economies

IMPACT FACTOR 2.1

CITESCORE 4.7

Article

One Policy Rate, Uneven Provincial Inflation: Shelter, Household Debt, and Provincial Structure in Canada

Constantin Colonescu

Special Issue
Monetary Policy and Inflation Dynamics

Edited by
Dr. Alessandro Piergallini

MDPI

https://doi.org/10.3390/economies14050180

economies

MDPI

Article

One Policy Rate, Uneven Provincial Inflation: Shelter, Household Debt, and Provincial Structure in Canada

Constantin Colonescu

Department of Economics, MacEwan University, Edmonton, AB T5J 4S2 Canada; colonescuc@macewan.ca

Abstract

This article studies why the same Bank of Canada tightening is reflected differently in provincial CPI inflation. It combines monthly provincial data from January 1991 to December 2024 with interacted local projections and public-data measures of common national monetary movements. The design estimates reduced-form provincial loadings in a common monetary environment, rather than structural responses to a single externally identified surprise. The main result is a housing-sensitive gap between headline inflation and inflation excluding shelter. Provinces with larger shelter weights and higher household debt–service exposure show a stronger headline response than non-shelter response after a common tightening. The evidence does not reduce to rent or to basket arithmetic alone: debt–service exposure is the more stable standalone component, while shelter weights tie the differential to measured CPI. Outside shelter, no single provincial characteristic dominates. Internal trade integration is associated with smaller baseline deviations from the national non-shelter response, but energy-related provincial composition is at least as informative in the competing-factor specifications. The paper therefore identifies shelter and household debt as the clearest sources of provincial incidence under one policy rate, while treating non-housing deviations from the national response as a broader provincial-structure result.

Keywords: monetary policy transmission; provincial inflation; shelter; household debt; Canada; local projections

JEL Classification: E31; E52; R12; R21

1. Introduction

A monetary union places many local economies under one monetary policy. Canada is a useful setting in which to study what that means in practice. It has one central bank, one policy rate, one inflation target, and a highly integrated internal market. At the same time, the provinces differ in housing markets, household debt burdens, industrial structure, commodity exposure, labour-market conditions, and the weights of different categories in the CPI basket. Those differences matter because a common change in the policy rate does not pass through every province in the same way. The same national monetary environment can affect local housing costs, household debt obligations, spending patterns, and measured inflation differently across provinces.

This article asks whether a small number of persistent provincial characteristics help explain those different inflation responses. The main comparison is between headline inflation and inflation excluding shelter. The strongest pattern in the data is tied to a joint shelter-and-household-debt measure that combines the shelter weight in the CPI with

Academic Editor: Alessandro Piergallini

Received: 27 March 2026
Revised: 5 May 2026
Accepted: 8 May 2026
Published: 14 May 2026

household debt–service exposure. Provinces higher on this measure show a larger response of headline inflation than of inflation excluding shelter after a common monetary tightening.

The paper also asks whether provincial inflation excluding shelter stays closer to, or farther from, the Canadian response after the same common movement. The evidence here is more mixed. In the baseline specifications, provinces with stronger internal trade links tend to show smaller deviations from the national non-shelter inflation response. Once other provincial characteristics are considered directly, however, energy-related provincial structure is often at least as important. The paper therefore does not argue that one provincial characteristic explains every outcome. Its clearest result concerns the housing-sensitive difference between headline inflation and inflation excluding shelter. The non-housing results point to a broader set of provincial characteristics.

The paper studies provincial inflation responses under common monetary conditions observed in the data. Its identification claim is disciplined and reduced form. On the theory side, the paper develops a simple regional framework that produces testable differences across provinces. On the empirical side, the strongest contribution is on the housing side: the same joint shelter-and-household-debt measure matters across several monetary-policy measures, remains informative when other provincial characteristics are introduced, and survives inference procedures built around the common time variation in the national monetary environment. The non-housing results are still useful but they are less settled and should be read at that more limited level.

This article also contributes to the Canadian literature on monetary transmission, housing, and inflation measurement. Champagne and Sekkel (2018) show that identified Canadian monetary-policy disturbances depend on the monetary regime, which is one reason this paper uses a family of public-data monetary measures. Bounajm and McWhirter (2024) show that mortgage debt and payment structures are central to how higher interest rates affect Canadian households. Statistics Canada documents both how the shelter component enters the CPI and how important shelter-related items have become in the aggregate basket (Statistics Canada, 2023, 2025). Together, these contributions support the paper's focus on a joint shelter-and-household-debt channel rather than on one narrowly defined shelter subcomponent.

Three features of the empirical design are especially important. First, the housing measure is broken down into its main parts rather than being treated as a black box. The evidence suggests that household debt–service exposure is the more stable source of the difference between headline inflation and inflation excluding shelter, while the shelter-weight component keeps that difference tied to the construction of measured CPI. Second, the non-housing results are interpreted at the level the data support. Interprovincial trade share is not treated as a pure structural measure of economic integration; it is read as one provincial characteristic among several related features of internal economic structure. Third, inference is part of the main result rather than an afterthought. The key 12-month interaction estimates are persuasive only if their sign and relative ordering survive both month-clustered standard errors and a time-cluster wild bootstrap.

The sample runs from January 1991 to December 2024, which broadly covers the inflation-targeting period and includes the post-2020 tightening. The coefficients should therefore be read as average reduced-form patterns over a long sample that contains a prominent recent tightening episode, rather than as parameters from one fixed regime or as evidence confined to that episode alone.

The paper proceeds as follows. Section 2 reviews the related literature and states the article's contribution. Section 3 presents a simple regional framework and links its parameters to the empirical proxies. Section 4 describes the data, the monetary-policy measures, the baseline local projections, the competing-factor design, and the inferential

approach. Section 5 presents the empirical results, beginning with the baseline housing and trade measures and then turning to decomposition of the housing measure, competing-factor specifications, lag sensitivity, robustness across monetary measures, and influence diagnostics. Section 6 interprets the findings and clarifies the scope of the contribution. Section 7 concludes.

2. Related Literature and Contribution

Three strands of literature are especially relevant. The first studies regional heterogeneity in monetary transmission within a common currency or policy regime. Beraja et al. (2019) show that housing finance and household balance sheets can generate materially different local responses to the same national monetary movement. Their work motivates the paper's focus on shelter intensity and debt–service exposure as provincial margins along which a common tightening can be transmitted unevenly.

The second work concerns identification and interpretation in macroeconomics and policy evaluation. Nakamura and Steinsson (2018) emphasize that macroeconomic identification depends on the informational content, timing, and institutional setting of the monetary measure itself. In the Canadian context, Champagne and Sekkel (2018) show that the identified monetary disturbance depends on the monetary regime. The broader policy-evaluation literature reaches a related conclusion in non-monetary settings: common policy changes can have different effects when they interact with the exposure margins and frictions created by the policy design. For example, Solodoha et al. (2023) use a natural experiment in angel-investment incentives and find that a tax policy intended to encourage seed-stage finance was associated with fewer angels investing in each targeted start-up. The context is entrepreneurial finance rather than monetary transmission, but the incidental lesson is useful here: a common policy instrument need not generate uniform outcomes across exposed units, and empirical work should study heterogeneous responses directly rather than infer them mechanically from the policy lever. Those points matter here because the article deliberately does not rest on one unquestioned external surprise series. Instead, it asks whether the ranking of provincial responses is stable across a family of public-data monetary measures and interprets the resulting coefficients as reduced-form evidence on heterogeneity under observed common monetary conditions.

The third work concerns the channels through which housing enters inflation and household adjustment in Canada. Statistics Canada documents both the composition of the shelter block in the CPI and the growing importance of shelter-related items in the aggregate basket (Statistics Canada, 2023, 2025). Bounajm and McWhirter (2024) show that mortgage debt and payment structures are central to how higher interest rates affect Canadian households. Those contributions make it natural to study shelter basket weights and debt–service exposure jointly rather than as unrelated provincial descriptors.

Against that background, the article makes three contributions. Economically, it shows that the clearest provincial difference after a common tightening is the gap between headline inflation and inflation excluding shelter, and that this gap is systematically related to shelter intensity and household debt exposure. Quantitatively, it studies that question with interacted local projections, province-specific trends, time-dimension inference, and a proxy-family robustness design that separates broad provincial rankings from dependence on one monetary measure. Conceptually, it links a tractable regional framework to empirical restrictions that distinguish a housing-sensitive headline wedge from broader non-shelter deviations from the national response. The result is a disciplined reduced-form account of how provincial exposure profiles shape measured inflation responses inside a common monetary regime.

3. A Tractable Regional Framework

The theoretical purpose of the framework is not to solve a fully specified Canadian multi-region equilibrium model: it is to organize the empirical interpretation. Consider a monetary union with provinces indexed by $i = 1, \ldots, N$. All provinces share the same nominal policy rate i_t, but local inflation is composed of two blocks. The first is inflation excluding shelter, denoted $\pi^c_{i,t}$. The second is shelter inflation, denoted $\pi^s_{i,t}$. A table of symbols is provided in Appendix A. Headline inflation is a weighted average of the two components,

$$\pi_{i,t} = (1 - \omega_i)\pi^c_{i,t} + \omega_i \pi^s_{i,t}, \qquad 0 < \omega_i < 1, \tag{1}$$

where ω_i is the provincial shelter share.

Local demand obeys

$$\begin{aligned} x_{i,t} = {} & E_t x_{i,t+1} - \sigma_i (i_t - E_t \pi_{i,t+1} - r^n_t) \\ & - \chi_i b_{i,t-1}(i_t - i_{t-1}) + \eta_i \sum_j w_{ij}(x_{j,t} - x_{i,t}) + u^x_{i,t}, \end{aligned} \tag{2}$$

where $b_{i,t-1}$ indexes debt–service exposure and η_i indexes the strength of internal integration. A higher $b_{i,t-1}$ implies that a policy-rate increase generates a stronger cash-flow squeeze. A higher η_i implies that expenditure switching and production linkages tie province i more tightly to the rest of the union.

Inflation excluding shelter follows

$$\pi^c_{i,t} = \beta E_t \pi^c_{i,t+1} + \kappa_i x_{i,t} + u^c_{i,t}, \tag{3}$$

while shelter inflation obeys

$$\pi^s_{i,t} = \beta E_t \pi^s_{i,t+1} + \kappa^s_i x_{i,t} + \lambda_i \zeta_i (i_t - i_{t-1}) + \xi_i \Delta q_{i,t} + u^s_{i,t}. \tag{4}$$

The term $\lambda_i \zeta_i (i_t - i_{t-1})$ is the financing-sensitive component. Mortgage interest cost is one empirically visible candidate for this term, but the model does not require the entire financing-sensitive response to be concentrated in one subcomponent.

Subtracting inflation excluding shelter from headline inflation yields the article's central differential object,

$$\Omega_{i,t} \equiv \pi_{i,t} - \pi^c_{i,t} = \omega_i \left(\pi^s_{i,t} - \pi^c_{i,t} \right). \tag{5}$$

The wedge is large when shelter carries a large weight and when shelter inflation moves differently from non-shelter inflation because the shelter block contains a direct financing term.

Three implications are especially relevant.

Proposition 1. *Let $m_t \equiv i_t - i_{t-1}$ denote the common monetary movement, with $m_t > 0$ corresponding to a tightening, and let $\Omega_{i,t} \equiv \pi_{i,t} - \pi^c_{i,t}$ denote the gap between headline inflation and inflation excluding shelter. Then*

$$\frac{\partial \Omega_{i,t}}{\partial m_t} = \omega_i \frac{\partial (\pi^s_{i,t} - \pi^c_{i,t})}{\partial m_t}. \tag{6}$$

Hence, whenever a tightening raises shelter inflation relative to inflation excluding shelter, that is, whenever

$$\frac{\partial (\pi^s_{i,t} - \pi^c_{i,t})}{\partial m_t} > 0,$$

a larger shelter weight increases the response of the headline inflation relative to the response of inflation excluding shelter:

$$\frac{\partial^2 \Omega_{i,t}}{\partial \omega_i \partial m_t} = \frac{\partial(\pi^s_{i,t} - \pi^c_{i,t})}{\partial m_t} > 0. \tag{7}$$

Moreover, on impact, holding fixed expected future inflation differentials, local demand, and local shelter-price changes, stronger financing-sensitive shelter pass-through also increases that differential response:

$$\frac{\partial^2 \Omega_{i,t}}{\partial(\lambda_i \zeta_i)\partial m_t} = \omega_i > 0. \tag{8}$$

Proof. From Equation (1), subtracting $\pi^c_{i,t}$ from both sides gives Equation (5). Subtracting Equation (3) from Equation (4) gives

$$\begin{aligned} \pi^s_{i,t} - \pi^c_{i,t} = \beta E_t(\pi^s_{i,t+1} - \pi^c_{i,t+1}) + (\kappa^s_i - \kappa_i)x_{i,t} \\ + \lambda_i \zeta_i m_t + \xi_i \Delta q_{i,t} + (u^s_{i,t} - u^c_{i,t}). \end{aligned} \tag{9}$$

Differentiating (5) with respect to m_t gives Equation (6). Differentiating once more with respect to ω_i yields (7). Therefore, whenever a tightening increases shelter inflation relative to inflation excluding shelter, a larger shelter weight amplifies the response of headline inflation relative to the response of inflation excluding shelter.

To show the role of financing-sensitive shelter pass-through more explicitly, differentiate (9) with respect to m_t. In general,

$$\begin{aligned} \frac{\partial(\pi^s_{i,t} - \pi^c_{i,t})}{\partial m_t} = \beta \frac{\partial E_t(\pi^s_{i,t+1} - \pi^c_{i,t+1})}{\partial m_t} + (\kappa^s_i - \kappa_i)\frac{\partial x_{i,t}}{\partial m_t} \\ + \lambda_i \zeta_i + \xi_i \frac{\partial \Delta q_{i,t}}{\partial m_t}. \end{aligned}$$

On impact, holding fixed expected future inflation differentials, local demand, and local shelter-price changes, this simplifies to

$$\frac{\partial(\pi^s_{i,t} - \pi^c_{i,t})}{\partial m_t} = \lambda_i \zeta_i.$$

Substituting into (5) gives

$$\frac{\partial \Omega_{i,t}}{\partial m_t} = \omega_i \lambda_i \zeta_i,$$

and therefore

$$\frac{\partial^2 \Omega_{i,t}}{\partial(\lambda_i \zeta_i)\partial m_t} = \omega_i > 0.$$

So, on impact, stronger financing-sensitive shelter pass-through increases the response of headline inflation relative to the response of inflation excluding shelter. □

Proposition 2. *When housing exposure combines a shelter-weight channel and a debt–service channel, the sign of the housing interaction in inflation excluding shelter is ambiguous, but the interaction in headline inflation exceeds the interaction in inflation excluding shelter whenever the shelter block contains a direct financing term.*

Proof. A higher shelter weight does not by itself determine the sign of the interaction in $\pi^c_{i,t}$. Debt–service sensitivity in Equation (2) can push local demand and non-shelter inflation downward, while correlated provincial characteristics can move the reduced-form interaction in the opposite direction. What is robust is the differential statement from

Equation (5): the response of headline inflation must exceed the response of inflation excluding shelter by the amount contributed by the housing-sensitive shelter block. □

Proposition 3. *Stronger interprovincial integration weakly reduces the cross-province dispersion of inflation excluding shelter following a common monetary movement, provided the linearized system is stable.*

Proof. The integration term in Equation (2) penalizes province-specific deviations in demand through network spillovers. Since inflation excluding shelter depends on the demand block rather than on the direct financing term in the shelter equation, stronger integration compresses heterogeneity in the response of the non-shelter inflation component. □

The framework is intentionally simple. Its purpose is to guide the empirical interpretation. A successful empirical version of the argument must therefore show more than generic provincial heterogeneity. It must show that housing exposure is associated with a larger headline-versus-ex-shelter differential response, that the broader shelter block rather than a single narrow component matters for that differential, and that internal integration is most visible in deviations of inflation excluding shelter from the national response. It must also show that these patterns are not artifacts of one proxy for the common monetary movement or of one omitted provincial characteristic. In the empirical work below, the non-shelter deviation prediction for trade integration is treated as a candidate explanation that competes directly with broader provincial structure rather than as a claim that trade is the unique predictor by construction.

The mapping from theory to empirics is direct but deliberately reduced form. The provincial shelter share ω_i corresponds to the shelter basket weight in the CPI. The financing-sensitive terms $\chi_i b_{i,t-1}$ and $\lambda_i \zeta_i$ are not separately observed, so the empirical work proxies their joint relevance with the household debt–service ratio and then combines that measure with the shelter weight in the composite housing index. The integration term η_i is not observed as a structural elasticity either, so interprovincial trade share is used as a persistent ranking variable that captures tighter internal-market links. The rival factors in the horse-race—average unemployment, pre-2020 inflation structure, and energy-related CPI weight—are interpreted as competing reduced-form provincial margins rather than as literal estimates of individual model parameters. This is why the empirical contribution is comparative: it asks whether the theoretical rankings survive in the data once plausible provincial alternatives are allowed to compete.

This mapping implies four empirical restrictions. First, the housing index should interact positively with a common tightening in the headline equation. Second, the same interaction should be smaller in inflation excluding shelter so that the headline-minus-ex-shelter wedge is positive. Third, the shelter block should display a stronger housing interaction than rent alone if the result is a broad measured-shelter channel rather than a narrow rent-channel claim. Fourth, the trade-integration ranking should be evaluated on deviations of inflation excluding shelter from the national response and should be allowed to lose explanatory power when richer provincial structure is included. The results section follows that sequence.

4. Data, Common Monetary Proxies, and Empirical Strategy

4.1. Data and Variable Construction

The empirical implementation uses official Canadian public data. Monthly provincial CPI indexes come from Statistics Canada Table 18-10-0004-01. Provincial CPI basket weights come from Table 18-10-0007-01. Provincial household-sector indicators, including debt–service ratios, come from Table 36-10-0226-01. Interprovincial trade exposure comes from

Table 36-10-0697-01. Provincial labour-market slack is measured with the unemployment rate from Table 14-10-0287-03. The common policy instrument is the Bank of Canada's policy rate through the Valet service (Statistics Canada, 2026a, 2026b, 2026c, 2026d, 2026e; Bank of Canada, 2026). The baseline estimating sample runs from January 1991 to December 2024. The Canada aggregate is retained only to construct national controls, national comparator series, and descriptive rows where useful. The local projections are estimated on the ten provinces, with Canada excluded as an estimation unit. The coefficients reported below therefore summarize average reduced-form provincial responses over the inflation-targeting era rather than invariant structural parameters from a single regime.

The main outcomes are constructed to match the theory. Headline inflation is the 12-month log change in all-items CPI. Shelter inflation is the corresponding log change in the shelter component, and rent inflation is constructed analogously from the relevant shelter subcomponent. The main comparator removes the entire shelter contribution with provincial basket weights,

$$\pi_{i,t}^{xS} \approx \frac{\pi_{i,t} - \omega_{i,t}^{S}\pi_{i,t}^{S}}{1 - \omega_{i,t}^{S}}, \tag{10}$$

with associated wedge

$$\Omega_{i,t}^{xS} \equiv \pi_{i,t} - \pi_{i,t}^{xS}. \tag{11}$$

In the descriptive tables, the average wedge refers to the sample mean of this provincial series over time. This is the article's main empirical counterpart to inflation excluding shelter. It is not presented as the one true measure of underlying inflation. It is a deliberately non-shelter comparison measure that removes the broad shelter block. The central empirical question is whether the differential between headline inflation and this broad non-shelter comparison is systematically related to provincial housing and household debt exposure.

Housing exposure is summarized by a composite index,

$$H_i \equiv \frac{1}{2}z(\text{debt-service ratio}_i) + \frac{1}{2}z(\text{shelter weight}_i), \tag{12}$$

where $z(\cdot)$ denotes standardization within the sample. This index is a summary measure, not a primitive structural parameter. The two ingredients serve different roles. Debt–service exposure proxies the cash-flow channel in the IS relation and therefore captures how sensitive households are to debt–service changes. The shelter weight captures the extent to which shelter prices enter measured headline CPI. The empirical question is whether provinces that are simultaneously more exposed to household debt–service changes and more shelter-intensive provinces respond differently to the same national monetary environment. The disaggregated results below suggest that the debt–service component is the more stable predictor of the headline-versus-non-shelter differential, while the shelter-weight component keeps that differential tied to the construction of measured CPI. The composite is retained because the strongest and easiest-to-interpret signal comes from the joint shelter-and-household-debt environment rather than from either component on its own.

Trade integration is measured as interprovincial trade as a share of total trade, standardized within the sample. This variable should be read as a summary measure rather than as a clean structural elasticity of integration. It likely bundles together internal trade ties, diversification, geography, and other features of provincial structure. For that reason, the empirical sections below treat trade as one possible explanation for smaller non-shelter deviations, not as a factor that is assumed in advance to dominate the competing provincial rankings. To make that interpretation more disciplined, a small set of rival provincial factors is constructed. They are chosen to be persistent and interpretable rather than exhaus-

tive. The first is the province-level average unemployment rate, which proxies a long-run cyclical position. The second is the province-level average headline inflation rate before 2020, which captures persistent differences in local price structure without mechanically incorporating the post-pandemic episode. The third is an energy-related CPI basket weight, which captures a leading alternative view that commodity-sensitive or energy-sensitive provinces may account for much of the heterogeneity. These rival factors are used directly in the single-factor scan and horse-race exercises so that the leading housing and trade rankings must show that they matter in the data.

Table 1 reports summary statistics for the monthly province-level panel, and Table 2 reports the province-level exposure measures used in the interaction analysis.

Table 1. Summary statistics.

Variable	Mean	Std. Dev.	Min	Max	Observations
Headline inflation	2.15	1.64	−2.52	10.52	4080
Inflation excluding shelter	2.05	1.64	−2.56	9.84	4080
Shelter inflation	2.44	2.64	−11.59	15.27	4080
Rent inflation	1.89	2.03	−5.32	15.29	4080
Headline minus ex-shelter wedge	0.10	0.61	−3.55	3.63	4080
Unemployment rate	4.90	4.38	1.12	22.07	4080
Policy rate	2.43	1.73	0.25	5.75	3480
Monetary movement proxy	0.00	0.20	−1.50	1.06	3470
Debt service ratio	7.22	1.52	4.14	11.39	4080
Interprovincial trade share	19.18	4.08	9.97	32.07	4080
Shelter basket share	0.25	0.02	0.22	0.29	4080

The panel is monthly and province level. Inflation rates are 12-month log changes in the corresponding CPI indexes. The monetary movement proxy is the preferred monthly policy-rate-change residual used in the baseline local projections.

Table 2. Province-level exposure measures.

Province	Average Headline Inflation	Average Wedge	Debt Service Ratio	Interprovincial Trade Share	Shelter Basket Share
AB	2.39	0.28	7.42	18.13	0.26
BC	2.01	−0.13	9.08	17.03	0.28
Canada	2.12	0.09	7.83	N/A	0.26
MB	2.19	0.14	6.54	22.55	0.24
NB	2.11	0.05	6.40	19.33	0.23
NL	2.09	0.12	6.11	18.33	0.22
NS	2.18	0.12	7.42	21.74	0.25
ON	2.13	0.11	8.38	12.12	0.29
PE	2.20	−0.01	7.14	26.37	0.24
QC	2.02	0.06	6.89	15.77	0.25
SK	2.23	0.26	6.88	20.46	0.24

Shelter basket share is the province-level mean from the available CPI basket-weight data. Debt–service ratio and trade share are annual indicators merged to the monthly panel. The average wedge is the sample mean of headline inflation minus inflation excluding shelter, so it measures the average gap between the two series in each province over the estimation period. Interprovincial trade share is not defined for Canada as a whole; the Canada row therefore reports N/A for that measure.

4.2. What the Common-Proxy Design Identifies

The common monetary movement is observed at the national level. That creates a substantive identification problem. A residual-based monthly policy-rate measure is easy to replicate, but it is not a fully external high-frequency surprise series. For that reason, the article does not ask one proxy to do more than the design can support.

Let $m_t^{(k)}$ denote the kth common monetary proxy in the public-data family

$$\mathcal{M} = \left\{ m_t^{\Delta i,\mathrm{res}}, m_t^{i,\mathrm{res}}, m_t^{\Delta i} \right\}, \tag{13}$$

where $m_t^{\Delta i,\text{res}}$ is the preferred residual from a monthly policy-rate-change equation with lagged national controls, $m_t^{i,\text{res}}$ is a residual from a policy-rate-level equation with the same kind of national information set, and $m_t^{\Delta i}$ is the raw monthly change in the policy rate. The baseline tables use the preferred change residual in its original units. The proxy-robustness table standardizes all three series to unit variance before estimation so that the table is read as a sign-stability check across related public-data measures rather than as a level-for-level comparison with the baseline coefficients.

This changes the empirical target in a useful way. The article is less concerned with one exact national impulse response than with whether the same provinces respond more or less strongly across reasonable measures of common national monetary conditions. If the interpretation is not driven by one proxy alone, the sign and general ordering of the interaction coefficients should remain stable across that family even when the scale of the coefficients changes. The resulting evidence is therefore best understood as reduced-form evidence on different provincial responses under observed national monetary movements. It is not a claim that the family of proxies removes every information effect or omitted policy determinant. The empirical question is whether the provincial ranking is stable enough across public monetary measures to deserve interpretation at all.

4.3. Baseline Local Projections, Competing Factors, and Influence

For each horizon $h \geq 0$ and each province-level outcome $y_{i,t}$, the baseline local projection is

$$\begin{aligned} y_{i,t+h} - y_{i,t-1} = \alpha_i^{(h)} + \delta_i^{(h)} t + \beta_h m_t^{(k)} + \theta_h m_t^{(k)} Z_{i,t-1} \\ + \Gamma_h' X_{i,t-1} + u_{i,t+h}, \end{aligned} \tag{14}$$

where $m_t^{(k)}$ is the common monetary proxy, $Z_{i,t-1}$ is either housing exposure or trade integration, and $X_{i,t-1}$ contains the lagged dependent variable, lagged local unemployment, lagged national inflation, lagged national unemployment, and the lagged policy rate. Province fixed effects are included, along with province-specific linear trends. Estimation follows the local-projections approach of Jordà (2005).

Equation (14) implies that the loading of province i on the common monetary movement is

$$\frac{\partial y_{i,t+h}}{\partial m_t^{(k)}} = \beta_h + \theta_h Z_{i,t-1}, \tag{15}$$

so the differential response between provinces i and j is

$$\frac{\partial y_{i,t+h}}{\partial m_t^{(k)}} - \frac{\partial y_{j,t+h}}{\partial m_t^{(k)}} = \theta_h \left(Z_{i,t-1} - Z_{j,t-1} \right). \tag{16}$$

This is the reason the pooled design is central rather than incidental. The paper is not estimating one national effect and then treating provincial heterogeneity as a nuisance. It is estimating whether provinces with systematically different exposure profiles load differently on the same common national environment.

The baseline dynamic specification uses one lag of the dependent variable and one lag of each local and national control. That choice is parsimonious by design. With only ten provinces and identification coming primarily from monthly common movements, adding many lags quickly absorbs the time variation on which the interacted design depends. The one-lag specification therefore serves as the baseline because it captures persistence while preserving comparability across horizons and monetary proxies. To check that the main conclusions do not hinge on that parsimonious choice, the results section reports a 12-month lag-sensitivity exercise that adds second and third lags of the dependent variable and of the control block. The central sign pattern survives those alternatives.

The core question in the paper is not only whether housing exposure and trade integration matter. It is whether they still matter once other persistent provincial characteristics are considered. The first step is a single-factor scan. Each candidate provincial factor enters Equation (14) one at a time. The scan is exploratory. Its purpose is to identify which characteristics are most informative for three empirical outcomes: headline inflation, the headline-versus-ex-shelter wedge, and absolute deviations of inflation excluding shelter from the national response.

The second step is a horse-race specification,

$$\begin{aligned} y_{i,t+h} - y_{i,t-1} = {} & \alpha_i^{(h)} + \delta_i^{(h)} t + \beta_h m_t^{(k)} \\ & + \theta_h^H m_t^{(k)} Z_{i,t-1}^H + \theta_h^T m_t^{(k)} Z_{i,t-1}^T \\ & + \Pi_h' \left(m_t^{(k)} W_{i,t-1} \right) + \Gamma_h' X_{i,t-1} + u_{i,t+h}, \end{aligned} \tag{17}$$

where $Z_{i,t-1}^H$ is housing exposure, $Z_{i,t-1}^T$ is trade integration, and $W_{i,t-1}$ contains the competing provincial factors. The horse-race does not solve the separate problem of common-movement exogeneity. Province-specific controls cannot turn a national residual into a fully external monetary-policy instrument. What they can do is ask whether the leading housing and trade measures still help explain provincial responses once plausible rivals are considered.

The third step is an influence diagnostic. The central pooled interactions are re-estimated while omitting one province at a time and then omitting Ontario and British Columbia jointly. That exercise does not substitute for a province-by-province transmission paper, but it does answer a narrower and important question: whether the main patterns are broad-based national results or are driven disproportionately by a few large housing-intensive provinces.

4.4. Inference

Inference is anchored in the time dimension. The identifying movement varies across months, not across provinces, so the baseline covariance estimator clusters by month. Because the dependent variables at horizons 6, 12, and 24 overlap and because the cross section is small, the article also reports time-cluster wild-bootstrap p-values for the key 12-month interactions. Those bootstrap results are treated as part of the main evidence, not as a decorative appendix check. The central sign pattern is viewed as persuasive only when the headline housing interaction, the ex-shelter housing interaction, the wedge interaction, and the baseline trade-dispersion interaction survive both month clustering and the time-cluster bootstrap.

5. Empirical Results

5.1. Baseline Different Responses and Time-Dimension Inference

Table 3 delivers the core result around which the rest of the paper is organized, and Table 4 shows immediately that the same qualitative ranking survives the time-cluster wild bootstrap.

This is the paper's central empirical claim and it should be stated narrowly. The result is not that housing exposure leaves non-shelter prices unchanged, nor is it that high-housing provinces necessarily experience weaker price adjustment excluding shelter. The result is differential. Provinces with greater shelter-and-household-debt exposure exhibit a larger headline response than non-shelter response, so measured headline inflation is more sensitive to that exposure than the broad non-shelter comparison measure is. A province can therefore appear more inflationary in headline CPI after a common tightening

without that fact alone implying proportionately stronger price pressure across the full consumption basket.

Table 3. Housing exposure and provincial inflation responses to a common monetary movement.

Outcome	h = 6	h = 12	h = 24
Headline inflation	0.218	0.656 ***	0.378 *
	(0.219)	(0.218)	(0.209)
Inflation excluding shelter	0.064	0.489 **	0.351 *
	(0.213)	(0.202)	(0.196)
Headline minus ex-shelter wedge	0.151 ***	0.165 ***	0.030
	(0.038)	(0.057)	(0.051)

Entries report the coefficient on the interaction between the common monetary movement and the composite housing-exposure index. The index is the simple average of the standardized debt–service ratio and the standardized provincial shelter basket share. Standard errors in parentheses are clustered by month. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 4. Key 12-month interactions and inference.

Interaction Term	Estimate	*p* (Month Cluster)	*p* (Wild Bootstrap)
Headline inflation × housing exposure	0.656 *** (0.218)	0.003	0.000
Inflation excluding shelter × housing exposure	0.489 ** (0.202)	0.016	0.003
Headline minus ex-shelter wedge × housing exposure	0.165 *** (0.057)	0.004	0.020
Absolute ex-shelter deviation × trade integration	−0.083 ** (0.038)	0.030	0.043

Month-clustered *p*-values are paired with time-cluster wild-bootstrap *p*-values. These rows summarize the paper's core differential-response restrictions. The deviation row is measured relative to the national Canadian response, which is omitted from the row label to save space. ** and *** denote statistical significance at the 5% and 1% levels, respectively.

The timing matters as well. The wedge is already visible by six months, is clearest at twelve months, and fades materially by twenty-four months. The evidence is therefore most consistent with a medium-run shelter-sensitive incidence pattern rather than with a permanently different provincial inflation process. The bootstrap results reinforce the central findings: the headline-versus-non-shelter differential and the ex-shelter deviation result remain visible when inference is anchored in the common time variation of the monetary environment.

5.2. Decomposing the Housing Measure

The housing result should not be read as evidence that one narrow shelter subcomponent alone explains provincial inflation heterogeneity. That interpretation is not supported by the data. Table 3 establishes the headline-versus-ex-shelter differential, and Table 5 shows that the broader shelter block matters.

The decomposition of the composite housing measure is therefore central to the paper's contribution. One part of the measure is shelter weight, which ties the result to the construction of measured CPI. The other part is debt–service exposure, which ties it to household cash-flow sensitivity and debt burdens. Table 6 makes that comparison systematic at the 12-month horizon. Debt–service exposure is the stronger standalone predictor of the headline response and of the headline-versus-ex-shelter wedge, while shelter basket share remains informative for headline inflation and keeps the wedge anchored in measured CPI composition. The composite performs best as a summary measure because

the two ingredients work together: one locates the result in the shelter share of measured inflation, and the other links it to a debt-sensitive transmission environment.

Table 5. Housing mechanism inside the shelter block.

Outcome	h = 6	h = 12	h = 24
Shelter inflation	0.574 **	0.963 ***	0.394
	(0.241)	(0.293)	(0.285)
Rent inflation	0.213	0.178	0.152
	(0.189)	(0.121)	(0.150)

The table compares the housing interaction across the broad shelter block and rent inflation. Together with the headline-versus-ex-shelter wedge, these estimates indicate whether the housing result reflects a broad shelter channel rather than only general service-price adjustment. ** and *** denote statistical significance at the 5% and 1% levels, respectively.

Table 6. Composite housing index and its components at the 12-month horizon.

Outcome	Composite Housing Index	Debt–Service Ratio	Shelter Basket Share
Headline inflation	0.656 ***	0.815 **	0.216 ***
	(0.218)	(0.342)	(0.058)
Inflation excluding shelter	0.489 **	0.625 *	0.145 ***
	(0.202)	(0.334)	(0.056)
Headline minus ex-shelter wedge	0.165 ***	0.193 **	0.064
	(0.057)	(0.082)	(0.043)

Entries report the 12-month interaction coefficient from separate interacted local projections. All exposure measures are standardized within the sample before estimation. The table makes the comparison between the joint housing index and its two ingredients explicit for the three main housing-related outcomes. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

That reading is more informative than either polar interpretation. The paper's main result is a joint shelter-and-household-debt result: not pure CPI arithmetic and not merely a household-cash-flow effect divorced from measured shelter intensity. Inflation excluding shelter is used as a non-shelter comparison measure, not as a uniquely privileged measure of latent inflation, and the headline-versus-non-shelter differential is read as a joint product of shelter composition and debt-sensitive incidence.

5.3. Internal Integration and Deviations Excluding Shelter

The baseline trade result is economically meaningful and remains useful, but it should be read narrowly. Provinces with stronger internal trade links deviate less from the national response of inflation excluding shelter, and the effect is weaker in the headline-deviation equation. That is the pattern the framework leads one to look for, and the 12-month trade interaction for the absolute ex-shelter deviation from the Canadian response equals −0.083. The baseline table therefore supports the idea that provinces more tightly linked to the Canadian internal market tend to display less divergence in prices outside shelter after a common tightening. Because the deviation outcome is defined relative to the Canada aggregate rather than to a leave-one-out national benchmark, the coefficient is best read as a compact benchmark-relative deviation result rather than as a precise structural elasticity of internal integration.

The baseline trade result does not, however, settle the broader ranking question by itself. Once rival provincial factors are considered, the evidence becomes more mixed. The internal-trade ranking remains relevant as a negative baseline predictor of ex-shelter deviations, but the competing-factor tables below show that energy-related provincial structure is at least as important and in the joint horse-race often stronger. The right

conclusion is therefore comparative and limited: non-shelter deviations are organized by provincial structure excluding housing, and internal trade ties are one part of that structure rather than a uniquely dominant factor in every specification. Table 7 reports how interprovincial trade integration is associated with deviations from the national Canadian inflation response.

Table 7. Trade integration and deviations from the national inflation response.

Outcome	h = 6	h = 12	h = 24
Absolute deviation of headline inflation	0.012	−0.072 *	−0.048
	(0.039)	(0.043)	(0.050)
Absolute deviation of ex-shelter inflation	0.022	−0.083 **	−0.053
	(0.031)	(0.038)	(0.041)

Entries report the coefficient on the interaction between the common monetary movement and the standardized interprovincial-trade share. The deviation outcomes are measured relative to the national Canadian response, which is omitted from the row labels to save space. Negative coefficients imply that more integrated provinces deviate less from that national response after a common tightening. * and ** denote statistical significance at the 10% and 5% levels, respectively.

5.4. Competing Provincial Factors

A natural concern is that housing exposure and trade integration may simply be standing in for other persistent provincial features. The single-factor scan and the horse-race specification therefore belong in the main results, not in a residual robustness section. The single-factor scan asks which persistent provincial characteristics matter most when each enters on its own. The outcomes are headline inflation, the headline-versus-ex-shelter wedge, and deviations of inflation excluding shelter. Throughout this discussion, "strongest" means the largest absolute interaction t-statistic in the relevant single-factor equation.

The answer differs by outcome. In the headline equation, both shelter basket share and the energy-related CPI weight are especially strong, with the energy-related CPI weight being only marginally stronger by the single-factor ranking metric. Shelter basket share nevertheless remains the more natural housing-side measure because it maps directly into the composition of headline CPI. In the wedge equation, the pre-2020 provincial inflation mean is the strongest single predictor. In the ex-shelter-deviation equation, the energy-related CPI weight is the strongest single predictor. The scan therefore does not point to one universal provincial characteristic. It points to different leading factors for different outcomes.

That pattern is informative rather than inconvenient. The housing result is best understood as a composite shelter-and-household-debt environment rather than as a claim that one housing ingredient dominates every single-factor comparison. By contrast, absolute deviations of inflation excluding shelter depend on a broader set of non-housing provincial characteristics.

The horse-race asks the harder question. Once housing exposure, trade integration, average unemployment, pre-2020 inflation, and an energy-related CPI weight all enter the same equation, which factors still help explain differential provincial responses? In the joint horse-race specifications, the housing measure remains informative in the headline and wedge equations, while the energy-related CPI weight is more informative than trade integration in the ex-shelter-deviation equation. The competing-factor evidence therefore supports the housing measure as the most interpretable and most recurrent predictor of the headline-versus-non-shelter differential, but not as a universal winner in every table.

Taken together, Tables 8 and 9 impose useful discipline on the interpretation of the results. Housing exposure remains informative where the framework says it should matter most: in the headline equation and, more clearly still, in the headline-versus-ex-shelter

differential. Even here, however, the competing-factor evidence is not a universal sweep. In the single-factor scan, pre-2020 inflation structure is the strongest predictor of the wedge, even though the housing measure remains positive, theoretically interpretable, and stable across the broader battery of exercises. The right summary is therefore not that housing dominates every rival in every table but that the shelter-and-household-debt measure is the most interpretable and most recurrent predictor of the differential the theory is built to explain.

Table 8. One-factor scan of candidate provincial organizers.

Provincial Factor	Headline × Factor	Wedge × Factor	Abs. Ex-Shelter Deviation × Factor
Composite housing index	0.656 ***	0.165 ***	0.012
	(0.218)	(0.057)	(0.029)
Debt–service ratio	0.815 **	0.193 **	−0.047
	(0.342)	(0.082)	(0.032)
Shelter basket share	0.216 ***	0.064	0.060 **
	(0.058)	(0.043)	(0.029)
Trade integration	−0.052	−0.013	−0.083 **
	(0.068)	(0.045)	(0.038)
Average unemployment	−0.288 ***	−0.154 ***	−0.067 *
	(0.078)	(0.055)	(0.039)
Pre-2020 inflation mean	0.179 **	0.279 ***	0.006
	(0.076)	(0.066)	(0.033)
Energy-related CPI weight	−0.329 ***	−0.138 **	−0.112 ***
	(0.086)	(0.055)	(0.042)

Each row reports a separate 12-month interacted local projection using one persistent provincial factor at a time. The scan is exploratory. Rows are ranked in the discussion by the absolute *t*-statistic on the interacted factor coefficient in each one-factor equation. Its purpose is to show whether the housing measure and trade integration remain informative when plausible alternatives are considered one at a time. The deviation outcome is measured relative to the national Canadian response. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 10 puts the caution around the horse-race on a measurable footing. In the ten-province cross section, the largest absolute pairwise correlation among the competing factors is below one and the largest variance inflation factor remains below five. That is consistent with meaningful overlap among provincial characteristics, but not with a purely mechanical collapse of the horse-race to one redundant regressor. The leave-one-province-out summary in the same table is equally important: in the ex-shelter-deviation horse-race, the energy-related interaction remains negative throughout the omission exercise, while the trade coefficient stays small and can move close to zero. The ranking is therefore not being created by one single province, but it should still be read as comparative evidence in a correlated small-cross-section setting rather than as a sharp structural contest with one permanent winner.

For ex-shelter deviations, the horse-race does not justify a claim that trade integration dominates every rival. In the joint equation, the energy-related CPI weight is larger and more precisely estimated, while the trade coefficient becomes small and statistically weak. The right reading is therefore narrower: the baseline negative trade interaction is real, but it belongs to a broader cluster of non-housing provincial characteristics in which energy-related composition is at least as important once the factors compete directly. That is still a useful result. It shows that the housing differential and the non-housing deviation pattern are not governed by the same provincial characteristic.

Table 9. Horse-race specification with rival provincial margins.

Outcome	Interaction Term	Estimate
Headline inflation	Housing exposure	0.864 **
		(0.407)
Headline inflation	Trade integration	0.150 *
		(0.090)
Headline inflation	Average unemployment	0.076
		(0.085)
Headline inflation	Pre-2020 inflation mean	0.313 **
		(0.128)
Headline inflation	Energy-related CPI weight	0.036
		(0.216)
Headline minus ex-shelter wedge	Housing exposure	0.213 **
		(0.101)
Headline minus ex-shelter wedge	Trade integration	0.030
		(0.053)
Headline minus ex-shelter wedge	Average unemployment	0.159 **
		(0.070)
Headline minus ex-shelter wedge	Pre-2020 inflation mean	0.368 ***
		(0.085)
Headline minus ex-shelter wedge	Energy-related CPI weight	−0.171
		(0.118)
Abs. ex-shelter deviation	Housing exposure	−0.148 **
		(0.061)
Abs. ex-shelter deviation	Trade integration	−0.026
		(0.042)
Abs. ex-shelter deviation	Average unemployment	0.056
		(0.052)
Abs. ex-shelter deviation	Pre-2020 inflation mean	0.018
		(0.037)
Abs. ex-shelter deviation	Energy-related CPI weight	−0.232 ***
		(0.081)

The common monetary movement is interacted simultaneously with housing exposure, trade integration, average unemployment, pre-2020 inflation, and an energy-related CPI weight. The question is not whether province-specific controls make the common monetary proxy externally identified. It is whether the leading housing and trade factors remain informative when plausible rivals are allowed to compete directly. The deviation outcome is measured relative to the national Canadian response, which is omitted from the row label to save space. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 10. Diagnostics for the horse-race specification.

Diagnostic	Value	Interpretation
Largest absolute pairwise correlation among horse-race factors	0.747	Energy-related CPI weight and average unemployment
Largest variance inflation factor	4.829	Energy-related CPI weight
Leave-one-province-out ranges in the ex-shelter-deviation horse-race		
Factor	Full sample	Omission range
Housing exposure	−0.148	[−0.202, −0.105]
Trade integration	−0.026	[−0.050, 0.007]
Energy-related CPI weight	−0.232	[−0.281, −0.164]

Pairwise correlations and variance inflation factors are calculated from the ten-province cross section used in the horse-race. The omission ranges re-estimate the 12-month ex-shelter-deviation horse-race after removing one province at a time. The table is intended as a compact diagnostic for multicollinearity and influential observations rather than as a separate identification strategy.

5.5. Sensitivity to Lag Structure

The baseline specification uses one lag because the common monetary movement is identified in the time dimension and the province cross section is small. Table 11 shows why that parsimonious choice is acceptable rather than decisive. The housing interaction in

the headline equation remains positive and of similar magnitude when the control block is expanded from one lag to two and three lags. The same is true for the headline-versus-ex-shelter wedge. The trade interaction in the ex-shelter deviation equation remains negative under all three lag choices. The main empirical pattern therefore does not depend on the particular one-lag baseline, even though the paper continues to prefer that specification for transparency and sample retention.

Table 11. Sensitivity of the key 12-month interactions to lag structure.

Outcome and Interacted Margin	One Lag	Two Lags	Three Lags
Headline inflation × housing	0.656 ***	0.686 ***	0.700 ***
	(0.218)	(0.216)	(0.214)
Wedge × housing	0.165 ***	0.144 **	0.126 **
	(0.057)	(0.058)	(0.058)
Abs. ex-shelter deviation × trade	−0.083 **	−0.079 **	−0.080 **
	(0.038)	(0.039)	(0.038)

The baseline specification uses one lag of the dependent variable and one lag of each local and national control. The additional columns add second and third lags of that same control block. The reported coefficient is the 12-month interaction estimate with month-clustered standard errors. ** and *** denote statistical significance at the 5% and 1% levels, respectively.

5.6. Robustness to the Measurement of the Common Monetary Movement

The proxy-family design addresses identification at the level this article needs. The question is whether the provincial response pattern is robust to reasonable alternative ways of measuring the observed common national monetary environment. Table 12 compares a preferred policy-change residual, a separate policy-level residual, and raw standardized policy change. The resulting coefficients are similar because the three series are highly correlated summaries of monthly policy-rate movements. Since each proxy is standardized to unit variance before estimation, the coefficients in Table 12 are not intended as level-for-level comparisons with the baseline residual table. The table is best read as a sign-stability exercise across related proxy definitions rather than as a comparison among independent external instruments. This does not turn the design into a fully external identification strategy but it does show that the main provincial response pattern does not depend on one proxy alone.

Table 12. Robustness to alternative public-data monetary proxies.

Monetary Proxy	Headline × Housing	Wedge × Housing	Abs. Ex-Shelter Deviation × Trade
Preferred change residual	0.102 ***	0.028 **	−0.014 **
	(0.037)	(0.012)	(0.007)
Policy-level residual	0.131 ***	0.033 ***	−0.017 **
	(0.044)	(0.011)	(0.008)
Raw policy change	0.050 *	0.029 ***	−0.012 *
	(0.027)	(0.011)	(0.007)

Each row reports the 12-month interaction estimate from a separate local projection. The preferred change residual, the distinct policy-level residual, and raw policy change are all standardized so magnitudes are comparable, but they are not identical series. The table is therefore read as a sign-stability exercise across closely related public-data proxies rather than as a horse-race among independent external instruments. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

5.7. Provincial Responses, Influence, and Inference

The pooled interaction design is central to the paper because the monetary movement is common across provinces. Still, two complementary questions remain. The first is

whether the main sign pattern is broad based or instead depends on one or two provinces with very large housing exposure. Table 13 answers that question directly. The baseline sign pattern survives the province-by-province omissions and is not driven solely by Ontario and British Columbia. The purpose of the exercise is to show that the main interaction pattern is broad based rather than the product of one or two highly exposed provinces.

The second question is how the pooled interactions map into concrete provincial responses. Table 14 provides that translation. It is not a second estimation step and it is not independent evidence separate from the interaction regressions. It is an accounting transformation of the 12-month interaction coefficients into province-specific relative loadings using province-average exposures. Positive headline and wedge increments indicate provinces that are more exposed to the housing measure. Negative ex-shelter-deviation increments indicate provinces in which internal integration compresses deviations from the Canadian response more strongly. The value of the table is interpretive rather than identificational: it makes the pooled interaction estimates easier to interpret at the province level.

The provincial-incidence and influence diagnostics therefore complete, rather than replace, the main empirical argument. The central sign pattern has already been shown to survive the time-cluster bootstrap and the competing-factor exercises. What remains here is to show that the same pattern is not being mechanically produced by one or two provinces and to translate the pooled coefficients into province-level relative loadings.

Table 13. Influence diagnostics: leave-one-unit-out at the 12-month horizon.

Excluded Unit	Headline × Housing	Wedge × Housing	Abs. Ex-Shelter Deviation × Trade
AB	0.585 ***	0.118 **	−0.078 **
	(0.202)	(0.051)	(0.038)
BC	0.829 ***	0.251 ***	−0.085 **
	(0.300)	(0.075)	(0.037)
MB	0.636 ***	0.158 ***	−0.092 **
	(0.201)	(0.057)	(0.043)
NB	0.663 ***	0.145 **	−0.085 **
	(0.211)	(0.058)	(0.038)
NL	0.684 ***	0.145 **	−0.082 **
	(0.246)	(0.064)	(0.040)
NS	0.621 ***	0.156 ***	−0.075 *
	(0.198)	(0.055)	(0.038)
ON	0.818 ***	0.253 ***	−0.107 **
	(0.279)	(0.078)	(0.053)
PE	0.576 ***	0.127 **	−0.067 *
	(0.191)	(0.053)	(0.035)
QC	0.636 ***	0.164 ***	−0.066 *
	(0.211)	(0.056)	(0.038)
SK	0.644 ***	0.180 ***	−0.093 **
	(0.212)	(0.055)	(0.042)
ON + BC	1.465 ***	0.577 ***	−0.113 **
	(0.569)	(0.158)	(0.055)

Each row re-estimates the 12-month interaction after removing one unit from the panel. Province rows are shown with abbreviations, and the final row removes ON and BC jointly. The deviation outcome is measured relative to the national Canadian response, which is omitted from the row label to save space. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 14. Provincial incidence implied by the 12-month interaction coefficients.

Province	Housing Index	Headline Increment	Wedge Increment	Trade Integration	Ex-Shelter Dev. Increment
AB	0.14	0.09	0.02	−0.26	0.02
BC	1.38	0.91	0.23	−0.53	0.04
Canada	0.49	0.32	0.08	N/A	N/A
MB	−0.47	−0.31	−0.08	0.83	−0.07
NB	−0.89	−0.58	−0.15	0.04	−0.00
NL	−1.16	−0.76	−0.19	−0.21	0.02
NS	−0.04	−0.03	−0.01	0.63	−0.05
ON	1.32	0.87	0.22	−1.73	0.14
PE	−0.29	−0.19	−0.05	1.76	−0.15
QC	−0.10	−0.07	−0.02	−0.84	0.07
SK	−0.39	−0.25	−0.06	0.31	−0.03

Entries translate the 12-month interaction coefficients into province-specific relative loadings using the province-average exposures. Because the exposures are standardized, the values are relative to the sample-average province and are not separate province-by-province regressions. Trade integration is a provincial ranking and is not defined for Canada as a whole; the Canada-row trade cells are therefore reported as N/A.

6. Interpretation and Scope

The results support a differentiated reading of provincial inflation data. The paper's central contribution is a difference between headline inflation and inflation excluding shelter linked to shelter intensity and household debt burdens: provinces higher on the shelter-and-household-debt measure exhibit a larger response of headline inflation than of inflation excluding shelter after a common tightening. That finding is best read jointly. Debt–service exposure is the more stable predictor across specifications, while the shelter-weight component ties the differential to measured headline CPI and to provincial basket composition.

This distinction helps explain what provincial headline inflation is capturing. A province with relatively persistent headline inflation after a tightening is not necessarily a province with proportionately stronger price pressure across the whole consumption basket. In these estimates, high-housing provinces show more positive responses in both headline and non-shelter inflation, but the headline response rises by more. Headline divergence therefore combines two pieces of information: broad price adjustment excluding shelter and the additional way in which the shelter block transmits the same monetary environment into measured CPI.

The mechanism evidence narrows the interpretation without collapsing it to one series. The data do not reduce the broader shelter result to simple basket arithmetic. The most coherent reading is that provinces that are both more shelter-intensive and more debt–service-sensitive show a larger response in headline inflation than in the non-shelter comparison measure.

The non-housing side should be read more cautiously. Interprovincial trade share is not interpreted here as a pure structural openness elasticity, and the horse-race estimates do not support such a claim. What the evidence supports is weaker but still informative: once the shelter block is netted out, deviations of inflation excluding shelter are associated with a broader cluster of provincial structure, with internal trade ties relevant in baseline specifications and energy-related composition often especially important when competing factors enter jointly.

Inflation excluding shelter is useful because it removes the broad shelter block that drives the headline differential and therefore provides a non-shelter comparison measure. It is not a uniquely privileged measure of latent underlying inflation. Its role in the paper is narrower and more concrete: it helps separate the headline inflation response from the broader provincial price response excluding shelter.

The paper's limits follow directly from that interpretation. The common monetary environment is measured with public-data monetary measures rather than with an exter-

nally identified instrument, so the coefficients should be read as disciplined reduced-form evidence on provincial differentials. The provincial measures are persistent and economically interpretable, but they are not structural primitives. The regional model is meant to discipline comparative empirical restrictions, not to deliver a complete welfare analysis. Those limits narrow the contribution, but they do not undo it. The article identifies a robust housing-related difference between headline inflation and inflation excluding shelter and a more qualified non-housing pattern in absolute deviations of inflation excluding shelter from the national response.

7. Conclusions

This article shows how the same common monetary tightening can produce different provincial inflation outcomes in Canada. The evidence points most clearly to shelter and household debt. Provinces with larger shelter weights and heavier household debt burdens show a larger response of headline inflation than of inflation excluding shelter after a common tightening. That difference is the paper's main result, and it is the part of the evidence that is most stable across specifications, decomposition exercises, and inference procedures.

The paper also studies whether inflation excluding shelter remains closer to, or farther from, the national response after the same common movement. Here the evidence is more mixed. Internal trade ties matter in baseline specifications, but once other provincial characteristics are introduced directly, energy-related provincial structure is often at least as important. The non-housing side of the paper is therefore best read as evidence on a broader set of provincial differences rather than as a claim that one non-housing factor dominates every specification.

The broader lesson is straightforward. Uneven provincial headline inflation after a common tightening does not necessarily imply proportionate divergence in the broader non-shelter price system. The estimates are informative about differential provincial incidence under observed common monetary conditions, but they do not claim to recover an externally identified monetary-policy effect. Part of that divergence reflects how shelter intensity, household debt exposure, and other persistent local features translate the same national monetary environment into different measured price indexes. Canada is especially informative because those provincial differences are visible in both the data and the institutional setting.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All data used in the article are publicly available from the Bank of Canada and Statistics Canada sources listed in the manuscript. The replication package includes the analysis code, the raw public-data inputs, and the generated LaTeX output files used to compile the manuscript. Intermediate files are created when the analysis is rerun. Package available by request.

Acknowledgments: During the preparation of this manuscript, the author used OpenAI ChatGPT, model GPT-5.5 Pro for language editing, clarification of exposition, and limited drafting and editing of analysis code used for data processing and table production. The author checked and edited the output and takes full responsibility for the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Table of Symbols

Table A1. Table of symbols used in the text.

Symbol	Meaning
Indices and horizons	
i, j	Province indexes.
t	Time index.
h	Local-projection horizon.
N	Number of provinces or regions in the theoretical monetary union.
Inflation blocks and local demand	
$\pi_{i,t}$	Headline inflation in province i at time t.
$\pi^{c}_{i,t}$	Inflation excluding shelter in province i at time t.
$\pi^{s}_{i,t}$	Shelter inflation in province i at time t.
ω_i	Provincial shelter share in headline CPI.
$x_{i,t}$	Local demand or local output-gap term in province i.
$\Omega_{i,t}$	Theoretical headline-versus-non-shelter wedge, $\Omega_{i,t} \equiv \pi_{i,t} - \pi^{c}_{i,t}$.
Policy rate, natural rate, and transmission parameters	
i_t	Common nominal policy rate set by the central bank.
r^{n}_{t}	Natural real interest rate.
σ_i	Province-specific sensitivity of local demand to the ex ante real interest-rate gap.
χ_i	Province-specific sensitivity of local demand to debt–service exposure.
$b_{i,t-1}$	Debt–service exposure carried into period t from period $t-1$.
η_i	Strength of internal integration for province i.
w_{ij}	Weight linking province i to province j in the internal-integration term.
Common monetary proxy and Phillips-curve parameters	
β	Forward-looking coefficient in the inflation equations.
κ_i	Slope coefficient linking local demand to inflation excluding shelter.
κ^{s}_{i}	Slope coefficient linking local demand to shelter inflation.
λ_i	Shelter pass-through parameter from policy-rate changes into the financing-sensitive shelter block.
ζ_i	Extent to which the financing-sensitive shelter pass-through enters measured shelter inflation.
ξ_i	Coefficient on local shelter-price changes in the shelter block.
$q_{i,t}$	Local shelter-price or housing-price term entering shelter inflation.
$u^{x}_{i,t}$	Demand disturbance in province i.
$u^{c}_{i,t}$	Non-shelter inflation disturbance in province i.
$u^{s}_{i,t}$	Shelter inflation disturbance in province i.
Empirical inflation measures	
$\pi^{S}_{i,t}$	Empirical shelter inflation.
$\omega^{S}_{i,t}$	Shelter basket weight used in the empirical ex-shelter construction.
$\pi^{xS}_{i,t}$	Inflation excluding the broad shelter block.
$\Omega^{xS}_{i,t}$	Headline-minus-ex-shelter wedge, $\Omega^{xS}_{i,t} \equiv \pi_{i,t} - \pi^{xS}_{i,t}$.
Exposure rankings and standardization	
H_i	Composite housing-exposure index, equal to the simple average of the standardized debt–service ratio and standardized shelter weight.
$z(\cdot)$	Standardization operator within the sample.
$Z_{i,t-1}$	Generic province-specific exposure entering the interacted local projection.
$Z^{H}_{i,t-1}$	Housing exposure ranking in the horse-race specification.
$Z^{T}_{i,t-1}$	Trade-integration ranking in the horse-race specification.
$W_{i,t-1}$	Vector of rival provincial factors in the horse-race specification.

Table A1. *Cont.*

Symbol	Meaning
Common monetary proxies	
$\mathcal{M}$	Family of common public-data monetary proxies used in the robustness exercises.
$m_t^{(k)}$	The kth common monetary proxy at time t.
$m_t^{\Delta i,\text{res}}$	Preferred residual from a monthly policy-rate-change equation with lagged national controls.
$m_t^{i,\text{res}}$	Residual from a policy-rate-level equation with lagged national controls.
$m_t^{\Delta i}$	Raw monthly change in the policy rate, standardized in the proxy-robustness table.
Local-projection notation	
$y_{i,t+h}$	Generic province-level outcome at horizon h.
$\alpha_i^{(h)}$	Province fixed effect in the horizon-h local projection.
$\delta_i^{(h)}$	Province-specific linear trend in the horizon-h local projection.
β_h	Coefficient on the common monetary proxy in the horizon-h local projection.
θ_h	Interaction coefficient between the common monetary proxy and a province-specific exposure in the horizon-h local projection.
θ_h^H	Housing-exposure interaction coefficient in the horse-race specification.
θ_h^T	Trade-exposure interaction coefficient in the horse-race specification.
Γ_h	Coefficient vector on baseline controls in the horizon-h local projection.
Π_h	Coefficient vector on competing interacted provincial factors in the horse-race specification.
$X_{i,t-1}$	Vector of baseline controls in the empirical specification.
$u_{i,t+h}$	Regression residual in the horizon-h local projection.

Notes. Roman i_t denotes the common nominal policy rate, whereas the subscript i indexes provinces. Superscript S denotes shelter, and xS means excluding shelter. The symbol β is the forward-looking coefficient in the theoretical inflation equations, while β_h is the horizon-specific coefficient on the common monetary proxy in the local-projection regressions.

References

Bank of Canada. (2026). *Valet API: Policy interest rate group*. Bank of Canada. Available online: https://www.bankofcanada.ca/valet/ (accessed on 4 May 2026).

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

Bounajm, F., & McWhirter, A. (2024). *Modelling Canadian mortgage debt and payments in a semi-structural model* [Bank of Canada Staff Analytical Note 2024–1]. Available online: https://www.bankofcanada.ca/2024/01/staff-analytical-note-2024-1/ (accessed on 4 May 2026).

Champagne, J., & Sekkel, R. (2018). Changes in monetary regimes and the identification of monetary policy shocks: Narrative evidence from Canada. *Journal of Monetary Economics*, *99*, 72–87. [CrossRef]

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

Nakamura, E., & Steinsson, J. (2018). Identification in macroeconomics. *Journal of Economic Perspectives*, *32*(3), 59–86. [CrossRef]

Solodoha, E., Rosenzweig, S., & Harel, S. (2023). Incentivizing angels to invest in start-ups: Evidence from a natural experiment. *Research Policy*, *52*(1), 104634. [CrossRef]

Statistics Canada. (2023). *Shelter in the Canadian CPI: An overview, 2023 update* (Prices Analytical Series). Statistics Canada.

Statistics Canada. (2025). *An analysis of the 2025 Consumer Price Index basket update, based on 2024 expenditures* (Prices Analytical Series). Statistics Canada.

Statistics Canada. (2026a). *Table 14-10-0287-03: Labour force characteristics by province, monthly, seasonally adjusted*. Available online: https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410028703 (accessed on 4 May 2026).

Statistics Canada. (2026b). *Table 18-10-0004-01: Consumer price index, monthly, not seasonally adjusted*. Available online: https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1810000401 (accessed on 4 May 2026).

Statistics Canada. (2026c). *Table 18-10-0007-01: Basket weights of the consumer price index, Canada, provinces, Whitehorse, Yellowknife and Iqaluit*. Available online: https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1810000701 (accessed on 4 May 2026).

Statistics Canada. (2026d). *Table 36-10-0226-01: Household sector, selected indicators, provincial and territorial*. Available online: https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610022601 (accessed on 4 May 2026).
Statistics Canada. (2026e). *Table 36-10-0697-01: Interprovincial trade, summary level*. Available online: https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=3610069701 (accessed on 4 May 2026).