1, the one-quarter-ahead Bankers' Acceptance Futures (BAX) contract surprise around scheduled Bank of Canada policy-rate announcements. BAX futures are the Canadian analogue of the Fed funds and Eurodollar futures used in the U.S. high-frequency literature: they are short-term interest-rate contracts whose price change in a short window around an announcement isolates the unanticipated component of the policy decision. The published Sekkel et al. (2025) dataset covers events from 1997 through 2023. Following the subsequent market transition from BAX to Canadian Overnight Repo Rate Average (CORRA) futures—as BAX liquidity declined and BAX contracts were converted to CORRA futures in 2024—the authors supplied us with an extension for 2024–25 that continues the same one-quarter-ahead short-rate surprise using three-month CORRA futures following the discontinuation of BAX futures.[1]

[1]Their dataset also separately identifies surprises around monetary policy report press conferences and

We note, following Nakamura and Steinsson (2018); Bauer and Swanson (2023a,b), that high-frequency monetary policy surprises can still be correlated with information the central bank reveals about its own economic outlook (the "Fed information effect"). Our controls (Section 4.1) may absorb part of the macroeconomic information associated with policy announcements, but they do not necessarily purge the surprise of central-bank information effects; we therefore interpret the estimated responses as effects of the measured announcement surprise, which may combine a pure policy component with information revealed by the Bank.

Combining the two steps. The two identification schemes enter the analysis sequentially rather than jointly. Classification of demand- and supply-driven inflation is performed first, using only price and quantity data and without reference to monetary policy; the identified monetary policy shock is introduced only in the second-stage local projections (Equation 10), estimated using the Jordà (2005) local projection method, which traces out impulse responses directly from a sequence of single-horizon regressions rather than through a VAR, avoiding the need to take a stand on the joint dynamics of demand- and supply-driven inflation. This sequencing means that any difference we find between the response of demand-driven and supply-driven inflation to a monetary policy shock cannot mechanically arise from the shock itself, since both components are constructed using an identical classification rule and only subsequently exposed to the same identified shock. A differential response is therefore informative about whether the response is concentrated in the demand- or supply-classified component—interpreted through the housing-search framework as the demand-side versus supply-side margin—through which monetary policy transmits to house prices.

speeches by Governing Council members, but we discard those non-rate communication surprises so that our measure of MP_t corresponds strictly to the policy-rate-announcement surprise.

2.2 Conceptual Framework: Demand and Supply in Housing Markets

Interpreting demand and supply in housing markets is better captured by departing from the standard goods-market framework, in which a frictionless, centralized, perfectly competitive (Walrasian) market clears instantaneously at a single price that equates the quantity demanded and supplied. Housing transactions occur through matching between buyers and available units, and the stock adjusts slowly. We therefore interpret demand as the flow of buyers entering the market and supply as the flow of available units: listings in resale markets and construction activity in new housing markets.

This interpretation differs by segment. In resale markets, quantities are completed transactions, and effective supply reflects listing behaviour rather than the physical housing stock. In new housing markets, quantities are construction-stage measures: starts, units under construction, and completions. Because a new home's price can be determined at any of these stages—at the start, during construction, or at completion—and no single construction-stage measure fully summarises new supply, we pair the NHPI price index with each measure in turn.

Two features of housing supply bear on how the decomposition should be read. First, the housing stock is durable and construction is largely irreversible, so supply is generally thought to respond more to positive than to negative demand shocks; in downturns, quantities may adjust little while prices bear most of the adjustment. Because our classification is built on price–quantity co-movements, this irreversibility bears on the first-stage classification itself—whether an observation is labelled demand- or supply-driven when quantities cannot easily fall—rather than only on the second-stage responses; for example, sluggish downward quantity adjustment would load more of a price movement onto the demand-driven category.[2]

[2]We note, too, that the sign of our high-frequency monetary policy surprise is a two-sided, mean-zero innovation and does not coincide with the direction of actual policy, so splitting the local projections by the sign of the surprise would speak to state-dependence in transmission rather than to the tightening-versus-easing asymmetry that irreversibility implies. A structural treatment of these asymmetries—which would require re-deriving the sign-restriction classification under asymmetric supply rather than re-estimating our

This mapping motivates the sign-restriction classification. Demand-driven movements raise prices and quantities together; supply-driven movements move them in opposite directions. As set out in Section 2.3, a simple search-and-matching model rationalises this restriction under the maintained assumptions of an upward-sloping supply curve and a downward-sloping demand curve, although it does not imply that each classified observation recovers a unique structural shock.

2.3 A Search-and-Matching Rationale

A simple search-and-matching model links the interpretation of demand and supply as market-participant flows to the sign restriction used in the empirical classification. With buyer inflow B_t, listings S_t, market tightness $\theta_t = B_t/S_t$, and a constant-returns matching function, a demand shock raises tightness, prices, and transactions together, while a negative listing-supply shock raises prices but lowers transactions. The same logic applies to new housing once the seller side is reinterpreted as a construction margin. Figure 1 visualizes this logic onto the two market segments we study: Panel A depicts the resale matching market, in which buyer and listing inflows determine transaction prices and units sold, while Panel B depicts the new-housing construction pipeline, in which buyer demand meets the sequence of starts, units under construction, and completions that furnishes the three quantity measures we pair with the NHPI. Appendix B gives the derivation. The model rationalises the restriction under the maintained assumptions of an upward-sloping supply curve and a downward-sloping demand curve; it does not prove that every classified observation corresponds to a unique structural shock.

reduced-form specification by subsample—is beyond the scope of this paper, but we flag it as a caveat when interpreting the pooled responses below.

3 Measuring Demand- and Supply-Driven House-Price Inflation

3.1 Decomposition of House-Price Inflation

3.1.1 Classification of Demand and Supply

The decomposition follows Shapiro (2026), using sign restrictions to classify movements in prices and quantities as either demand- or supply-driven.[3] The methodology assumes an upward-sloping supply curve and a downward-sloping demand curve, so that demand shocks move prices and quantities in the same direction while supply shocks move them in opposite directions.

Classification as demand- or supply-driven relies on the sign and magnitude of reduced-form residuals. Let $p_{i,t}$ denote log prices and $q_{i,t}$ denote log quantities. Because the classification exploits the high-frequency co-movement of prices and quantities, seasonal patterns common to both series could otherwise be misread as genuine demand- or supply-driven innovations. We therefore remove seasonality directly within the first stage: each reduced-form equation includes a full set of calendar-month dummies, so that the residuals used for classification are purged of the common within-year seasonal variation.[4] We estimate, over a rolling window and separately for each region, the reduced-form price and quantity equations

$$q_{i,t} = \sum_{j=1}^{12} \gamma_j^{qp} p_{i,t-j} + \sum_{j=1}^{12} \gamma_j^{qq} q_{i,t-j} + \sum_{m=2}^{12} \delta_m^q D_{m,t} + c + \nu_{i,t}^q, \tag{1}$$

$$p_{i,t} = \sum_{j=1}^{12} \gamma_j^{pp} p_{i,t-j} + \sum_{j=1}^{12} \gamma_j^{pq} q_{i,t-j} + \sum_{m=2}^{12} \delta_m^p D_{m,t} + c + \nu_{i,t}^p, \tag{2}$$

[3] The decomposition identifies the *type* of each price movement—demand- or supply-driven—but not the magnitude of the underlying shock. Moreover, because observed quantities are equilibrium outcomes, the economic content of the supply classification differs across market segments (listing behaviour in the resale market versus construction activity in the new housing market); these distinctions should be borne in mind when interpreting the results.

[4] Including month dummies in the reduced-form equations is equivalent to deseasonalising the price and quantity series with month fixed effects before extracting the innovations. All classification results reported in the paper are based on these seasonally adjusted residuals.

where $D_{m,t}$ are calendar-month dummies (January omitted as the reference month), and obtain the reduced-form residuals. Let W_T denote the rolling window of data ending in period T, where we use the capital letter T for the final (most recent) period of a window to distinguish it from the general period index t running over the observations within the window. Following Shapiro (2026), we retain from each window only the residual of its final period: $\nu^p_{i,T}$ and $\nu^q_{i,T}$ are the price and quantity residuals in period T, the last observation of window W_T, so that the subscript T on ν refers to this single final period. We then roll the window forward one month and re-estimate the regressions, generating one pair of residuals per period. We standardise each final-period residual by the standard deviation of all residuals estimated within the same window to express each innovation in standard-deviation units,

$$\tilde{\nu}^p_{i,T} = \frac{\nu^p_{i,T}}{\sigma^p_{i,W_T}}, \qquad \tilde{\nu}^q_{i,T} = \frac{\nu^q_{i,T}}{\sigma^q_{i,W_T}}, \tag{3}$$

where σ^p_{i,W_T} and σ^q_{i,W_T} denote the standard deviations of the price and quantity residuals over the entire window W_T. The subscripts thus play different roles: the numerator $\nu_{i,T}$ is the innovation in the single final period T, whereas the denominator σ_{i,W_T} summarises the variability of residuals across the whole window W_T used to scale it. The standardised residuals $\tilde{\nu}^p_{i,T}$ and $\tilde{\nu}^q_{i,T}$ measure how large each price and quantity innovation is relative to its own recent variability. Observations are classified as demand-driven, supply-driven, or

ambiguous according to the sign and magnitude of these standardised residuals:

$$1_{i\in D,T} = \begin{cases} 1 & \text{if } \operatorname{sign}(\tilde{\nu}^p_{i,T}) = \operatorname{sign}(\tilde{\nu}^q_{i,T}) \text{ and } \min\left(|\tilde{\nu}^p_{i,T}|, |\tilde{\nu}^q_{i,T}|\right) \geq \tau \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

$$1_{i\in S,T} = \begin{cases} 1 & \text{if } \operatorname{sign}(\tilde{\nu}^p_{i,T}) \neq \operatorname{sign}(\tilde{\nu}^q_{i,T}) \text{ and } \min\left(|\tilde{\nu}^p_{i,T}|, |\tilde{\nu}^q_{i,T}|\right) \geq \tau \\ 0 & \text{otherwise,} \end{cases} \tag{5}$$

$$1_{i\in A,T} = \begin{cases} 1 & \text{if } \min\left(|\tilde{\nu}^p_{i,T}|, |\tilde{\nu}^q_{i,T}|\right) < \tau \\ 0 & \text{otherwise,} \end{cases} \tag{6}$$

where τ is the magnitude threshold below which a movement is treated as ambiguous, set to $\tau = 0.3$ standard deviations in the baseline classification.[5]

For each region, we estimate rolling price and quantity regressions using a 10-year window and 12 monthly lags in the baseline specification. A regional implementation is important because local economic conditions can generate substantial differences in house-price growth (Feng et al., 2023). Estimating the first-stage equations separately by region removes each market's own predictable price and quantity dynamics before the innovations are classified and aggregated.

These choices match the default implementation in the replication code accompanying Shapiro (2026).[6] We adopt these values as a benchmark to preserve comparability with Shapiro (2026) for three reasons. First, the window and lag are tuning parameters of a purely statistical procedure, governing how much historical data is used to estimate the

[5] Shapiro (2026)'s baseline classification imposes no such threshold: it labels an observation purely by the signs of the two reduced-form residuals, so that every period with a determinate sign pattern is assigned to the demand- or supply-driven category, and the magnitude filter that relabels small-residual observations as ambiguous is treated only as a robustness check. We instead adopt that filter—with $\tau = 0.3$—as our baseline, so that only pronounced co-movements are classified. The magnitude filter makes the baseline classification more conservative by leaving small residual pairs unclassified; it does not by itself establish that the remaining observations correspond to structural demand or supply shocks. The sign-only case, which corresponds to Shapiro's baseline, is reported as a robustness check in Appendix A.3.

[6] Alternative windows and lags—including shorter and longer rolling windows and 3-lag and 24-lag variants—are reported below as robustness checks.

reduced-form price and quantity equations and how persistent the estimated dynamics are assumed to be, rather than parameters that encode economically meaningful assumptions specific to housing markets; although housing differs from the consumption categories studied by Shapiro (2026) in its adjustment speed and persistence, these remain tuning choices of a statistical procedure, and we assess their influence directly in the robustness analysis rather than assuming the PCE-based values carry over unchanged. Second, because Shapiro (2026) already subjects this classification procedure to extensive robustness analysis in the original paper—including alternative lag lengths—adopting the same baseline preserves direct comparability between our results and the broader literature using this classification method. Third, and most importantly, we do not treat this choice as beyond scrutiny: Appendix A.2 reports the same local projection coefficients across a grid of windows (3, 5, 10, and 15 years) and lags (3, 12, and 24 months). The 10-year window and the three lag lengths correspond to the benchmark and robustness specifications in Shapiro (2026), while the shorter and longer windows are our own additional sensitivity analysis; the sensitivity of our conclusions to this specification choice is thus reported transparently rather than assumed away.

The paper uses two classifications. Our *baseline* classification requires both standardised residuals to exceed a conservative magnitude threshold of 0.3 standard deviations; all tables and figures described as "baseline" below use it. We contrast this with the more inclusive sign-only, or "threshold-zero," classification—closer to Shapiro (2026)'s main sign-based categorization—which drops the magnitude requirement and classifies a period by the sign of its price and quantity innovations alone. This robustness exercise is reported in Appendix A.3.

3.1.2 Aggregation

To reconstruct Canadian housing price inflation from its regions we define aggregate inflation as:

$$\pi_t^D = \sum_i 1_{i \in D,t}\, \omega_{i,t-1}\, \pi_{i,t}, \tag{7}$$

$$\pi_t^S = \sum_i 1_{i \in S,t}\, \omega_{i,t-1}\, \pi_{i,t}. \tag{8}$$

Here, $\pi_{i,t}$ denotes the change in price in region i, and $\omega_{i,t-1}$ is the lagged weight of that region. Equations 7 and 8 are the demand- and supply-driven components of the aggregate decomposition in Equation 15 of Shapiro (2026), applied here at the regional rather than the product-category level. Because the weight is lagged, the aggregation is Laspeyres-type (base-weighted): current-period classification is combined with weights fixed in period $t-1$. Whereas Shapiro (2026) weights product categories by their expenditure shares, our regional application has no expenditure analogue; we instead weight each region by its share of total transaction value—the dollar volume of housing transactions (price times units sold)—in period $t-1$, so that $\omega_{i,t-1}$ is region i's lagged transaction-value share.[7][8]

Total inflation is:

$$\pi_t = \pi_t^D + \pi_t^S + \pi_t^A. \tag{9}$$

While the classification rule is identical across datasets, the interpretation of quantities differs between resale and new housing markets.

[7]For resale housing, the weight is the region's lagged share of CREA dollar volume (average sales price times units sold). For new housing, where the NHPI is published as an index, we construct the analogous value weight from the regional NHPI level—scaled to the CREA average sales price in December 2016 (Appendix C)—multiplied by the relevant construction quantity. Because the regional NHPI is an index, this December 2016 scaling determines its cross-regional level and therefore affects these weights, although it does not affect regional NHPI inflation, which depends only on within-region changes.

[8]Some weights are subject to reweighting because the presence of zeros in the quantity series must be converted to missing values when we take logs of quantities and prices. To prevent the rolling windows from being of inconsistent sizes across regions, we drop the regions where a missing value is present in the data, and reweight by renormalizing the weights across the valid regions for each period t.

3.2 Data and Housing-Market Measures

3.2.1 Resale Housing: CREA Data

We use monthly regional data from the Canadian Real Estate Association (CREA) for resale housing markets (Canadian Real Estate Association, 2026). These series are aggregated at the level of real estate board regions rather than reported only as a single national figure, which is what allows us to exploit variation across local housing markets.[9] The data begin in 1988 and cover 91 regions.[10] For our study, we use CREA's residential series, which excludes industrial and commercial properties, vacant land, farms, and multi-family properties with five or more units. Because the classification requires observing prices and quantities jointly, the CREA data are particularly well suited to our exercise: unlike publicly available house-price indices—including the NHPI, which reports prices only—the CREA series provide matched average sales prices and unit sales at a monthly frequency for a broad panel of regions since 1988. This granular regional panel, which is distributed through CREA's subscription service, has to date seen relatively limited use in the academic literature. A complete list of data sources, variable definitions, and the regions used in the analysis is provided in Appendix C.

For our study, we use CREA's average sales price series as $p_{i,t}$. Average sales price is constructed as the ratio of total dollar volume to units sold in a given region. $q_{i,t}$ is CREA's total units sold.[11] In this setting, quantities reflect transaction volumes from matching buyers

[9] One limitation of the CREA data is that regions are bound to historical and current real estate board regions. This means the data cannot perfectly map to common Canadian regional descriptors like the CMA.

[10] Constructing the regional panel requires a few aggregation choices—most notably dropping the largest overlapping aggregates in favour of their components and building a synthetic Toronto region—which we detail in Appendix C.

[11] We use the average sales price rather than CREA's MLS Home Price Index (HPI) because it exhibits greater month-to-month variation, which makes monthly price shocks within regions easier to detect. The MLS HPI is a quality-controlled index: within each market, it is built up from fine-grained subareas (neighbourhoods defined by the participating real estate boards), and a hybrid repeat-sales and hedonic regression is used in each subarea to estimate the price of a "benchmark home" whose attributes are typical of homes traded there. These subarea benchmark prices are then aggregated into market, provincial, and national benchmark prices as weighted averages, with weights given by each subarea's share of benchmark-category sales activity over a three-year window, so that the weights themselves change only gradually rather than from month to month (Canadian Real Estate Association, 2023). Because it holds the quality of the home fixed and smooths the composition of transactions, the HPI varies less month to month than the average

and sellers. Demand-classified movements are interpreted as reflecting buyer activity, while supply-classified movements are interpreted through listing behaviour and the availability of homes for sale. This interpretation is closely related to Anenberg and Ringo (2024), who use listings data and an estimated housing-search model to conclude that demand accounts for most short-run variation in U.S. home sales and price growth, while variation in newly offered for-sale supply plays a more limited role.

Figure 2 illustrates our price and quantity series for a selection of regions from the CREA data, both shown on a seasonally adjusted basis. As expected, units and prices are positively correlated over time. The series clearly capture episodes such as the Great Financial Crisis, when prices and units sold both fell, and the pandemic period, when both rose. This co-movement motivates decomposing price changes into demand- and supply-driven components. Seasonal adjustment matters here: in the raw data, prices and quantities share pronounced within-year seasonal patterns that, left in place, would generate spurious same-sign co-movements and be misclassified as genuine demand innovations. The classification removes these common seasonal patterns in the first stage (Section 3.1.1); for the figure, we display seasonally adjusted series so that the plotted co-movement reflects non-seasonal variation.

Figure 3 displays the total units of resales (blue) and completions (yellow), as well as the share that resales account for in the sum of resales and completions (red). This figure illustrates that resale unit sales are substantially larger than new-housing completions, underscoring the importance of the resale market for understanding house-price inflation in the overall market.

sales price. The two series also differ in coverage: the HPI is available only for participating HPI markets, whereas the average sales price is computed from sales reported through all (or nearly all) CREA member MLS® Systems in the monthly national statistics. For our purposes, the HPI is therefore less suitable than the average-price series, because it displays less month-to-month variation in prices and spans a narrower set of markets.

3.2.2 New Housing: NHPI and Construction Activity

For the new housing market, we use Statistics Canada and Canada Mortgage and Housing Corporation (CMHC) data.[12] Prices are measured using the New Housing Price Index (NHPI), which reports separate indexes for the structure, land, and the two components combined. The NHPI is a monthly, current-market, matched-model index of contractors' selling prices for new residential dwellings, holding the specification of a representative home fixed so that the index isolates price changes from changes in quality (Statistics Canada, 2022). Two features of its construction matter for our exercise. First, it covers low-rise new housing—single, semi-detached, and row (town) houses—and excludes high-rise apartment and condominium units. Second, it records the price at which a representative model would sell in the reference month rather than a contract price locked in at an earlier pre-sale stage, so the index reflects current-market pricing and is not tied to any particular point in the construction pipeline.[13] Statistics Canada reports the index separately for the structure ("house"), the land, and the two combined, which is what allows us to decompose new house-price inflation into its structure and land components. The classification results for the structure and land components are reported separately in Appendix D. Quantities are measured using three construction-based variables from CMHC: housing starts, units under construction, and housing completions (Canada Mortgage and Housing Corporation, 2026).

These variables capture different stages of the supply pipeline. Starts reflect investment decisions, units under construction capture ongoing supply, and completions reflect additions to the housing stock. Regions for the NHPI and CMHC series are constructed from Canadian census metropolitan areas (CMAs). Because the NHPI is published as an index, we convert it to dollar (level) space by scaling each regional index to the corresponding CREA average

[12] CMHC is Canada's national housing agency and produces the Starts and Completions Survey used in this analysis. See https://www.cmhc-schl.gc.ca/.

[13] Because the index is a current-market rather than a contract price, pairing it with each construction stage isolates variation in the *quantity* measure rather than in the timing of the price concept; in particular, it is what allows us to interpret the larger supply-classified response under the completions pairing without attributing it mechanically to prices having been fixed one to two years earlier at the start stage (Section 5).

sales price. This requires a CREA price anchor for every region and therefore restricts the panel to regions common to CREA and the CMAs. As with the CREA data, a small number of synthetic regions are constructed to align the two sources. These construction details, together with the full region list, are provided in Appendix C.

In this setting, demand corresponds to demand for newly built homes, while supply reflects developers' construction decisions and constraints. Figure 4 illustrates the underlying price and quantity series for the new housing market across a selection of major regions, plotting units under construction—the baseline quantity measure—against the year-over-year change in the NHPI and its structure and land components, with all series shown on a seasonally adjusted basis; the analogous figures for housing starts and completions are reported in Appendix C. As in the resale market, the construction-quantity series exhibit strong within-year seasonality in the raw data, which would otherwise generate spurious same-sign co-movements; the classification removes these common seasonal patterns in the first stage (Section 3.1.1), and for the figure we display seasonally adjusted series.

3.3 Decomposition of demand and supply

Before turning to the results, we emphasise what this decomposition is and is not. Following the reduced-form approach of Shapiro (2026), each month is *classified* as demand-driven, supply-driven, or ambiguous according to the sign of the co-movement between the price and quantity innovations under the maintained sign restriction (upward-sloping supply, downward-sloping demand). The labels therefore describe classifications of reduced-form innovations rather than realised structural demand or supply shocks, and the exercise does not uniquely recover the underlying structural drivers.

Figure 5 presents the Shapiro-style decomposition of year-over-year house-price inflation into demand-driven, supply-driven, and ambiguous contributions (using the 0.3 standard-deviation threshold) for the resale and new-housing markets. Because classification is performed at the region–month level and then aggregated across regions, each month carries

demand-, supply-, and ambiguous-driven contributions simultaneously rather than a single label; the panels therefore plot the relative size of these three contributions over time. Panel A uses CREA resale data: the demand-driven contribution is large in most months, while the supply-driven contribution is smaller but consistently present. The demand-driven contribution is especially prominent during and after large price movements, such as the global financial crisis and the post-pandemic period, whereas the supply-driven contribution is more evenly distributed over time and becomes relatively larger in the run-up to those episodes. Panels B through D use the combined NHPI price series paired, respectively, with housing starts, units under construction, and completions as the quantity measure. Relative to the resale panel, the new-housing panels show a smaller demand-driven contribution and a larger supply-driven contribution, and this difference is most visible during periods of high new-housing price inflation. The combined new-housing decomposition shown here aggregates the structure and land price components; Appendix D (Figure D1) reports the two components separately under each quantity measure.

The resulting demand- and supply-driven inflation series are monthly. For the resale (CREA) market, they begin in December 1997, once the ten-year rolling window used in the first-stage classification becomes available; for the new housing (NHPI) market they begin earlier, as the underlying price and construction series extend back to the 1980s. In the local projections, however, both series are used only over their overlap with the monetary policy shock series (Section 4.1), so the effective estimation sample runs from December 1997 to December 2025 in both market segments. The two segments nonetheless enter the regressions with a slightly different number of usable months at the shortest horizon (329 in the resale market versus 335 in new housing): the resale classified aggregate has a few more missing early months, because the ten-year first-stage window is applied to CREA data beginning only in 1988 and the region-level dropping and reweighting of missing observations (Appendix C) leaves a handful of early months without a defined aggregate, whereas the new-housing series rest on construction data extending back to the early 1980s. The difference

is small relative to the overall sample and reflects data availability rather than a different intended estimation period.

4 The Effects of Monetary Policy on Demand and Supply

4.1 Local Projection Specification

We estimate:

$$\pi^k_{t+h,\,t-1} = \alpha^k_h + \beta^k_h MP_t + \Gamma^k_h X_t + \varepsilon^k_{t+h}, \tag{10}$$

where $k \in \{D, S\}$ indexes the demand-driven (D) and supply-driven (S) components of house-price inflation defined in Section 3.1.1; we use "Demand" and "Supply" as equivalent labels for these components in the tables and figures below. The dependent variable $\pi^k_{t+h,\,t-1}$ is the *cumulative* house-price inflation of component k between $t-1$ and $t+h$, obtained by compounding the single-period (month-over-month) inflation rates from the impact month through horizon h; the subscript $t-1$ makes explicit that cumulation starts from the pre-impact month, and at $h=0$ the variable reduces to the one-month inflation rate, so β^k_h traces out the cumulative response of component-k inflation to the shock. The control vector X_t contains six lags of the monetary policy shock and six lags of the single-period (month-over-month) house-price inflation rate of component k. It also includes one to six lags of the log of real GDP, the log of the commodity price index, the log of CPI, the Bank Rate, and the five-year rate, as well as month fixed effects; the contemporaneous shock and response are excluded from X_t. In the estimation files, GDP, the commodity price index, and CPI enter as 100 times log levels, while the Bank Rate and five-year rate enter in percentage points. Each horizon-h regression in Equation 10 is estimated by OLS.[14]

[14] Because the horizon-h dependent variable is a cumulative response, the projection residuals in Equation 10 are serially correlated up to order h under overlapping forecast windows. Our baseline specification already includes six lags of the single-period house-price inflation rate—the flow variable whose forward

4.2 Monetary Policy Shocks

Monetary policy shocks are identified using high-frequency surprises around policy announcements, following Kuttner (2001), Gürkaynak et al. (2005), and Shapiro (2026). Specifically, we use the one-quarter-ahead Bank of Canada policy-rate-announcement surprise supplied directly by Sekkel et al. (2025), as described in Section 2.1. In the original Sekkel et al. (2025) sample this is the BAX1 contract surprise, which covers events from 1997 through 2023; in the extended vintage used here, the same one-quarter-ahead short-rate surprise is continued for 2024–25 with three-month CORRA futures, using an extension supplied by the authors following the BAX-to-CORRA market transition. The estimation files use this supplied series directly. The shock variable is divided by 100 before estimation, so it is measured in percentage points: one unit equals a 100-basis-point surprise. We treat these announcement surprises as predetermined within the monthly housing-price equation, while recognising that they may combine a pure policy component with information revealed by the Bank.

Figure 6 plots the resulting monetary policy shock series, MP_t, used in the local projections. The series is available at a monthly frequency from June 1997 to December 2025 and is expressed in percentage points, where a value of one corresponds to a 100-basis-point surprise. The 1997 start reflects the coverage of the underlying high-frequency dataset: the surprises are measured from short-term interest rate futures around Bank of Canada announcements, for which Sekkel et al. (2025) report coverage from January 1997 onward. Because the shock enters the local projections both contemporaneously and with six lags, and its sample is

cumulation defines the dependent variable—in X_t as lag augmentation; this follows the general logic emphasised by Montiel Olea and Plagborg-Møller (2021), although the choice of six lags is specific to our application. As a more conservative alternative, we have also re-estimated Equation 10 with Newey–West heteroskedasticity- and autocorrelation-consistent standard errors, using a bandwidth chosen to accommodate the overlap-induced serial correlation through horizon h (Jordà, 2005; Newey and West, 1987). The point estimates and implied supply shares are unchanged by construction. The HAC inference yields the same significance conclusions as the classical OLS inference except for the resale demand-driven response at the one-year ($h = 12$) horizon, which is significant at the 5% level under OLS but only at the 10% level under HAC; this is not central to the paper's main findings, since the resale demand-driven response is significant at the 1% level at shorter horizons under both covariance estimators. Accordingly, the substantive conclusions do not depend on the choice of covariance estimator.

the shortest of the series involved, it binds the estimation window: the local projections are estimated over the overlapping sample from December 1997 to December 2025. The shocks fluctuate around zero with no systematic trend, consistent with their interpretation as announcement surprises that are predetermined with respect to contemporaneous housing-price movements, and the largest surprises coincide with well-known episodes of monetary policy repricing, such as the global financial crisis and the post-pandemic tightening cycle.

4.3 Impulse Responses of Demand- and Supply-Driven House-Price Inflation

This section reports the local projection estimates of the effect of a monetary policy shock on demand- and supply-driven house-price inflation. Section 4.3.1 covers the resale market, using the CREA-based decomposition, while Section 4.3.2 covers the new housing market, using the NHPI- and CMHC-based decomposition. The full horizon-by-horizon local projection coefficient estimates and standard errors summarised by the impulse-response figures in this section are available from the authors upon request.

4.3.1 Resale Market: CREA

Figure 7 reports the local projection estimates for the monthly cumulative responses of demand- and supply-driven house-price inflation to a monetary policy shock, measured from $t-1$. Panel (a) shows the response of demand-driven house-price inflation, while Panel (b) shows the corresponding response of supply-driven house-price inflation.

The figure highlights a clear asymmetry in the transmission of monetary policy across the two components. The response of demand-driven house-price inflation becomes statistically significant at the 95% confidence level by the fourth month following the shock and remains significant at that level through much of the first year, with month 10 significant only at the 90% level, whereas the response of supply-driven house-price inflation is statistically insignificant at every horizon across the full three-year window of the analysis. The effect

is also economically large: a one-percentage-point monetary policy shock lowers cumulative demand-driven house-price inflation by roughly 10 percentage points over the six-to-twenty-four-month horizon.[15]

4.3.2 New Housing Market: NHPI

Baseline window–lag specification. Figures 8–10 report the baseline new-housing responses for the combined, structure, and land NHPI components, using units under construction as the quantity measure. For all three baseline NHPI components, the demand-driven response becomes significantly negative over economically relevant horizons. As in the resale market, monetary policy therefore transmits to new house-price inflation through the demand channel.[16] The new-housing results are in fact closest to the resale market for the land component: for land, the supply-driven response is statistically insignificant at every horizon—just as the resale supply-driven response is—so that the same demand-only asymmetry documented for resale prices carries over directly to land (Figure 10). This pattern for land also holds when supply is measured by housing starts (Figure 13).

The combined and structure components are where the two markets diverge. For these components, the supply-driven response is sizable and frequently significant—in contrast to the resale market, where it is statistically indistinguishable from zero at every horizon—so supply is an active margin through which new house-price inflation adjusts. Moreover, the prominence of this supply channel is sensitive to how the quantity margin is measured. For the structure component, the supply-driven response is significant under both alternative quantity measures: under housing starts (Figure 11) the supply-driven response is as large as its demand-driven counterpart, and under completions (Figure 12) both the demand- and supply-driven components react significantly. For land, the supply-driven response—weak

[15] The underlying horizon-by-horizon coefficients, β_h^k in Equation 10 for $k \in \{D, S\}$, confirm the asymmetry seen in the figure: the demand-driven coefficient peaks at $h = 12$ (-9.248, significant at the 5% level), whereas the supply-driven coefficients are an order of magnitude smaller and never significant at conventional levels. The full set of estimates is available from the authors upon request.

[16] The full horizon-by-horizon coefficients and standard errors underlying these figures are available from the authors upon request.

under starts and units under construction—likewise becomes larger when supply is measured by completions (Figure 14).

Robustness over the window–lag specification. As in the resale market, we assess the sensitivity of these NHPI results to the choice of first-stage lag length by holding the window fixed at the baseline's 10 years and re-estimating the local projections at 3- and 24-month lags, for each of the combined, structure, and land components. Figures 15 and 16 report this comparison for the combined structure-and-land component, Figures 17 and 18 for the structure component, and Figures 19 and 20 for the land component, in each case relative to the baseline (10-year window, 12-month lag) results already shown in Figures 8–10. Consistent with Tables A1 and A2, the supply-driven response is more prominent for the NHPI components than for the resale market across all three lag choices, and its relative importance continues to vary more with the choice of quantity measure and NHPI component than with the lag length itself, particularly for the land component. This supports the conclusion that, unlike in the resale market, supply plays a materially larger role in new housing and that this pattern is not an artifact of the baseline lag choice.

5 Discussion: The Demand and Supply Channels of Monetary Policy

The central question of this section is how much housing supply matters for house-price inflation and, correspondingly, through which channel monetary policy is transmitted to house prices. We answer this question using the impulse responses of demand- and supply-driven house-price inflation estimated in Section 4.3.1 and Section 4.3.2. The main evidence comes from the contrast in the point estimates and confidence bands across the two classified components. When demand-driven house-price inflation responds significantly to a monetary policy shock while the supply-driven response is statistically indistinguishable from zero, the

pattern is consistent with monetary policy moving house-price inflation through the demand channel and with supply playing, at most, a limited role. When the supply-driven response is itself significant, supply is an active margin through which house-price inflation adjusts, and the demand channel is no longer the whole story. Before turning to this evidence, we set out what the demand- and supply-driven components represent in each market segment.

5.1 Interpretation of Demand and Supply

The interpretation of demand- and supply-driven inflation differs across market segments.

Resale markets. In the resale-market interpretation, demand corresponds to the buyer-entry margin, while supply corresponds to listings or seller participation. The empirical classification itself uses transaction prices and completed sales rather than direct measures of buyer entry or listings. Demand-classified inflation reflects periods in which increased buyer activity is consistent with higher prices and transaction volumes. Supply-classified inflation reflects periods in which the co-movement is consistent with tighter effective supply—fewer listings or reduced seller participation—raising prices while reducing sales. In the search framework, a monetary policy shock then feeds into the demand-classified component through buyer entry and borrowing conditions, and into the supply-classified component through sellers' listing decisions.

New housing markets. In new housing markets, demand reflects demand for newly constructed homes, while supply reflects construction activity and capacity constraints. Demand-classified inflation corresponds to periods where increased demand drives both prices and construction activity (starts, units under construction, or completions). Supply-classified inflation reflects constraints on construction or delays in supply expansion, leading to higher prices but weaker construction activity. Here a monetary policy shock can operate on both margins: on demand for newly built homes and on the pace at which the housing stock expands.

5.2 A Demand Channel in Both Markets and an Active Supply Channel in New Construction: Evidence from the Impulse Responses

The impulse responses trace how a monetary policy shock is transmitted to each component of house-price inflation, and—following the logic set out at the start of this section—the contrast between a significant demand-driven response and an insignificant supply-driven one is informative about whether the response is concentrated in the demand- or supply-classified component in each segment.

This exercise speaks directly to a long-standing presumption in the housing and monetary economics literature. Macro-housing models and empirical studies emphasise collateral constraints, mortgage borrowing conditions, and credit supply as important drivers of housing demand and house prices (Iacoviello, 2005; Justiniano et al., 2019; Kaplan et al., 2020; Greenwald and Guren, 2025), and the transmission literature has related the aggregate price response to local supply elasticities (Aastveit and Anundsen, 2022). Yet the presumption that policy acts on prices primarily by moving housing *demand*, rather than by shifting effective *supply*, has rarely been tested directly: the two forces are typically not separated in the price response itself. By classifying each period's price movement as demand- or supply-driven and then estimating the response of each component, we provide evidence on this question—and we do so even though our classification uses observed transaction volumes in the resale market (interpreting the supply-classified component through a listing-flow mechanism) and starts, units under construction, and completions in the new housing market, rather than latent structural shocks. The results below support the literature's presumption in the resale market, where the demand channel is strong and the supply channel weak, while qualifying it for new construction.

Resale market. In the resale market, the baseline impulse responses display the significant-demand/insignificant-supply pattern expected under a demand channel. The baseline results

therefore indicate that the resale response is concentrated in the demand-classified component rather than the supply-classified component; in the search framework, this pattern is consistent with an adjustment operating mainly through buyer demand and borrowing conditions rather than through listings. A natural reason for this muted supply response is that resale supply is a *stock-reallocation* margin rather than a production margin: listings are drawn from the existing stock, and the decision to list is governed largely by idiosyncratic and life-cycle considerations—household mobility, downsizing, and equity or lock-in positions (Genesove and Mayer, 2001; Ngai and Sheedy, 2020)—that need not respond to monetary-policy shocks in the same way or at the same speed as buyer demand. If anything, the mortgage lock-in channel reinforces this interpretation: when current mortgage rates rise relative to owners' locked-in rates, existing owners become less likely to move (Quigley, 1987; Ferreira et al., 2010; Fonseca and Liu, 2024), and to the extent that moves involve selling the existing home, this mechanism is consistent with weaker listing supply rather than an increase in listings following a tightening. A positive, contractionary policy-rate surprise therefore need not raise resale supply, and the small, statistically insignificant supply-driven response we estimate is consistent with—rather than contradicted by—this literature.

This does not mean that resale supply is fixed or that it merely mirrors demand: because our CREA data are regional, an owner who sells in one region in order to relocate to another adds to local listings without a matching local buyer, so effective supply is a genuine and independently varying flow. The point is that our estimates provide little evidence of a systematic response in the supply-classified component at the horizons considered, so the identified supply-driven variation is small. Construction, by contrast, is a production margin financed on interest-rate-sensitive terms, which is consistent with the larger supply-driven response we find in the new-housing market.

Beyond its role in assessing the demand channel, the resale market is quantitatively large: resale unit sales are substantially larger than new-housing completions over the period shown in Figure 3. This relative scale makes the resale market central to understanding house-price

inflation in the overall market.

New housing market. The new housing market only partly departs from this clean demand-channel pattern, and where it does, the departure is confined to the structure-intensive margins. For the land component, the new-housing market in fact mirrors the resale market: the supply-driven response is small and statistically insignificant when supply is measured by housing starts or units under construction, so monetary policy transmits to land-price inflation through demand alone, exactly as in resale. The structure and combined components are where the supply channel becomes active: there the supply-driven response is materially larger and often significant, consistent with monetary policy also acting on supply-side margins such as construction activity and the pace of supply expansion. The land component joins this group only when quantity is measured by completions; under that pairing, as discussed below, its supply-classified response becomes larger.

A plausible interpretation of the completions result rests on the timing of the construction pipeline. Starts and units under construction are forward-looking stages that developers can accelerate or postpone in response to current demand and financing conditions, so at these stages quantities tend to co-move with prices in the manner of a demand shock. Completions, by contrast, largely reflect projects initiated in earlier periods and therefore adjust less readily to contemporaneous demand conditions. Pairing current-market prices with completions may consequently place greater weight on pipeline constraints, construction delays, carrying costs, and lagged project decisions than pairings based on starts or units under construction. This interpretation is suggestive rather than structural. It is especially relevant for land, whose effective supply can be constrained by geography and regulation (Saiz, 2010; Glaeser and Gyourko, 2018): the larger supply-classified land response under completions is consistent with pipeline constraints becoming more visible when current deliveries are relatively predetermined, in line with the large land supply shares we find for completions in Table A2. The decomposition does not identify which particular constraint is responsible.[17]

[17]Pre-selling does not threaten this interpretation. Although large condominium projects typically pre-

Taken together, the pattern of point estimates and statistical precision delivers a clear answer to the section's central question. Supply matters little for resale house-price inflation, where monetary policy transmits primarily through demand, but it matters materially in the new housing market, where the construction margin makes the supply-driven response an economically and statistically meaningful part of the transmission. The relative scale of the two segments helps put these results in perspective. Resale units represent roughly 70% of the sum of resale transactions and new-housing completions in Figure 3, and the resale response is concentrated in the demand-classified component. This suggests that demand is likely to account for a substantial share of monetary-policy transmission to house-price inflation overall, even though supply plays a materially larger role within new housing. This inference is suggestive rather than a formal aggregation. Resale transactions and completions are different flow concepts, and the NHPI covers single-detached, semi-detached and row homes rather than apartment buildings, whereas the CMHC construction quantities cover a broader range of dwelling types.

We can also quantify *how much* of each response is attributable to supply rather than simply whether supply is significant. At the two-year horizon, the supply share—the magnitude of the supply-driven response relative to the combined magnitude of the demand- and supply-driven responses—is small in the resale market (21% at the baseline specification, and below 30% in most window–lag combinations), confirming that resale transmission is demand-led. In the new housing market, the supply share is substantially higher: it clusters around 30–55% for the combined and structure components, where demand and supply are both classified as distinguishable from zero under the descriptive 32% criterion of Appendix A.1 in essentially every specification, and it is strongly stage-dependent for land:

sell a substantial share of units before securing construction financing—which could in principle fix prices well before completion—the NHPI covers low-rise new housing (single, semi-detached, and row houses) and excludes high-rise apartment and condominium units, and it is a current-market, matched-model index that records the price at which a representative model would sell in the reference month rather than a contract price locked in at an earlier pre-sale date (Statistics Canada, 2022). The CMHC quantity measures do span a broader set of dwelling types than the NHPI, so the price and quantity series are not perfectly aligned in scope; the direction of any resulting measurement error is ambiguous, and the completions results should be interpreted with this limitation in mind.

the supply share is generally lower under starts and units under construction than under completions, reaching single digits in some under-construction specifications while exceeding 50% in most completions specifications.[18]

6 Conclusion

The response of house prices to monetary policy blends two economically distinct forces—shifts in demand, which move prices and quantities together, and shifts in supply, which move them apart—that carry very different implications for construction, transaction volumes, and the housing stock. This paper separates the two using the sign-restriction classification of Shapiro (2026) together with local projections. The classification labels each period's price–quantity co-movement using only price and quantity data, while the local projections trace the dynamic response of each classified component to high-frequency-identified monetary policy shocks. Because the classification is constructed before and independently of the identified shock, the contrast between the demand and supply responses is informative about the channel through which policy transmits rather than an artefact of how the shock is measured.

We implemented this approach on two segments of the Canadian housing market. For the resale market, we used monthly average sales prices and unit sales from the Canadian Real Estate Association (CREA). In this setting, quantities are completed transactions, and supply-classified movements are interpreted through listing or seller-participation behaviour. For the new housing market, we used the New Housing Price Index (NHPI)—separately for its combined, structure, and land components—together with construction-activity measures from CMHC: housing starts, units under construction, and completions. Monetary policy shocks were identified from high-frequency surprises around Bank of Canada policy-rate

[18] Appendix A.1 reports the full window–lag grid together with a descriptive D/S/B summary of each response. That summary is a one-standard-error description evaluated at a weak, 32% cutoff rather than conventional hypothesis-test inference, which we state explicitly because the D/S/B labelling—though not the magnitude shares, which do not depend on this cutoff—depends on it.

announcements.

Our central finding is a pronounced asymmetry that differs across housing-market segments. In the resale market, monetary policy transmits to house-price inflation primarily through the demand-driven component: a contractionary shock significantly reduces demand-driven inflation, while the supply-driven response is small and statistically insignificant at every horizon in the baseline specification. In the new housing market, the supply-driven component responds more systematically, and for the price of new structures it is comparable in magnitude to the demand-driven response, while the land component is demand-dominated when supply is measured by starts or units under construction, mirroring the resale market, but often supply-dominated when supply is measured by completions. These patterns are stable across rolling windows at the baseline lag length and under the sign-only threshold-zero classification, while the resale supply share is more sensitive under the longest 24-month first-stage lag.

Taken together, the results support a long-standing but rarely tested presumption in the housing and monetary literature: that monetary policy cools resale house-price inflation primarily by acting on demand. At the same time, the new-housing evidence shows that this presumption is not universal—where prices reflect the cost of newly built structures, the supply side is an active margin of transmission, consistent with monetary policy affecting construction activity. This distinction matters for interpreting how policy shapes not only house prices but also the flow of new housing supply. Given the relative scale of resale activity, its demand-dominated response is likely to account for a substantial share of monetary-policy transmission to house-price inflation overall, although the paper does not construct a formal aggregate response across market segments. For policy, it implies that interest-rate changes affect the large resale segment mainly through the demand-classified component—and hence prices relatively quickly—whereas in new housing both demand- and supply-classified components respond, with construction-stage quantities introducing a slower pipeline through which monetary policy can affect additions to the housing stock; the

two segments therefore differ in both the speed of transmission and the balance of demand- and supply-classified responses.

A few limitations qualify these findings and point to future work. Our evidence is specific to Canada and to the sample period, and the estimated resale supply contribution is sensitive to specifications that allow for longer price–quantity dynamics. The analysis is also national in scope and treats the resale and new-housing segments separately rather than modelling their interaction. Extending the decomposition to regional data, to a longer or cross-country sample, and to a structural setting that links the two segments are natural directions for future work.

References

Aastveit, K. A. and A. K. Anundsen (2022). Asymmetric effects of monetary policy in regional housing markets. *American Economic Journal: Macroeconomics 14*(4), 499–529.

Anenberg, E. and D. Ringo (2024). Volatility in home sales and prices: Supply or demand? *Journal of Urban Economics 139*, 103610.

Bauer, M. D. and E. T. Swanson (2023a). An alternative explanation for the "Fed information effect". *American Economic Review 113*(3), 664–700.

Bauer, M. D. and E. T. Swanson (2023b). A reassessment of monetary policy surprises and high-frequency identification. In *NBER Macroeconomics Annual*, Volume 37, pp. 87–155. University of Chicago Press.

Canada Mortgage and Housing Corporation (2026). Starts and completions survey: Starts, completions and units under construction by geography. Monthly data. Accessed June 2026.

Canadian Real Estate Association (2023). MLS home price index methodology. Methodology document.

Canadian Real Estate Association (2026). Canadian MLS residential market statistics: Average price, unit sales and dollar volume by real estate board. Monthly subscription data. Proprietary data obtained under subscription from CREA; accessed June 2026.

Cerutti, E., S. Claessens, and L. Laeven (2017). The use and effectiveness of macroprudential policies: New evidence. *Journal of Financial Stability 28*, 203–224.

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

Feng, X., N. Jaimovich, K. Rao, S. J. Terry, and N. Vincent (2023). Location, location, location: Manufacturing and house price growth. *The Economic Journal 133*(653), 2055–2067.

Ferreira, F., J. Gyourko, and J. Tracy (2010). Housing busts and household mobility. *Journal of Urban Economics 68*(1), 34–45.

Fonseca, J. and L. Liu (2024). Mortgage lock-in, mobility, and labor reallocation. *Journal of Finance 79*(6), 3729–3772.

Genesove, D. and C. Mayer (2001). Loss aversion and seller behavior: Evidence from the housing market. *Quarterly Journal of Economics 116*(4), 1233–1260.

Glaeser, E. L. and J. Gyourko (2018). The economic implications of housing supply. *Journal of Economic Perspectives 32*(1), 3–30.

Greenwald, D. L. and A. M. Guren (2025). Do credit conditions move house prices? *American Economic Review 115*(10), 3559–3596.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking 1*(1), 55–93.

Gyourko, J., A. Saiz, and A. Summers (2008). A new measure of the local regulatory environment for housing markets. *Urban Studies 45*(3), 693–729.

Head, A., H. Lloyd-Ellis, and H. Sun (2014). Search, liquidity, and the dynamics of house prices and construction. *American Economic Review 104*(4), 1172–1210.

Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review 95*(3), 739–764.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review 95*(1), 161–182.

Justiniano, A., G. Primiceri, and A. Tambalotti (2019). Credit supply and the housing boom. *Journal of Political Economy 127*(3), 1317–1350.

Kaplan, G., K. Mitman, and G. L. Violante (2020). The housing boom and bust: Model meets evidence. *Journal of Political Economy 128*(9), 3285–3345.

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

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics 47*(3), 523–544.

Lee, J. (2024). What factors drive house prices in the USA? Sign restricted VAR approach. *Empirical Economics 66*(6), 2533–2556.

Montiel Olea, J. L. and M. Plagborg-Møller (2021). Local projection inference is simpler and more robust than you think. *Econometrica 89*(4), 1789–1823.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: The information effect. *Quarterly Journal of Economics 133*(3), 1283–1330.

Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*(3), 703–708.

Ngai, L. R. and K. D. Sheedy (2020). The decision to move house and aggregate housing-market dynamics. *Journal of the European Economic Association 18*(5), 2487–2531.

Ngai, L. R. and S. Tenreyro (2014). Hot and cold seasons in the housing market. *American Economic Review 104*(12), 3991–4026.

Ortalo-Magné, F. and S. Rady (2006). Housing market dynamics: On the contribution of income shocks and credit constraints. *Review of Economic Studies 73*(2), 459–485.

Piazzesi, M. and M. Schneider (2016). Housing and macroeconomics. In J. B. Taylor and H. Uhlig (Eds.), *Handbook of Macroeconomics*, Volume 2, pp. 1547–1640. Elsevier.

Quigley, J. M. (1987). Interest rate variations, mortgage prepayments and household mobility. *Review of Economics and Statistics 69*(4), 636–643.

Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics 125*(3), 1253–1296.

Sekkel, R., H. Stern, and X. Zhang (2025). Money talks: How foreign and domestic monetary policy communications move financial markets. Staff Working Paper 2025-33, Bank of Canada.

Shapiro, A. H. (2026). Decomposing supply- and demand-driven inflation. *Journal of Money, Credit and Banking 58*(2), 365–388.

Sheremirov, V. (2022). Are the demand and supply channels of inflation persistent? Evidence from a novel decomposition of PCE inflation. Current policy perspectives, Federal Reserve Bank of Boston.

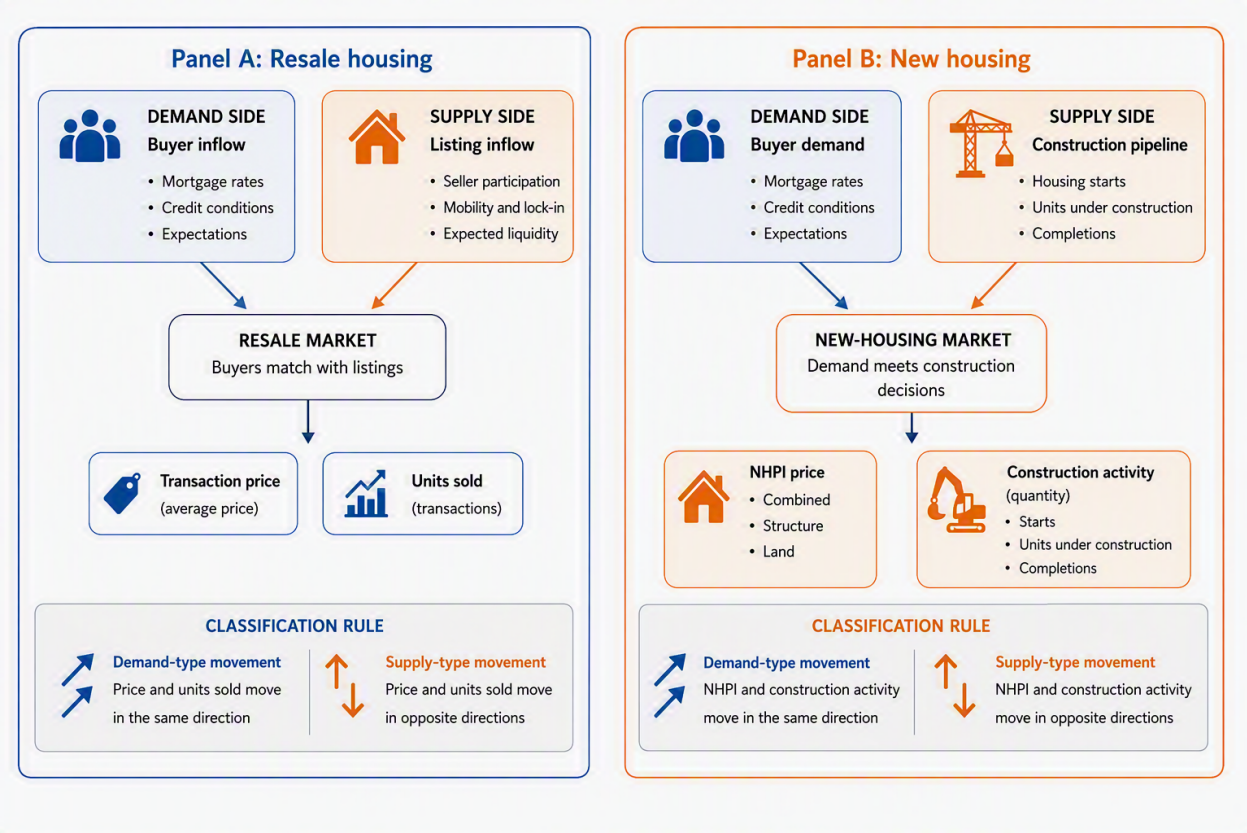
Statistics Canada (2022). Technical guide for the new housing price index. Prices Analytical Series Catalogue no. 62F0014M2022003, Statistics Canada, Ottawa. Available at https://www150.statcan.gc.ca/n1/pub/62f0014m/62f0014m2022003-eng.htm.

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

Williams, J. C. (2015). Measuring monetary policy's effect on house prices. FRBSF Economic Letter 2015-28, Federal Reserve Bank of San Francisco.

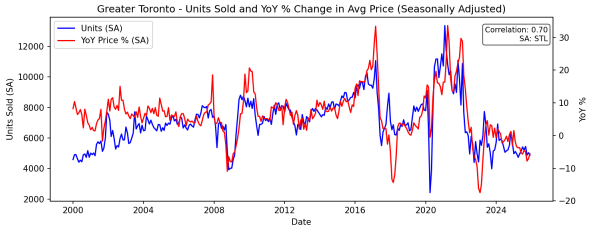
Figures and Tables

Figure 1: Demand, supply, and observed prices and quantities in housing markets

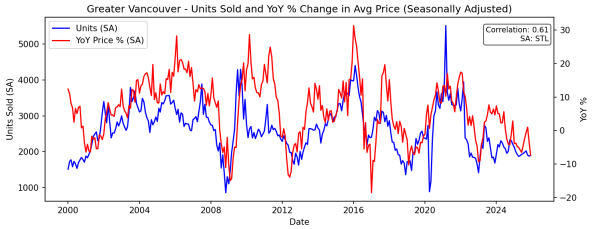


Note: The figure maps the reduced-form classification onto the two market segments studied in the paper; the underlying rationale is developed in Section 2.3 and derived in Appendix B. In the resale market (Panel A), buyer inflows and listing inflows interact through a matching process that determines transaction prices and units sold, so that opposite-sign price–sales innovations are interpreted as reflecting the listing or seller-participation margin. In the new-housing market (Panel B), buyer demand interacts with the construction pipeline—housing starts, units under construction, and completions—that supplies the three quantity measures we pair in turn with the NHPI, so that supply is interpreted at successive stages of construction. In both segments, under the maintained assumptions of upward-sloping supply and downward-sloping demand, same-sign price–quantity innovations are classified as demand-driven and opposite-sign innovations as supply-driven. The classification does not uniquely identify structural shocks.

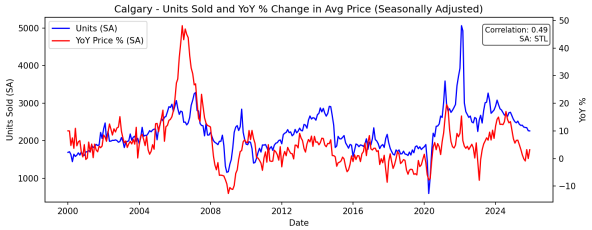
Figure 2: Price and quantity series for a selection of CREA regions.



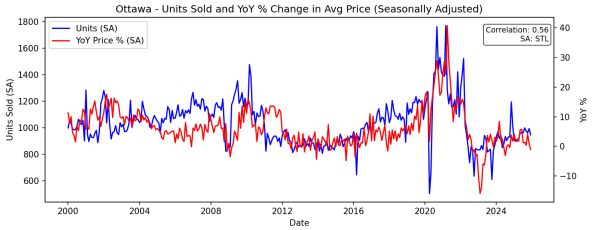
(a) Greater Toronto



(b) Greater Vancouver



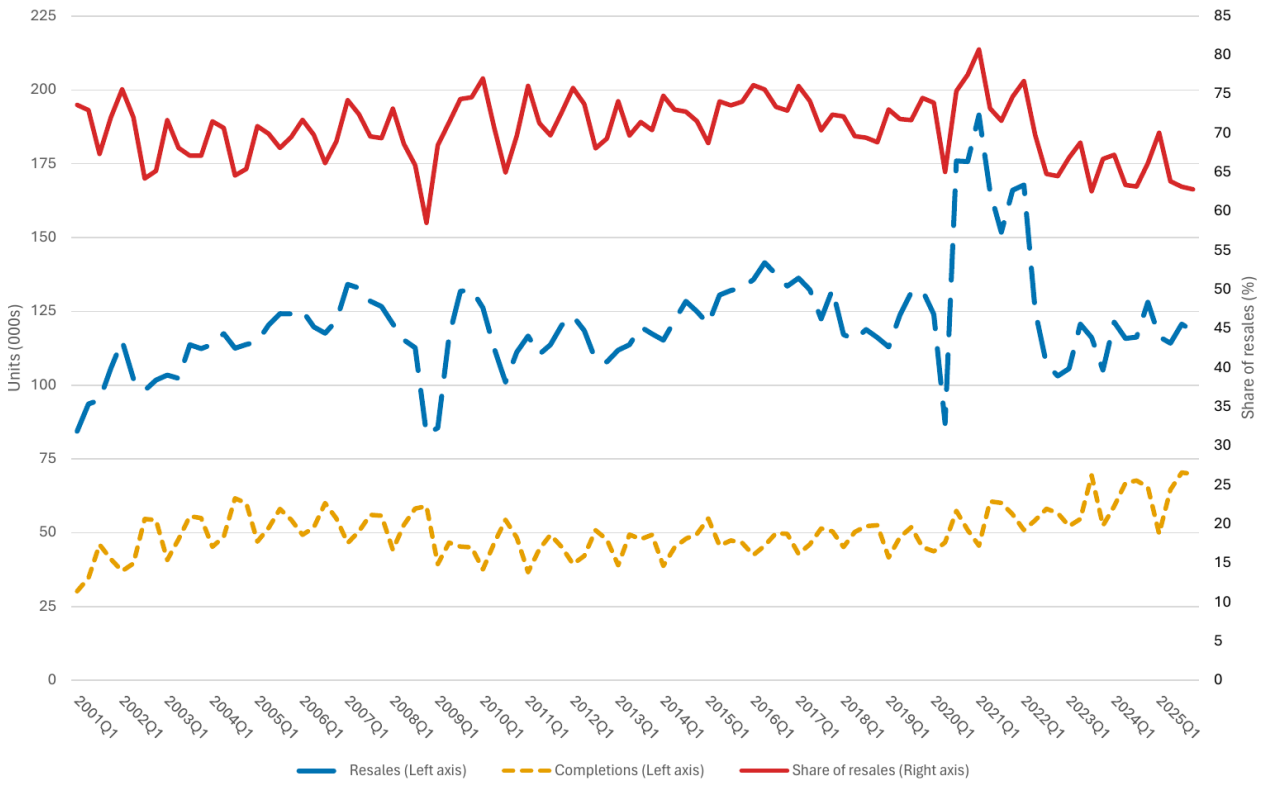
(c) Calgary



(d) Ottawa

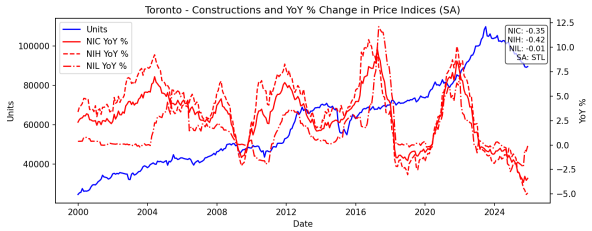
Note: Each panel plots, for the indicated region, the number of residential units sold (left axis, in units) against the year-over-year percentage change in the average resale price (right axis). Both series are seasonally adjusted, and the price series is expressed as the year-over-year change of the seasonally adjusted price level. The figure reported in the upper-right corner of each panel is the correlation between the two seasonally adjusted series. Data are monthly and are drawn from CREA, retrieved via Haver Analytics; authors' calculations.

Figure 3: The Relative Scale of Resale Transactions and New-Housing Completions

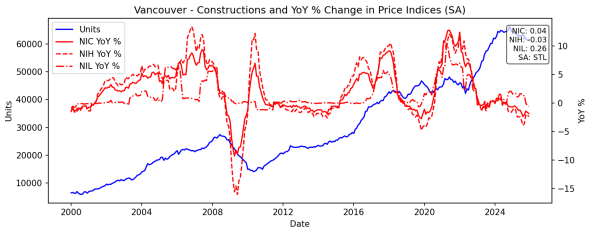


Note: Sources: CREA and CMHC, retrieved via Haver Analytics; authors' calculations. Quarterly data, 2001Q1–2025Q4. Resales (left axis) are Canadian residential unit sales (CREA, series CACERUSQ@CREA); completions (left axis) are total housing completions (CMHC). Share of resales (right axis) is the ratio of resale units to the sum of resale units and completions.

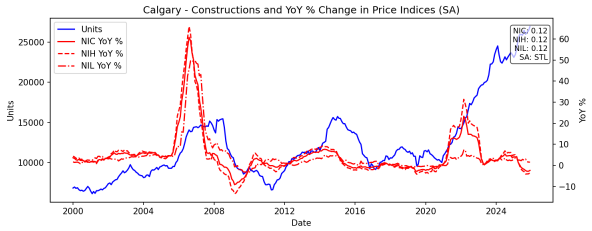
Figure 4: Price and quantity series for new housing across a selection of regions: units under construction (NHPI/CMHC).



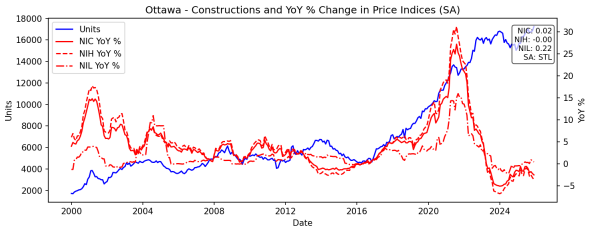
(a) Toronto, units under construction



(b) Vancouver, units under construction



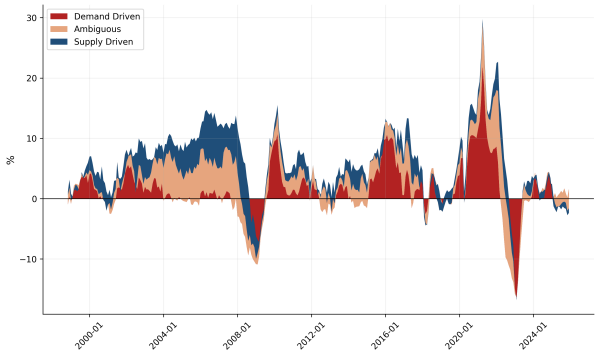
(c) Calgary, units under construction



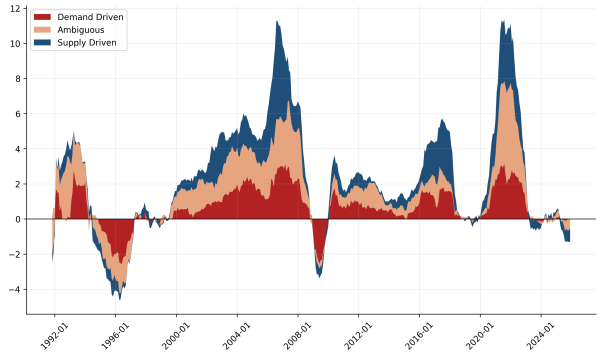
(d) Ottawa, units under construction

Note: Each panel plots, for the indicated region, units under construction (left axis, in units) against the year-over-year percentage change in the New Housing Price Index and its components (right axis): the combined structure-and-land index (NIC), the structure component (NIH), and the land component (NIL). All series are seasonally adjusted, and each price series is expressed as the year-over-year change of the seasonally adjusted price level. The figures reported in the upper-right corner of each panel are the correlations between the seasonally adjusted quantity series and each of the three seasonally adjusted NHPI components (NIC, NIH, and NIL). Price data are from the Statistics Canada NHPI and quantity data are from CMHC, retrieved via Haver Analytics; authors' calculations. The analogous figures for housing starts and completions are reported in Appendix C (Figures C1 and C2).

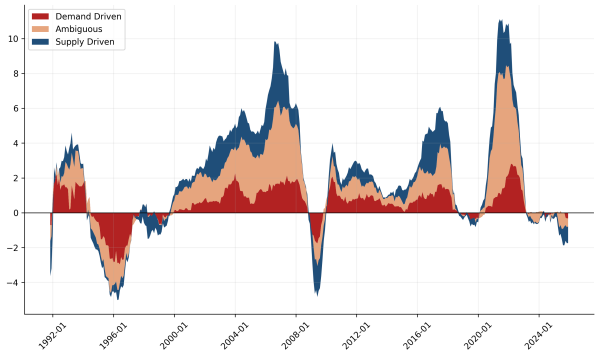
Figure 5: Shapiro-style decomposition of house-price inflation into demand-driven, supply-driven, and ambiguous components: resale (CREA) and new housing (NHPI/CMHC) markets.



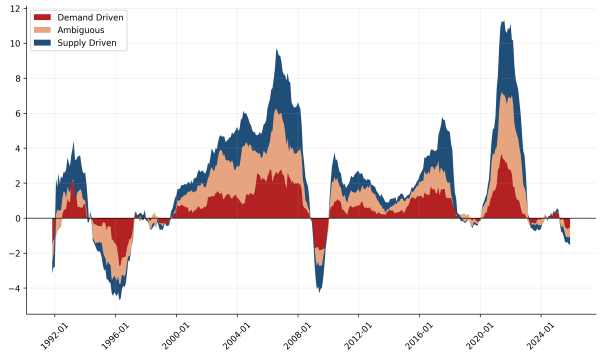
(a) Resale market



(b) New housing market, housing starts



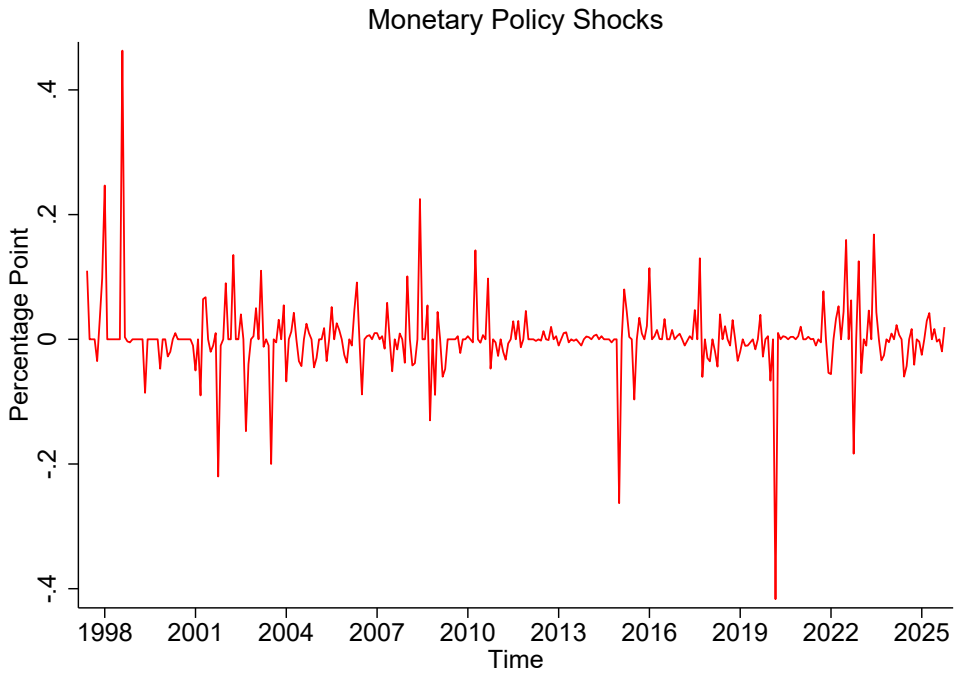
(c) New housing market, units under construction



(d) New housing market, completions

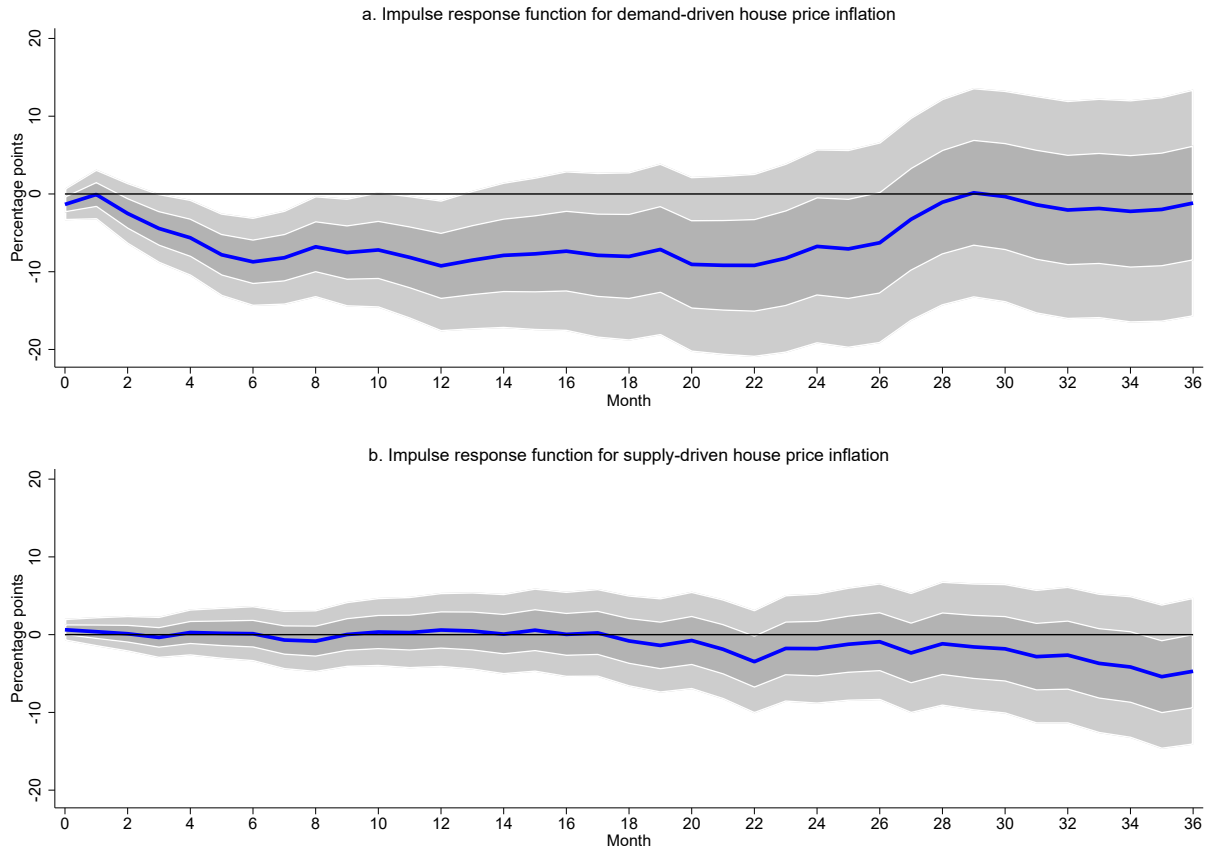
Note: Each panel decomposes year-over-year house-price inflation into demand-driven, supply-driven, and ambiguous components following Shapiro (2026). Panel A uses CREA resale data (average price and units sold). Panels B through D use the NHPI Land and Housing price series paired, respectively, with housing starts, units under construction, and completions as the quantity measure. Classification is based on the co-movement of the price and quantity innovations: an observation is labelled *demand-driven* when the two innovations have the same sign and *supply-driven* when they have opposite signs. To avoid classifying very small co-movements in the baseline specification, both innovations must exceed a minimum magnitude of 0.3 standard deviations in absolute value; if either falls below this threshold, the observation is labelled *ambiguous* regardless of sign. The 0.3 threshold therefore partitions each panel into three regions: demand, supply, and ambiguous.

Figure 6: Monetary policy shock series



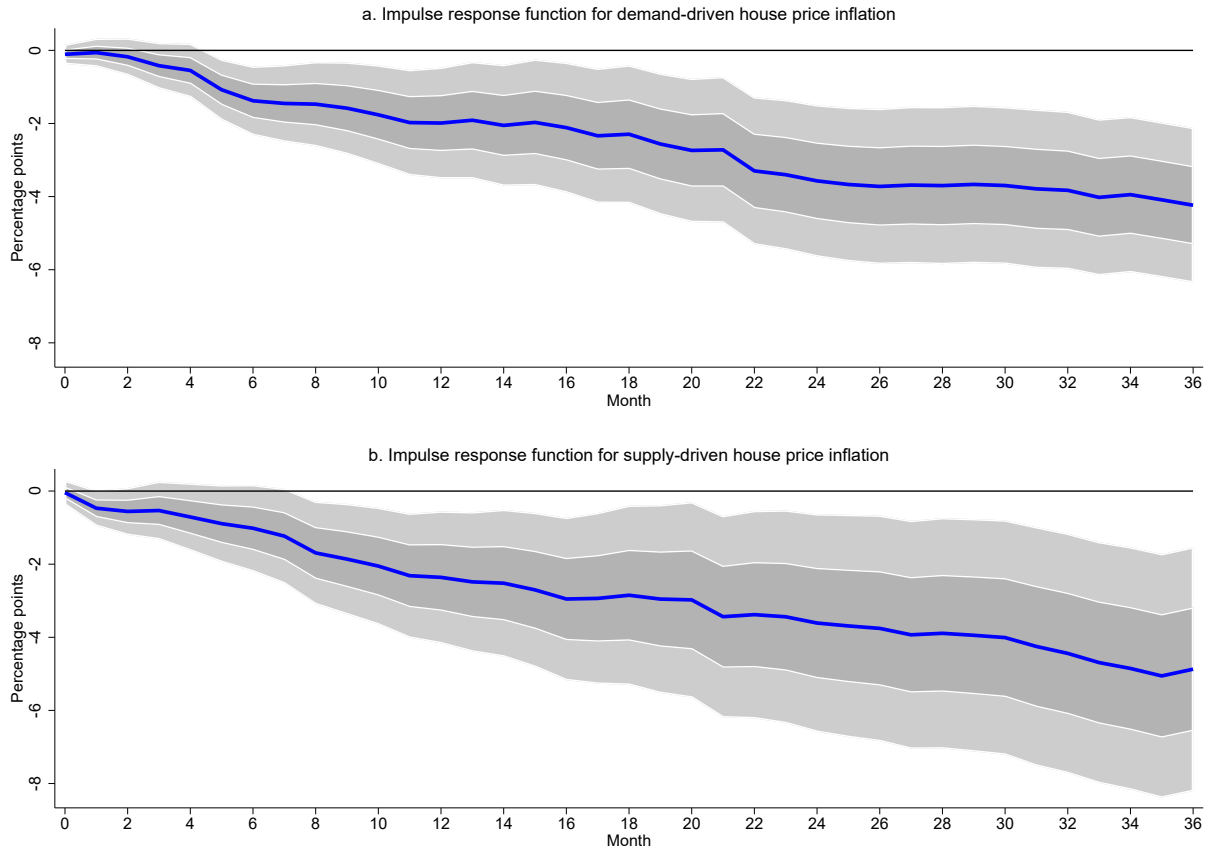
Note: The figure plots the monetary policy shock series MP_t used in the local projections of Section 4.1. Shocks are the one-quarter-ahead Bank of Canada rate-announcement surprises supplied directly by Sekkel et al. (2025): BAX1 in the original 1997–2023 sample and, for 2024–25, an author-supplied extension that continues the same surprise using three-month CORRA futures. The shock is expressed in percentage points, so a value of one equals a 100-basis-point surprise.

Figure 7: Impulse responses of demand- and supply-driven house-price inflation to a monetary policy shock: resale markets.



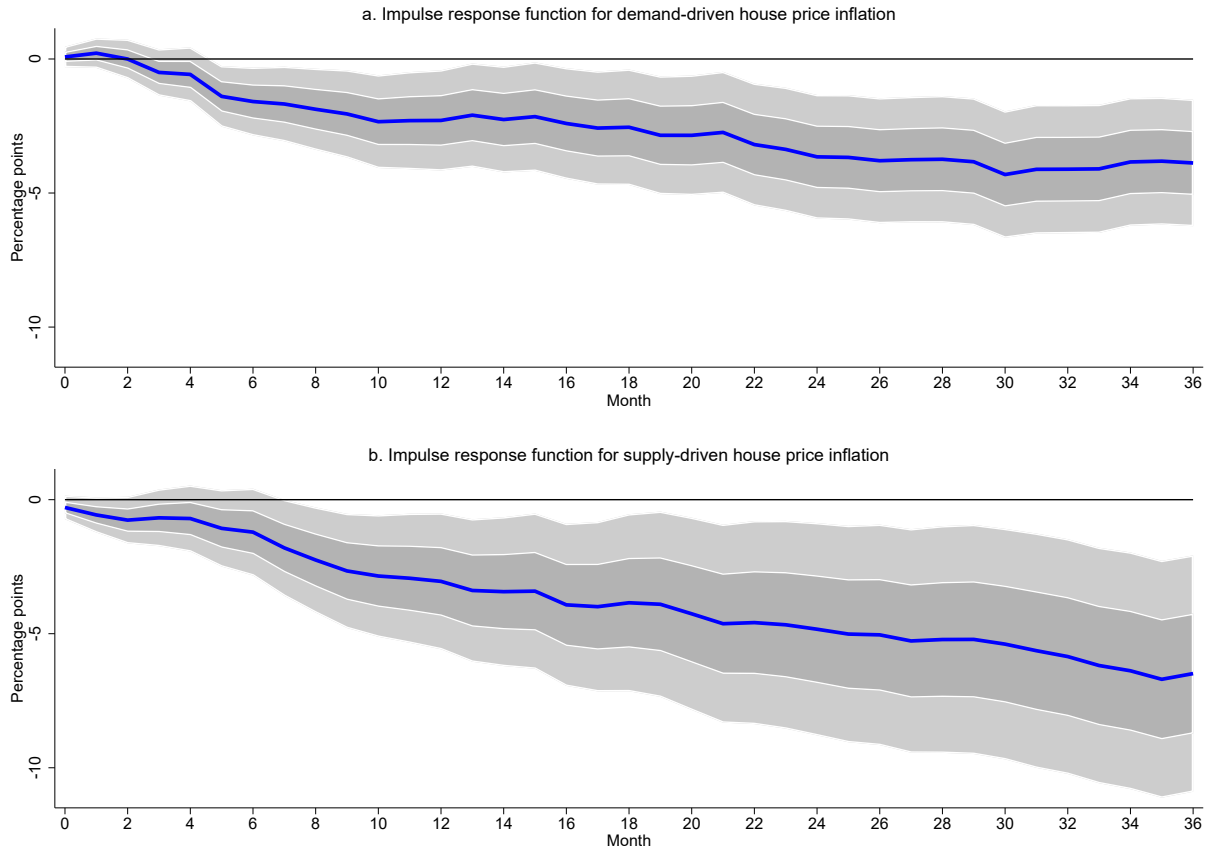
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points). Panel (a) shows the response of demand-driven house-price inflation, and panel (b) shows the response of supply-driven house-price inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 8: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: combined structure-and-land component of the NHPI.



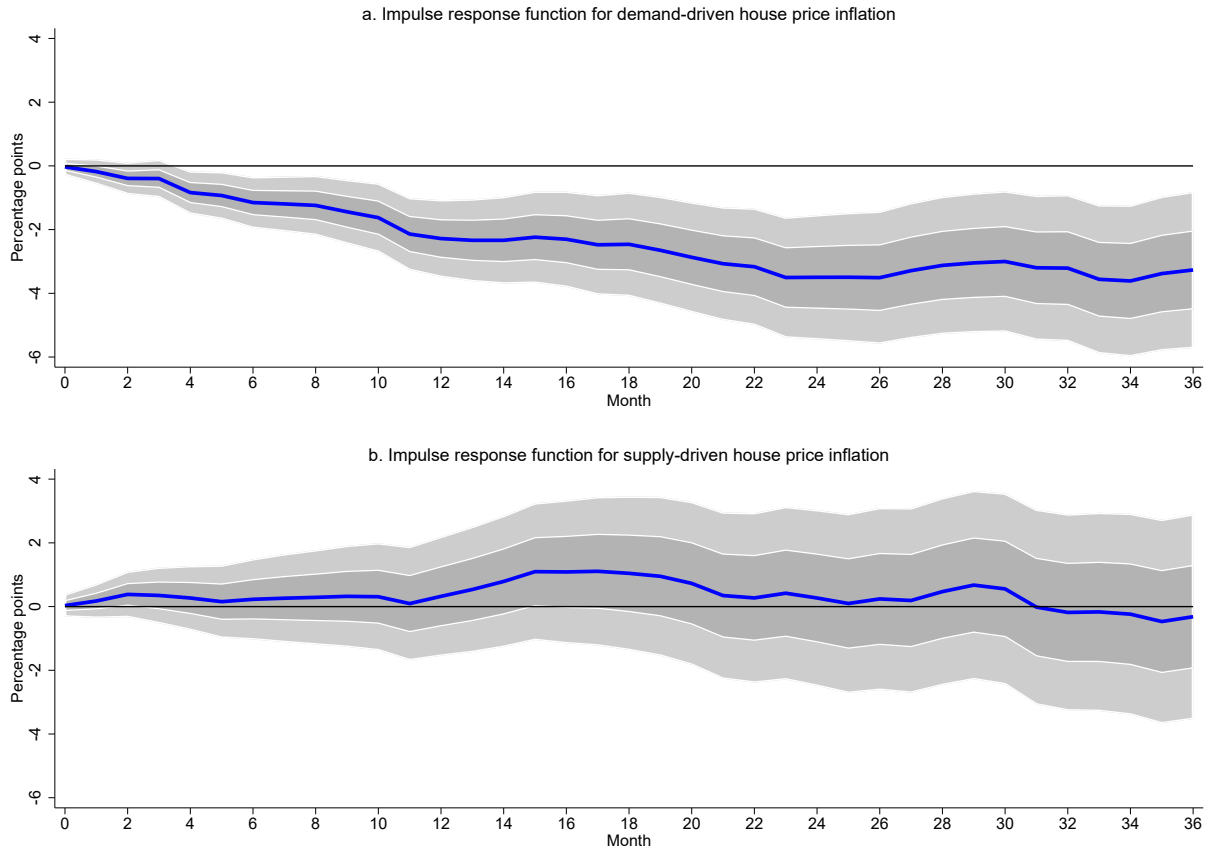
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 9: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: structure component of the NHPI.



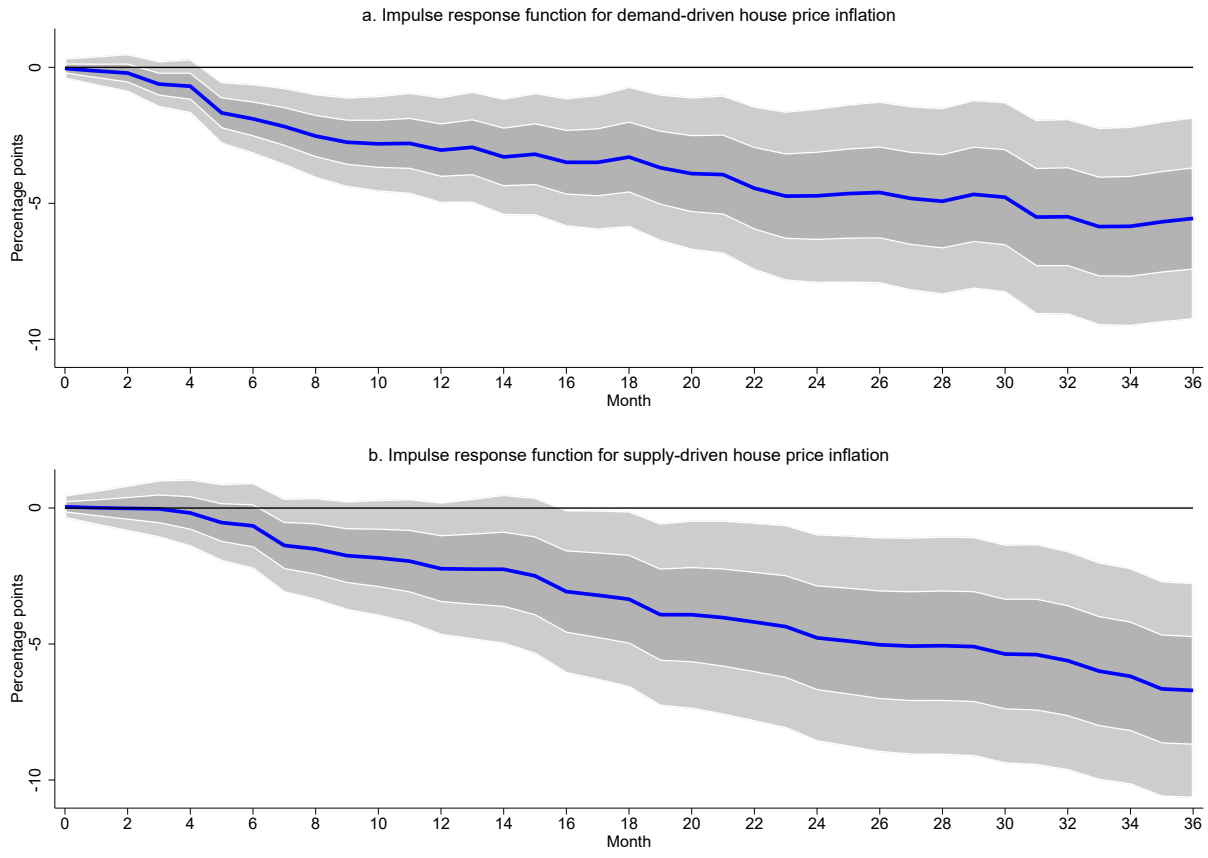
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 10: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: land component of the NHPI.



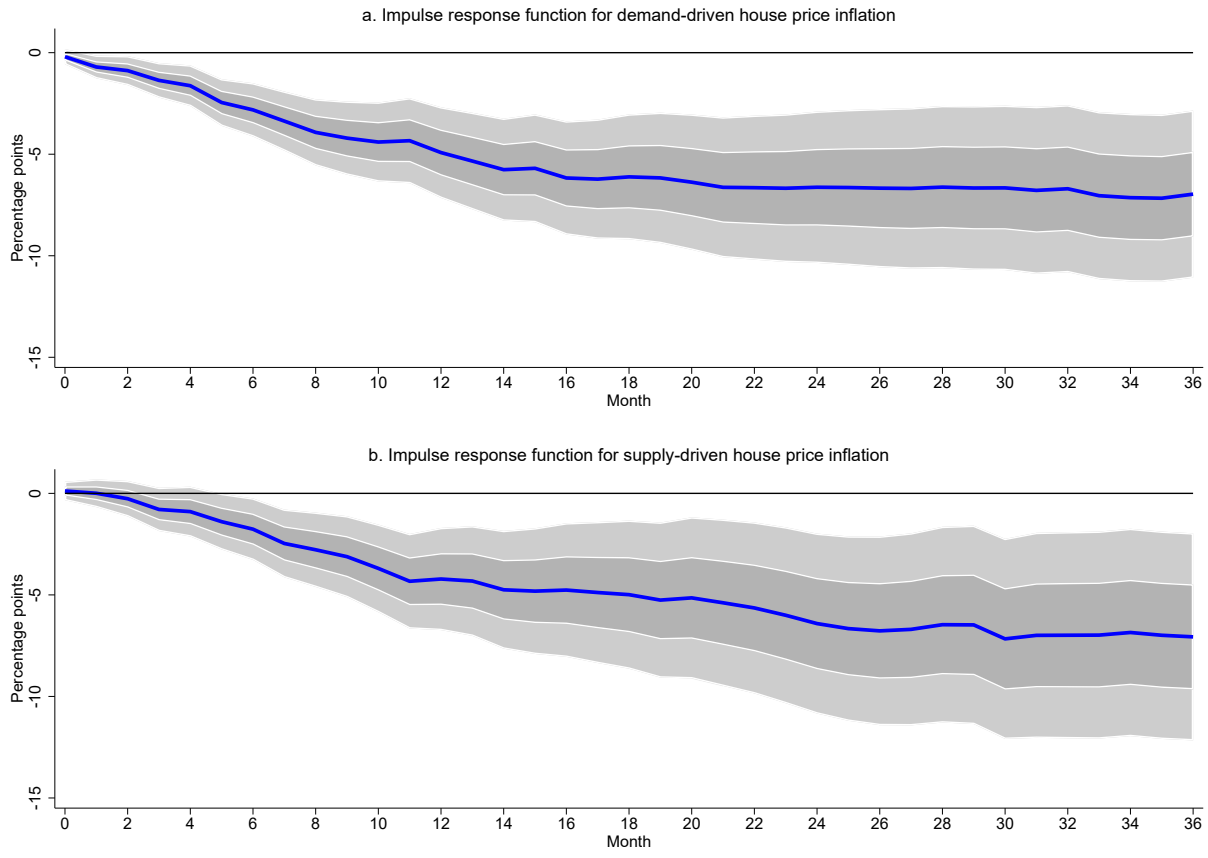
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 11: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: structure component of the NHPI, housing starts as the quantity measure.



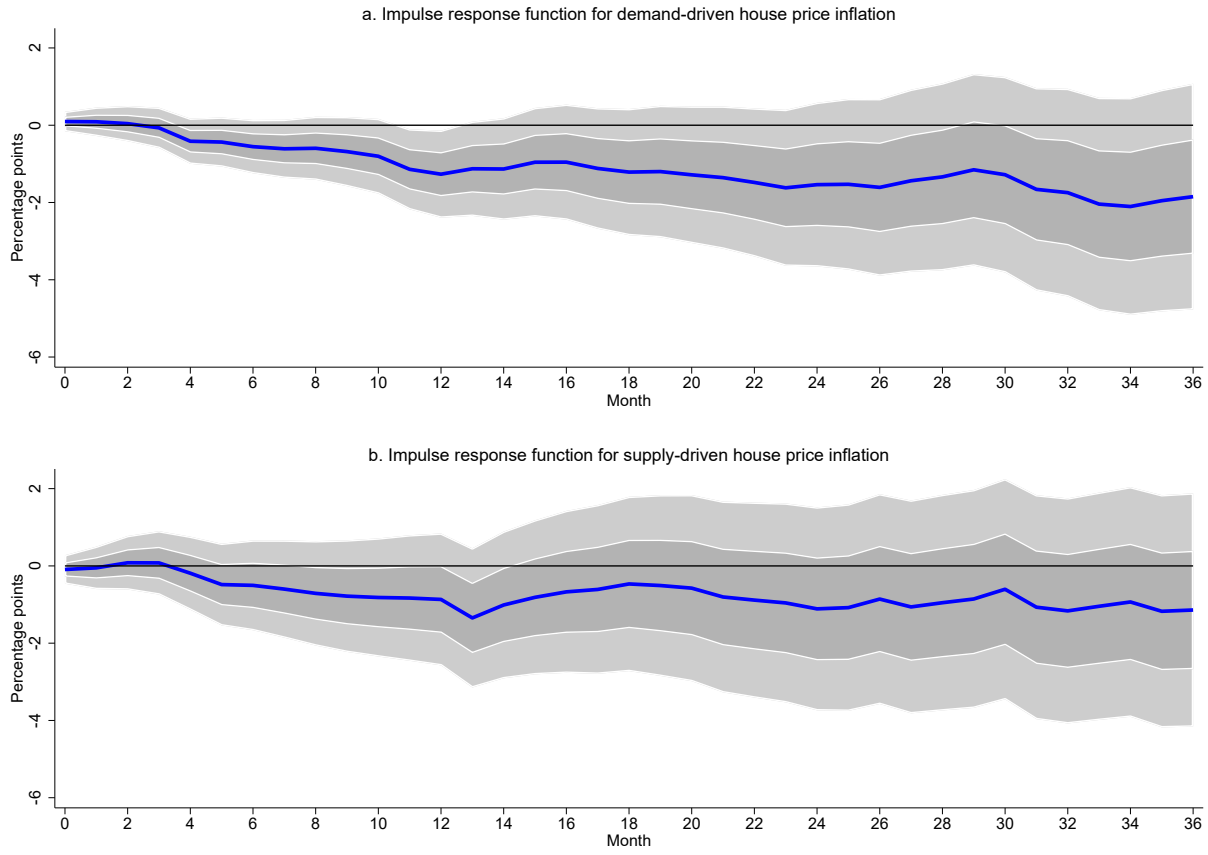
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using housing starts as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 12: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: structure component of the NHPI, completions as the quantity measure.



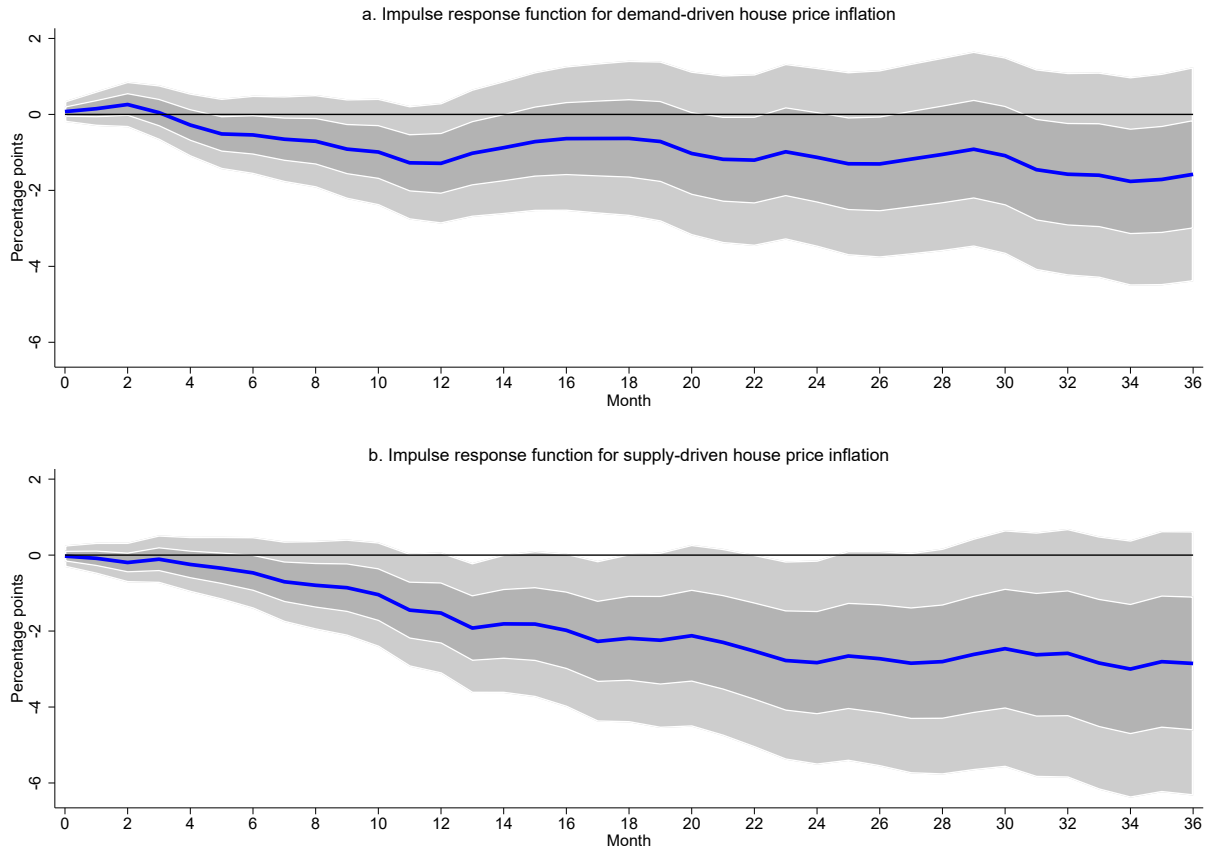
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using completions as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 13: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: land component of the NHPI, housing starts as the quantity measure.



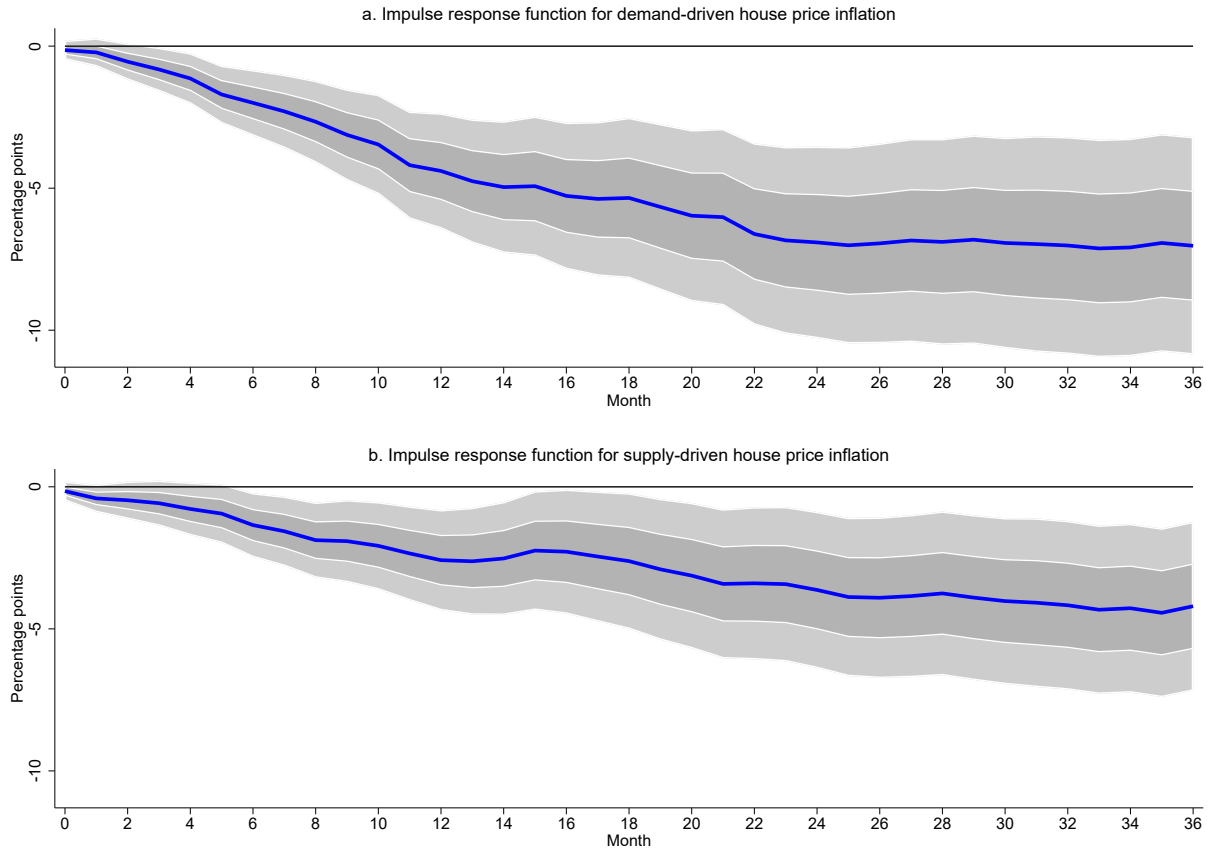
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using housing starts as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 14: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: land component of the NHPI, completions as the quantity measure.



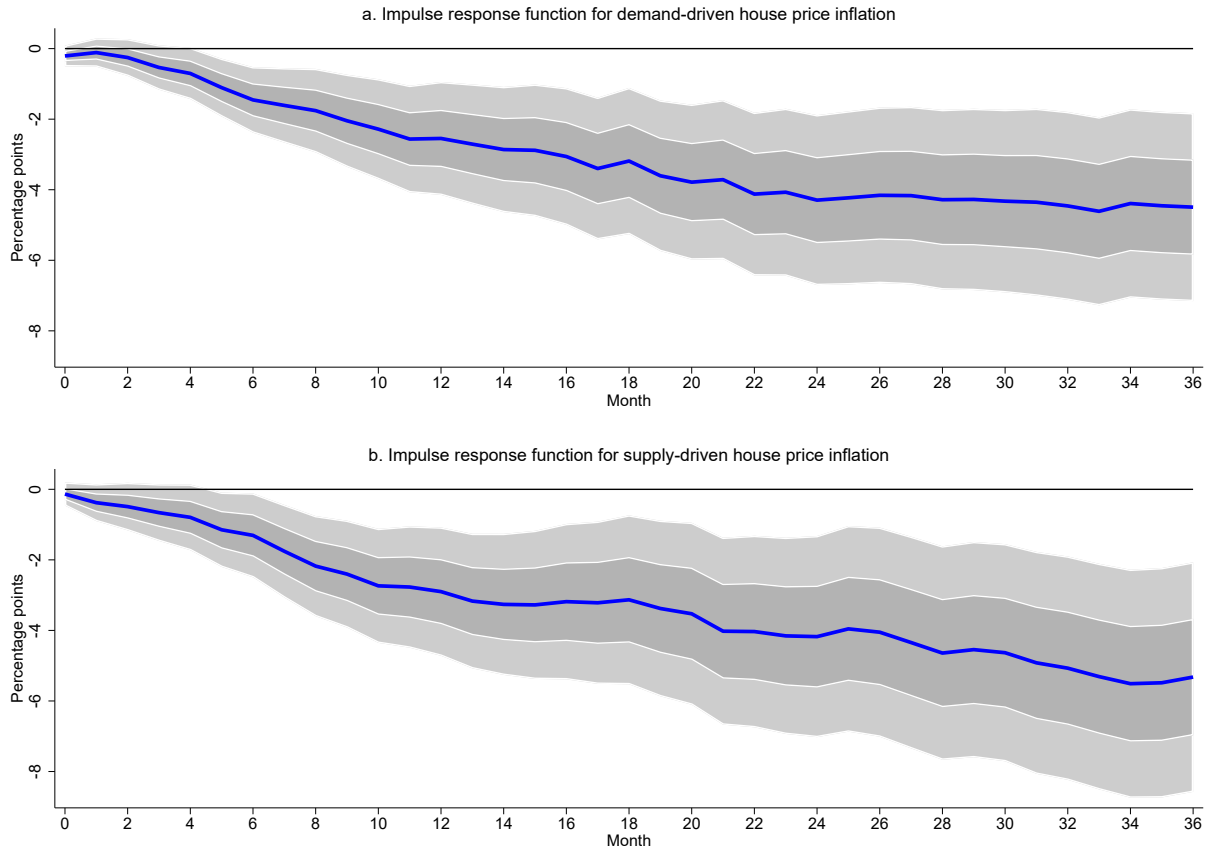
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using completions as the quantity measure. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 15: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: combined structure-and-land component of the NHPI, 10-year window, 3-month lag.



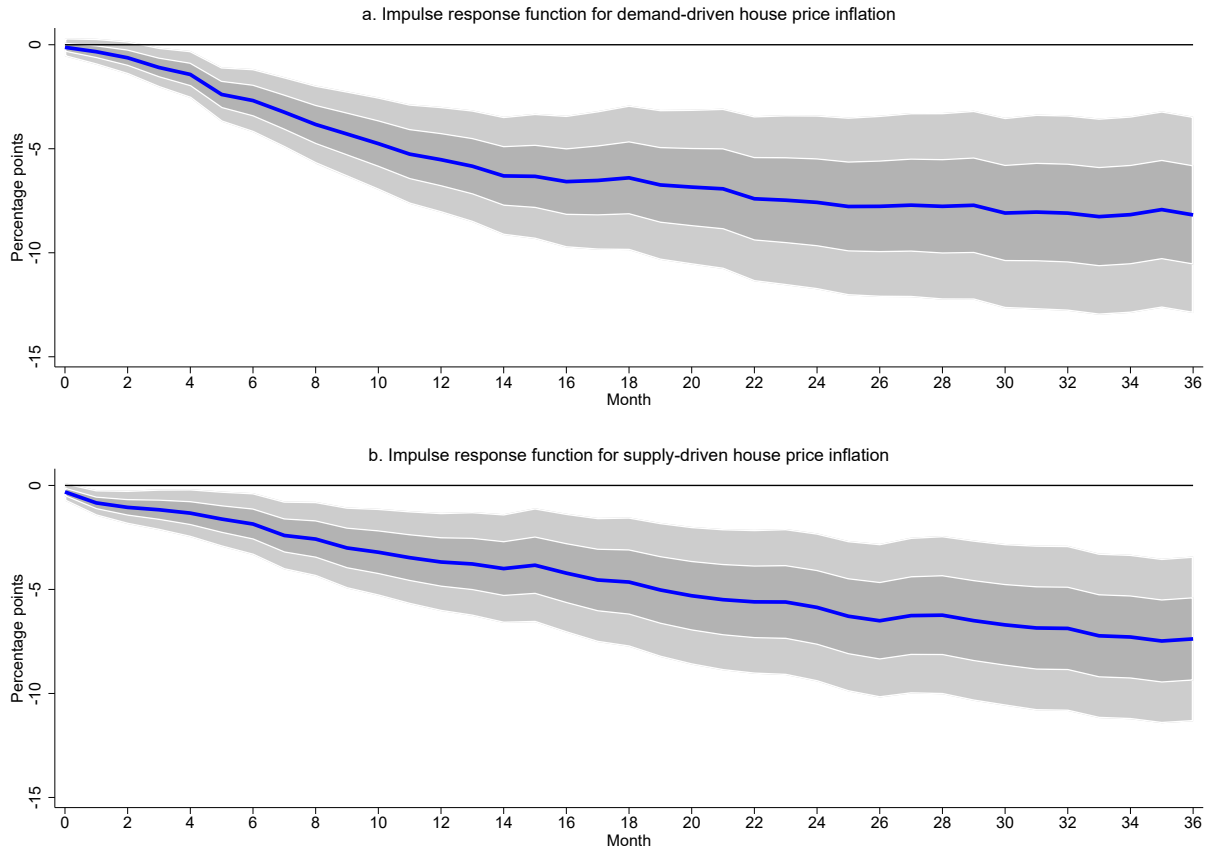
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure and a first-stage classification with a 10-year rolling window and a 3-month lag. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 16: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: combined structure-and-land component of the NHPI, 10-year window, 24-month lag.



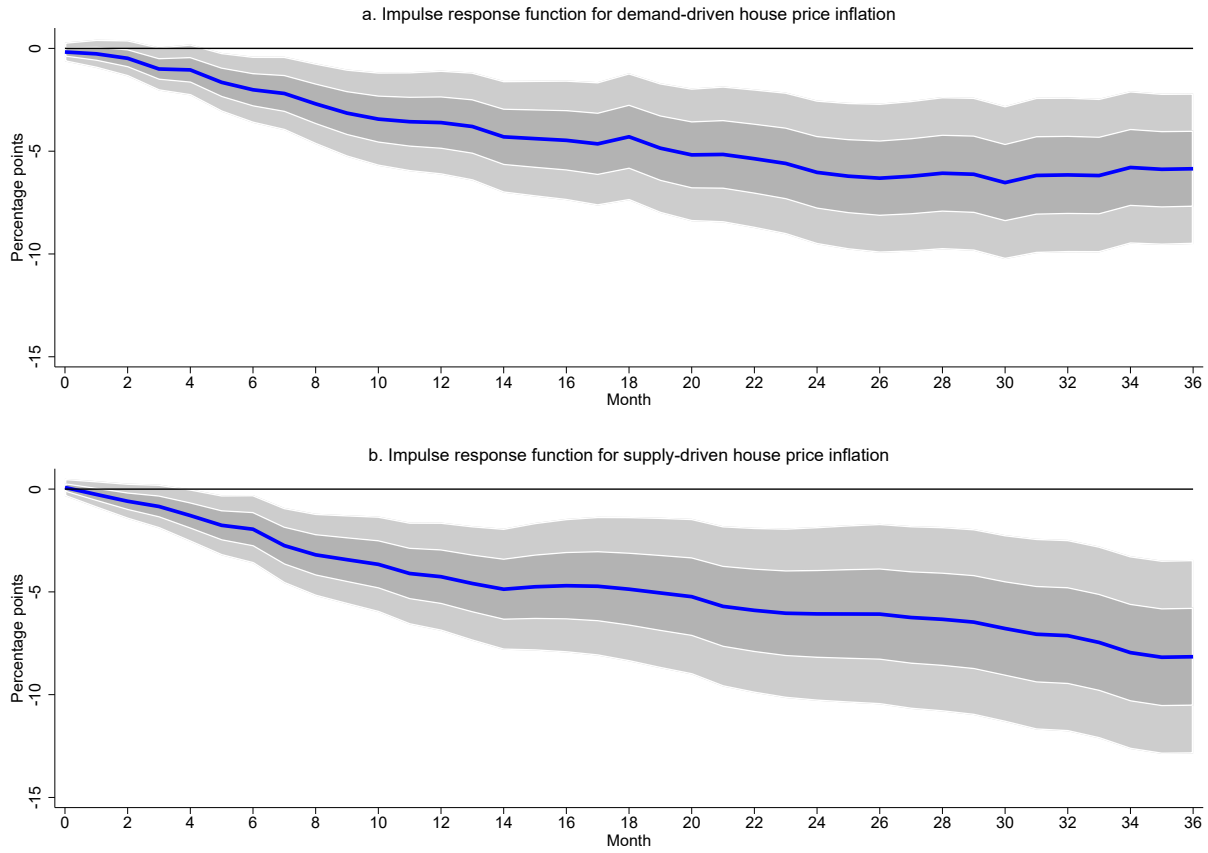
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure and a first-stage classification with a 10-year rolling window and a 24-month lag. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 17: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: structure component of the NHPI, 10-year window, 3-month lag.



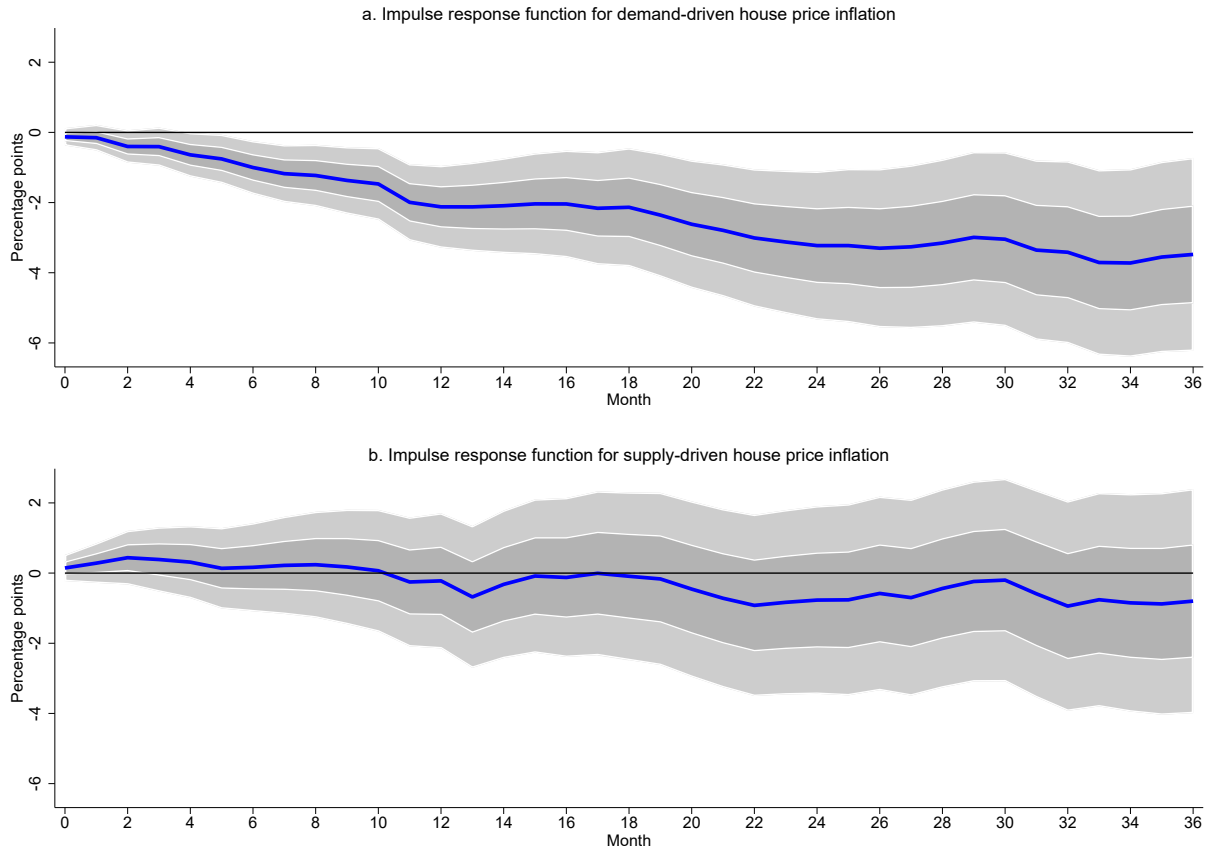
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure and a first-stage classification with a 10-year rolling window and a 3-month lag. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 18: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: structure component of the NHPI, 10-year window, 24-month lag.



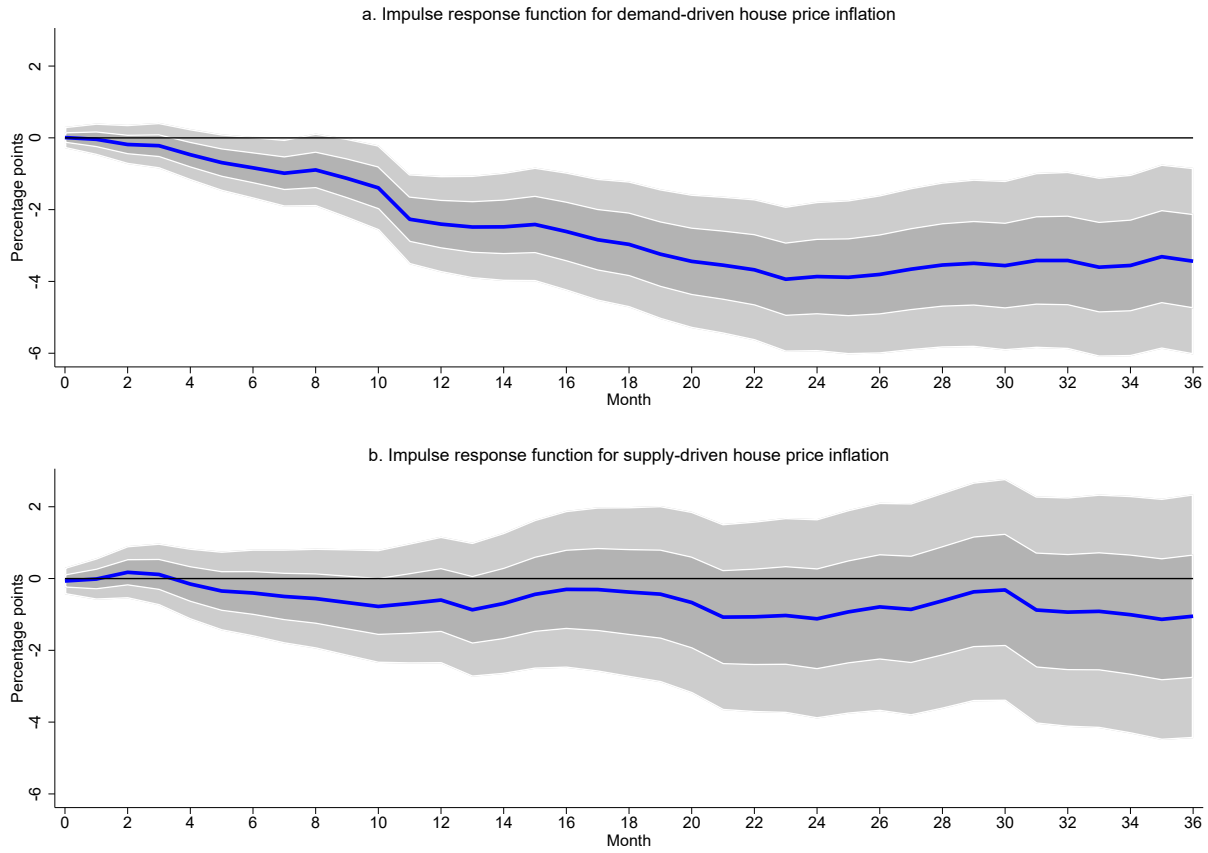
Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure and a first-stage classification with a 10-year rolling window and a 24-month lag. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 19: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: land component of the NHPI, 10-year window, 3-month lag.



Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure and a first-stage classification with a 10-year rolling window and a 3-month lag. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure 20: Impulse responses of demand- and supply-driven new house-price inflation to a monetary policy shock: land component of the NHPI, 10-year window, 24-month lag.



Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using units under construction as the quantity measure and a first-stage classification with a 10-year rolling window and a 24-month lag. Panel (a) shows demand-driven inflation and panel (b) shows supply-driven inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Appendix

A Robustness of the Local Projection Estimates

This appendix collects the robustness analysis for the local projection estimates. Section A.1 reports, for every first-stage window and lag, the share of the two-year response attributable to supply together with the descriptive D/S/B summary of each response; Section A.2 reports how the estimated coefficients vary across the full grid of first-stage windows and lags; and Section A.3 verifies that the main findings survive replacing the baseline magnitude threshold with a more inclusive sign-only classification.

A.1 Supply Shares Across Specifications

This subsection reports, for every first-stage window and lag combination, the share of the two-year house-price-inflation response attributable to supply, together with a descriptive D/S/B summary of the demand- and supply-driven responses. Tables A1 and A2 present the resale and new-housing results, respectively. The supply share is defined as $100 \times |\hat{\beta}^S_{24}| / (|\hat{\beta}^S_{24}| + |\hat{\beta}^D_{24}|)$, so a value below 50% indicates that the demand-driven response is larger in magnitude than the supply-driven one. The accompanying descriptive D/S/B summary labels each specification "D" (only the demand-driven response is distinguishable from zero under the 32% cutoff), "S" (only supply), or "B" (both). *We emphasise that this is a descriptive one-standard-error summary rather than conventional hypothesis-test inference, evaluated at a weak,* 32% *cutoff.* The magnitude shares themselves do not depend on the 32% cutoff, but the D/S/B labelling would change under a stricter conventional level; we adopt the weak level only as a compact way to summarise which responses are jointly distinguishable from zero at the two-year horizon, and we read it together with the stricter, conventional-level evidence in Section A.2 and in the main-text impulse responses.

In the resale market (Table A1), the supply share is below 30% in most specifications and close to zero in several, and the descriptive D/S/B summary is "D" or "B" in all but one case, supporting the interpretation that resale house-price inflation responds primarily through demand. The baseline specification (10-year window, 12-month lag) yields a supply share of only 21% with a demand-only ("D") classification. The single case in which the supply share exceeds 50%—the 15-year window paired with a 24-month lag—arises at the longest first-stage lag, consistent with the lag sensitivity documented in Section A.2.

The new housing market (Table A2) also points to the dominance of demand in magnitude, but with greater heterogeneity and a more active supply channel. For the combined structure-and-land index, supply shares cluster around 30–55% and the classification is "B" in essentially every specification, so demand and supply are both active under the descriptive 32% criterion even where demand is larger in magnitude. The structure component is similar—supply shares generally between 28% and 57%, with a "B" classification throughout—so there are no clearly demand-only cases. The land component exhibits by far the largest variation: supply shares fall below 10% under units under construction and some starts specifications, where the classification is correspondingly demand-led ("D"), but rise well above 50% under completions, where the classification turns "S". This stage dependence—demand-led at the forward-looking starts and under-construction stages, supply-led at the predetermined completions stage—mirrors the interpretation in Section 5.2.

Overall, reading the shares and the descriptive D/S/B summary together, monetary policy affects house-price inflation primarily through the demand channel, with supply playing a secondary role in resale but a materially larger and—for the structure component and the completions-stage land component—jointly active role in new housing.

A.2 Sensitivity to the Window and Lag Choice

Section A.1 summarises robustness in terms of the *share* of the response attributable to supply at a single horizon. This subsection complements that evidence by reporting how the

Table A1: Relative importance of supply-driven house-price inflation responses at the 2-year horizon: resale market (CREA, baseline classification)

Window	Lags	Supply share (%)	D/S/B
3	3	42.1	B
5	3	12.9	D
10	3	21.6	D
15	3	10.4	D
3	12	30.3	B
5	12	14.1	D
10	12	21.0	D
15	12	2.9	D
3	24	22.6	B
5	24	34.0	B
10	24	34.2	B
15	24	62.4	S

Notes: The supply share at the 24-month horizon is computed as $100 \times |\hat{\beta}^S_{24}|/(|\hat{\beta}^S_{24}| + |\hat{\beta}^D_{24}|)$; a value below 50% indicates a demand-driven response that is stronger in magnitude than the supply-driven one. The final column provides a descriptive one-standard-error summary: D indicates that only the demand response is distinguishable from zero under the 32% cutoff, S indicates only supply, and B indicates both. This column is not conventional hypothesis-test inference.

underlying local projection coefficients themselves—the same objects plotted in Figure 7 for the baseline (10-year window, 12-month lag) specification—vary across a grid of first-stage windows (3, 5, 10, and 15 years) and lags (3, 12, and 24 months)—the 10-year window and the 3-, 12-, and 24-month lags following the benchmark and robustness specifications of Shapiro (2026), with the shorter and longer windows added as our own sensitivity analysis—holding the resale (CREA) market and the baseline classification fixed.

Table A3 reports, for each of the 12 window–lag combinations, the horizon h at which the demand-driven and supply-driven coefficients reach their largest magnitude over $h = 0, \ldots, 36$, together with the corresponding coefficient value and significance level. Two features stand out. First, the demand-driven peak is consistently large and negative across all 12 specifications, ranging from -7.1 to -25.8, and is statistically significant at conventional levels in 11 of the 12 cases; this supports the baseline finding that monetary policy transmits to resale house-price inflation primarily through the demand-driven component. Second, the

Table A2: Relative importance of supply-driven house-price inflation responses at the 2-year horizon: new housing market (baseline classification)

Starts				Under construction				Completions			
Window	Lags	Share	D/S/B	Window	Lags	Share	D/S/B	Window	Lags	Share	D/S/B
Panel A: Combined (structure and land)											
3	3	43.7	B	3	3	30.8	B	3	3	43.6	B
5	3	46.1	B	5	3	31.1	B	5	3	45.1	B
10	3	40.8	B	10	3	34.4	B	10	3	49.4	B
15	3	43.3	B	15	3	43.9	B	15	3	56.1	B
3	12	38.5	B	3	12	38.1	B	3	12	34.3	B
5	12	52.5	B	5	12	46.2	B	5	12	48.7	B
10	12	37.8	B	10	12	50.3	B	10	12	50.8	B
15	12	35.9	B	15	12	36.3	B	15	12	54.6	B
3	24	48.8	B	3	24	31.8	B	3	24	29.9	B
5	24	41.7	B	5	24	42.2	B	5	24	39.4	B
10	24	31.5	B	10	24	49.3	B	10	24	42.1	B
15	24	38.1	B	15	24	39.2	B	15	24	67.2	B
Panel B: Land											
3	3	28.1	D	3	3	37.1	B	3	3	70.6	S
5	3	45.3	B	5	3	18.3	D	5	3	71.6	S
10	3	51.8	B	10	3	19.2	D	10	3	87.4	S
15	3	42.0	B	15	3	30.4	B	15	3	99.2	S
3	12	37.4	B	3	12	24.9	B	3	12	37.1	B
5	12	53.5	S	5	12	23.1	B	5	12	54.8	B
10	12	42.0	D	10	12	7.1	D	10	12	71.5	S
15	12	45.4	B	15	12	9.7	D	15	12	81.0	S
3	24	71.2	S	3	24	22.3	D	3	24	60.9	S
5	24	6.0	D	5	24	29.9	B	5	24	22.1	B
10	24	43.8	B	10	24	22.5	D	10	24	59.2	S
15	24	27.6	D	15	24	21.2	D	15	24	83.8	S
Panel C: Structure											
3	3	42.1	B	3	3	35.1	B	3	3	41.2	B
5	3	44.8	B	5	3	36.3	B	5	3	42.3	B
10	3	49.4	B	10	3	43.6	B	10	3	49.3	B
15	3	47.1	B	15	3	44.4	B	15	3	49.1	B
3	12	35.0	B	3	12	33.7	B	3	12	34.8	B
5	12	44.7	B	5	12	44.4	B	5	12	44.3	B
10	12	50.3	B	10	12	57.0	B	10	12	49.2	B
15	12	37.8	B	15	12	40.9	B	15	12	54.4	B
3	24	34.6	B	3	24	28.2	B	3	24	28.7	B
5	24	40.4	B	5	24	27.8	B	5	24	34.4	B
10	24	27.9	B	10	24	50.2	B	10	24	43.9	B
15	24	37.7	B	15	24	50.9	B	15	24	55.3	B

Notes: The supply share at the 24-month horizon is computed as $100 \times |\hat{\beta}^S_{24}|/(|\hat{\beta}^S_{24}| + |\hat{\beta}^D_{24}|)$; a value below 50% indicates a demand-driven response that is stronger in magnitude than the supply-driven one. The final column provides a descriptive one-standard-error summary: D indicates that only the demand response is distinguishable from zero under the 32% cutoff, S indicates only supply, and B indicates both. This column is not conventional hypothesis-test inference. Panels report the combined, land, and structure NHPI price components; within each panel, columns group the three CMHC quantity measures (housing starts, units under construction, and completions).

supply-driven coefficient is more sensitive to the choice of lag length than to the choice of window: at a 24-month lag, the supply-driven peak is negative and statistically significant (at least at the 5% level) for all four windows, at horizons of roughly two to three years; at a 12-month lag it is significant for the 3- and 5-year windows but not for the baseline 10-year window or the 15-year window; and at a 3-month lag it is significant only for the 3-year window. The supply-driven peak therefore tends to grow in magnitude and statistical significance as the first-stage lag lengthens, whereas at the baseline lag and window it remains statistically insignificant.

This pattern indicates that the paper's central conclusion—that demand dominates the transmission of monetary policy to resale house-price inflation—is robust to the choice of rolling window at the baseline lag length, but is sensitive to extending the first-stage lag length to 24 months, which materially increases the estimated role of supply. We treat the 24-month-lag version as a demanding sensitivity case rather than the preferred monthly classification. It permits substantially more persistent price–quantity dynamics than the baseline and produces a materially larger supply response. Nonetheless, we report this sensitivity transparently here rather than presenting only the baseline result, consistent with the discussion in Section 4.3.1.

Figures A1 and A2 illustrate this lag sensitivity directly by holding the window fixed at the baseline's 10 years and varying only the lag length, at 3 and 24 months respectively, for comparison against the baseline (10-year window, 12-month lag) result already shown in Figure 7. At a 3-month lag (Figure A1), the qualitative pattern mirrors the baseline: the demand-driven response is large and significant, while the supply-driven response remains comparatively small. At a 24-month lag (Figure A2), however, the supply-driven response becomes visibly larger and, as reported in Table A3, statistically significant at conventional levels.

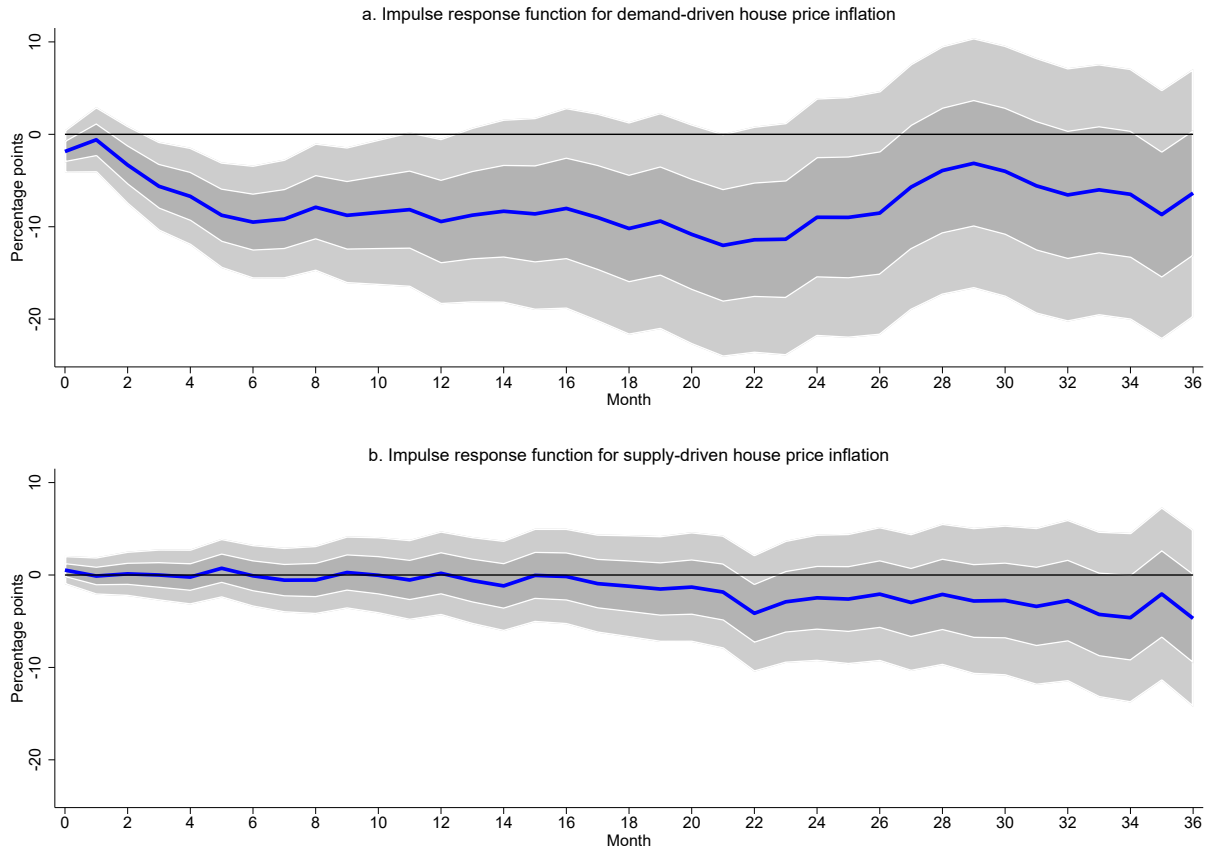
Tables A4–A6 repeat this peak-sensitivity exercise for the new housing market, reporting one table per quantity measure—housing starts, units under construction, and completions—

each with separate panels for the combined, structure, and land components. The demand-driven peak is again large, negative, and statistically significant across essentially all window–lag combinations, components, and quantity measures, mirroring the resale-market result. The supply-driven peak, however, behaves very differently from the resale market. For the combined and structure components, it is statistically significant across the great majority of specifications and at all three lag lengths, rather than only at the 24-month lag as in the resale market; for the structure component it is frequently comparable to, and in several specifications exceeds, the demand-driven peak in magnitude. The land component is again the exception, with a supply-driven peak that is small in magnitude, often statistically insignificant, and in some specifications positive. Taken together, the peak-sensitivity tables indicate that the larger role of supply in the new housing market is not confined to the baseline specification or to a particular lag choice, but is instead a robust feature of the combined and structure components across the window–lag grid, while the land component is demand-dominated when quantity is measured by starts or units under construction but often supply-dominated when quantity is measured by completions.

A.3 Robustness to the Sign-Only Classification (Threshold Zero)

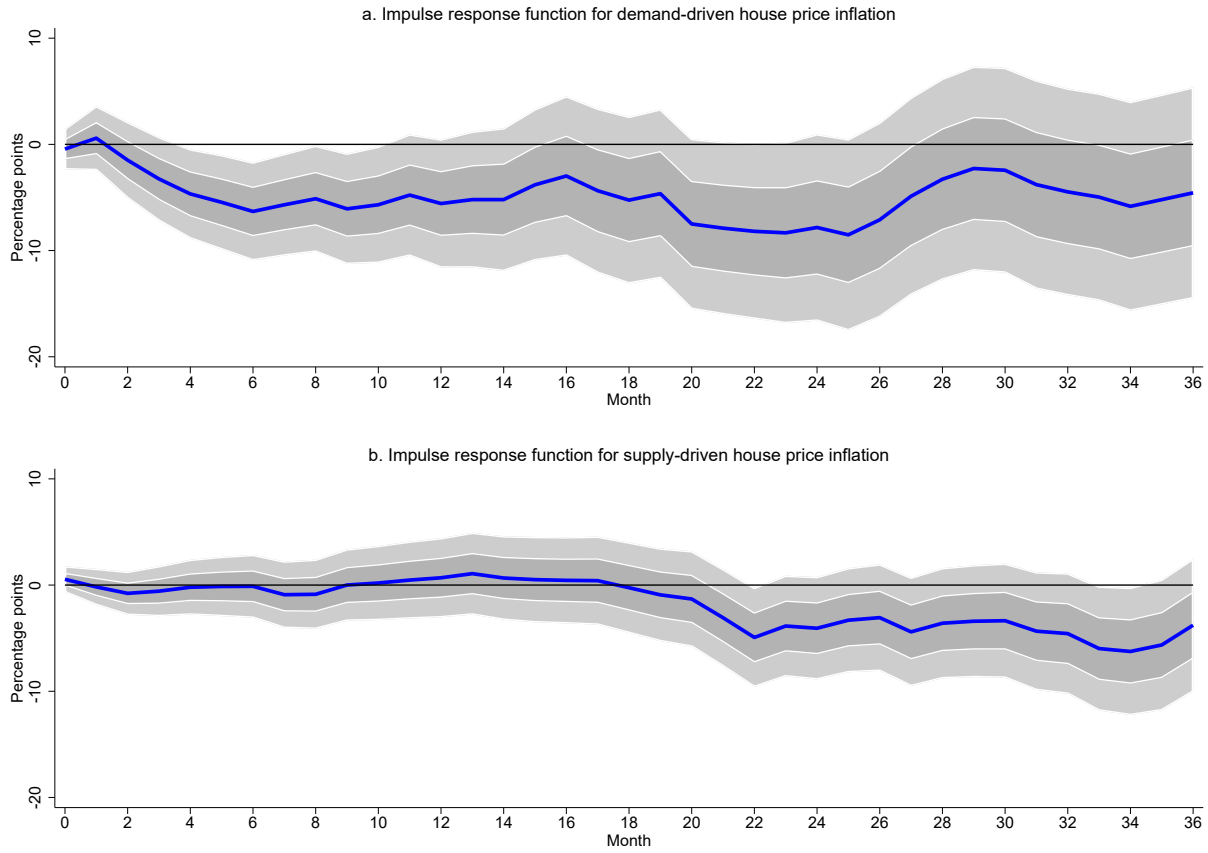
The classification used throughout the main text retains only those months whose price and quantity innovations exceed a strict magnitude threshold of 0.3 standard deviations—the paper's *baseline* classification—so that each period is assigned to the demand- or supply-driven category only when the underlying co-movement is pronounced. Shapiro (2026)'s main classification is sign-based, assigning a period to demand when the two innovations share the same sign and to supply when they have opposite signs; the paper also considers precision filters that treat very small residuals as ambiguous. Our threshold-zero rule implements the sign-only version, without requiring either innovation to exceed any magnitude threshold. This rule is the most inclusive classification, since it assigns every period with a determinate sign pattern to one of the two categories. This appendix re-estimates the entire decomposi-

Figure A1: Impulse responses of demand- and supply-driven house-price inflation to a monetary policy shock: resale markets, 10-year window, 3-month lag.



Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using a first-stage classification with a 10-year rolling window and a 3-month lag. Panel (a) shows the response of demand-driven house-price inflation, and panel (b) shows the response of supply-driven house-price inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Figure A2: Impulse responses of demand- and supply-driven house-price inflation to a monetary policy shock: resale markets, 10-year window, 24-month lag.



Note: The figure reports cumulative monthly inflation responses measured from $t-1$ to a one-unit monetary policy shock (100 basis points), using a first-stage classification with a 10-year rolling window and a 24-month lag. Panel (a) shows the response of demand-driven house-price inflation, and panel (b) shows the response of supply-driven house-price inflation. The solid line is the local projection point estimate and the shaded areas are pointwise 68% and 95% confidence bands.

Table A3: Sensitivity of the peak local projection coefficient to the first-stage window and lag choice: resale market (CREA, baseline classification)

Window (yrs)	Lag (months)	Demand-driven peak		Supply-driven peak	
		h	Coefficient	h	Coefficient
3	3	20	-10.27**	36	-8.298**
3	12	22	-16.25***	22	-7.471**
3	24	22	-19.53***	34	-6.478**
5	3	12	-9.510**	36	-5.075
5	12	21	-10.62**	34	-7.315**
5	24	22	-15.31***	22	-7.851**
10	3	21	-12.02*	36	-4.693
10	12 (baseline)	12	-9.248**	35	-5.403
10	24	25	-8.514*	34	-6.252**
15	3	22	-25.82***	33	-6.108
15	12	22	-20.47***	15	4.396
15	24	21	-7.076	33	-10.27**

Notes: For each first-stage window (rows) and lag (rows) combination, the table reports the monthly horizon h at which the estimated coefficient on the contemporaneous monetary policy shock, β_h^k in Equation 10 for $k \in \{D, S\}$, reaches its largest magnitude over $h = 0, \ldots, 36$, together with the coefficient value at that horizon. All coefficients are estimated on CREA data using the baseline classification. *** p<0.01, ** p<0.05, * p<0.1.

tion and local projection exercise under threshold zero, holding fixed the baseline first-stage specification (10-year rolling window, 12-month lag), to verify that the paper's two central findings do not depend on the baseline magnitude threshold.

Figure A3 reports the resale-market (CREA) results. The pattern is essentially identical to the baseline classification: the demand-driven coefficient is large and negative, statistically significant from the fourth month onward, and reaches its largest magnitude of -16.88 (significant at the 5% level) at a horizon of roughly twenty-two months, whereas the supply-driven coefficient is small in magnitude and not statistically significant at any horizon. Under the sign-only classification, as under the baseline classification, the resale response is concentrated in the demand-driven component.

Figures A4–A6 report the corresponding results for the new housing market, using the under-construction quantity measure to match the baseline new-housing figures in the main text. For the combined and structure components (Figures A4 and A5), the supply-driven response is large and negative—comparable to, and at longer horizons exceeding, the demand-

Table A4: Sensitivity of the peak local projection coefficient to the first-stage window and lag choice: new housing market (housing starts, baseline classification).

Window (yrs)	Lag (months)	Demand-driven peak		Supply-driven peak	
		h	Coefficient	h	Coefficient
Panel A: Combined (house and land)					
3	3	36	-5.725***	27	-3.869***
3	12	36	-10.52***	33	-6.021**
3	24	36	-8.565**	35	-6.730***
5	3	35	-6.620***	33	-5.478***
5	12	33	-4.534***	35	-5.409**
5	24	34	-9.593***	36	-7.216***
10	3	34	-6.023***	33	-4.359**
10	12 (baseline)	35	-5.489***	35	-4.387**
10	24	33	-5.999***	34	-3.045**
15	3	34	-5.796***	35	-4.932**
15	12	35	-6.303***	35	-4.172**
15	24	35	-6.299***	36	-4.512**
Panel B: House					
3	3	36	-8.038***	35	-5.609**
3	12	35	-13.64***	36	-7.377***
3	24	35	-12.63***	33	-7.306***
5	3	35	-6.770***	35	-6.099***
5	12	33	-6.109***	36	-5.583**
5	24	35	-12.75***	33	-8.374***
10	3	35	-6.741***	36	-7.482***
10	12 (baseline)	33	-5.855***	36	-6.706***
10	24	35	-9.896***	34	-4.037**
15	3	33	-7.132***	35	-7.792***
15	12	35	-8.382***	36	-6.736**
15	24	33	-8.670***	36	-6.257***
Panel C: Land					
3	3	36	-3.609*	27	-1.193
3	12	35	-4.679	34	-2.613
3	24	36	-2.573	34	-3.715*
5	3	35	-3.697*	13	-2.188**
5	12	36	-1.582*	23	-1.070
5	24	34	-6.332*	36	-1.599
10	3	35	-3.358*	23	-2.455
10	12 (baseline)	34	-2.104	13	-1.346
10	24	34	-3.079*	35	-2.204
15	3	35	-3.365	34	-1.870
15	12	34	-2.759*	36	-2.102
15	24	34	-3.236*	35	-1.552

Notes: For each first-stage window and lag combination, the table reports the monthly horizon h at which the estimated coefficient on the contemporaneous monetary policy shock, β_h^k in equation (10) for $k \in \{D, S\}$, reaches its largest magnitude over $h = 0, \ldots, 36$, together with the coefficient value at that horizon. All coefficients are estimated on the New Housing Price Index using housing starts as the quantity measure, under the baseline classification. *** p<0.01, ** p<0.05, * p<0.1.

Table A5: Sensitivity of the peak local projection coefficient to the first-stage window and lag choice: new housing market (units under construction, baseline classification).

Window (yrs)	Lag (months)	Demand-driven peak		Supply-driven peak	
		h	Coefficient	h	Coefficient
Panel A: Combined (house and land)					
3	3	25	-6.466***	35	-4.592***
3	12	33	-9.757***	36	-7.109***
3	24	36	-10.97***	34	-6.261***
5	3	25	-7.127***	35	-4.090***
5	12	29	-3.863***	34	-4.274***
5	24	36	-9.673***	34	-6.993***
10	3	33	-7.121***	35	-4.435***
10	12 (baseline)	36	-4.235***	35	-5.056***
10	24	33	-4.612***	34	-5.511***
15	3	36	-5.101***	30	-4.202***
15	12	33	-5.865***	35	-4.755**
15	24	33	-5.781***	35	-4.781**
Panel B: House					
3	3	29	-9.089***	35	-7.096***
3	12	36	-15.05***	35	-7.558***
3	24	33	-15.22***	35	-7.058***
5	3	30	-8.053***	35	-6.735***
5	12	30	-6.761***	35	-6.833***
5	24	35	-14.46***	36	-7.030***
10	3	33	-8.264***	35	-7.478***
10	12 (baseline)	30	-4.311***	35	-6.699***
10	24	30	-6.527***	35	-8.180***
15	3	30	-6.572***	35	-6.473***
15	12	30	-7.463***	35	-6.532***
15	24	33	-5.678***	35	-7.268***
Panel C: Land					
3	3	34	-3.010**	35	-1.667
3	12	34	-5.912*	35	-1.857
3	24	36	-5.125*	34	-1.830
5	3	34	-3.186*	26	-0.752
5	12	23	-3.776***	35	-1.564
5	24	35	-4.885	34	-1.995
10	3	34	-3.721***	32	-0.940
10	12 (baseline)	34	-3.612***	17	1.107
10	24	23	-3.938***	35	-1.136
15	3	34	-3.527**	22	-1.525
15	12	34	-3.384**	35	-0.979
15	24	26	-3.626***	35	-1.633

Notes: For each first-stage window and lag combination, the table reports the monthly horizon h at which the estimated coefficient on the contemporaneous monetary policy shock, β_h^k in equation (10) for $k \in \{D, S\}$, reaches its largest magnitude over $h = 0, \ldots, 36$, together with the coefficient value at that horizon. All coefficients are estimated on the New Housing Price Index using units under construction as the quantity measure, under the baseline classification. *** p<0.01, ** p<0.05, * p<0.1.

Table A6: Sensitivity of the peak local projection coefficient to the first-stage window and lag choice: new housing market (completions, baseline classification).

Window (yrs)	Lag (months)	Demand-driven peak		Supply-driven peak	
		h	Coefficient	h	Coefficient
Panel A: Combined (house and land)					
3	3	35	-5.667***	33	-3.972**
3	12	36	-10.42***	35	-5.730**
3	24	35	-11.19***	31	-4.619***
5	3	33	-5.806***	32	-4.749**
5	12	35	-5.032***	25	-3.937***
5	24	36	-10.28***	35	-6.693***
10	3	35	-5.545***	30	-5.305***
10	12 (baseline)	33	-5.155***	30	-5.098***
10	24	33	-5.561***	36	-5.030**
15	3	35	-4.643**	31	-5.801**
15	12	34	-5.388***	30	-6.083***
15	24	34	-3.822***	36	-7.555***
Panel B: House					
3	3	36	-8.268***	30	-5.390***
3	12	36	-14.39***	35	-6.994***
3	24	36	-15.70***	30	-6.327***
5	3	35	-8.084***	33	-5.752***
5	12	35	-7.758***	30	-5.733***
5	24	35	-15.02***	30	-6.870***
10	3	35	-6.774***	33	-7.028**
10	12 (baseline)	35	-7.168***	30	-7.165***
10	24	34	-6.110***	36	-5.619***
15	3	35	-6.622***	36	-7.098**
15	12	35	-6.617***	33	-7.921***
15	24	35	-6.683***	36	-8.380***
Panel C: Land					
3	3	35	-1.153	34	-1.919
3	12	35	-4.524	34	-2.903
3	24	35	-3.415	36	-3.336**
5	3	35	-1.920	36	-2.952
5	12	35	-2.602**	34	-2.402
5	24	35	-5.747	34	-1.549
10	3	34	-1.610	34	-4.446**
10	12 (baseline)	34	-1.763	34	-2.999*
10	24	34	-2.256	36	-1.677
15	3	19	0.422	34	-3.721*
15	12	34	-1.303	34	-2.618
15	24	34	-1.198	33	-2.961**

Notes: For each first-stage window and lag combination, the table reports the monthly horizon h at which the estimated coefficient on the contemporaneous monetary policy shock, β_h^k in equation (10) for $k \in \{D, S\}$, reaches its largest magnitude over $h = 0, \ldots, 36$, together with the coefficient value at that horizon. All coefficients are estimated on the New Housing Price Index using completions as the quantity measure, under the baseline classification. *** p<0.01, ** p<0.05, * p<0.1.

driven response in magnitude—so the substantially larger role of supply in the new housing market survives at threshold zero. The land component (Figure A6) again behaves differently: the demand-driven response is negative and significant while the supply-driven response hovers around zero and is not distinguishable from it, leaving land demand-dominated. Both central findings of the paper—demand dominance in the resale market and a larger supply role for the structure-intensive components of new housing but not for land—are thus robust to adopting Shapiro (2026)'s sign-only (threshold-zero) classification.

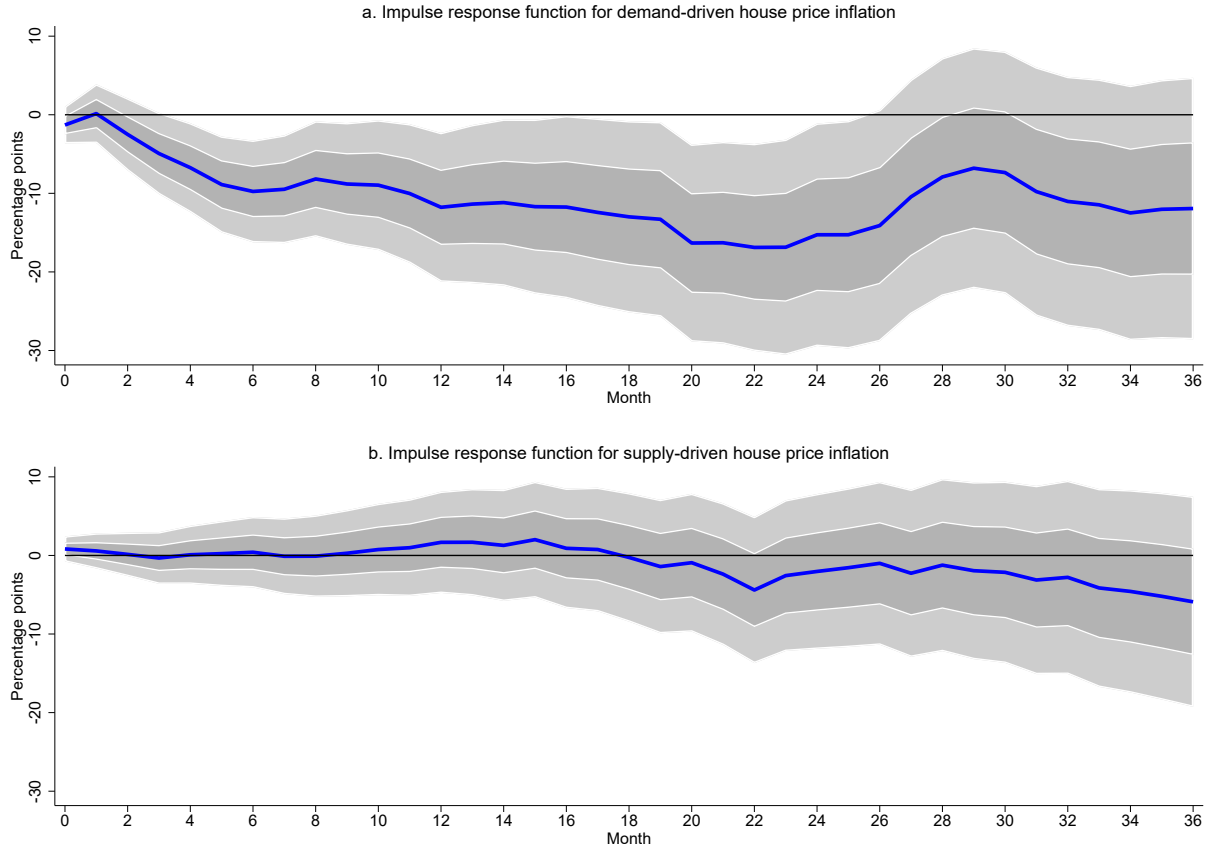


Figure A3: Resale market (CREA): impulse responses of demand- and supply-driven house-price inflation to a monetary policy shock under the sign-only classification (threshold zero, 10-year window, 12-month lag).

Note: The blue line is the local projection point estimate of the response of demand-driven (Panel a) and supply-driven (Panel b) house-price inflation to a monetary policy shock; shaded areas are 68% and 95% confidence bands. The decomposition uses Shapiro (2026)'s sign-only classification (threshold zero) with a 10-year rolling window and 12-month lag.

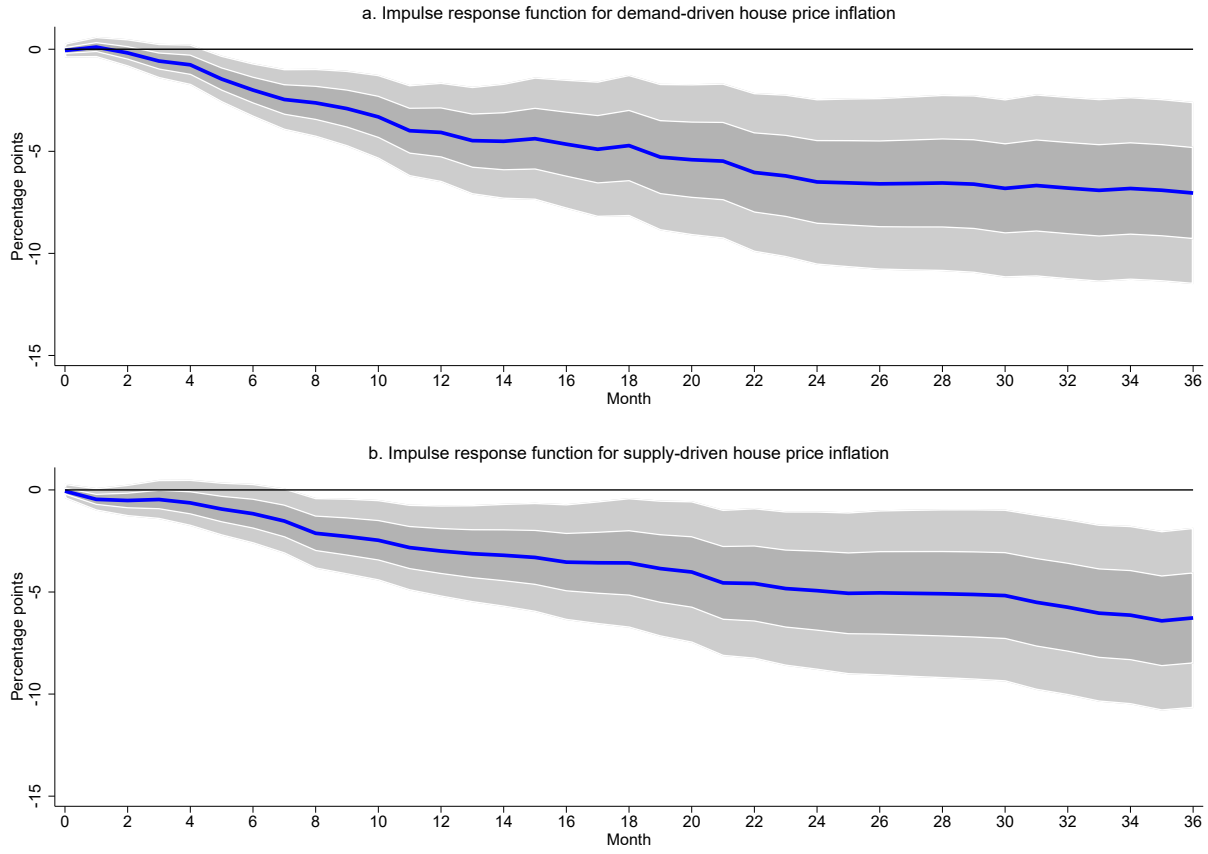


Figure A4: New housing market, combined price component (NHPI, under-construction quantity): impulse responses under the sign-only classification (threshold zero, 10-year window, 12-month lag).

Note: See the note to Figure A3. The decomposition uses the combined (structure and land) NHPI price component and the under-construction quantity measure.

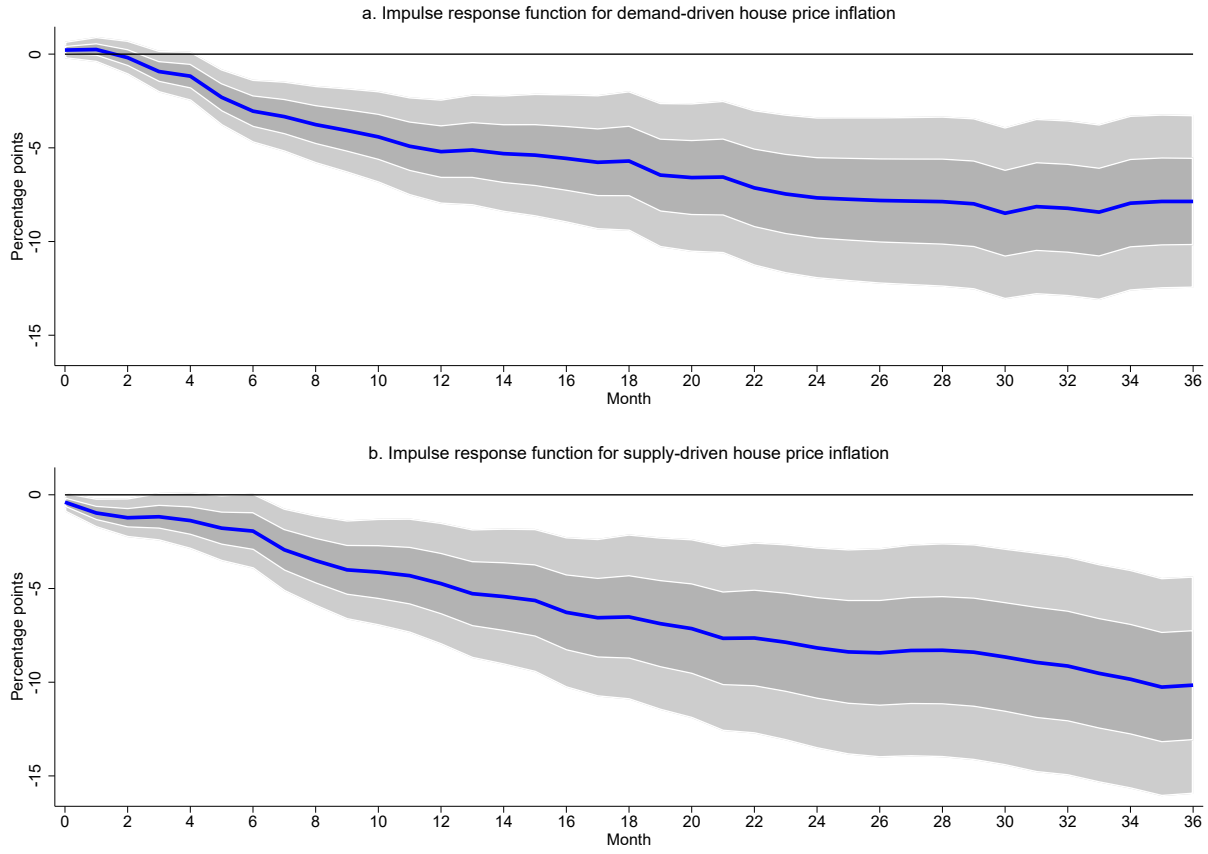


Figure A5: New housing market, structure price component (NHPI, under-construction quantity): impulse responses under the sign-only classification (threshold zero, 10-year window, 12-month lag).

Note: See the note to Figure A3. The decomposition uses the structure NHPI price component and the under-construction quantity measure.

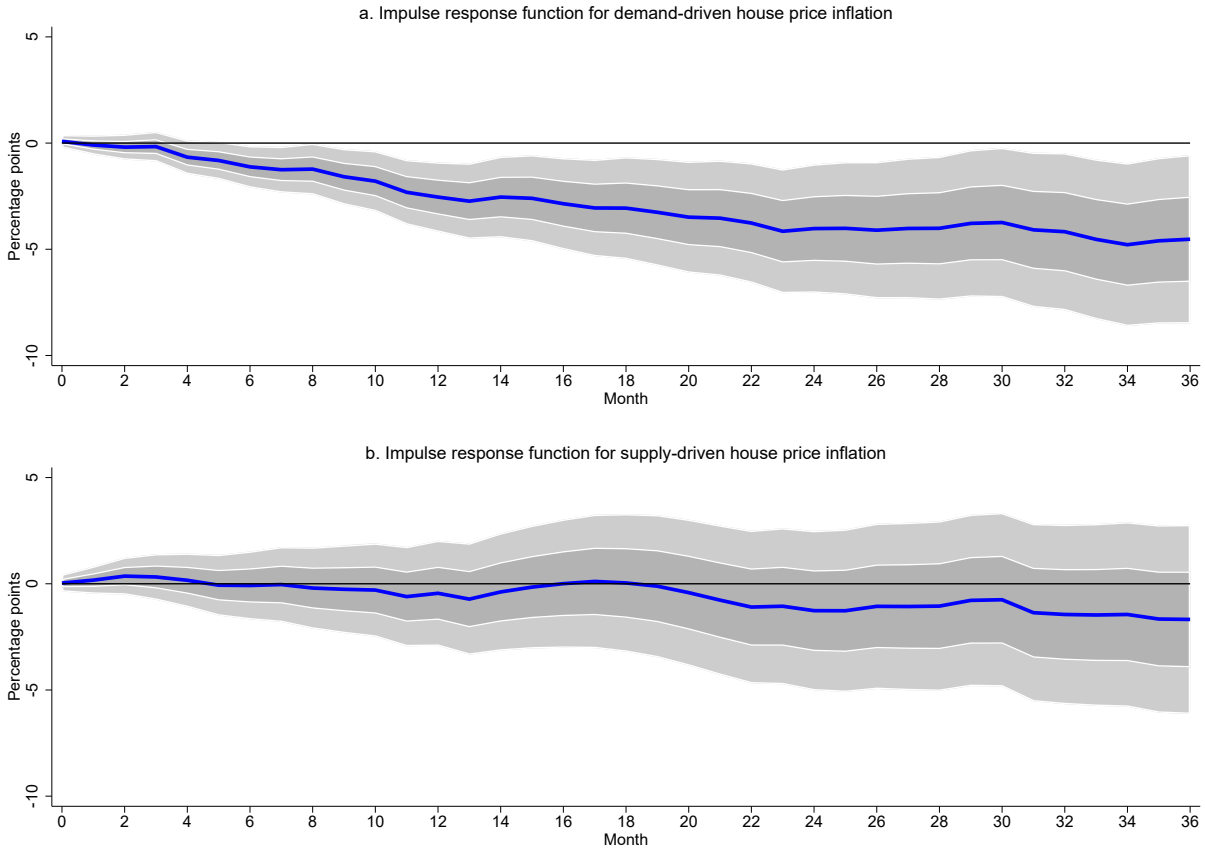


Figure A6: New housing market, land price component (NHPI, under-construction quantity): impulse responses under the sign-only classification (threshold zero, 10-year window, 12-month lag).

Note: See the note to Figure A3. The decomposition uses the land NHPI price component and the under-construction quantity measure.

B Search-and-Matching Rationale for the Sign Restriction

The interpretation of demand and supply as flows of participants, and of observed quantities as equilibrium outcomes of a matching process, corresponds to a well-developed body of theory on housing search and matching. We sketch the relevant environment here to clarify the identifying restriction used in the empirical decomposition; we do not estimate the structural model.

Environment. Consider the resale market for existing homes, following the housing matching tradition of Wheaton (1990) and its dynamic extensions (Krainer, 2001; Díaz and Jerez, 2013; Head et al., 2014; Ngai and Sheedy, 2020). At each date, the market is described by a stock of listings S_t—houses actively offered for sale—and a flow of searching buyers B_t. Meetings between buyers and listings are generated by a constant-returns-to-scale matching function

$$M_t = M(B_t, S_t), \qquad M_B > 0, \quad M_S > 0, \tag{11}$$

so that the number of transactions is increasing in both the mass of buyers and the mass of listings. Market *tightness*

$$\theta_t \equiv \frac{B_t}{S_t} \tag{12}$$

summarises the relative scarcity of the two sides. Because M has constant returns, the transaction probabilities on each side depend on θ_t alone. Writing $q(\theta_t) \equiv M(B_t, S_t)/S_t = M(\theta_t, 1)$ for the matches generated per listing, the seller's selling probability and the buyer's purchase probability are

$$\pi^s(\theta_t) = q(\theta_t), \qquad \pi^b(\theta_t) = \frac{q(\theta_t)}{\theta_t}, \tag{13}$$

with $\pi^{s\prime}(\theta_t) > 0$ and $\pi^{b\prime}(\theta_t) < 0$: a tighter market raises a seller's chance of selling—equivalently, it shortens expected time-on-market, the standard empirical measure of liquidity (Krainer, 2001)—while lowering a buyer's chance of purchasing. The elasticity of matches per listing with respect to tightness,[19]

$$\varepsilon_q \equiv \frac{q'(\theta_t)\,\theta_t}{q(\theta_t)} \in (0,1), \tag{14}$$

lies strictly between zero and one under constant returns. The number of completed transactions is

$$T_t = M(B_t, S_t) = S_t\, q(\theta_t). \tag{15}$$

When a buyer and a seller meet, the transaction price p_t is determined by bargaining over the match surplus and therefore lies between the buyer's and the seller's reservation values. Because a tighter market improves the seller's outside option relative to the buyer's (Wheaton, 1990; Head et al., 2014), the bargained price is increasing in tightness,

$$p_t = p(\theta_t), \qquad \eta_p \equiv \frac{p'(\theta_t)\,\theta_t}{p(\theta_t)} > 0. \tag{16}$$

Participation and the two flows. Both sides of the market participate endogenously. Let J_t^b and J_t^s denote the expected surplus a buyer and a seller obtain from a completed match, and let κ_b and κ_s denote the per-period costs of searching and of listing. Buyers enter until the expected value of search is driven to zero,

$$\pi^b(\theta_t)\, J_t^b = \kappa_b, \tag{17}$$

[19] The bounds $\varepsilon_q \in (0,1)$ follow from constant returns to scale. Since $q(\theta_t) = M(\theta_t, 1)$, we have $q'(\theta_t) = M_B(\theta_t, 1) > 0$, so $\varepsilon_q > 0$. Euler's theorem for the degree-one homogeneous matching function gives $M = M_B B + M_S S$, and dividing by S yields $q(\theta_t) = M_B(\theta_t, 1)\,\theta_t + M_S(\theta_t, 1)$. Hence $\varepsilon_q = M_B\theta_t/(M_B\theta_t + M_S) < 1$, with the gap from one equal to the share of matches attributable to the listing margin. This is equivalent to the buyer's purchase probability $\pi^b(\theta_t) = q(\theta_t)/\theta_t$ being decreasing in tightness.

which pins down the buyer inflow B_t and makes it respond to credit conditions, borrowing costs, and expectations about future prices through J_t^b (Ortalo-Magné and Rady, 2006). Sellers list whenever the value of selling now exceeds the option value of waiting,

$$\pi^s(\theta_t)\, J_t^s \;\geq\; \kappa_s, \tag{18}$$

so that listing decisions respond to equity positions, prior purchase prices, and expected liquidity through J_t^s; this can make effective supply sticky or even backward-bending in downturns (Genesove and Mayer, 2001; Ngai and Sheedy, 2020). In this environment, "demand" is naturally identified with the buyer inflow B_t and "supply" with the seller/listing inflow S_t.

Why the sign restriction holds in equilibrium. The comparative statics of this market deliver the identifying restriction directly. Log-differentiating the tightness definition (12), the price schedule (16), and the transaction identity (15), and using the elasticity (14), gives

$$\mathrm{d}\ln\theta_t = \mathrm{d}\ln B_t - \mathrm{d}\ln S_t, \tag{19}$$

$$\mathrm{d}\ln p_t = \eta_p\, \mathrm{d}\ln\theta_t, \tag{20}$$

$$\mathrm{d}\ln T_t = \mathrm{d}\ln S_t + \varepsilon_q\, \mathrm{d}\ln\theta_t. \tag{21}$$

Consider a positive *demand* shock that raises the buyer inflow, $\mathrm{d}\ln B_t > 0$, holding listings fixed ($\mathrm{d}\ln S_t = 0$). From (19) tightness rises, $\mathrm{d}\ln\theta_t = \mathrm{d}\ln B_t > 0$, so by (20)–(21)

$$\mathrm{d}\ln p_t = \eta_p\, \mathrm{d}\ln B_t > 0, \qquad \mathrm{d}\ln T_t = \varepsilon_q\, \mathrm{d}\ln B_t > 0 : \tag{22}$$

price and quantity move in the *same* direction. Now consider a negative *supply* shock that withdraws listings, $\mathrm{d}\ln S_t < 0$, holding buyer entry fixed ($\mathrm{d}\ln B_t = 0$). Then $\mathrm{d}\ln\theta_t =$

$-\,\mathrm{d}\ln S_t > 0$, and substituting into (20)–(21),

$$\mathrm{d}\ln p_t = -\,\eta_p\,\mathrm{d}\ln S_t > 0, \qquad \mathrm{d}\ln T_t = (1-\varepsilon_q)\,\mathrm{d}\ln S_t < 0, \tag{23}$$

where the sign of $\mathrm{d}\ln T_t$ follows because $\varepsilon_q \in (0,1)$ by (14): the direct loss of listings dominates the offsetting rise in matches per listing, so price and quantity move in *opposite* directions. Equations (22)–(23) are precisely the sign patterns that Shapiro (2026) imposes to classify demand- and supply-driven movements. The requirement $\varepsilon_q < 1$—that the matching technology exhibit diminishing returns to tightness, as any constant-returns matching function does—is exactly what guarantees that a supply-driven price increase is accompanied by falling transactions.

New housing and construction. In the new housing market, the same logic applies once the seller side is reinterpreted as developer behaviour. Head et al. (2014) embed construction in a search model in which developer entry, the flow of new units, buyer entry, prices, and sales are jointly determined in equilibrium; the listing margin S_t in (18) is then replaced by a construction margin, and supply adjusts through starts, units under construction, and completions rather than through listings of the existing stock. The elasticity with which demand shocks translate into prices versus quantities—governed in (22) by η_p relative to ε_q—depends, as in the reduced-form literature, on how readily this construction margin responds, which is in turn shaped by geography and regulation (Saiz, 2010; Glaeser and Gyourko, 2018).

C Data Sources and Regions

This appendix documents the data underlying the analysis. Table C1 lists every series used—house prices and quantities for both market segments, the monetary policy shock, and the macroeconomic controls that enter the local projections—together with its coverage period, type, source institution, retrieval provider, and series code. Tables C2 and C3 then list the regions in the two regional panels underlying the first-stage classification: the granular CREA resale-market panel used to classify demand- and supply-driven inflation in the resale market (Section 4.3.1), and the smaller, census-metropolitan-area (CMA) NHPI/CMHC panel used for the new housing market (Section 4.3.2). The two panels differ in coverage because the CREA series provide matched prices and quantities for a broad set of local markets, whereas the NHPI/CMHC panel is restricted to CMAs that can also be matched to CREA prices for the index-to-level conversion.

CREA resale-market regions. The CREA data are reported at the level of real estate board regions and are aggregated within each market, province, and the nation. We depart from these aggregates in two places to maximize the number of usable regions. First, we drop the two largest overlapping aggregates—Greater Toronto and Edmonton Board Total—in favour of their component regions: Edmonton Board Total is replaced by its components Edmonton and North Eastern Alberta, and Greater Toronto (which subsumes Durham, Mississauga, Orangeville, and York Region) is dropped in favour of those components. Second, to retain a Toronto observation we construct a synthetic Toronto region by subtracting the component regions from the Greater Toronto total unit-sales and total-dollar-volume series, and then form an average price for Toronto as the ratio of the resulting total dollar volume to total units sold.

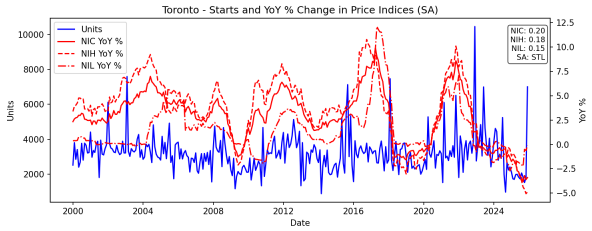
NHPI/CMHC new-housing regions. The NHPI and CMHC series are reported for census metropolitan areas (CMAs). Because the NHPI is a price index, we convert it to level space by scaling the index to the CREA average sales price in December 2016; this restricts the new-housing panel to regions found in both the CREA data and the CMAs. As with the CREA data, a small number of synthetic regions are constructed to align the two sources. To increase the number of CMA-like regions within the CREA data, St. Catharines–Niagara is formed by aggregating St. Catharines, Niagara Falls–Fort Erie, and Welland; Saint John/Fredericton/Moncton by combining the Saint John, Fredericton, and Moncton markets; and Kitchener-Cambridge-Waterloo from Kitchener–Waterloo and Cambridge. Two further synthetic regions are constructed from the NHPI and CMHC data to match CREA aggregates: a Saint John/Fredericton/Moncton aggregate combining the Saint John and Moncton regions, and a province-level Quebec aggregate combining Montreal, Quebec City, Sherbrooke, Trois-Rivières, Saguenay, and Gatineau.

Table C1: Data sources, coverage, and series codes for all price, quantity, monetary-policy-shock, and control variables used in the analysis

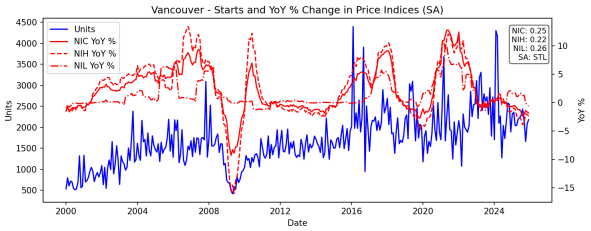
Variable	Period used in analysis	Series definition / unit	Producing institution	Retrieval provider	Retrieval code / official identifier
Average House Price	Jan 1988 – Dec 2025	Residential resale average price, Canadian dollars	CREA	HAVER Analytics	CACERPSM@CREA
Unit Sales	Jan 1988 – Dec 2025	Residential resale sales, number of units	CREA	HAVER Analytics	CACERUSM@CREA
New Housing Price Index: House and Land	Jan 1981 – Dec 2025	NHPI, house and land, index	Statistics Canada	HAVER Analytics	V1U55442@CANADAR
New Housing Price Index: House Only	Jan 1981 – Dec 2025	NHPI, house only, index	Statistics Canada	HAVER Analytics	V1U55443@CANADAR
New Housing Price Index: Land Only	Jan 1981 – Dec 2025	NHPI, land only, index	Statistics Canada	HAVER Analytics	V1U55444@CANADAR
Housing Starts	Jan 1990 – Dec 2025	Housing starts, units, seasonally adjusted annual rate (SAAR)	CMHC	HAVER Analytics	GS00001@CANADAR
Housing Completions	Jan 1981 – Dec 2025	Housing completions, all dwelling types, number of units	CMHC	HAVER Analytics	BR0011@CANADAR
Units Under Construction	Jan 1981 – Dec 2025	Housing units under construction, all dwelling types, number of units	CMHC	HAVER Analytics	BD0014@CANADAR
Monetary Policy Shock	Jun 1997 – Dec 2025	One-quarter-ahead policy-rate announcement surprise, percentage points	Sekkel et al. (2025)	See Sekkel et al. (2025)	BAX1/CORRA
Commodity Price Index	Jan 1997 – Jan 2026	Total Bank of Canada Commodity Price Index, Fisher index, converted from US-dollar to Canadian-dollar terms using the USD/CAD exchange rate	Bank of Canada	Bank of Canada	M.BCPI; legacy vector v52673496
USD/CAD Exchange Rate	Jan 1997 – Jan 2026	Monthly average exchange rate, Canadian dollars per US dollar	Bank of Canada	Bank of Canada	IEXM0102_AVG (legacy, pre-2017); FXMUSDCAD (2017 onward)
Gross Domestic Product	Jan 1997 – Jan 2026	Real GDP at basic prices, all industries, seasonally adjusted, chained 2017 dollars	Statistics Canada	Statistics Canada	Table 36-10-0434-01, All industries
Consumer Price Index	Jan 1997 – Jan 2026	All-items CPI, seasonally adjusted, 2002=100	Statistics Canada	Statistics Canada	Table 18-10-0006-01, All-items
Bank Rate	Jan 1997 – Jan 2026	Bank Rate, monthly, per cent	Bank of Canada	HAVER Analytics	CIRD@CANADA; underlying official series V39078
Five-year rate	Jan 1997 – Jan 2026	Government of Canada benchmark bond yield, 5 year, monthly, per cent	Bank of Canada	Bank of Canada	V122540

Note: The table lists every raw series entering the analysis, separating the institution that produces the series from the provider through which it was retrieved and the retrieval code. The first block reports the price and quantity series used in the first-stage demand–supply classification (CREA average prices and unit sales for the resale market; the combined, structure, and land NHPI components and the CMHC construction measures for the new housing market); the remaining rows report the monetary policy shock and the macroeconomic controls entering the local projections. In the local projections, GDP, the commodity price index, and CPI enter as 100 times log levels with lags 1–6; the Bank Rate and five-year rate enter in percentage points with lags 1–6. The monetary policy shock splices the published one-quarter-ahead BAX1 surprise of Sekkel et al. (2025) for 1997–2023 with an author-supplied CORRA-based extension for 2024–25 (Section 2.1). The Bank of Canada Commodity Price Index is published in US-dollar terms and is converted to Canadian-dollar terms using the monthly USD/CAD exchange rate (Canadian dollars per US dollar). StatCan labels the three NHPI price series "House and Land," "House Only," and "Land Only"; we keep these official series names in the table but, throughout the rest of the paper, refer to the "House Only" component as the *structure* component and the "House and Land" aggregate as the *combined* (structure-and-land) component. We adopt this terminology because elsewhere we use "house price" to mean the price of the entire dwelling (structure and land together), so naming the building component "structure" avoids ambiguity when the structure and land components are discussed separately.

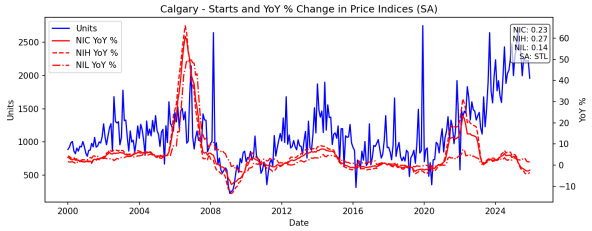
Figure C1: Price and quantity series for new housing across a selection of regions: housing starts (NHPI/CMHC).



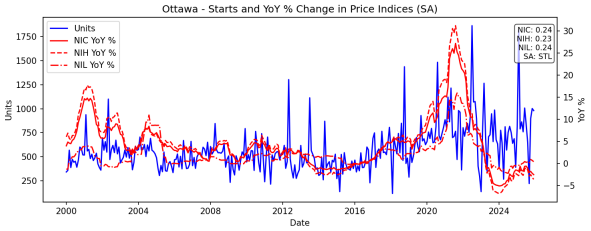
(a) Toronto, housing starts



(b) Vancouver, housing starts



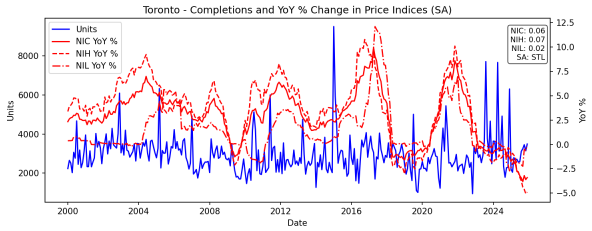
(c) Calgary, housing starts



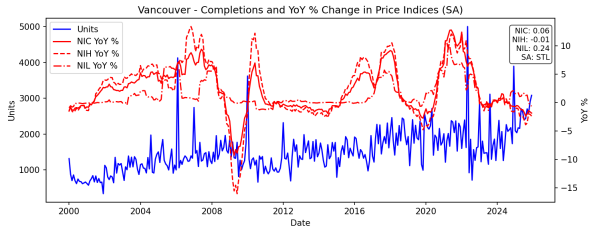
(d) Ottawa, housing starts

Note: Each panel plots, for the indicated region, housing starts (left axis, in units) against the year-over-year percentage change in the New Housing Price Index and its components (right axis): the combined structure-and-land index (NIC), the structure component (NIH), and the land component (NIL). All series are seasonally adjusted, and each price series is expressed as the year-over-year change of the seasonally adjusted price level. The figures reported in the upper-right corner of each panel are the correlations between the seasonally adjusted quantity series and each of the three seasonally adjusted NHPI components (NIC, NIH, and NIL). Price data are from the Statistics Canada NHPI and quantity data are from CMHC, retrieved via Haver Analytics; authors' calculations.

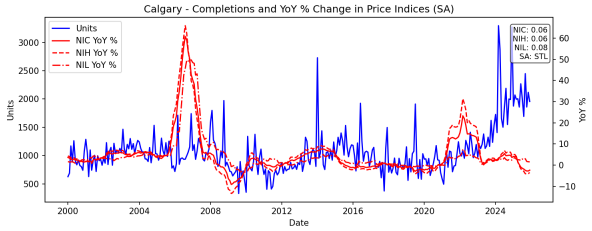
Figure C2: Price and quantity series for new housing across a selection of regions: completions (NHPI/CMHC).



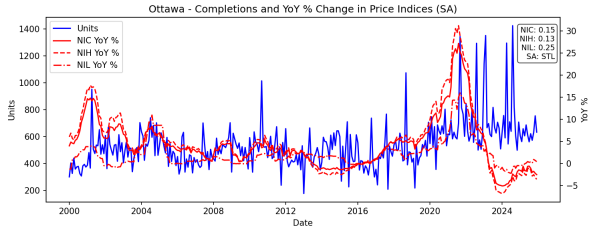
(a) Toronto, completions



(b) Vancouver, completions



(c) Calgary, completions



(d) Ottawa, completions

Note: Each panel plots, for the indicated region, completions (left axis, in units) against the year-over-year percentage change in the New Housing Price Index and its components (right axis): the combined structure-and-land index (NIC), the structure component (NIH), and the land component (NIL). All series are seasonally adjusted, and each price series is expressed as the year-over-year change of the seasonally adjusted price level. The figures reported in the upper-right corner of each panel are the correlations between the seasonally adjusted quantity series and each of the three seasonally adjusted NHPI components (NIC, NIH, and NIL). Price data are from the Statistics Canada NHPI and quantity data are from CMHC, retrieved via Haver Analytics; authors' calculations.

Table C2: CREA resale-market regions used in the first-stage classification

Alberta West	Grey Bruce Owen Sound	North Bay	Saint John
Annapolis Valley	Guelph	North Eastern Alberta	Sarnia-Lambton
Bancroft	Halifax-Dartmouth	Northern Lights	Saskatoon
Barrie	Hamilton-Burlington	Northern New Brunswick	Sault Ste Marie
Battlefords	Highland	Northern Nova Scotia	Simcoe
BC Northern	Huron Perth	Northumberland Hills	South Central Alberta
Brandon	Kamloops	Oakville-Milton	South Okanagan
Brantford	Kawartha Lakes	Okanagan-Mainline	South Shore
Calgary	Kingston	Orangeville	Thunder Bay
Cambridge	Kitchener-Waterloo	Orillia	Tillsonburg
Cape Breton	Kootenay	Ottawa	Timmins
Central Alberta	Lethbridge	Parry Sound	Toronto
Chatham Kent	Lloydminster [AB]	Peterborough & the Kawarthas	Vancouver Island
Chilliwack	Lloydminster [SK]	Portage La Prairie	Victoria
Cornwall	London and St Thomas	Powell River	Welland
Durham Region	Medicine Hat	Prince Albert	Windsor-Essex
Edmonton	Mississauga	Prince Edward Island	Winnipeg
Fort McMurray	Moncton	Quinte	Woodstock-Ingersoll
Fraser Valley	Moose Jaw	Regina	Yarmouth
Fredericton	Muskoka & Haliburton	Renfrew	York Region
Grande Prairie	Newfoundland & Labrador	Rideau-St Lawrence	Yorkton
Greater Vancouver	Niagara Falls-Fort Erie		

Note: Regions in the CREA resale-market panel, for which matched average sales prices and unit sales are observed at monthly frequency. These are the regions over which demand- and supply-driven inflation is classified in the first stage for the resale market (Section 4.3.1). A few entries are synthetic regions constructed by aggregating adjacent markets to align with census metropolitan areas, as described in the main text.

Table C3: NHPI/CMHC new-housing regions used in the first-stage classification

Calgary
Edmonton
Greater Sudbury
Guelph
Halifax
Hamilton
Kitchener
London
Ottawa
Quebec
Regina
Saint John/Fredericton/Moncton
Saskatoon
St. Catharines-Niagara
Toronto
Vancouver
Victoria
Windsor
Winnipeg

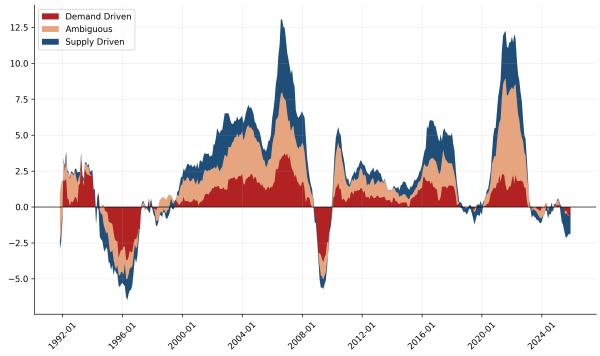
Note: Census metropolitan area (CMA) regions in the NHPI/CMHC new-housing panel, for which NHPI prices (combined, structure, and land) and CMHC construction quantities (housing starts, units under construction, and completions) are observed. These are the regions over which demand- and supply-driven inflation is classified in the first stage for the new housing market (Section 4.3.2). The panel is restricted to CMAs that can also be matched to CREA data, because the NHPI is converted from index to level space using CREA prices; some entries are synthetic aggregates, as described in the main text.

D Decomposition of the Structure and Land Price Components

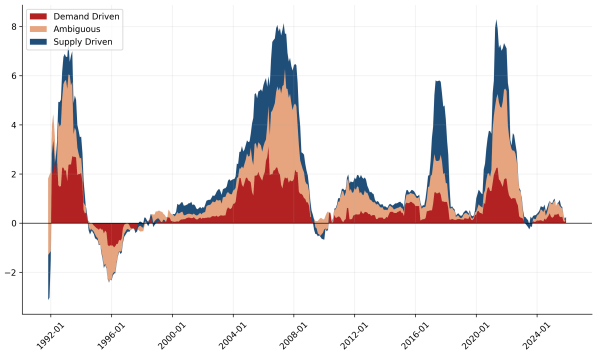
Figure 5 in the main text decomposes the *combined* NHPI price series (structure and land together). This appendix reports the analogous Shapiro-style decompositions for the two NHPI subcomponents separately—the structure component and the land component—under each of the three quantity measures (housing starts, units under construction, and completions). Figure D1 arranges the six decompositions with the structure component in the left column and the land component in the right column, and with the quantity measure varying down the rows from housing starts to units under construction to completions.

Two patterns stand out, and both are consistent with the local projection results in Section 4.3.2. First, the structure component displays a substantial incidence of supply-driven periods across all three quantity measures, and especially at the completions stage, where supply-driven episodes become more frequent and pronounced. Second, the land component is more heavily demand-driven at the starts and under-construction stages, mirroring the resale market, but exhibits a markedly larger role for supply once the quantity measure is completions. Together these panels show that the aggregate contrast between the structure and land components documented in the main text is visible directly in the underlying classification, and that the stage of the construction pipeline at which quantities are measured matters for the land component in particular. The historical incidence of supply-classified observations is conceptually distinct from the estimated monetary-policy response of the supply-classified component; the former describes how often and how strongly such observations contribute to inflation, while the latter is estimated in the second-stage local projections.

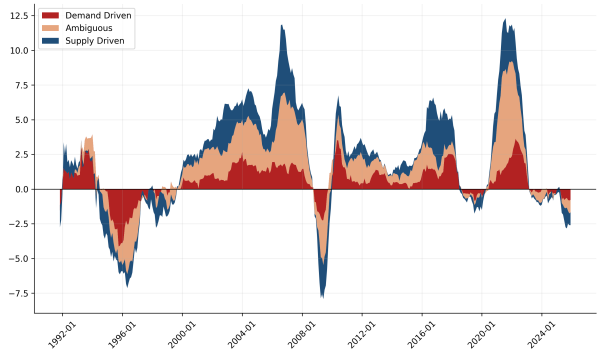
Figure D1: Shapiro-style decomposition of new-housing price inflation into demand-driven, supply-driven, and ambiguous components: structure and land NHPI price components (NHPI/CMHC).



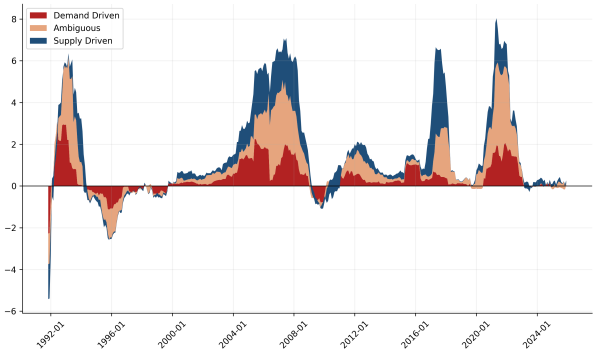
(a) Structure, housing starts



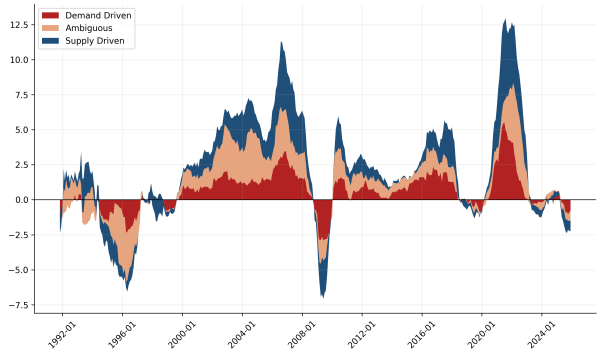
(b) Land, housing starts



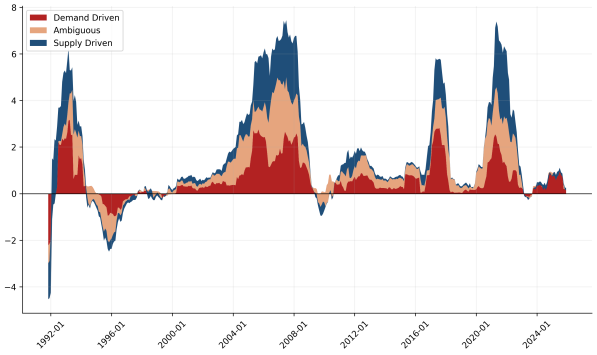
(c) Structure, units under construction



(d) Land, units under construction



(e) Structure, completions



(f) Land, completions

Note: Each panel decomposes year-over-year new-housing price inflation into demand-driven, supply-driven, and ambiguous components following Shapiro (2026), using the baseline first-stage specification (10-year rolling window, 12-month lag). The left column uses the structure NHPI price component and the right column uses the land NHPI price component; the quantity measure varies down the rows as housing starts, units under construction, and completions. Classification is based on the co-movement of the price and quantity innovations: an observation is labelled *demand-driven* when the two innovations have the same sign and *supply-driven* when they have opposite signs. Both innovations must exceed a minimum magnitude of 0.3 standard deviations in absolute value; if either falls below this threshold the observation is labelled *ambiguous*, regardless of sign.