Regional Phillips Curve in a Small Open Economy: Evidence from Bolivia

Roger Mario Lopez Justiniano*

Hugo Vaca Pereira Rocha†

Abstract

We estimate New Keynesian Phillips Curve relationships using quarterly data for Bolivia's nine departments over 1993–2019, using specifications that control for common shocks and allow for heterogeneous regional slopes. Despite substantial subnational variation, the estimated association between inflation and the local output gap is small and statistically indistinguishable from zero across subsamples and slack measures. The regional panel also lets us evaluate competing explanations. While output-gap slopes vary across departments, this heterogeneity does not aggregate into a meaningful national relationship, suggesting that flat aggregate estimates are not driven by aggregation bias. More tellingly, even non-tradable inflation—the component most plausibly determined by local market conditions—shows no sensitivity to departmental slack, offering little support for a compositional story. Taken together, the evidence indicates that local prices do not move systematically with local slack. We interpret this pattern as consistent with an inflation process dominated by common forces in a setting shaped by Bolivia's fixed exchange rate regime and weak monetary transmission—conditions shared by many emerging markets where slack-based stabilization frameworks may have limited operational traction.

JEL Classification: E31, E52, C23, R11, O54

Keywords: Phillips Curve, Inflation, Bolivia, Regional Heterogeneity

---

*PricewaterhouseCoopers. Email: roger.lopez.justiniano@pwc.com

†University of Wisconsin - La Crosse. Email: hvacapereirarocha@uwlax.edu. We thank Walter Morales Carrasco, Carlos Gustavo Machicado and Luiggi Donayre for helpful comments and suggestions.

1 Introduction

A central question for monetary policy in small open economies is whether domestic slack meaningfully shapes inflation once exchange-rate movements, import prices, and other common forces are accounted for. The answer has direct implications for the operational reliability of stabilization frameworks that rely on output-gap measures: if inflation is driven predominantly by common forces rather than local demand pressure, slack-based rules provide limited traction regardless of how carefully the output gap is measured. An emerging literature exploits regional variation within countries to identify the Phillips Curve slope more cleanly than national time-series estimates allow, by absorbing common shocks through time fixed effects and isolating within-date cross-sectional variation (McLeay and Tenreyro, 2020; Hazell et al., 2022). We apply this strategy to Bolivia — a small open economy with a prolonged fixed-exchange-rate episode, weak interest-rate transmission, and a highly informal labor market — and find that even with rich subnational variation and demanding specifications, the inflation-slack relationship remains flat. We argue that Bolivia is not exceptional in these features but representative of a broader class of emerging market economies, and that the conditions driving our findings can be assessed portably across countries

National-level Phillips Curve estimates often rely on limited observations and can obscure regional variation. Coen et al. (1999) and Kapetanios et al. (2021) argue that aggregation in the presence of heterogeneous regional dynamics can lead to biased parameter estimates. A growing literature therefore emphasizes the advantages of using regional data, which provide richer variation to capture inflation dynamics and to disentangle demand- and supply-driven shocks. Moreover, endogeneity arising from nationally determined monetary policy can contaminate aggregate Phillips Curve estimates. Regional panels help mitigate this concern by exploiting cross-sectional variation. This yields more stable inflation–output relationships and helps distinguish demand versus supply effects (Fitzgerald et al., 2024). In a similar spirit, McLeay and Tenreyro (2020) stress that properly accounting for policy reactions and other common forces is essential for interpreting observed inflation dynamics.

A key insight from the New Keynesian Phillips Curve literature is that inflation dynamics depend on both real activity and inflation expectations, with the latter anchoring the forward-looking component of the NKPC and playing a central role in monetary policy transmission. However, inflation expectations are difficult to estimate reliably, given their strong dependence on model assumptions and identification strategy (Mavroeidis et al., 2014). Hazell et al. (2022) demonstrate that time fixed effects in regional panels can absorb shared beliefs about the monetary regime, thereby controlling for expectation-driven inflation dynamics that are common across regions. In our setting, department-level measures of inflation expectations are not directly observed. We therefore treat expectation-driven movements as predominantly common across departments at a given date and absorb them using time fixed effects — and, in our common-factor specifications, additional controls for aggregate components. Our estimand is thus the sensitivity of local inflation to local slack, conditional on these common forces, which is the relevant object for understanding how real activity shapes inflation across regional components of an economy.

This regional identification follows the logic emphasized by McLeay and Tenreyro (2020): once aggregate components driven by the monetary regime and policy responses are netted out, the local-slack slope is identified from within-date cross-department variation. We interpret our results primarily as reflecting an inflation process dominated by common forces in Bolivia — particularly the post-2011 fixed exchange rate and pass-through environment, together with administered prices and weak monetary transmission — rather than by local demand conditions. Other channels, including expectation formation, may matter, but in our setting they likely operate mainly through time-varying common components absorbed by time fixed effects.

One identification caveat is nonetheless worth noting. Residual within-quarter cross-department variation in the output gap may reflect not only local demand pressure but also department-specific supply shocks, which move output and prices jointly and therefore survive the fixed-effects transformation. To the extent such shocks are present — particularly in commodity-dependent departments such as Potosí and Oruro — the estimated slope captures a mixture of demand and supply variation rather than isolating local demand slack in the strict sense. The consistency of near-zero estimates across specifications, subsamples, and inflation measures suggests, however, that the result is not driven by any particular source of local variation.

Our central question is whether inflation responds systematically to local economic slack once nationwide and external forces are accounted for — and what institutional conditions determine the answer. We argue that three features, individually common and jointly prevalent across a wide range of emerging markets, are sufficient to flatten the empirical Phillips Curve: (i) a fixed or heavily managed exchange rate that constrains monetary policy responses to domestic fluctuations; (ii) weak interest-rate transmission arising from dollarization, shallow financial markets, or fiscal dominance; and (iii) a highly informal labor market that absorbs demand expansions without generating wage pressure, thereby muting the pass-through from output gaps to inflation. Bolivia during 1993–2019 exhibits all three, making it a stringent and well-documented test case. But the same configuration characterizes economies across Latin America, sub-Saharan Africa, and parts of Asia, and our findings carry direct implications for how slack-based stabilization frameworks should be evaluated in those settings. In what follows, we therefore place greatest weight on this three-condition interpretation as the most portable mechanism behind a weak local-slack slope.

Most of this literature has focused on states within the United States and the Euro Area (Berk and Swank, 2011; Beraja et al., 2019; Fitzgerald et al., 2020; Hooper et al., 2020; Hazell et al., 2022; Schuffels et al., 2022), while emerging market economies remain underexplored due to limited regional data availability. Nonetheless, recent studies have begun to fill this gap. El-Shagi and Tochkov (2024) document significant provincial heterogeneity in China, driven by differences in industrial structure and market development. Behera et al. (2018) find that India's state-level Phillips Curve remains robust despite global structural changes. Orlov and Postnikov (2022) estimate separate Phillips Curves for 80 Russian regions using quarterly panel data, while Aginta (2023) confirm the existence of the Phillips Curve in Indonesia, showing stronger inflation responses when mining provinces are excluded, along with regional variation in inflation persistence and exchange rate pass-through.

Despite the growing body of regional Phillips Curve research, much less is known about how these relationships operate in smaller, data-constrained contexts like Bolivia. In many Latin American countries, monetary policy is set nationally and is informed primarily by aggregate indicators, which may not fully reflect regional variation in economic conditions and inflation dynamics. While some central banks systematically incorporate regional information into their assessment of the economy, comparable evidence on subnational inflation dynamics is more limited in much of Latin America. In this setting, clarifying whether (and in what components of inflation) prices co-move with local cyclical conditions can improve the interpretation of aggregate signals within a single national policy framework. This paper contributes to this emerging literature by examining the New Keynesian Phillips Curve (NKPC) across Bolivia's nine departments from 1993 to 2019. We estimate both standard and hybrid Phillips Curve models using regional and time fixed effects to assess whether inflation responds to economic slack in a regionally heterogeneous setting, and we additionally estimate the Phillips Curve separately for tradable and non-tradable inflation, in the spirit of McLeay and Tenreyro (2020), since nationally priced tradables can mechanically attenuate regional Phillips Curve estimates. More broadly, Bolivia provides a stringent test case for slack-based inflation models in a small open economy with limited monetary transmission and substantial exposure to common shocks. Our results therefore speak to a wider class of emerging-market settings in which inflation may be driven predominantly by nationwide and external forces. We view the paper primarily as an empirical contribution with direct implications for how slack-based stabilization frameworks should be evaluated in such environments.

The paper makes five contributions. First, it estimates regional Phillips Curves using methods designed for panels with pervasive common shocks and heterogeneous slopes, clarifying how common forces affect inference on local slack. Second, it exploits CPI microdata to decompose inflation into tradable and non-tradable components, providing an open-economy diagnostic for whether locally set prices exhibit stronger slack sensitivity. Third, it conducts an unusually extensive slack-measurement robustness exercise to distinguish stable patterns from artifacts of output-gap construction. Fourth, it assembles and harmonizes a long quarterly department-level dataset on inflation and slack, expanding the empirical basis for studying inflation dynamics in a data-constrained setting. Fifth, by mapping Bolivia's institutional configuration onto a typology of three conditions — exchange-rate regime constraint, weak interest-rate transmission, and an informal labor market — the paper provides a portable diagnostic for other emerging market economies. Where these conditions jointly prevail, policymakers and researchers should anticipate flat empirical Phillips Curve slopes and treat slack-based inflation projections with corresponding caution.

To investigate whether inflation responds similarly to cyclical conditions across departments, we test for spatial heterogeneity in the slope of the New Keynesian Phillips Curve. The evidence indicates cross-department variation in estimated slopes, suggesting that a single pooled relationship may mask differences in how regional inflation co-moves with local activity. By comparing pooled and heterogeneous-slope specifications and by accounting for unobserved common factors, the paper also assesses how well aggregate-style restrictions fit subnational data and whether common shocks materially affect inference.

Across our baseline TWFE, hybrid, and common-factor specifications, the estimated sensitivity of both headline and core inflation to the departmental output gap is small and typically statistically indistinguishable from zero. Allowing for slope heterogeneity and cross-sectional dependence does not materially change this conclusion, which points to inflation dynamics in Bolivia being driven largely by common forces—potentially including regime-wide expectations—rather than local demand pressure. These results illustrate how subnational panels, paired with estimators that accommodate common shocks, can sharpen inference on the empirical relevance of local slack for inflation.

The paper is structured as follows. Section 2 describes Bolivia's macroeconomic context. Section 3 introduces the department-level dataset and details the construction of inflation and output-gap measures. Section 4 presents the NKPC specifications and estimation strategy for headline and core inflation. Section 5 extends the analysis by estimating the Phillips curve separately for tradable and non-tradable inflation. Section 6 reports a simulation-based robustness exercise that evaluates the sensitivity of the headline and core results to alternative output-gap constructions. Section 7 concludes by summarizing the findings and discussing implications for monetary policy analysis.

2 The Phillips Curve in Bolivia

Understanding inflation dynamics in Bolivia requires attention to the country's institutional and policy evolution following the stabilization of the late 1980s. After the hyperinflation episode of the mid-1980s (Sachs, 1986; Morales and Sachs, 1989), Bolivia pursued successive waves of reform, with the 1995 Central Bank Law representing an important milestone within these broader "second-generation" changes. The law sought to strengthen the central bank's formal mandate and autonomy, with price stability placed at the core of monetary policy objectives (Morales, 2005). Nevertheless, persistent dollarization and shallow financial markets continued to limit the effectiveness of interest-rate-based transmission (Morales, 2003), increasing the relative importance of expectations, administered prices, and external conditions in shaping inflation outcomes.

Bolivia's institutional setting motivates why the regional Phillips Curve may be weak. Unlike advanced economies, Bolivia has historically relied on non-traditional policy tools such as exchange rate management and credit controls (Requena et al., 2002). These features are not unique to Bolivia but are common among emerging markets, where limited credibility, fiscal dominance, and vulnerability to external shocks constrain conventional monetary frameworks (Frankel, 2010). The adoption of a fixed exchange rate regime in 2011 further restricted the scope for interest-rate-based stabilization and limited the central bank's ability to respond to domestic fluctuations. Together with recurrent fiscal deficits and a narrow export base concentrated in hydrocarbons and minerals, these characteristics imply that inflation may track common nationwide and external forces more closely than local demand conditions—making Bolivia a stringent test of slack-based inflation models and complicating the correspondence between domestic slack measures and inflation outcomes.

Consistent with this macroeconomic setting, existing studies for Bolivia tend to find limited evidence of a strong inflation–output gap relationship. Valdivia (2008), Murillo (2014), and Mora-Barrenechea (2021) emphasize the importance of inflation expectations—particularly backward-looking ones—in driving price dynamics, while estimated output-gap effects remain limited. Figure 1 plots the relationship between quarter-on-quarter inflation (both core and headline) and the output gap across four distinct macroeconomic subperiods, defined for historical context: (i) Recovery and Slow Growth (1987–1998), (ii) Financial Crisis (1999–2002), (iii) Nationalization and Growth (2003–2018), and (iv) the full sample period (1993–2019). These periods follow the long-run growth phase classifications developed by Machicado (2018), which are based on structural shifts in Bolivia's economic model, external conditions, and macro policy regimes. For instance, the 1987–1998 period reflects the post-hyperinflation stabilization and liberalization era; 1999–2002 captures the macroeconomic instability and banking crisis; 2003–2018 marks a commodity boom phase driven by natural gas and mineral exports under state-led development; and the period beginning in 2019 reflects rising fiscal pressures and the onset of economic stagnation, which intensified with the COVID-19 shock after the end of our sample.

Despite these distinct phases, the inflation–output gap relationship remains weak or flat across all subperiods. The scatterplots in Figure 1 show no strong correlation between output gaps and inflation, even during the 2003–2018 expansionary period when capacity pressures might have been expected to drive prices. The lack of sensitivity is especially evident in core inflation, suggesting that even underlying price trends were decoupled from cyclical demand pressures. This weak relationship is consistent with a setting in which inflation is influenced by factors beyond domestic slack, including supply-side conditions and inflation expectations, and may also reflect institutional and pricing features such as the post-2011 exchange rate arrangement, the prevalence of informal labor markets, and the role of administered prices.

These persistent patterns underscore the value of using subnational data to document heterogeneity in real activity and to test whether local slack can explain inflation once common shocks are controlled for. As we show below, however, even with department-level variation and demanding specifications that allow for common factors and heterogeneous slopes, the inflation–slack relationship remains weak and unstable in this setting.

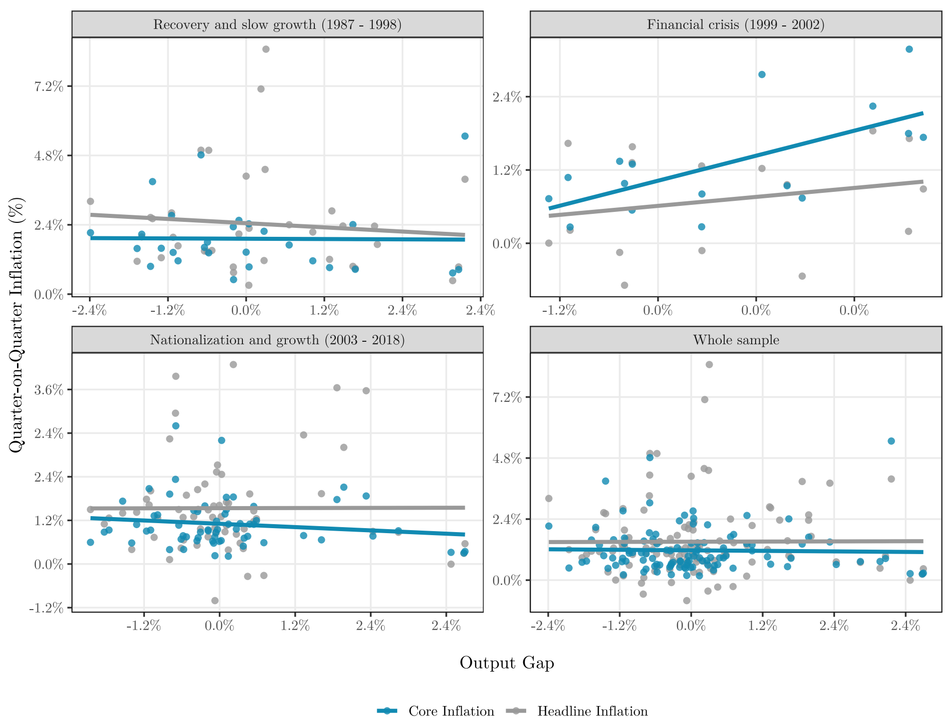

Figure 1: Relationship Between Output Gap and Quarterly Inflation Across Macroeconomic Phases in Bolivia. Scatterplots display core and headline inflation against the output gap across three macroeconomic subperiods and the full sample period (1993-2019): (i) Recovery and Slow Growth (1987–1998), (ii) Financial Crisis (1999–2002) and (iii) Nationalization and Growth (2003–2018). The plots show little to no correlation, especially for core inflation, indicating a weak Phillips Curve relationship across all periods. Source: National Institute of Statistics (INE) and authors' calculations. Subperiod classification follows Machicado (2018)

3 Data

3.1 Inflation and Output Data for Headline and Core Measures

We use an unbalanced panel covering Bolivia's nine departments from 1993:Q3 to 2019:Q4. Complete quarterly data are available for La Paz, Cochabamba, and Santa Cruz over the full period. Data for the other six departments—Beni, Chuquisaca, Oruro, Pando, Potosí, and Tarija—begin in 2008:Q2, reflecting the timing of data collection by INE (Table 1).

We obtain inflation data from INE, including both headline and core measures. Headline inflation reflects the quarterly change in the Consumer Price Index (CPI), while core inflation excludes volatile components such as food and energy. Quarterly inflation rates were computed as the average of monthly inflation rates and seasonally adjusted using the Census X-11 procedure.

Regional output data are drawn from Chalup and Escobar (2023), who estimate quarterly GDP estimates at the departmental level. The output gap is calculated as the percentage deviation of actual output from its trend, with potential output estimated using the Hodrick-Prescott filter ($\lambda = 1600$), consistent with standard practice for quarterly macroeconomic series. This measure carries a conceptual caveat worth acknowledging. In departments whose output is driven primarily by commodity extraction — notably Potosí and Oruro — fluctuations in regional GDP may reflect resource production cycles or centralized rent flows that do not necessarily correspond to local absorption or labor market tightness. To the extent that the output gap captures these dynamics rather than genuine demand slack, the estimated inflation-slack slope may be attenuated toward zero through measurement error rather than reflecting a true absence of a demand channel. Additionally, for the six departments whose series begin in 2008Q2, the shorter sample increases the susceptibility of the HP filter to endpoint artifacts. These limitations reinforce our reliance on the simulation-based robustness exercise as a check on whether the weak slope is an artifact of any particular output gap construction.

Table 1: Inflation and Output Gap by Department

Department	**N**	**Date Range**	**Headline Inflation**	**Core Inflation**	**Output Gap**
Beni	47	2008Q2–2019Q4	0.681 (0.106)	0.657 (0.073)	0.122 (0.130)
Cochabamba	106	1993Q3–2019Q4	1.494 (0.124)	1.237 (0.093)	0.035 (0.099)
La Paz	106	1993Q3–2019Q4	1.224 (0.106)	1.057 (0.076)	0.007 (0.147)
Oruro	47	2008Q2–2019Q4	1.060 (0.204)	0.825 (0.125)	0.993 (0.355)
Pando	47	2008Q2–2019Q4	0.581 (0.206)	0.568 (0.121)	0.318 (0.232)
Potosi	47	2008Q2–2019Q4	1.007 (0.212)	0.745 (0.109)	1.135 (0.581)
Santa Cruz	106	1993Q3–2019Q4	1.313 (0.111)	1.199 (0.084)	0.088 (0.179)
Chuquisaca	47	2008Q2–2019Q4	1.076 (0.177)	0.747 (0.077)	0.024 (0.290)
Tarija	47	2008Q2–2019Q4	1.496 (0.317)	0.827 (0.127)	-0.332 (0.519)

Table 1 summarizes department-level inflation and output gap statistics with standard deviations in parentheses, while Table 2 highlights their correlations with national aggregates. Only Pando (0.718) and Santa Cruz (0.634) show output gap correlations above 0.6, while most departments—such as La Paz (0.047) and Potosí (0.190)—exhibit weak comovement. In contrast, inflation is highly correlated nationwide, suggesting that prices move together even when real activity varies considerably across regions. Figures 2 and 3 illustrate these patterns.

Table 2: Correlation of Regional Headline Inflation, Core Inflation, and Output Gap with National Aggregates

Department	N	Date Range	Headline Inflation	Core Inflation	Output Gap
Beni	47	2008Q2–2019Q4	0.842	0.547	0.574
Cochabamba	104	1993Q3–2019Q4	0.929	0.837	0.554
La Paz	104	1993Q3–2019Q4	0.951	0.922	0.047
Oruro	47	2008Q2–2019Q4	0.919	0.871	0.392
Pando	47	2008Q2–2019Q4	0.677	0.675	0.718
Potosi	47	2008Q2–2019Q4	0.899	0.877	0.190
Santa Cruz	104	1993Q3–2019Q4	0.940	0.906	0.634
Chuquisaca	47	2008Q2–2019Q4	0.929	0.899	0.339
Tarija	47	2008Q2–2019Q4	0.579	0.824	0.260
National average	–	–	**0.852**	**0.818**	**0.412**

Figure 2 displays quarterly output gaps from 1993 to 2019. The blue line shows the output gap estimated using the mean of the series provided by Chalup and Escobar (2023), while the shaded gray area reflects the range of variation when the gap is recalculated using randomly selected series based on the 84$^{\text{th}}$ and 16$^{\text{th}}$ percentiles, also derived from Chalup and Escobar (2023). Departments such as Potosí and Oruro exhibit persistently positive and volatile gaps, reflecting their exposure to commodity cycles. In contrast, Chuquisaca and Tarija display more muted or even negative cyclical positions. These patterns underscore the value of using subnational data to document heterogeneity in real activity and to test whether local slack can explain inflation once common shocks are controlled for.

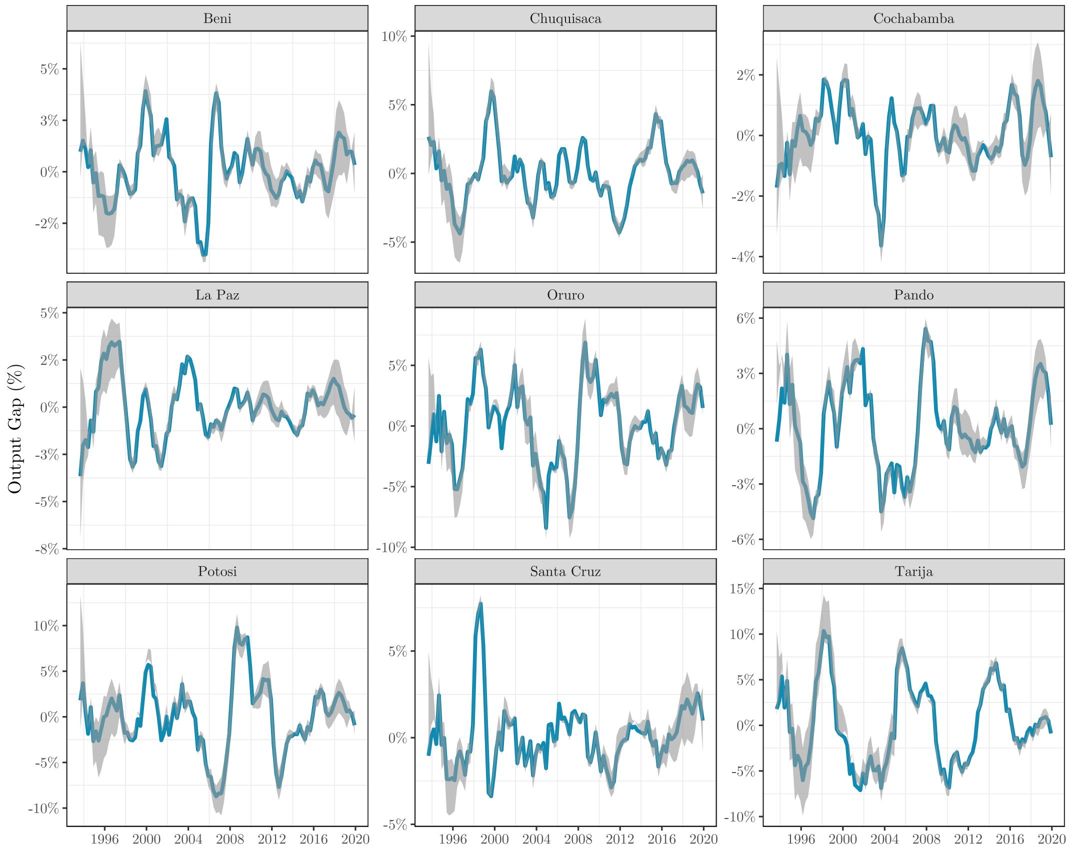

Figure 2: Estimated Output Gap by Department (1993–2019). Each panel shows the quarterly output gap for one department. Volatility is especially high in mining-dependent regions such as Potosí and Oruro.

Figure 3 plots core and headline inflation against output gaps, illustrating the contemporaneous relationship between these variables. Most departments exhibit no clear relationship, consistent with the weak or absent Phillips Curve evidence found in the descriptive tables. The lack of consistent correlation across departments reinforces earlier findings and casts doubt on the applicability of a uniform Phillips Curve framework at the regional level.

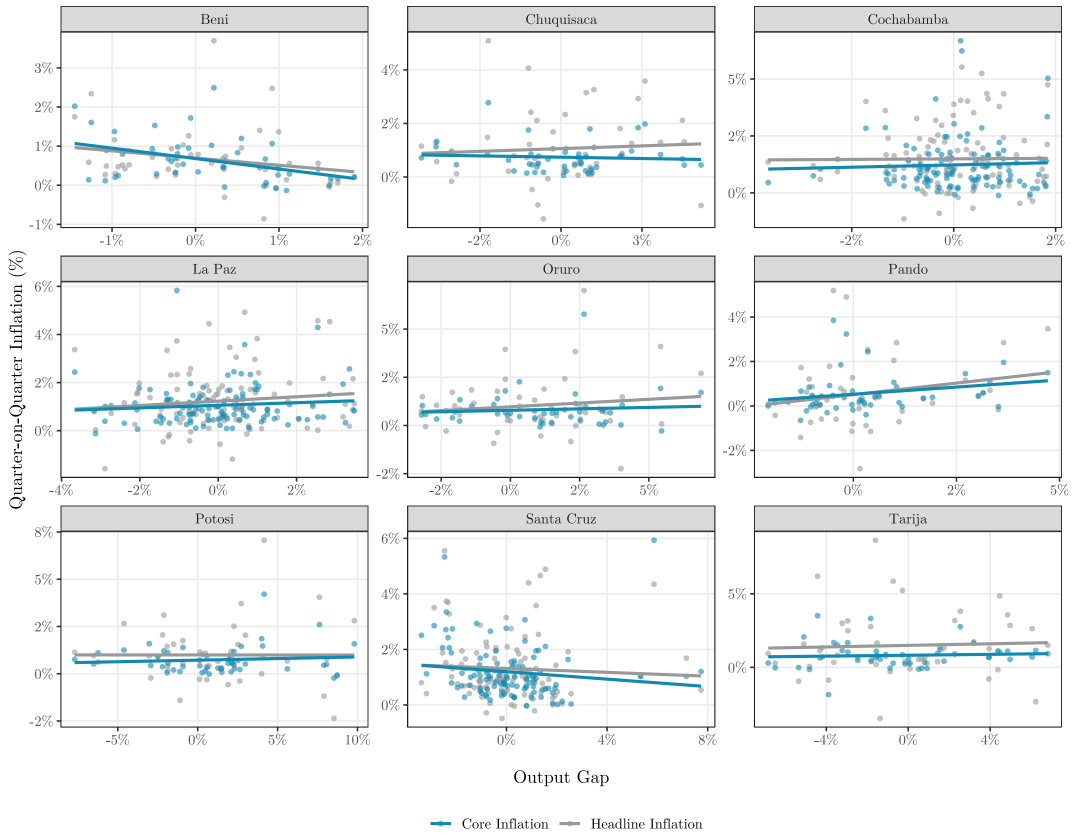

Figure 3: Scatterplots of Output Gap vs. Inflation (Core and Headline) by Department. Each point represents a quarterly observation. The weak correlation highlights regional variation in inflation dynamics.

These descriptive patterns confirm substantial heterogeneity in Bolivia's inflation-output dynamics. The lack of consistent relationships across departments motivates the regional panel estimation strategy developed in the next section.

3.2 Constructing a Non-Tradables Inflation Indicator

To isolate the component of inflation that is most plausibly priced locally, we construct a department-level non-tradables (NT) inflation series from CPI microdata. The underlying CPI microdata come from the Bolivian Institute of National Statistics (INE) and span two official CPI baskets: Base 2007 (January 2008–December 2017) and Base 2017 (January 2018 onward). Because weights and index availability differ across baskets, the construction proceeds in two steps within each basket and then merges the resulting NT series into a single continuous quarterly panel. Since consistent product-level CPI microdata are only available from 2008 onward, the tradable and non-tradable inflation series start in 2008Q1. All estimations using these series therefore rely on this shorter balanced sample window.

First, within each basket we assign CPI products to tradable (T) and non-tradable (NT) groups using a product-level concordance. For the Base 2007 basket, microdata are available as variety-level average prices, which we aggregate to product-level indices before applying product weights. For the Base 2017 basket, product-level indices and expenditure weights are available directly at the department level. Using the resulting product indices and weights, we compute department-level group indices as weighted averages of product indices within the NT set (and analogously for tradables), and then splice the Base 2007 and Base 2017 group series to obtain a consistent NT index over time. Appendix B documents the classification, aggregation, and merging steps in detail and reports validation exercises.

Second, we convert monthly group indices into quarterly inflation rates using a three-month change and seasonally adjust the resulting quarterly series using the Census X-11 procedure, consistent with the treatment of headline and core inflation in Section 3.1. The final outcome is a department-level panel of quarterly non-tradables inflation, which we use in the sectoral Phillips Curve estimations in Section 5 to assess whether locally priced inflation exhibits a stronger relationship with local slack.

Figure 4 plots quarter-on-quarter non-tradables inflation against the output gap by department and illustrates that, even for NT inflation, the contemporaneous relationship with slack is weak in most departments. This is especially informative: non-tradables are the component most plausibly priced locally, so a flat relationship here suggests the weak Phillips Curve is not merely an artifact of CPI composition or tradables masking local dynamics. In other words, even the "most local" inflation measure does not rescue the Phillips Curve. This conclusion should nonetheless be read alongside a data construction caveat. The NT series is assembled through several processing steps — basket splicing, limited outlier replacement, boundary imputation, and X-11 seasonal adjustment — each of which has the potential to smooth out the specific regional variation needed to identify the slope. The reconstruction fit is also weaker for Cochabamba and Tarija, as reported in Table A.2. We cannot fully exclude the possibility that the NT-null result is partly attributable to attenuation introduced by the pipeline rather than reflecting the underlying economics alone, and we encourage readers to weight the NT evidence alongside the headline and core results, which rest on a simpler and more direct data construction.

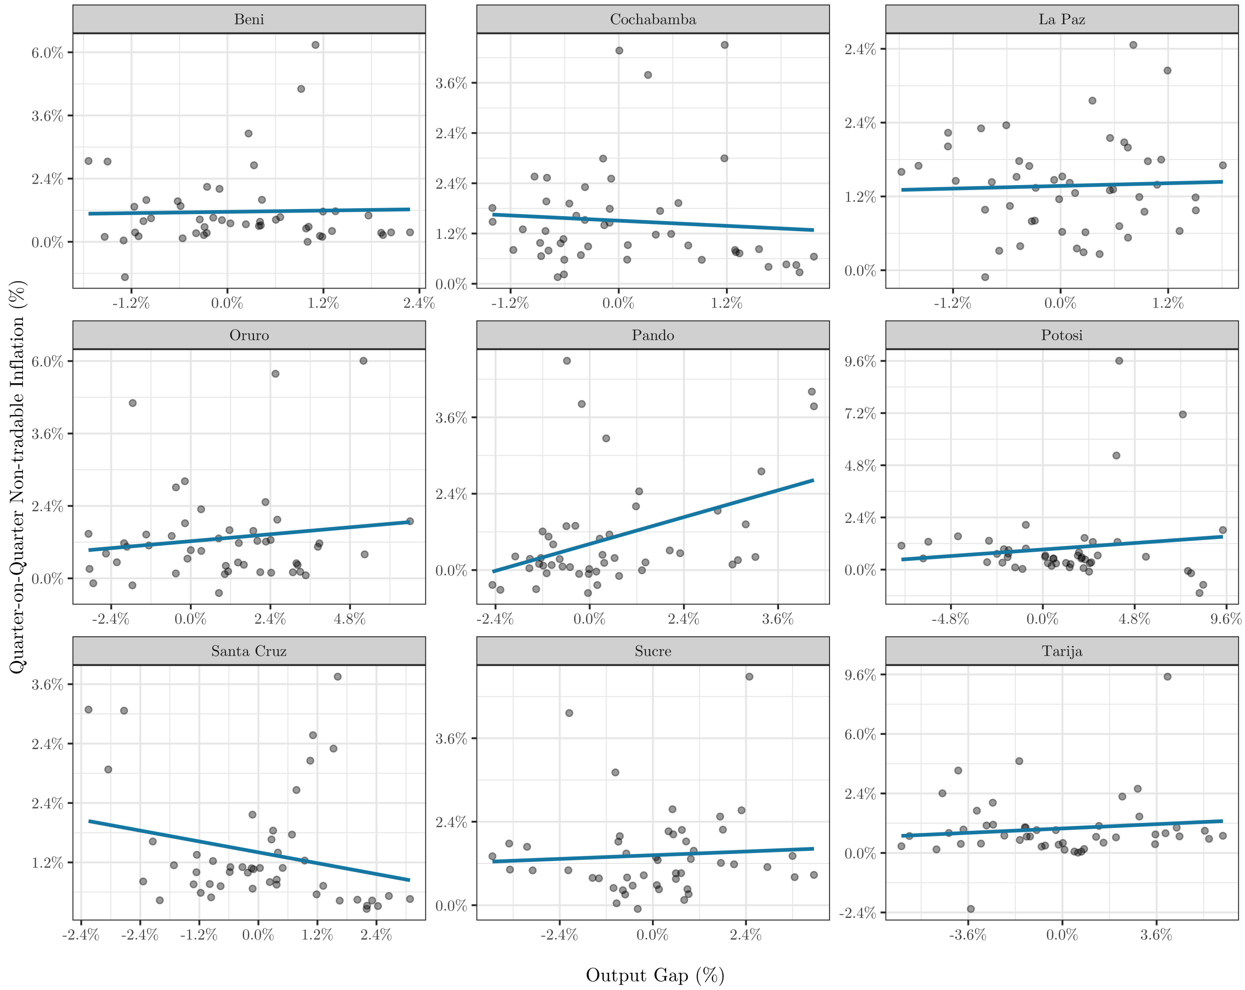

Figure 4: Quarterly (QoQ) non-tradable inflation vs. the output gap by department. Each dot is a quarter; the blue line is a linear fit. Common y-axis limits across panels facilitate cross-department comparison.

4 Methodological Framework

4.1 Two-Way Fixed Effects Estimation

To study the relationship between department-level inflation and local economic slack, we begin with a two-way fixed effects (TWFE) panel regression specified as

$$\pi_{i,t} = \alpha_i + \delta_t + \kappa \left(y_{i,t} - y_{i,t}^n \right) + u_{i,t},$$

where $\pi_{i,t}$ denotes the inflation rate—either headline or core—for department i at time t; α_i captures department-specific unobserved heterogeneity through fixed effects; and δ_t represents time fixed effects, controlling for common national shocks and shared inflation expectations. The output gap, represented by $\left(y_{i,t} - y_{i,t}^n \right)$, is measured as the percentage deviation of observed output

$y_{i,t}$ from its estimated potential $y^n_{i,t}$ (throughout, we use the terms "output gap" and "slack" interchangeably). The parameter of interest, κ, reflects the sensitivity of inflation to cyclical economic conditions. The error term $u_{i,t}$ encapsulates residual variation at the department level.

We estimate the TWFE specification over the full sample and over subperiods motivated by data coverage and a policy-relevant pre-period ending in 2011Q3. Specifically, we report results for the pre-period ($\leq$ 2011Q3) and for a balanced pre-window starting when department coverage becomes complete (2008Q2–2011Q3). To mitigate sample-composition concerns, we also estimate the model over a balanced long window (2008Q2–2019Q4) as well as over the full available sample. This structure allows us to separate pre-period dynamics from coverage-driven variation while maintaining comparability across specifications.

Table 3: Baseline Regression Results

	$\leq$ 2011 Q3	**2008 Q2 - 2011 Q3**	**2008 Q2 - 2019 Q4**	**Full Sample**
Headline Inflation				
Output Gap	-0.019	-0.069	0.002	0.002
	(0.025)	(0.052)	(0.015)	(0.012)
R^2	0.87	0.85	0.68	0.74
N	303	129	426	600
Core Inflation				
Output Gap	-0.019	0.003	0.006	-0.002
	(0.020)	(0.061)	(0.006)	(0.006)
R^2	0.80	0.76	0.70	0.78
N	303	129	426	600

Notes: Robust department-clustered standard errors in parentheses.

Significance levels: $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

While the baseline TWFE specification provides a transparent benchmark for the contemporaneous inflation-slack relationship, it abstracts from inflation persistence. To capture these dynamics, we extend the baseline model by adding lagged inflation, yielding a hybrid NKPC that allows inflation to depend on both current slack and backward-looking components. The next section presents the specification and estimates.

4.2 Hybrid New Keynesian Phillips Curve

This hybrid specification allows us to measure inflation persistence through the parameter γ, which captures how much past inflation influences current price changes.

$$\pi_{i,t} = \alpha_i + \delta_t + \kappa \left(y_{i,t} - y^n_{i,t}\right) + \gamma \pi_{i,t-1} + u_{i,t}, \tag{1}$$

where γ captures the degree of inflation persistence within each department. A statistically significant estimate of γ would indicate that past inflation plays a systematic role in current price dynamics, consistent with models incorporating rule-of-thumb price setters or staggered contracts. The department fixed effects α_i absorb time-invariant differences across departments (such as long-run price levels, sectoral composition, or persistent measurement differences), while the time fixed effects δ_t absorb shocks and policy-related movements that are common to all departments in a given quarter. This is important because nationwide forces can generate inflation movements that are correlated with slack, confounding the interpretation of the inflation–slack slope; accounting for these common components is central to interpreting Phillips-curve comovement (McLeay and Tenreyro, 2020). Accordingly, κ should be read as the local-slack slope conditional on these fixed effects, not as a fully structural NKPC parameter.

Table 4 reports the results of this hybrid specification for both headline and core inflation across the same sample windows described above. As with the baseline model, the estimated coefficients on the output gap remain small and statistically insignificant across all periods. In contrast, lagged inflation is statistically significant for core inflation in the full sample and in the broader post-2008 window. Taken together, these estimates indicate that inflation persistence is a salient feature of regional core inflation dynamics, which is consistent with an important role for expectations formation and backward-looking behavior in price setting, while local slack adds limited explanatory power in these specifications.

Two finite-sample limitations of the dynamic specifications deserve acknowledgment. First, inference relies on department-clustered standard errors with only nine clusters, a number too small to guarantee that asymptotic justifications for cluster-robust variance estimation hold; the reported standard errors and significance levels should therefore be read with appropriate caution, and the null results are best interpreted alongside the broader pattern of near-zero point estimates across all specifications rather than on any single test statistic. Second, the inclusion of lagged inflation alongside department fixed effects introduces the possibility of Nickell bias, which in short panels can distort estimates of the persistence parameter γ toward zero. Given that our primary inferential interest is the output-gap slope rather than the persistence parameter, and that the output-gap slope estimates remain near zero even in the TWFE specifications that do not include lagged inflation, we do not believe this bias materially affects our main conclusions — but readers should treat the magnitude of the estimated persistence coefficients with corresponding caution.

Table 4: Hybrid New Keynesian Phillips Curve

	≤ 2011 Q3	**2008 Q2 - 2011 Q3**	**2008 Q2 - 2019 Q4**	**Full Sample**
Headline Inflation				
Output Gap	-0.020	-0.078	0.001	0.001
	(0.025)	(0.049)	(0.016)	(0.013)
Lagged Inflation	0.076	0.072	-0.047	-0.035
	(0.057)	(0.075)	(0.077)	(0.077)
R^2	0.87	0.86	0.68	0.74
N	302	128	425	599
Core Inflation				
Output Gap	-0.015	0.004	0.006	-0.001
	(0.019)	(0.059)	(0.005)	(0.005)
Lagged Inflation	0.200^{*}	0.196	0.224^{*}	0.246^{**}
	(0.071)	(0.115)	(0.073)	(0.062)
R^2	0.81	0.77	0.72	0.79
N	302	128	425	599

Notes: Robust department-clustered standard errors in parentheses.
Significance levels: $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

4.3 Pooled and Mean Group Estimators

To relax pooling restrictions and account for cross-sectional dependence, we estimate heterogeneous-slope versions of the hybrid NKPC using common correlated effects (CCE) (Pesaran, 2006). Given evidence of cross-sectional dependence in Table A.1, we implement the Mean Group CCE estimator (MG-CCE), which allows κ_i and γ_i to vary across departments while controlling for unobserved common factors via cross-sectional averages of inflation and slack (and their lags) (Pesaran and Smith, 1995; Pesaran, 2006). As a parsimonious benchmark, we also consider a Pooled Mean Group CCE variant (PMG-CCE), which imposes pooling restrictions on longer-run responses while allowing short-run dynamics and intercepts to differ across departments (Pesaran, 2006; Ditzen, 2021).

Table 5 reports Hausman tests comparing MG-CCE and PMG-CCE for headline and core inflation. For both inflation measures, the test statistics ($\chi^2 = 22.33$ for headline and $\chi^2 = 7.84$ for core) are statistically significant at conventional levels, rejecting coefficient homogeneity and favoring the heterogeneous-slope specification.[1]

[1] While we report a Hausman test to compare MG and PMG estimates, we do not apply formal dispersion-based slope homogeneity tests (Pesaran and Yamagata, 2008) due to the short time dimension of our panel. These tests rely on asymptotic properties that may not hold when T is small, leading to biased inference.

Table 5: Hausman Test Results: Headline vs Core Inflation (PMG vs MG)

		MG (b)	PMG (B)	Diff (b - B)		Std. Err.
Headline Inflation						
	Lagged Inflation	0.0902	0.1161	-0.0259		0.0101
	Output Gap	0.0179	0.0121	0.0058		0.0406
	$\chi^2(2)$			22.33	p-value = 0.0000	
Core Inflation						
	Lagged Inflation	0.1298	0.2074	-0.0776		0.0281
	Output Gap	-0.0200	0.0052	-0.0252		0.0322
	$\chi^2(2)$			7.84	p-value = 0.0198	

Note: MG estimates are consistent under both H_0 and H_a; PMG is efficient under H_0 but inconsistent under H_a.

Figure 5 complements the Hausman evidence by displaying department-level MG-CCE estimates of the output-gap slope κ_i for headline and core inflation (with 95% confidence intervals). The dispersion in point estimates across departments—including variation in sign—is consistent with heterogeneous regional inflation-slack comovement, even though many individual coefficients are imprecisely estimated.

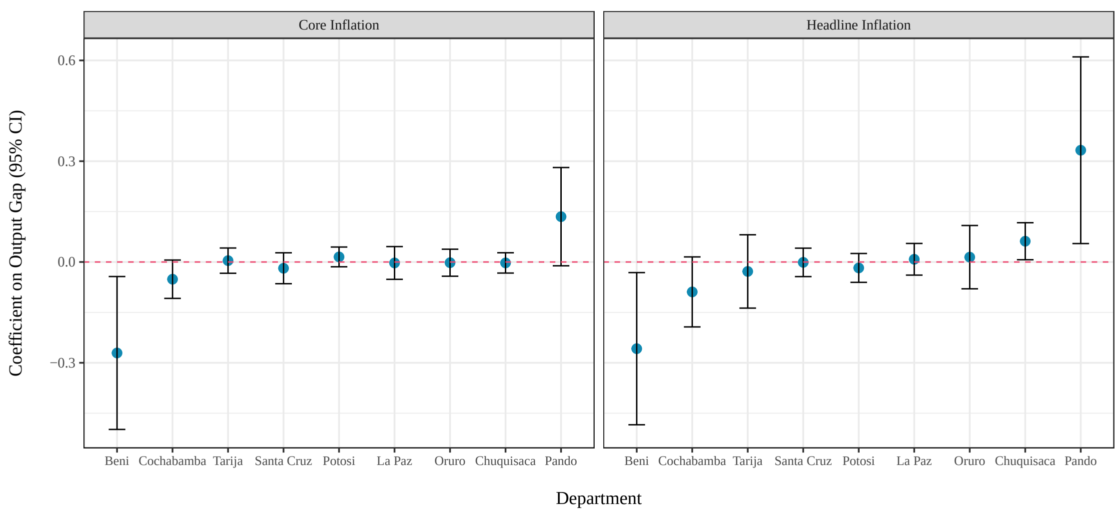

Figure 5: Estimated Department-Level Phillips Curve Slopes. This figure displays MG-CCE estimates of the output-gap coefficient κ_i for each of Bolivia's nine departments, shown separately for core and headline inflation with 95% confidence intervals. The results indicate substantial cross-department variation in the inflation-output gap relationship, particularly for core inflation.

Turning to the estimated coefficients, Table 6 reports MG-CCE and PMG-CCE estimates for headline and core inflation, focusing on lagged inflation and the output gap. For headline inflation, both estimators yield positive but statistically insignificant output-gap coefficients, consistent with a weak average relationship between cyclical conditions and inflation dynamics. For core inflation, lagged inflation is statistically significant under both estimators and larger under PMG-CCE, highlighting persistence in core inflation across regions. By contrast, the MG-CCE output-gap coefficient is negative, suggesting that heterogeneous regional responses may partially offset one another when long-run pooling restrictions are imposed.

Table 6: Comparison of PMG and MG Estimations: Dynamic Common Correlated Effects (CCE) Model

		Pooled Mean Group (PMG)		**Mean Group (MG)**	
		Coef.	**Std. Err.**	**Coef.**	**Std. Err.**
Headline Inflation					
	Lagged Inflation	0.116^{+}	(0.067)	0.090	(0.068)
	Output Gap	0.012	(0.018)	0.018	(0.044)
Core Inflation					
	Lagged Inflation	0.207^{***}	(0.054)	0.130^{**}	(0.059)
	Output Gap	0.005	(0.021)	-0.020	(0.038)
R^2		0.88		0.82	
N		582		582	

Notes: All models estimated using ‘xtdcce2‘ with one cross-sectional lag.

Cross-sectional averages of Inflation and Output Gap included.

PMG constrains coefficients to be homogeneous across panels; MG allows heterogeneity.

Significance levels: $^{+}p < 0.1$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

5 Non-Tradables Phillips Curve Dynamics

5.1 Two-Way Fixed Effects Estimation

In this section, we revisit our baseline Phillips-curve evidence through the lens of tradable and non-tradable inflation. We use tradable vs non-tradable inflation as an open-economy diagnostic for whether the weak regional Phillips Curve reflects pricing that is largely common across locations. In an open economy, a large share of the CPI basket is influenced by forces that are common across locations—import prices, exchange-rate movements, and nationally integrated supply chains (Edwards and Cabezas, 2022)—so observed regional inflation may respond only weakly to regional economic conditions even if aggregate slack matters in principle. If tradable prices are dominated by these common forces, then pooling them into headline inflation can mechanically attenuate regional slack coefficients. By contrast, non-tradable prices are more likely to be set locally and to reflect domestic marginal costs, suggesting that non-tradable inflation should be a cleaner object for detecting a relationship between domestic slack and inflation. At the same time, if tradable inflation is periodically shaped by regime-dependent pass-through or other common shocks, slack sensitivity—if detectable at all—may appear episodically in tradables in specific subperiods.

The key result is that this "cleaner" non-tradables object still shows essentially no slack sensitivity, strengthening the conclusion that local demand pressure plays a limited role in Bolivia's inflation dynamics. This interpretation is consistent with the identification issues emphasized by McLeay and Tenreyro (2020), who show that observed inflation–slack correlations can be muted or distorted when policy and other aggregate forces respond endogenously to the economy, underscoring the need to separate common components from locally driven variation when assessing Phillips-curve slopes.

Motivated by this logic, we construct department-level tradable (T) and non-tradable (NT) inflation from CPI microdata (Appendix B) and re-estimate the baseline specification separately by type of item. We focus on the balanced sample (2008Q1–2019Q4) and report results by subperiods split at 2012Q1 (i.e., 2008Q1–2011Q4 and 2012Q1–2019Q4) to assess whether the inflation-slack relationship is stronger for non-tradables and whether it differs across subsamples.

Table 7: Two-Way Fixed Effects Estimation

	$<$ 2012 Q1	**$\geq$ 2012 Q1**	**Full sample**
Non-Tradables			
Output Gap	-0.050	0.021	0.013
	(0.061)	(0.029)	(0.018)
R^2	0.69	0.31	0.60
N	144	288	432
Tradables			
Output Gap	0.017	0.044*	0.001
	(0.067)	(0.014)	(0.015)
R^2	0.84	0.59	0.78
N	144	288	432

Robust department-clustered standard errors reported in parentheses.
Significance levels: $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 7 reports TWFE estimates of the output-gap slope for non-tradable and tradable inflation. For non-tradables, the estimated slopes are small and statistically indistinguishable from zero across subsamples, providing little evidence that locally priced inflation responds systematically to local slack. For tradables, the slope is close to zero in the pre-2012 and full samples but becomes positive and statistically significant in the post-2012 fixed-regime period, suggesting that any detectable sensitivity to local slack is episodic and concentrated in tradable inflation.

5.2 Hybrid New Keynesian Phillips Curve

Table 8 adds lagged inflation to the previous TWFE regressions and leaves the slack results essentially unchanged, while clarifying differences in persistence by tradability. For non-tradables, the output-gap coefficient remains small and statistically insignificant across all subsamples, and lagged inflation is also imprecisely estimated, providing little evidence of either slack sensitivity or persistence in NT inflation. For tradables, the output-gap slope is positive and statistically significant only in the post-2012 fixed-regime period, while lagged inflation is statistically significant in the pre-2012 period and in the full sample, consistent with greater persistence in tradable inflation.

Table 8: Hybrid New Keynesian Phillips Curve

	$<$ 2012 Q1	**$\geq$ 2012 Q1 (Fixed)**	**Full sample**
Non-Tradables			
Output Gap	-0.039	0.024	0.013
	(0.053)	(0.029)	(0.016)
Lagged Inflation	0.089	-0.032	0.072
	(0.059)	(0.040)	(0.061)
R^2	0.69	0.32	0.60
N	143	287	431
Tradables			
Output Gap	0.032	0.044*	0.003
	(0.075)	(0.016)	(0.015)
Lagged Inflation	0.118*	-0.025	0.079*
	(0.041)	(0.040)	(0.025)
R^2	0.85	0.59	0.78
N	143	287	431

Notes: Robust department-clustered standard errors reported in parentheses. Significance levels: $^+p < 0.10$, $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

5.3 Pooled and Mean Group Estimators

Using the DCCE framework from Section 4.3, we re-estimate the tradables/non-tradables specifications to assess whether common shocks and slope heterogeneity differ by item type. We report MG estimates alongside PMG and pooled benchmarks, and use these comparisons to gauge whether pooling restrictions are empirically defensible and whether the weak inflation-slack relationship reflects a genuinely flat Phillips Curve or instead masks heterogeneous regional responses.

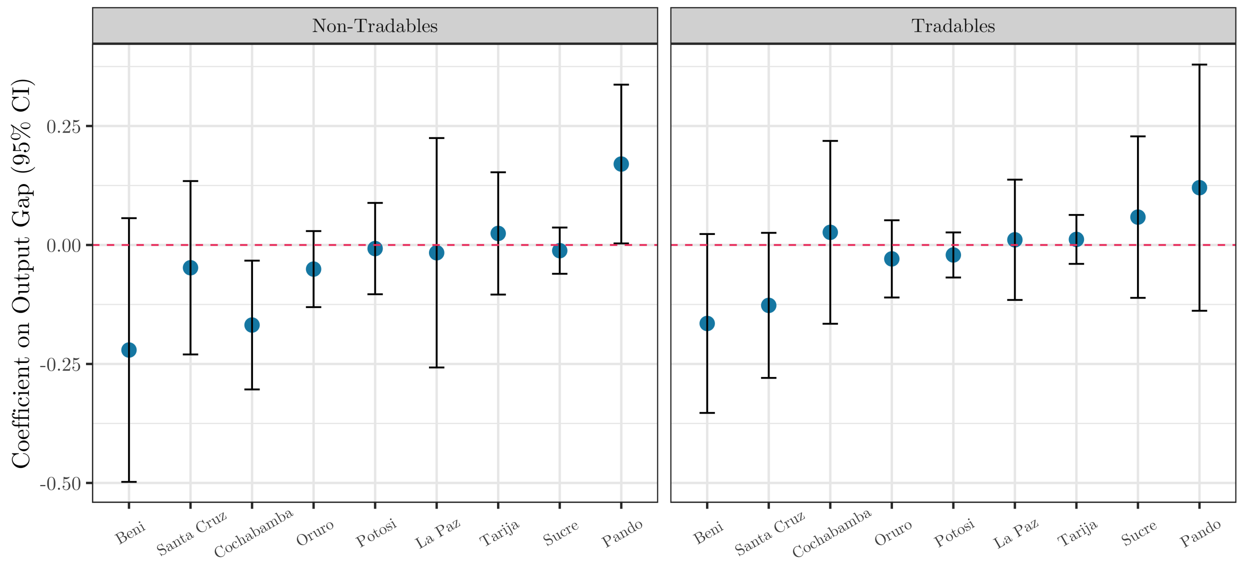

Figure 6: **Dynamic Common Correlated Effects Mean Group Estimates by Department (Full Sample)**
We estimate heterogeneous Phillips Curves using the Dynamic Common Correlated Effects Mean Group estimator. For each department i, we estimate a time-series regression of inflation on the output gap and lagged inflation, augmented with cross-sectional averages of inflation and the output gap to control for unobserved common factors. We include one lag of the cross-sectional averages ($K = 1$). The estimated model is:

$$\pi_{i,t} = \alpha_i + \kappa_i(y_{i,t} - y^n_{i,t}) + \gamma_i \pi_{i,t-1} + \lambda_{i,0}\bar{\pi}_t + \lambda_{i,1}\bar{\pi}_{t-1} + \theta_{i,0}\overline{(y-y^n)}_t + \theta_{i,1}\overline{(y-y^n)}_{t-1} + u_{i,t}.$$

where $\pi_{i,t}$ is inflation, $(y_{i,t} - y^n_{i,t})$ is the output gap, $\bar{\pi}_t$ is the cross-sectional average of inflation, and $\overline{(y-y^n)}_t$ is the cross-sectional average of the output gap (computed across departments in each quarter). The plotted coefficients are the estimated $\hat{\kappa}_i$ from each department-specific regression, with heteroskedasticity-and-autocorrelation-consistent (Newey–West) standard errors and 95% confidence intervals. The dashed line at zero aids visual comparison.

Table 9: DCCE Phillips Curve estimates by period and tradability

	Non-tradable			Tradable		
	$< 2012Q1$	$\geq 2012Q1$ (Fixed)	Full sample	$< 2012Q1$	$\geq 2012Q1$ (Fixed)	Full sample
Panel A: Pooled Mean Group						
Lagged inflation	0.074	-0.046	0.074	0.124	-0.033	0.073^{**}
	(0.244)	(0.070)	(0.055)	(0.104)	(0.067)	(0.025)
Output gap	-0.132	0.021	0.008	0.020	0.035	0.001
	(0.180)	(0.110)	(0.036)	(0.135)	(0.026)	(0.009)
Panel B: Mean Group						
Lagged inflation	-0.240^{***}	-0.034	-0.005	-0.108	-0.075	0.062
	(0.059)	(0.074)	(0.041)	(0.073)	(0.069)	(0.040)
Output gap	-0.083	-0.031	-0.024	-0.017	0.007	-0.008
	(0.068)	(0.072)	(0.047)	(0.163)	(0.045)	(0.030)
Observations	135	279	423	135	279	423

Notes: Standard errors are in parentheses. $< 2012Q1$: 2008Q1–2011Q4; $\geq 2012Q1$: 2012Q1–2019Q4; full sample: 2008Q1–2019Q4. Each regression includes cross-sectional averages of inflation and the output gap (and their lags) to control for shocks common to all departments; coefficients on these controls are omitted from the table. Observation counts are lower because lagged terms remove initial periods. Significance levels: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Once common factors are controlled for in the DCCE framework, output-gap coefficients are consistently close to zero across tradables and non-tradables and across regimes, indicating a limited role for local slack on average. Inflation dynamics are more informative: tradable inflation exhibits some persistence in the full-sample PMG estimates, while non-tradables show evidence of short-run adjustment in the pre-2012 MG results; however, this estimate should be interpreted with caution, as the pre-2012 window provides only approximately fifteen usable time periods after lags are applied, a span too short to rule out finite-sample bias and overfitting in a parameter-rich dynamic specification. We therefore do not place structural weight on this particular persistence estimates. Overall, the results point to modest heterogeneity in persistence rather than robust differences in slack sensitivity.

6 Robustness

6.1 Sensitivity Analysis

A key concern in estimating the Phillips Curve relationship is the accuracy of the output gap as a proxy for economic slack (Gagliardone et al., 2023). In this paper, the output gap is derived from department-level GDP estimates, which are subject to measurement error—particularly in emerging market contexts with limited subnational data. To assess whether the baseline results are sensitive to plausible variation in output gap measurement, we conduct a simulation-based robustness analysis using alternative specifications.

Building upon Chalup and Escobar (2023), we construct alternative output gap series using the 16th and 84th percentiles of department-level real GDP growth rates. These percentiles are used to bound the plausible range of variation in regional output dynamics. We then simulate 500,000 alternative output gap series within this range and re-estimate the baseline regression model for each draw.

Table 10 presents the proportion of simulations in which the estimated t-statistic on the output gap coefficient exceeds the critical value of 1.645, corresponding to a one-sided significance test at the 5% level, consistent with testing for a positive Phillips Curve slope. Across all time periods and inflation measures, the output gap remains statistically insignificant in the vast majority of simulations. For headline inflation, significance is observed in less than 1% of simulations in all cases. For core inflation, the rejection rate in the post-2008 subperiods is approximately 5.3–5.5%, which is indistinguishable from the nominal size of the test and therefore provides no evidence of a positive relationship; in the full sample, the rate falls to 0.39%, well below the 5% threshold.

Table 10: Proportion of t-statistics exceeding the critical value of 1.645

	Inflation	
Regression	Headline	Core
$\leq$ 2011 Q3	0.06%	0.25%
2008 Q2 - 2011 Q3	0.14%	5.53%
2008 Q2 - 2019 Q4	0.19%	5.35%
Full sample	0.20%	0.39%

Figure 7 displays the distribution of estimated output gap coefficients for each subperiod and inflation measure. The distributions are centered around zero, and in all cases the mass of the coefficients lies in the statistically insignificant range. For subperiods with slightly higher proportions of significant results (notably core inflation in 2008–2019), the distributions appear wider, suggesting greater model uncertainty but no robust directional effect.

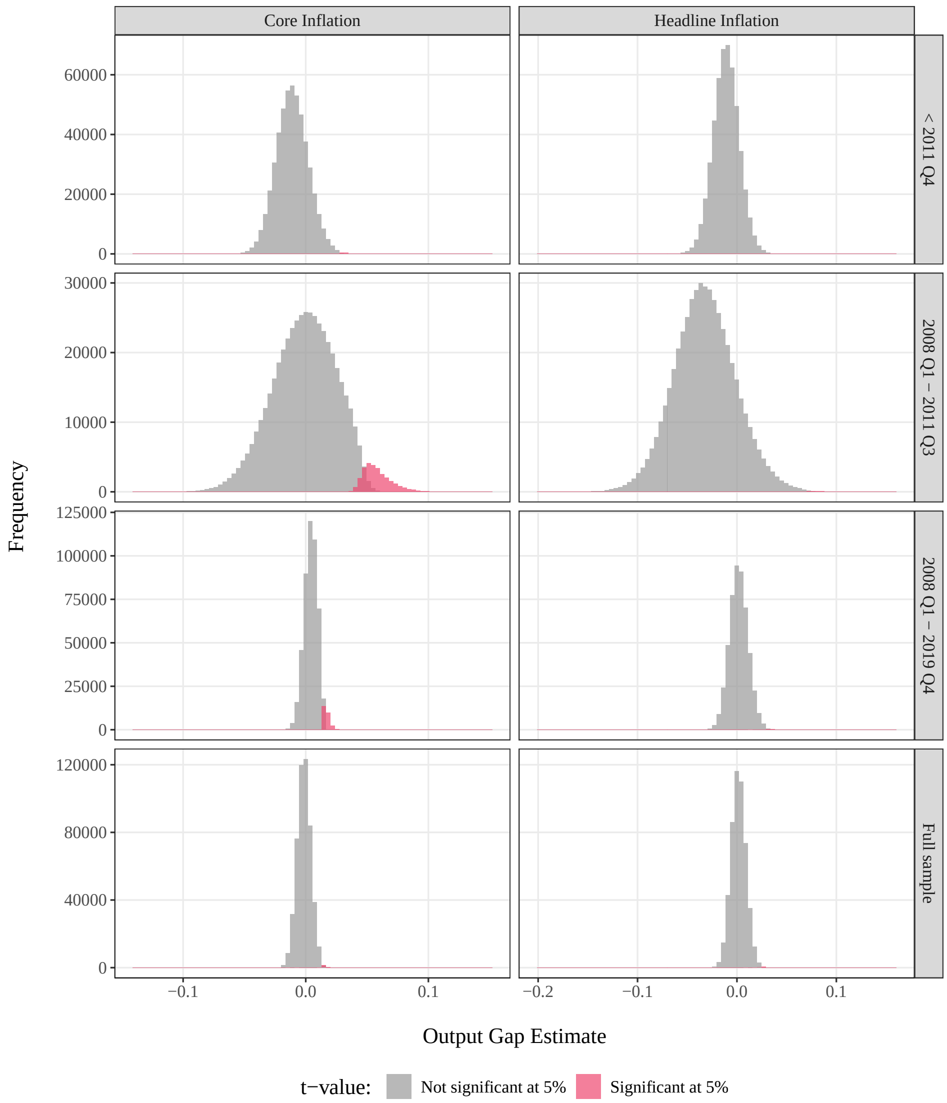

Figure 7: Distribution of Estimated Coefficients from Robustness Checks. Simulated output gap estimates were used to re-estimate the baseline model 500,000 times. Bars show frequency distributions by inflation type and subperiod.

Taken together, these results suggest that the baseline finding of a weak or nonexistent inflation-slack relationship in Bolivia remains broadly consistent across reasonable alternative constructions of the output gap. While no single specification can definitively rule out cyclical influences on inflation, the overall pattern points to a limited and unstable relationship between inflation and output fluctuations—particularly in the case of headline inflation.

7 Conclusion

This paper examines whether inflation responds systematically to local economic slack in a small open economy, using quarterly department-level data for Bolivia from 1993 to 2019. We find little evidence that it does. Across two-way fixed effects, hybrid, and dynamic common correlated effects specifications, the output-gap coefficient is consistently small — ranging from -0.069 to 0.006 across baseline estimates — and statistically indistinguishable from zero for both headline and core inflation. This conclusion holds across subsamples, survives 500,000 simulation draws of alternative output-gap series (yielding significant t-statistics in fewer than 1% of draws for headline inflation and under 0.4% in the full sample for core), and is not restored when inflation is decomposed into tradable and non-tradable components: even non-tradable inflation — the component most plausibly priced locally — shows no systematic sensitivity to departmental slack. The one partial exception is tradable inflation in the post-2012 fixed-rate period, where a positive slope of 0.044 is statistically significant in the TWFE and hybrid specifications, but this result does not survive the DCCE framework and should be read as episodic rather than structural. Core inflation does exhibit meaningful persistence — lagged core inflation enters significantly at the 0.20–0.25 range across the hybrid and CCE specifications — consistent with backward-looking price-setting behavior, though this persistence operates independently of local slack.

Two methodological results sharpen the interpretation. First, the regional panel rules out aggregation bias as an explanation: even when identification comes entirely from within-date cross-department variation and slopes are allowed to differ across departments, the average relationship remains flat. Second, the tradability decomposition rules out a compositional explanation: if the weak headline result merely reflected nationally priced tradables masking a genuine local response, non-tradables should show a stronger slope. They do not. The most plausible interpretation combines three institutional features that jointly suppress the local slack channel: Bolivia's fixed exchange rate since late 2011, which constrains monetary responses to domestic fluctuations; persistent dollarization and shallow financial markets, which limit interest-rate transmission; and a highly informal labor market, which we hypothesize acts as a shock absorber that decouples local demand pressure from wage and price dynamics. Unlike the first two conditions, which are documented directly in Bolivia's institutional record, the informality channel is inferred from the structure of Bolivia's labor market and from the persistence of the near-zero slope across specifications and inflation measures; it should therefore be read as a candidate mechanism consistent with the evidence rather than an independently established finding. Under these conditions, time fixed effects absorb the bulk of inflation variation, leaving little residual for local output gaps to explain.

Bolivia is not unusual in exhibiting this configuration. The same three conditions — fixed or heavily managed exchange rate, weak interest-rate transmission, and labor market informality — are jointly present across a wide range of emerging markets. In the CFA franc zone, where monetary policy is set regionally and administered prices are pervasive, regional Phillips Curve slopes should be expected to be similarly flat. Across Latin America, fully dollarized economies such as Ecuador, El Salvador, and Panama face an identical identification environment. Gulf Cooperation Council members with currency pegs and energy price subsidies satisfy all three conditions as well. In each case, the operational case for output-gap-based inflation forecasting is subject to the same challenge documented here. The practical implication is not that the Phillips Curve is irrelevant in these economies, but that its empirical slope — even when identified cleanly from subnational variation — is too flat and too unstable to anchor stabilization policy. Where the three conditions we identify are present, policymakers and forecasters should place greater weight on nominal anchors, expectations management, and external price indicators, and should treat output-gap projections with corresponding caution. More broadly, our results suggest a concrete research agenda: systematic cross-country comparison of Phillips Curve slopes against a small set of institutional indicators — exchange-rate regime flexibility, a measure of interest-rate pass-through, and the administered-price share of CPI — could yield a parsimonious diagnostic for where slack-based frameworks are likely to perform well and where they are not. Bolivia provides one well-identified data point for that comparison; the typology we propose can organize many others.

References

Aginta, H. (2023). Revisiting the phillips curve for indonesia: What can we learn from regional data? *Journal of Asian Economics*, 85:101592.

Behera, H., Wahi, G., and Kapur, M. (2018). Phillips curve relationship in an emerging economy: Evidence from india. *Economic Analysis and Policy*, 59:116–126.

Beraja, M., Hurst, E., and Ospina, J. (2019). The aggregate implications of regional business cycles. *Econometrica*, 87(6):1789–1833.

Berk, J. M. and Swank, J. (2011). Price level convergence and regional phillips curves in the us and emu. *Journal of International Money and Finance*, 30(5):749–763.

Chalup, M. S. and Escobar, L. F. (2023). Efectos macroeconǿmicos de la política fiscal durante la crisis del covid-19: evidencia de bolivia a nivel regional. *Revista de Economía del Rosario*, 26(1):1–39.

Coen, R., Eisner, R., Marlin, J., and Shah, S. (1999). The nairu and wages in local labor markets. *American Economic Review*, 89(2):52–57.

Ditzen, J. (2021). Estimating long-run effects and the exponent of cross-sectional dependence: An update to xtdcce2. *The Stata Journal*, 21(4):982–1007.

Edwards, S. and Cabezas, L. (2022). Exchange rate pass-through, monetary policy, and real exchange rates: Iceland and the 2008 crisis. *Open Economies Review*, 33:197–230.

El-Shagi, M. and Tochkov, K. (2024). Regional heterogeneity and the provincial phillips curve in china. *Economic Analysis and Policy*, 81:1036–1044.

Fan, J., Liao, Y., and Yao, J. (2015). Power enhancement in high-dimensional cross-sectional tests. *Econometrica*, 83(4):1497–1541.

Fitzgerald, T., Jones, C., Kulish, M., and Nicolini, J. P. (2020). Is there a stable relationship between unemployment and future inflation? Staff Report 614, Federal Reserve Bank of Minneapolis.

Fitzgerald, T., Jones, C., Kulish, M., and Nicolini, J. P. (2024). Is there a stable relationship between unemployment and future inflation? *American Economic Journal: Macroeconomics*, 16(4):114–42.

Frankel, J. A. (2010). Monetary policy in emerging markets: A survey. Working Paper 16125, National Bureau of Economic Research.

Gagliardone, L., Gertler, M., Lenzu, S., and Tielens, J. (2023). Anatomy of the phillips curve: Micro evidence and macro implications. NBER Working Paper 31382, National Bureau of Economic Research.

Hazell, J., Herreño, J., Nakamura, E., and Steinsson, J. (2022). The Slope of the Phillips Curve: Evidence from U.S. States. *The Quarterly Journal of Economics*, 137(3):1299–1344.

Hooper, P., Mishkin, F. S., and Sufi, A. (2020). Prospects for inflation in a high pressure economy: Is the phillips curve dead or is it just hibernating? *Research in Economics*, 74(1):26–62.

Juodis, A. and Reese, S. (2021). Cross-sectional dependence, factor structure and demeaning. *Journal of Econometrics*, 220(2):325–348.

Kapetanios, G., Price, S., Tasiou, M., and Ventouri, A. (2021). State-level wage phillips curves. *Econometrics and Statistics*, 18:1–11.

Machicado, C. G. (2018). De las causas próximas a las causas profundas del crecimiento económico de bolivia entre 1950 y 2015. Technical Report 09/2018, Institute for Advanced Development Studies.

Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical evidence on inflation expectations in the new keynesian phillips curve. *Journal of Economic Literature*, 52(1):124–188.

McLeay, M. and Tenreyro, S. (2020). Optimal inflation and the identification of the phillips curve. *NBER Macroeconomics Annual*, 34(1):199–255.

Mora-Barrenechea, M. (2021). Una revisiǿn a la Curva de Phillips en Bolivia. *Revista Latinoamericana de Desarrollo Econǿmico*, pages 159 – 188.

Morales, J. A. (2003). Dollarization of assets and liabilities: Problem or solution?. The case of Bolivia. *Revista de Análisis del BCB*, 6(1):7–39.

Morales, J. A. (2005). Las principales políticas del banco central en el marco de la ley 1670. In Banco Central de Bolivia, editor, *Historia Monetaria Contemporánea de Bolivia: Siete momentos capitales en los 77 años de historia del Banco Central de Bolivia*, pages 303–330. Banco Central de Bolivia, La Paz, Bolivia.

Morales, J. A. and Sachs, J. D. (1989). *Bolivia's Economic Crisis*, pages 57–80. University of Chicago Press, Chicago.

Murillo, A. (2014). Estimación de una curva de phillips neokeynesiana para bolivia. Documento de trabajo 05/2014, Banco Central de Bolivia.

Orlov, D. and Postnikov, E. (2022). Phillips curve: Inflation and nairu in the russian regions. *Journal of the New Economic Association*, 55(3):61–80. In Russian, with English summary.

Pesaran, M. H. (2006). Estimation and inference in large heterogeneous panels with a multifactor error structure. *Econometrica*, 74(4):967–1012.

Pesaran, M. H. (2015). Testing weak cross-sectional dependence in large panels. *Econometric Reviews*, 34(6-10):1089–1117.

Pesaran, M. H. and Smith, R. (1995). Estimating long-run relationships from dynamic heterogeneous panels. *Journal of Econometrics*, 68(1):79–113.

Pesaran, M. H. and Xie, Y. (2021). Diagnostic tests for cross-section dependence in panel data models. *Oxford Bulletin of Economics and Statistics*, 83(1):40–61.

Pesaran, M. H. and Yamagata, T. (2008). Testing slope homogeneity in large panels. *Journal of Econometrics*, 142(1):50–93.

Requena, J., Mendoza, R., Lora, O., and Escobar, F. (2002). La Política Monetaria del Banco Central de Bolivia. *Revista de Análisis del Banco Central de Bolivia*, 5:9 – 45.

Sachs, J. (1986). The bolivian hyperinflation and stabilization. Working Paper 2073, National Bureau of Economic Research.

Schuffels, J., Kool, C., Lieb, L., and van Veen, T. (2022). Is the slope of the euro area phillips curve steeper than it seems? heterogeneity and identification. Working Paper 10103, CESifo.

Valdivia, D. (2008). ¿es importante la fijación de precios para entender la dinámica de la inflación en bolivia? Development Research Working Paper Series 02/2008, Institute for Advanced Development Studies (INESAD), La Paz.

A Output-Gap Construction and Cross-Sectional Dependence Tests

The INE does not publish quarterly GDP data at the departmental level. Consequently, we employ the estimates of Chalup and Escobar (2023) to construct the regional output gap. Actual output is the departmental real GDP series provided by those authors; potential output is obtained with the Hodrick–Prescott filter using the conventional quarterly smoothing parameter $\lambda = 1600$. The output gap is defined as the difference between actual and potential output, expressed as a percentage of potential output.

Figure A.1 plots quarter-on-quarter real-GDP growth for each of Bolivia's nine departments. The grey band marks the 16th to 84th percentiles, highlighting the dispersion of regional growth rates.

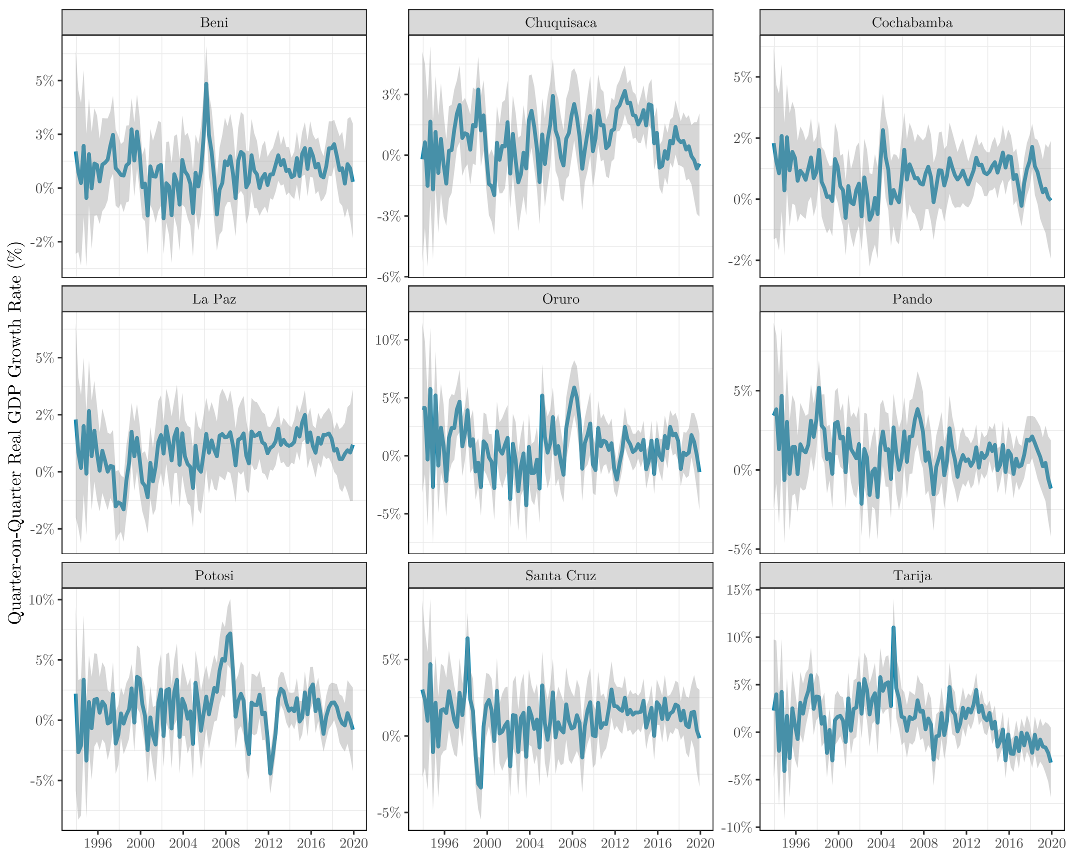

Figure A.1: Quarter-on-quarter real GDP growth by department

Table A.1 reports four different statistics that test for cross-sectional dependence (CSD) among departments in the main variables used later in the paper: the original CD test of Pesaran (2015), the bias-corrected CD_w of Juodis and Reese (2021), its power-enhanced version CD_w^+ (Fan et al., 2015), and the principal-component diagnostic CD^* of Pesaran and Xie (2021). All four tests reject the null of weak CSD for headline inflation, core inflation, their lags, and (in most cases) the output gap, indicating that common shocks or national factors shape regional inflation dynamics. These results motivate our use of the Mean-Group estimator with Common Correlated Effects (CCE), which controls for unobserved common factors while allowing slope heterogeneity across departments.

Table A.1: Cross-sectional dependence test results

Group	Variable	CD	CD_w	CD_w^+	CD^*
Headline Inflation	Inflation	44.63 (0.000)	-5.45 (0.000)	262.34 (0.000)	-7.16 (0.000)
	Lagged Inflation	44.73 (0.000)	-5.44 (0.000)	262.94 (0.000)	-7.08 (0.000)
	Output Gap	16.57 (0.000)	-1.93 (0.053)	127.85 (0.000)	-1.75 (0.080)
Core Inflation	Inflation	42.89 (0.000)	11.95 (0.000)	269.30 (0.000)	-7.24 (0.000)
	Lagged Inflation	42.89 (0.000)	11.91 (0.000)	269.25 (0.000)	-7.24 (0.000)
	Output Gap	16.57 (0.000)	5.69 (0.000)	135.47 (0.000)	-1.75 (0.080)

Notes: p-values in parentheses. CD = Pesaran (2015); Pesaran and Xie (2021); CD_w = Juodis and Reese (2021); CD_w^+ = Fan et al. (2015); CD^* = Pesaran and Xie (2021) using four principal components. Null hypothesis H_0: weak CSD; alternative H_1: strong CSD.

B Construction of Tradable and Non-Tradable Inflation Indices

The objective of this appendix is to document the construction of department-level inflation indicators for *tradable* (T) and *non-tradable* (NT) items. These series are the main inputs for the empirical analysis in the paper, which studies the relationship between NT inflation and the department-level output gap.

B.1 Conceptual requirements

Constructing a price index for non-tradables requires two ingredients:

1. A list of CPI items and their expenditure weights (at the most disaggregated level at which weights are available); and
2. A mapping that classifies each CPI item as *tradable* or *non-tradable*.

In Bolivia, both elements must be built and harmonized across two official CPI baskets (Base 2007 and Base 2017).

B.2 Tradability classification and national validation against ECLAC

We classify CPI products into tradables and non-tradables using expert judgment, producing two concordances: one for the Base 2007 product list and one for the Base 2017 product list. Let $g(p) \in \{T, NT\}$ denote the tradability status assigned to product p.

To assess the plausibility of this classification, we construct national tradable and non-tradable inflation series using product-level CPI indices and official national weights, and compare them with the corresponding tradable/non-tradable CPI series published by CEPAL/ECLAC. The resulting co-movement is close enough to support the use of the expert-based mapping in the subsequent department-level construction. Figure A.2 reports the national comparison.

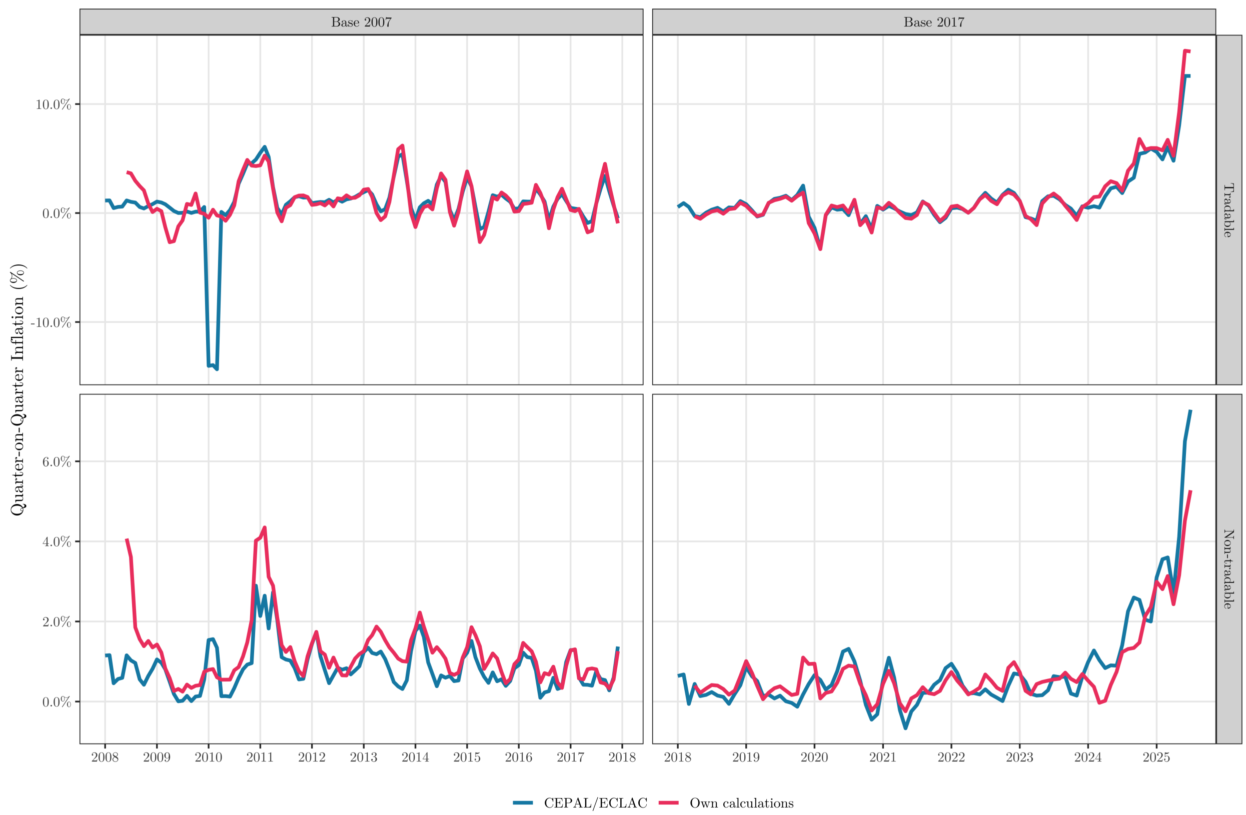

Figure A.2: National validation of the tradability classification: internal vs. ECLAC tradable/non-tradable inflation.

B.3 Data sources and CPI baskets

The construction relies on CPI microdata provided by INE under two baskets:

- **Base 2007:** January 2008 to December 2017. Prices are available at the *variety* level as average prices, and must be aggregated up to the *product* level to match the level at which weights are provided.
- **Base 2017:** January 2018 onward. Data are already reported as *product-level indices* by departament, together with product expenditure weights at the department level.

Because weights are available at the product level, the Base 2007 data require an explicit aggregation from varieties to products prior to constructing T/NT indices.

B.4 Base 2007: from variety prices to product indices

This subsection describes the construction of product indices for the Base 2007 basket, following the methodology outlined in INE's *Documento Metodológico – Índice de Precios al Consumidor: Cambio de año base 2007*.

B.4.1 Variety-level relatives across informants

Let P_{ijt} denote the observed price quote for variety i reported by informant j in month t. The elementary price relative is

$$R_{ijt} = \frac{P_{ijt}}{P_{ij,t-1}}. \tag{2}$$

B.4.2 Variety-level index (geometric mean across informants)

For a given variety i and month t, we aggregate informant-level relatives using a geometric mean:

$$I_{it} = \left(\prod_{j=1}^{n_{it}} R_{ijt} \right)^{1/n_{it}}, \tag{3}$$

where n_{it} is the number of informants reporting variety i at month t.

B.4.3 Implementation with INE-provided average prices

In practice, INE's microdata provide *average* prices at the variety level rather than the full set of informant-level quotes. Therefore, the implemented variety index corresponds to

$$I_{it} = \frac{\overline{P}_{it}}{\overline{P}_{i,t-1}}, \tag{4}$$

where $\overline{P}_{it}$ is the reported average price for variety i at month t. If INE's average is geometric, this expression coincides with the geometric-mean formula above; if it is arithmetic, the expression remains a coherent month-to-month relative based on the best available information.

B.4.4 Product-level aggregation

Let p denote a CPI product composed of a set of varieties $\mathscr{I}(p)$. Since CPI weights are available at the product level but not at the variety level, we aggregate varieties into a product index using an equal-weight geometric mean:

$$I_t^p = \exp\left(\frac{1}{|\mathscr{I}(p)|} \sum_{i \in \mathscr{I}(p)} \ln I_{it} \right). \tag{5}$$

Equivalently, this is a geometric mean with uniform variety weights. The absence of within-product variety weights implies that any alternative weighting scheme would introduce additional assumptions. Empirically, slightly more than one-third of products contain two or more varieties, making this step quantitatively relevant.

B.5 Validation via CPI reconstruction and limited outlier adjustments

To validate that the Base 2007 aggregation (varieties $\rightarrow$ products) is internally consistent, we reconstruct the overall department-level CPI and compare the implied inflation rates with the official series published by INE. This diagnostic assesses whether the chaining and aggregation steps reproduce the main dynamics of official inflation.

During this process, we implement a technical outlier-screening procedure at the variety–department series level, consisting of: (i) excluding effectively constant series; (ii) applying STL decomposition to log prices and extracting residuals; and (iii) flagging extreme residuals using a z-score threshold (absolute value above 10). While the screening can flag a broader set of candidates, in the final dataset we modify only a minimal set of observations based on economic interpretability and to prevent isolated spikes from mechanically distorting aggregate indices.

Specifically, we replace 11 monthly observations across 5 departments by the series-specific mean computed excluding flagged values, and we apply one additional rescaling (anchoring) adjustment for a regulated item to preserve continuity after a clear level shift. The number of modified observations is negligible relative to the size of the dataset.

Figure A.3 compares recalculated versus official inflation at the three-month horizon.

Figure A.3: Validation of Base 2007 aggregation: recalculated vs. official three-month inflation by department.

Overall, while the reconstructed CPI is not identical to the published series in every department, it reproduces the key inflation dynamics sufficiently well to support the construction of tradable and non-tradable indices at the department level. As shown in Table A.2, the goodness of fit is high for most cities, with R^2 values above 0.8 in the majority of cases, although the fit is weaker in Cochabamba and Tarija.

Table A.2: Goodness of Fit by Department

Department	**Observations (n)**	R^2
Pando	115	0.963
La Paz	115	0.927
Potosí	115	0.898
Oruro	115	0.893
Sucre	115	0.832
Beni	115	0.804
Santa Cruz	115	0.804
Cochabamba	115	0.557
Tarija	115	0.442

B.6 Merging baskets and applying the T/NT classification

Once product-level indices are available for Base 2007, we merge them with the product indices from Base 2017 and append the tradability classification $g(p)$ for each product in the corresponding basket.

At this stage, the dataset contains, for each basket separately, product indices by department and month and the associated product weights, together with the T/NT tag.

B.7 Department-level tradable and non-tradable indices

For each department c, month t, and group $g \in \{T, NT\}$, we compute a group-specific index as a weighted average of product indices within the group:

$$I_{c,g,t} = \sum_{p:g(p)=g} w_{c,p} I_{c,p,t}, \tag{6}$$

where $I_{c,p,t}$ is the product index and $w_{c,p}$ is the official CPI expenditure weight for product p in department c under the relevant basket. This aggregation is implemented *within basket* to avoid mixing weight systems across base years.

B.8 Quarterly inflation and imputation of boundary missing values

The empirical application uses a quarterly inflation measure. For each basket, department, and group g, we compute quarter-on-quarter inflation from monthly indices using a three-month change:

$$\pi^{QoQ}_{c,g,t} = \frac{I_{c,g,t} - I_{c,g,t-3}}{I_{c,g,t-3}}. \tag{7}$$

By construction, the first three months of each basket-specific series do not have a QoQ growth rate. To avoid losing these initial observations in downstream steps (particularly seasonal adjustment, which benefits from contiguous series), we impute the missing QoQ values using an automatic univariate ARIMA model estimated separately for each department$\times$group series (AutoARIMA, `fable` in R). Imputed points are flagged and used only to ensure continuity of the transformation pipeline. Specifically, these imputed QoQ values are used solely to run the X-11 seasonal adjustment on a contiguous series and are excluded from the estimation sample once the seasonally adjusted series has been constructed; the loss of the first quarter per department in the regression tables therefore reflects the inclusion of lagged terms only, not the removal of imputed observations. Figure A.4 illustrates observed versus imputed segments.

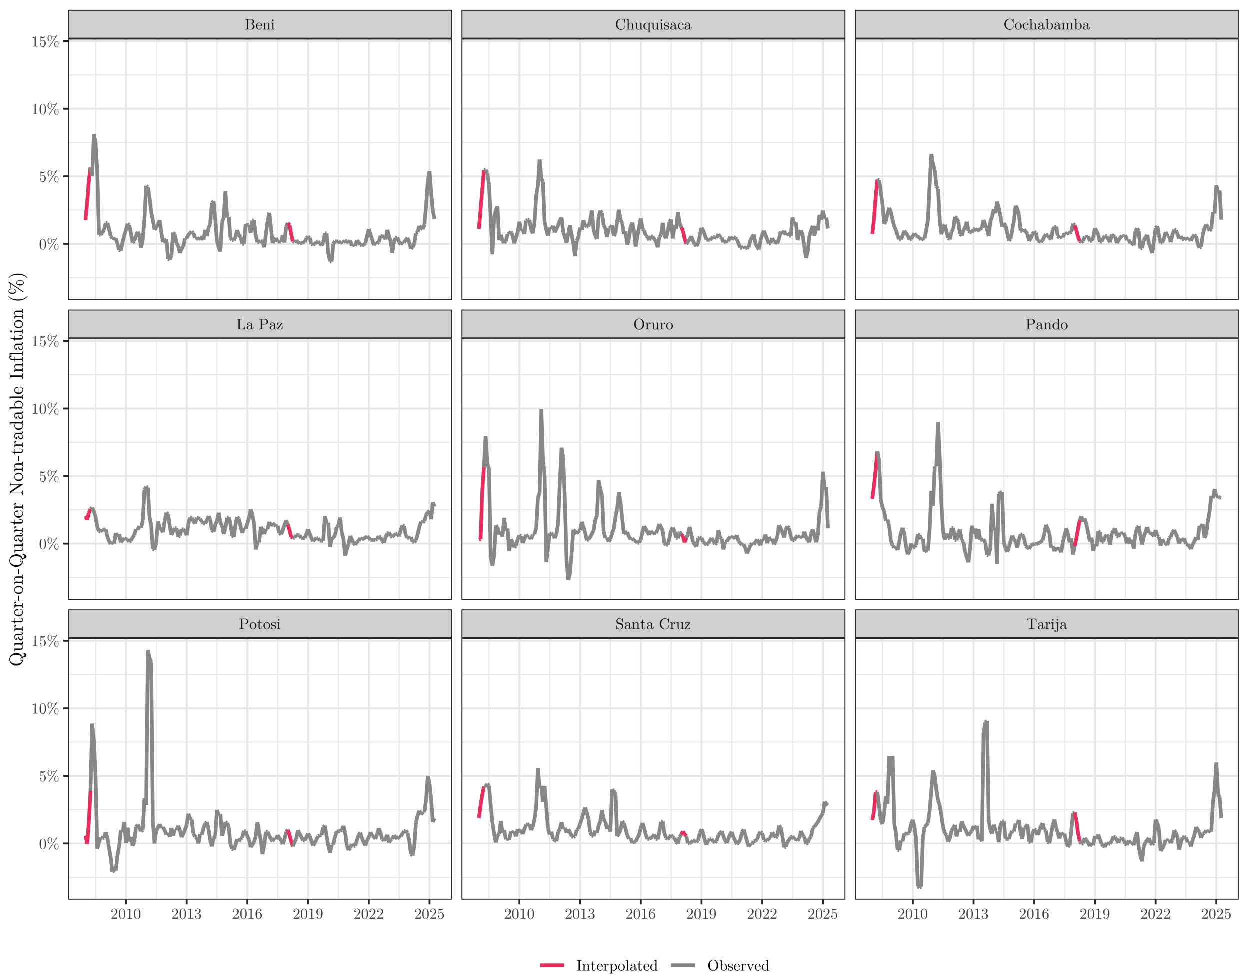

Figure A.4: QoQ inflation series: observed values and AutoARIMA imputations used to fill boundary missing observations.

B.9 Seasonal adjustment

Finally, we seasonally adjust the QoQ inflation series using Census X-11-ARIMA seasonal adjustment, applied to each department×group series. Let $\pi_{c,g,t}^{QoQ,SA}$ denote the seasonally adjusted series. Figure A.5 compares raw and seasonally adjusted QoQ inflation for non-tradables.

Figure A.5: Seasonal adjustment of QoQ NT inflation using X11: raw vs. seasonally adjusted series. The large spikes visible in some departments — most notably Potosí in the early part of the sample — reflect identifiable localized supply disruptions, most likely associated with civic conflicts or regional blockades, and are confirmed to be expressed in the same QoQ units defined in equation 7; no additional outlier treatment beyond the variety-level screening described in Appendix B.5 was applied to the NT series prior to the Phillips-curve regressions.

The final output is a department-level panel of QoQ inflation (raw and seasonally adjusted) for tradables and non-tradables. For the regression analysis in Section 5, we convert this to a quarterly frequency by averaging the monthly observations within each quarter for each department and group, aligning the inflation measure with the quarterly output gap.

Table A.3: Phillips Curve estimates with common-shock controls (pooled specification)

	Non-tradable			Tradable		
	$< 2012Q1$	$\geq 2012Q1$ (Fixed)	Full sample (2008Q1–2019Q4)	$< 2012Q1$	$\geq 2012Q1$ (Fixed)	Full sample (2008Q1–2019Q4)
Lagged inflation	0.134	0.045	0.100	0.159	0.027	0.100***
	(0.195)	(0.072)	(0.067)	(0.108)	(0.072)	(0.027)
Output gap	-0.015	0.032	0.004	-0.022	0.027	-0.000
	(0.064)	(0.039)	(0.022)	(0.027)	(0.023)	(0.010)
Observations	135	279	423	135	279	423

Notes: The dependent variable is quarterly inflation. Each regression includes cross-sectional averages of inflation and the output gap (and their lags) to account for sh departments; coefficients on these controls are not reported. Standard errors are in parentheses. Pre: 2008Q1–2011Q4; Post: 2012Q1–2019Q4; Full: 2008Q1–2019Q4. are lower because lagged terms remove initial periods. Significance levels: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.