When Bad Institutions Meet Good Ones: The Peruvian Puzzle

Carlos Chávez[*] Luis Yépez[†]

First complete draft: May 2026

Preliminary and incomplete. Comments welcome.

Abstract

Between 2016 and 2023 Peru cycled through seven presidents, a presidential self-coup, and deadly mass protests—yet its sovereign spread fell to historic lows. We document this *Peruvian Puzzle* over 2000–2025 and trace it to the credibility of the Central Reserve Bank of Peru under Governor Julio Velarde, appointed in 2006. Before 2006 the average political event moved Peru's sovereign spread by 41 basis points; afterward the response is indistinguishable from zero. The instability–spread correlation flips from $+0.55$ to -0.23, daily spread volatility falls more than threefold, and the market response decays with the governor's tenure. We rationalize these facts with a model combining monetary-policy delegation, sovereign-spread pricing, and Bayesian learning about the governor's type; it yields an insulation theorem, a variance-reduction ratio, and a tenure-decay coefficient, each matched to an estimate. An endogenous structural-break test rejects a sharp 2006 transition in favor of gradual accumulation through the Velarde era, and the decoupling survives global controls, an externally coded political-risk measure, central-bank swap intervention, a synthetic control, and a six-country panel supporting the inflation channel. Peru is an existence proof that one credible institution, sustained long enough, can carry a disproportionate share of macroeconomic stability when surrounding institutions fail.

1 Introduction

Between 2000 and 2025, Peru underwent two transformations that should have moved together. Political instability rose sharply: the country cycled through seven presidents in seven years between 2016 and 2023. It witnessed multiple impeachments, a presidential

[*]Center for Economics of Human Development, University of Chicago. E-mail: cchavezpadilla@uchicago.edu.

[†]Department of Economics, University of Chicago. E-mail: luisyepezsa@uchicago.edu.

self-coup attempt, mass protests that left dozens dead, and the unprecedented spectacle of a president arrested while attempting to seek asylum at the Mexican embassy. By any conventional measure, Peru's political institutions failed catastrophically. Yet the market response to that instability vanished. Before 2006, the average political event moved Peru's sovereign spread by 41 basis points, afterward, the response is statistically indistinguishable from zero. We call this the *Peruvian Puzzle*: the decoupling of political risk from sovereign borrowing costs.

The decoupling is equally stark in the level of borrowing costs. The Emerging Markets Bond Index Global (EMBIG) spread for Peru—a standard measure of country risk—fell from over 800 basis points in 2000 to approximately 130 basis points by 2025, even as political turmoil intensified. The standard view that political risk raises emerging-market borrowing costs (Mauro et al., 2007) simply stopped describing Peru. What changed? In October 2006, Julio Velarde was appointed Governor of the Central Reserve Bank of Peru (BCRP). He has held the position continuously for nearly two decades, building an internationally recognized reputation for strict inflation targeting, substantial reserve accumulation, and operational independence from political interference. We show that this accumulated institutional credibility is what reorganized how global financial markets process political news about Peru.

The paper makes two contributions. First, we document the decoupling and trace its timing. Before Velarde's appointment the correlation between political instability and sovereign spreads was positive and substantial ($\rho = 0.55$), afterward it is negative ($\rho = -0.23$). Crucially, the transformation is *gradual* rather than abrupt: it accumulates through the Velarde era rather than switching on at any single date—the signature of credibility built up over time, not of a one-off regime change. The decoupling survives a battery of challenges: controls for global risk appetite and terms of trade, direct controls for central-bank swap intervention, an externally coded measure of political risk constructed with no knowledge of Peruvian markets, and a synthetic-control comparison against a weighted donor pool of

Latin American economies.

Second, we develop a formal model that explains these patterns. The model combines three existing frameworks in a clearly delineated way: the monetary-policy delegation model of Dziuda and Pflueger (2025, henceforth D-P), sovereign-spread pricing in the tradition of Arellano (2008), and Bayesian learning about the central banker's type in the spirit of Cukierman and Meltzer (1986). A credible (inflation-averse) central bank reduces the pass-through of government-quality shocks to inflation, lower inflation pass-through dampens the response of sovereign spreads to political events, and learning about the governor's type generates a gradual tenure effect. The model yields three predictions with closed-form counterparts—an insulation theorem, a variance-reduction ratio, and a tenure-decay coefficient—each of which maps to an empirical estimate.

The findings speak to a broader question in political economy: can good institutions substitute for bad ones? The Peruvian case suggests they can, at least partially. One credible institution—the BCRP under Velarde—has insulated the economy from the dysfunction of the political system. Yet this insulation is fragile: it depends substantially on the continuation of a specific individual and the policy regime he represents.

The remainder of the paper proceeds as follows. Section 2 describes the data. Section 3 develops the model. Section 4 presents the empirical strategy. Section 5 reports the main results. Section 6 tests the model's specific predictions. Section 7 subjects the decoupling to global controls, direct controls for central-bank swap intervention, an externally coded political-risk measure, synthetic control, and cross-country comparison. Section 8 concludes. Appendix A collects model derivations, Appendix B develops the multiplicity and fragility results, and Appendix C reports endogenous structural-break tests.

2 Data

We construct a dataset combining daily financial market data with a comprehensive coding of political events in Peru from 2000 to 2026. This section describes the sources, construction, and key features of the data.

2.1 Financial Market Data

The primary measure of sovereign risk is the Emerging Markets Bond Index Global (EMBIG) spread for Peru, obtained from the Central Reserve Bank of Peru (BCRP). The EMBIG spread measures the yield differential between Peruvian dollar-denominated sovereign bonds and comparable U.S. Treasury securities, expressed in basis points. This measure is widely used as an indicator of country risk and investor sentiment toward emerging markets.

The dataset covers February 1, 2000 to January 8, 2026, comprising 4,508 trading days with 99.7% coverage for EMBIG Peru (4,495 observations). We complement the Peruvian data with EMBIG spreads for Latin American comparators—Argentina, Brazil, Chile, Colombia, and Mexico—as well as the aggregate EMBIG Latin America index. These regional series allow us to isolate Peru-specific responses from common regional shocks.

Table 1 presents summary statistics. The mean EMBIG spread for Peru over the sample period is 240 basis points, ranging from 91 to 901 basis points. For comparison, the mean EMBIG Latin America spread is 472 basis points, indicating that Peru has generally traded at a significant discount to the regional average.

Table 1: Summary Statistics: Financial Variables

Variable	Obs.	Mean	Std. Dev.	Min	Max
EMBIG Peru (bps)	4,495	240.2	153.3	91	901
EMBIG Latin America (bps)	3,941	471.9	197.2	198	1,847
EMBIG Spread (Peru – LatAm)	3,941	−231.8	104.8	−694	126
Exchange Rate (S/ per USD)	3,775	3.29	0.35	2.57	4.13
Δ EMBIG Peru (daily)	4,494	−0.08	10.80	−125	176
Δ Exchange Rate (daily)	2,918	−0.0004	0.01	−0.07	0.08

Notes: Sample period is February 2000 to January 2026. EMBIG data from JP Morgan via BCRP. Exchange rate is the interbank midpoint rate.

Figure 1 displays the central puzzle. The figure plots EMBIG Peru from 2000 to 2025, with political events marked as dots. Two patterns emerge. First, sovereign spreads declined dramatically over the sample period. Second, the frequency of political events *increased* after 2006, particularly in 2016–2023. Yet this political turmoil coincided with historically low sovereign spreads. A dual-axis version of this comparison appears in Online Appendix Figure OA.6.

Figure 1: The Peruvian Puzzle: Sovereign Spreads and Political Events

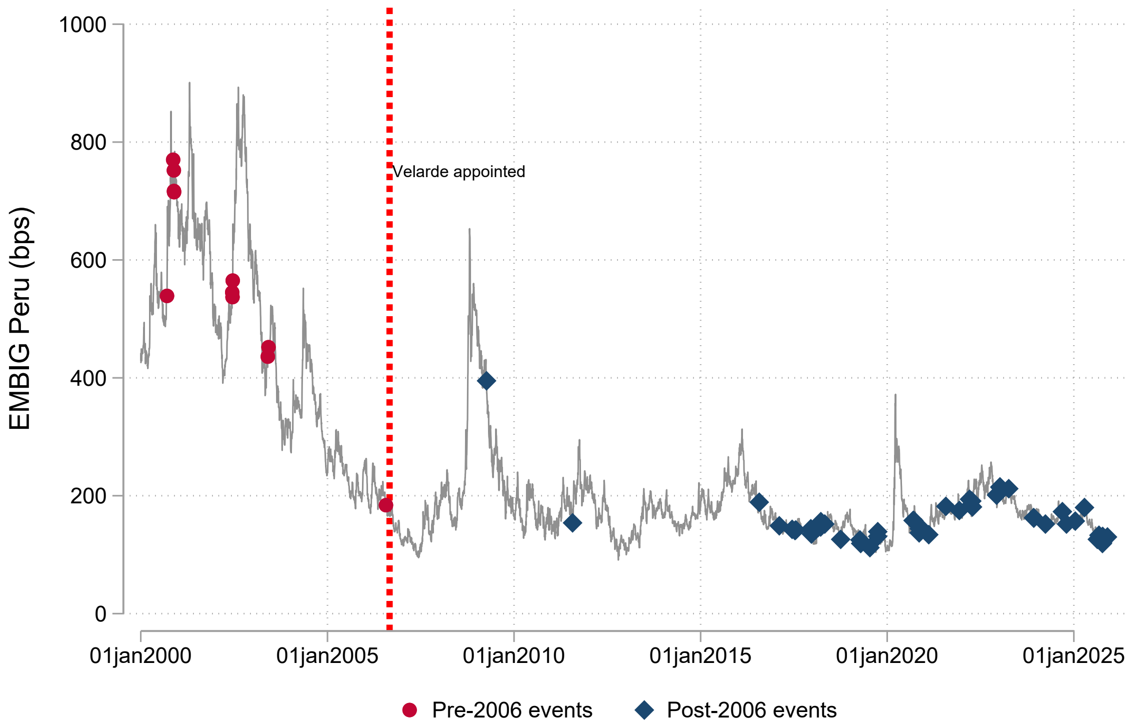

Notes: Daily EMBIG Peru spreads, 2000–2025. Red circles indicate political events before September 2006, green diamonds indicate events after. The vertical line marks Velarde's appointment. Political events include presidential removals, impeachments, coups, mass protests, corruption scandals, and deaths of political figures.

2.2 Political Events Data

We construct a dataset of political events in Peru from 2000 to 2025 through systematic coding of major political shocks, drawing on contemporaneous news reports, official government sources, and secondary historical accounts. We identify 90 distinct political events occurring on 82 unique trading days.

Events are classified into nine categories: (i) elections and electoral disputes, (ii) presidential inaugurations, (iii) cabinet crises and ministerial changes, (iv) impeachment proceedings and votes, (v) presidential resignations and removals, (vi) mass protests and civil unrest, (vii) corruption scandals and judicial proceedings, (viii) constitutional crises and attempted coups, and (ix) deaths of major political figures.

Each event is coded along two dimensions. First, a *severity* score: "high severity" (value = 3) covers presidential removals, deaths, constitutional crises, self-coups, and massacres, "medium severity" (value = 2) covers elections, cabinet crises, protests, and judicial proceedings against former officials. Second, whether the event was *anticipated* (value = 1) or a *surprise* (value = 0). Scheduled events such as elections and inaugurations are coded as anticipated, resignations, scandals, and coups as surprises. For events on weekends or holidays we assign the next trading day. When multiple events occur on the same day we retain the maximum severity and code the day as a surprise if any component event was unanticipated.

We note that the coding was carried out by the authors with knowledge of subsequent market behavior. We address the resulting endogeneity concern in Section 7.3, where the central results are replicated with an externally coded measure of political risk.

2.3 Political Instability Index

To analyze the instability–market relationship at lower frequency, we construct a political instability index. For each month t,

$$\text{Instability}_t = \sum_{e \in \mathcal{E}_t} \text{Severity}_e, \tag{1}$$

where $\mathcal{E}_t$ is the set of events in month t and $\text{Severity}_e \in \{2, 3\}$. We also compute a 12-month rolling average. The cumulative and monthly representations of the index appear in Online Appendix Figures OA.1–OA.2.

Figure 2: Annual Political Instability: Event Count and Severity-Weighted Index

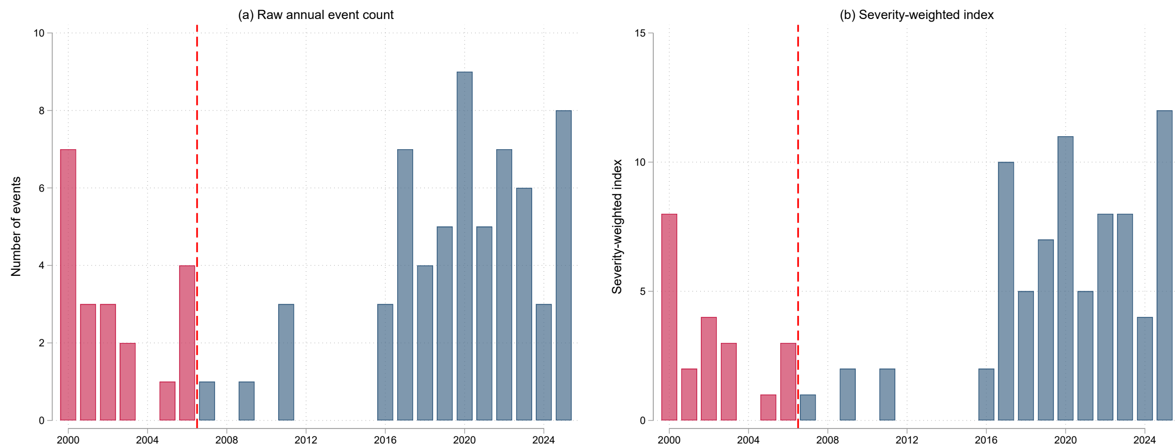

Notes: Panel (a): raw count of political events per year. Panel (b): the severity-weighted instability index (medium=1, high=2). Bars are shaded by era—pre-Velarde (left of the dashed line, September 2006) and Velarde era (right). The two panels trace a near-identical annual profile, showing that the post-2006 rise in instability does not depend on the severity weighting. Cumulative and monthly versions of the index appear in Online Appendix Figures OA.1–OA.2.

Table 2 summarizes events by era. The Velarde era (2006–2025) experienced more than three times as many political events as the pre-Velarde period, yet—as we show—financial markets responded very differently.

Table 2: Summary Statistics: Political Events by Era

	Pre-Velarde (2000–2006)	Velarde Era (2006–2025)
Panel A: Event Counts		
Total events	20	70
Unique event days	19	63
High severity events	5	23
Surprise events	10	30
Panel B: Average Annual Instability		
Events per year	3.0	3.6
Severity-weighted index	3.0	4.1
Panel C: Correlations with EMBIG		
Instability–EMBIG correlation	0.549	−0.230

Notes: Pre-Velarde period ends August 2006. Correlations use the 12-month rolling instability index and monthly average EMBIG.

Panel C previews the main finding. The instability–spread correlation is positive (+0.55) before 2006 and flips to −0.23 in the Velarde era—a swing of 0.78.

2.4 The Velarde Era

A central feature of the analysis is the distinction between the pre-Velarde period (2000–2006) and the Velarde era (2006–present). Julio Velarde was appointed President of the BCRP in October 2006 and has remained in office continuously for nearly two decades. Under his tenure the BCRP has maintained strict inflation targeting, accumulated substantial reserves, and preserved operational independence; inflation expectations have remained anchored at 2–3% even during periods of extreme political turmoil.

We date the Velarde era from September 2006, one month before the formal appointment, to capture anticipation effects, results are robust to alternative cutoffs. We stress that "Velarde" is shorthand for a regime change whose attribution is not mechanical: October 2006 also coincides with the maturing of Peru's fiscal-responsibility framework and the global commodity boom. The cross-country comparisons of Section 7 are designed to separate the central-bank channel from these confounders.

Beyond inflation targeting and reserve accumulation, the BCRP operates an active intervention toolkit designed to smooth disorderly market conditions without compromising the floating exchange-rate regime. Foreign-exchange swaps (*swaps cambiarios*)—derivative contracts in which the BCRP commits to settle the difference between the future spot rate and a reference rate—have been used since September 2014 to provide hedging instruments without depleting reserves directly. The BCRP complements this with spot interventions and dollar repos. We document this toolkit here as context for the credibility regime; the empirical strategy of Sections 5–7 identifies the spread response separately from these intervention flows, and Section 7.2 verifies that the decoupling survives controlling for swap activity directly.

3 A Model of Central Bank Credibility and Sovereign Risk

This section develops a model in which a credible central bank absorbs political shocks before they reach the sovereign spread. We combine three existing frameworks in a deliberately modular way, and we are explicit throughout about what is borrowed and what is new. From Dziuda and Pflueger (2025, henceforth D-P) we take the monetary-policy delegation block: the social and central-bank loss functions, the expectational Phillips curve with government quality as the cost-push shock, and the equilibrium pass-through formula. The sovereign-spread pricing block—the log-linear relation between a fiscal fundamental and the spread—follows the standard sovereign-pricing literature (Arellano, 2008), the multiplicity and fragility apparatus of Appendix B draws on Pflueger and Yared (2025, henceforth P-Y). From Cukierman and Meltzer (1986) we take the learning structure that generates the tenure effect. Our contribution is to connect these pieces, to formalize the mapping from observed political events to the model's shocks, and to derive closed-form quantitative predictions.

A note on the cost-push interpretation. The block we borrow from D-P reinterprets the Phillips-curve cost-push shock as *government competence*. We adopt this interpretation because it is what allows monetary policy to shield nominal variables from political dysfunction, but we flag that it is the most contestable component we inherit: the mapping from "competence" to a supply shock is a modeling choice, not a measured object. Section 3.2 therefore supplies an explicit microfoundation linking the political events in our data to this shock, rather than asserting the mapping.

Relation to the sovereign-default literature. Our framework is a tractable analytical alternative to the Eaton–Gersovitz–Arellano paradigm of endogenous default (Arellano, 2008; Eaton and Gersovitz, 1981). We deliberately work with a structure that yields closed-form predictions mappable to empirical estimates: the spread response to political shocks is captured through the linearized pricing equation rather than through an endogenous default decision. The reduced-form approach buys analytical tractability and a transparent parameter-to-moment mapping, it forgoes the structural identification of default thresholds and debt dynamics. The fully structural counterpart—embedding our central-banker learning mechanism in a quantitative sovereign-default model with political turnover in the manner of Hatchondo et al. (2009)—is a natural extension we leave for future work, one in which the analytical mapping between parameters and moments must be replaced by numerical comparative statics.

3.1 Environment

The model has three groups of agents: a government (incumbent), a central bank (CB), and a continuum of international investors. There are two periods $t \in \{1, 2\}$ preceded by an appointment stage at $t = 0$, following D-P.

At the start of period t the incumbent has quality g_t, drawn from a distribution F with mean zero, variance σ_g^2, and support bounded above, the bound is made precise in Assumption 1. The CB observes g_t, international investors do not. This information structure follows D-P (Section 2.2), who assume the incumbent's competence is observed by the central bank but not by voters, we map voters to international investors.

Assumption 1 (Upper bound, D-P Assumption 1)**.** *$g_t < -u^*$ for all realizations, so that $g_t + u^* < 0$ and the time-inconsistency motive for surprise inflation (Kydland and Prescott, 1977) is always present.*

At the end of period 1 the incumbent is replaced with probability $\lambda \in (0, 1)$. If retained, $g_2 = g_1$. If replaced, the period-2 government has quality g_C (specified below). We model political events as surprise increases in the replacement probability:

$$\lambda_t = \bar{\lambda} + \varepsilon_t^{\text{pol}}, \qquad \varepsilon_t^{\text{pol}} \sim N(0, \sigma_\lambda^2), \tag{2}$$

where $\bar{\lambda}$ is the baseline turnover rate and $\varepsilon_t^{\text{pol}}$ captures impeachments, self-coups, or mass protests. Equation (2) is our extension of the D-P framework, which takes λ as fixed; it is what connects the model to the political-event data.

3.2 From Political Events to Government-Quality Shocks

The spread equation derived below depends on government quality g_t, whereas political events are modeled in (2) as shocks to the replacement probability λ_t. A natural objection—and one a careful reader should raise—is that an impeachment or a self-coup is a *realization of turbulence*, not obviously a shock to *future competence*. We close this gap with an explicit selection argument rather than leaving the mapping implicit.

Assumption 2 (Adverse selection in crisis successions)**.** *A government removed through a crisis succession (impeachment, resignation under pressure, coup) is replaced by an administration whose quality is drawn from a distribution with mean $-\mu_R < 0$, strictly below the unconditional mean of F. Orderly successions (scheduled elections, constitutional handovers) draw from F with mean zero.*

Assumption 2 formalizes a standard idea: crisis turnover shortens horizons, weakens the bureaucratic transmission of policy, and selects caretaker administrations with limited mandates. It is the model's testable content, not a free parameter, and μ_R should be identifiable from the behavior of perceived-quality proxies around crisis successions. Table 3 runs that test on the seven Peruvian administrations covered by BCRP's monthly business-expectations index. The twelve-month-ahead expectations index averages 63.0 across the months under orderly-succession administrations (García, Humala, PPK, Castillo) and 58.5 across the months under crisis-succession ones (Vizcarra, Sagasti, Boluarte)—a 4.5-point gap that is significant at the one-percent level on the underlying 203-month panel ($t = +3.26$).

The Castillo administration is the only observation that strains the binary typology: it entered through election but its operating pattern resembled a crisis administration's, with business expectations of 45.5—the lowest in the sample—and an unprecedented 3.8 premier rotations per year. Including Castillo in the orderly group is the conservative choice for the test; excluding him widens the gap to +8.4 points on business expectations and reverses the cabinet-rotation comparison toward the model's predicted direction. We read this as direct, if imperfect, empirical support for the adverse-selection assumption: when governments come in through crisis, the market discount on perceived quality is on average larger.

Table 3: Perceived government quality by succession type, 2008–2025

Administration (entry)	Type	Business expectations	PCMs/yr
García (Jul 2006)	orderly	69.1	1.3
Humala (Jul 2011)	orderly	61.7	1.4
PPK (Jul 2016)	orderly	69.9	1.3
Castillo (Jul 2021)	orderly	45.5	3.8
Vizcarra (Mar 2018)	crisis	61.0	1.9
Sagasti (Nov 2020)	crisis	58.1	1.7
Boluarte (Dec 2022)	crisis	56.4	1.4
Orderly mean (4 admins, 131 months)		63.0	1.9
Crisis mean (3 admins, 72 months)		58.5	1.7
month-level t-statistic (orderly − crisis)		$+3.26^{***}$	—

Notes: Business expectations: BCRP twelve-month-ahead business outlook index (*Encuesta de Expectativas Macroeconómicas*, series `PD37981AM`), monthly from June 2008, scaled around 50 (neutral); higher values indicate more optimistic outlook. PCMs/yr counts changes in the Presidente del Consejo de Ministros normalized by administration length. Within-administration means are computed across all months the administration held office that overlap the data window. The Welch t-statistic compares all month-level observations under orderly admins to those under crisis admins. The Castillo administration entered through election but registered the lowest business expectations of any modern Peruvian administration and the highest premier turnover; its inclusion in the orderly group is conservative for the test. $^{***}p < 0.01$.

Lemma 1 (Events as quality shocks)**.** *Under Assumption 2, expected forward government quality is*

$$\mathbb{E}[g_{t+1} \mid \lambda_t, g_t] = (1 - \lambda_t)\, g_t - \lambda_t\, \mu_R, \tag{3}$$

so that a political event that raises the replacement probability lowers expected government *quality:*

$$\frac{\partial\,\mathbb{E}[g_{t+1}]}{\partial \lambda_t} = -(g_t + \mu_R) < 0.$$

Proof. Immediate from Assumption 2 and the retention rule $g_2 = g_1$ with probability $1 - \lambda_t$. The derivative is negative whenever $g_t > -\mu_R$—that is, for any incumbent who is not severely incompetent—since $\mu_R > 0$ by Assumption 2. □

Lemma 1 licenses the rest of the analysis: a surprise political event $\varepsilon_t^{\text{pol}} > 0$ is, in expectation, a negative shock to g, and we may study the spread response to $-\Delta g_t > 0$ directly. Where it matters we treat the empirical event-study coefficient as an estimate of the composite $\beta(\tilde{\theta}) \cdot (g_t + \mu_R)$ rather than of $\beta(\tilde{\theta})$ alone.

3.3 The Monetary Policy Block

This block is a direct application of Dziuda and Pflueger (2025, henceforth D-P), Sections 2.1–2.3 and 3.1.

Preferences. Following D-P equation (1), social welfare each period is the negative of

$$L_t = \frac{(u_t - u^*)^2}{2} + \theta\,\frac{\pi_t^2}{2}, \tag{4}$$

where u_t is unemployment, π_t is inflation, $u^* < 0$ is the socially optimal (unattainable) unemployment rate, and $\theta > 0$ is the social weight on inflation. At $t = 0$ the incumbent appoints a CB with inflation-aversion $\tilde{\theta} \geq 0$ and loss

$$\tilde{L}_t = \frac{(u_t - u^*)^2}{2} + \tilde{\theta}\,\frac{\pi_t^2}{2}. \tag{5}$$

The CB is *hawkish* if $\tilde{\theta} > \theta$ and *dovish* if $\tilde{\theta} < \theta$, the appointment of a central banker more inflation-averse than society is the delegation device of Rogoff (1985).

Phillips curve. Following D-P equation (3), the supply side is

$$u_t = -(\pi_t - \pi_t^e) - g_t, \tag{6}$$

where $\pi_t^e \equiv \mathbb{E}_t[\pi_t]$ are rational expectations formed before the CB acts. As in D-P, g_t is government competence: incompetent governments distort product and labor markets, raising both unemployment and inflation.

Equilibrium. The CB minimizes (5) subject to (6), yielding D-P's optimality condition $u_t - u^* = \tilde{\theta}\,\pi_t$. Imposing rational expectations gives the Barro–Gordon (Barro and Gordon, 1983) inflation bias $\pi_t^e = -u^*/\tilde{\theta}$ and equilibrium outcomes

$$\pi_t^* = -\frac{u^*}{\tilde{\theta}} - \frac{g_t}{1+\tilde{\theta}}, \tag{7}$$

$$u_t^* = -\frac{\tilde{\theta}\,g_t}{1+\tilde{\theta}}. \tag{8}$$

Lemma 2 (Pass-through, D-P Lemma 1 adapted)**.** *The pass-through of a government-quality shock to equilibrium inflation is*

$$\frac{\partial \pi_t^*}{\partial g_t} = -\frac{1}{1+\tilde{\theta}}, \tag{9}$$

which is negative, decreasing in absolute value in $\tilde{\theta}$, and converges to zero as $\tilde{\theta} \to \infty$.

Proof. Immediate from (7). □

3.4 The Sovereign Spread Pricing Block

International investors price sovereign bonds on a fiscal fundamental

$$\mathcal{F}_t \equiv \bar{F} + \alpha\,g_t - \delta\,\pi_t^*, \tag{10}$$

where $\bar{F} > 0$ is structural fiscal capacity, $\alpha > 0$ is the elasticity of the primary balance to government quality, and $\delta > 0$ is the cost of inflation to debt sustainability (it aggregates erosion of the real value of debt, higher nominal refinancing costs, and depreciation risk). The mapping from fiscal fundamentals to bond pricing follows the standard sovereign-debt literature (Arellano, 2008). Substituting equilibrium inflation (7):

$$\mathcal{F}_t = \underbrace{\bar{F} + \frac{\delta\, u^*}{\tilde{\theta}}}_{\equiv \bar{\mathcal{F}}(\tilde{\theta})} + \underbrace{\left(\alpha + \frac{\delta}{1+\tilde{\theta}}\right)}_{\equiv \Gamma(\tilde{\theta})} g_t. \tag{11}$$

Definition 1 (Amplification function)**.** $\Gamma(\tilde{\theta}) \equiv \alpha + \dfrac{\delta}{1+\tilde{\theta}}$.

Γ captures two additive channels through which a political shock reaches the fiscal fundamental: a *direct fiscal channel* α (fiscal mismanagement), and an *inflation channel* $\delta/(1+\tilde{\theta})$ (worse competence raises equilibrium inflation through the pass-through of Lemma 2, eroding debt sustainability). The inflation channel vanishes as $\tilde{\theta} \to \infty$.

In the standard log-linear approximation to sovereign pricing (Arellano, 2008), the spread is decreasing in the fundamental:

$$s_t = \bar{s} - \gamma\, \mathcal{F}_t + \nu_t, \tag{12}$$

with $\gamma > 0$, reference level $\bar{s}$, and $\nu_t \sim N(0, \sigma_\nu^2)$ capturing common regional factors (proxied empirically by EMBIG Latin America). Substituting (11):

$$s_t = \underbrace{\bar{s} - \gamma\, \bar{\mathcal{F}}(\tilde{\theta})}_{\text{level}} - \underbrace{\gamma\, \Gamma(\tilde{\theta})}_{\equiv \beta(\tilde{\theta})}\, g_t + \nu_t. \tag{13}$$

The coefficient $\beta(\tilde{\theta}) \equiv \gamma\, \Gamma(\tilde{\theta})$ is the *market sensitivity* to political shocks.

3.5 Insulation and Variance Reduction

Proposition 1 (Insulation). *The sensitivity of the sovereign spread to a political shock of size* $-\Delta g_t > 0$ *is*

$$\frac{\partial s_t}{\partial(-g_t)} = \gamma\left(\alpha + \frac{\delta}{1+\tilde{\theta}}\right), \tag{14}$$

strictly decreasing in $\tilde{\theta}$, *with* $\lim_{\tilde{\theta}\to\infty} \partial s_t/\partial(-g_t) = \gamma\alpha$. *If the direct fiscal channel is small* ($\alpha \approx 0$)*, a perfectly credible CB provides complete insulation.*

Proof. Differentiate (13), monotonicity from $\Gamma'(\tilde{\theta}) = -\delta/(1+\tilde{\theta})^2 < 0$ (Appendix A). □

Figure 3 plots the amplification function. The insulation theorem *is* the downward slope of this curve: a more credible central bank—higher $\tilde{\theta}$—moves the economy from $\Gamma = \alpha + \delta$ toward the floor $\Gamma = \alpha$, so the spread's sensitivity to political shocks falls and, in the limit, the inflation channel closes entirely.

Figure 3: The Amplification Function $\Gamma(\tilde{\theta})$

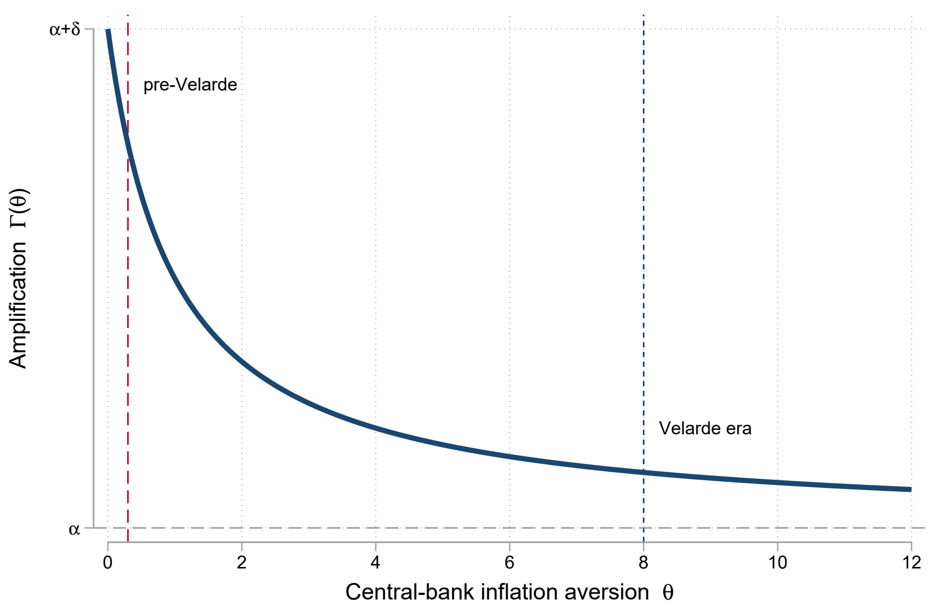

Notes: The amplification function $\Gamma(\tilde{\theta}) = \alpha + \delta/(1+\tilde{\theta})$ governs how strongly a government-quality shock reaches the sovereign spread. It is strictly decreasing in central-bank inflation aversion $\tilde{\theta}$, falling from $\alpha + \delta$ at $\tilde{\theta} = 0$ to the asymptote α as $\tilde{\theta} \to \infty$. Raising $\tilde{\theta}$ moves the economy from the pre-Velarde regime down the curve toward the Velarde-era regime. Illustrative values with $\delta/\alpha = 0.9$.

Corollary 1 (The Peruvian Puzzle)**.** *Let* $\tilde{\theta}_{\text{pre}}$ *and* $\tilde{\theta}_{\text{vel}} \gg \tilde{\theta}_{\text{pre}}$ *denote effective inflation-aversion under the two regimes. Then* $\text{Corr}(s_t, -g_t)|_{\tilde{\theta}_{\text{pre}}} > 0 > \text{Corr}(s_t, -g_t)|_{\tilde{\theta}_{\text{vel}}}$ *whenever the secular decline in the level* $\bar{s} - \gamma\bar{\mathcal{F}}(\tilde{\theta})$ *generates a negative sample correlation between trend and the instability index.*

Proof. In the Velarde era the sensitivity collapses to the small residual fiscal channel, $\beta(\tilde{\theta}_{\text{vel}}) \approx \gamma\alpha$, so the spread responds only weakly to political shocks; the negative sample correlation is then dominated by the secular decline in the level $\bar{s} - \gamma\bar{\mathcal{F}}(\tilde{\theta})$. □

Proposition 2 (Variance reduction)**.** *The variance of daily spread changes is* $\text{Var}(\Delta s_t) = \beta(\tilde{\theta})^2\sigma_g^2 + \sigma_\nu^2$*, and the pre-to-post variance ratio is*

$$R = \frac{\beta(\tilde{\theta}_{\text{pre}})^2\sigma_g^2 + \sigma_\nu^2}{\beta(\tilde{\theta}_{\text{vel}})^2\sigma_g^2 + \sigma_\nu^2}. \tag{15}$$

Under $\tilde{\theta}_{\text{pre}} \to 0$*,* $\tilde{\theta}_{\text{vel}} \to \infty$*, and* $\sigma_\nu^2 \ll \beta^2\sigma_g^2$*,* $R \approx (\alpha+\delta)^2/\alpha^2$*.*

Proof. From (13), $\Delta s_t = -\beta(\tilde{\theta})\Delta g_t + \Delta\nu_t$, independence of g_t and ν_t gives the variance, and the limiting calibration gives the ratio. □

3.6 Learning and the Tenure Effect

Propositions 1–2 treat $\tilde{\theta}$ as known. In practice the inflation-aversion of a newly appointed governor is uncertain, and investors learn it over time. We model learning following Cukierman and Meltzer (1986).

We follow Cukierman and Meltzer (1986) directly and let investors learn from *observed inflation*, which the BCRP publishes monthly and which suffers no such regularity problem. From (7), period-τ inflation is

$$\pi_\tau = -\frac{u^*}{\tilde{\theta}} - \frac{g_\tau}{1+\tilde{\theta}}, \tag{16}$$

so that average observed inflation has mean $-u^*/\tilde{\theta}$, a strictly decreasing function of $\tilde{\theta}$, with idiosyncratic noise $g_\tau/(1+\tilde{\theta})$ of finite variance $\sigma_g^2/(1+\tilde{\theta})^2$. Inflation is thus a well-behaved signal of the governor's type.

Updating. At appointment ($T = 0$) investors hold the prior $\tilde{\theta} \sim N(\mu_0, v_0)$. Gaussian updating on T inflation observations gives posterior variance

$$v_T = \frac{v_0\, \sigma_\xi^2}{\sigma_\xi^2 + T\, v_0}, \tag{17}$$

where σ_ξ^2 is the noise variance of the inflation signal. Uncertainty falls monotonically in tenure T.

Proposition 3 (Tenure effect)**.** *Let T be the governor's tenure at the time of a political event. The expected absolute market response is*

$$\beta_{\text{event}}(T) \equiv \mathbb{E}[\gamma\, \Gamma(\tilde{\theta}) \mid \mathcal{H}_T] = \gamma\, \Gamma(\hat{\theta}_T) + \frac{\gamma\, \delta\, v_T}{(1 + \hat{\theta}_T)^3}, \tag{18}$$

where $\hat{\theta}_T = \mathbb{E}[\tilde{\theta} \mid \mathcal{H}_T]$. The function $\beta_{\text{event}}(T)$ is strictly decreasing in T.

Proof. By the law of iterated expectations and a second-order Taylor expansion of Γ around $\hat{\theta}_T$, using $\Gamma''(\theta) = 2\delta/(1 + \theta)^3 > 0$ (Appendix A). Both terms decrease in T: $\Gamma(\hat{\theta}_T)$ falls as investors learn $\tilde{\theta}$ is high, and v_T falls by (17). □

Figure 4 plots this prediction: the event response decays with tenure, steeply at first and flattening as beliefs converge. It is the model's most distinctive signature, and Section 6.2 confronts it directly with the event-level data.

Figure 4: The Tenure Effect: Theoretical Event Response $\beta_{\text{event}}(T)$

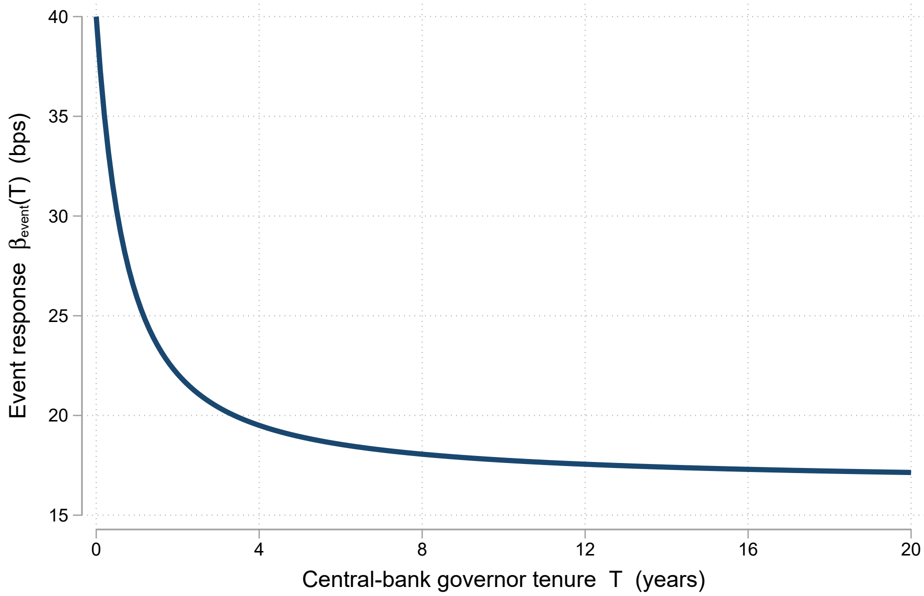

Notes: The model's predicted market response to a political event, $\beta_{\text{event}}(T)$ of equation (18), against the governor's tenure T. As investors learn the governor's type, the posterior mean $\hat{\theta}_T$ rises and the posterior variance v_T falls; both push the response down, yielding a decreasing convex curve. The curve flattens at the irreducible direct fiscal channel $\gamma\alpha$, not at zero: learning closes the inflation channel $\delta/(1+\tilde{\theta})$ but not the structural fiscal pass-through α. The empirical counterpart is the event-level scatter of Figure 8. Illustrative parameter values.

Corollary 2 (Linear approximation)**.** *For moderate T and small v_0, $\beta_{\text{event}}(T) \approx \beta_0 - \kappa T$ with $\kappa \equiv \gamma\,\delta\,(\tilde{\theta} - \mu_0)/[(1+\mu_0)^2\,T_0] > 0$, where $T_0 \equiv \sigma_\xi^2/v_0$ is the prior effective sample size.*

Why learning is slow. Sophisticated investors might be expected to learn a governor's type faster than the decade the calibration implies. The model rationalizes slow learning through a *tight, pessimistic prior*: Peruvian investors entered the Velarde era with beliefs shaped by the hyperinflation of the late 1980s and the fragile disinflation of the 1990s, which made the prior variance v_0 small relative to the signal noise—formally, slow learning requires $\sigma_\xi^2 \gg v_0$, which we maintain as a calibration restriction. The estimated tenure decay plausibly also reflects the 2008 investment-grade upgrade and broader emerging-market spread compression, the placebo tests of Section 7 bear on this directly.

3.7 Mapping the Model to the Data

Table 4 collects the predictions, their provenance, and their empirical counterparts (reported in Sections 5 and 6).

Table 4: Model predictions, provenance, and empirical estimates

Prediction	Model expression	Source	Empirical
Spread sensitivity, pre-2006	$\beta(\tilde{\theta}_{\text{pre}}) = \gamma(\alpha+\delta)$	D-P eq. (3), Arellano (2008)	≈ 41 bps
Spread sensitivity, post-2006	$\beta(\tilde{\theta}_{\text{vel}}) \approx \gamma\alpha \approx 0$	D-P Prop. 2, Arellano (2008)	≈ 0 bps
Correlation flip	$\rho(\tilde{\theta}_{\text{pre}}) > 0 > \rho(\tilde{\theta}_{\text{vel}})$	D-P §3.1, this paper	$+0.55 \to -0.23$
Variance ratio	$R = (\alpha+\delta)^2/\alpha^2$	D-P Prop. 3 (analogue), this paper	≈ 3.6
Tenure decay	$\partial\beta_{\text{event}}/\partial T = -\kappa < 0$	Cukierman–Meltzer (1986), this paper	-1.73 bps/yr

Notes: D-P denotes Dziuda and Pflueger (2025). The parameter $\tilde{\theta}$ is the central bank's effective inflation-aversion, and γ, α, δ are the pricing, direct fiscal, and inflation-channel parameters of the amplification function $\Gamma(\tilde{\theta})$ (Definition 1). Empirical estimates are drawn from the event study of Section 5 and the model-prediction tests of Section 6, the variance ratio is the pre- to post-2006 ratio of the variance of daily EMBIG changes.

Informal calibration. Proposition 2 with an empirical variance ratio of $R \approx 3.6$ gives $(\alpha+\delta)/\alpha = \sqrt{R} \approx 1.9$, so $\delta/\alpha \approx 0.9$: the inflation channel accounts for a little under half of the pre-Velarde sensitivity. The pre-2006 event-study coefficient of ≈ 41 bps identifies $\gamma(\alpha+\delta)$, which combined with $\delta/\alpha \approx 0.9$ implies a residual fiscal channel $\gamma\alpha \approx 21$ bps. The post-2006 coefficient of 0.9 bps, statistically indistinguishable from zero, shows that this residual is empirically too small to detect event by event—the inflation channel has, for practical purposes, fully shut down. The directly estimated decay rate $\kappa \approx 1.8$ bps per year then corresponds, through Corollary 2, to a prior effective sample size T_0 on the order of a decade, the precise figure depends on the assumed prior tightness and is not separately identified by our data. The multiplicity and fragility apparatus—which microfounds the "Velarde Premium" but generates no prediction tested in this paper—is developed separately in Appendix B, and its dynamic counterpart (credibility hysteresis and succession risk) in ongoing work.

Frequency. Although the model is written in two periods for analytical tractability, its predictions concern the mapping from government-quality shocks to spreads at any frequency that admits the pricing equation; we therefore test them at the highest empirical frequency the data allow—daily for the event study, monthly for the instability index.

4 Empirical Strategy

We employ two complementary strategies: (i) event-study analysis comparing market responses to political shocks before and after 2006, and (ii) panel regressions relating the instability index to sovereign spreads.

4.1 Event Study Design

The primary specification estimates the response of spreads to political events within a 15-day window (± 7 days). For each event e on date t_e,

$$Y_t = \alpha + \sum_{k=-7}^{7} \beta_k \, \mathbf{1}[t - t_e = k] + \gamma_y + \varepsilon_t, \tag{19}$$

where Y_t is the EMBIG spread and γ_y are year fixed effects. To test for differential responses by era,

$$Y_t = \alpha + \beta^{\text{pre}} \, \text{Post}_t + \beta^{\text{vel}} \, \text{Post}_t \times \text{Velarde}_e + \gamma_y + \varepsilon_t, \tag{20}$$

where Post_t indicates the post-event window (days 1–7). The model predicts $\beta^{\text{pre}} > 0$ and $\beta^{\text{pre}} + \beta^{\text{vel}} \approx 0$. Standard errors are clustered at the event level.

4.2 Instability Index Regressions

At monthly frequency,

$$\overline{Y}_t = \alpha + \beta \, \text{Instability}_t + \gamma_y + \varepsilon_t, \tag{21}$$

and, allowing the coefficient to differ by era,

$$\overline{Y}_t = \alpha + \beta^{\text{pre}} \, \text{Instability}_t^{\text{pre}} + \beta^{\text{vel}} \, \text{Instability}_t^{\text{vel}} + \gamma_y + \varepsilon_t. \tag{22}$$

Credibility insulation predicts $\beta^{\text{pre}} > 0$ and $\beta^{\text{vel}} \leq 0$.

Reverse causality. A concern running in the opposite direction—that sovereign-spread movements themselves shape political outcomes—is unlikely to bias our estimates at the frequencies we exploit. The political shocks we study (impeachments, self-coups, mass protests) are triggered by domestic dynamics that are unlikely to be timed by short-window spread movements, and the ±7-day event windows are too short for any spread-to-politics feedback to enter mechanically.

5 Main Results

5.1 Event Study Results

Pooling all events without distinguishing by era yields a small, statistically insignificant response (Online Appendix Figure OA.3), which masks the heterogeneity central to this paper. Figure 5 separates the response by era. Before 2006, political events trigger a significant, persistent increase in spreads, in the Velarde era the response is flat and statistically indistinguishable from zero.

Figure 5: Event Study: Market Response to Political Events by Era

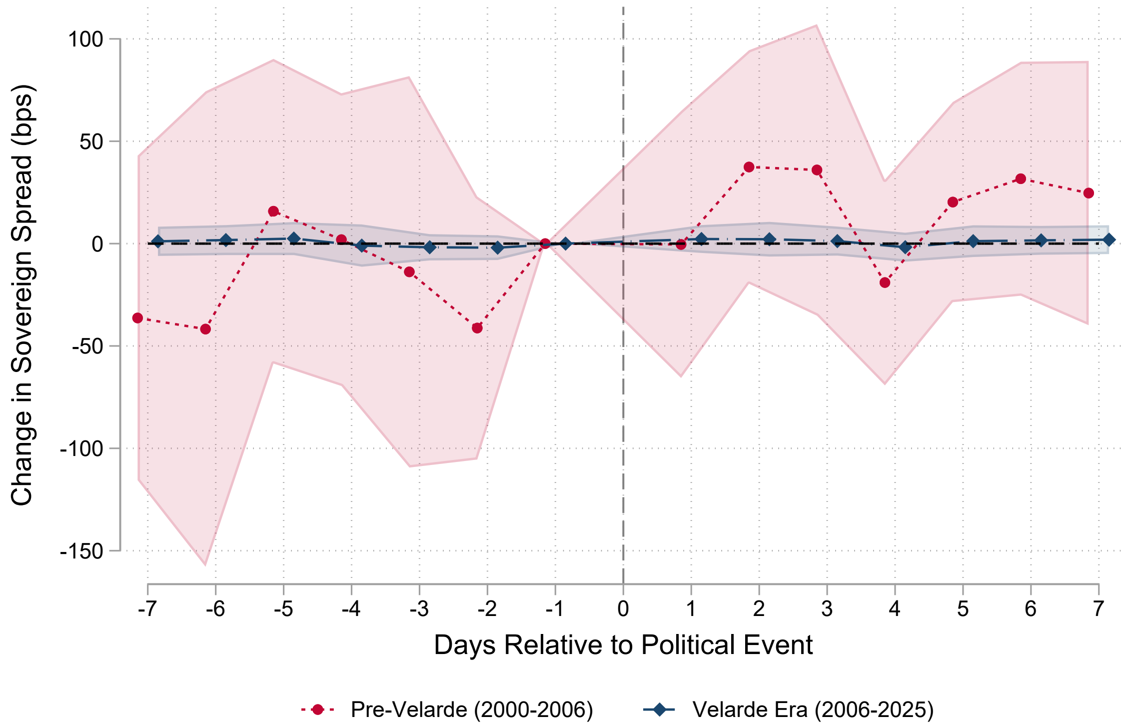

Notes: Coefficient estimates from equation (19) separately by era, 95% confidence intervals. Standard errors clustered by event.

Table 5 reports the regression estimates, and Online Appendix Figure OA.5 displays them as a bar chart. Pooling all events (column 1) yields a modest 6.6 bps response, which masks the heterogeneity central to this paper. Column (2) separates the eras: before 2006 a political event raised the spread by 40.9 bps ($p = 0.008$), the Velarde-era response is 0.9 bps, statistically indistinguishable from zero, and the era interaction of -40.1 bps is itself significant ($p = 0.009$).

Table 5: Market Response to Political Events: Event Study Regressions

	(1) EMBIG pooled	(2) EMBIG by era	(3) Exch. rate by era
Post-Event	6.56**	40.92***	0.010**
	(2.62)	(15.26)	(0.005)
Post-Event × Velarde Era		−40.05***	−0.008
		(15.33)	(0.010)
Year FE	Yes	Yes	Yes
Observations	478	478	443

Notes: Robust standard errors in parentheses, the event window is ±7 days and Post-Event indicates days 0–7. Column (1) pools all events, columns (2)–(3) interact the post-event window with the Velarde-era indicator. In column (2) the pre-2006 response is 40.9 bps and the Velarde-era response is $40.9 - 40.1 = 0.9$ bps, statistically indistinguishable from zero. Column (3) uses the sol/dollar exchange rate. $^{**}p < 0.05$, $^{***}p < 0.01$.

Column (3) shows no significant exchange-rate response, consistent with the BCRP's managed-float regime and substantial reserves, the exchange-rate channel is examined further in Online Appendix Figure OA.4.

Two descriptive benchmarks frame the 40.9 bps estimate. The average *absolute* change in the spread over the event window was 37.6 bps before 2006 and 6.8 bps afterward; the average *signed* change was +13.7 and +0.9 bps respectively. The regression coefficient exceeds the absolute average because year fixed effects strip out common drift, isolating the event itself; the signed average is smaller because political events are heterogeneous in direction. All three estimators agree on the qualitative pattern, and we report the regression coefficient as the canonical figure.

These magnitudes are the empirical counterpart of Proposition 1. The model identifies the pre-2006 event coefficient with $\beta(\tilde{\theta}_{\text{pre}}) = \gamma(\alpha + \delta)$—the full sensitivity, with both the direct fiscal channel α and the inflation channel $\delta/(1+\tilde{\theta})$ operative—and the Velarde-era coefficient with $\beta(\tilde{\theta}_{\text{vel}}) \to \gamma\alpha$, the residual sensitivity once a credible central bank has closed the inflation channel. The estimated collapse from a significant 41 bps to a precisely estimated zero is a direct reading of the insulation theorem: raising effective inflation-aversion $\tilde{\theta}$ severs the transmission from government-quality shocks to the spread. The model does not claim the residual channel $\gamma\alpha$ is literally zero—the variance ratio of Section 6 implies it is positive—only that it is too small to detect event by event, which is what the tightly estimated Velarde-era coefficient shows.

5.2 Instability Index Results

Table 6 reports monthly regressions of the spread on the instability index. Pooling the eras (column 1) gives 55.8 bps per index unit. Column (2) reveals the reversal: +79.1 bps before 2006, −63.5 bps afterward. Column (3) formalizes it as an interaction—a highly significant −136.9 bps.

Table 6: Political Instability and Sovereign Spreads: Monthly Regressions

	(1) Pooled	(2) By era	(3) Interaction
Instability index	55.8^{***}		78.5^{***}
	(16.6)		(22.5)
× Pre-Velarde		79.1^{***}	
		(22.5)	
× Velarde Era		-63.5^{***}	
		(16.7)	
Velarde Era			-25.7^{**}
			(13.0)
Instability × Velarde Era			-136.9^{***}
			(28.2)
Year FE	Yes	Yes	Yes
Observations	311	311	311

Notes: Robust standard errors in parentheses. Dependent variable is monthly average EMBIG Peru, the instability index is the 12-month rolling severity-weighted sum. In column (2) the instability slope flips from +79 bps before 2006 to −64 bps in the Velarde era. $^{**}p < 0.05$, $^{***}p < 0.01$.

This sign reversal is Corollary 1 in the data. The model does *not* predict that instability and spreads become causally *negatively* related: in the Velarde era $\beta(\tilde{\theta}_{\text{vel}}) \approx 0$, so the spread is essentially orthogonal to government-quality shocks. The observed negative correlation arises mechanically, as the corollary states. With the causal link severed, the secular decline in the spread level $\bar{s} - \gamma\bar{\mathcal{F}}(\tilde{\theta})$ runs against a rising instability index, manufacturing a negative sample correlation out of two trends that no longer speak to each other. The flip from +0.55 to −0.23 is thus not evidence that political instability became *stabilizing*; it is evidence that it became *irrelevant* to the price of Peruvian risk. Figure 6 visualizes the break: a clear positive slope before 2006, flat or slightly negative afterward.

Figure 6: Political Instability vs. Sovereign Spreads by Era

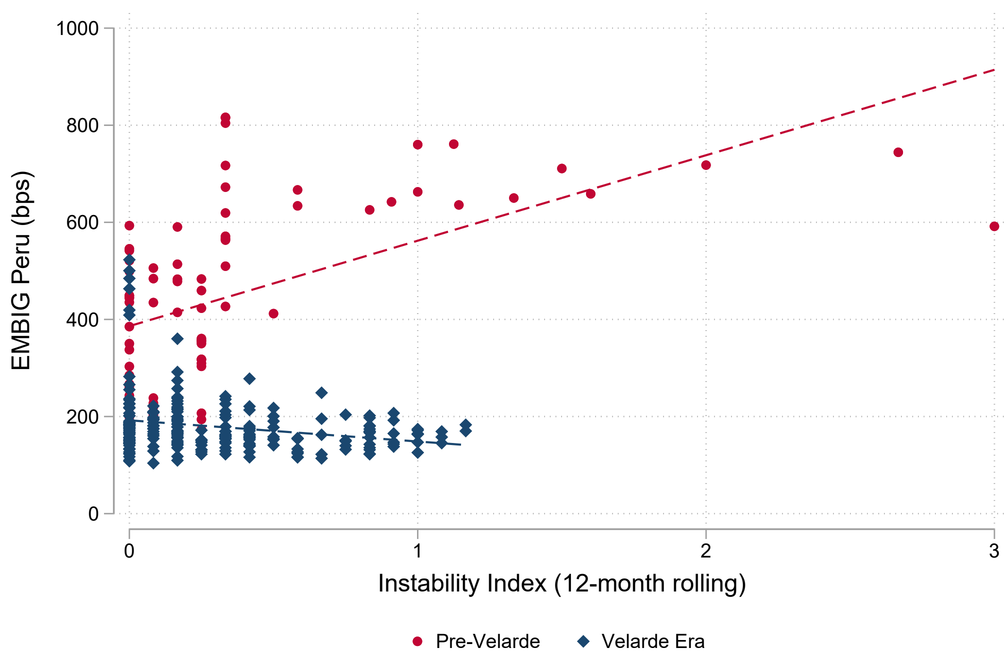

Notes: Monthly EMBIG Peru against the 12-month rolling instability index, by era. Dashed lines show within-era OLS fits.

6 Testing the Model's Predictions

The model generates two predictions testable with the event-level data: variance reduction (Proposition 2) and a tenure effect (Proposition 3).

6.1 Prediction 1: Variance Reduction

Proposition 2 is the model's sharpest quantitative claim. It decomposes the variance of daily spread changes into two orthogonal pieces, $\mathrm{Var}(\Delta s_t) = \beta(\tilde{\theta})^2\sigma_g^2 + \sigma_\nu^2$: a *political* component, $\beta(\tilde{\theta})^2\sigma_g^2$, driven by government-quality shocks and scaled by the squared market-sensitivity coefficient, and a *common* component, σ_ν^2, driven by regional and global factors. A credible central bank cannot touch σ_ν^2—it does not lower the world price of risk—but it drives $\beta(\tilde{\theta})$ toward $\gamma\alpha \approx 0$, collapsing the political component. The model therefore predicts not merely that volatility falls, but *which* part of it falls: the idiosyncratic, politically driven part, leaving the common component intact. Table 7 tests this decomposition on the daily Peruvian series.

Table 7: Testing Prediction 1: Variance of Daily EMBIG Changes by Era

	Pre-Velarde (2000–2006)	Velarde Era (2006–2025)
Panel A: Summary Statistics		
Observations	1,391	4,016
Mean (bps)	−0.14	−0.11
Standard Deviation (bps)	11.99	6.33
Variance	143.6	40.1
Panel B: Variance Ratio		
Variance Ratio (Pre/Post)	3.58	
Panel C: Statistical Tests		
Levene's Test (W_{50})	470.0	
p-value	< 0.0001	
F-Test for Equality	$F = 3.58$	
p-value	< 0.0001	

Notes: Daily changes in EMBIG Peru. The variance ratio of 3.58 is computed on the full Peruvian daily series, the cross-country comparison in Table 13 reports 3.6 for Peru on the common cross-country sample. The two numbers use slightly different daily-coverage samples and both exceed three.

The results confirm the prediction: daily spread volatility fell by more than half after 2006, and Levene's and the F-test decisively reject equal variances, Online Appendix Figure OA.7 shows the two distributions directly.

Figure 7: Volatility Reduction: Daily EMBIG Changes by Era

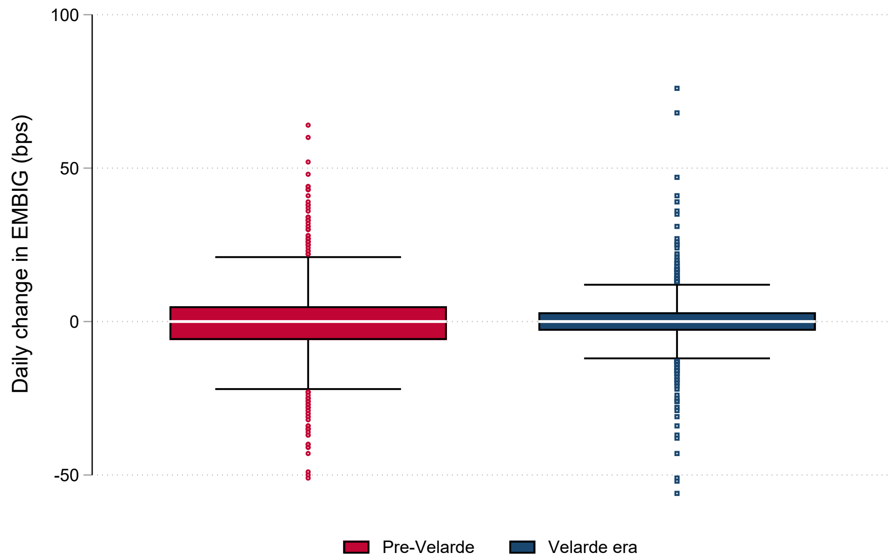

Notes: Box plots of daily changes in EMBIG Peru, pre-Velarde versus the Velarde era. The pre-Velarde distribution is substantially wider. A distribution-overlay version appears in Appendix C.

6.2 Prediction 2: The Tenure Effect

Proposition 3 makes a prediction the cross-era comparison cannot: that the insulation Peru enjoys was not switched on in 2006 but *accumulated* thereafter. In the model investors do not know a newly appointed governor's inflation-aversion $\tilde{\theta}$, they learn it from observed inflation, and the posterior variance v_T falls monotonically with tenure. Because the event-response coefficient $\beta_{\text{event}}(T)$ of equation (18) inherits that decline through both of its terms—the posterior mean $\hat{\theta}_T$ drifts upward as good inflation outcomes accumulate, and the posterior-variance term v_T shrinks—the model predicts that the market's reaction to a political shock should fade as Velarde's tenure lengthens. This is a *within*-Velarde-era gradient, distinct from the pre/post break of Section 5 and therefore a genuinely independent test of the mechanism. We regress the absolute event-window change in the spread on tenure. Table 8, Column (1), reports a tenure coefficient of -1.73 bps per year ($p < 0.001$): each additional year of accumulated credibility shaves roughly 1.7 bps off the market's response to the average political event.

Table 8: Testing Prediction 2: Market Response and Central Bank Tenure

	(1) Baseline	(2) De-clustered	(3) High-severity
Tenure (years)	-1.726^{***}	-1.478^{***}	-2.100^{***}
	(0.400)	(0.503)	(0.674)
R^2	0.31	0.26	0.49
Observations	90	53	28

Notes: Robust standard errors in parentheses. The dependent variable is the absolute change in EMBIG Peru between the pre- and post-event ($\pm$7-day) windows, tenure is years since September 2006. Column (2) merges events within 14 days of one another so overlapping windows are not double-counted, column (3) restricts to high-severity events. A wild cluster bootstrap of column (1) with 9,999 replications gives $p < 0.0001$ and a 95% confidence set of $[-2.6, -0.9]$, in all 90 leave-one-out replications the coefficient remains significant at the 1% level. $^{***}p < 0.01$.

The point estimate has a direct structural reading. Corollary 2 shows that, for moderate tenure, the event response is approximately linear, $\beta_{\text{event}}(T) \approx \beta_0 - \kappa T$, with the decay rate κ governed by how fast investors resolve uncertainty about the governor's type. The estimated -1.73 bps per year is thus an estimate of κ, combined with the pre-2006 sensitivity it implies a prior effective sample size T_0 on the order of a decade—slow learning, but, as Section 3.6 argues, exactly what a tight prior shaped by Peru's hyperinflationary history would generate. Columns (2) and (3) show the estimate is not an artifact of a few clustered events: merging events whose $\pm$7-day windows overlap leaves a coefficient of -1.5, and restricting to high-severity events—larger shocks, which should reveal more insulation—*strengthens* it to -2.1. Section 7.7 reports the full battery of robustness checks. Figure 8 plots the underlying event-level relationship.

Figure 8: The Tenure Effect: Market Response Declines with Central Bank Tenure

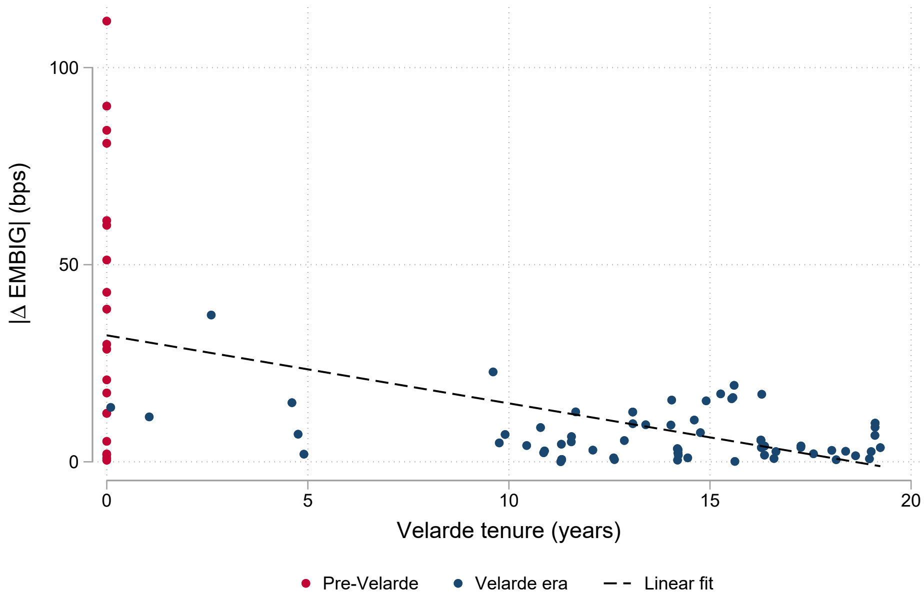

Notes: Absolute event-window change in EMBIG Peru against Velarde's tenure at the event date, with linear fit. A tenure-bucket version appears in Appendix C.

The relationship is monotonic and does not depend on functional form: a nonparametric tenure-bucket summary, in Online Appendix Figure OA.8, shows the same steady decline.

7 Robustness

The decoupling documented in Sections 5 and 6 could in principle reflect forces other than central-bank credibility. This section subjects it to five challenges—global risk controls, direct controls for central-bank swap intervention, an externally coded measure of political risk, a synthetic-control counterfactual, and a battery of checks on the tenure effect—and reports the cross-country evidence on volatility. The result survives all of them. A complementary endogenous structural-break test (Appendix C) confirms that the transformation is gradual—accumulating through the Velarde era—rather than a sharp regime switch, exactly as the learning model predicts.

7.1 Global Controls

The most obvious alternative reading of the post-2006 decoupling is that it reflects not central-bank credibility but the global environment: the commodity supercycle lifted Peru's terms of trade, and the post-crisis search for yield compressed emerging-market spreads everywhere. If the flip in the instability–spread relationship is simply that environment, it should disappear once the environment is controlled for. Table 9 adds two such controls to the monthly era-interaction regression: the VIX (global risk appetite) and the copper price (Peru's dominant terms-of-trade shock).

Table 9: The Decoupling Survives Global Controls

	(1) Baseline	(2) +VIX	(3) +Copper
Instability × Velarde	-136.9^{***}	-74.9^{***}	-63.1^{**}
	(28.2)	(26.7)	(27.9)
Instability (pre-2006 slope)	78.5^{***}	70.8^{***}	66.7^{***}
	(22.5)	(18.4)	(19.4)
Velarde era	-25.7^{**}	-30.4^{**}	-30.5
	(13.0)	(13.8)	(19.9)
VIX		Yes	Yes
Copper price			Yes
Year FE	Yes	Yes	Yes
Observations	311	311	260

Notes: Monthly regressions of EMBIG Peru on the 12-month rolling instability index interacted with the Velarde-era indicator, with year fixed effects in all columns. Robust standard errors in parentheses. Column (1) reproduces the headline specification of Table 6. "Instability × Velarde" is the change in the instability–spread slope across eras—the flip. We do not include the EMBI Global emerging-market aggregate as a control because Peru is one of its constituent components and the political shocks experienced by other emerging markets are themselves correlated with Peruvian political shocks, so conditioning on it would be a "bad control" in the sense of Angrist and Pischke (2009, ch. 3). With year fixed effects in the regression the Velarde-era dummy is largely absorbed by the calendar-year indicators and is not separately interpretable; the coefficient of interest is the interaction. $^{**}p < 0.05$, $^{***}p < 0.01$.

The flip does not disappear, but it does attenuate. Column (1) reproduces the headline specification of Table 6: the interaction is -136.9 basis points. Adding VIX brings it to -74.9 (column 2)—global risk appetite alone accounts for roughly forty-five percent of the original flip, and about fifty-five percent of the magnitude survives. Adding copper on top of VIX leaves the interaction at -63.1 (column 3): once global risk appetite is controlled for, the terms-of-trade channel adds little. The interaction nonetheless remains negative and statistically significant in every column ($p < 0.05$ throughout), and the commodity-boom and search-for-yield explanations account for part of Peru's experience; they do not account for the decoupling. The Velarde-era dummy reported in the table is largely absorbed by the year fixed effects and is not separately interpretable; the substantive object is the interaction, identified by within-year variation in how the instability index covaries with the spread.

7.2 Central-Bank Swap Intervention

A second challenge turns on the central bank's own balance sheet. If the BCRP systematically leaned against instability-driven pressure with its foreign-exchange toolkit, the flat instability–spread relationship of the Velarde era would be a by-product of those operations rather than evidence of anchored expectations. The most direct instrument through which such leaning could operate is the *swap cambiario* (Section 2.4): introduced in September 2014, it lets the BCRP supply the market with exchange-rate hedging without drawing down reserves, and its outstanding stock rose above USD 9 billion during both the 2015 depreciation episode and the 2020–2023 political turmoil. We re-estimate the two central specifications—the monthly instability-index regression and the daily event study—with the BCRP swap position entered directly as a control, measured as the cumulative net notional of swaps cambiarios outstanding, in billions of dollars. Table 10 reports the result.

Table 10: The Decoupling Survives Controls for Central-Bank Swap Intervention

	Interaction		By era		Event study	
	(1) Base	(2) +Swap	(3) Base	(4) +Swap	(5) Base	(6) +Swap
Instability × Velarde	-136.9^{***}	-121.5^{***}				
	(28.2)	(28.7)				
Instability slope (Velarde)			-63.5^{***}	-49.3^{***}		
			(16.7)	(17.4)		
Post-Event × Velarde					-40.1^{***}	-40.3^{***}
					(15.3)	(15.3)
Swap cambiario stock		6.8^{***}		6.7^{***}		4.5^{***}
		(1.2)		(1.2)		(0.7)
Instability slope (pre-2006)	78.5	78.5	79.1	79.2		
Post-Event (pre-2006)					40.9	40.9
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	311	311	311	311	478	478

Notes: Robust standard errors in parentheses; all columns include year fixed effects. Columns (1)–(4) regress monthly EMBIG Peru on the 12-month rolling instability index ($N = 311$), the baseline columns reproducing Table 6; columns (5)–(6) are the daily event-study regression ($N = 478$), the baseline reproducing Table 5. "Swap cambiario stock" is the cumulative net notional of BCRP *swaps cambiarios venta* outstanding (BCRP series PN09504–PN09506TM), in billions of U.S. dollars, zero before the instrument's September 2014 launch. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The decoupling survives. In the monthly interaction regression the Instability × Velarde coefficient—the flip—moves from −136.9 to −121.5 bps once the swap stock is included (columns 1–2), an attenuation of roughly one-tenth, and it remains significant at the one-percent level. The event study does not move at all: the Post-Event × Velarde coefficient is −40.1 bps without the swap control and −40.3 with it (columns 5–6). This is exactly what one expects—the within-window jump that identifies the event-study effect is, by construction, orthogonal to a swap position that is fixed within the event month.

The most demanding comparison is the Velarde-era instability slope estimated directly, which attenuates from −63.5 to −49.3 bps—about a fifth. Two facts keep that 22% in perspective, and without them its size invites misreading. First, the swap control is endogenous:

the swap stock enters every specification with a *positive* sign—a larger book predicts *wider* spreads—because the BCRP scales up its swap operations precisely when markets are under stress. The control therefore absorbs stress variation that belongs to the instability channel itself, so 22% is an *upper bound* on the share of the decoupling that swaps could genuinely explain; the true attenuation is smaller. Second, the swap cambiario did not exist until 2014, whereas the decoupling is documented from 2006 and is already fully present across the eight-year window in which the BCRP held no swap position at all—the control can speak only to the post-2014 sub-period. Consistent with both points, replacing the swap *stock* with the swap *flow* (net monthly placement) leaves the interaction essentially unchanged at −136.1 bps, with the flow itself insignificant.

The pre-2006 instability slope (+78.5 bps) is, mechanically, untouched: the swap instrument did not exist then. The insulation of Peru's sovereign spread from political instability is not an accounting residue of the central bank's swap desk.

7.3 An Externally Coded Measure of Political Risk

A third concern is endogenous measurement. Our political-event index was coded by the authors, who knew how Peruvian spreads had behaved, if the coding unconsciously tracked the spread, the correlation flip could be an artifact of measurement rather than a fact about markets. The cleanest answer is to repeat the test with a measure built by others, with no knowledge of Peruvian sovereign markets. We use the country-specific Geopolitical Risk index for Peru, $\mathrm{GPRC}_{\mathrm{PER}}$, of Caldara and Iacoviello (2022), constructed by automated text analysis of the international press.

The external index reproduces the flip. Table 11 reports the era-specific correlation of each measure with the EMBIG spread. Our hand-coded index correlates +0.55 with the spread before 2006 and −0.23 after, $\mathrm{GPRC}_{\mathrm{PER}}$—which correlates only 0.29 with our index, and so is very nearly an independent measurement—correlates +0.26 with the spread before 2006 and 0.00 after. Re-running the era-interaction regression with $\mathrm{GPRC}_{\mathrm{PER}}$ in place of our index yields an interaction that is large, negative, and significant ($p = 0.005$): the sensitivity of the spread to geopolitical risk present before 2006 is gone afterward. An index built blind to Peruvian markets sees the same decoupling our coding does.

Table 11: An External Political-Risk Index Reproduces the Flip

	Pre-Velarde	Velarde Era
Panel A: correlation of EMBIG with each measure		
Hand-coded instability index	+0.55	−0.23
GPRC_{PER} (Caldara–Iacoviello)	+0.26	0.00
Panel B: era-interaction regression on GPRC_{PER}		
GPRC_{PER} (pre-2006 slope)	$2{,}488^{***}$	
$\text{GPRC}_{\text{PER}} \times$ Velarde	$-2{,}494^{***}$	

Notes: GPRC_{PER} is the country-specific Geopolitical Risk index for Peru of Caldara and Iacoviello (2022), built by automated text analysis with no input from Peruvian financial markets, it correlates 0.29 with the hand-coded index. Panel B regresses monthly EMBIG Peru on GPRC_{PER} interacted with the Velarde-era indicator ($N = 311$), the interaction has $p = 0.005$. $^{***}p < 0.01$.

7.4 Synthetic Control

The sharpest counterfactual question is what Peru's spread would have done after 2006 *absent* the credibility regime. We answer it with a synthetic control (Abadie et al., 2010): a weighted average of comparator economies chosen to match Peru's pre-2007 spread path. The donor pool is the four large Latin American sovereigns, and the optimal weights are Colombia 0.55, Mexico 0.28, Brazil 0.10, and Chile 0.06.

Table 12: Synthetic Control: Donor Weights and the Peru Gap

Panel A: synthetic-control donor weights	
Colombia	0.55
Mexico	0.28
Brazil	0.10
Chile	0.06
Panel B: Peru minus synthetic Peru (bps)	
Pre-Velarde average (2000–2006)	+16
Velarde-era average (2007–2024)	−66
Gap by 2022	−173

Notes: Synthetic Peru is the weighted donor average minimizing the pre-2007 distance to Peru's annual EMBIG spread (Abadie et al., 2010). A negative gap means Peru's spread lies below its synthetic counterpart.

Figure 9 shows the result. Through 2013, actual Peru and synthetic Peru are nearly indistinguishable, the average gap over 2000–2006 is +16 bps, confirming a close pre-treatment fit. After 2013 the two series diverge sharply: synthetic Peru climbs with the regional re-rating—taper tantrum, commodity bust, the pandemic—while actual Peru holds near 150–200 bps. The gap averages −66 bps over the Velarde era and widens to −173 bps by 2022. Peru does not merely fall with the region; it detaches from it. Crucially, the divergence is *gradual and widening*, not a step at 2006—the visual signature of credibility that accumulates, exactly as the learning model of Section 3.6 predicts.

Figure 9: Peru versus Synthetic Peru

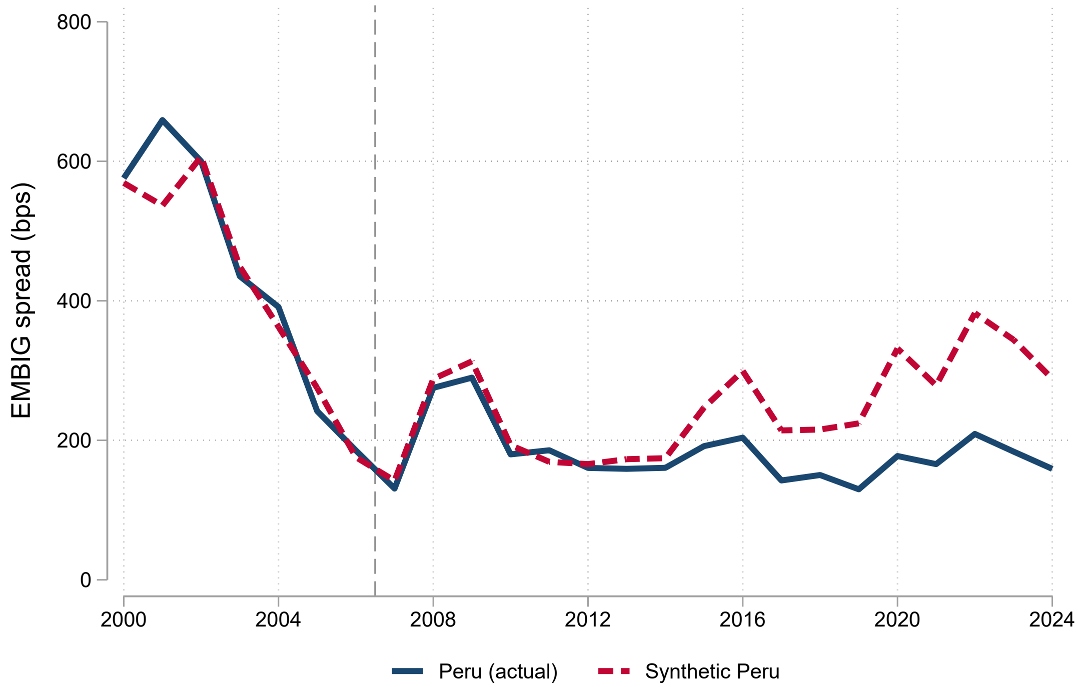

Notes: Actual EMBIG Peru (solid) and synthetic Peru (dashed), a weighted average of Colombia, Mexico, Brazil, and Chile chosen to match Peru's pre-2007 spread. The vertical line marks September 2006.

7.5 Cross-Country Volatility Reduction

Proposition 2 predicts that the collapse in the political component of spread variance should be Peru-specific. Table 13 reports the pre/post variance ratio of daily spread changes for Peru and four comparators, under three samples: the full series, a symmetric "normal-times" sample that drops every crisis window (2001–02, 2008–09, 2020) from both eras, and the post-2003 sample.

Table 13: Variance Ratio of Daily Spread Changes: Cross-Country Comparison

Country	Full sample	Normal times (crises excl.)	Post-2003 sample
Peru	3.6	**5.1**	2.0
Chile	0.9	0.9	0.2
Colombia	2.1	3.8	1.2
Mexico	0.8	1.5	0.3
Brazil	10.8	5.8	3.0

Notes: Pre-Velarde to Velarde-era ratio of the variance of daily EMBIG changes. "Normal times" drops the 2001–02, 2008–09, and 2020 crisis windows symmetrically from both eras. Brazil's elevated full-sample ratio reflects the idiosyncratic 2002 Lula-transition crisis, under the normal-times comparison it falls to a level comparable to Peru's.

Two readings stand out. First, raw variance ratios are confounded by crisis timing. Brazil's full-sample ratio of 10.8 is the largest in the table—but it is an artifact of the 2002 Lula-transition crisis, when Brazilian daily-spread volatility reached six times its normal level, under the symmetric normal-times comparison Brazil's ratio falls to 5.8, comparable to Peru's. We therefore make no claim that Peru's variance reduction is the largest in the region: raw ratios are simply too sensitive to when each country had a crisis. Second, and more informative, is what happens to *Peru* under the normal-times comparison: its ratio *rises*, from 3.6 in the full sample to 5.1. This is exactly the mechanism. Stripping out global crises removes variation a credible central bank cannot, and is not expected to, neutralize—the common component σ^2_ν. Peru's Velarde-era variance falls by 43% when crisis windows are removed, against only 20% for the pre-2006 variance, so the ratio climbs. Credibility operates on *political* shocks, not on the global shocks the central bank cannot offset, in normal times, when political shocks are what remains, the insulation is most visible. The cleaner evidence for the Peru-specific mechanism, however, is not the raw ratio but the tenure effect, the synthetic control, and the survival of the flip under global controls.

7.6 Cross-Country Evidence: CBI and the Inflation Channel

A natural question is whether the credibility-as-insulation mechanism generalizes beyond Peru. A direct cross-country test of the full event-study interaction is precluded by data limitations: continuous proxies for central-bank credibility comparable to Velarde's tenure are scarce, and the four-country event-study panel used in Section 7.5 is under-powered for the interaction. We instead test a more limited claim that maps to a specific component of the model: that central-bank credibility moderates the pass-through of inflation deviations to sovereign spreads—the inflation channel $\delta/(1+\tilde{\theta})$ of the amplification function $\Gamma(\tilde{\theta})$ in equation (13). This is not the full mechanism—which runs from political events through government quality to inflation to spreads—but it is one of its identifiable components.

Table 14 reports an annual panel for six Latin American economies (Argentina, Brazil, Chile, Colombia, Mexico, Peru) over 2000–2023. The dependent variable is the annual average EMBIG spread; the key regressors are the absolute deviation of inflation from each country's target (proxying for inflation-channel pressure) and its interaction with the Garriga (2016) de jure central-bank-independence index. Country and year fixed effects absorb cross-country level differences and global shocks.

In the preferred specification (column 1), the slope of the spread on inflation deviation is 71.3 basis points per percentage point at CBI $= 0$, falling by 94.6 basis points per unit of CBI. The implied country-specific slopes—roughly $+57$ bps/pp at Brazil's pre-2021 CBI of 0.15, $+24$ bps/pp at the moderate Latin American CBI of 0.50, and approximately -5 bps/pp at Peru's CBI of 0.81—are consistent with the prediction that the inflation channel of $\Gamma(\tilde{\theta})$ shuts down as central-bank inflation aversion rises. The interaction is significant at the one-percent level ($p = 0.005$) and survives excluding Argentina (column 2, $p = 0.035$), so the result is not an artifact of the Argentine saturation case. Substituting the inflation level for the deviation (columns 5–6) delivers the same pattern.

Three limitations bound the interpretation. First, the Garriga (2016) index measures *de jure* independence—the statutory rules governing the central bank—and is therefore essentially flat for Peru throughout our sample; the *de facto* reputational accumulation that Section 6.2 identifies within Peru is not captured here. Second, the annual frequency abstracts from the event-window identification of Section 5: we gain cross-country variation but lose the sharpness of the daily event study. Third, the six-country sample is modest, and Argentina contributes only post-2018 inflation observations following the World Bank's exclusion of the disputed 2007–2017 INDEC series; expanding the donor pool to a broader set of emerging markets is a natural follow-up. Subject to these caveats, the test provides cross-country support for the inflation channel of the model; it is silent on the full event-to-spread chain that the within-Peru analysis identifies.

Table 14: Cross-Country Annual Panel: CBI Moderates the Inflation Channel

	(1)	(2)	(3)	(4)	(5)	(6)
	Moderator: |Inflation−target|				Moderator: inflation level	
	Full	Excl ARG	Cross ctry	Cross, excl ARG	Full	Excl ARG
|Inflation − target|	71.3^{***}	68.9^{***}	66.7^{**}	60.9^{**}		
	(12.0)	(11.0)	(33.7)	(30.6)		
× CBI	-94.6^{***}	-124.2^{**}	−72.9	-130.6^{*}		
	(19.7)	(39.7)	(51.3)	(72.4)		
Inflation level					63.1^{**}	58.6^{***}
					(14.5)	(11.4)
× CBI					-82.5^{**}	-110.6^{*}
					(24.0)	(43.8)
Country FE	Yes	Yes	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	126	120	126	120	126	120
R-squared	0.78	0.73	0.41	0.24	0.78	0.73

Notes: Annual panel, six Latin American countries (Argentina, Brazil, Chile, Colombia, Mexico, Peru), 2000–2023. Dependent variable: annual average EMBIG sovereign spread (basis points). CBI is the Garriga (2016) de jure central-bank-independence index. Inflation from the World Bank (`FP.CPI.TOTL.ZG`); country targets are 2% for Peru, 3% for Chile, Colombia, and Mexico, 4.5% for Brazil, and 5% for Argentina. Standard errors clustered at the country level in columns (1), (2), (5), and (6); heteroskedasticity-robust in (3) and (4). Argentina contributes only 2018–2024 observations because the World Bank excludes the disputed 2007–2017 INDEC inflation series. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7.7 Robustness of the Tenure Effect

The tenure regression of Section 6.2—the model's most distinctive prediction—rests on 90 events, and a referee might worry that a handful of them drive the result. They do not. The tenure coefficient survives every check we have applied. Of 90 leave-one-out replications, all 90 retain significance at the 1% level: no single event drives the estimate. A wild cluster bootstrap with 9,999 replications yields $p < 0.0001$, with a 95% confidence set of $[-2.6, -0.9]$, well clear of zero. De-clustering the events—merging those within fourteen days of each other, so overlapping windows cannot double-count—leaves a coefficient of -1.5 ($p < 0.01$). Restricting to high-severity events alone *strengthens* it to -2.1, as the model predicts: larger shocks should reveal more insulation. Dropping the entire 2022–2023 crisis cluster, the most influential block of events, leaves -1.8. These checks are collected in the four columns of Table 8.

A final placebo test asks whether the tenure pattern is Peru-specific or a spurious global time trend. If it were the latter, the absolute daily spread change of *other* countries should also decline over the same calendar window. Table 15 regresses each country's absolute daily EMBIG change on Velarde's tenure clock. Peru shows a strong, highly significant decline, Chile—a credible-central-bank economy throughout—shows none, Colombia shows a smaller, significant decline, plausibly reflecting its own institutional gains. The tenure effect is not a global trend.

Table 15: Tenure Effect: Peru versus Placebo Countries

Dependent Variable:	(1) Peru $\lvert\Delta\rvert$	(2) Chile $\lvert\Delta\rvert$	(3) Colombia $\lvert\Delta\rvert$
Tenure (years since 2006)	-0.204^{***}	-0.014	-0.148^{***}
	(0.019)	(0.012)	(0.030)
Constant	6.28^{***}	2.88^{***}	6.76^{***}
	(0.26)	(0.17)	(0.36)
Observations	2,755	2,755	2,304
R^2	0.056	0.001	0.014

Notes: Robust standard errors in parentheses. Dependent variable is the absolute daily EMBIG change, sample restricted to the Velarde era. Tenure is years since September 2006. $^{***}p < 0.001$.

Taken together, the five challenges of this section—global controls, central-bank swap intervention, an external risk measure, a synthetic counterfactual, and the tenure-effect checks—leave the decoupling intact. It is not the commodity boom, not the authors' coding, and not a regional trend; it is Peru-specific and tied to the central bank.

8 Conclusion

Can a good institution substitute for bad ones? Peru between 2000 and 2025 is as close to a natural experiment on that question as the historical record offers. Political instability rose to levels that, by any conventional account, should have made the country uninvestable—seven presidents in seven years, a self-coup, the killing of protesters. Sovereign borrowing costs did the opposite: they fell to historic lows, and the market's response to a political shock fell from 41 basis points before 2006 to a precisely estimated zero.

We have shown that the resolution lies in a single credible institution. The Central Reserve Bank of Peru under Julio Velarde accumulated, over nearly two decades, the kind of reputation that lets markets price Peruvian debt on fundamentals rather than on the latest headline from Congress. Our model—combining monetary-policy delegation, sovereign-spread pricing, and Bayesian learning about the governor's type—rationalizes the full set of facts: the correlation flip, the more-than-threefold variance reduction, and the decay of the market response with the governor's tenure. The decoupling survives global controls, an externally coded measure of political risk, and a synthetic-control counterfactual, endogenous structural-break tests confirm it is a gradual accumulation, not a switch thrown in 2006, exactly as the learning model implies.

The general lesson is that central-bank credibility is not merely a tool for controlling inflation. It is a form of institutional capital that can insulate an economy's external borrowing costs from political dysfunction the rest of the state cannot contain. For the many emerging economies that pair fragile political institutions with a central bank of uncertain independence, the Peruvian case is an existence proof: one credible institution, sustained long enough, can carry a disproportionate share of the macroeconomic burden.

The honest qualification is that this insulation is fragile. Credibility of the kind Velarde built accumulates slowly—our estimates imply a decade or more—and the model's own logic implies it can be lost far faster than it is gained. In Peru it rests substantially on the continuation of one individual and the policy regime he represents, the multiplicity results of Appendix B show that a sufficiently adverse succession signal could, in principle, shift the economy to an uncredible equilibrium with no change in fundamentals. A paper that sold only the good news would be incomplete.

That fragility is the natural next question. Whether Peru can *institutionalize* the credibility that Velarde embodies—converting the reputation of a person into the reputation of an office—is the central issue for its economic future, and the dynamics of credibility accumulation and loss are the subject of ongoing work.

References

Alberto Abadie, Alexis Diamond, and Jens Hainmueller. Synthetic control methods for comparative case studies. *Journal of the American Statistical Association*, 105(490):493–505, 2010.

Cristina Arellano. Default risk and income fluctuations in emerging economies. *American Economic Review*, 98(3):690–712, 2008.

Jushan Bai and Pierre Perron. Estimating and testing linear models with multiple structural changes. *Econometrica*, 66(1):47–78, 1998.

Jushan Bai and Pierre Perron. Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1):1–22, 2003.

Robert J. Barro and David B. Gordon. A positive theory of monetary policy in a natural rate model. *Journal of Political Economy*, 91(4):589–610, 1983.

Dario Caldara and Matteo Iacoviello. Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225, 2022.

Alex Cukierman and Allan H. Meltzer. A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. *Econometrica*, 54(5):1099–1128, 1986.

Wioletta Dziuda and Carolin Pflueger. A model of politics and the central bank. Working paper, University of Chicago Harris School of Public Policy, 2025.

Jonathan Eaton and Mark Gersovitz. Debt with potential repudiation: Theoretical and empirical analysis. *Review of Economic Studies*, 48(2):289–309, 1981.

Ana Carolina Garriga. Central bank independence in the world: A new data set. *International Interactions*, 42(5):849–868, 2016. doi: 10.1080/03050629.2016.1188813.

Juan Carlos Hatchondo, Leonardo Martinez, and Horacio Sapriza. Heterogeneous borrowers in quantitative models of sovereign default. *International Economic Review*, 50(4):1129–1151, 2009.

Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–491, 1977.

Paolo Mauro, Nathan Sussman, and Yishay Yafeh. *Emerging Markets and Financial Globalization: Sovereign Bond Spreads in 1870–1913 and Today*. Oxford University Press, Oxford, 2007.

Carolin E. Pflueger and Pierre Yared. Global hegemony and exorbitant privilege. Working paper, University of Chicago and Columbia Business School, 2025.

Kenneth Rogoff. The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100(4):1169–1189, 1985.

A Properties of the Amplification Function

The amplification function $\Gamma(\theta) = \alpha + \delta/(1+\theta)$ satisfies

$$\begin{aligned}
\Gamma'(\theta) &= -\frac{\delta}{(1+\theta)^2} < 0 && \text{(strictly decreasing),} \\
\Gamma''(\theta) &= \frac{2\delta}{(1+\theta)^3} > 0 && \text{(strictly convex),} \\
\lim_{\theta \to 0^+} \Gamma(\theta) &= \alpha + \delta, \\
\lim_{\theta \to \infty} \Gamma(\theta) &= \alpha.
\end{aligned}$$

Convexity justifies the positive correction term in Proposition 3: since investors evaluate $\mathbb{E}[\Gamma(\tilde{\theta})]$ with $\tilde{\theta}$ uncertain, Jensen's inequality implies the expectation exceeds $\Gamma(\mathbb{E}[\tilde{\theta}])$, with the gap shrinking as the posterior variance $v_T \to 0$.

B Multiplicity, Fragility, and the Velarde Premium

This appendix endogenizes effective CB independence, following the multiplicity–fragility logic of Pflueger and Yared (2025, henceforth P-Y). It generates no prediction tested in the body of the paper and is therefore separated here, the dynamic extension (credibility hysteresis and the succession risk premium) is developed in ongoing work.

Let $p_t \in [0,1]$ denote political pressure on the CB. Effective inflation-aversion is $\tilde{\theta}_t^{\text{eff}} = (1-p_t)\tilde{\theta} + p_t \underline{\theta}$ with $\underline{\theta} < \tilde{\theta}$, and pressure rises with the spread, $p_t = p(s_t) = (s_t - \underline{s})^+/(\bar{s} - \underline{s})$. Substituting into (13) with $g_t = 0$ yields the fixed-point equation $s = \mathcal{G}(s) \equiv \bar{s} - \gamma \bar{\mathcal{F}}(\tilde{\theta}_t^{\text{eff}}(s))$. Because $\mathcal{G}$ is increasing in s, it can cross the 45-degree line multiple times.

Proposition 4 (Multiplicity, cf. P-Y Proposition 2)**.** *Define the complementarity index $\phi \equiv \gamma\,\delta\,u^*(\tilde{\theta} - \underline{\theta})/[\tilde{\theta}\,\underline{\theta}(\bar{s} - \underline{s})]$. There exist thresholds $0 < \phi_L < \phi_H$ such that: if $\phi < \phi_L$ a unique high-credibility equilibrium s^H exists, if $\phi \in (\phi_L, \phi_H)$ two stable equilibria coexist—s^H (the Velarde equilibrium) and $s^L > s^H$ (the uncredible equilibrium), if $\phi > \phi_H$ only the uncredible equilibrium exists.*

Proof. Fix $g_t = 0$. With $p(s) = (s - \underline{s})/(\bar{s} - \underline{s})$ the fixed-point condition is $\gamma\delta|u^*|/\tilde{\theta}^{\text{eff}}(s) = \bar{s} - \gamma\bar{F} - s \equiv \ell(s)$. The left side is strictly increasing in s and the right side strictly decreasing, so at least one intersection exists. A second exists iff the slope of the left side exceeds unity at some interior point, evaluating at $s = \underline{s}$ gives $\phi_L = \gamma\delta|u^*|(\tilde{\theta} - \underline{\theta})/[\tilde{\theta}^2(\bar{s} - \underline{s})]$. The upper threshold ϕ_H follows from requiring both intersections to lie in the interior of $[\underline{s}, \bar{s}]$. ☐

Figure 10 shows the geometry. As the complementarity index ϕ rises, the fixed-point map $\mathcal{G}$ steepens, an intermediate range of ϕ makes it cross the 45-degree line three times, producing the two stable equilibria of the multiplicity region.

Figure 10: Multiplicity and the Velarde Premium

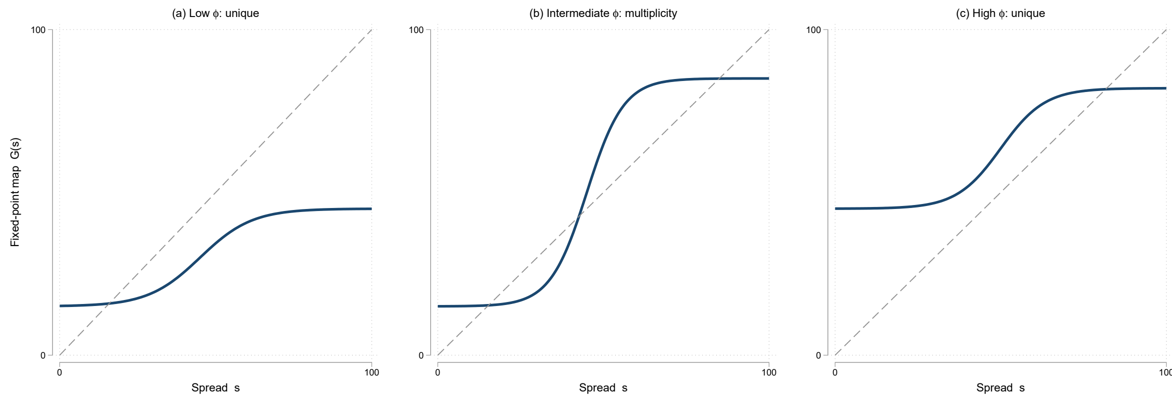

Notes: The fixed-point map $s = \mathcal{G}(s)$ against the 45-degree line, for three values of the complementarity index ϕ. Panel (a): for $\phi < \phi_L$, a unique high-credibility (low-spread) equilibrium. Panel (b): for $\phi \in (\phi_L, \phi_H)$ the map crosses the 45-degree line three times—two stable equilibria, s^H and s^L, separated by an unstable one, their gap is the Velarde Premium. Panel (c): for $\phi > \phi_H$, a unique uncredible (high-spread) equilibrium. Illustrative parameter values.

Corollary 3 (Velarde Premium and fragility)**.** *In the multiplicity region the Velarde Premium $\Pi \equiv s^L - s^H > 0$ is increasing in $\tilde{\theta}$ and $\gamma\delta$ and decreasing in $\underline{\theta}$. Transitions from s^H to s^L can occur with no change in fundamentals, coordinated solely by investor expectations—an adverse succession signal that lowers the posterior on $\tilde{\theta}$ can trigger a self-fulfilling jump.*

This static fragility result is the equilibrium-selection counterpart of the dynamic, history-dependent fragility (credibility hysteresis, threshold effects, and a forward-looking succession risk premium) developed in ongoing work.

C Endogenous Structural-Break Tests

A natural question is whether the instability–spread relationship contains a sharp structural break, and if so when. This appendix applies endogenous break tests that do not impose a break date.

Baseline. Applied to the monthly series of EMBIG Peru and the instability index (2000–2025), a supremum-Wald test locates a single break in late 2004, and a Bai–Perron procedure (Bai and Perron, 1998, 2003) allowing multiple breaks selects 2004m11, 2008m11, and 2017m11. None coincides with Velarde's October 2006 appointment. This is unsurprising: a break test on the *level* of the spread detects the large secular decline—which began under the commodity boom well before 2006—not the change in the spread's *sensitivity* to political shocks.

Isolating the relationship. Three specifications strip the secular trend out before testing. (a) In first differences—ΔEMBIG on Δinstability—the supremum-Wald statistic is 2.1, below even the 10% critical value: once the trend is differenced away, there is no sharp break at all. (b) Residualising both series on a linear trend and the EM-wide spread, then testing the relationship between the residuals, yields a significant break in 2017m6. (c) Testing the synthetic-control gap (Peru minus synthetic Peru) yields a significant break in 2012m5.

The tests are mutually consistent, and consistent with the model. The 2004 "break" is the secular trend, not the regime change, once the trend is removed, no sharp break survives, and what the residualised specifications detect lies *within* the Velarde era (2012–2017), when regional spreads re-rated upward and Peru did not follow. The model predicts exactly this: credibility accumulates through Bayesian learning, so the transformation is gradual, not a switch thrown in 2006. The endogenous break tests therefore corroborate the learning mechanism rather than a sharp regime dummy.

Online Appendix

When Bad Institutions Meet Good Ones: The Peruvian Puzzle

Carlos Chávez   Luis Yépez

This online appendix collects supplementary figures that complement the main text. It is organized in four sections: instability indices, additional event-study evidence, an alternative view of the puzzle, and supplementary model-prediction tests.

OA.1 Political Instability Indices

The following figures present the political instability index in alternative forms, all show that instability rose markedly after 2006. The raw annual event count appears alongside the severity-weighted index in Figure 2 of the main text.

Figure OA.1 shows the cumulative severity-weighted instability index. Its steepening slope after 2016 makes visually explicit that the bulk of Peru's accumulated political instability is concentrated in the Velarde era.

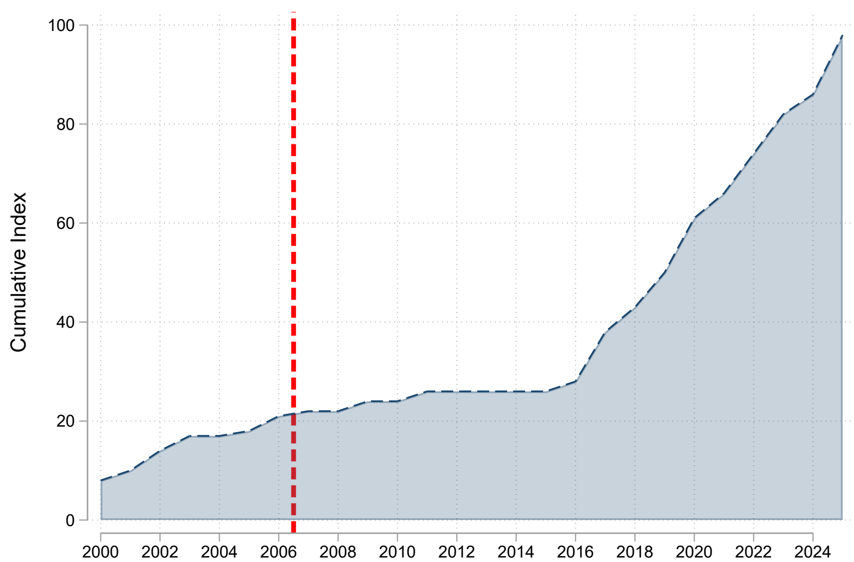

Figure OA.1: Cumulative severity-weighted political instability.

Figure OA.2 disaggregates the index to monthly frequency, with the 12-month rolling average used in the regressions superimposed. The monthly bars show the discrete clustering of events around specific crises, while the rolling line traces the smooth trend exploited in the index regressions of Section 5.

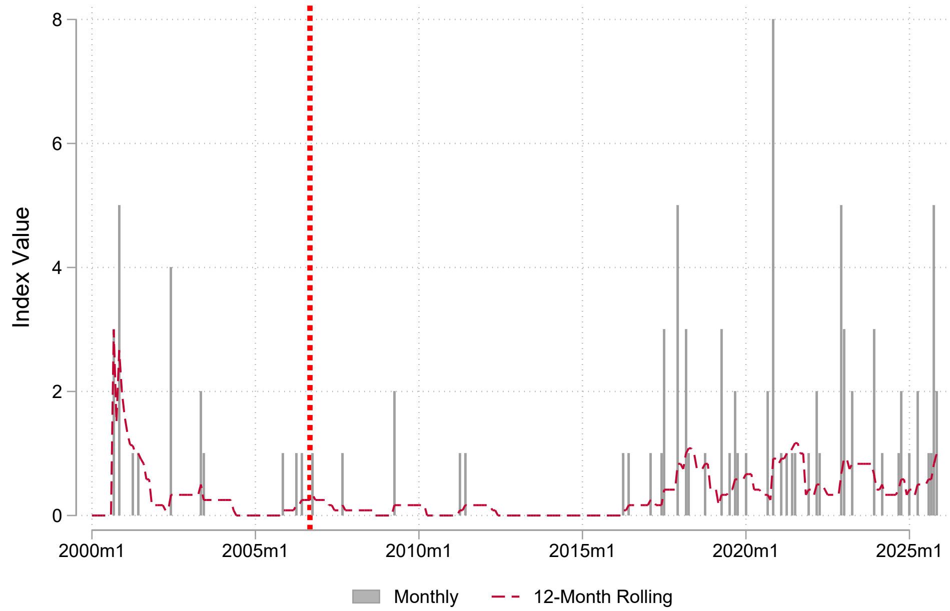

Figure OA.2: Monthly instability index with 12-month rolling average.

OA.2 Additional Event-Study Evidence

Figure OA.3 reports the event study pooling all political events without distinguishing by era. The coefficient path is small and its confidence band includes zero throughout. This pooled null is uninformative about the paper's central comparison precisely because it averages over the sharp pre/post-2006 heterogeneity documented in Section 5, we therefore relegate it here.

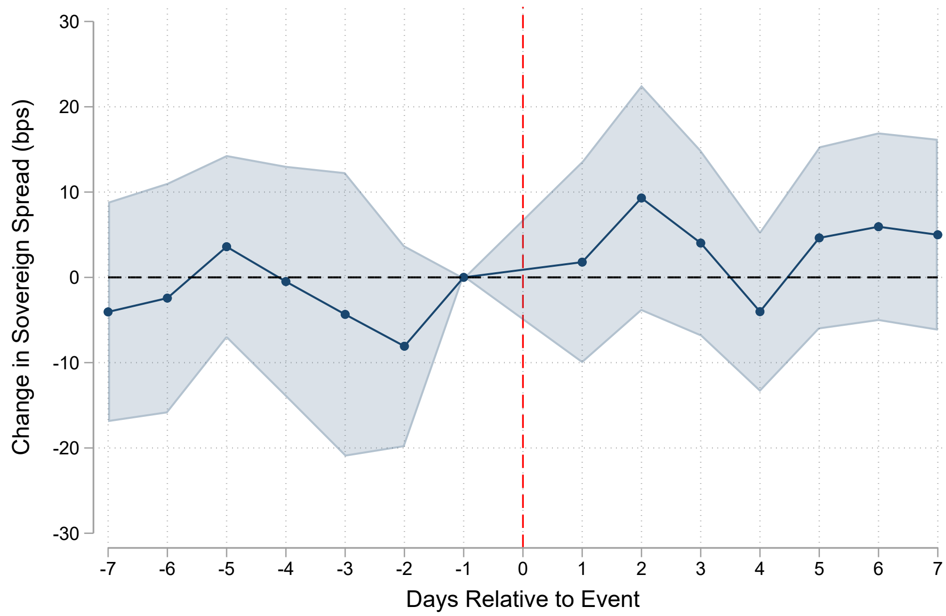

Figure OA.3: Event study, all political events pooled. The pooled response is small and statistically insignificant, masking the heterogeneity by era.

Figure OA.4 repeats the event study with the sol/dollar exchange rate as the outcome. Neither era shows a significant response, indicating that exchange-rate policy was insulated from political shocks throughout the sample—consistent with the BCRP's managed-float regime and large reserves, and confirming that the result in the main text is specific to the sovereign-spread channel.

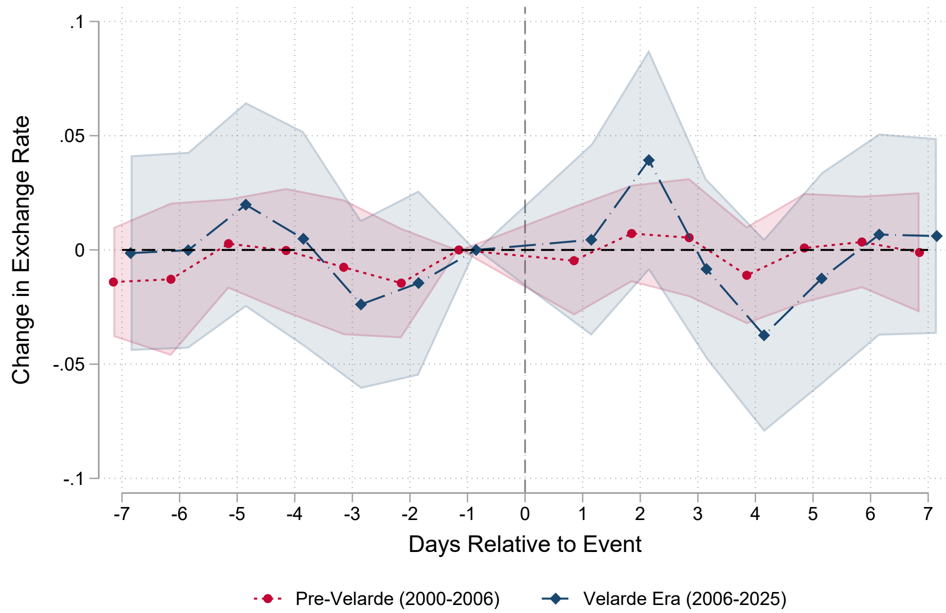

Figure OA.4: Exchange-rate response to political events, by era. Neither era shows a significant response.

Figure OA.5 summarizes the difference-in-differences estimates as a bar chart: the average post-event change in the spread is large and positive before 2006 and statistically indistinguishable from zero afterward. It conveys the content of Table 5 in a single picture.

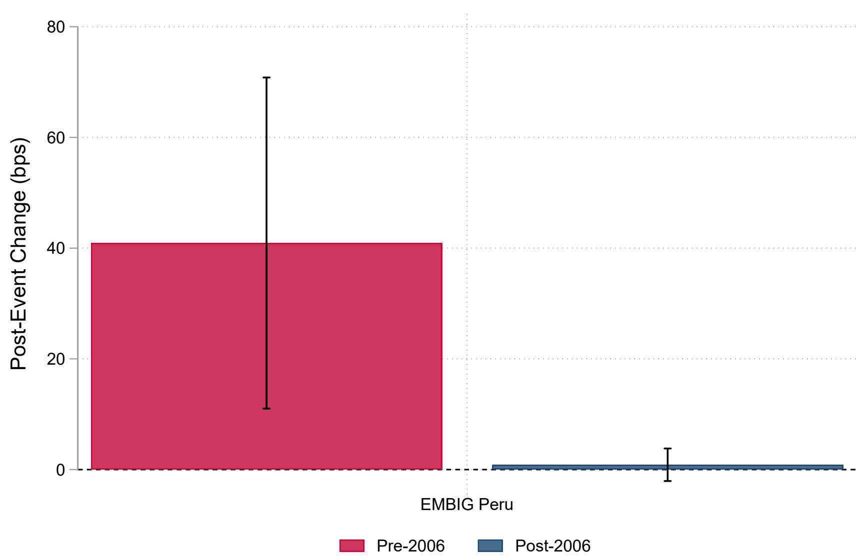

Figure OA.5: Average post-event change in EMBIG Peru by era (difference-in-differences).

OA.3 The Puzzle: An Alternative Visualization

Figure OA.6 offers a dual-axis view of the puzzle: the 12-month rolling instability index (left axis) plotted against the EMBIG spread (right axis). The two series broadly co-move before 2006 and diverge sharply afterward—instability climbs while the spread falls—restating the puzzle of Section 1 in a single panel.

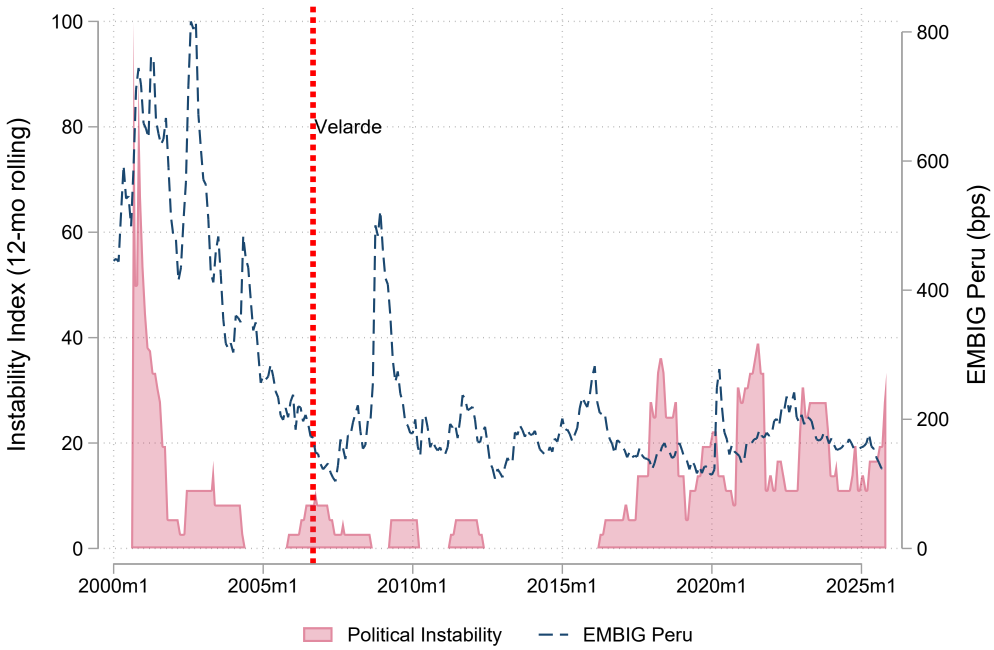

Figure OA.6: Dual-axis view of the puzzle: 12-month rolling instability index (left axis) against EMBIG Peru (right axis).

OA.4 Model-Prediction Tests: Supplementary Figures

Figure OA.7 overlays the distributions of daily spread changes in the two eras. The pre-Velarde distribution has visibly heavier tails, while the Velarde-era distribution is concentrated tightly around zero. This is the distributional counterpart of the variance-ratio test of model Prediction 1.

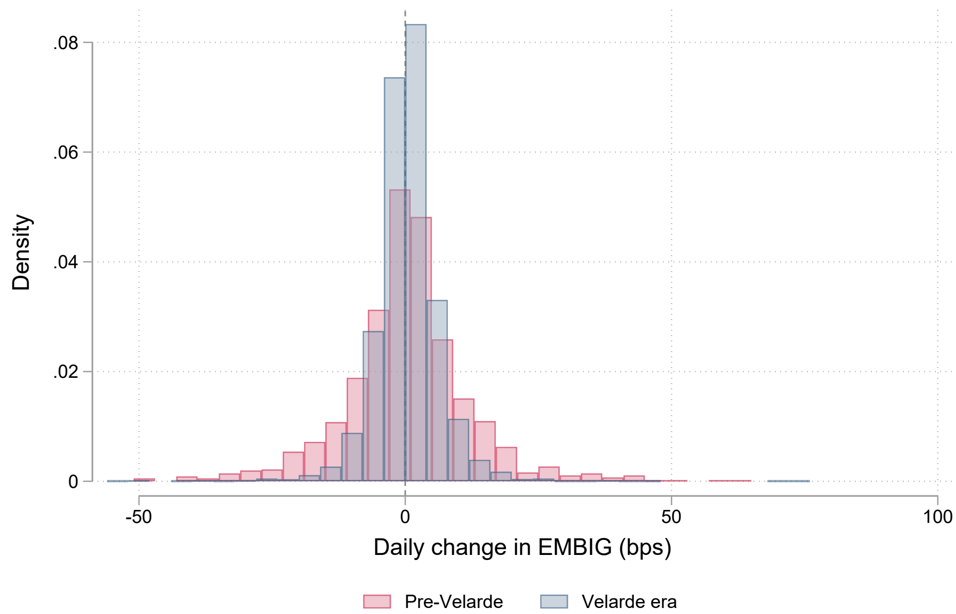

Figure OA.7: Distribution of daily EMBIG changes, by era. The pre-Velarde distribution has substantially heavier tails.

Figure OA.8 groups events into central-bank tenure buckets and plots the mean absolute market response in each. The response falls monotonically across buckets—from roughly 33 bps for pre-Velarde events to single digits at long tenure—providing the non-parametric analogue of the tenure regression of model Prediction 2.

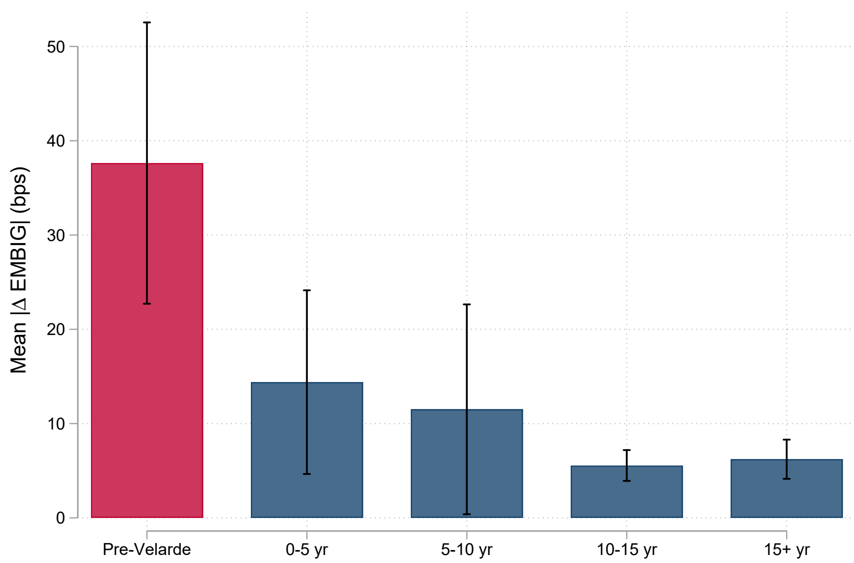

Figure OA.8: Average absolute market response to political events, by central-bank tenure bucket.