Geoeconomic Pressure, Sovereign Risk, and Asset Prices: Evidence from Mexico

Stefano Lord-Medrano*

Job Market Paper

University of Wisconsin-Madison

This version: September 2, 2026

Abstract

A powerful country can raise a smaller sovereign's borrowing costs without imposing a single tariff or sanction, simply by threatening to, repricing its default risk well before any of that risk is realized. I study this channel using the pressure that Donald Trump directed at Mexico during his 2016 campaign and first year in office, a setting that isolates the geoeconomic shock from Mexico's own fiscal position. Using a synthetic-control estimator on daily credit default swap data, I find that this pressure raised Mexico's sovereign spread by 24 basis points on average. A textual analysis of Trump's social-media posts corroborates this interpretation: the spread responds to his threats, more to trade threats than migration threats. I then trace how this Mexico-specific repricing transmits to asset prices: a one-basis-point increase in the spread depreciates the peso by 5.3 cents and lowers Mexican equity returns by 0.11 percent, with financial and tradable-goods sectors most exposed. A quantitative sovereign-default model, in which the pressure tilts the income process toward adverse states, rationalizes the repricing as the price of a genuine shift toward adverse income states and implies a welfare cost of 0.02 percent of lifetime consumption for the transitory episode, and an average consumption contraction of 0.13 percent. As complementary, higher-frequency evidence that avoids time aggregation, filtering and sufficient-statistic methods put the consumption drop at 0.35 percent.

JEL classification: C32, E44, F13, F31, F34, F50, F51, O16

Keywords: Sovereign default, geoeconomic pressure, sovereign risk, synthetic controls, high-frequency identification, trade uncertainty, asset pricing, filtering methods

---

*Department of Economics, University of Wisconsin-Madison, 1180 Observatory Dr, Madison, WI 53706. Email: lordmedrano@wisc.edu. I am extremely grateful to my advisors Louphou Coulibaly, Lydia Cox, Kim Ruhl, and Ken West for numerous discussions regarding this work. I would also like to express my gratitude to Miguel Acosta, Hengjie Ai, Naoki Aizawa, Manuel Amador, Yan Bai, Scott Baker, Panle Barwick, Job Boerma, Carter Braxton, Yong Cai, Diego Cid, Jack Collison, Dean Corbae, Duong Dang, Charles Engel, Cutberto Frias, Agustín Gutiérrez, Rafael Jiménez-Durán, Tim Kehoe, Rishabh Kirpalani, Jennifer La'O, Jose Lara, Fernando Lopes, Fabrizio López-Gallo, João Monteiro, Matias Moretti, Michael Nattinger, Pascal Noel, Zachary Orlando, Pablo Ottonello, Alberto Romero, Jesse Schreger, Yuya Shimizu, Walter Steingress, Alexander von Hafften, Mengqi Wang, and Yeonggyu Yun for their helpful comments as well as participants at the 2026 Econometric Society NASM, the 18th FIW Research Conference 'International Economics', the 2026 and 2024 UMN-UW Workshop, the 2024 JIE Summer School, the 2023 Midwest Macroeconomic Conference, the Graduate Student Workshop in Macroeconomics and International Economics at UW-Madison, and seminar participants of the Financial Stability Division at Banco de México. This work received support from the Institute for Humane Studies under grant no. IHS021029. Refine.ink was used to proofread the paper for consistency and clarity. All remaining errors are my own.

1 Introduction

A powerful country can raise a smaller government's borrowing costs without imposing a single tariff or sanction, simply by threatening to. Threats to restrict trade, investment, or migration can shift the distribution of the smaller economy's future income toward adverse states and thereby reprice its default risk. The 2016 U.S. presidential campaign and the first year of the Trump administration, marked by a sharp rise in anti-Mexican rhetoric and explicit policy threats, provide a particularly useful setting in which to study this form of geoeconomic pressure. In particular it raises two important questions. First, how much did foreign political pressure change investors' perceptions of a country's sovereign risk, and which threats mattered most? Second, how does the resulting repricing affect the exchange rate, stock market, and aggregate consumption?

To answer these questions, I proceed in three empirical steps and then quantify the cost of the episode in two ways. First, I quantify how much the geoeconomic pressure raised Mexico's perceived default risk. Using a synthetic-control framework with daily credit-default-swap (CDS) data, I construct a counterfactual trajectory of Mexico's sovereign risk in the absence of the pressure. Second, I use a textual analysis of Trump's social-media posts, combined with the synthetic-control estimates, to show that this repricing lines up with Trump's threats, and to identify which elements of the "America First" agenda, trade or migration threats, mattered most for markets. Third, I estimate how this Mexico-specific repricing transmits to financial markets, the exchange rate and equity prices, using local projections and high-frequency event data. To rationalize the repricing of sovereign risk and quantify its cost, I build a structural sovereign default model in which the pressure shifts the income process toward adverse states. As a complementary measure, I derive a sufficient statistic from the same model's intertemporal efficiency condition, under a lighter set of assumptions, and combine it with high-frequency asset prices and the synthetic-control estimates. The two are complementary rather than competing: the structural model delivers the normative welfare cost, while the data-driven approach recovers the high-frequency consumption response the model cannot.

I first measure how much the pressure raised Mexico's sovereign risk. Using a synthetic-control estimator on daily 5-year CDS spreads, I build a counterfactual for Mexico from a donor pool of emerging market economies with comparable credit ratings, and read the gap between Mexico's observed CDS and this synthetic control as the component of Mexico's default risk attributable to the pressure. The gap opens right after Trump's nomination in July 2016 and widens around the first presidential debate, the election, the inauguration, and the rounds of NAFTA renegotiation, peaking near 60 basis points and averaging 24. It is robust to trimming the donor pool, to leave-one-out removal, and to purging the CDS data from U.S. monetary-policy movements, and additional exercises reveal no anticipation effects and negligible spillovers. Permutation-based tests confirm that the premium is statistically significant and not an artifact of the persistence in the CDS series.

Having measured the repricing, I ask whether the gap behaves as the coercion interpretation requires. If the wedge reflects pressure on Mexico, it should display two properties: sensitivity, rising when the U.S. president threatens the Mexican economy and easing when he signals willingness to preserve the trade agreement, and targeting, reacting to threats aimed at Mexico rather than at other economies. I test both using a textual analysis of President Trump's posts on 𝕏 (formerly Twitter), a high-frequency, precisely timestamped record of his public statements. The premium rises on days when Trump threatens Mexico and responds to trade threats rather than migration-related ones, and it moves only with pressure aimed at Mexico, not with threats directed at other economies, the sole exception being China, whose threats lower the premium, just as the targeting logic predict, since pressure on China diverts global value chains and investment toward Mexico. Both patterns tie the premium directly to the geoeconomic pressure.

Having established that the gap is a geoeconomic repricing of Mexico's sovereign risk, I trace how it transmits to asset prices. Feeding the premium into local projections, I find that a one-basis-point rise depreciates the peso by 5.3 cents, about 0.26 percent, and lowers Mexican equities by 0.11 percent on impact, with firms in the financial and tradable-goods sectors the most exposed. At the peak of the episode, these movements are comparable in magnitude to those around the collapse of Lehman Brothers and the onset of Covid-19. To confirm that the results are not an artifact of the synthetic-control method, I re-identify the responses with high-frequency instruments built around "Trump days", days carrying threats against Mexico or major U.S. electoral events during the 2016-2017 period, and recover almost identical magnitudes. A further battery of robustness checks leaves the effects on the peso and equities unchanged.

To interpret the sovereign-risk repricing and quantify its welfare consequences, I build a quantitative sovereign-default model with risk-averse lenders. Geoeconomic pressure enters as a regime shift that tilts the distribution of future income toward adverse states. The pressure is therefore not modeled as noise or a change in beliefs divorced from fundamentals; it represents a genuine, though transitory, deterioration in the economic environment that the government and lenders price in common. Calibrated to Mexico, the model reproduces an average increase in sovereign spreads close to the synthetic-control estimate as well as standard business-cycle features of a small open economy. The model implies a consumption-equivalent welfare cost of 0.02 percent for the observed transitory episode. If the pressure were instead permanent, the cost would rise to 0.06 percent. These magnitudes fall within the range reported in the quantitative sovereign-default literature and are comparable, for example, to the welfare cost of world-interest-rate uncertainty estimated by Johri et al. (2022).

Because the model is calibrated at business-cycle frequency, time-aggregating the daily pressure into quarterly data can attenuate the consumption response it recovers. I therefore complement it with two higher-frequency measures. The first is a sufficient statistic that iterates the model's own intertemporal optimality condition, under isoelastic preferences and a few simplifications, and reads the consumption loss directly from asset prices. The second is a purely data-driven daily consumption series, constructed by Kalman-filtering a monthly proxy for aggregate consumption together with daily equity returns, on which I estimate the marginal response to the premium. These approaches imply an overall consumption drop of 0.35 percent over the episode and a marginal response of 0.02 percent per basis point. The overall figure is the empirical counterpart to the model's average quarterly contraction of 0.13 percent; that it is larger reflects several forces I discuss in Section 5, within-quarter aggregation among them, so I read the data-driven estimates as an upper bound on the consumption cost.

Taken together, the results identify geoeconomic pressure as a distinct source of sovereign risk: a large economy can impose meaningful financial and real costs on a smaller country through the threat of policy, before those policies are implemented. To my knowledge, this paper is among the first to provide causal, high-frequency evidence that foreign geoeconomic pressure reprices the default risk of an emerging market economy and to trace that repricing through the exchange rate, equity prices, and aggregate consumption. More broadly, the findings show that sovereign risk need not originate in domestic fiscal or political conditions. In a policy environment in which major powers increasingly use access to trade, finance, and migration as instruments of coercion, targeted economies can bear sizable costs from pressure alone. The paper also shows how combining synthetic controls, high-frequency identification, and a structural sovereign-default model can measure those costs as they unfold.

Related literature. This paper is most closely related to the emerging literature on geoeconomics, which studies how states use economic linkages, trade, finance, and sanctions, to pursue political and strategic goals. Mohr and Trebesch (2025) survey the field and organize it around five themes: (i) international trade and politico-economic power, (ii) sanctions, (iii) domestic spillovers of geopolitical shocks, (iv) international finance and geopolitics, and (v) the economics of war and military spending. Clayton et al. (2023) and Clayton et al. (2025) formalize how a dominant economy exerts pressure on others through their reliance on its markets and financial system. Separately, and predating this field, a large literature studies how political frictions shape sovereign default, but almost entirely from the perspective of domestic instability (e.g., Chang (2007), Cuadra and Sapriza (2008), Yu (2016), Fourakis (2022)). This paper bridges the two: I show that foreign geoeconomic pressure, and not only domestic politics, moves sovereign default risk.

Second, the paper contributes to both the empirical and structural literature on sovereign default and its costs. This work has largely measured the economic cost of default, but because default and economic activity are jointly determined, most estimates are correlations rather than causal effects (e.g., Rose (2005), Borensztein and Panizza (2009), Yeyati and Panizza (2011), Martinez and Sandleris (2011), Zarazaga (2012), Tomz and Wright (2013), Cruces and Trebesch (2013)). The closest antecedents use high-frequency financial data to identify exogenous movements in sovereign risk: Hébert and Schreger (2017) use the 2014 U.S. court rulings on Argentina's 2001 default, together with high-frequency CDS data, to estimate the causal effect of default risk on Argentine equities and the exchange rate, and Sagawa (2026) identifies exogenous sovereign liquidity shocks from high-frequency movements in Argentine CDS spreads around IMF announcements. I follow this high-frequency logic but exploit a different source of variation, a geoeconomic repricing of Mexico's CDS. One difference is worth stating up front: Hébert and Schreger's ruling plausibly moves asset prices only through default risk, whereas geoeconomic pressure can also act directly through trade and cash-flow channels. I therefore read my estimand as the effect of a Mexico-specific repricing shock and, in Section 3.2, isolate the component that operates through the CDS. I extend this literature in two directions. First, I estimate the response of aggregate consumption to this type of foreign political pressure. Second, I connect the reduced-form evidence to the quantitative sovereign-default tradition (see Eaton and Gersovitz (1981), Bulow and Rogoff (1988), Aguiar and Gopinath (2006), Arellano (2008), Uribe and Schmitt-Grohé (2017), Aguiar and Amador (2021)) by embedding geoeconomic pressure directly in a default model, which prices the shock and quantifies its welfare cost.

Third, the paper relates to work on the global effects of President Trump's economic strategy, often called "Trumponomics." Most of this research studies the source economy, the United States. For example, Filippou et al. (2025) document that the dollar appreciates when Trump posts on social media, Caldara et al. (2020) estimate that higher trade-policy uncertainty cut U.S. investment by about 1%, and Amiti et al. (2021) put the welfare loss from the 2018 trade war at roughly 3% for the United States (see also Bown (2017)). I instead study a targeted economy and ask how an EME absorbs this foreign pressure.

Fourth, a handful of papers study Mexico directly. They mostly estimate how the peso and stock returns respond to Trump's threats, with mixed magnitudes and little on the mechanism. Beard et al. (2017) study how a higher probability of Trump winning the 2016 election moves the peso, the 5-year CDS, and stock returns; they find a sizable depreciation but small, short-lived effects on default risk and equities. Using different identification and a sample that spans Trump's first year in office, I find larger and more persistent effects on the peso and equities operating through default risk. Mao (2019) finds that the peso reacts strongly to migration-related tweets at first but that the response fades with repetition. In contrast, and in line with Matveev and Ruge-Murcia (2024), I find that trade threats, not migration threats, drive the response of financial markets.

Finally, the paper contributes to the applied synthetic-control literature. Since Abadie and Gardeazabal (2003), the method has been used to evaluate policies and shocks across many settings: Abadie et al. (2010) on California's tobacco-control program, Abadie et al. (2015) on German reunification, and Born et al. (2019) on the Brexit referendum. These applications almost always use low-frequency data; I contribute by applying the method to high-frequency financial data such as daily CDS spreads. This choice is supported by Goldsmith-Pinkham and Lyu (2025), who show that the factor-model approach standard in financial event studies is biased when events coincide with volatile markets or with market-correlated timing, and that synthetic-control counterfactuals, which match the pre-event path without imposing a factor structure, are the robust alternative. My setting, a politically turbulent episode in which event timing lines up with market conditions, is exactly the case they flag.

Outline. The paper proceeds as follows. Section 2 measures the geoeconomic repricing of Mexico's sovereign risk, describing the data and the synthetic-control method and presenting the estimated premium. Section 3 shows that the premium responds to President Trump's observable pressure and traces its transmission to the exchange rate and equity prices. Section 4 develops a structural model of sovereign default that rationalizes the increase in spreads and quantifies the associated welfare cost. Section 5 recovers the consumption response through a complementary, data-driven approach at high frequency. Section 6 concludes.

2 Measuring the geoeconomic repricing of Mexico's sovereign risk

2.1 Data description

I use daily data from Markit on the 5-year CDS for Mexico and 18 other economies with comparable S&P credit ratings as the outcome variable.[1] I use the 5-year maturity because it is the most liquid and most actively traded point on the sovereign CDS curve (see Longstaff et al. (2011)), which also maximizes data availability across the donor pool. The whole sample runs from January 2, 2013 to February 28, 2018, split into pre- and post-intervention periods. The treatment date is July 19, 2016, the day the Republican National Convention nominated Donald Trump as the party's presidential candidate, which I take as the onset of the geoeconomic pressure episode. I limit the post-intervention sample to avoid contami-

[1] By "comparable" I mean sovereigns whose S&P credit rating over the pre-treatment sample lay within 4 notches, above or below, of Mexico's pre-treatment rating of BBB+.

nation from Mexico's 2018 electoral process.[2] This yields a pre-intervention sample of 925 observations and a post-intervention sample of 422.

TABLE 1: Descriptive statistics of sovereign 5-year CDS (basis points)

Country	Rating		Mean		Std. Dev.		Min		Max	
	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Brazil	BB	BB-	239.97	227.56	111.40	43.96	102.07	142.63	534.56	320.72
Bulgaria	BB+	BBB-	143.06	117.13	25.38	22.29	91.47	61.65	207.87	156.35
Chile	AA-	A+	90.25	70.26	20.73	13.73	60.29	42.24	153.35	107.45
Colombia	BBB	BBB-	149.66	141.10	58.32	27.90	74.25	89.49	326.44	210.09
Indonesia	BB+	BBB-	188.24	126.28	42.10	26.47	123.97	76.92	285.57	198.85
Kazakhstan	BBB-	BBB-	217.05	155.33	58.26	31.04	127.62	91.4	379.6	230.04
Malaysia	A-	A-	124.37	100.60	38.39	30.12	72.73	50.84	242.36	172.11
Mexico	BBB+	BBB+	118.51	129.55	35.15	25.63	64.17	95.44	231.99	193.01
Panama	BBB	BBB	127.40	106.29	34.41	26.79	73.94	55.58	223.43	157.13
Peru	BBB+	BBB+	131.90	93.19	33.41	16.35	76.76	65.79	230.49	131.31
Philippines	BBB	BBB	103.39	83.71	14.15	19.77	78.51	51.87	154.25	134.81
Poland	BBB+	BBB+	76.16	66.36	11.86	10.95	53.18	48.72	132.29	89.76
Russia	BB+	BBB-	264.56	174.39	104.84	39.67	119.03	106.55	628.39	241.33
South Africa	BBB-	BB	221.65	204.92	53.81	33.82	135.79	140.74	391.65	270.34
South Korea	AA-	AA	63.24	52.83	11.75	9.48	45.49	39.33	108.14	75.49
Thailand	BBB+	BBB+	116.30	65.47	21.41	17.23	80.84	40.36	177.00	104.83
Trinidad and Tobago	A-	BBB+	133.03	116.83	14.36	17.72	105.1	72.89	323.12	300.43
Turkey	BB+	BB	212.24	221.60	47.25	41.13	110.95	154.92	326.92	306.33
Uruguay	BBB	BBB	220.40	164.20	86.14	17.77	117.59	93.4	543.46	199.35

Note: Daily data on sovereign CDS from Markit. The data set covers the period from January 2, 2013, to February 28, 2018. The sample is divided into a pre-treatment period (i.e. January 2, 2013, up to July 18, 2016) and a post-treatment period (i.e. July 19, 2016, up to February 28, 2018.). Ratings correspond to the last one observed in each sub-sample. Statistics are expressed in basis points.

Using credit ratings allows me to construct a donor pool of countries with macroeconomic conditions comparable to those of Mexico, which makes building a synthetic control easier (see Abadie and Gardeazabal (2003)). The S&P credit ratings range from BB- to AA in both the pre-treatment and post-treatment samples.[3] Analyzing the ratings of the donor-pool countries before and after the treatment, there are no large downgrades or upgrades, signaling a lack of major structural changes. This is essential, since a good comparison group is what allows the method to avoid extrapolation or bias. Table 1 displays descriptive statistics for the 5-year CDS spreads of Mexico and the donor pool.[4]

Regarding the predictors of the outcome variable (CDS), I use the average annual percentage change in the nominal exchange rate, the average log real GDP per capita, the average log

[2]According to the National Electoral Institute (*INE*, for its abbreviation in Spanish), the electoral process started with the pre-campaign from December 14, 2017, to February 11, 2018, and the campaign lasted from March 30, 2018, to June 27, 2018. The election day was July 1, 2018.

[3]Credit ratings by other credit rating agencies such as Moody's and Fitch are very similar to those of S&P.

[4]Since roughly 39% of the donor pool consists of oil-exporting countries, the larger pre-intervention mean and standard deviation mainly reflects the 2015–2016 collapse in oil prices, which raised the spreads of oil exporters over the pre-treatment sample (the WTI and Brent mix bottomed at $26.19 and $28.82 per barrel on February 11, 2016). Furthermore, the fact that the post-intervention mean and standard deviation is smaller for almost all countries, with the exception of Mexico and Turkey, could reflect an absence of large idiosyncratic shocks in the donor pool, which is relevant according to Abadie (2021) when building the comparison group. In particular, among Latin American countries, the smaller post-treatment first and second moments could signal no spillover effects and little evidence of attenuation bias in the estimated treatment effect.

real consumption per capita, and the log capital stock per capita for the years 2013 to 2015. Exchange-rate data come from Bloomberg, while the aggregate-demand components come from the Penn World Table. In addition to these predictors, I include the level of the 5-year CDS at the start of each year from 2013 to 2016 as matching variables, a common practice in synthetic-control studies to tighten the pre-treatment fit (e.g., Abadie et al. (2010)).

2.2 Empirical strategy: the synthetic control method

The challenge in measuring how geoeconomic pressure moves lenders' perceptions is that I do not observe the counterfactual, what Mexico's CDS would have done absent the pressure.[5] To construct it, I use the synthetic-control method of Abadie and Gardeazabal (2003) and Abadie et al. (2010). Let $\mathbb{Y}_1$ be a $(T \times 1)$ vector containing the 5-year CDS for Mexico over the full sample, and let $\mathbb{Y}_0$ be a $(T \times J)$ matrix with the 5-year CDS of the donor-pool countries. Let $\mathbb{X}_1$ be a $(K \times 1)$ vector of K predictors of Mexico's daily CDS. As predictors, I use pre-treatment observations of the Mexican CDS, the average log real GDP per capita, the average log real consumption per capita, the average log capital stock per capita, and the average annual percentage change in the peso-dollar exchange rate for 2013 to 2015. Let $\mathbb{X}_0$ be a $(K \times J)$ matrix with the same predictors for the J donor countries, $\mathbb{W}$ a $(J \times 1)$ vector of weights, and $\mathbb{V}$ a $(K \times K)$ diagonal matrix with non-negative entries.

The method finds the weights $\mathbb{W}^\star$ that best match the pre-treatment characteristics of the treated unit, solving:

$$\min_{\mathbb{W} \in \mathcal{W}} ||\mathbb{X}_1 - \mathbb{X}_0 \mathbb{W}||^2_{\mathbb{V}} = (\mathbb{X}_1 - \mathbb{X}_0 \mathbb{W})^T \, \mathbb{V} \, (\mathbb{X}_1 - \mathbb{X}_0 \mathbb{W}) \, , \tag{1}$$

where

$$\mathcal{W} = \left\{ (w_2, \ldots, w_{J+1}) \middle| w_j \geq 0 \;\text{ and }\; \sum_{j=2}^{J+1} w_j = 1 \right\},$$

is the set of non-negative weights that sum to one. The solution for the weight vector depends on the choice of the symmetric, positive semidefinite matrix $\mathbb{V}$. As in Abadie and Gardeazabal (2003), I choose $\mathbb{V}$ to minimize the mean squared prediction error of the outcome variable over the pre-intervention period:

$$\mathbb{V}^\star = \arg\min_{\mathbb{V} \in \mathcal{V}} \left(\mathbb{Z}_1 - \mathbb{Z}_0 \mathbb{W}^\star(\mathbb{V})\right)^T \left(\mathbb{Z}_1 - \mathbb{Z}_0 \mathbb{W}^\star(\mathbb{V})\right), \tag{2}$$

where $\mathbb{Z}_1$ is a $(T_0 \times 1)$ vector of pre-treatment Mexican CDS observations, $\mathbb{Z}_0$ a $(T_0 \times J)$ matrix of the donor pool's pre-intervention CDS, with $T_0 < T$, and $\mathcal{V}$ the set of non-negative diagonal matrices. The optimal weights are given by $\mathbb{W}^\star = \mathbb{W}^\star(\mathbb{V}^\star)$, the synthetic counterfactual is $\mathbb{Y}_0 \mathbb{W}^\star$, and the estimator of the treatment effect is:

$$\hat{\alpha}_t = Y_{1,t} - \sum_{j=2}^{J+1} w_j^\star Y_{j,t} \quad \text{for} \;\; t \in \{T_0 + 1, \ldots, T\}, \tag{3}$$

where $Y_{1,t}$ represents the Mexican 5-year CDS, $w_j^\star$ is country j weight and $Y_{j,t}$ is country j 5-year CDS. This treatment effect is what I call the repricing of sovereign risk due to the geoeconomic pressure.

[5] A very early application of this approach, along with other variants, was developed with Diego Cid, Fabrizio López-Gallo, and Alberto Romero, and described in my master's thesis at ITAM (see Lord-Medrano (2020)); the implementation in this paper is my own.

2.3 The geoeconomic repricing of Mexico's sovereign risk

Only three countries in the donor pool receive a positive weight in the synthetic control for the Mexican 5-year CDS: Panama ($w_2^\star = 0.3771$), Colombia ($w_3^\star = 0.3279$), and Poland ($w_4^\star = 0.2950$); the remaining 15 donors receive zero weight.[6] All three carry almost the same sovereign credit rating as Mexico (BBB+), so the synthetic Mexico shares macroeconomic fundamentals close to Mexico's over the pre-treatment sample. Intuitively, the weighted donors span the dimensions along which Mexico's sovereign risk is priced: all three are investment-grade emerging-market sovereigns whose spreads lie in Mexico's range and co-move with it through common global factors. Colombia and Panama supply the regional, commodity-and-trade-linked Latin American dimension, while Poland proxies the feature that most defines Mexico's external profile, a manufacturing exporter tightly integrated with a large advanced-economy market (the EU, as Mexico is with the United States), but from outside the reach of the pressure I study. Table 2 compares the pre-treatment characteristics of Mexico, the synthetic control, and the simple average of the 18 donor countries: the synthetic control matches Mexico's pre-treatment characteristics far more closely than the simple average does, so it provides a better control group than an arbitrary equal-weighting of the donor pool.

Figure 1A plots the observed 5-year CDS for Mexico against the synthetic counterfactual, and Figure 1B plots the dynamic treatment effect, $\hat{\alpha}_t$, the gap between the two series. The pre-intervention fit is close, with an average gap of -0.74 basis points. After the intervention the two series diverge sharply, with the largest gaps on the day of the first presidential debate (37 bp), the day Trump was declared the winner of the election (47 bp), and the day he took office (59 bp); the gap averages 24 basis points over the post-treatment sample.

TABLE 2: CDS predictor means

Variables	Mexico	Synthetic	Donor average
Nominal exchange rate depreciation in 2013	-2.91	0.52	2.59
Nominal exchange rate depreciation in 2014	4.35	2.22	7.54
Nominal exchange rate depreciation in 2015	19.31	17.49	17.92
Log GDP per capita	9.80	9.94	9.82
Log consumption per capita	9.59	9.69	9.54
Log capital stock per capita	11.26	10.95	11.01
CDS at the beginning of 2013	97.23	91.28	105.95
CDS at the beginning of 2014	92.32	104.55	145.86
CDS at the beginning of 2015	103.45	107.66	158.87
CDS at the beginning of 2016	169.55	166.88	209.90

Note: 'Mexico' is the observed Mexican series, 'Synthetic' is the synthetic control, and 'Donor average' is the simple mean across the 18 donor countries. The nominal exchange rate depreciation corresponds to the average depreciation observed that year. Log GDP per capita, log consumption per capita, and log capital stock per capita correspond to the average using data from 2013 to 2015.

Following Hébert and Schreger (2017) and assuming a constant hazard rate, I compute the one-year risk-neutral default probability as

$$\mathbb{P}_t^{\mathbb{Q}}\left(Default \leq 1 \text{ year}\right) = 1 - \exp\left(-\frac{CDS_t}{1-R}\right), \tag{4}$$

[6] That only a few donors receive positive weights is a well-known feature of the method, which keeps the counterfactual transparent and easy to interpret.

where CDS_t is the observed CDS and R is the recovery rate. I set $R = 0.395$, the estimate from Benjamin and Wright (2009) for Mexico's 1982 default episode. Figure 2A plots the one-year risk-neutral default probability and Figure 2B plots the estimated wedge under alternative recovery rates from the literature. Under this approximation, the geoeconomic pressure raised the one-year risk-neutral default probability from 1.67% to 2.02% on average, and the estimate is robust to the choice of R. Because the default probability is a monotone transformation of the CDS, Figure 2A has the same shape as Figure 1A and differs only in the vertical scale.

FIGURE 1. Observed CDS, synthetic control and treatment effect

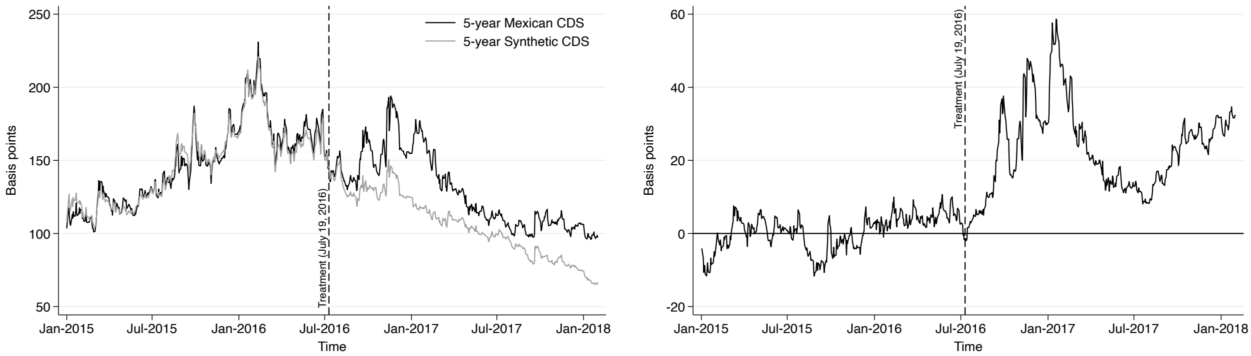

(A) Observed and synthetic CDS

(B) Premium $\hat{\alpha}_t$

Note: For the estimation, I use daily data on sovereign 5-year CDS for Mexico and 18 other emerging market economies with comparable sovereign credit ratings. The estimated weights are: Panama ($w_2^\star = 0.3771$), Colombia ($w_3^\star = 0.3279$), and Poland ($w_4^\star = 0.2950$), respectively. Panel (b) plots the estimated premium, $\hat{\alpha}_t$, the observed-minus-synthetic gap defined in equation (3).

FIGURE 2. Observed and synthetic one-year risk-neutral default probabilities

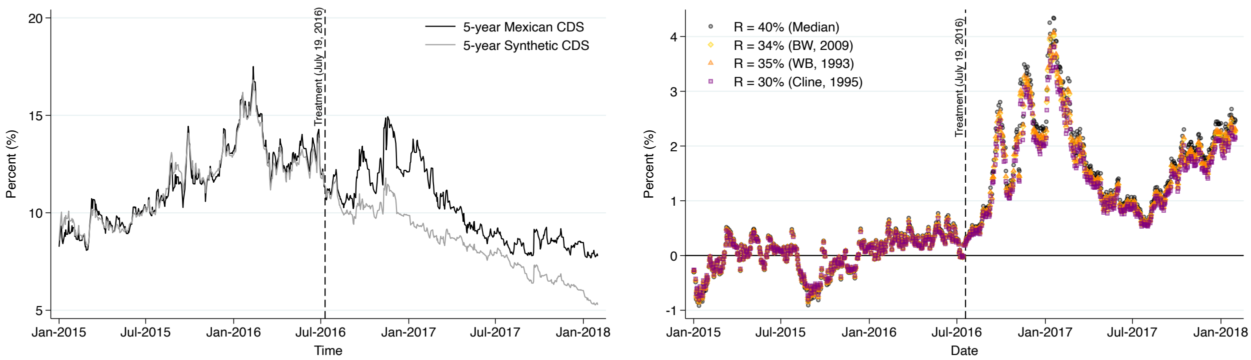

(A) Observed and synthetic probabilities

(B) Premium for different R

Note: On the left, Figure 2A displays the implied one-year risk-neutral default probabilities from the observed and synthetic CDS, respectively. To compute the implied one-year risk-neutral default probabilities, I follow Hébert and Schreger (2017) assuming a constant hazard rate and a recovery rate of 39.5%. On the right, Figure 2B plots the treatment effect in one-year risk-neutral default-probability units (percentage points) for different values of the recovery rate reported in Benjamin and Wright (2009); the baseline $R = 0.395$ curve is the probability-scale counterpart of the premium in Figure 1B.

I read $\hat{\alpha}_t$ as the component of Mexico's default risk attributable to the geoeconomic pressure: the wedge between Mexico's CDS and that of comparable sovereigns that opens once the pressure begins, which the synthetic control isolates from the common global and regional factors that move these sovereigns together. Two features of the post-treatment period bear on this reading. Banco de México (2017) judged Mexico's fundamentals strong through mid-2017, aided by structural reforms, notably the energy reform, that were drawing in foreign direct investment, conditions that if anything point to lower, not higher, spreads. Over these quarters, the only material risk it flagged to the outlook, beyond ordinary structural factors, was the effect of the U.S. administration's threats on trade and remittances. From the last two quarters of 2017 into early 2018, the central bank began flagging additional risks alongside the NAFTA renegotiation: the upcoming 2018 election, rising U.S. interest rates, and the September 19, 2017 earthquake. I therefore take the full post-treatment sample as the baseline but flag the post-July-2017 window (shaded in Figure 3) as potentially contaminated. I keep it because it contains the rounds of NAFTA renegotiation, which are central to the trade-related component of the pressure that I isolate below. Reassuringly, the effect is already present in the early post-treatment window and builds gradually (the backdating and block-placebo exercises in Appendix A), so it does not hinge on the potentially contaminated tail.

FIGURE 3. Potential contamination

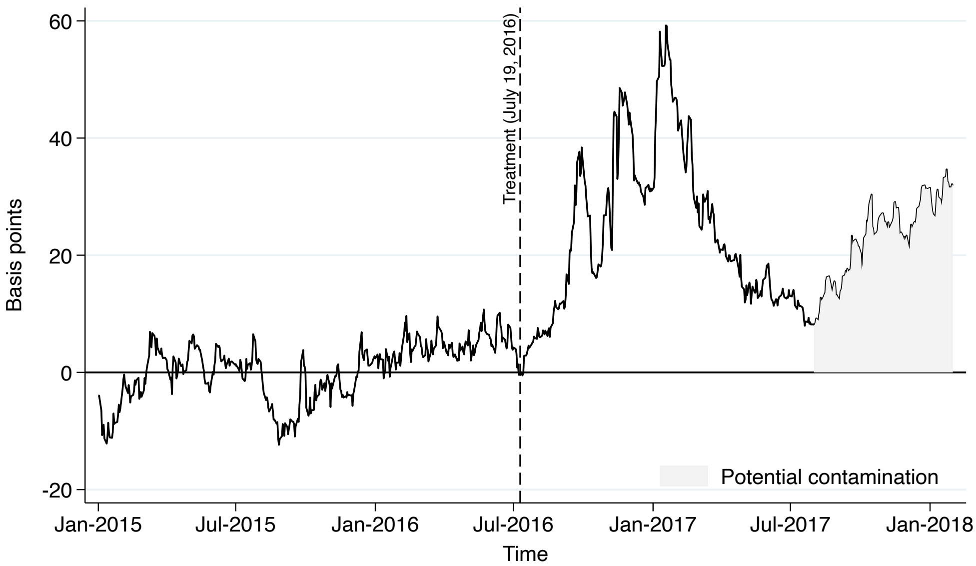

Note: In this Figure, I display the estimated treatment effect (i.e., the difference between the observed 5-year CDS and the 5-year synthetic CDS) where the gray area corresponds to the period in which there might be some contamination given the presence of other "shocks" different from Trump.

I subject the baseline premium to a battery of validation and robustness checks, following the requirements for synthetic-control designs in Abadie (2021) and the diagnostics in Abadie and Vives-i-Bastida (2022). Table 3 summarizes them; the details, figures, and trimmed-donor-pool estimates are in Appendix A. The premium survives trimming the donor pool to Mexico's rating class, leave-one-out donor removal, smoothing the volatile CDS series, and residualizing out U.S. monetary-policy changes; a permutation-based test ranks Mexico's post-to-pre-treatment fit ratio first among all 19 units, the most extreme outcome attainable with this donor pool ($p = 1/19 \approx 0.053$); prediction intervals put the observed CDS above the synthetic 95% band; backdating confirms the effect emerges only after the nomination; and a block-placebo procedure shows that the wedge is not an artifact of serial correlation of CDS data. Excluding donors most exposed to the same shocks (Latin America, high U.S.-trade exposure, and oil exporters) raises the estimated average premium to 27–33 basis points, so the 24 bp baseline is, if anything, a lower bound.

TABLE 3: Validation and robustness of the estimated premium

Exercise	Result
Baseline (18-country donor pool)	Average gap 24 bp; peaks 37–59 bp
Trimmed donor pool (BBB)	Weights and gap almost identical
Leave-one-out	All estimates positive, centered on baseline
Exclude Latin America	Average gap of 27.44 bp
Exclude high U.S. trade exposure	Average gap of 33 bp
Exclude oil exporters	Average gap of 28 bp
Smoothing (30/60/90-day MA)	Gap unchanged (no overfitting)
Placebo permutation	Mexico's gap most extreme; $p \approx 5.3\%$
Prediction intervals	Observed CDS above the 95% interval
Backdating (1–6 months)	Effect emerges only after the nomination
Block placebo (20–120 days)	Effect in the right tail at all horizons
Fed-funds residualization	Gap unchanged

Note: Each row summarizes a validation or robustness exercise for the estimated premium $\hat{\alpha}_t$. Full details, figures, and the trimmed-donor-pool predictor means are in Appendix A.

3 Geoeconomic pressure and its transmission to asset prices

3.1 Does the premium respond to observable pressure?

So far I have interpreted the estimated gap, $\hat{\alpha}_t$, as a repricing of Mexican default risk driven by coercive pressure from the United States. This section asks whether the gap behaves as that interpretation requires. If the premium reflects pressure on Mexico, it should respond to observable realizations of that pressure, rising when the U.S. president threatens the Mexican economy and easing when he signals willingness to preserve the trade agreement, and it should react to threats aimed at Mexico rather than to unrelated foreign-policy disputes. I test both properties, sensitivity and targeting, using President Trump's activity on 𝕏, formerly Twitter, a high-frequency and precisely timestamped record of his public statements. According to Factba.se (2023), between 2016 and 2018 he posted roughly 10,500 times, of which 38 referred directly to Mexico, mostly threats tied to migration and to the trade agreement. I then ask which dimension of those threats, trade or migration, matters most to foreign lenders.

Two caveats frame the exercise. First, the tweets are one observable signal of a broader pressure process that also travels through institutional channels, such as White House statements and party platforms, and through the market's anticipation that pressure will arrive even on days when the president is silent. A premium that moves with the visible signal corroborates the pressure interpretation, but because much of the pressure is anticipated or communicated through other channels, the tweet-day response captures only the news revealed on those days, the salient increment of the pressure, and therefore understates its full contribution to the gap. The exercise validates the reading of $\hat{\alpha}_t$; it does not claim that social media alone generates it. Second, the number of Mexico-related tweets is small, so the estimates below are statistically underpowered and I read them as suggestive rather than as precise causal effects.

To construct the dataset on Trump tweets, I use the information available at Factba.se (2023) for the period corresponding to the post-treatment sample (i.e., July 2016 to February 2018). The main benefit of using this source is that it contains detailed information about the tweets, such as the content, timestamp, and author. According to the source, of those 38 tweets, 26 are attributed to Trump himself and the remaining 12 to his staff members. Based on the timestamp and the trading hours of the NYSE, I assign each tweet to a given day. For example, if a tweet was made during the weekend, I record it as the next trading day and, if it was made after 4:00 PM ET on any day of the week, I also record it as the next trading day. Given the content of the tweets, I classify them as "positive", "negative", or "neutral". They are classified as negative if there is any explicit threat against the Mexican economy, positive if there is any signal of willingness to cooperate or continue with the trade agreement, and neutral if there is no clear threat or desire to deepen commercial ties. Here are some particular examples of such tweets and their classification:

> **Positive:** *I received calls from the President of Mexico and the Prime Minister of Canada asking to renegotiate NAFTA rather than terminate. I agreed*
>
> @realDonaldTrump (04/27/17 7:12 AM)
>
> **Negative:** *Toyota Motor said will build a new plant in Baja, Mexico, to build Corolla cars for U.S. NO WAY! Build plant in U.S. or pay big border tax*
>
> @realDonaldTrump (01/05/17 6:14 PM)
>
> **Neutral:** *Former President Vicente Fox, who is railing against my visit to Mexico today, also invited me when he apologized for using the "f-bomb"*
>
> @realDonaldTrump (08/31/16 1:07 PM)

In Table B1, I list the 26 tweets attributed to Trump with their date, timestamp, and classification. Because several fall on the same day, the estimation sample contains 20 distinct trading days with a Mexico-related tweet, 15 "negative", 2 "positive", and 3 "neutral".

To measure the impact of these announcements on gap $\hat{\alpha}_t$, I conduct two exercises. In the first, I estimate the effect of positive, negative, and neutral tweets on the premium through the following linear equation:

$$\hat{\alpha}_t = \iota + \delta t + \beta \mathbb{1}\{\text{Bad Tweet}_t\} + \rho \mathbb{1}\{\text{Good Tweet}_t\} + \psi \mathbb{1}\{\text{Neutral Tweet}_t\} + \gamma' \boldsymbol{x}_t + \eta_t, \quad (5)$$

where $\hat{\alpha}_t$ is the difference between the observed and synthetic CDS, t accounts for linear trends, $\mathbb{1}\{\text{Good Tweet}_t\}$ is an indicator function that takes the value of 1 whenever Trump tweeted anything "good" for Mexico, $\mathbb{1}\{\text{Bad Tweet}_t\}$ is an indicator function that takes the value of 1 whenever Trump made a threat to Mexico, $\mathbb{1}\{\text{Neutral Tweet}_t\}$ accounts for days in which the President posted tweets classified as "neutral", $\boldsymbol{x}_t$ is a vector of controls that includes the VIX index, oil prices, changes in the Fed funds rate, and lagged changes in domestic monetary policy, and η_t is the error term.

In the second, I decompose the "negative" tweets by their content. In particular, I classify them based on the content of the threat into two topics, "trade" and "migration", and estimate the following linear model:

$$\hat{\alpha}_t = \iota + \delta t + \lambda \mathbb{1}\{\text{Trade}_t\} + \theta \mathbb{1}\{\text{Migration}_t\} + \boldsymbol{\phi}' \tilde{x}_t + \epsilon_t, \quad (6)$$

where $\hat{\alpha}_t$ is the difference between the observed and synthetic CDS, t accounts for linear trends, $\mathbb{1}\{\text{Trade}_t\}$ is an indicator function that takes the value of 1 if Trump made a trade-related threat against Mexico, $\mathbb{1}\{\text{Migration}_t\}$ is an indicator function that takes the value of 1 if Trump made a migration-related threat to Mexico in day t, $\tilde{x}_t$ is a vector of controls that accounts for the other tweets classified as "good" and "neutral", global factors, and lagged domestic monetary policy, and ϵ_t is the error term. In both specifications, to account for autocorrelation and heteroskedasticity, I compute HAC robust standard errors.

FIGURE 4. Effect of Trump announcements via Twitter

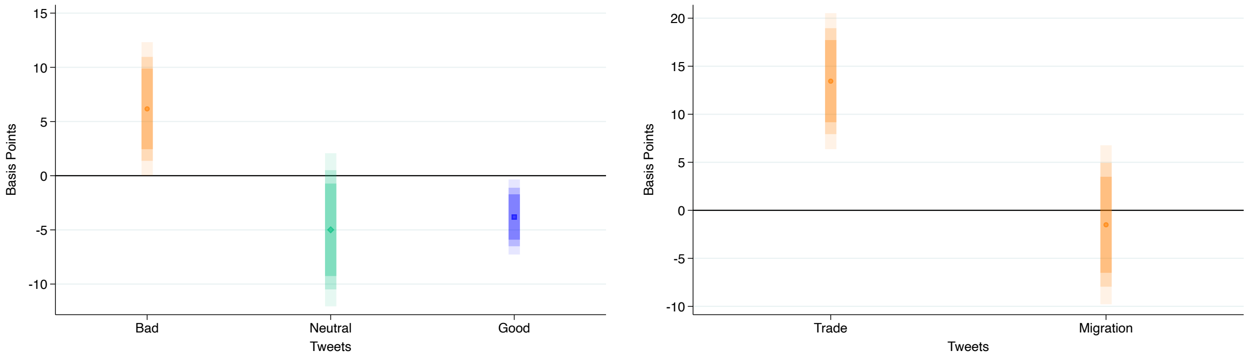

(A) Average effect

(B) By topic

Note: On the left, Figure 4A displays the point estimates for β, ρ, and ψ with the respective CI at 90%, 80%, and 68%. On the right, 4B displays the point estimates of "bad" tweets when they are separated by topic (i.e., "trade" λ and "migration" θ) with the corresponding CI at 90%, 80%, and 68%, respectively. I use HAC robust standard errors for both econometric specifications.

For the first specification, I estimate an average increase in the risk premium of 6.17 basis points in days when Trump threatened the Mexican economy, an average decrease of 3.81 basis points whenever he showed good intentions, both significant at the 10% level, and no statistically significant response in days with tweets that convey neutral information (see Figure 4A). From the second specification, tweets that contained trade-related threats (e.g., pulling out of NAFTA or imposing large trade tariffs) increased the premium by 13.44 basis points on average, while tweets that contained migration-related threats such as funding the border wall had no statistically significant effect.[7] That the response loads on the trade dimension rather than migration is consistent with the interpretation that lenders price the coercion most likely to reach Mexico's external accounts, and with the broader evidence on the costs of trade-policy uncertainty during this period (see Handley and Limão (2017)).

3.1.1 Alternative classification method

A natural concern is that the results reflect my own subjective coding. To address it, I reclassify the tweets with a large language model, ChatGPT-4o, and compare its labels to mine.[8] I query the model at temperature zero with a fixed prompt, so that sampling is

[7]Table B5 in Appendix B displays the point estimates for all linear specifications. Figure A1 displays the distribution of President Trump's tweets over time.

[8]Querying the model at temperature zero disables sampling, so the same tweet and prompt return the same label on repeated runs. This makes the classification replicable, though not mechanical: the model generates the labels rather than applying hand-coded rules, which is precisely why it serves as an independent check

disabled and the labeling is reproducible across runs, and I supply the tweets with the following instruction:

```
Prompt: I am going to give you some ``Tweets'' made by President Trump
regarding the Mexican economy during 2016-2018. Based on the content of
such tweets and its potential impact for the Mexican economy, classify
them as ``positive'', ``negative'', and ``neutral.'' Once you finish,
please provide the definitions that you used to make the classification.
```

The model returned the following definitions for the three categories:

```
Positive: Tweets that express goodwill, praise, or acknowledge mutual
benefits for Mexico, its people, or its economy, potentially fostering a
constructive relationship.
Negative: Tweets that criticize, blame, or undermine Mexico's economy,
trade policies, or societal issues, or that propose actions harmful to
its interests.
Neutral: Tweets that state facts, describe events, or focus on U.S.
actions without directly praising or criticizing Mexico or its economy.
```

Furthermore, I provide the following prompt to request that the model classify the "negative" posts into "trade" and "migration" related threats:

```
Prompt: From the tweets that you classified as ``negative'', now
classify them as ``trade related'' and ``migration related'' based on
their content.
```

With the new classification, I re-estimate equations (5) and (6). The estimates are displayed in Figure 5. The results are almost identical to those of the baseline classification, the only difference being in the estimated effect of "positive" tweets by President Trump. This is explained by the fact that one of the tweets that I classified as "neutral" is classified by the model as "positive".[9]

3.1.2 Placebos

This brings me to the targeting property. If the premium reflects pressure on Mexico specifically, threats the president directed at unrelated economies should leave it unmoved. Mexico's largest trading rival, China, is the natural exception, since U.S. pressure on China can divert trade and investment toward Mexico. I test this by estimating the effect of Trump's tweets targeting China, North Korea, and Russia over the same period, of which there are 33, 28, and 61, respectively. Some examples are displayed below:

> **China:** *The meeting next week with China will be a very difficult one in that we can no longer have massive trade deficits...*

on my coding. Large language models have been used in the recent geoeconomics literature to classify policy statements at scale (see Clayton et al. (2025)).

[9]The tweet is: "God bless the people of Mexico City. We are with you and will be there for you." This is in the context of Mexico City experiencing a 7.1 magnitude earthquake on September 19, 2017.

FIGURE 5. Effect of Trump announcements via Twitter using AI classification

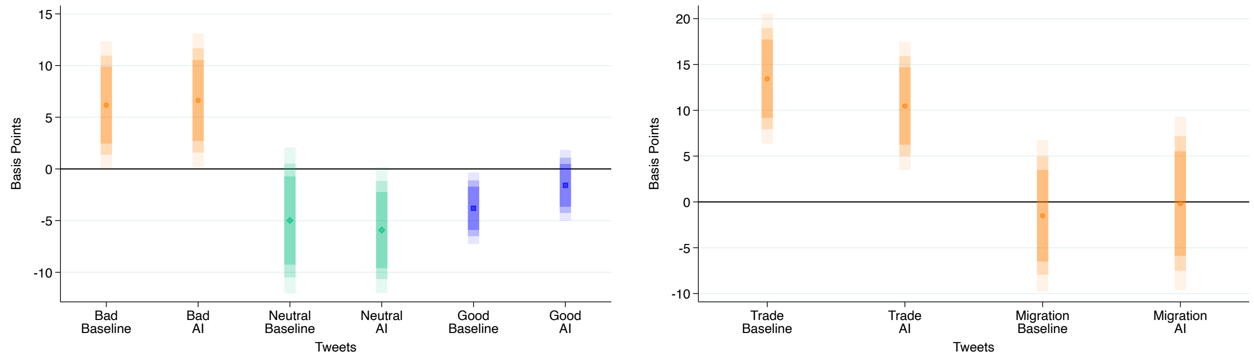

(A) Average effect

(B) By topic

Note: On the left, Figure 5A displays the point estimates with the respective CI at 90%, 80%, and 68%. On the right, 5B displays the point estimates of "bad" tweets when they are separated by topic (i.e., "trade" and "migration") with the corresponding CI at 90%, 80%, and 68%, respectively, with HAC robust standard errors.

@realDonaldTrump (03/30/17 6:16 PM)

> **North Korea:** *Military solutions are now fully in place,locked and loaded,should North Korea act unwisely. Hopefully Kim Jong Un will find another path!*

@realDonaldTrump (08/11/17 7:29 AM)

> **Russia:** *Having a good relationship with Russia is a good thing, not a bad thing. Only "stupid" people, or fools, would think that it is bad! We have enough problems around the world without yet another one. When I am President, Russia will respect us far more than they do now and...*

@realDonaldTrump (01/07/17 10:02 AM)

For the detailed list, see Appendix B. To estimate the effect of these comments on the premium, I estimate the following linear model:

$$\hat{\alpha}_t = \iota + \delta t + \beta_0 \mathbb{1}\{\text{China}_t\} + \beta_1 \mathbb{1}\{\text{N.Korea}_t\} + \beta_2 \mathbb{1}\{\text{Russia}_t\} + \boldsymbol{\xi}' z_t + \varepsilon_t, \tag{7}$$

where $\hat{\alpha}_t$ represents the difference between the observed and synthetic Mexican 5-year CDS, $\mathbb{1}\{\text{Country}_t\}$ is an indicator function that takes the value of 1 if in day t there was a tweet directed to "Country", z_t is a vector that includes the tweets about Mexico and other continuous controls such as the VIX index, oil prices, changes in the Fed funds rate, and lagged changes in domestic monetary policy, and ε_t represents the error term.[10]

Figure 6 reports the estimates. Threats aimed at North Korea or Russia have no statistically significant effect on $\hat{\alpha}_t$, so the premium does not simply track the president's foreign-policy rhetoric in general. The only exception is China, and its effect is negative and statistically significant, the opposite of what a threat to Mexico would produce. Rather than signaling contamination, this is the exception the targeting logic anticipates: pressure on China plausibly benefits Mexico by diverting global value chains and investment toward it, consistent with the recent evidence on nearshoring in the wake of the U.S.-China trade war (e.g., Utar

[10]In Table B6, the point estimates for linear model (7) are displayed.

FIGURE 6. Estimated effect of tweets directed to other economies

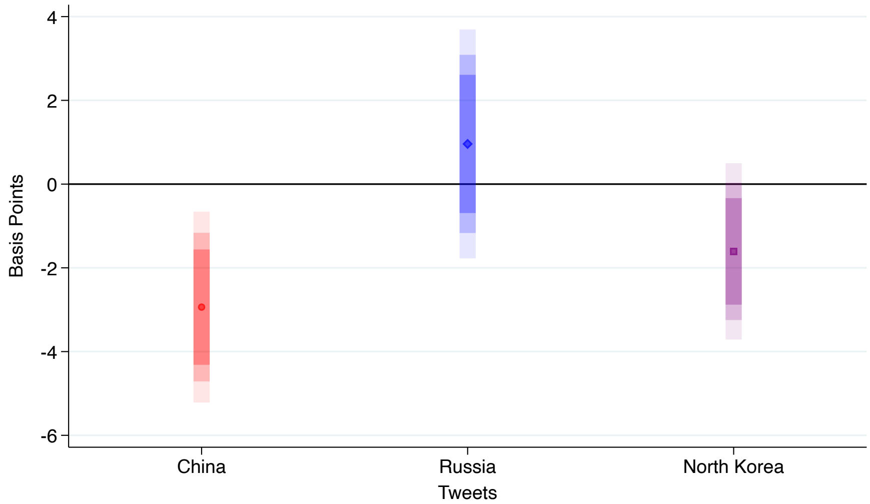

Note: Shaded areas represent confidence intervals at 90%, 80%, and 68%, respectively. Standard errors are robust to heteroskedasticity and autocorrelation.

et al. (2023) and Gopinath et al. (2025)). As Chinese firms face higher or more uncertain tariffs, some relocate production to tariff-free trade partners such as Mexico, and the associated foreign direct investment signals stronger fundamentals and a lower spread.

3.2 Transmission to asset prices

Having established that coercive pressure repriced Mexico's default risk, I now quantify how that repricing transmitted to the USD/MXN nominal exchange rate (pesos per U.S. dollar) and to the log-returns of the Mexican stock index (*IPC*, for its abbreviation in Spanish). My main estimator uses local projections to trace the dynamic response of each asset price to the estimated premium $\hat{\alpha}_t$. I then subject these responses to a battery of robustness checks, the most demanding of which discards the synthetic-control premium entirely and re-identifies the effect with high-frequency instruments around "Trump days". That both approaches deliver the same magnitudes is the central validation of this section.

3.2.1 Local projections

I use local projections à-la Jordà (2005) to estimate the response functions of the exchange rate and the daily *IPC* index to an increase in the estimated premium. The linear specification is given by:

$$y_{t+h} - y_{t-1} = \gamma^h + \hat{\alpha}_t \beta^h + \boldsymbol{x}_t \psi^h + u_{t+h}, \quad \text{with } h = 0, \ldots, 20, \tag{8}$$

where y_{t+h} is either the exchange rate or the daily log-*IPC* stock index, $\hat{\alpha}_t$ is the difference between the observed CDS and the synthetic CDS, $\boldsymbol{x}_t$ contains controls such as lagged observations for both the outcome variable and $\hat{\alpha}_t$, lagged daily changes in the domestic monetary policy rate, and global volatility controls such as the VIX index, oil prices, and changes in the Fed fund rates, and u_{t+h} is the error term. Because $\hat{\alpha}_t$ is highly persistent, I include lags of the premium to absorb autocorrelation, as suggested by Ramey (2016).[11] The cumulative impulse-response function corresponds to the estimates of the dynamic multiplier β^h for $h = 0, \ldots, 20$. The identifying assumption is that, conditional on these controls, $\hat{\alpha}_t$ is predetermined with respect to the contemporaneous outcomes. This is not innocuous, since a trade threat could in principle reach the peso and equities through channels other than sovereign risk; I probe both the predeterminedness of $\hat{\alpha}_t$ and this exclusion restriction in the robustness analysis below.[12]

FIGURE 7. Cumulative impulse-response functions

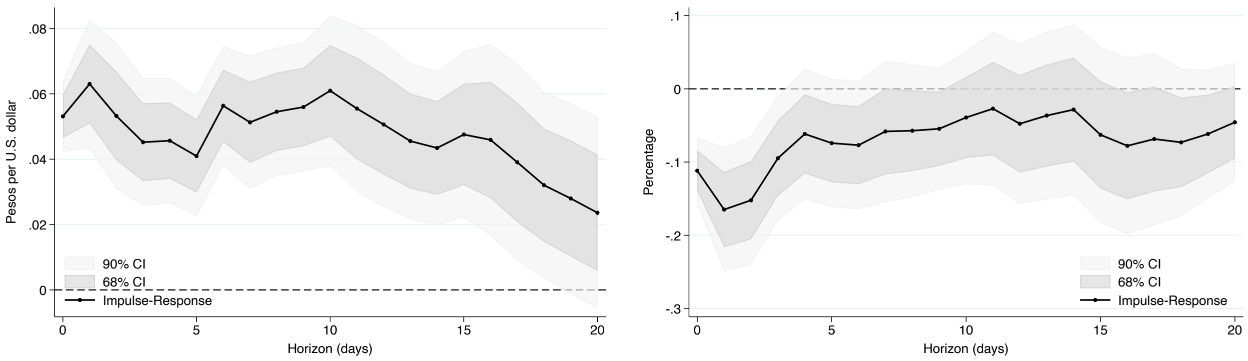

(A) Exchange rate

(B) Stock returns

Note: Figure 7A displays the cumulative impulse-response function of the first difference in the exchange rate to an increase of 1 basis point in the premium while Figure 7B presents the cumulative impulse-response function of the daily log-return of the Mexican stock index to an increase of 1 basis point in $\hat{\alpha}_t$. The shaded areas correspond to the confidence intervals at 68% and 90%, respectively, computed using HAC robust standard errors.

Figure 7 reports the responses. On impact, the peso depreciates by 5.3 cents per dollar, approximately 0.26 percent, and the stock index falls by 0.11% in response to a 1 basis point increase in $\hat{\alpha}_t$.[13] The effect on the peso takes about 17 days to dissipate, while the effect on equities fades within 5 days. These estimates imply that, for an increase in $\hat{\alpha}_t$ of a size comparable to the one observed when Trump was declared the winner of the U.S. Presidential Election on November 9th, 2016, the exchange rate depreciates by \$1.55 pesos per U.S. dollar and the Mexican stock index falls by 3.41%, approximately.To provide some perspective, the USD/MXN exchange rate depreciated by \$0.92 pesos less than a month after the collapse of Lehman Brothers and by \$1.12 pesos a week after the Covid-19 outbreak was declared a pandemic by the World Health Organization (WHO). For the same events, stock returns sank by 7.27% and 6.64%, respectively.

[11] This is isomorphic to estimating the local projections using the residuals of a first-stage in which one regresses the variables on their own lags (see Miranda-Agrippino and Ricco (2021)).

[12] Because $\hat{\alpha}_t$ is a linear function of the 5-year Mexican CDS, endogeneity is a potential concern; as discussed in the previous section, the dynamic treatment effect captures the exogenous variation in the CDS generated by the foreign political pressure exerted by President Trump. For more details, see Appendix C.

[13] Following Montiel Olea and Plagborg-Møller (2021), I also compute robust to heteroskedasticity standard errors using the Huber-White estimator; the results are almost identical to those estimated with HAC robust standard errors.

3.2.2 High-frequency identification

The most demanding checks discard the synthetic-control premium altogether. One might worry that $\hat{\alpha}_t$ reflects not only U.S. pressure but also idiosyncratic factors, especially late in the sample, such as the onset of Mexico's 2018 election. I therefore re-identify the asset-price response without using $\hat{\alpha}_t$ at all, using two high-frequency strategies that exploit the same "Trump events" through different identifying variation: a high-frequency external instrument built from the event-window CDS surprise, in the spirit of the monetary-policy and oil-market literatures (Nakamura and Steinsson (2018), Miranda-Agrippino and Ricco (2021), Känzig (2021)), and identification through the heteroskedasticity of the CDS around those events (Rigobon (2003), Rigobon and Sack (2004), Hébert and Schreger (2017)), using daily information from July 2016 to February 2018. These "Trump events" correspond to days in which Donald Trump tweeted something about Mexico, the three presidential debates, the vice-presidential debate, the election day and electoral result, and the rounds of renegotiation of the Free Trade Agreement. For the classification of dates regarding tweets, I follow the same procedure outlined in Section 3.1. To construct the data set, I compute the first difference of the Mexican 5-year CDS, the exchange rate, and the logarithm of the *IPC*, and classify each trading day as an event or a non-event day, giving 41 event days and 376 non-event days (417 observations in total). The high-frequency-instrument, heteroskedasticity, and OLS estimators use the full sample; the event-study estimator uses only the 41 event days.

Consider the simultaneous system:

$$\Delta y_t = \beta \Delta CDS_t + \theta_1 X_{1,t} + \theta_2 X_{2,t} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), \tag{9}$$

$$\Delta CDS_t = \varphi \Delta y_t + \delta_1 X_{1,t} + \delta_2 X_{2,t} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{10}$$

where Δy_t is the first difference in the USD/MXN exchange rate or the log of the *IPC* index, ΔCDS_t is the first difference of the 5-year Mexican CDS, and $X_{1,t}$ and $X_{2,t}$ are the VIX index and the oil price (WTI). The key identifying assumption is an exclusion restriction: U.S. pressure moves the outcome only through the CDS, so that Trump events act as idiosyncratic CDS shocks η_t rather than as direct shocks to the peso or equities. This is a strong assumption, and it is the same one that underlies the causal reading of $\hat{\alpha}_t$ throughout the paper; the high-frequency strategies make it explicit and testable.

The first strategy uses the event-window CDS surprise as an external instrument and recovers the effect of the CDS on the outcome as

$$\hat{\beta}^{HF} = \left(z'\Delta CDS\right)^{-1} \left(z'\Delta y\right), \text{ where } z_t = \mathbb{1}\{t \in E\} \times \Delta CDS_t, \tag{11}$$

with E the set of events, and ΔCDS and Δy vectors that combine the event and non-event observations. The second implements identification through heteroskedasticity in the sense of Rigobon (2003), using the instrument

$$z'_t = \frac{|E| + |NE|}{|E|} \mathbb{1}\{t \in E\} \times \Delta CDS_t - \frac{|E| + |NE|}{|NE|} \mathbb{1}\{t \in NE\} \times \Delta CDS_t, \tag{12}$$

where E and NE denote the event and non-event days;[14] it exploits the rise in the variance of the CDS shock on event days relative to non-event days (Hébert and Schreger (2017)), which

[14]The regime weights in z'_t are evaluated on a balanced partition ($|E| = |NE| = 41$), while the estimator uses the full sample.

I verify with tests of equal variances reported in Appendix C. In the same spirit as Cochrane and Piazzesi (2002), restricting the sample to event days I also compute the event-study estimator:

$$\hat{\beta}^{ES} = \left(\Delta CDS_E' \Delta CDS_E\right)^{-1} \left(\Delta CDS_E' \Delta y_E\right).$$

TABLE 4: Estimates of the effect of the CDS on outcome

Outcome variable	HF (1)	HET (2)	ES (3)	OLS (4)	LP (5)
First difference in Mexican peso	0.05***	0.06**	0.05***	0.03***	0.05***
	(0.02)	(0.03)	(0.02)	(0.004)	(0.01)
Daily log-return of stock exchange	-0.12***	-0.14***	-0.11***	-0.08***	-0.11***
	(0.02)	(0.04)	(0.02)	(0.02)	(0.03)
Sector 2: Materials	-0.05	-0.06	-0.05	-0.05**	-0.05
	(0.08)	(0.12)	(0.07)	(0.02)	(0.04)
Sector 3: Industrial	-0.16***	-0.19***	-0.15***	-0.10***	-0.15***
	(0.04)	(0.06)	(0.04)	(0.02)	(0.03)
Sector 4: Non-basic G&S	-0.11**	-0.09	-0.10**	-0.09***	-0.12***
	(0.05)	(0.09)	(0.05)	(0.02)	(0.03)
Sector 5: Frequent consumption	-0.07***	-0.08*	-0.07***	-0.06***	-0.08***
	(0.02)	(0.04)	(0.02)	(0.02)	(0.03)
Sector 6: Health	-0.15***	-0.15**	-0.12***	-0.05***	-0.09***
	(0.06)	(0.06)	(0.04)	(0.01)	(0.02)
Sector 7: Financial services	-0.19***	-0.21***	-0.19***	-0.11***	-0.17***
	(0.05)	(0.08)	(0.05)	(0.02)	(0.04)
Sector 9: Telecommunications	-0.10**	-0.11	-0.08**	-0.04**	-0.04
	(0.04)	(0.07)	(0.03)	(0.02)	(0.03)
Obs.	417	417	41	417	417
F-stat	61.090	15.884			
F-stat (robust)	1,028.316	4.899			

Note: (1) is the high-frequency external-instrument estimate (event-window CDS surprise, in the spirit of Nakamura and Steinsson (2018), Miranda-Agrippino and Ricco (2021), and Känzig (2021)); (2) is the heteroskedasticity-based estimate following Rigobon (2003), Rigobon and Sack (2004), and Hébert and Schreger (2017); (3) is the event-study estimate; (4) is the OLS estimate over the full sample; and (5) is the contemporaneous ($h = 0$) local-projection response from the previous subsection. For stock returns, I also display the estimated effect by sector using the S&P (2023) Sectoral Analytical Indices. Each specification uses the daily VIX index and the WTI oil price as controls. ***, **, * represent statistical significance at 1%, 5%, and 10%, respectively. Robust standard errors are reported in parenthesis. *F*-stat corresponds to the Cragg and Donald (1993) *F*-statistic and *F*-stat (robust) to the Kleibergen and Paap (2006) *rk F*-statistic allowing for heteroskedasticity and autocorrelation.

Table 4 reports the estimates from the high-frequency instrument (HF), the heteroskedasticity-based instrument (HET), the event study (ES), OLS, and, for comparison, the contemporaneous local-projection response that uses the $\hat{\alpha}_t$ premium (LP). A 1 basis point increase in the CDS translates into roughly a 6-cent depreciation of the peso and a 0.12–0.14% decline in equities, with the sharpest sectoral effects in financial services, industrials, and health. Two points stand out. First, OLS is attenuated by endogeneity. Second, and more important, the strategies that discard the synthetic-control premium deliver estimates essentially identical to the local-projection response built from the premium; because they share no common construction, their agreement is strong evidence that $\hat{\alpha}_t$ captures the geoeconomic pressure rather than an artifact of the synthetic control. The high-frequency instrument has a strong first stage F-stat (Montiel Olea and Pflueger (2013) effective F-statistic of 1,028), whereas the heteroskedasticity-based instrument, which relies on the more demanding change in second moments, has a weaker first stage F-stat (Montiel Olea and Pflueger (2013) effective F-statistic of 4.899).

Extending the approach to local projections with the high-frequency instrument yields dynamic responses that closely track the baseline (Appendix C): a persistent depreciation of the peso and a decline in equities following an increase in sovereign risk. Taken together, the two strategies capture a common Mexico-specific geoeconomic repricing shock, which reinforces that the baseline results are not driven by the construction of the synthetic-control premium.

Robustness and validation I subject the baseline responses to a battery of checks, all of which leave the estimates in Figure 7 essentially unchanged; Table 5 summarizes them and Appendix C reports the full results. First, because $\hat{\alpha}_t$ is a function of the CDS, which itself responds to the peso and equity prices, I ask whether the premium is predetermined with respect to the outcomes. Granger-causality tests in a VAR show that the premium forecasts the exchange rate and stock returns while the outcomes do not forecast the premium, consistent with weak exogeneity, though this is a statement about predeterminedness and not a guarantee of contemporaneous exogeneity. Second, I added additional controls such as the forward curve from one month to five years and the trade-policy-uncertainty index of Caldara et al. (2020) to guard against omitted variables. Third, I re-estimate the peso response against the Euro and the Canadian dollar and recover a similar dynamic as the one for the USD/MXN, consistent with a flight-to-safety interpretation. Fourth, I assess estimation uncertainty in the generated regressor $\hat{\alpha}_t$ with two independent checks: an LP-IV approach that instruments the Mexican EMBI spread with the premium, and a block wild bootstrap of the synthetic control. Each yields nearly identical responses. Finally, I estimate the equity response in a sector-level panel with Driscoll and Kraay (1998) standard errors, finding dynamics that match the aggregate and the sharpest declines for firms within the financial services, manufacturing, and non-basic goods and services industries.

TABLE 5: Robustness of the asset-price responses

Exercise	Result
Granger / VAR predeterminedness	Premium forecasts outcomes, not the reverse
Additional controls (forward curve, TPU)	Responses unchanged
Euro and Canadian-dollar peso	Same depreciation (flight to safety)
EMBI instrumented by the premium	Near-identical responses
Block wild bootstrap	Near-identical CI to HAC
Sector-level panel (Driscoll–Kraay)	Dynamics match aggregate

Note: Each row summarizes a robustness exercise for the baseline local-projection responses in Figure 7. Full results are reported in Appendix C.

4 A model of geoeconomic pressure and sovereign risk

The empirical evidence so far establishes that geoeconomic pressure against Mexico caused a statistically significant and economically meaningful repricing of Mexico's sovereign risk.

I use a quantitative model of sovereign default for two purposes: to check whether the repricing can be rationalized as the price of a heavier left tail in future income, and to translate the episode into a welfare cost. The framework extends the long-term debt model of Hatchondo and Martinez (2009) to incorporate a regime switch in the income process through which geoeconomic pressure shifts the distribution of future tradable income toward adverse states, with the size and shape of the shift disciplined by the SCM premium and the observed change in the skewness of Mexican income.

The regime switch is designed to capture the elevated downside risk created by the *threats* that the Trump administration directed at Mexico (e.g., withdrawal from NAFTA, a border-adjustment tax, and blanket tariffs on Mexican exports) rather than any realized change in policy. During the episode, these threats raised the probability that Mexican tradable income would fall into adverse states, and this heightened risk was priced by the government and international lenders alike.

Importantly, this is not a change in beliefs about a fixed income process (as in, e.g., Pouzo and Presno (2016)). Here, before the geoeconomic shock, income is governed by one Markov process, and after it by a different process with a thicker left tail, with *both* the government and international lenders observing and operating under the prevailing regime. I study two experiments that differ in how long the adverse regime is in force. The permanent switch treats the thicker-tailed process as the new, enduring environment (i.e., as if the geoeconomic pressure reflected a durable realignment in U.S.–Mexico relations) and delivers an upper bound on the welfare cost, since the economy is exposed to the heavier left tail in every future period. The transitory switch instead treats the adverse regime as a temporary episode: the pressure is tied to a particular political moment and lasts only for a limited spell, after which income reverts to its pre-shock process. This second case is the empirically relevant one, as it mirrors the Mexican experience during the period studied, in which the threats elevated downside risk for a time but were largely never carried out.

4.1 Environment

Households. A small open economy is populated by a representative household that receives a stochastic tradable endowment x_t each period. Following Chatterjee and Eyigungor (2012), the endowment is the sum of a persistent and a transitory component,

$$x_t = y_t + m_t, \tag{13}$$

where the persistent component y_t is the exponential of a first-order autoregressive process in logs,

$$\log y_{t+1} = \rho \log y_t + \sigma \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim \mathcal{N}(0,1), \tag{14}$$

with persistence $\rho \in (0,1)$ and standard deviation σ, and the transitory component m_t is drawn independently each period from a mean-zero truncated normal distribution on $[-\bar{m}, \bar{m}]$ with scale σ_m and bound $\bar{m} = 3\sigma_m$.[15] Households have time-separable preferences

[15] The transitory component follows Chatterjee and Eyigungor (2012): it is a computational device with a continuous cumulative distribution that convexifies the government's problem and delivers a continuous bond-price schedule, which greatly improves the numerical robustness of the long-term-debt model. m_t is assigned a small variance and is not intended to be a quantitatively important source of income fluctuations; its role is to smooth the value and price functions.

over consumption sequences,

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t\frac{c_t^{1-\gamma}}{1-\gamma}\right], \quad \text{with } \gamma>0,\ \beta\in(0,1). \tag{15}$$

Government. A benevolent government maximizes household welfare subject to a sequential budget constraint. The government has access to an international credit market where it trades non-contingent long-term bonds. Following Hatchondo and Martinez (2009), outstanding debt is characterised by a coupon that decays geometrically at rate $\delta \in (0,1)$. In each period, a bond issued at time t delivers one unit of the consumption good at $t+1$, $(1-\delta)$ units at $t+2$, $(1-\delta)^2$ units at $t+3$, and so on. The parameter $\delta = 1$ reduces to the standard one-period bond of Arellano (2008).

Let b_t denote the government's bond position, where negative values indicate borrowing. Under repayment, the sequential budget constraint is given by

$$c_t + q_t\big[b_{t+1} - (1-\delta)b_t\big] = x_t + \kappa(\delta)\, b_t, \tag{16}$$

where $x_t = y_t + m_t$ is the total endowment, q_t is the price of a new bond, and $[b_{t+1} - (1-\delta)b_t]$ is net new issuance. The coupon received on existing holdings each period is $\kappa(\delta) \equiv (r+\delta)/(1+r)$ where r is the world risk-free rate.

Each period the government chooses between repayment and default. Under default it is excluded from international capital markets and, during exclusion, loses a fraction of its persistent output through the direct cost $\phi(y_t) = \max\{0,\, d_0 y_t + d_1 y_t^2\}$, which captures the disruption to economic activity that accompanies sovereign crises; autarky consumption is therefore $c_t = y_t - \phi(y_t) + m_t = x_t - \phi(y_t)$. Following Chatterjee and Eyigungor (2012), in the period of default the transitory component takes its lowest value $-\bar{m}$, such that $y - \phi(y) - \bar{m} > 0$ so that consumption is positive for every (y, m).[16] With probability θ per period the government regains access to capital markets; otherwise exclusion continues.[17]

International lenders. A continuum of competitive foreign lenders prices sovereign bonds. Following Johri et al. (2022), lenders are risk-averse: they value future payoffs with a stochastic discount factor (SDF) that loads on the small open economy's persistent-income innovation,

$$\mathcal{M}(y_t, y_{t+1}) = \exp\Big\{-r - \chi\Big(\varepsilon^y_{t+1} + \tfrac{1}{2}\chi\,\sigma^2\Big)\Big\}, \qquad \varepsilon^y_{t+1} \equiv \log y_{t+1} - \rho\log y_t, \tag{17}$$

where ε^y_{t+1} is the AR(1) income innovation from equation (14) and $\chi \geq 0$ is the price of income risk. Because sovereign bonds pay less precisely in low-income states, where default is more likely and the SDF $\mathcal{M}$ is high, lenders demand a risk premium over and above expected default losses; the convexity term $\frac{1}{2}\chi\sigma^2$ normalizes $\mathbb{E}[\mathcal{M}] = e^{-r}$ so that the risk-free bond is priced correctly. Setting $\chi = 0$ recovers the risk-neutral benchmark in which lenders discount at the rate r. Lenders earn zero expected profits under $\mathcal{M}$, and the bond price

[16] As Chatterjee and Eyigungor (2012) note, having the transitory shock take its lowest value in the default period is a technical device that speeds computation and is not quantitatively important (setting $m = 0$ in the default period yields nearly identical results).

[17] The re-entry probability $\theta = 1/12$ is consistent with the average duration of market exclusion for Mexico following its 1982 default episode (see Gelos et al. (2011)).

schedule $q(b_{t+1}, y_t)$ is determined by the pricing condition of Section 4.3 below. Because the transitory shock m is i.i.d., knowledge of the current m_t does not help forecast future income or default; the price therefore depends only on the next-period debt b_{t+1} and the persistent component y_t, not on m_t.

4.2 Geoeconomic pressure as an income-process regime switch

In normal times, the persistent component of income is approximated by a finite-state Markov chain with N states and transition matrix $\Pi^{(0)}(y_{t+1}|y_t)$, obtained by discretizing equation (14) using the method of Tauchen (1986) with $N = 21$ equally spaced grid points in log-income space. The transitory component m is discretized separately on a grid over $[-\bar{m}, \bar{m}]$; it is symmetric and common across regimes, so the regime switch acts only on the persistent chain $\Pi^{(0)}(y_{t+1}|y_t)$.

The key innovation of this paper is to model geoeconomic pressure as a regime switch in the income process: the Markov transition matrix reallocates probability mass toward a set of adverse income states $\mathcal{L} \subset \mathcal{Y}$. Before the geoeconomic shock, income evolves according to $\Pi^{(0)}$; after the shock, it evolves according to $\Pi^{(1)}$, which places additional probability mass on low-income realizations. This captures the idea that Trump's America-First rhetoric and policy threats (e.g., withdrawal from NAFTA, border taxes, tariffs on Mexican exports) raised the likelihood of states associated with restricted market access, reduced export revenues, and disrupted supply chains. Both the government and foreign lenders observe the prevailing regime and use the corresponding transition matrix. This is a rational-expectations equilibrium under each regime: agents are not making mistakes or holding biased beliefs about a fixed process; rather, the underlying stochastic environment itself has changed.

I construct the post-shock transition matrix as a convex combination of the baseline and a target distribution that concentrates mass on bad income realizations:

$$\Pi^{(1)}(y_{t+1}|y_t;\, \omega, \zeta) = (1-\omega)\, \Pi^{(0)}(y_{t+1}|y_t) + \omega\, \Pi^{\text{target}}(y_{t+1}|y_t;\, \zeta), \tag{18}$$

where $\omega \in [0, 1)$ governs the intensity of the regime switch (at $\omega = 0$ the post-shock process coincides with the baseline $\Pi^{(0)}$, i.e., the standard long-term-debt framework) and $\Pi^{\text{target}}(y_{t+1}|y_t;\, \zeta)$ concentrates mass on the bottom tercile of the income grid, $\mathcal{L} = \{y_1, \ldots, y_7\}$ (states 1 through 7 of the 21-state grid). The within-tercile shape parameter $\zeta > 0$ controls how that additional mass is distributed across the seven states: lower ζ concentrates it on the worst realizations (a fatter left tail) and $\zeta = 1$ spreads it uniformly. The explicit tilt weights $\lambda_j(\zeta)$ and the transparent $\zeta = 1$ special case are described in Appendix D.

Figure 8 illustrates the two margins along which the regime switch operates. Panel A plots the ergodic distribution of the persistent income component under the baseline chain $\Pi^{(0)}$ in blue and under the post-shock chain $\Pi^{(1)}$ at the calibrated intensity ω^* (with $\zeta = \zeta^* = 0.513$) in red: the parameter ω governs how much probability mass is reallocated, shifting weight from the center and right of the distribution toward the low-income states in the left tail and raising the likelihood of the adverse realizations that drive default risk. Panel B shows how that additional mass is allocated within the bottom tercile: it plots the tilt weights $\lambda_j(\zeta)$ of equation (D.2) across the seven states of the tercile. At $\zeta = 1$ the mass is spread uniformly; $\zeta < 1$ concentrates it on the very worst states (a fatter left tail), while $\zeta > 1$ shifts it toward the least-bad states within the tercile.

FIGURE 8. The two margins of the regime switch: intensity ω and within-tercile shape ζ.

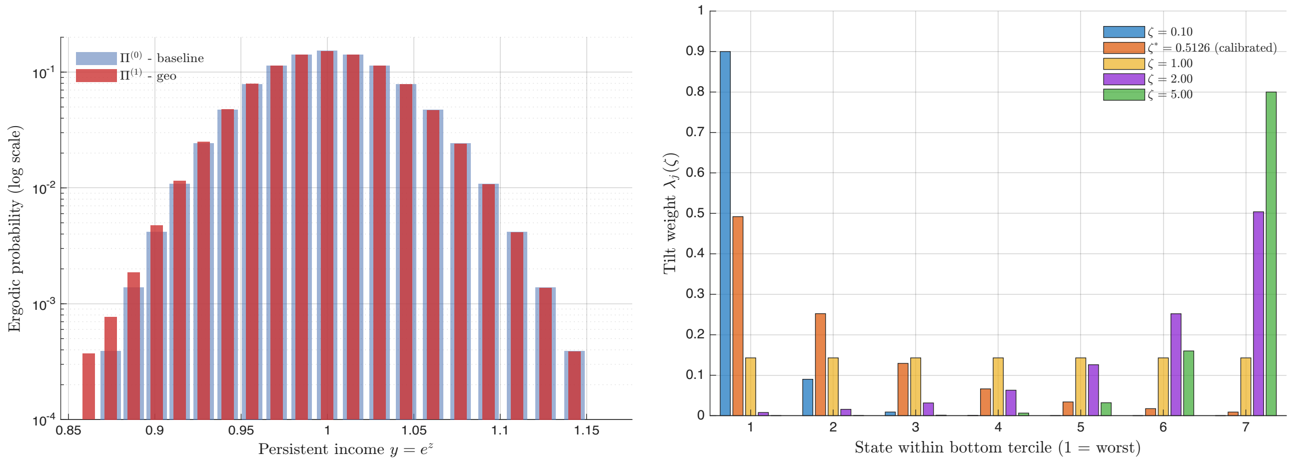

(A) Intensity ω: mass shifted to the left tail

(B) Shape ζ: allocation within the tercile

Note: Panel A: ergodic distribution of the persistent income component $y = e^z$ under the baseline transition matrix $\Pi^{(0)}$ (blue) and under the post-shock matrix $\Pi^{(1)}$ (red) at $\omega = \omega^* = 0.0003$ and $\zeta = \zeta^* = 0.513$; ω controls how much mass is reallocated toward the left tail. Panel B: tilt weights $\lambda_j(\zeta)$ across the seven states of the bottom tercile for $\zeta \in \{0.1, 0.51, 1, 2, 5\}$; $\zeta < 1$ concentrates the additional mass on the worst states, $\zeta = 1$ spreads it uniformly, and $\zeta > 1$ shifts it toward the least-bad states.

The intensity ω and the within-tercile shape ζ parameters are both calibrated in Stage 2 of the estimation procedure (Section 4.4). The intensity ω is disciplined by the 24 bp spread increase identified by the synthetic control method, while the shape ζ is disciplined by the extent to which the left tail of income thickened around the treatment, specifically the change in the skewness of the cyclical component of Mexican real GDP between the pre- and post-pressure samples.[18] Because the transitory shock m is common to both regimes, the welfare comparison below isolates the effect of the persistent-income tilt.

The tilted transition matrix in equation (D.3) yields a mathematical object similar to the one that arises in the robustness and ambiguity aversion framework of Pouzo and Presno (2016), in which a sovereign with concerns about model misspecification slants its beliefs toward worst-case distributions, producing an exponentially tilted measure over low-income states. The economic interpretation here is different. In my framework, there is no belief distortion, no ambiguity aversion, and no disagreement between agents: the tilted matrix $\Pi^{(1)}$ represents the actual income process after the geoeconomic shock, which is common knowledge to the government and lenders alike.[19] The regime switch is calibrated to an externally identified shock rather than derived from a preference for robustness, which makes the welfare calculation transparent and directly comparable to the reduced-form evidence.

[18]The cyclical component of Mexican real GDP becomes more negatively skewed during the pressure episode, its skewness falls from -1.060 before to -1.097 after, direct evidence of a thicker left tail. I target the change in skewness (-0.036); matching it returns $\zeta^* = 0.513$. Because the calibrated intensity ω^* is small, reproducing even this modest skewness change requires concentrating the reallocated mass on the worst states of the bottom tercile ($\zeta < 1$, a fatter left tail).

[19]In Pouzo and Presno (2016), only the lenders fear misspecification while the sovereign has perfect information of the true endowment process. In my setting, both parties know the prevailing regime, and there is no informational asymmetry.

4.3 Recursive equilibrium

I solve for a stationary Markov-perfect equilibrium under a given regime, i.e., in which all agents use the prevailing transition matrix ($\Pi^{(0)}$ before the shock, $\Pi^{(1)}$ after). The state of the economy is $\boldsymbol{s} \equiv (b, y, m)$: the bond position b, the persistent income component y, and the transitory shock m. Let $V(\boldsymbol{s})$ denote the value function of the government and $d(\boldsymbol{s}) \in \{0, 1\}$ the default indicator.

Value functions. The value of repayment solves

$$V^R(\boldsymbol{s}) = \max_{b' \leq 0} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta \, \mathbb{E}_{y',m'|y} \, V(\boldsymbol{s}') \right\}, \tag{19}$$

subject to the sequential budget constraint $c = y + m + \kappa(\delta)\, b - [b' - (1-\delta)b]\, q(b', y)$.

The value of financial autarky (i.e., market exclusion) depends on the current income (y, m) and solves

$$X(y, m) = \frac{[y - \phi(y) + m]^{1-\gamma}}{1-\gamma} + \beta \, \mathbb{E}_{y',m'|y} \left[\theta \, V(0, y', m') + (1-\theta)\, X(y', m') \right], \tag{20}$$

where $\phi(y) = \max\{0,\, d_0 y + d_1 y^2\}$ is the output cost and, upon re-entry (with probability θ), the sovereign starts with no debt. Following Chatterjee and Eyigungor (2012), the transitory shock takes its lowest value $-\bar{m}$ in the period of default, so the value of defaulting at state $\boldsymbol{s}$ is $X(y, -\bar{m})$, independent of the current m.

The government defaults if and only if $X(y, -\bar{m}) > V^R(\boldsymbol{s})$:

$$\begin{aligned} d(\boldsymbol{s}) &= \mathbb{1}\big\{X(y, -\bar{m}) > V^R(\boldsymbol{s})\big\}, \\ V(\boldsymbol{s}) &= \max\big\{V^R(\boldsymbol{s}),\, X(y, -\bar{m})\big\}. \end{aligned} \tag{21}$$

Bond pricing. Risk-averse foreign lenders price bonds with the SDF $\mathcal{M}_{t+1} \equiv \mathcal{M}(y_t, y_{t+1})$ of equation (17). A bond issued at t pays, at $t+1$, the coupon $\kappa(\delta)$ plus the continuation value $(1-\delta)\, q_{t+1}$ if the government repays ($d_{t+1} = 0$) and zero if it defaults, so its price satisfies the recursive pricing equation

$$q_t = \mathbb{E}_t \big[\mathcal{M}_{t+1} \, (1 - d_{t+1}) \big(\kappa(\delta) + (1-\delta)\, q_{t+1} \big) \big], \tag{22}$$

where $q_t \equiv q(b_{t+1}, y_t)$, $d_{t+1} \equiv d(b_{t+1}, y_{t+1}, m_{t+1})$, and the expectation is taken over (y_{t+1}, m_{t+1}). Because the transitory shock is i.i.d., the current m does not enter q_t, so the price depends only on (b_{t+1}, y_t).

Equilibrium definition. A recursive equilibrium consists of value functions $\{V, V^R, X\}$, a default rule $d(\boldsymbol{s})$, a borrowing policy $g(\boldsymbol{s})$, and a bond price schedule $q(b', y)$ such that: (i) V, V^R, X satisfy (19)-(21); (ii) $g(\boldsymbol{s})$ attains the maximum in (19); and (iii) $q(b', y)$ satisfies (22). The solution method is described at a high level below; full algorithmic details are provided in Appendix D.

4.4 Numerical solution, calibration, and model fit

The model is calibrated to Mexico at a quarterly frequency over the pre-treatment sample 2013Q1–2016Q2. The calibration proceeds in two stages. Table 6 summarizes all parameters. The risk aversion coefficient $\gamma = 2$ is standard in the sovereign debt literature. The quarterly risk-free rate $r = 0.01$ comes from Johri et al. (2022). The persistent income process parameters $\rho = 0.922$ and $\sigma = 0.0144$ are estimated by OLS on the cyclical component of log Mexican real GDP over 1993Q1–2016Q4, obtained by a linear detrend. The coupon decay $\delta = 0.0742$ implies a bond duration of $D = (1+r)/(\delta+r) \approx 12$ quarters, consistent with the average maturity of Mexico's external sovereign debt (see Johri et al. (2022)). The re-entry probability $\theta = 1/12$ matches the average duration of market exclusion documented by Gelos et al. (2011) for Mexico. The price of income risk is fixed at $\chi = 8.75$, the value estimated by Johri et al. (2022). The size of the tilt ω and the within-tercile shape parameter ζ are calibrated in Stage 2 (described below).

TABLE 6: Calibration

Parameter	Description	Value	Target/Source	Method
Fixed externally				
γ	Risk aversion	2	Standard	Literature
r	Risk-free rate (qtrly.)	0.010	Johri et al. (2022)	Literature
ρ	AR(1) persistence	0.922	Mexican GDP	Estimated
σ	Innovation std. dev.	0.0144	Mexican GDP	Estimated
δ	Coupon decay	0.0742	$D \approx 12$q duration	Literature
θ	Re-entry probability	0.0833	3-yr exclusion	Literature
χ	Price of income risk	8.75	Johri et al. (2022)	Literature
Calibrated (Stage 1)				
β	Discount factor	0.968	Mean debt/GDP = 12.1%	SMM
d_0	Output cost (intercept)	-0.253	Mean spread = 2.23%	SMM
d_1	Output cost (curvature)	0.309	Std(spread) = 55 bp	SMM
σ_m	Transitory shock scale	0.0089	Default prob = 1.67%	SMM
Calibrated (Stage 2)				
ω	Tilt intensity	0.0003	$\Delta s = 24$ bp	SCM
ζ	Within-tercile shape	0.513	Δ skew. = -0.036	SMM

In Stage 1, the discount factor β, the output cost parameters d_0 and d_1, and the transitory-shock scale σ_m are jointly calibrated by simulated minimum distance to four data targets: (i) mean external debt-to-GDP of 12.13%; (ii) mean sovereign spread of 2.23% per year; (iii) standard deviation of the spread of 55 bp; and (iv) an annual default probability of 1.67%, the risk-neutral implied default probability for Mexico over the pre-treatment sample. In Stage 2, the tilt intensity ω and the within-tercile shape ζ are calibrated jointly, holding the Stage 1 parameters fixed, to two targets: the 24 bp spread increase identified by the synthetic control estimator, and the change in the skewness of the cyclical component of Mexican real GDP between the pre- and post-pressure samples (-0.036). The robustness analysis additionally varies ζ over a grid as a sensitivity check on the tail shape. The calibrated values are $\beta = 0.968$, $d_0 = -0.253$, $d_1 = 0.309$, $\sigma_m = 0.0089$, $\omega^* = 0.0003$, and $\zeta^* = 0.513$.[20]

[20] I solve the model by value function iteration (VFI) on a discretized state space. The income grid has $N = 21$ points and the debt grid has $N_b = 200$ points on $[b_{\min}, 0]$. To handle the fixed-point nature of the bond pricing equation (22), I adopt the finite-horizon initialization of Hatchondo and Martinez (2009) as im-

Table 7 compares model moments to data moments over the calibration sample. The model matches the four Stage-1 targets closely as well as the Stage-2 targets.[21] On the untargeted dimensions, the model produces more volatile consumption than output, a high correlation between consumption and output, and countercyclical spreads, features that are standard for small open economies.

TABLE 7: Model fit

Moment	Data	Model
Targeted moments		
Mean external debt/GDP	12.13%	13.2%
Mean spread (ann.)	2.23%	2.04%
Std(spread, ann.)	55 bp	84 bp
Default prob. (ann.)	1.67%	1.16%
Spread increase	24 bp	23 bp
Skewness change	−0.036	−0.036
Untargeted moments		
$\text{Std}(C)/\text{Std}(Y)$	1.06	1.11
$\text{Corr}(C, Y)$	0.91	0.94
$\text{Corr}(\text{spread}, Y)$	−0.26	−0.69

Note: Model is simulated for 15,000 quarters after a 2,000-quarter burn-in. All moments are averages over non-default periods. The mean spread and standard deviation targets of 2.23% and 55 bp correspond to the EMBI Global Mexico average/standard deviation over 2013Q1–2016Q2. The default-probability target of 1.67% is the one-year risk-neutral implied annual default probability over the same window using CDS data. The mean external debt/GDP target corresponds to external debt of the federal public sector from *SHCP Estadísticas Oportunas de Finanzas Públicas*.

4.5 Results

The model's policy functions and equilibrium spread schedule behave as in the standard long-term-debt model: default is concentrated at high debt and low persistent income, and the regime switch shifts the entire spread schedule upward (most at low income, where the additional mass is concentrated), as lenders price in the elevated probability of adverse states. These objects are computed under the ergodic, permanent regime; the transitional dynamics studied next instead use the transitory MIT shock. Appendix D plots the debt policy, the default region, and the spread schedule under $\Pi^{(0)}$ and $\Pi^{(1)}$.

4.5.1 Impulse responses to the geoeconomic regime switch

To trace the transitional dynamics, I simulate the model's response to an unexpected switch from $\Pi^{(0)}$ to $\Pi^{(1)}(\omega^*, \zeta^*)$ at date $t = 0$ (an MIT shock). The experiment uses $S = 10{,}000$ independent paths, each initialized by a 200-quarter burn-in under $\Pi^{(0)}$ so that the starting state (b_0, y_0, m_0) is drawn from the ergodic distribution of the pre-shock economy. At $t = 0$

plemented in Johri et al. (2022): beginning from a terminal period in which the government cannot borrow, the algorithm iterates backward until value functions and price schedules converge. Full algorithmic details are in Appendix D.

[21] The long-term-debt model's fit on the spread level is substantially better than the short-term alternative: the one-period bond model of Arellano (2008) generates a mean spread of only 0.41%, well below the 2.23% target, because matching high spreads with short-term debt requires a counterfactually high default frequency. A full comparison appears in Appendix D.

the economy switches to the policy functions and bond price schedule associated with $\Pi^{(1)}$, and I track four variables over the seven post-treatment quarters (2016Q3–2018Q1): output Y_t, consumption C_t, external debt-to-GDP, and the sovereign spread. To place the model beside the data, I overlay each model path average on the corresponding Mexican series. Detailed procedures are in Appendix D.

FIGURE 9. Impulse responses to a geoeconomic regime switch

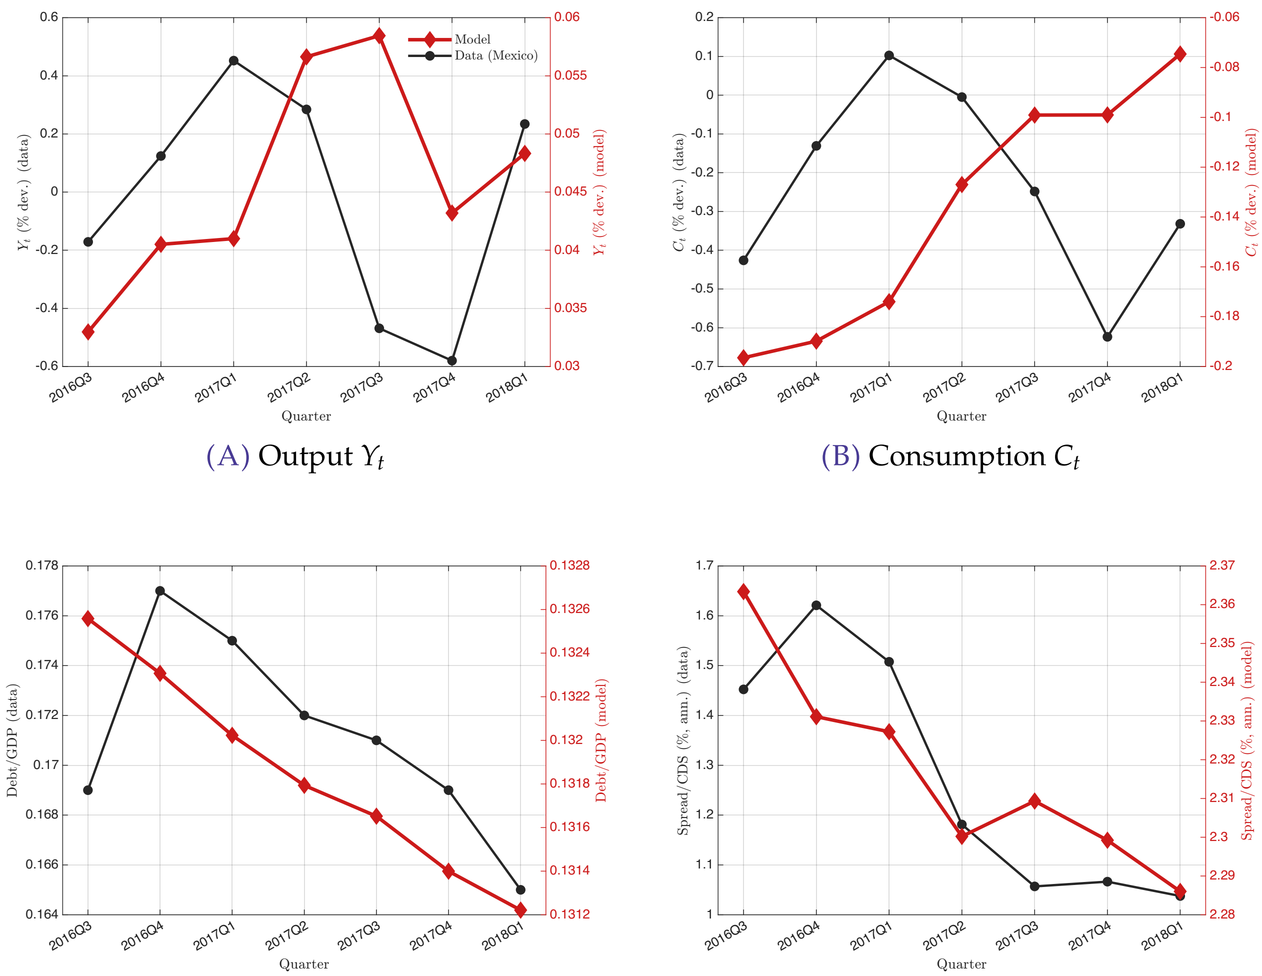

(A) Output Y_t

(B) Consumption C_t

(C) External debt/GDP

(D) Sovereign spread (CDS)

Note: Impulse responses to the regime switch, conditioning on non-default periods. Red lines with diamond markers: model average over $S = 10{,}000$ paths (right axis). Black lines with circle markers: Mexican data (left axis). Output and consumption are in percent deviation from the pre-shock mean (two-quarter moving average); external debt/GDP and the CDS spread are in levels. The window is 2016Q3–2018Q1.

Figure 9 displays the responses. Three features stand out. First, consumption falls after the shock, driven by the rising cost of debt: as lenders price in the higher probability of adverse income, bond prices fall and the government reduces borrowing and compresses consumption, before gradually recovering. Output (panel A) rises during the first 5 quarters and then falls; the increase in output reflects the non-default conditioning underlying the figure, since the paths retained at each date are selected toward higher income, and it falls in the all-periods version (Appendix D, Figure D1). The model captures the sign and comovement of the consumption and output responses, and account for a portion of their observed amplitude. Second, external debt-to-GDP (panel C) declines monotonically as the government deleverages when debt becomes more expensive; the data show a similarly declining path. Third, the spread (panel D) rises on impact and then gradually reverts toward a higher steady-state level, a pattern present in both the model and the data.

4.5.2 Welfare cost: permanent versus transitory regime switch

I quantify the welfare cost of the regime switch through two experiments: a permanent switch, which treats the geoeconomic shock as lasting forever and provides an upper bound, and a transitory switch, which lasts only for the seven-quarter post-treatment window and is the empirically relevant object. For each I compute a consumption-equivalent variation (CEV); for the transitory experiment I additionally report the average consumption contraction over the seven quarters, which is directly comparable to the sufficient statistic of Section 5.

The permanent CEV is the proportional change in consumption that makes a household indifferent between living forever under the baseline regime ($\Pi^{(0)}$) and living forever under the post-shock regime ($\Pi^{(1)}$); its closed form and the averaging over the stationary distribution are given in Appendix D. At the calibrated parameters ($\omega^* = 0.0003$, $\zeta^* = 0.513$), the welfare cost of the permanent regime switch is 0.06% of permanent consumption.

The permanent CEV overstates the welfare cost of the geoeconomic pressure episode, since the shock was not permanent: most of the threatened measures did not materialize and the pressure receded. To construct the empirically relevant counterpart, I compute the CEV of a transitory switch: the economy starts under $\Pi^{(0)}$, switches unexpectedly to $\Pi^{(1)}$ at $t = 0$, remains under $\Pi^{(1)}$ for $T^\star = 7$ quarters (the length of the post-treatment window), and then reverts to $\Pi^{(0)}$.[22] Exposed to the adverse regime for only seven quarters rather than in perpetuity, the transitory welfare cost is smaller than the permanent CEV: it equals 0.02% of consumption, about 1/3 of the permanent value. The average consumption contraction over the seven-quarter window is 0.13%, which is the object most directly comparable to the reduced-form sufficient-statistic estimate of 0.35% (Section 5).

The welfare cost is robust to how the additional mass is allocated within the tercile. Holding $\omega = \omega^*$ fixed and varying $\zeta \in \{0.25, \ldots, 5\}$, the permanent CEV ranges from 0.077% (fattest left tail, $\zeta = 0.25$) to 0.027% (least-adverse allocation, $\zeta = 5$), and remains economically meaningful throughout. Appendix D reports the full schedule (Table D3).

4.5.3 Model-implied premium and the SCM gap

Because the 24 bp average spread increase is itself a Stage-2 calibration target, the model matches the level of the SCM premium by construction; what the following exercise adds is a comparison of the dynamics of the model-implied premium against the SCM gap over the episode. To construct the model-implied premium, I simulate the economy under $\Pi^{(1)}$ and $\Pi^{(0)}$ and compute the associated spreads. Specifically, for each of the $S = 10{,}000$ paths, I

[22] This transitory experiment involves two MIT shocks: an unanticipated switch to $\Pi^{(1)}$ at $t = 0$ and an unanticipated reversion to $\Pi^{(0)}$ at $t = T^\star$. Operationally, I solve for the policy functions and bond price schedules under each regime, then simulate paths that use the $\Pi^{(1)}$ policies for $t = 0, \ldots, T^\star - 1$ and the $\Pi^{(0)}$ policies thereafter, with the bond price schedule reverting to its $\Pi^{(0)}$ value at $t = T^\star$. The transitory CEV then solves (D.4) with $\{c_t^{\Pi^{(1)}}\}$ replaced by the transitory-shock consumption stream.

simulate two trajectories simultaneously from the same starting state and the same sequence of income draws: a treated path that switches to $\Pi^{(1)}$ at $t = 0$, and a counterfactual path that remains under $\Pi^{(0)}$ forever. Denoting by $\mathcal{S}_t^{\text{nd}}$ the set of paths that are in repayment at date t, the model-implied spread premium is given by

$$\widehat{\Delta s}_t^{\text{model}} = \frac{1}{|\mathcal{S}_t^{\text{nd}}|} \sum_{s \in \mathcal{S}_t^{\text{nd}}} \left(z_t^{\Pi^{(1)},s} - z_t^{\Pi^{(0)},s} \right) \times 10{,}000 \text{ bp}, \tag{23}$$

where $z_t^{\Pi^{(1)},s}$ and $z_t^{\Pi^{(0)},s}$ are the annualized spreads of the treated and counterfactual paths s at date t, respectively. The difference is taken path-by-path before averaging, and both paths share the same sequence of income draws, so that the premium reflects only the effect of the regime switch, not idiosyncratic income fluctuations.

FIGURE 10. Model-implied CDS spread premium vs. SCM estimate.

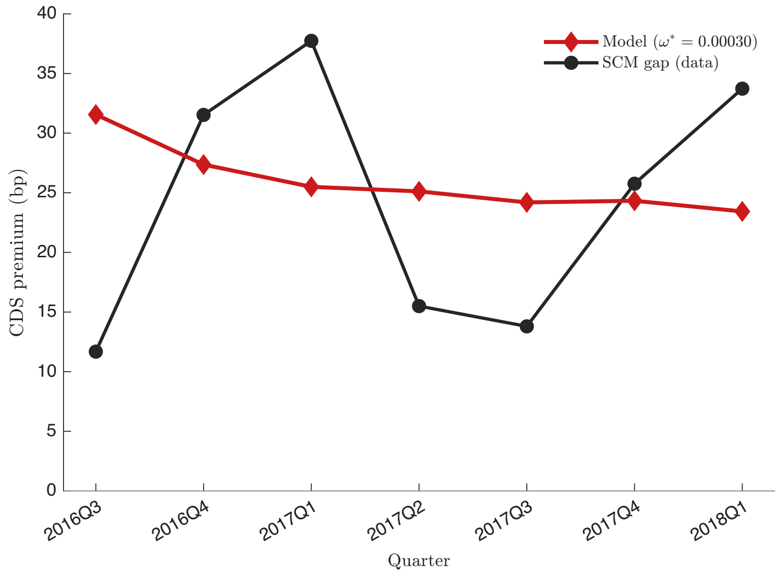

Note: Black circles display the SCM quarterly premium (basis points), while red diamonds correspond to the model-implied premium from matched-pairs simulation under constant $\omega^* = 0.0003$ and $\zeta^* = 0.513$.

Figure 10 compares the model-implied quarterly spread premium to the SCM estimates from Section 2.3 over 2016Q3–2018Q1. The quarterly premium corresponds to the average of daily observations within a quarter. By construction, the model reproduces the average SCM premium: the joint Stage-2 calibration delivers a +23 bp ergodic premium against the 24 bp SCM point estimate. The informative content is instead the shape: the model's transition path overshoots in the first two quarters and then gradually declines toward this ergodic gap. This pattern reflects the MIT shock structure: on impact, the bond price schedule reprices the entire outstanding debt stock under the new, more adverse income process, generating a large instantaneous premium. As the government adjusts its borrowing toward the new ergodic distribution, the premium reverts. The empirical SCM gap is non-monotone, which the constant-ω model cannot fully replicate. Overall, the mechanism is quantitatively plausible: a single tail-risk channel, disciplined by two targets, delivers a repricing of the right magnitude and the qualitatively correct dynamics of spreads, debt, output, and consumption.

4.6 Discussion

The model abstracts from the exchange rate and equity markets by focusing on a single tradable endowment good. However, the mechanism embedded in the regime switch has natural implications for these asset prices that are consistent with the empirical evidence. When geoeconomic pressure raises the cost of funding, the government is forced to reduce its net external liabilities, which in a richer model with a non-tradable sector would require a real depreciation of the domestic currency to restore current account balance, a twin pressure mechanism akin to the twin-defaults phenomenon documented by Na et al. (2018). Similarly, the compression of future consumption paths implied by a higher sovereign spread translates, through standard asset pricing, into lower equity valuations. Both the peso depreciation and the contraction of Mexican equity prices documented in Section 3.2, are therefore qualitatively consistent with the model, even though the current framework does not quantify these channels explicitly. Incorporating a non-tradable good or a production sector into the model would allow a more precise quantification, and is left for future work.

The post-shock transition matrix $\Pi^{(1)}$ provides a parsimonious and tractable way to capture the shift in the income environment induced by geoeconomic pressure; two other standard devices for tail risk are less suited to this episode. Stochastic volatility (Fernández-Villaverde et al. (2011)) is a symmetric change in the second moment (i.e., it thickens both tails at once) so its effect on default risk is state-dependent and generally ambiguous (see Seoane (2019)); geoeconomic pressure, by contrast, is a purely downside shift, which the left-tail tilt in $\Pi^{(1)}$ isolates. Disaster risk (Mallucci (2022)) disciplines rare output drops by their realized frequency and severity and embeds them in the ergodic distribution; here the threatened disruptions never materialized, so there is no realized-event distribution to calibrate, and the object of interest is instead the welfare cost of a tail that was priced but did not occur, with the intensity of the tilt ω identified from the causal spread response and its within-tercile shape ζ from the change in the skewness of income.

5 Transmission to consumption: A data-driven approach

In Section 4, I used a structural model of sovereign default to quantify how the geoeconomic pressure exerted by President Trump affected Mexican households. Is there a complementary, data-driven approach that can answer a similar question while exploiting the identified estimates from the synthetic control method together with high-frequency financial data? Pursuing such a route is attractive for two reasons. First, it pairs with the structural model: the model delivers the normative welfare cost, while this approach recovers the consumption response directly from the data. Second, it can recover variation that the model and low-frequency data cannot: there are no high-frequency measures of consumption, and averaging the daily treatment effect into monthly or quarterly observations attenuates much of the variation generated by President Trump, so that low-frequency estimates yield only a lower bound of the true effect.

To address such a challenge, I use two distinct methods. The first one is a sufficient statistic approach or back-of-the-envelope calculation that relies on the household's intertemporal efficiency condition and estimates from the synthetic control method to measure the overall effect on the level of consumption. The second one consists of filtering out a daily measure of consumption by combining daily information on stock returns and the Kalman filter.

Once the measure is constructed, I estimate the dynamic response of daily consumption to changes in the estimated treatment effect $\hat{\alpha}_t$ via local projections. The first method enables me to estimate the overall effect on the level of consumption, whereas the second method provides estimates of the marginal effect on the level of consumption. Both methods recover a positive measure of the consumption response (i.e., the depth of the transitory contraction due to geoeconomic coercion) which is the data-driven counterpart to the structural consumption impulse response and the average consumption contraction of Section 4, while the consumption-equivalent variation (CEV) reported there remains the model's normative, welfare-based measure.

Both approaches share the same principle as the one outlined in Amiti et al. (2021). They use high-frequency financial data to measure the welfare cost of the 2018 U.S.-China trade war. In particular, they show that they can measure the welfare impact associated with the imposition of tariffs by analyzing the policy-induced movements in the present value of firm cash flows. The principle is shared (i.e., asset prices as a sufficient statistic for the welfare impact of a policy shock) though the approach differs along two axes: they measure firm value and investor wealth through equity, whereas I measure household consumption through sovereign credit and equity; and their inference rests on the present-value identity for firm cash flows, whereas mine rests on a consumption Euler equation and a filtered daily consumption series. In that same sense, I argue that I can identify the consumption loss due to Trump's rhetoric and actions by analyzing the changes in default risk and stock returns at high-frequency.[23] From the first approach, I find that the treatment led to an overall contraction of consumption of 0.35% relative to a counterfactual path. Using the second approach, I find that Mexican households decreased their consumption by 0.02% for any 1 basis point increase in the shock measure (i.e., $\hat{\alpha}_t$) and such effect takes up to four weeks to materialize.

5.1 Sufficient statistic approach

I recover the effect of Trump's rhetoric on the level of consumption from the household's intertemporal efficiency condition of the Section 4 government problem, evaluated under a small set of simplifications that make the calculation parsimonious and let the consumption loss be read directly from asset prices. The calculation inherits CRRA preferences, the discount factor β, and the constant world risk-free rate; it also keeps the sovereign bond priced by the same risk-averse lenders, whose premium enters through the observed price q_t that I recover from the CDS data. Two features of the model outlined in the previous section are relaxed, one because it would otherwise preclude a well-defined Euler equation, the other for tractability. First, endogenous default makes the value function kinked; I take the default probability as exogenous, priced in the CDS, so default is no longer strategic and the household's optimality condition is well-defined along the repayment path. Second, long-term debt would require tracking the outstanding stock and the issuance policy, making the task of tracking these at high-frequency quite hard; I therefore use short-term debt. A single distributional assumption, that consumption growth is conditionally log-normal, then delivers the closed form. The resulting Euler equation is the Section 4 household's optimality

[23] If one considers an agent that only values the consumption good and nothing else (i.e., the only input to the instantaneous utility function is C_t), then there is a one-to-one mapping between changes in consumption and changes in welfare.

condition under these relaxations, which under CRRA preferences takes the form:

$$C_t = \left(\frac{\beta}{q_t} \mathbb{E}_t \left[C_{t+1}^{-\gamma} \right] \right)^{-\frac{1}{\gamma}}, \tag{24}$$

where C_t stands for the level of consumption, q_t represents the one-period discount price on foreign debt, β is the subjective discount factor, and $\mathbb{E}_t[\cdot]$ is the expectation operator conditional on information available up to period t. Log-linearizing the previous expression and iterating forward T periods yields:

$$c_t = \frac{1}{\gamma} \sum_{j=0}^{T} \mathbb{E}_t \left[\log(q_{t+j}) \right] - \frac{(T+1)}{\gamma} \log(\beta) - \frac{\gamma}{2} \sum_{j=0}^{T} \mathbb{E}_t \left[\mathrm{Var}_{t+j}(\Delta c_{t+j+1}) \right] + \mathbb{E}_t \left[c_{t+T+1} \right], \tag{25}$$

where $\log(C_t) \equiv c_t$, γ stands for the coefficient of risk aversion, q_{t+j} is the discount price of debt in period $t+j$, β the discount factor, $\mathrm{Var}_t(\Delta c_{t+j})$ is the conditional variance of consumption growth, and $\mathbb{E}_t[c_{t+T+1}]$ represents a terminal condition. Now, let c_t represent the level of consumption associated with the treatment, and let $\tilde{c}_t$ represent the counterfactual level of consumption in the absence of the treatment which will have the same expression as (25) but with counterfactual observations for the one-period discount price, log consumption growth, and terminal condition.[24] Then the overall effect of the repricing of sovereign risk due to geoeconomic pressure on consumption can be decomposed into the following three channels or components:

$$c_t - \tilde{c}_t = \underbrace{\frac{1}{\gamma} \left(\sum_{j=0}^{T} \mathbb{E}_t \left[\log q_{t+j} \right] - \sum_{j=0}^{T} \mathbb{E}_t \left[\log \tilde{q}_{t+j} \right] \right)}_{\text{(1) Pricing component}}$$

$$\underbrace{- \frac{\gamma}{2} \left(\sum_{j=0}^{T} \mathbb{E}_t \left[\mathrm{Var}_{t+j} \left(\Delta c_{t+j+1} \right) \right] - \sum_{j=0}^{T} \mathbb{E}_t \left[\mathrm{Var}_{t+j} \left(\Delta \tilde{c}_{t+j+1} \right) \right] \right)}_{\text{(2) Uncertainty component}} + \underbrace{\left(\mathbb{E}_t \left[c_{t+T+1} \right] - \mathbb{E}_t \left[\tilde{c}_{t+T+1} \right] \right)}_{\text{(3) Transitory vs. permanent component}} . \tag{26}$$

The pricing component is measured directly from the observed and synthetic 5-year CDS. From the observed spread I recover the daily risk-neutral default probability π_{t+j}, and hence the bond price $q_{t+j} = (1 - \pi_{t+j})/(1 + r_f)$, and from the synthetic-control spread its counterpart $\tilde{\pi}_{t+j}$ (equivalently $\tilde{q}_{t+j}$), transforming each spread under a constant-hazard assumption following Hébert and Schreger (2017).[25] The uncertainty channel or component is hard to measure as there is no daily data on both the observed and counterfactual consumption that I can use to recover the conditional variances; however, following the literature of consumption-based asset pricing (see Campbell (2003) and Ludvigson (2013) for a detailed exposition), I posit that the variance of daily consumption growth is proportional to the variance of stock returns:

$$\mathrm{Var}_t(\Delta c_{t+1}) \propto \Psi \mathrm{Var}_t(r_{t+1}), \tag{27}$$

[24]For a detailed derivation, see Appendix E.

[25]Given a constant hazard rate $\lambda = CDS/(1-R)$, the daily risk-neutral default probability is $1 - \exp(-5\lambda/1{,}300) = 1 - \exp(-\lambda/260)$. The recovery rate R enters only to extract the probability from the spread; the pricing identity $q_t = (1 - \pi_t)/(1 + r_f)$ underlying the derivation is that of a zero-recovery one-period claim, so only the probability, not the recovery assumption, is carried across, and the two are mutually consistent.

where Ψ is a constant of proportionality and $\text{Var}_t(r_{t+1})$ is the conditional variance of log-returns. In this particular case, I set $\Psi = 2\%$ and $\text{Var}_t(r_{t+1})$ is estimated using a standard GARCH(1,1) model.[26] Regarding the conditional variance of the counterfactual growth rate of consumption, I approximate it by using the convex combination of volatilities of stock returns for the countries that got positive weights in the synthetic control for the 5-year CDS. In other words, using the variances of stock returns for Panama, Colombia, Poland, and the estimated weights from the synthetic control $w_i^\star$, I approximate the counterfactual variance as:

$$\text{Var}_t(\Delta\tilde{c}_{t+1}) \propto \Psi \left[\sum_{i=2}^{4} w_i^\star \text{Var}_t(r_{i,t+1}) \right]. \tag{28}$$

Figure 11 illustrates the dynamics of the conditional variances of consumption growth, estimated from both observed and counterfactual stock returns.[27] Prior to the treatment, the two series are relatively similar. Following the treatment, I observe that the variance of consumption growth for the treated unit is consistently higher compared to the counterfactual. It is worth emphasizing that the weights used, $\{w_i^\star\}_{i=2}^4$, correspond to the synthetic control estimates for the CDS data; that is, they are not re-estimated to match the pre-treatment dynamics of the volatility of returns $\text{Var}_t(r_{t+1})$. I observe that the treatment resulted in heightened uncertainty regarding consumption growth for Mexican households.

FIGURE 11. Counterfactual variance of stock returns

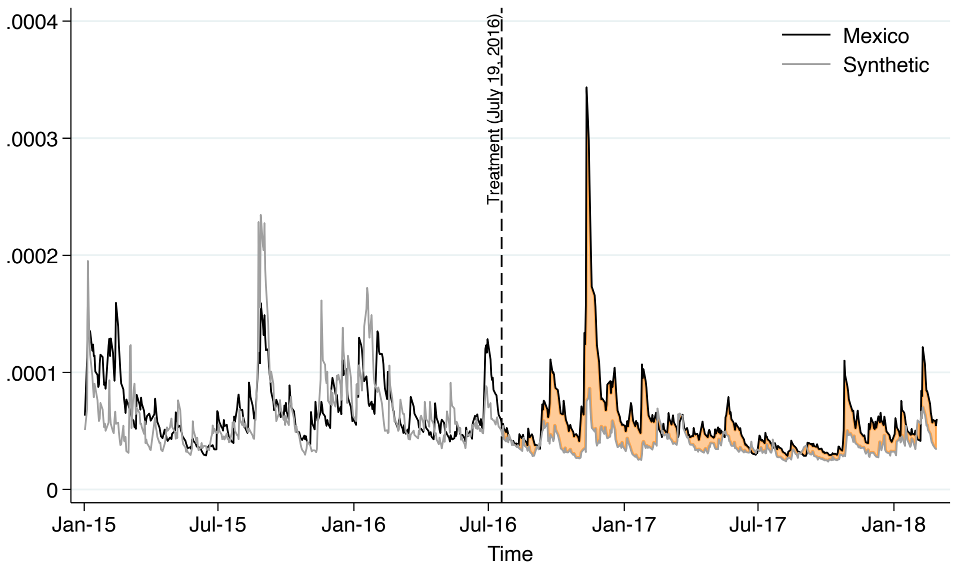

Note: The black line represents the estimated variance of Mexican stock returns using a GARCH(1,1) model. The red line corresponds to the convex combination of estimated variances of stock returns for Panama, Colombia, and Poland. The weights are those of the baseline estimation of the synthetic control using 5-year CDS data.

Finally, I set the horizon T to the length of the treatment window (i.e., the trading days

[26] The empirical evidence at observed lower frequencies (e.g., monthly, quarterly, semiannual, and yearly) on the ratio $\frac{\text{Var}(\Delta c_t)}{\text{Var}(r_t)}$ is consistently 2%. Alternatively, if I estimate a linear model as in Campbell (2003) where I regress Δc_t on r_t at low frequencies, I obtain, on average, a point estimate of the slope coefficient of 0.04.

[27] To estimate the conditional variance of stock returns for Mexico, I utilize daily data from the S&P/BMV IPC index covering the period from January 2014 to February 2018. For Panama, Colombia, and Poland, I use the BVPS index, the COLCAP index, and the WIG index, respectively, for the same time frame.

spanning President Trump's first campaign and initial year in office) beyond which the observed and counterfactual economies coincide, so that $q_{t+j} = \tilde{q}_{t+j}$ and $\text{Var}_{t+j}(\Delta c_{t+j+1}) = \text{Var}_{t+j}(\Delta \tilde{c}_{t+j+1})$ for any $t+j$ past the window. The infinite sums in the first two channels therefore truncate to the treatment window, and under the assumption that the geoeconomic pressures were transitory to the Mexican economy, so both economies converge to the same long-run level, the third channel vanishes.[28] Using a standard value for the coefficient of risk aversion of $\gamma = 2$, I estimate that the overall effect of Trump's rhetoric and actions on consumption is -0.35%. That is, consumption fell 0.35% relative to the counterfactual path as a consequence of the treatment. The two channels reduce consumption through distinct mechanisms. The pricing channel operates through the intertemporal price of resources: geoeconomic pressure raises the market-implied default probability, so the sovereign bond price q_t falls and external borrowing becomes more expensive; facing a higher effective cost of moving resources to the present, the household borrows less and cuts consumption today. The uncertainty channel operates through precautionary saving: the heavier left tail raises the conditional variance of future consumption growth, inducing households to save more and consume less today. Both push consumption down; their relative weight is governed by γ and Ψ. At the conservative baseline ($\gamma = 2$, $\Psi = 2\%$) the pricing channel accounts for the bulk of the contraction (94% vs. 6%), because the repricing of default risk is first-order while the precautionary term scales with γ and the small consumption–return variance ratio; as either γ or Ψ rises, the precautionary channel gains weight (see Figure 12).

FIGURE 12. Overall effect on consumption

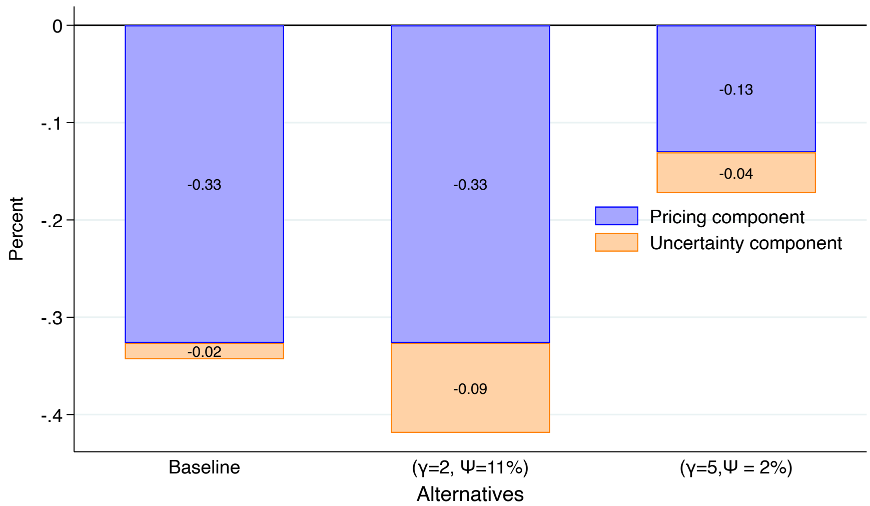

Note: The blue shaded bars correspond to the first component of the effect on log consumption, while the orange shaded bars represent the second component. The calculation relies on the assumption that Trump's actions and rhetoric are a transitory shock to the Mexican economy. The first bar displays the result for the baseline calibration, while the second and third bars correspond to the contraction of log-consumption for different values for Ψ and γ, respectively.

[28] Given that most of the threats made by President Trump did not materialize and that he did not get re-elected immediately after his first mandate, it is most likely that the shock was transitory.

5.2 Filtering a daily consumption measure

The previous approach allows for estimating the overall effect of the treatment on the level of consumption; however, it is silent about the marginal response of household's consumption to an increase in the premium $\hat{\alpha}_t$ (i.e., a semi-elasticity). This second approach will allow me to analyze how consumption changed in response to a Trump-induced increase in risk premia of 1 basis point by using the local-projection framework outlined in Section 3.2. This is useful as it will allow me to estimate the response of Mexican households to particular events or threats by President Trump (e.g., trade-related threats).

I construct the daily series by combining daily Mexican equity returns with the monthly *INEGI* private-consumption index through a Kalman filter.[29] Following the Long-Run Risk model of Bansal and Yaron (2004), daily consumption growth and stock returns share an unobservable, persistent component, which provides an economically-motivated mapping from the daily return signal to daily consumption growth; a monthly aggregation constraint anchors the amplitude of the inferred series to the observed data, up to measurement error. The Long-Run Risk structure is used here as an interpolation device, not as a fully-fledged asset-pricing model. I estimate the resulting mixed-frequency state-space system by maximum likelihood and recover the smoothed daily series with a Rauch-Tung-Striebel smoother; Appendix F details the state-space representation, the augmented-state formulation, the estimation, and the smoothing algorithm, and reports the smoothed series against the observed monthly index.

With the smoothed daily series in hand, I estimate the dynamic response of consumption to changes in $\hat{\alpha}_t$ with the following local projection:[30]

$$c_{t+h} - c_{t-1} = a^h + \hat{\alpha}_t \theta^h + \boldsymbol{w}_t \psi^h + u_{t+h}, \quad \text{with } h = 0, \ldots, 20, \tag{29}$$

where $c_{t+h} - c_{t-1}$ is the long-difference in the log-level of consumption, $\hat{\alpha}_t$ is the difference between the observed CDS and the synthetic CDS, $\boldsymbol{w}_t$ contains controls such as lagged observations for $\hat{\alpha}_t$, lagged daily changes in the exchange rate, equity returns, and domestic monetary policy rate, global volatility controls such as the VIX index, oil prices and changes in the Fed fund rates, and u_{t+h} is the error term. Figure 13 displays the cumulative response $\{\hat{\theta}^h\}_{h=0}^{20}$ to a 1 bp increase in the shock variable. I find that consumption falls around 0.02% after a 1 bp increase in the shock measure, and that response takes some time to build up, almost four weeks after the shock occurred.

Combining these estimates with the tweet responses of Section 3.1, a negative announcement implies a consumption contraction of about 0.12% after two weeks, rising to roughly 0.27% for trade-related threats; a one-standard-deviation increase in $\hat{\alpha}_t$ implies about 0.23%, and the 24 bp average premium implies about 0.48%. These marginal estimates are broadly consistent with the 0.35% overall effect from the sufficient-statistic approach. As a further, out-of-sample check, Appendix F uses the same smoothed series to gauge the consumption

[29] *INEGI* reports the monthly index of private consumption, which tracks household expenditures on goods and services and is the highest-frequency indicator available for the largest demand-side component of GDP. I use the seasonally-adjusted series; both the equity index (Bloomberg) and the consumption index run from January 2014 to February 2018.

[30] The Rauch-Tung-Striebel smoother is a two-sided filter, so the reconstructed series at t embeds full-sample information. Re-estimating the local projection on the one-sided filtered series (conditioning only on information through t) delivers very similar dynamic responses, so the estimates are not an artifact of the two-sided smoother.

FIGURE 13. Cumulative impulse response function to a 1 bp increase in $\hat{\alpha}_t$

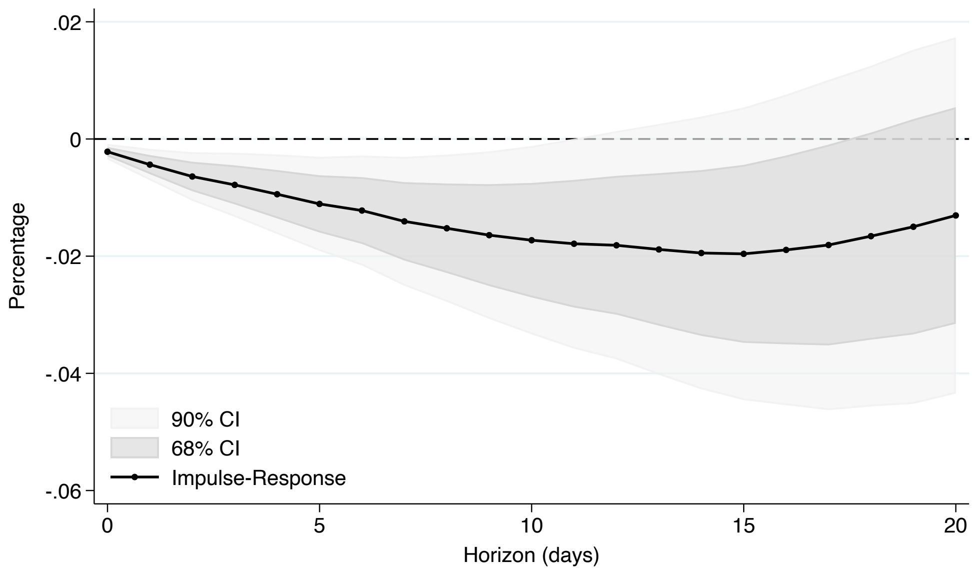

Note: The black markers represent the point estimates $\{\hat{\theta}^h\}_{h=0}^{20}$. The shaded areas represent the 68% and 90% confidence intervals, respectively, constructed using HAC robust standard errors.

cost of an outright default, and finds it comparable to Mexico's 1994–1995 crisis and to past sovereign defaults in Argentina, Ecuador, and Russia.

Relation to the literature and the structural model These magnitudes are in line with the literature on country spreads and business cycles. Neumeyer and Perri (2005) find that a one-standard-deviation shock to the country spread contracts Argentine consumption by about 1% on impact, and Fernández-Villaverde et al. (2011) show that shocks to spread volatility generate consumption contractions of about 0.21% on average, with hump-shaped dynamics; my estimates share both the sign and the delayed build-up. The closest antecedent is Sagawa (2026), who recovers exogenous sovereign-risk shocks from high-frequency Argentine CDS movements and, like me, traces them to consumption with local projections; scaled to a common 100-basis-point move, our responses are of the same order of magnitude.

Relative to the structural model of Section 4, the data-driven responses are larger. Three forces, all operating in the same direction, account for the gap. First, the structural model isolates a single default-risk channel and omits the real amplification mechanisms that the spreads-and-business-cycle literature relies on and that the high-frequency data absorb. Second, the model delivers a quarterly-average response, whereas the high-frequency methods recover the effect before within-quarter aggregation attenuates it (see Heaton (1993)). Third, the sufficient-statistic calculation prices a one-period bond and so refinances the entire debt stock each period, whereas the structural model carries long-term debt that reprices only a fraction; this short-term assumption amplifies the consumption response.[31] Because all three forces push the same way, the data-driven estimates are naturally read as an upper bound.

[31] I document a larger welfare cost under short-term than under long-term debt in Short-term debt model.

6 Conclusion

Coercive geoeconomic pressure from a large economy can reprice a smaller sovereign's default risk by shifting the distribution of its future income toward adverse states, well before any of that risk is realized. This paper uses the pressure that the Trump administration directed at Mexico during the 2016 campaign and his first year in office as a laboratory to measure that repricing and trace its transmission. Treating the episode as a natural experiment, I recover the exogenous variation in Mexico's country-risk premium and use it to quantify the response of the exchange rate, equity prices, and aggregate consumption.

I find that Mexico's 5-year CDS rose 24 basis points on average relative to a synthetic control built from countries with comparable fundamentals, with the largest gaps on the day of the first presidential debate, the day Trump was declared the winner of the U.S. presidential election, and the January 2017 inauguration. The premium is sensitive to his threats, especially trade-related ones, and it responds to pressure aimed at Mexico rather than at other economies. This geoeconomic repricing had sizable real effects: a 1 basis point increase in the premium translates into a 5.3 cent depreciation of the peso (approximately 0.26%), a 0.11% decline in the Mexican stock index, an overall fall of 0.35% in household consumption, and a marginal consumption response of -0.02% per basis point. To read these magnitudes through the lens of sovereign risk, I embed the episode in a quantitative default model as a regime switch that shifts income toward adverse states. The model produces an average consumption contraction of 0.13% over the episode, the counterpart to the 0.35% data estimate, and a welfare cost, expressed as a consumption equivalent, of 0.02% for the transitory episode (the empirically relevant case), rising to 0.06% if the pressure were permanent.

The design has a scope condition worth making explicit. Synthetic control recovers the repricing cleanly when pressure is aimed at a single economy, so that comparable, unpressured economies form a valid counterfactual. That condition fits the 2016–2018 episode but not the current conjuncture: as the second Trump administration directs pressure at many economies at once and increasingly carries out its threats, the donor pool is itself contaminated and the synthetic counterfactual erodes. Extending this agenda to the present period therefore calls for high-frequency identification that does not rely on a clean control group, such as the heteroskedasticity-based approach used here.

More broadly, the episode shows that coercive geoeconomic pressure can move sovereign risk and pass through to asset prices and consumption even when the threatened measures are never carried out, broadening the sources of sovereign risk beyond the domestic fundamentals that standard default models emphasize. A natural agenda treats the pressure itself as an equilibrium object: on the sending side, a strategic external actor could choose the intensity of coercion by weighing its geopolitical objective against the cost of exerting it, in the spirit of Clayton et al. (2023), with sovereign credit markets not merely pricing the threat but amplifying it, since the repricing of default risk raises the target's cost of external finance and can itself make the country more susceptible to coercion. On the receiving side, lenders could be uncertain whether the threats reflect resolve or are largely cheap talk, and learn about it as the administration does or does not act; unlike the worst-case beliefs over a fixed income process in Pouzo and Presno (2016), here the object of learning is the type of an external actor, and Bayesian updating could endogenize the rise and gradual reversion of the premium documented here. Bringing the coercer's choice and the market's inference into a single framework, linked by the sovereign-risk channel this paper isolates, is a promising direction for future work.

References

A. Abadie. Using Synthetic Controls: Feasibility, Data Requirements, and Methodological Aspects. *Journal of Economic Literature*, 59(2):391–425, 2021.

A. Abadie and J. Gardeazabal. The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132, 2003.

A. Abadie and J. Vives-i-Bastida. Synthetic Controls in Action. *arXiv preprint arXiv:2203.06279*, 2022.

A. Abadie, A. Diamond, and J. Hainmueller. Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(490):493–505, 2010.

A. Abadie, A. Diamond, and J. Hainmueller. Comparative Politics and the Synthetic Control Method. *American Journal of Political Science*, 59(2):495–510, 2015.

M. Aguiar and M. Amador. *The Economics of Sovereign Debt and Default*, volume 1. Princeton University Press, 2021.

M. Aguiar and G. Gopinath. Defaultable Debt, Interest Rates and the Current Account. *Journal of International Economics*, 69(1):64–83, 2006.

M. Amiti, S. H. Kong, and D. E. Weinstein. Trade Protection, Stock-Market Returns, and Welfare. Technical report, National Bureau of Economic Research, 2021.

C. Arellano. Default Risk and Income Fluctuations in Emerging Economies. *American Economic Review*, 98(3):690–712, 2008.

Banco de México. Quarterly Reports, 2017. URL `https://www.banxico.org.mx/publications-and-press/quarterly-reports/quarterly-reports-prices-banc.html`.

R. Bansal and A. Yaron. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.

T. R. Beard, H. Kim, and M. L. Stern. Is Good News for Donald Trump Bad News for the Peso? *Applied Economics Letters*, 24(19):1363–1368, 2017.

D. Benjamin and M. L. Wright. Recovery before Redemption: A Theory of Delays in Sovereign Debt Renegotiations. *Available at SSRN 1392539*, 2009.

E. J. Blanchard. Renegotiating NAFTA: The Role of Global Supply Chains. In *Economics and Policy in the Age of Trump*, chapter 16, pages 175–183. CEPR Press London, UK, 2017.

E. Borensztein and U. Panizza. The Costs of Sovereign Default. *IMF Staff Papers*, 56(4): 683–741, 2009.

B. Born, G. J. Müller, M. Schularick, and P. Sedláček. The Costs of Economic Nationalism: Evidence from the Brexit Experiment. *The Economic Journal*, 129(623):2722–2744, 2019.

C. Bown. *Economics and Policy in the Age of Trump*. CEPR Press London, UK, 2017.

M. B. Brown and A. B. Forsythe. Robust Tests for the Equality of Variances. *Journal of the American Statistical Association*, 69(346):364–367, 1974.

J. I. Bulow and K. S. Rogoff. Sovereign Debt: Is to Forgive to Forget?, 1988.

D. Caldara, M. Iacoviello, P. Molligo, A. Prestipino, and A. Raffo. The Economic Effects of Trade Policy Uncertainty. *Journal of Monetary Economics*, 109:38–59, 2020.

J. Y. Campbell. Consumption-based Asset Pricing. *Handbook of the Economics of Finance*, 1: 803–887, 2003.

M. D. Cattaneo, Y. Feng, and R. Titiunik. Prediction Intervals for Synthetic Control Methods. *Journal of the American Statistical Association*, 116(536):1865–1880, 2021.

M. D. Cattaneo, Y. Feng, F. Palomba, and R. Titiunik. `scpi`: Uncertainty Quantification for Synthetic Control Methods. *Journal of Statistical Software*, 2025. Forthcoming.

R. Chang. Financial Crises and Political Crises. *Journal of Monetary Economics*, 54(8):2409–2420, 2007.

S. Chatterjee and B. Eyigungor. Maturity, Indebtedness, and Default Risk. *American Economic Review*, 102(6):2674–2699, 2012.

C. Clayton, M. Maggiori, and J. Schreger. A Framework for Geoeconomics. Technical report, National Bureau of Economic Research, 2023.

C. Clayton, A. Coppola, M. Maggiori, and J. Schreger. Geoeconomic Pressure. *Columbia Business School Research Paper Forthcoming, Stanford University Graduate School of Business Research Paper Forthcoming*, 2025.

J. H. Cochrane and M. Piazzesi. The Fed and Interest Rates — A High-Frequency Identification. *American Economic Review*, 92(2):90–95, 2002.

T. F. Cooley and L. J. White. Potential Changes in Financial Regulation in the Age of Trump. In *Economics and Policy in the Age of Trump*, chapter 11, pages 121–128. CEPR Press London, UK, 2017.

Council on Foreign Relations. 2017-2021: Trump's Foreign Policy Moments, 2024. URL `https://www.cfr.org/timeline/trumps-foreign-policy-moments`.

J. G. Cragg and S. G. Donald. Testing Identifiability and Specification in Instrumental Variable Models. *Econometric Theory*, 9(2):222–240, 1993.

J. J. Cruces and C. Trebesch. Sovereign Defaults: The Price of Haircuts. *American Economic Journal: Macroeconomics*, 5(3):85–117, 2013.

G. Cuadra and H. Sapriza. Sovereign Default, Interest Rates and Political Uncertainty in Emerging Markets. *Journal of International Economics*, 76(1):78–88, 2008.

J. C. Driscoll and A. C. Kraay. Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data. *Review of Economics and Statistics*, 80(4):549–560, 1998.

J. Eaton and M. Gersovitz. Debt with Potential Repudiation: Theoretical and Empirical Analysis. *The Review of Economic Studies*, 48(2):289–309, 1981.

Factba.se. Donald Trump, Speeches, Tweets, Policy, Unedited, Unfiltered, Instantly, 2023. URL `https://factba.se/trump/`.

R. C. Feenstra, R. Inklaar, and M. P. Timmer. The Next Generation of the Penn World Table. *American Economic Review*, 105(10):3150–3182, 2015.

J. Fernández-Villaverde, P. Guerrón-Quintana, J. F. Rubio-Ramirez, and M. Uribe. Risk Matters: The Real Effects of Volatility Shocks. *American Economic Review*, 101(6):2530–2561, 2011.

I. Filippou, A. E. Gozluklu, M. T. Nguyen, and G. Viswanath-Natraj. Signal in the Noise: Trump Tweets and the Currency Market. *Journal of International Money and Finance*, page 103343, 2025.

S. Fourakis. *Essays on Sovereign Default with Information Frictions*. PhD thesis, University of Minnesota, 2022.

R. G. Gelos, R. Sahay, and G. Sandleris. Sovereign Borrowing by Developing Countries: What Determines Market Access? *Journal of International Economics*, 83(2):243–254, March 2011. doi: 10.1016/j.jinteco.2010.11.007.

P. Goldsmith-Pinkham and T. Lyu. Causal Inference in Financial Event Studies. Working paper, Yale University, 2025.

G. Gopinath, P.-O. Gourinchas, A. F. Presbitero, and P. Topalova. Changing Global Linkages: A New Cold War? *Journal of International Economics*, 153:104042, 2025.

K. Handley and N. Limão. Trade under T.R.U.M.P. Policies. In *Economics and Policy in the Age of Trump*, chapter 13, pages 141–152. CEPR Press London, UK, 2017.

J. C. Hatchondo and L. Martinez. Long-duration Bonds and Sovereign Defaults. *Journal of Iinternational Economics*, 79(1):117–125, 2009.

J. Heaton. The Interaction Between Time-Nonseparable Preferences and Time Aggregation. *Econometrica*, pages 353–385, 1993.

B. Hébert and J. Schreger. The Costs of Sovereign Default: Evidence from Argentina. *American Economic Review*, 107(10):3119–3145, 2017.

A. Johri, S. Khan, and C. Sosa-Padilla. Interest Rate Uncertainty and Sovereign Default Risk. *Journal of International Economics*, 139:103681, 2022.

Ò. Jordà. Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182, 2005.

D. R. Känzig. The Macroeconomic Effects of Oil Supply News: Evidence from OPEC Announcements. *American Economic Review*, 111(4):1092–1125, 2021.

F. Kleibergen and R. Paap. Generalized Reduced Rank Tests Using the Singular Value Decomposition. *Journal of Econometrics*, 133(1):97–126, 2006.

H. Levene. Robust Tests for Equality of Variances. *Contributions to Probability and Statistics*, pages 278–292, 1960.

D. Li, M. Plagborg-Møller, and C. K. Wolf. Local Projections vs. VARs: Lessons from Thousands of DGPs. Technical report, National Bureau of Economic Research, 2022.

F. A. Longstaff, J. Pan, L. H. Pedersen, and K. J. Singleton. How Sovereign Is Sovereign Credit Risk? *American Economic Journal: Macroeconomics*, 3(2):75–103, 2011.

S. Lord-Medrano. The Effect of the U.S. Presidential Election and Political Uncertainty on the Mexican CDS. Master's thesis, Instituto Tecnológico Autónomo de México (ITAM), 2020.

S. C. Ludvigson. Advances in Consumption-Based Asset Pricing: Empirical Tests. *Handbook of the Economics of Finance*, 2:799–906, 2013.

A. Lusompa. Local Projections, Autocorrelation, and Efficiency. *Quantitative Economics*, 14 (4):1199–1220, 2023.

E. Mallucci. Natural Disasters, Climate Change, and Sovereign Risk. *Journal of International Economics*, 139:103672, 2022.

S. Mao. Actions Speak Louder than Tweets. *University of California San Diego Department of Economics, San Diego, CA USA*, 2019.

J. V. Martinez and G. Sandleris. Is it Punishment? Sovereign Defaults and the Decline in Trade. *Journal of International Money and Finance*, 30(6):909–930, 2011.

D. Matveev and F. Ruge-Murcia. Tariffs and the Exchange Rate: Evidence from Twitter. *IMF Economic Review*, 72(3):1185–1211, 2024.

S. Miranda-Agrippino and G. Ricco. The transmission of Monetary Policy Shocks. *American Economic Journal: Macroeconomics*, 13(3):74–107, 2021.

C. Mohr and C. Trebesch. Geoeconomics. *Annual Review of Economics*, 17, 2025.

J. L. Montiel Olea and C. Pflueger. A Robust Test for Weak Instruments. *Journal of Business & Economic Statistics*, 31(3):358–369, 2013.

J. L. Montiel Olea and M. Plagborg-Møller. Local Projection Inference is Simpler and More Robust than you Think. *Econometrica*, 89(4):1789–1823, 2021.

K. M. Murphy and R. H. Topel. Estimation and Inference in Two-Step Econometric Models. *Journal of Business & Economic Statistics*, 20(1):88–97, 2002.

S. Na, S. Schmitt-Grohé, M. Uribe, and V. Yue. The Twin D's: Optimal Default and Devaluation. *American Economic Review*, 108(7):1773–1819, 2018.

E. Nakamura and J. Steinsson. High-Frequency Identification of Monetary Non-Neutrality: The Information Effect. *The Quarterly Journal of Economics*, 133(3):1283–1330, 2018.

P. A. Neumeyer and F. Perri. Business Cycles in Emerging Economies: The Role of Interest Rates. *Journal of Monetary Economics*, 52(2):345–380, 2005.

A. Pagan. Econometric Issues in the Analysis of Regressions with Generated Regressors. *International Economic Review*, pages 221–247, 1984.

D. Pouzo and I. Presno. Sovereign Default Risk and Uncertainty Premia. *American Economic Journal: Macroeconomics*, 8(3):230–266, 2016.

V. A. Ramey. Macroeconomic Shocks and their Propagation. *Handbook of Macroeconomics*, 2: 71–162, 2016.

R. Rigobon. Identification through Heteroskedasticity. *Review of Economics and Statistics*, 85 (4):777–792, 2003.

R. Rigobon and B. Sack. The Impact of Monetary Policy on Asset Prices. *Journal of Monetary Economics*, 51(8):1553–1575, 2004.

A. K. Rose. One Reason Countries Pay their Debts: Renegotiation and International Trade. *Journal of Development Economics*, 77(1):189–206, 2005.

T. Sagawa. Sovereign Liquidity Shocks. Working paper, Northwestern University, February 2026. URL `https://takaakisagawa.com/files/papers/sovereign_liquidity_shocks.pdf`.

H. D. Seoane. Time-varying Volatility, Default, and the Sovereign Risk Premium. *International Economic Review*, 60(1):283–301, 2019.

S&P. S&P/BMV Sectoral Analytical Indices, 2023. URL `https://www.spglobal.com/spdji/en/indices/equity/sp-bmv-consumer-discretionary-sector-index/#overview`.

G. Tauchen. Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions. *Economics letters*, 20(2):177–181, 1986.

M. Tomz and M. L. Wright. Empirical Research on Sovereign Debt and Default. *Annu. Rev. Econ.*, 5(1):247–272, 2013.

M. Uribe and S. Schmitt-Grohé. *Open Economy Macroeconomics*. Princeton University Press, 2017.

H. Utar, L. B. Torres Ruiz, and A. C. Zurita. The US-China Trade War and the Relocation of Global Value Chains to Mexico. *CESifo Working Paper*, 2023.

B. C. Waterhouse. "Donald Trump: Foreign Affairs," University of Virginia Miller Center, 2022. URL `https://millercenter.org/president/trump/foreign-affairs`.

E. L. Yeyati and U. Panizza. The Elusive Costs of Sovereign Defaults. *Journal of Development Economics*, 94(1):95–105, 2011.

S. Yu. The Effect of Political Factors on Sovereign Default. *Review of Political Economy*, 28(3): 397–416, 2016.

C. E. Zarazaga. Default and Lost Opportunities: A Message from Argentina for Euro-Zone Countries. *Economic Letter*, 7(4), 2012.

Appendix Contents

Appendix A: Synthetic control - validation and robustness

This appendix reports the validation and robustness exercises summarized in Table 3. Synthetic-control estimates can be threatened by bias, poor pre-treatment fit, or insufficient pre- and post-treatment data, and a credible design must also rule out anticipation and spillovers. I therefore check the requirements for synthetic-control designs in Abadie (2021): (i) an outcome with a large treatment size and limited volatility, (ii) a comparison group with characteristics similar to the treated unit, (iii) no anticipation effect, (iv) no spillovers from the treated unit to the controls, (v) a convex-hull condition (the treated unit's pre-treatment predictors lie within the convex hull of the donor pool's), (vi) enough pre-intervention observations, and (vii) enough post-treatment observations to trace the full effect. Following Abadie and Vives-i-Bastida (2022), I address these through trimming the donor pool, leave-one-out estimators, placebo runs, smoothing the outcome, backdating, and residualizing the outcome, described in turn below.

TABLE A1: CDS predictor means with trimmed donor pool

Variables	Observed	Synthetic	Average
Nominal exchange rate depreciation in 2013	-2.91	0.55	2.75
Nominal exchange rate depreciation in 2014	4.35	2.58	7.53
Nominal exchange rate depreciation in 2015	19.31	19.42	17.09
Log GDP per capita	9.80	9.92	9.68
Log consumption per capita	9.59	9.68	9.44
Log capital stock per capita	11.26	10.93	10.78
CDS at the beginning of 2013	97.23	91.06	108.33
CDS at the beginning of 2014	92.32	104.31	139.21
CDS at the beginning of 2015	103.45	107.92	144.37
CDS at the beginning of 2016	169.55	166.93	202.08

Note: The trimmed donor pool is constructed by keeping countries with a credit rating between BBB-, BBB, and BBB+ during the pre-treatment sample (i.e., Colombia, Kazakhstan, Panama, Peru, Philippines, Poland, South Africa, Thailand, and Uruguay). The nominal exchange rate depreciation corresponds to the average depreciation observed that year. Log GDP per capita, log consumption per capita, and log capital stock per capita correspond to the average using data from 2013 to 2015.

Trimming the donor pool

The method requires a donor pool unaffected by the treatment. According to Waterhouse (2022) and Council on Foreign Relations (2024), Trump's "America First" agenda directly targeted three economies during the campaign and his first year in office, China, Canada, and Mexico; he also had frictions with North Korea, several Middle Eastern countries, and Russia, but there is no evidence of direct trade or immigration pressure on the donor-pool countries.[32] To be safe, I trim the pool to countries with sovereign ratings in the BBB range, which leaves 9 countries (Colombia, Kazakhstan, Panama, Peru, Philippines, Poland, South Africa, Thailand, and Uruguay). The estimation again loads only on Colombia ($w_2^\star = 0.36$), Panama ($w_3^\star = 0.33$), and Poland ($w_4^\star = 0.31$), with pre-treatment characteristics in Table A1.

[32]During the first month of his government, President Trump signed an executive order to ban citizens from Iran, Iraq, Lybia, Somalia, Sudan, Syria, and Yemen. Those countries were already identified as "countries of concern" during President Obama's administration.

Because the weights are close to the baseline, $\hat{\alpha}_t$ is almost identical to the baseline in Figure 1B.

Leave-one-out and spillovers

Figure A1A re-estimates the synthetic control while dropping, one at a time, each donor with positive baseline weight (Colombia, Panama, Poland). Every leave-one-out estimate tracks the Mexican CDS closely before the intervention, and all post-treatment effects are positive and centered on the baseline. A related concern is that the treatment spills over

FIGURE A1. Robustness

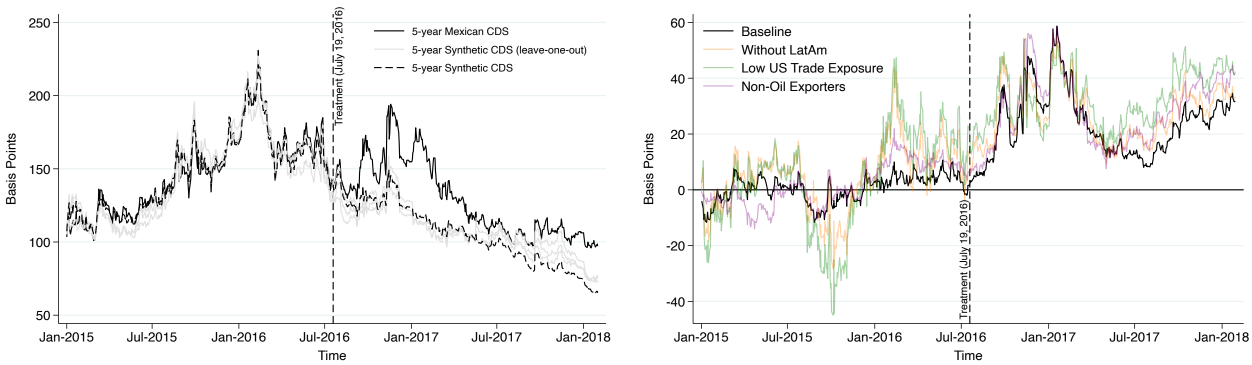

(A) Leave-one-out estimator

(B) Spillovers

Note: On the left, Figure A1A displays the leave-one-out estimator that comes from estimating the synthetic control for Mexico using a donor pool that drops only one of the countries with positive weights in the baseline estimation. On the right, I display the dynamic treatment effect from applying the synthetic control method to a donor pool that excludes LatAm countries (i.e., orange), countries with above-average trade exposure to the US (i.e., green), and oil exporters (i.e., purple), respectively. The estimated weights for each restriction on the donor pool are: (Poland - $w_2^\star = 0.3181$, Kazakhstan - $w_3^\star = 0.0023$, South Africa - $w_4^\star = 0.1196$, and Malaysia -$w_5^\star = 0.5600$), (Malaysia - $w_2^\star = 0.7655$ and Philippines - $w_3^\star = 0.2345$), and (Panama - $w_2^\star = 0.8731$ and Poland - $w_3^\star = 0.1269$).

to the donor pool. If some donors are exposed to the same geoeconomic shocks, through geographic proximity, trade linkages, or common global factors, their sovereign risk co-moves with Mexico's CDS, so the synthetic control partially internalizes the treatment and biases the estimate toward zero. To assess this, I build alternative donor pools that drop the donors most likely to be exposed: (i) all Latin American economies, (ii) countries with above-average export exposure to the United States, and (iii) oil exporters.[33] Figure A1B shows the resulting effects. In every case the pre-treatment fit stays comparable to the baseline while the post-treatment gap is systematically larger, with similar dynamics: the average effect rises from 24 bp in the baseline to 27.44 bp excluding Latin America, 33 bp excluding high U.S.-trade-exposure countries, and 28 bp excluding oil exporters. This is exactly the pattern attenuation from spillovers would produce, since partially treated donors absorb part of

[33]Exposure to the United States is measured as the share of a country's total exports destined to the U.S., using data from the World Bank's World Integrated Trade Solution (WITS), based on UN Comtrade. For each country, I compute the average exposure over the period 2013–2016. Across the sample of countries considered, the mean exposure is approximately 12%. Countries with above-average exposure include Chile, Colombia, Panama, Peru, Poland, South Korea, and Trinidad and Tobago.

the shock. Removing them raises the estimate without changing the timing or shape of the response, so the baseline likely understates the true effect.

Overfitting

Because the CDS is volatile, a natural concern is overfitting to noise. In Figure A2 I smooth the series with 30-, 60-, and 90-day moving averages and re-estimate; the treatment effects are very close to the baseline, which alleviates the concern.[34]

FIGURE A2. Smoothing outcome variables

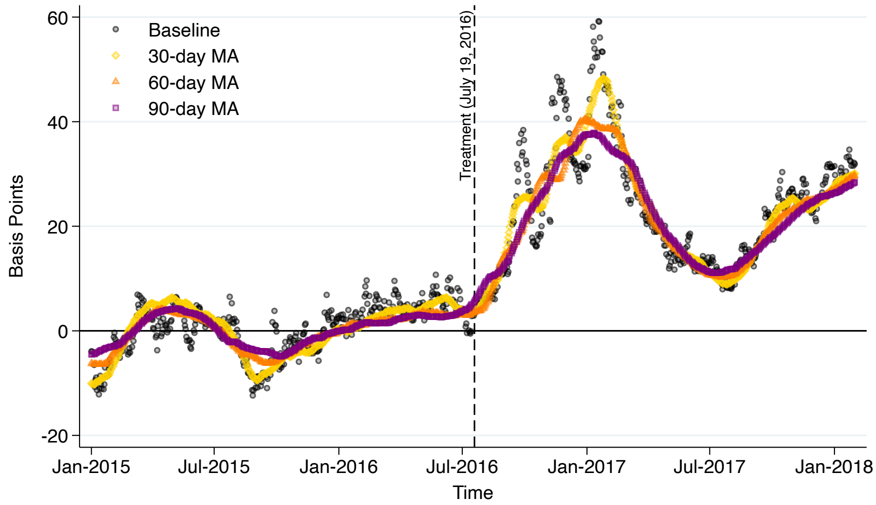

Note: this figure displays the treatment effect from applying the synthetic control method to smoothed data using a 30, 60, and 90 moving average, respectively.

Statistical significance

I assess significance with placebo (permutation) tests following Abadie et al. (2010), moving Mexico into the donor pool and estimating the synthetic control for every other country. The placebo runs measure how rare Mexico's gap is relative to the donor distribution. In Figure A3A, the black line is Mexico's estimated effect and the gray lines are the placebo effects for donors with a pre-intervention MSPE of at most five times Mexico's.[35] Figure A3B plots the treatment distribution over time using end-of-month observations and the same MSPE cutoff, with a black dot for Mexico. Both make clear that the geoeconomic pressure produced a large gap in Mexico's risk premium. Figure A4 reports the post-to-pre-treatment MSPE ratio for Mexico and every donor: Mexico has the highest ratio, so the probability of a gap this large under a random permutation of the intervention date is about 5.3%.

[34]Results are robust to other ways of "filtering" the CDS data. Abadie and Vives-i-Bastida (2022) documents that other sources of overfitting in the context of synthetic controls where the outcome variable is volatile are small pre-treatment samples and a large donor pool. In this application, the pre-treatment sample is large (925

FIGURE A3. Placebo runs and treatment distribution

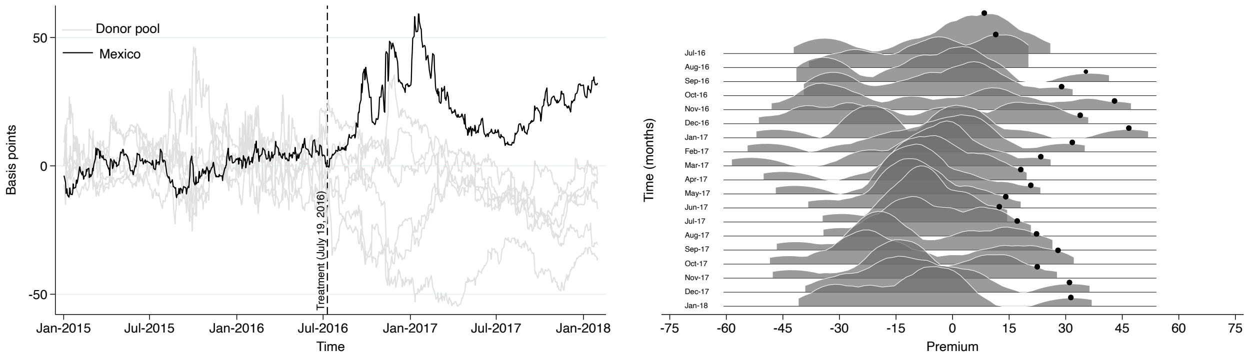

(A) Placebo runs

(B) Treatment distribution

Note: On the left, Figure A3A displays in gray the estimated premiums for the other countries in the donor pool while the superimposed black line is the premium for Mexico. To construct the gray lines, I move Mexico to the donor pool and estimate the synthetic control for all the countries. Following the literature, I only display those countries with a Mean Squared Prediction Error (MSPE) of, at most, κ times the MSPE of Mexico in the pre-treatment sample. I arbitrarily set $\kappa = 5$; however, the results are robust to different cutoff values. On the right, Figure A3B plots the treatment distribution over time using end-of-month observations and the same MSPE cutoff as in the first panel.

FIGURE A4. Post-Pre MSPE ratio

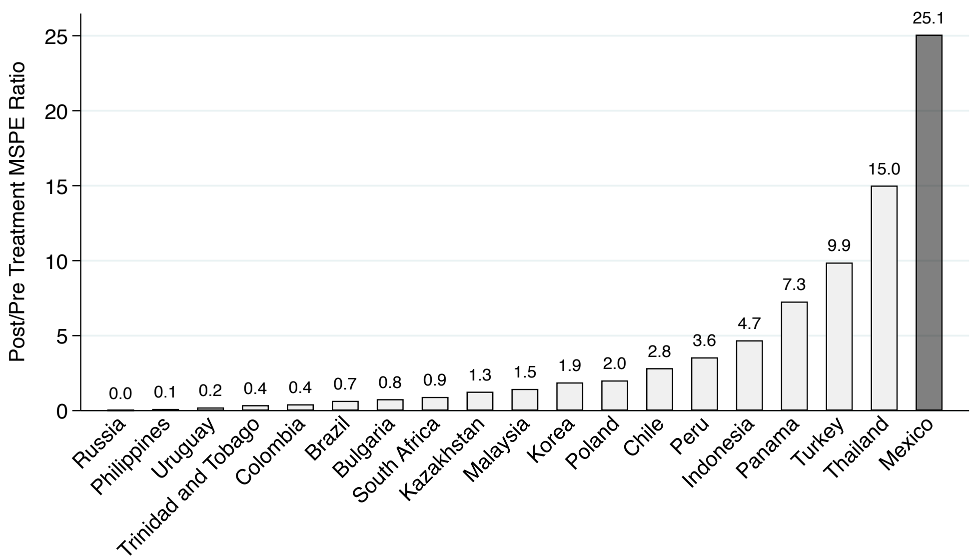

Note: This Figure displays the post-pre treatment MSPE ratio for all countries in the donor pool and Mexico.

Prediction intervals

As an alternative to the permutation test, I follow Cattaneo et al. (2021) and Cattaneo et al. (2025) to construct prediction intervals for $\hat{\alpha}_t$ that rely only on the outcome variable and

observations), and the results are robust to trimming the donor pool.

[35]The MSPE is defined as $MSPE = \frac{1}{T_0}\sum_{t=1}^{T_0}\left(Y_{j,t} - \hat{Y}_{j,t}\right)^2$, where $\hat{Y}_{j,t}$ is the outcome variable produced by the synthetic control for country j in period t.

allow for non-stationarity.[36] Figure A5A plots the Mexican CDS and the synthetic counter-

FIGURE A5. Prediction intervals

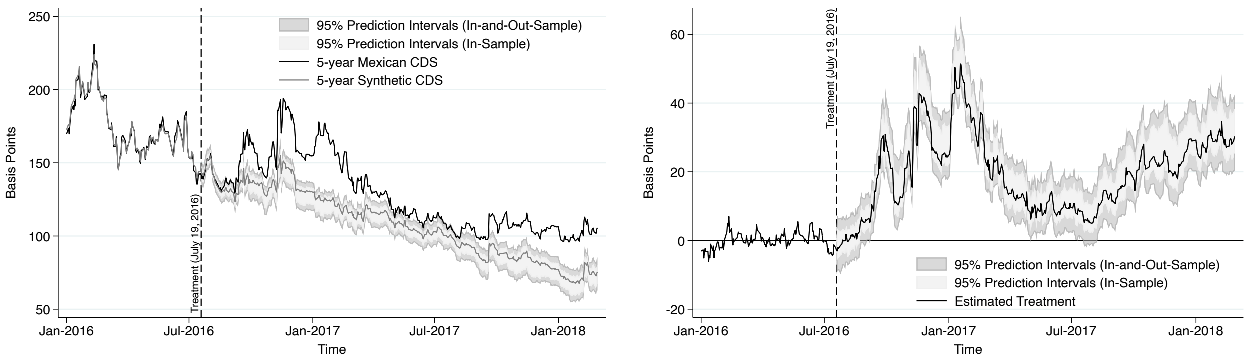

(A) Observed and synthetic CDS

(B) Treatment effect

Note: On the left, Figure A5A displays the daily evolution of the outcome variable and the synthetic control along with the 95% prediction intervals. On the right, Figure A5B presents the treatment effect and the 95% prediction intervals taking into account the in-sample uncertainty and the in-sample and out-of-sample uncertainty.

factual with their 95% prediction intervals, and Figure A5B plots the treatment effect with its own. The observed CDS sits above the synthetic control's prediction interval for most of the post-treatment sample, again pointing to a significant effect.

Backdating

I reassign the treatment to dates up to six months before Trump's nomination and re-estimate. Figure A6 shows that the out-of-sample fit is good over each hold-out window, that the gap opens only after the actual nomination, and that the magnitudes are very similar to the baseline. This is what one expects if the nomination, not a pre-existing trend, drives the effect.

Serial correlation

To address serial dependence in the daily CDS, I run a time-based placebo that randomly assigns treatment dates within the pre-treatment period and recomputes the synthetic control for each draw, building the distribution of effects over 20-, 40-, 60-, and 120-trading-day horizons. Figure A7 shows that the true effect (red dashed line) sits in the right tail at every horizon. At short horizons it is not always the most extreme draw, but it becomes more pronounced as the horizon lengthens, consistent with a gradual buildup. The effect is therefore unlikely to be an artifact of serial correlation or spurious persistence.

[36] To compute the prediction intervals for computational reasons, I re-estimated the synthetic control using a smaller pre-intervention sample, which spans from January 1, 2016, to July 18, 2016. This analysis relies solely on the outcome variable, which is the CDS data, without including any other predictors. The countries assigned strictly positive weights are similar to those in the baseline estimation: Panama ($w_2^\star = 0.46$), Colombia ($w_3^\star = 0.25$), and Poland ($w_4^\star = 0.22$). However, Turkey ($w_5^\star = 0.04$), South Korea ($w_6^\star = 0.02$), and South Africa ($w_7^\star = 0.01$) now contribute to replicating the pre-intervention dynamics of the outcome variable.

FIGURE A6. Backdating

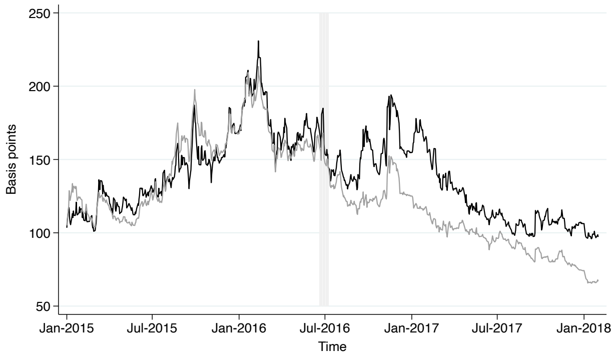

(A) 1 month

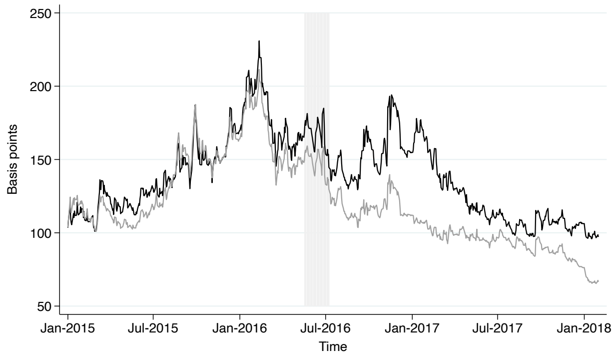

(B) 2 months

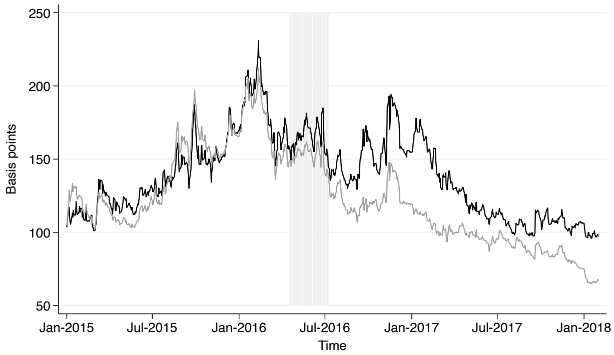

(C) 3 months

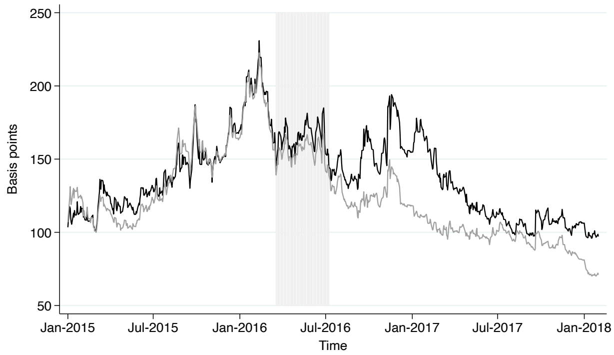

(D) 4 months

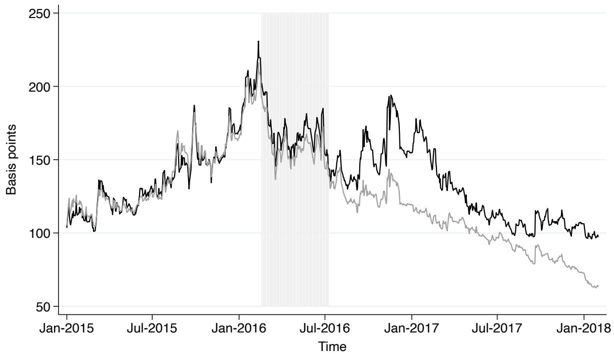

(E) 5 months

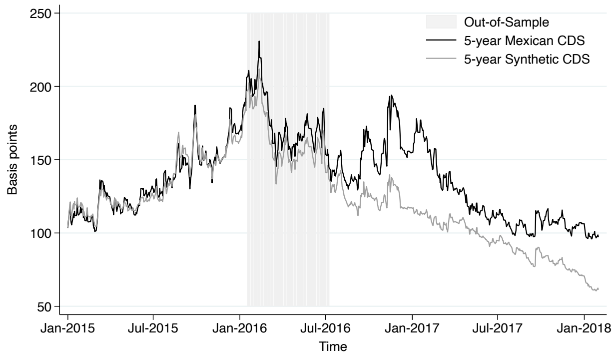

(F) 6 months

Note: I reassigned the treatment up to six months before the nomination of Donald Trump as the official candidate of the Republican Party. I get the same set of countries to build the synthetic 5-year CDS with almost identical weights.

U.S. monetary policy

In November 2015 the Federal Reserve began raising its policy rate, from 0.11% in January 2015 to 1.42% in February 2018. Since emerging-market spreads are sensitive to U.S. monetary policy, this could contaminate the estimate. To check, I estimate, for each country,

$$CDS_{i,t} = \alpha_i + \beta_i \Delta \text{Fed}_t + \eta_{i,t}, \tag{1}$$

FIGURE A7. Size of average $\hat{\alpha}_t$ relative to alternative/random gaps

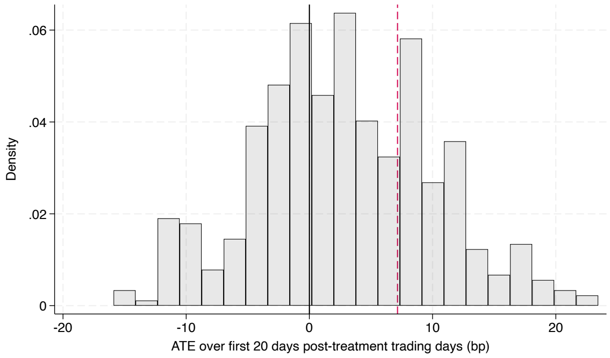

(A) 1 month

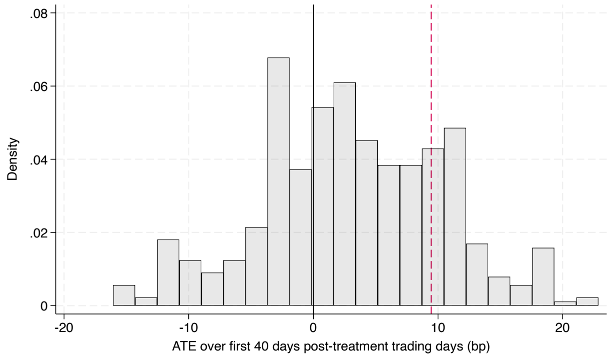

(B) 2 months

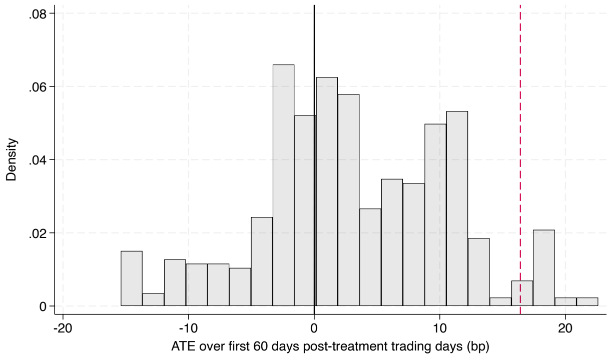

(C) 3 months

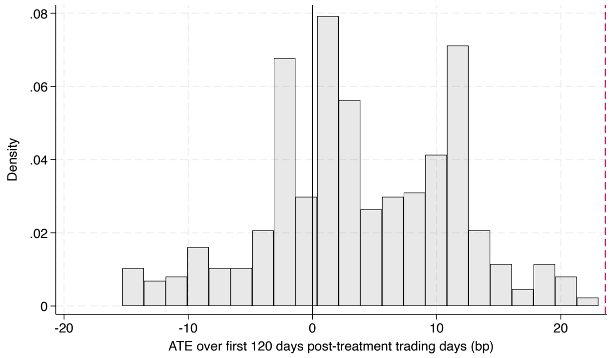

(D) 6 months

Note: This displays the histograms associated with the average treatment effects for 20, 40, 60, and 90 days, respectively, from the placebo runs using 500 simulations for each. The red-dashed line corresponds the ATE for the same number of days using the true treatment date.

where $CDS_{i,t}$ is the 5-year CDS, ΔFed_t the daily change in the Fed funds rate, and $\eta_{i,t}$ the error term, and I re-run the synthetic control on the residualized series $\hat{\eta}_{i,t}$ (the CDS variation left unexplained by U.S. policy) in place of the raw CDS, residualizing Mexico and the 18 donor economies alike. Figure A8 shows the premium on residualized data is practically identical to the baseline.[37]

[37] Since all countries were exposed to such changes in the US monetary policy, the effect would cancel out when taking the difference between the observed and synthetic outcome variable, as the synthetic unit is mathematically constructed to match Mexico's unobserved factor loadings (i.e., its specific heterogeneous exposure to the shock).

FIGURE A8. Residualizing CDS from changes in US Monetary Policy

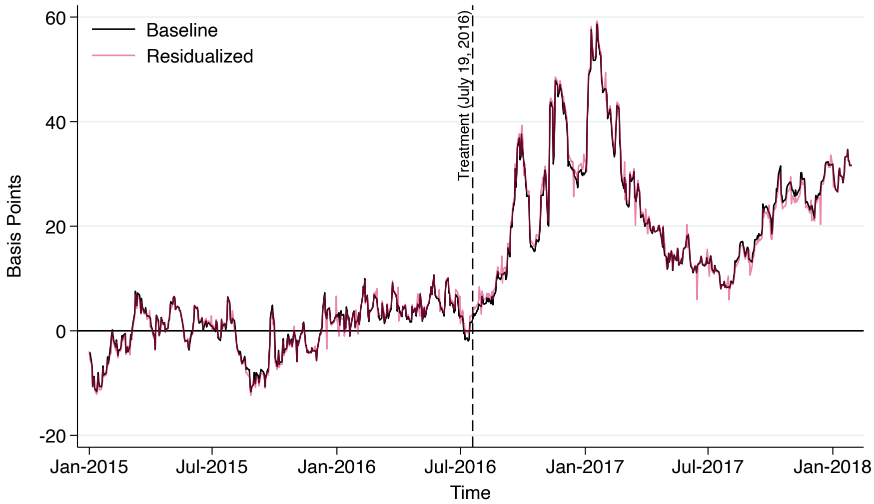

Note: I residualize the CDS data for each country before estimating the synthetic control using the model: $CDS_{i,t} = \alpha_i + \beta_i \Delta \text{Fed}_t + \eta_{i,t}$, where $CDS_{i,t}$ is the 5-year CDS for country i at time t, and ΔFed_t is the daily change in the Fed funds rate. For the synthetic control estimation, I use $\hat{\eta}_{i,t}$ instead of $CDS_{i,t}$. The black line shows the baseline estimate, and the red line displays the estimates using the residualized data.

Appendix B: President Trump Tweets

In the following Table B1, the Tweets made by President Trump during the post-treatment sample for the exercise regarding the sensitivity of the premium are displayed.

Date	Timestamp	Content	Classification	Id
07/30/16	10:12 PM	"The "Rust Belt" was created by politicians like the Clintons who allowed our jobs to be stolen from us by other countries like Mexico. END!"	Negative	T1
08/31/16	1:07 PM	"Former President Vicente Fox, who is railing against my visit to Mexico today, also invited me when he apologized for using the "f-bomb"."	Neutral	
09/01/16	12:43 AM	"Great trip to Mexico today - wonderful leadership and high-quality people! Look forward to our next meeting."	Neutral	
09/01/16	10:31 AM	"Mexico will pay for the wall!"	Negative	T2
11/18/16	2:01 AM	"Just got a call from my friend Bill Ford, Chairman of Ford, who advised me that he will be keeping the Lincoln plant in Kentucky - no Mexico."	Negative	T3
12/03/16	3:06 AM	"Rexnord of Indiana is moving to Mexico and rather viciously firing all of its 300 workers. This is happening all over our country. No more!"	Negative	T4
12/20/16	8:27 PM	"Yes, it is true - Carlos Slim, the great businessman from Mexico, called me about getting together for a meeting. We met, HE IS A GREAT GUY!"	Neutral	
01/04/17	1:19 PM	"Thank you to Ford for scrapping a new plant in Mexico and creating 700 new jobs in the U.S. This is just the beginning - much more to follow."	Negative	T5
01/05/17	6:14 PM	"Toyota Motor said will build a new plant in Baja, Mexico, to build Corolla cars for U.S. NO WAY! Build plant in U.S. or pay big border tax."	Negative	T6
01/06/17	11:19 AM	"The dishonest media does not report that any money spent on building the Great Wall (for sake of speed), will be paid back by Mexico later!"	Negative	T6
01/09/17	4:05 AM	"Dishonest media says Mexico won't be paying for the wall if they pay a little later so the wall can be built more quickly. Media is fake!"	Neutral	

01/09/17	2:16 PM	"Ford said last week that it will expand in Michigan and U.S. instead of building a BILLION dollar plant in Mexico. Thank you Ford & Fiat C!"	Negative	T7
01/26/17	1:51 PM	"The U.S. has a 60 billion dollar trade deficit with Mexico. It has been a one-sided deal from the beginning of NAFTA with massive numbers..."	Negative	T8
01/26/17	1:55 PM	"...of jobs and companies lost. If Mexico is unwilling to pay for the badly needed wall, then it would be better to cancel the upcoming meeting."	Negative	T8
01/27/17	1:19 PM	"Mexico has taken advantage of the U.S. for long enough. Massive trade deficits & little help on the very weak border must change, NOW!"	Negative	T9
04/23/17	11:44 AM	"Eventually, but at a later date so we can get started early, Mexico will be paying, in some form, for the badly needed border wall."	Negative	T10
04/27/17	7:12 AM	"I received calls from the President of Mexico and the Prime Minister of Canada asking to renegotiate NAFTA rather than terminate. I agreed."	Positive	T11
05/07/17	6:58 PM	"Rexnord of Indiana made a deal during the Obama Administration to move to Mexico. Fired their employees. Tax product big that's sold in U.S."	Negative	T12
06/22/17	6:15 PM	"Mexico was just ranked the second deadliest country in the world, after only Syria. Drug trade is largely the cause. We will BUILD THE WALL!"	Negative	T13
06/29/17	8:27 AM	"New Sugar deal negotiated with Mexico is a very good one for both Mexico and the U.S. Had no deal for many years which hurt U.S. badly."	Positive	T14
08/27/17	9:44 AM	"With Mexico being one of the highest crime Nations in the world, we must have THE WALL. Mexico will pay for it through reimbursement/other."	Negative	T15
08/27/17	9:51 AM	"We are in the NAFTA (worst trade deal ever made) renegotiation process with Mexico & Canada. Both being very difficult, may have to terminate?"	Negative	T15
09/19/17	4:05 PM	"God bless the people of Mexico City. We are with you and will be there for you."	Neutral	

01/11/18	9:49 PM	"More great news as a result of historical Tax Cuts and Reform: Fiat Chrysler announces plan to invest more than $1 BILLION in Michigan plant, relocating their heavy-truck production from Mexico to Michigan, adding 2,500 new jobs and paying $2000 bonus to U.S. employees!"	Negative	T16
01/11/18	9:53 PM	"Chrysler is moving a massive plant from Mexico to Michigan, reversing a years long opposite trend. Thank you Chrysler, a very wise decision. The voters in Michigan are very happy they voted for Trump/Pence. Plenty of more to follow!"	Negative	T16
01/18/18	6:25 AM	"The Wall will be paid for, directly or indirectly, or through longer term reimbursement, by Mexico, which has a ridiculous $71 billion dollar trade surplus with the U.S. The $20 billion dollar Wall is "peanuts" compared to what Mexico makes from the U.S. NAFTA is a bad joke!"	Negative	T17

TABLE B1: Tweets by Donald Trump regarding Mexico

Date	Timestamp	Content	Id
09/06/16	7:12 AM	"China wouldn't provide a red carpet stairway from Air Force One and then Philippines President calls Obama "the son of a whore." Terrible!"	$T1^{China}$
11/16/16	7:17 AM	"I have received and taken calls from many foreign leaders despite what the failing @nytimes said. Russia, U.K., China, Saudi Arabia, Japan, ..."	$T2^{China}$
12/04/16	5:23 PM	"Did China ask us if it was OK to devalue their currency (making it hard for our companies to compete), heavily tax our products going into their country (the U.S. doesn't tax them) or to build a massive military complex in the middle of the South China Sea? I don't think so!"	$T3^{China}$
12/17/16	8:57 PM	"China steals United States Navy research drone in international waters - rips it out of water and takes it to China in unprecedented act. We should tell China that we don't want the drone they stole back.- let them keep it!"	$T4^{China}$
01/02/17	6:47 PM	"China has been taking out massive amounts of money & wealth from the U.S. in totally one-sided trade, but won't help with North Korea. Nice!"	$T5^{China}$
02/10/17	8:35 AM	"The failing @nytimes does major FAKE NEWS China story saying "Mr.Xi has not spoken to Mr. Trump since Nov.14." We spoke at length yesterday!"	$T6^{China}$

03/17/17	9:07 AM	"North Korea is behaving very badly. They have been "playing" the United States for years. China has done little to help!"	$T7^{China}$
03/30/17	6:16 PM	"The meeting next week with China will be a very difficult one in that we can no longer have massive trade deficits. . ."	$T8^{China}$
04/08/17	10:50 AM	"It was a great honor to have President Xi Jinping and Madame Peng Liyuan of China as our guests in the United States. Tremendous. . ."	$T9^{China}$
04/11/17	7:59 AM	"I explained to the President of China that a trade deal with the U.S. will be far better for them if they solve the North Korean problem!"	$T10^{China}$
04/11/17	8:03 AM	"North Korea is looking for trouble. If China decides to help, that would be great. If not, we will solve the problem without them! U.S.A."	$T10^{China}$
04/12/17	8:22 AM	"Had a very good call last night with the President of China concerning the menace of North Korea."	$T11^{China}$
04/13/17	9:08 AM	"I have great confidence that China will properly deal with North Korea. If they are unable to do so, the U.S., with its allies, will! U.S.A."	$T12^{China}$
04/16/17	8:18 AM	"Why would I call China a currency manipulator when they are working with us on the North Korean problem? We will see what happens!"	$T13^{China}$
04/21/17	9:04 AM	"China is very much the economic lifeline to North Korea so, while nothing is easy, if they want to solve the North Korean problem, they will"	$T14^{China}$
04/28/17	7:26 PM	"North Korea disrespected the wishes of China & its highly respected President when it launched, though unsuccessfully, a missile today. Bad!"	$T15^{China}$
05/12/17	9:20 AM	"China just agreed that the U.S. will be allowed to sell beef, and other major products, into China once again. This is REAL news!"	$T16^{China}$
05/29/17	8:18 AM	"North Korea has shown great disrespect for their neighbor, China, by shooting off yet another ballistic missile...but China is trying hard!"	$T17^{China}$
06/20/17	2:38 PM	"While I greatly appreciate the efforts of President Xi & China to help with North Korea, it has not worked out. At least I know China tried!"	$T18^{China}$
07/03/17	10:24 PM	"...and Japan will put up with this much longer. Perhaps China will put a heavy move on North Korea and end this nonsense once and for all!"	$T19^{China}$
07/05/17	7:21 AM	"Trade between China and North Korea grew almost 40% in the first quarter. So much for China working with us - but we had to give it a try!"	$T20^{China}$

07/29/17	7:29 PM	"I am very disappointed in China. Our foolish past leaders have allowed them to make hundreds of billions of dollars a year in trade, yet... they do NOTHING for us with North Korea, just talk. We will no longer allow this to continue. China could easily solve this problem!"	T21$^{\text{China}}$
08/05/17	6:44 PM	"The United Nations Security Council just voted 15-0 to sanction North Korea. China and Russia voted with us. Very big financial impact!"	T22$^{\text{China}}$
09/03/17	7:39 AM	"..North Korea is a rogue nation which has become a great threat and embarrassment to China, which is trying to help but with little success."	T23$^{\text{China}}$
09/13/17	10:22 PM	"China has a business tax rate of 15%. We should do everything possible to match them in order to win with our economy. Jobs and wages!	T24$^{\text{China}}$
10/25/17	3:24 PM	"Spoke to President Xi of China to congratulate him on his extraordinary elevation. Also discussed NoKo & trade, two very important subjects! Melania and I look forward to being with President Xi, Madame Peng Liyuan in China in two weeks for what will hopefully be a historic trip!"	T25$^{\text{China}}$
11/07/17	8:00 PM	"Getting ready to make a major speech to the National Assembly here in South Korea, then will be headed to China where I very much look forward to meeting with President Xi who is just off his great political victory."	T26$^{\text{China}}$
11/08/17	10:40 AM	"Looking forward to a full day of meetings with President Xi and our delegations tomorrow. THANK YOU for the beautiful welcome China! @FLOTUS Melania and I will never forget it!"	T26$^{\text{China}}$
11/09/17	8:58 AM	"In the coming months and years ahead I look forward to building an even STRONGER relationship between the United States and China."	T27$^{\text{China}}$
11/09/17	6:39 PM	"I don't blame China, I blame the incompetence of past Admins for allowing China to take advantage of the U.S. on trade leading up to a point where the U.S. is losing $100's of billions. How can you blame China for taking advantage of people that had no clue? I would've done same!"	T28$^{\text{China}}$
11/09/17	8:17 PM	"I am leaving China for #APEC2017 in Vietnam. @FLOTUS Melania is staying behind to see the zoo, and of course, the Great WALL of China before going to Alaska to greet our AMAZING troops."	T28$^{\text{China}}$
11/11/17	6:32 PM	"President Xi of China has stated that he is upping the sanctions against #NoKo. Said he wants them to denuclearize. Progress is being made."	T29$^{\text{China}}$
11/15/17	11:35 AM	"The failing @nytimes hates the fact that I have developed a great relationship with World leaders like Xi Jinping, President of China..."	T30$^{\text{China}}$

11/29/17	9:40 AM	"Just spoke to President XI JINPING of China concerning the provocative actions of North Korea. Additional major sanctions will be imposed on North Korea today. This situation will be handled!	$T31^{\text{China}}$
11/30/17	7:25 AM	"The Chinese Envoy, who just returned from North Korea, seems to have had no impact on Little Rocket Man. Hard to believe his people, and the military, put up with living in such horrible conditions. Russia and China condemned the launch."	$T32^{\text{China}}$
12/28/17	11:24 AM	"Caught RED HANDED - very disappointed that China is allowing oil to go into North Korea. There will never be a friendly solution to the North Korea problem if this continues to happen!"	$T33^{\text{China}}$

TABLE B2: Tweets by Donald Trump regarding China

Date	Timestamp	Content	Id
07/25/16	7:31 AM	"The new joke in town is that Russia leaked the disastrous DNC e-mails, which should never have been written (stupid), because Putin likes me"	$T1^{\text{Russia}}$
07/26/16	6:47 PM	"In order to try and deflect the horror and stupidity of the Wikileakes disaster, the Dems said maybe it is Russia dealing with Trump. Crazy!"	$T2^{\text{Russia}}$
07/26/16	6:50 PM	"For the record, I have ZERO investments in Russia."	$T2^{\text{Russia}}$
07/27/16	5:43 AM	"Funny how the failing @nytimes is pushing Dems narrative that Russia is working for me because Putin said "Trump is a genius." America 1st!"	$T2^{\text{Russia}}$
07/27/17	12:16 PM	"If Russia or any other country or person has Hillary Clinton's 33,000 illegally deleted emails, perhaps they should share them with the FBI!"	$T2^{\text{Russia}}$
08/01/16	9:03 AM	"So with all of the Obama tough talk on Russia and the Ukraine, they have already taken Crimea and continue to push. That's what I said!"	$T3^{\text{Russia}}$
09/04/16	12:04 PM	"Crooked Hillary's V.P. pick said this morning that I was not aware that Russia took over Crimea. A total lie - and taken over during O term!"	$T4^{\text{Russia}}$
09/13/16	11:21 PM	"Russia took Crimea during the so-called Obama years. Who wouldn't know this and why does Obama get a free pass?"	$T5^{\text{Russia}}$
09/26/16	10:28 PM	"Russia has more warheads than ever, N Korea is testing nukes, and Iran got a sweetheart deal to keep theirs. Thanks, @HillaryClinton."	$T6^{\text{Russia}}$
11/16/16	7:17 AM	"I have received and taken calls from many foreign leaders despite what the failing @nytimes said. Russia, U.K., China, Saudi Arabia, Japan,. . ."	$T7^{\text{Russia}}$
12/12/16	8:17 AM	"Can you imagine if the election results were the opposite and WE tried to play the Russia/CIA card. It would be called conspiracy theory!"	$T8^{\text{Russia}}$

12/15/16	9:24 AM	“If Russia, or some other entity, was hacking, why did the White House wait so long to act? Why did they only complain after Hillary lost?”	T9[Russia]
01/05/17	7:30 PM	“The Democratic National Committee would not allow the FBI to study or see its computer info after it was supposedly hacked by Russia”	T10[Russia]
01/07/17	10:02 AM	“Having a good relationship with Russia is a good thing, not a bad thing. Only ”stupid” people, or fools, would think that it is bad! We have enough problems around the world without yet another one. When I am President, Russia will respect us far more than they do now and...”	T11[Russia]
01/11/17	7:13 AM	“Russia just said the unverified report paid for by political opponents is ”A COMPLETE AND TOTAL FABRICATION, UTTER NONSENSE.” Very unfair!”	T12[Russia]
01/11/17	7:31 AM	“Russia has never tried to use leverage over me. I HAVE NOTHING TO DO WITH RUSSIA - NO DEALS, NO LOANS, NO NOTHING!”	T12[Russia]
01/13/17	6:11 AM	“Totally made up facts by sleazebag political operatives, both Democrats and Republicans - FAKE NEWS! Russia says nothing exists. Probably...”	T13[Russia]
02/07/17	7:11 AM	“I don’t know Putin, have no deals in Russia, and the haters are going crazy - yet Obama can make a deal with Iran, #1 in terror, no problem!”	T14[Russia]
02/15/17	7:19 AM	“Information is being illegally given to the failing @nytimes & @washingtonpost by the intelligence community (NSA and FBI?).Just like Russia”	T15[Russia]
02/15/17	7:42 AM	“Crimea was TAKEN by Russia during the Obama Administration. Was Obama too soft on Russia?”	T15[Russia]
02/16/17	9:39 AM	“The Democrats had to come up with a story as to why they lost the election, and so badly (306), so they made up a story - RUSSIA. Fake news!”	T16[Russia]
02/26/17	1:16 PM	“Russia talk is FAKE NEWS put out by the Dems, and played up by the media, in order to mask the big election defeat and the illegal leaks!”	T17[Russia]
03/03/17	12:54 PM	“We should start an immediate investigation into @SenSchumer and his ties to Russia and Putin. A total hypocrite!”	T18[Russia]
03/03/17	4:02 PM	“I hereby demand a second investigation, after Schumer, of Pelosi for her close ties to Russia, and lying about it.”	T19[Russia]
03/07/17	8:13 AM	“For eight years Russia “ran over” President Obama, got stronger and stronger, picked-off Crimea and added missiles. Weak! @foxandfriends”	T20[Russia]
03/20/17	6:35 AM	“James Clapper and others stated that there is no evidence Potus colluded with Russia. This story is FAKE NEWS and everyone knows it!”	T21[Russia]

03/23/17	8:18 AM	"Just watched the totally biased and fake news reports of the so-called Russia story on NBC and ABC. Such dishonesty!"	$T22^{\text{Russia}}$
03/27/17	9:26 PM	"Why isn't the House Intelligence Committee looking into the Bill & Hillary deal that allowed big Uranium to go to Russia, Russian speech money to Bill, the Hillary Russian "reset," praise of Russia by Hillary, or Podesta Russian Company. Trump Russia story is a hoax. #MAGA!	$T23^{\text{Russia}}$
03/28/17	7:16 AM	"Watch @foxandfriends now on Podesta and Russia!"	$T23^{\text{Russia}}$
03/28/17	6:41 PM	"Why doesn't Fake News talk about Podesta ties to Russia as covered by @FoxNews or money from Russia to Clinton - sale of Uranium?"	$T24^{\text{Russia}}$
04/01/17	8:43 AM	"When will Sleepy Eyes Chuck Todd and @NBCNews start talking about the Obama SURVEILLANCE SCANDAL and stop with the Fake Trump/Russia story?"	$T25^{\text{Russia}}$
04/01/17	9:02 AM	"It is the same Fake News Media that said there is "no path to victory for Trump" that is now pushing the phony Russia story. A total scam!"	$T25^{\text{Russia}}$
04/01/17	1:02 PM	"Not associated with Russia. Trump team spied on before he was nominated." If this is true, does not get much bigger. Would be sad for U.S.	$T25^{\text{Russia}}$
04/03/17	7:16 AM	"Was the brother of John Podesta paid big money to get the sanctions on Russia lifted? Did Hillary know?"	$T25^{\text{Russia}}$
04/13/17	9:16 AM	"Things will work out fine between the U.S.A. and Russia. At the right time everyone will come to their senses & there will be lasting peace!"	$T26^{\text{Russia}}$
05/02/17	11:06 PM	"...Trump/Russia story was an excuse used by the Democrats as justification for losing the election. Perhaps Trump just ran a great campaign?"	$T27^{\text{Russia}}$
05/07/17	7:15 AM	"When will the Fake Media ask about the Dems dealings with Russia & why the DNC wouldn't allow the FBI to check their server or investigate?"	$T28^{\text{Russia}}$
05/08/17	6:41 PM	"Director Clapper reiterated what everybody, including the fake media already knows- there is "no evidence" of collusion w/ Russia and Trump."	$T29^{\text{Russia}}$
05/08/17	6:46 PM	"The Russia-Trump collusion story is a total hoax, when will this taxpayer funded charade end?"	$T29^{\text{Russia}}$
05/11/17	4:34 PM	"Russia must be laughing up their sleeves watching as the U.S. tears itself apart over a Democrat EXCUSE for losing the election."	$T30^{\text{Russia}}$
05/16/17	7:03 AM	"As President I wanted to share with Russia (at an openly scheduled W.H. meeting) which I have the absolute right to do, facts pertaining..."	$T31^{\text{Russia}}$
05/16/17	7:13 AM	"...to terrorism and airline flight safety. Humanitarian reasons, plus I want Russia to greatly step up their fight against ISIS & terrorism"	$T31^{\text{Russia}}$

05/31/17	6:37 AM	"So now it is reported that the Democrats, who have excoriated Carter Page about Russia, don't want him to testify. He blows away their..."	T32$^{\text{Russia}}$
06/15/17	3:43 PM	"Why is that Hillary Clintons family and Dems dealings with Russia are not looked at, but my non-dealings are?"	T33$^{\text{Russia}}$
06/22/17	9:18 AM	"Former Homeland Security Advisor Jeh Johnson is latest top intelligence official to state there was no grand scheme between Trump & Russia.	T34$^{\text{Russia}}$
06/22/17	9:22 AM	"By the way, if Russia was working so hard on the 2016 Election, it all took place during the Obama Admin. Why didn't they stop them?"	T34$^{\text{Russia}}$
06/23/17	8:43 PM	"Just out: The Obama Administration knew far in advance of November 8th about election meddling by Russia. Did nothing about it. WHY?"	T35$^{\text{Russia}}$
06/26/17	8:37 AM	"The reason that President Obama did NOTHING about Russia after being notified by the CIA of meddling is that he expected Clinton would win.."	T35$^{\text{Russia}}$
06/26/17	8:59 AM	"The real story is that President Obama did NOTHING after being informed in August about Russian meddling. With 4 months looking at Russia."	T35$^{\text{Russia}}$
06/27/17	6:33 AM	"Wow, CNN had to retract big story on "Russia," with 3 employees forced to resign. What about all the other phony stories they do? FAKE NEWS!	T36$^{\text{Russia}}$
07/09/17	7:37 AM	"We negotiated a ceasefire in parts of Syria which will save lives. Now it is time to move forward in working constructively with Russia!"	T37$^{\text{Russia}}$
07/22/17	7:47 AM	"...What about all of the Clinton ties to Russia, including Podesta Company, Uranium deal, Russian Reset, big dollar speeches etc."	T38$^{\text{Russia}}$
07/24/17	6:52 AM	"After 1 year of investigation with Zero evidence being found, Chuck Schumer just stated that "Democrats should blame ourselves,not Russia."	T38$^{\text{Russia}}$
07/24/17	8:49 AM	"So why aren't the Committees and investigators, and of course our beleaguered A.G., looking into Crooked Hillarys crimes & Russia relations?	T38$^{\text{Russia}}$
07/24/17	9:12 AM	"Sleazy Adam Schiff, the totally biased Congressman looking into "Russia," spends all of his time on television pushing the Dem loss excuse!	T38$^{\text{Russia}}$
07/29/17	7:07 AM	"In other words, Russia was against Trump in the 2016 Election - and why not, I want strong military & low oil prices. Witch Hunt!"	T39$^{\text{Russia}}$
08/03/17	8:18 AM	"Our relationship with Russia is at an all-time & very dangerous low. You can thank Congress, the same people that can't even give us HCare!	T40$^{\text{Russia}}$
08/05/17	6:44 PM	"The United Nations Security Council just voted 15-0 to sanction North Korea. China and Russia voted with us. Very big financial impact!"	T41$^{\text{Russia}}$

08/15/17	6:54 AM	“According to report just out, President Obama knew about Russian interference 3 years ago but he didn’t want to anger Russia!”	T42$^{\text{Russia}}$
09/22/17	6:44 AM	“The Russia hoax continues, now it’s ads on Facebook. What about the totally biased and dishonest Media coverage in favor of Crooked Hillary?”	T43$^{\text{Russia}}$
10/19/17	7:17 AM	“Uranium deal to Russia, with Clinton help and Obama Administration knowledge, is the biggest story that Fake Media doesn’t want to follow!”	T44$^{\text{Russia}}$
10/19/17	7:56 AM	“Workers of firm involved with the discredited and Fake Dossier take the 5th. Who paid for it, Russia, the FBI or the Dems (or all)?”	T44$^{\text{Russia}}$
10/27/17	9:33 AM	“It is now commonly agreed, after many months of COSTLY looking, that there was NO collusion between Russia and Trump. Was collusion with HC!”	T45$^{\text{Russia}}$
10/27/17	10:13 PM	“WHAT HAPPENED” “How Team Hillary played the press for fools on Russia””	T46$^{\text{Russia}}$
10/29/17	10:02 AM	“...the Uranium to Russia deal, the 33,000 plus deleted Emails, the Comey fix and so much more. Instead they look at phony Trump/Russia,...”	T46$^{\text{Russia}}$
10/29/17	10:48 AM	“All of this ”Russia” talk right when the Republicans are making their big push for historic Tax Cuts & Reform. Is this coincidental? NOT!”	T46$^{\text{Russia}}$
11/11/17	7:16 PM	“Met with President Putin of Russia who was at #APEC meetings. Good discussions on Syria. Hope for his help to solve, along with China the dangerous North Korea crisis. Progress being made.”	T47$^{\text{Russia}}$
11/11/17	7:18 PM	“When will all the haters and fools out there realize that having a good relationship with Russia is a good thing, not a bad thing. There always playing politics - bad for our country. I want to solve North Korea, Syria, Ukraine, terrorism, and Russia can greatly help!”	T47$^{\text{Russia}}$
11/11/17	7:43 PM	“Does the Fake News Media remember when Crooked Hillary Clinton, as Secretary of State, was begging Russia to be our friend with the misspelled reset button? Obama tried also, but he had zero chemistry with Putin.”	T47$^{\text{Russia}}$
11/26/17	4:29 PM	“Since the first day I took office, all you hear is the phony Democrat excuse for losing the election, Russia, Russia,Russia. Despite this I have the economy booming and have possibly done more than any 10 month President. MAKE AMERICA GREAT AGAIN!”	T48$^{\text{Russia}}$
11/30/17	7:25 AM	“The Chinese Envoy, who just returned from North Korea, seems to have had no impact on Little Rocket Man. Hard to believe his people, and the military, put up with living in such horrible conditions. Russia and China condemned the launch.”	T49$^{\text{Russia}}$

12/02/17	9:22 PM	“Congratulations to ABC News for suspending Brian Ross for his horrendously inaccurate and dishonest report on the Russia, Russia, Russia Witch Hunt. More Networks and “papers” should do the same with their Fake News!”	T50[Russia]
12/12/17	7:10 AM	“Despite thousands of hours wasted and many millions of dollars spent, the Democrats have been unable to show any collusion with Russia - so now they are moving on to the false accusations and fabricated stories of women who I don’t know and/or have never met. FAKE NEWS!”	T51[Russia]
12/26/17	8:24 AM	“WOW, @foxandfrlends “Dossier is bogus. Clinton Campaign, DNC funded Dossier. FBI CANNOT (after all of this time) VERIFY CLAIMS IN DOSSIER OF RUSSIA/TRUMP COLLUSION. FBI TAINTED.” And they used this Crooked Hillary pile of garbage as the basis for going after the Trump Campaign!”	T52[Russia]
12/29/17	7:46 AM	“While the Fake News loves to talk about my so-called low approval rating, @foxandfriends just showed that my rating on Dec. 28, 2017, was approximately the same as President Obama on Dec. 28, 2009, which was 47%...and this despite massive negative Trump coverage & Russia hoax!	T53[Russia]
01/05/18	9:32 AM	“Well, now that collusion with Russia is proving to be a total hoax and the only collusion is with Hillary Clinton and the FBI/Russia, the Fake News Media (Mainstream) and this phony new book are hitting out at every new front imaginable. They should try winning an election. Sad!”	T54[Russia]
01/10/18	10:00 AM	“The fact that Sneaky Dianne Feinstein, who has on numerous occasions stated that collusion between Trump/Russia has not been found, would release testimony in such an underhanded and possibly illegal way, totally without authorization, is a disgrace. Must have tough Primary!”	T55[Russia]
01/10/18	10:14 AM	“The single greatest Witch Hunt in American history continues. There was no collusion, everybody including the Dems knows there was no collusion, & yet on and on it goes. Russia & the world is laughing at the stupidity they are witnessing. Republicans should finally take control!”	T55[Russia]
02/03/18	7:26 PM	“Great jobs numbers and finally, after many years, rising wages- and nobody even talks about them. Only Russia, Russia, Russia, despite the fact that, after a year of looking, there is No Collusion!”	T56[Russia]
02/16/18	3:18 PM	“Russia started their anti-US campaign in 2014, long before I announced that I would run for President. The results of the election were not impacted. The Trump campaign did nothing wrong - no collusion!”	T57[Russia]

02/17/18	11:22 PM	“General McMaster forgot to say that the results of the 2016 election were not impacted or changed by the Russians and that the only Collusion was between Russia and Crooked H, the DNC and the Dems. Remember the Dirty Dossier, Uranium, Speeches, Emails and the Podesta Company!”	T58[Russia]
02/18/18	7:33 AM	“I never said Russia did not meddle in the election, I said “it may be Russia, or China or another country or group, or it may be a 400 pound genius sitting in bed and playing with his computer.” The Russian “hoax” was that the Trump campaign colluded with Russia - it never did!”	T58[Russia]
02/18/18	8:11 AM	“If it was the GOAL of Russia to create discord, disruption and chaos within the U.S. then, with all of the Committee Hearings, Investigations and Party hatred, they have succeeded beyond their wildest dreams. They are laughing their asses off in Moscow. Get smart America! ”	T58[Russia]
02/20/18	7:24 AM	“Thank you to @foxandfriends for the great timeline on all of the failures the Obama Administration had against Russia, including Crimea, Syria and so much more. We are now starting to win again!”	T59[Russia]
02/20/18	8:38 AM	“I have been much tougher on Russia than Obama, just look at the facts. Total Fake News!”	T59[Russia]
02/20/18	8:08 PM	“Bad ratings @CNN & @MSNBC got scammed when they covered the anti-Trump Russia rally wall-to-wall. They probably knew it was Fake News but, because it was a rally against me, they pushed it hard anyway. Two really dishonest newscasters, but the public is wise!”	T60[Russia]
02/24/18	6:44 PM	““Russians had no compromising information on Donald Trump” @FoxNews Of course not, because there is none, and never was. This whole Witch Hunt is an illegal disgrace...and Obama did nothing about Russia!”	T61[Russia]

TABLE B3: Tweets by Donald Trump regarding Russia

Date	Timestamp	Content	Id
09/26/16	10:28 PM	“Russia has more warheads than ever, N Korea is testing nukes, and Iran got a sweetheart deal to keep theirs. Thanks, @HillaryClinton.”	T1[N. Korea]
10/04/16	10:19 PM	“CLINTON IS WEAK ON NORTH KOREA: #VPDebate”	T2[N. Korea]
01/02/17	6:05 PM	“North Korea just stated that it is in the final stages of developing a nuclear weapon capable of reaching parts of the U.S. It won’t happen!”	T3[N. Korea]
06/20/17	12:41 PM	“The U.S. once again condemns the brutality of the North Korean regime as we mourn its latest victim.”	T4[N. Korea]

06/29/17	9:44 PM	"Just finished a very good meeting with the President of South Korea. Many subjects discussed including North Korea and new trade deal!"	$T5^{\text{N. Korea}}$
06/30/17	4:55 PM	"The era of strategic patience with the North Korea regime has failed. That patience is over. We are working closely with..."	$T6^{\text{N. Korea}}$
07/03/17	10:19 PM	"North Korea has just launched another missile. Does this guy have anything better to do with his life? Hard to believe that South Korea."	$T7^{\text{N. Korea}}$
08/05/17	6:44 PM	"The United Nations Security Council just voted 15-0 to sanction North Korea. China and Russia voted with us. Very big financial impact!"	$T8^{\text{N. Korea}}$
08/05/17	7:14 PM	"United Nations Resolution is the single largest economic sanctions package ever on North Korea. Over one billion dollars in cost to N.K."	$T8^{\text{N. Korea}}$
08/06/17	9:22 PM	"Just completed call with President Moon of South Korea. Very happy and impressed with 15-0 United Nations vote on North Korea sanctions."	$T8^{\text{N. Korea}}$
08/07/17	4:15 PM	"The Fake News Media will not talk about the importance of the United Nations Security Council's 15-0 vote in favor of sanctions on N. Korea!"	$T9^{\text{N. Korea}}$
08/08/17	7:17 AM	"After many years of failure, countries are coming together to finally address the dangers posed by North Korea. We must be tough & decisive!"	$T9^{\text{N. Korea}}$
08/11/17	7:29 AM	"Military solutions are now fully in place,locked and loaded,should North Korea act unwisely. Hopefully Kim Jong Un will find another path!"	$T10^{\text{N. Korea}}$
08/16/17	7:39 AM	"Kim Jong Un of North Korea made a very wise and well reasoned decision. The alternative would have been both catastrophic and unacceptable!"	$T11^{\text{N. Korea}}$
08/30/17	8:47 AM	"The U.S. has been talking to North Korea, and paying them extortion money, for 25 years. Talking is not the answer!"	$T12^{\text{N. Korea}}$
09/03/17	7:46 AM	"South Korea is finding, as I have told them, that their talk of appeasement with North Korea will not work, they only understand one thing!"	$T13^{\text{N. Korea}}$
09/03/17	12:07 PM	"I will be meeting General Kelly, General Mattis and other military leaders at the White House to discuss North Korea. Thank you."	$T13^{\text{N. Korea}}$
09/03/17	12:14 PM	"The United States is considering, in addition to other options, stopping all trade with any country doing business with North Korea."	$T13^{\text{N. Korea}}$
09/17/17	7:53 AM	"I spoke with President Moon of South Korea last night. Asked him how Rocket Man is doing. Long gas lines forming in North Korea. Too bad!"	$T14^{\text{N. Korea}}$
09/20/17	6:40 AM	"After allowing North Korea to research and build Nukes while Secretary of State (Bill C also), Crooked Hillary now criticizes."	$T15^{\text{N. Korea}}$

09/21/17	1:58 PM	"Today, I announced a new Executive Order with re: to North Korea. We must all do our part to ensure the complete denuclearization of #NoKo."	T16[N. Korea]
09/22/17	6:28 AM	"Kim Jong Un of North Korea, who is obviously a madman who doesn't mind starving or killing his people, will be tested like never before!"	T17[N. Korea]
09/23/17	5:59 PM	"Iran just test-fired a Ballistic Missile capable of reaching Israel. They are also working with North Korea. Not much of an agreement we have!"	T18[N. Korea]
09/23/17	11:08 PM	"Just heard Foreign Minister of North Korea speak at U.N. If he echoes thoughts of Little Rocket Man, they won't be around much longer!"	T18[N. Korea]
09/26/17	7:14 AM	"Great interview on @foxandfriends with the parents of Otto Warmbier: 1994 - 2017. Otto was tortured beyond belief by North Korea."	T19[N. Korea]
10/07/17	3:30 PM	"Presidents and their administrations have been talking to North Korea for 25 years, agreements made and massive amounts of money paid..."	T20[N. Korea]
10/09/17	6:50 AM	"Our country has been unsuccessfully dealing with North Korea for 25 years, giving billions of dollars & getting nothing. Policy didn't work!"	T20[N. Korea]
11/07/17	10:12 PM	"The North Korean regime has pursued its nuclear & ballistic missile programs in defiance of every assurance, agreement, & commmitment it has made to the U.S. and its allies. It's broken all of those commitments..."	T21[N. Korea]
11/07/17	10:43 PM	"Together, we dream of a Korea that is free, a peninsula that is safe, and families that are reunited once again!"	T21[N. Korea]
11/28/17	8:45 PM	"After North Korea missile launch, it's more important than ever to fund our gov't & military! Dems shouldn't hold troop funding hostage for amnesty & illegal immigration. I ran on stopping illegal immigration and won big. They can't now threaten a shutdown to get their demands."	T22[N. Korea]
12/22/17	3:47 PM	"The United Nations Security Council just voted 15-0 in favor of additional Sanctions on North Korea. The World wants Peace, not Death!"	T23[N. Korea]
01/02/18	9:08 AM	"Sanctions and "other" pressures are beginning to have a big impact on North Korea. Soldiers are dangerously fleeing to South Korea. Rocket man now wants to talk to South Korea for first time. Perhaps that is good news, perhaps not - we will see!"	T24[N. Korea]
01/02/18	7:49 PM	"North Korean Leader Kim Jong Un just stated that the "Nuclear Button is on his desk at all times." Will someone from his depleted and food starved regime please inform him that I too have a Nuclear Button, but it is a much bigger & more powerful one than his, and my Button works!"	T25[N. Korea]

01/04/17	6:32 AM	"With all of the failed "experts" weighing in, does anybody really believe that talks and dialogue would be going on between North and South Korea right now if I wasn't firm, strong and willing to commit our total "might" against the North. Fools, but talks are a good thing!"	$T26^{N.\ Korea}$
01/14/18	7:58 AM	"The Wall Street Journal stated falsely that I said to them "I have a good relationship with Kim Jong Un" (of N. Korea). Obviously I didn't say that. I said "I'd have a good relationship with Kim Jong Un," a big difference. Fortunately we now record conversations with reporters..."	$T27^{N.\ Korea}$
02/08/18	1:44 PM	"I will be meeting with Henry Kissinger at 1:45pm. Will be discussing North Korea, China and the Middle East."	$T28^{N.\ Korea}$

TABLE B4: Tweets by Donald Trump regarding North Korea

TABLE B5: Estimates of linear models

Variables	Baseline	AI Classification	Trade vs. Migration Baseline	Trade vs. Migration AI Classification
Linear trend	-0.03	-0.03	-0.03	-0.03
	(0.02)	(0.02)	(0.02)	(0.02)
$\mathbb{1}\{\text{Bad}\}$	6.17*	6.63*		
	(3.73)	(3.93)		
$\mathbb{1}\{\text{Trade}\}$			13.44***	10.47**
			(3.91)	(4.25)
$\mathbb{1}\{\text{Migration}\}$			-1.51	-0.18
			(5.02)	(5.73)
$\mathbb{1}\{\text{Positive}\}$	-3.81*	-1.59	-3.92*	-1.60
	(2.10)	(2.08)	(2.09)	(2.12)
$\mathbb{1}\{\text{Neutral}\}$	-4.99	-5.91	-4.92	-5.85
	(4.28)	(3.70)	(4.23)	(3.65)
VIX index	0.65***	0.65***	0.65***	0.65***
	(0.14)	(0.14)	(0.14)	(0.14)
Oil prices (WTI)	0.73**	0.74**	0.72**	0.73**
	(0.30)	(0.30)	(0.29)	(0.30)
Change in Fed Funds	0.004	0.001	0.03	0.01
	(0.06)	(0.06)	(0.06)	(0.06)
Surprise measure	-0.02	-0.02	-0.03	-0.02
	(0.04)	(0.04)	(0.04)	(0.04)
Lagged change in MP	0.14*	0.14*	0.13**	0.14**
	(0.08)	(0.08)	(0.06)	(0.07)
Constant	Yes	Yes	Yes	Yes
Obs.	421	421	421	421
R-squared	0.20	0.20	0.21	0.21

Note: ***, **, * denotes statistical significance at 1%, 5%, and 10%, respectively. In parenthesis, HAC robust standard errors are displayed. The surprise measure corresponds to the difference between official released statistic and Bloomberg specialists forecast, and MP stands for domestic monetary policy rate.

TABLE B6: Estimates for placebo

Variables	Baseline	Disjoint sets
Linear trend	-0.03	-0.03
	(0.02)	(0.02)
$\mathbb{1}\{\text{China}\}$	-2.89**	-3.46**
	(1.40)	(1.37)
$\mathbb{1}\{\text{N. Korea}\}$	-1.58	-2.42*
	(1.28)	(1.28)
$\mathbb{1}\{\text{Russia}\}$	0.98	0.96
	(1.65)	(1.75)
$\mathbb{1}\{\text{Bad}\}$	5.88	5.90
	(3.69)	(3.80)
$\mathbb{1}\{\text{Positive}\}$	-3.99*	-4.06*
	(2.13)	(2.11)
$\mathbb{1}\{\text{Neutral}\}$	-5.18	-5.25
	(4.30)	(4.31)
VIX index	0.63***	0.63***
	(0.14)	(0.14)
Oil prices (WTI)	0.74**	0.74**
	(0.29)	(0.29)
Change in Fed Funds	-0.005	-0.01
	(0.06)	(0.06)
Surprise measure	-0.02	-0.01
	(0.04)	(0.04)
Lagged change in MP	0.15**	0.15**
	(0.08)	(0.08)
Constant	Yes	Yes
Obs.	421	421
R-squared	0.21	0.21

Note: ***, **, * denotes statistical significance at 1%, 5%, and 10%, respectively. In parenthesis, HAC robust standard errors are displayed. The surprise measure corresponds to the difference between official released statistic and Bloomberg specialists forecast, and MP stands for domestic monetary policy rate. The third column presents the estimates for the same linear model; however, the indicator functions are constructed such that there are no days in which they overlap.

FIGURE A1. Distribution of “Tweets” over time

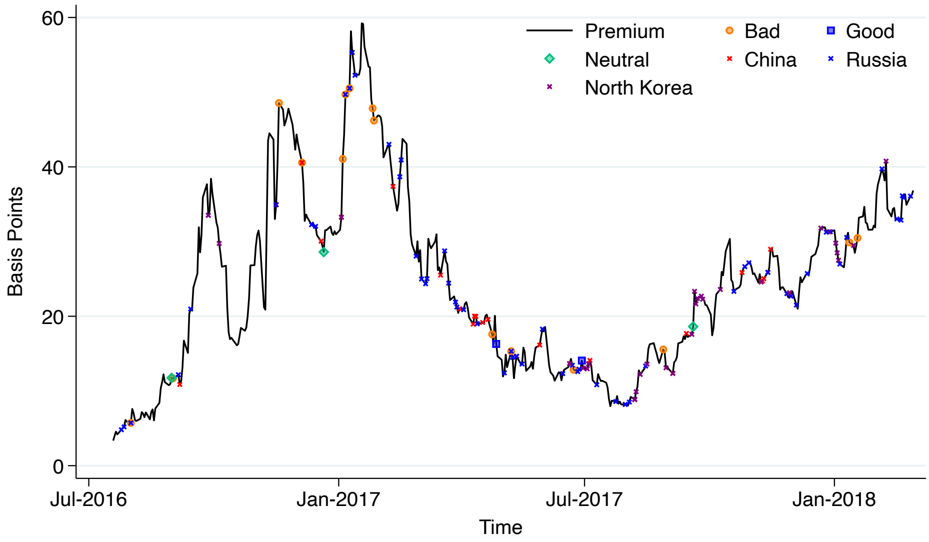

Note: The black line represents the difference between the 5-year Mexican CDS and the counterfactual, while the markers corresponds to the days in which President Trump “tweeted”.

Appendix C: Asset prices - robustness and derivations

This appendix collects the robustness exercises for the asset-price responses of Section 3.2.1, together with the derivations underlying the identification-through-heteroskedasticity estimates.

Additional impulse-response results

Adding more controls. To alleviate concerns of the dynamic responses being biased by omitted variables, I incorporate additional controls into the baseline specification. In particular, following Matveev and Ruge-Murcia (2024), I include the forward rate from 1 month up to 5 years to account for variations in the expected exchange rate. Also, to account for additional trade-policy uncertainty, I include the TPU index developed by Caldara et al. (2020). Figure C1 displays the corresponding results. The black markers correspond to the baseline estimates while the blue markers display the dynamic response when adding more controls.

FIGURE C1. Cumulative impulse-response functions: Additional controls

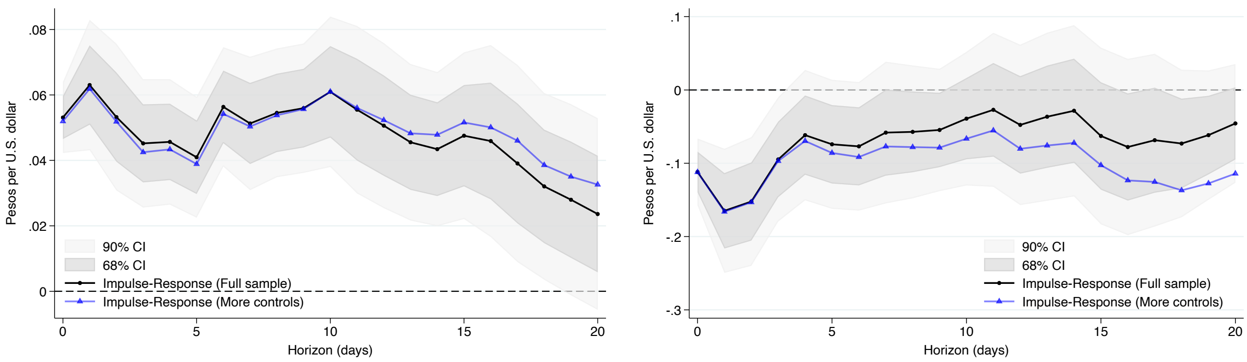

(A) Exchange rate

(B) Stock returns

Note: Figure C1A displays the cumulative impulse-response function of the exchange rate to an increase of 1 basis point in the premium while Figure C1B presents the cumulative impulse-response function of the Mexican stock index to an increase of 1 basis point in $\hat{\alpha}_t$. The shaded areas correspond to the confidence intervals at 68% and 90%, respectively, computed using HAC robust standard errors. Following Montiel Olea and Plagborg-Møller (2021), I also compute robust to heteroskedasticity standard errors using the Huber-White estimator.

Response of the peso relative to other currencies. I also estimate how the Mexican peso vis-à-vis the Euro and the Canadian dollar, respectively, responds to the estimated treatment effect $\hat{\alpha}_t$. Figure C2 displays the corresponding cumulative impulse-responses. In both cases, I observe a depreciation of the peso. On impact, it depreciates by 3 cents against the Euro and 1 cent against the Canadian dollar. Also, in both cases, the effect takes almost a month to disappear. This result is consistent with a "flight-to-safety" response. As foreigners perceive Mexico as a riskier destination for their investments due to Trump's actions and rhetoric, they sell their positions in local currency, exerting downward pressure on the value of the peso.

Role of estimation uncertainty. As discussed in Pagan (1984) and Murphy and Topel (2002),

FIGURE C2. Cumulative impulse-response: Peso vis-à-vis the Euro and Canadian Dollar

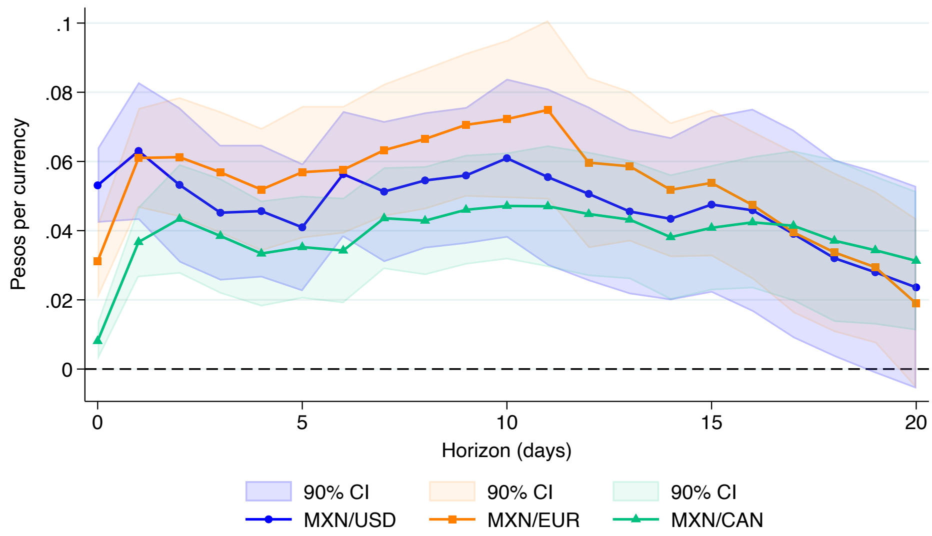

Note: The blue markers correspond to the response of the USD/MXN exchange rate to a 1 bp increase in $\hat{\alpha}_t$, the orange markers correspond to the estimated response of the EUR/MXN exchange rate, and the green markers correspond to the CAD/MXN exchange rate. The shaded areas represent the 90% confidence intervals computed using HAC robust standard errors.

there might be some issues when estimating the impulse responses given possible estimation uncertainty in $\hat{\alpha}_t$. To assess the potential role of estimation uncertainty, I estimate a variant of equation (8) where, instead of using $\hat{\alpha}_t$ as regressor, I use the Mexican EMBI spread instrumented with the premium. The results for the cumulative impulse responses are displayed in Figure C3, where both point estimates and confidence intervals are very similar to those of the baseline specification. This suggests that estimation uncertainty in the premium does not play a crucial role in driving the dynamic responses of the exchange rate and stock returns to a 1 bp increase in $\hat{\alpha}_t$.

An alternative to the previous approach is to treat $\hat{\alpha}_t$ explicitly as a generated regressor and propagate the uncertainty of the first stage into the second. In each bootstrap replication I perturb the pre-treatment fit of the synthetic control, re-estimate the donor weights, rebuild the premium over the post-treatment window, and re-run the local projections in equation (8); the variability of the resulting coefficients across replications is then added to the HAC standard errors. Intuitively, this asks how much the impulse responses would change had the synthetic control been estimated on a slightly different pre-treatment sample. Figure C4 overlays the resulting 90% confidence interval on the baseline bands for both asset prices. The two are nearly coincident: incorporating estimation uncertainty widens the standard errors by less than three percent at every horizon, so the sampling variability of the synthetic-control premium is second order for the dynamic responses.

Panel estimation for stock returns. Regarding the effect over stock prices, one can analyze the effect of Trump's pressure by sector using the S&P (2023) Sectoral Analytical Indices, which measure the performance of the economic sectors within the Mexican equity markets, where the index constituents are weighted by market capitalization. In particular, I analyze the performance of stock returns of firms within 7 sectors: materials, industrial, non-basic goods and services, frequent consumption, health, financial services, and telecommunica-

FIGURE C3. Cumulative impulse-response functions: LP-IV approach

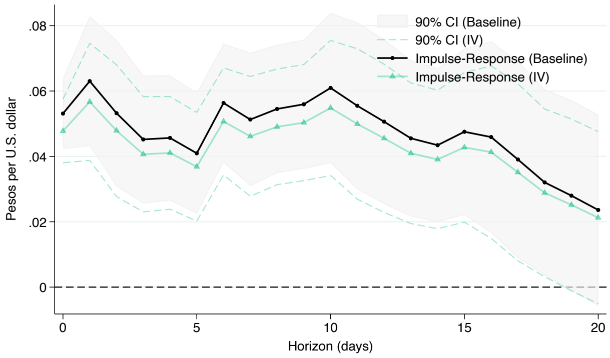

(A) Exchange rate

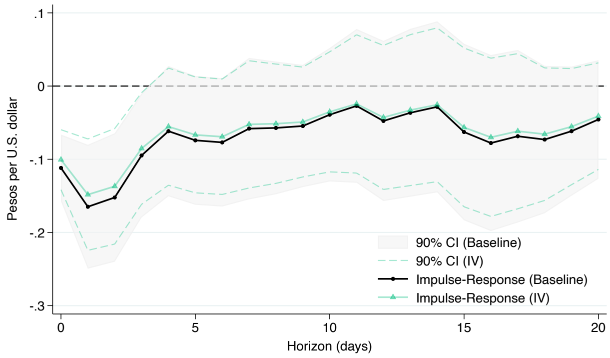

(B) Stock returns

Note: Figure C3A displays the cumulative impulse-response function of the first difference in the exchange rate to an increase of 1 basis point in the instrumented EMBI spread while Figure C3B presents the cumulative impulse-response function of the daily log-return of the Mexican stock index to an increase of 1 basis point in the instrumented EMBI spread. The black markers correspond to the baseline estimates while the green markers are the estimates of the IV approach. The first-stage F statistic for the linear equation between the EMBI spread and $\hat{\alpha}_t$ is 20.96. The shaded areas correspond to the confidence intervals of the baseline estimation at 68% and 90%, respectively, while the green dashed line represents the 90% confidence intervals for the IV approach.

FIGURE C4. Cumulative impulse-response functions: Block wild bootstrap

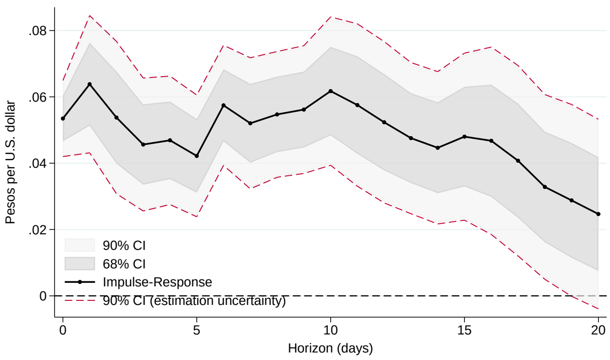

(A) Exchange rate

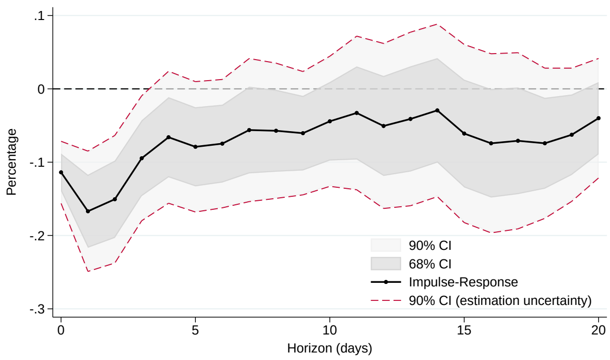

(B) Stock returns

Note: Figure C4A displays the cumulative impulse-response function of the first difference in the exchange rate to a 1 basis point increase in $\hat{\alpha}_t$, and Figure C4B the response of the daily log-return of the Mexican stock index. The shaded areas are the baseline 68% and 90% confidence intervals computed with HAC robust standard errors; the dashed line is the 90% confidence interval that incorporates estimation uncertainty in $\hat{\alpha}_t$, obtained from a block wild bootstrap that re-estimates the synthetic control and re-runs the local projections ($B = 2{,}000$ replications). Across horizons the two intervals differ by less than three percent.

tions. For each sector, I have daily information on the index for which I compute the daily log return and estimate the following series of panel regressions for horizons $h = 0, \ldots, 20$:

$$y_{i,t+h} - y_{i,t-1} = \gamma_i^h + \hat{\alpha}_t \beta^h + \boldsymbol{x}_t \psi^h + u_{i,t+h}, \tag{C.1}$$

where $y_{i,t+h}$ is the log of sector i index at time $t+h$, γ_i^h represents the sector fixed effect, $\hat{\alpha}_t$ is the estimated premium, $\boldsymbol{x}_t$ is a vector of controls that includes lagged observations of y_t and $\hat{\alpha}_t$ to account for autocorrelation, lagged daily changes in the domestic monetary policy rate, the VIX index, oil prices, and the daily change in Fed funds rates, and $u_{i,t+h}$ is the error term.

FIGURE C5. Average effect of Trump shock

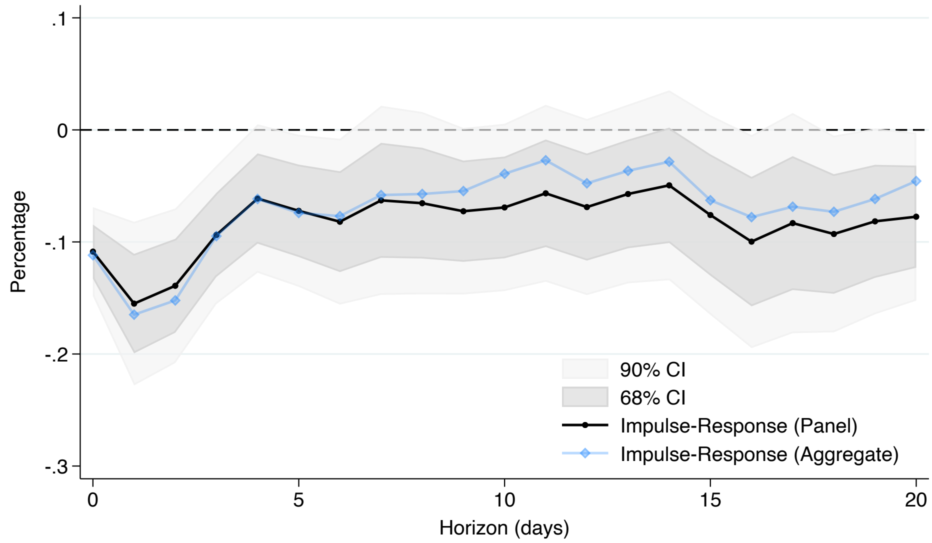

Note: The black markers correspond to the average effect while the blue markers display the aggregate effect estimated using only the daily log-returns of the IPC. The shaded areas correspond to the confidence intervals at 68% and 90%, respectively, computed using Driscoll and Kraay (1998) standard errors.

Figure C5 displays the cumulative impulse-response functions to a 1 bp increase in the dynamic treatment effect. Since $\hat{\alpha}_t$ is the same for each sector, the error term is potentially cross-sectionally correlated and, to account for that, the confidence intervals are constructed using Driscoll and Kraay (1998) standard errors. The black markers correspond to the average effect, and the blue markers present the aggregate effect estimated using only the daily log-returns of the IPC. Once I account for the unobserved fixed sector characteristics, the average effect takes more time to dissipate than the aggregate estimates; however, the overall dynamics are very similar.

Decomposition by sector. Figure C6 displays the cumulative response to a positive shock of 1 bp in the premium. According to the estimates, all sectors were affected on impact by a 1 bp increase in the premium; however, the most sensitive sectors are the financial services sector with a negative response of -0.17%, followed by the industrial sector with a contraction of -0.15%, and the non-basic goods and services sector with a drop of -0.12%. These results are consistent with higher trade policy uncertainty as documented by Handley and Limão (2017), possible issues with financial deregulation according to Cooley and White (2017), and potential interruption of global supply chains following Blanchard (2017).

Impulse responses from a VAR-X model

As a standard in the literature of estimation of impulse responses, I also estimate the impulse response for the daily change in the Mexican peso and the daily log-return of the *IPC* using the following VAR with exogenous variables:

$$\boldsymbol{y}_t = \boldsymbol{C} + \boldsymbol{B}_1 \boldsymbol{y}_{t-1} + \ldots + \boldsymbol{B}_p \boldsymbol{y}_{t-p} + \boldsymbol{\Theta} \boldsymbol{x}_t + \boldsymbol{u}_t, \quad \boldsymbol{u}_t \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}) \tag{C.2}$$

FIGURE C6. Cumulative impulse-response functions by sector

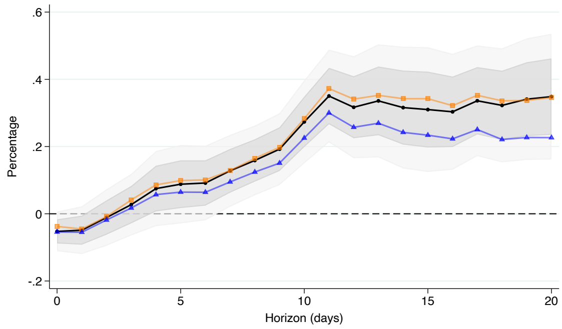

(A) Materials

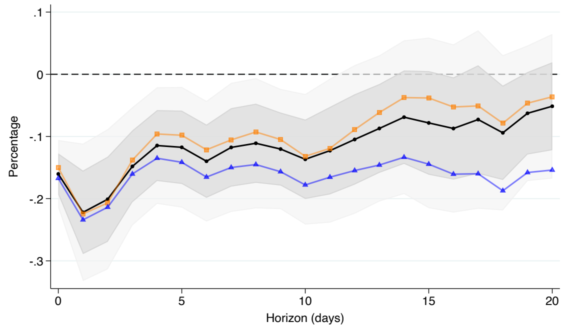

(B) Industrial

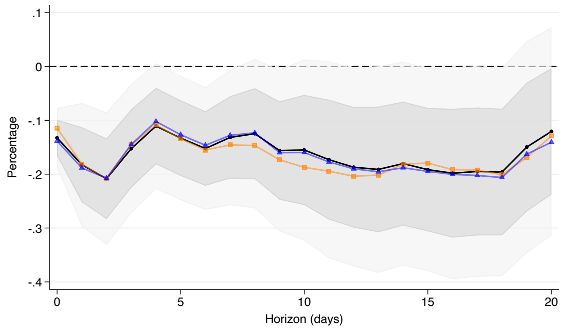

(C) Non-basic G&S

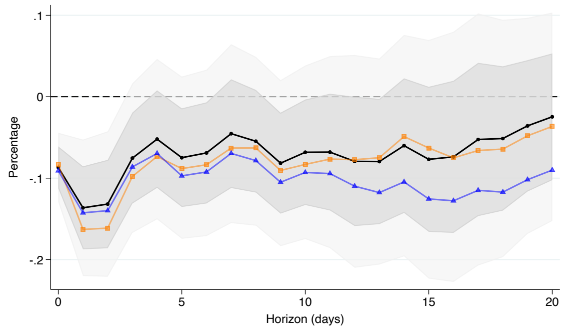

(D) Frequent consumption

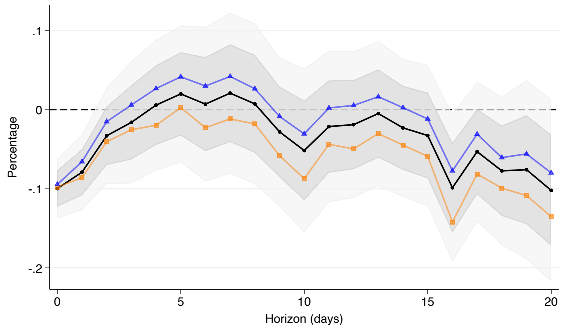

(E) Health

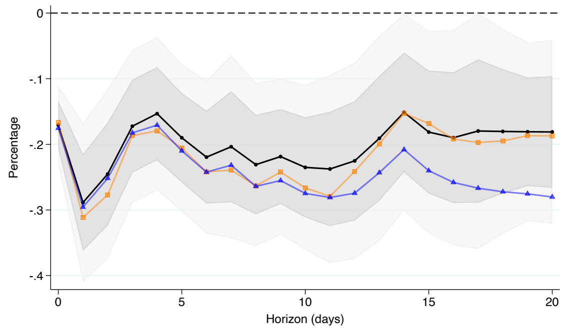

(F) Financial services

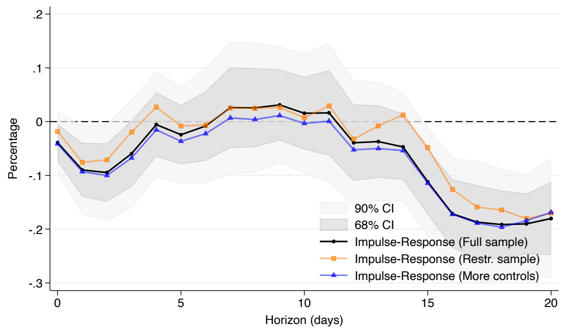

(G) Telecommunications

Note: This Figure displays the estimated CIRFs using local projections (LP) for the daily log return of the corresponding sectoral index. The shaded areas correspond to the 68% and 90% confidence intervals, respectively, computed using HAC robust standard errors.

with

$$\boldsymbol{y}_t = \begin{bmatrix} \hat{\alpha}_t \\ \Delta er_t \\ \Delta \log(IPC_t) \\ \Delta MP_t \end{bmatrix}, \quad \boldsymbol{x}_t = \begin{bmatrix} VIX_t \\ Oil_t \\ \Delta Fed_t \end{bmatrix}, \quad \text{and} \quad \boldsymbol{u}_t = \begin{bmatrix} u_{1,t} \\ u_{2,t} \\ u_{3,t} \\ u_{4,t} \end{bmatrix},$$

where $\boldsymbol{C}$ is a (4×1) vector of constants, $\boldsymbol{B}_j$ with $j = 1, \ldots, p$ are (4×4) matrices of coefficients, $\boldsymbol{\Theta}$ is a (4×3) matrix of coefficients for the exogenous variables, and $\boldsymbol{u}_t$ is a (4×1) vector of error terms. The ordering of the VAR implies that the premium is the "most" exogenous variable in the system followed by the daily changes in the exchange rate, log return of the stock exchange, and the daily change in the domestic monetary policy rate (in basis points). I choose $p = 5$ to account for a week.

The Granger causality test suggests that the premium $\hat{\alpha}_t$ is "weakly" exogenous in the sense that I cannot reject the null hypothesis that the coefficients of the other variables in the premium equation are equal to zero (see Table C1); however, I do reject the null hypothesis for the coefficients of the premium on the other equations. Finally, for identification purposes, I use zero contemporaneous restrictions such that the premium reacts with lags to changes in the other variables in the system.

TABLE C1: Granger-causality Tests

Variable	p-value
Exchange rate	0.145
IPC Index	0.276
Monetary Policy	0.427
All	0.119

Note: This table displays the p-value associated with the null-hypothesis that the coefficients of variable i in the premium equation are zero. For all variables, I take the first difference to ensure stationarity. I allow for up to 5 lags.

Figure C7 displays the comparison between the impulse-response functions for the exchange rate and stock returns to a 1 basis point increase in the premium while Figure C8 presents the comparison for the cumulative responses. The IRF and CIRF from the VAR are similar in magnitude and smoother than those from local projections. These results are consistent with the bias-variance trade-off discussed in Li et al. (2022) and Lusompa (2023).

Identification through heteroskedasticity

The simultaneous system of equations presented in Section 3.2.2 is given by:

$$\Delta y_t = \beta \Delta CDS_t + \theta_1 X_{1,t} + \theta_2 X_{2,t} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), \tag{C.3}$$

$$\Delta CDS_t = \varphi \Delta y_t + \delta_1 X_{1,t} + \delta_2 X_{2,t} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \tag{C.4}$$

where Δy_t is the first difference in the USD-MXN exchange rate or the log of the IPC index, ΔCDS_t is the first difference of the 5-year Mexican CDS, $X_{1,t}$ and $X_{2,t}$ represent the VIX index and the oil prices (WTI), respectively, and ε_t and η_t are error terms. By solving the system of

FIGURE C7. Comparison of Impulse-Response Functions

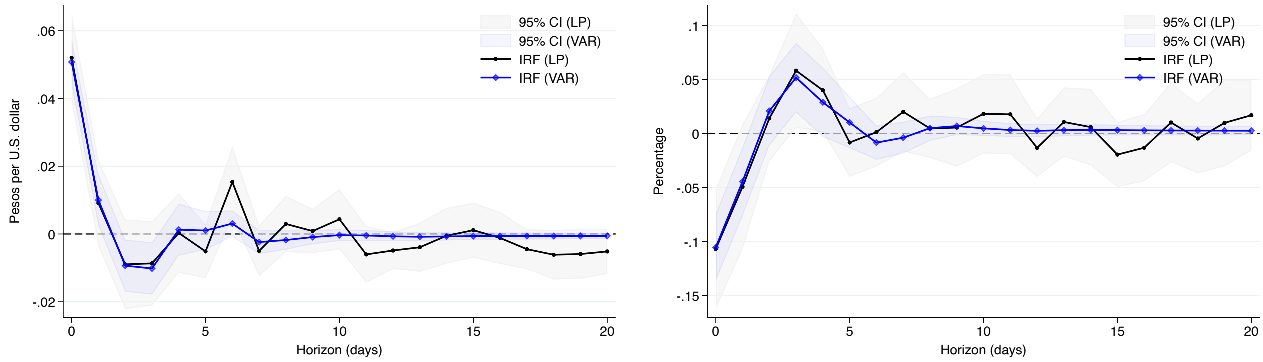

(A) IRF exchange rate

(B) IRF stock returns

Note: This figure displays the comparison between the estimated impulse response functions using local projections (LP) and vector autoregressions (VAR). The shaded area corresponds to the 95% confidence interval. The blue line and marker correspond to the VAR estimates while the black line and markers correspond to the LP estimates.

FIGURE C8. Comparison of Cumulative Impulse-Response Functions

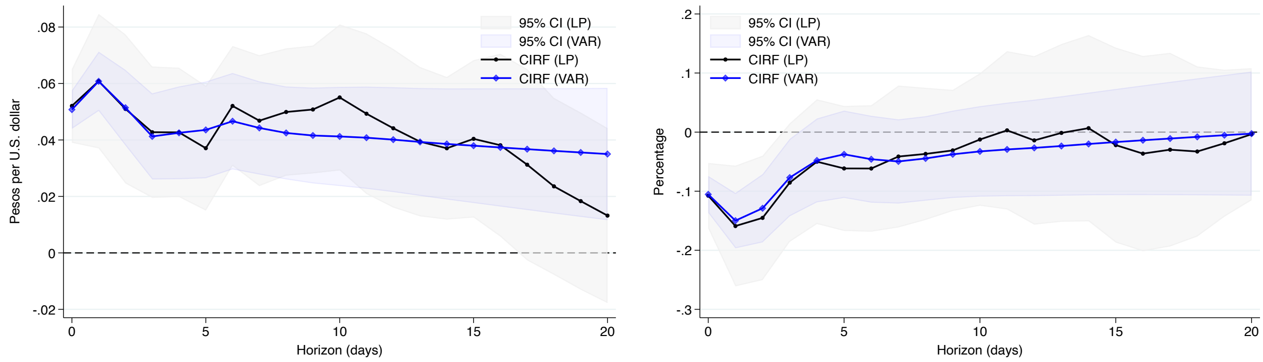

(A) IRF exchange rate

(B) IRF stock returns

Note: This figure displays the comparison between the estimated cumulative impulse response functions using local projections (LP) and vector autoregressions (VAR). The shaded area corresponds to the 95% confidence interval. The blue line and marker correspond to the VAR estimates while the black line and markers correspond to the LP estimates.

equations as functions of the error terms ε_t, η_t, and exogenous predictors $X_{1,t}$ and $X_{2,t}$, we can derive the following reduced-form expressions for Δy_t and ΔCDS_t, respectively:

$$\Delta y_t = \left[\frac{1}{1-\beta\varphi}\right]\left[(\theta_1+\beta\delta_1)\,X_{1,t} + (\theta_2+\beta\delta_2)\,X_{2,t} + \varepsilon_t + \beta\eta_t\right],$$

$$\Delta CDS_t = \left[\frac{1}{1-\beta\varphi}\right]\left[(\varphi\theta_1+\delta_1)\,X_{1,t} + (\varphi\theta_2+\delta_2)\,X_{2,t} + \varphi\varepsilon_t + \eta_t\right].$$

Consider estimating β in equation (C.3) by OLS, including $X_{1,t}$ and $X_{2,t}$ as controls. Under the assumptions that $\mathbb{E}[\varepsilon_t\varepsilon_{t-j}] = \mathbb{E}[\eta_t\eta_{t-j}] = \mathbb{E}[\varepsilon_t\eta_t] = \mathbb{E}[\varepsilon_t X_{j,t}] = \mathbb{E}[\eta_t X_{j,t}] = 0$ for $j = 1,2$, the Frisch-Waugh-Lovell theorem implies that the OLS estimator of β is equivalent to a bivariate regression of Δy_t on the residualized regressor $\widetilde{\Delta CDS}_t$ obtained from projecting

ΔCDS_t onto $X_{1,t}$ and $X_{2,t}$. From the reduced-form expression for ΔCDS_t, this residualized regressor is given by:

$$\widetilde{\Delta CDS}_t = \left[\frac{1}{1-\beta\varphi}\right][\varphi\varepsilon_t + \eta_t],$$

which is orthogonal to $X_{1,t}$ and $X_{2,t}$ by construction. Consequently, the probability limit of the multivariate OLS estimator of β is given by:

$$\text{plim}\,\hat{\beta}_{OLS} = \beta + \underbrace{(1-\beta\varphi)\left[\frac{\varphi\sigma_\varepsilon^2}{\varphi^2\sigma_\varepsilon^2 + \sigma_\eta^2}\right]}_{\equiv\Omega}, \tag{C.5}$$

where Ω represents the asymptotic bias due to endogeneity. Importantly, by the Frisch-Waugh-Lovell theorem, the bias does not depend on $\sigma_{X_1}^2$, $\sigma_{X_2}^2$, or $\sigma_{X_1X_2}$: including $X_{1,t}$ and $X_{2,t}$ as controls partials out their contribution, leaving only the component of the bias attributable to the simultaneity between Δy_t and ΔCDS_t. The bias is therefore driven by $\varphi\sigma_\varepsilon^2$—the product of the feedback coefficient from Δy_t to ΔCDS_t and the variance of the structural shock in (C.3)—relative to the total variance of the residualized regressor.

As described by Rigobon and Sack (2004), there are some strong identifying assumptions required, for example, in the event-study approach that will allow us to eliminate the bias in the estimation of β. These strong assumptions are that, in the limit, the ratio $\frac{\sigma_\eta^2}{\sigma_\varepsilon^2}$ tends to infinity. In other words, the variance of the CDS shock becomes infinitely larger relative to the variance of the shock to the outcome equation. If this ratio is finite, the bias in the estimation of the parameter of interest will be far from zero.

Rigobon and Sack (2004) shows that we can use a much weaker set of assumptions to identify the parameter of interest. Identification through heteroskedasticity does not require that the ratio of variances of the shocks become infinitely larger but instead relies on the change in the covariance between Δy_t and ΔCDS_t when the variance of the CDS shock increases. In particular, after solving for the reduced form of the previous system of equations, we have that the covariance matrix for type $i = E, NE$ window is given by:

$$\Sigma_i = \left(\frac{1}{1-\beta\varphi}\right)^2 \begin{bmatrix} f\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) + \beta^2(\sigma_\eta^i)^2 + (\sigma_\varepsilon^i)^2 & h\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) + \beta(\sigma_\eta^i)^2 + \varphi(\sigma_\varepsilon^i)^2 \\ h\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) + \beta(\sigma_\eta^i)^2 + \varphi(\sigma_\varepsilon^i)^2 & g\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) + (\sigma_\eta^i)^2 + \varphi^2(\sigma_\varepsilon^i)^2 \end{bmatrix},$$

with

$$\begin{aligned} f\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) &= (\theta_1+\beta\delta_1)^2(\sigma_{X_1}^i)^2 + (\theta_2+\beta\delta_2)^2(\sigma_{X_2}^i)^2 + 2(\theta_1+\beta\delta_1)(\theta_2+\beta\delta_2)\sigma_{X_1X_2}^i, \\ h\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) &= (\varphi\theta_1+\delta_1)(\theta_1+\beta\delta_1)(\sigma_{X_1}^i)^2 + (\varphi\theta_1+\delta_1)(\theta_2+\beta\delta_2)\sigma_{X_1X_2}^i \\ &\quad + (\varphi\theta_2+\delta_2)(\theta_1+\beta\delta_1)\sigma_{X_1X_2}^i + (\varphi\theta_2+\delta_2)(\theta_2+\beta\delta_2)(\sigma_{X_2}^i)^2, \\ g\left(\sigma_{X_1}^i, \sigma_{X_2}^i, \sigma_{X_1X_2}^i\right) &= (\varphi\theta_1+\delta_1)^2(\sigma_{X_1}^i)^2 + (\varphi\theta_2+\delta_2)^2(\sigma_{X_2}^i)^2 + 2(\varphi\theta_1+\delta_1)(\varphi\theta_2+\delta_2)\sigma_{X_1X_2}^i. \end{aligned}$$

The weak identifying assumptions are that $\sigma_{X_j}^E = \sigma_{X_j}^{NE}$ for $j = 1, 2$, $\sigma_{X_1X_2}^E = \sigma_{X_1X_2}^{NE}$, $\sigma_\varepsilon^E = \sigma_\varepsilon^{NE}$, and $\sigma_\eta^E > \sigma_\eta^{NE}$, such that when taking the difference between Σ_E and Σ_{NE} the functions $f(\cdot)$, $h(\cdot)$, $g(\cdot)$, and the variance of the error term in equation (C.3) will cancel out giving us the following expression for the difference in covariance matrices:

$$\Delta\Sigma \equiv \Sigma_E - \Sigma_{NE} = \frac{(\sigma_\eta^E)^2 - (\sigma_\eta^{NE})^2}{(1-\beta\varphi)^2}\begin{bmatrix} \beta^2 & \beta \\ \beta & 1 \end{bmatrix},$$

which implies that we can recover the effect of the CDS on the exchange rate or stock returns as:

$$\hat{\beta}^{HET} = \frac{\Delta\hat{\Sigma}_{12}}{\Delta\hat{\Sigma}_{22}} = \frac{\hat{Cov}_E(\Delta CDS_t, \Delta y_t) - \hat{Cov}_{NE}(\Delta CDS_t, \Delta y_t)}{\hat{Var}_E(\Delta CDS_t) - \hat{Var}_{NE}(\Delta CDS_t)}.$$

As discussed in Rigobon and Sack (2004), "this approach can be thought of as estimating β from the change in the bias in equation (C.5) as the variance of the CDS shock changes, rather than requiring that the level of the bias goes to zero."

Tests of differences in variances

To test the validity of the identification assumptions, I conduct tests of differences in variances of the 5-year CDS, exchange rates, and stock returns between event and non-event windows. Table C2 displays Levene (1960) and Brown and Forsythe (1974) F-statistics for the test of equality of variances. I reject the null hypothesis of equal variances only for the CDS spread at standard levels of 5%. For two estimators, I reject the null for stock returns at 10%.[38]

TABLE C2: Tests of differences in variances

Test	CDS	Exchange Rate	Stock Returns
Levene	4.14**	1.21	2.94*
Brown-Forsythe (median)	4.27**	1.17	2.76
Brown-Forsythe (trimmed mean)	4.26**	1.16	2.80*

Note: "Test" column displays the F-statistic being computed for the test of equal variances. The sample used corresponds to the balanced sample of 41 observations for event windows and 41 observations for non-event windows. ***, **, * denotes statistical significance at 1%, 5%, and 10%, respectively.

Common factor during event windows?

Following Hébert and Schreger (2017), I display two scatter plots to see if there are signs of a common factor inducing co-movements between the 5-year CDS of Mexico and the stock returns of Mexico and Panama during event days. Ideally, during event windows, there should only be evidence of co-movement between Mexico's sovereign risk measure and the Mexico's stock index. Figure C9 displays both scatter plots. For Mexico and Panama, I observe a negative relationship between stock returns and daily changes in Mexico's 5-year CDS spread during non-event days. However, during event days, that relationship disappears for Panama and gets stronger for Mexico. This may be indicative that omitted common factors are not very important during event days.

Dynamic responses via the heteroskedasticity-based instrument

I extend the heteroskedasticity-based identification approach to estimate dynamic responses using local projections. Specifically, I replace the synthetic-control premium with daily CDS

[38]The null hypothesis of equal variances for the control variables (i.e., oil prices and VIX index) is not rejected.

FIGURE C9. Stock returns and Mexico's sovereign risk indicator

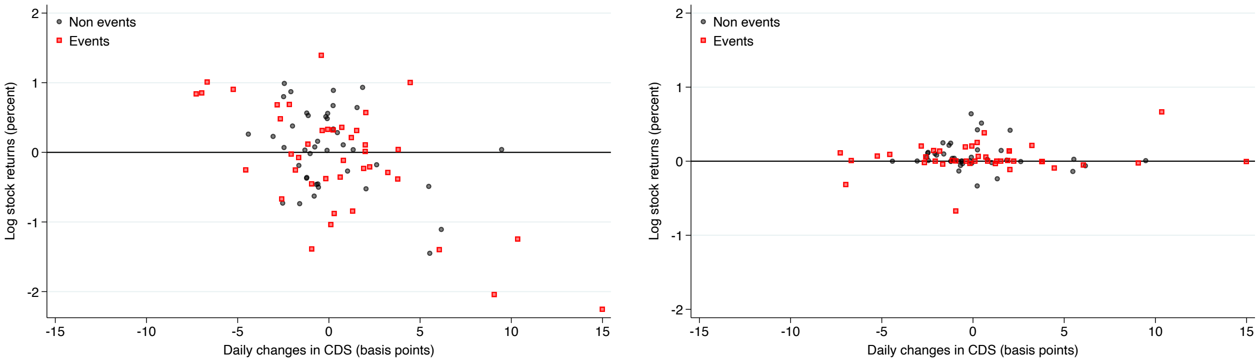

(A) Mexico stock index (B) Panama stock index

Note: The vertical axis displays the percentage change in log returns for Mexico and Panama computed using the *IPC* and *BVPS* indexes, respectively. The horizontal axis displays the changes in Mexico's 5-year CDS spread in basis points. Black markers corresponds to non-event days, while red markers correspond to event days.

innovations instrumented by the event-based heteroskedasticity instrument, while keeping the remaining specification unchanged. Figure C10 displays the resulting impulse responses, which closely mirror those obtained in the baseline specification. Both approaches imply a persistent depreciation of the peso and a decline in equity prices following an increase in sovereign risk. While the IV-based responses are somewhat larger in magnitude, their overall shape and timing are remarkably similar.

FIGURE C10. Cumulative impulse-response functions using IV heteroskedasticity

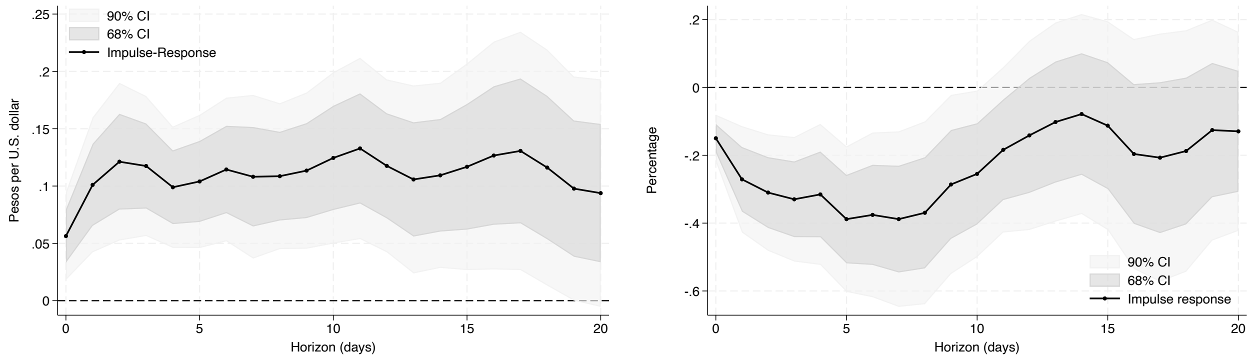

(A) Exchange rate (B) Stock returns

Note: Figure C10A displays the cumulative impulse-response function of the first difference in the exchange rate to an increase of 1 basis point in the instrumented first difference of the CDS while Figure C10B presents the cumulative impulse-response function of the daily log-return of the Mexican stock index. The shaded areas correspond to the confidence intervals at 68% and 90%, respectively, computed using HAC robust standard errors. The first stage F-stat is given by 54.92.

Appendix D: Details on default model

Numerical algorithm

I solve the model by value function iteration on a discretised state space, following the finite-horizon initialisation of Hatchondo and Martinez (2009) as implemented in Johri et al. (2022). The persistent-income grid has $N = 21$ equally spaced points in log-income space, constructed using the method of Tauchen (1986); the transitory shock is discretised on a grid of $N_m = 11$ points over $[-\bar{m}, \bar{m}]$ with probabilities from the truncated normal. A fine transitory grid is important: it approximates the continuous transitory CDF of Chatterjee and Eyigungor (2012) and delivers a smooth bond-price schedule, which is what makes the Stage-2 calibration of ω well-behaved. The debt grid has $N_b = 200$ points on $[b_{\min}, 0]$, where $b_{\min}$ is set sufficiently negative that the borrowing constraint never binds in equilibrium.

The algorithm proceeds as follows. Define the coupon per outstanding bond as $\kappa(\delta) = (r+\delta)/(1+r)$. Initialise with a terminal period T in which no borrowing is possible, so the terminal bond price is $q_T = 0$ and the terminal value function is $V_T(b, y, m) = u(y + m + \kappa(\delta)\, b)$ for all (b, y, m).

Algorithm 1 Value function iteration for the sovereign default model

---

Require: Grids $\{b_i\}_{i=1}^{N_b}$, $\{y_j\}_{j=1}^{N}$, $\{m_l\}_{l=1}^{N_m}$ with transitory weights $\{p_l\}$; transition matrix $\Pi^{(1)}$; parameters $(r, \delta, \beta, \gamma, \theta, \chi, d_0, d_1, \sigma_m)$ (with χ the lenders' price of income risk); tolerance $\varepsilon > 0$.

Ensure: Value functions $\{V, V^R, X\}$, default rule $d(b, y, m)$, borrowing policy $g(b, y, m)$, bond price schedule $q(b', y)$.

Initialise: $X^0(y_j, m_l) = u(y_j - \phi(y_j) + m_l)$, $V^{R,0}(b_i, y_j, m_l) = u(y_j + m_l + \kappa b_i)$, $V^0 = \max\{V^{R,0}, X^0(y_j, -\bar{m})\}$, $q^0 \equiv 0$, $\Delta \leftarrow \infty$.

while $\Delta \geq \varepsilon$ **do**

  Bond prices: iterate $q^{\text{new}}(b_i, y_j) = \sum_{j'} \Pi^{(1)}(y_{j'}|y_j)\, \mathcal{M}(y_j, y_{j'}) \sum_{l'} p_{l'}[1 - d(b_i, y_{j'}, m_{l'})]\big[\kappa(\delta) + (1-\delta)\, q\big(g(b_i, y_{j'}, m_{l'}), y_{j'}\big)\big]$ until $\sup_{i,j} |q^{\text{new}} - q| < 10^{-12}$, where $\mathcal{M}(y_j, y_{j'}) = \exp\{-r - \chi(\varepsilon'_y + \frac{1}{2}\chi\sigma^2)\}$ is the risk-averse lenders' stochastic discount factor and the repaid continuation is priced at the next-period debt policy $g(b_i, y_{j'}, m_{l'})$, reflecting debt dilution; set $q \leftarrow q^{\text{new}}$.

  Autarky value: $X^{\text{new}}(y_j, m_l) = u(y_j - \phi(y_j) + m_l) + \beta \sum_{j'} \Pi^{(1)}(y_{j'}|y_j) \sum_{l'} p_{l'}[\theta V(0, y_{j'}, m_{l'}) + (1-\theta) X(y_{j'}, m_{l'})]$.

  Repayment value: for each (b_i, y_j, m_l), $V^{R,\text{new}}(b_i, y_j, m_l) = \max_{b_k \leq 0}\{u(y_j + m_l + \kappa b_i - (b_k - (1-\delta) b_i) q(b_k, y_j)) + \beta \sum_{j'} \Pi^{(1)}(y_{j'}|y_j) \sum_{l'} p_{l'} V(b_k, y_{j'}, m_{l'})\}$; store $g(b_i, y_j, m_l) \leftarrow b_k^*$.

  Default rule: $d^{\text{new}}(b_i, y_j, m_l) = \mathbf{1}\{X^{\text{new}}(y_j, -\bar{m}) > V^{R,\text{new}}(b_i, y_j, m_l)\}$; $V^{\text{new}} = \max\{V^{R,\text{new}}, X^{\text{new}}(y_j, -\bar{m})\}$.

  Check convergence: $\Delta \leftarrow \max(\sup |V^{\text{new}} - V|, \sup |q^{\text{new}} - q|)$.

  Update: $V \leftarrow V^{\text{new}}$, $V^R \leftarrow V^{R,\text{new}}$, $X \leftarrow X^{\text{new}}$, $d \leftarrow d^{\text{new}}$, $q \leftarrow q^{\text{new}}$.

end while

Output: $\{V, V^R, X, d, g, q\}$.

The outer loop iterates until the supremum norm of changes in both the value function and the bond price schedule falls below $\varepsilon = 10^{-7}$. The repayment maximization (Step 4) uses a grid search over all N_b candidate next-period debt positions, vectorized over the transitory grid; no interpolation is used. Both the bond price (Step 2) and the autarky continuation (Step 3) are pre-integrated over the transitory shock, so the price schedule remains a function of (b', y) alone. Because m is i.i.d. it never propagates through the persistent state; dilution changes only *what* is pre-integrated. As the repaid continuation $q(g(b', y', m'), y')$ is evaluated at the sovereign's m'-specific debt choice, the inner sum over m' must return the joint object $\bar{q}(b', y') \equiv \mathbb{E}_{m'}[(1 - d(b', y', m'))\, q(g(b', y', m'), y')]$, rather than an m-integrated default probability $\bar{d}(b', y')$ multiplying a fixed continuation $q(b', y')$ as in the buy-and-hold case. Following Johri et al. (2022), the default set is updated state by state within each sweep, so that the price responds immediately to the current default decision, which keeps the dilution fixed point well-behaved. Convergence is typically achieved in 500–600 outer iterations.

Data and calibration details

Data sources and transformations

Table D1 provides a complete description of all variables used in the calibration, including sources, frequency, and transformations.

TABLE D1: Data sources and transformations

Variable	Source	Frequency	Transformation
Real GDP (y_t)	FRED (NGDPRSAXDCMXQ)	Quarterly	Log, linear detrend
Private consumption (c_t)	INEGI / FRED	Quarterly	Log, linear detrend
EMBI spread	J.P. Morgan / Bloomberg	Daily	Quarterly average, annualized
5-year CDS spread	Markit	Daily	Quarterly average, bp
External debt/GDP	SHCP (see below)	Annual	Linear interpolation to quarterly
US risk-free rate	FRED (TB3MS)	Monthly	Quarterly average, decimal

Income process. The persistence $\rho = 0.922$ and innovation standard deviation $\sigma = 0.0144$ of the persistent component are estimated by OLS on the cyclical component of log Mexican real GDP (FRED series NGDPRSAXDCMXQ) over 1993Q1–2016Q4, where the cyclical component is obtained by a linear detrend. The sample starts in 1993Q1 to avoid the 1982 and 1988–1994 debt crises, which may have generated structural breaks in the income process. The transitory component is not estimated separately from the data: following Chatterjee and Eyigungor (2012), its scale $\sigma_m = 0.0089$ is treated as a computational parameter and is disciplined by the average default probability in Stage 1 of the calibration below.

Sovereign spreads. The primary spread measure is the EMBI Global Mexico, which tracks the yield spread of Mexican sovereign external bonds over US Treasury bonds of comparable maturity. Daily values are averaged to obtain quarterly series. The mean spread target of 2.23% per year is the EMBI average over 2013Q1–2016Q2. The standard deviation target of 55 bp is computed from the same series. The 5-year CDS spread (denominated in basis points) is used in the model-data comparison of the impulse responses and the SCM premium figure, having been converted to annual percentage by dividing by 100.

External debt-to-GDP. The debt-to-GDP target is constructed from the SHCP publication *Estadísticas Oportunas de Finanzas Públicas*, specifically the series on the external debt of the federal public sector. This series measures the stock of outstanding debt issued by the federal government and public enterprises that is owed to foreign creditors. It includes federal government bonds placed in international markets, loans from multilateral institutions (World Bank, IADB), and bilateral official debt. It excludes domestically held public debt and subnational borrowing.[39]

The SHCP series is reported at annual frequency; I convert it to quarterly frequency using linear interpolation (e.g., the 2017Q1 value is $\frac{3}{4}(2016) + \frac{1}{4}(2017)$). The calibration target of 12.13% is the simple average over 2013Q1-2016Q2 of the linearly interpolated series.

Risk-free rate. The quarterly risk-free rate is set to $r = 0.010$, the standard value used in the long-term-debt literature (see Chatterjee and Eyigungor (2012)).

Two-stage calibration procedure

Stage 1. The discount factor β, the output cost parameters d_0 and d_1, and the transitory-shock scale σ_m are jointly determined by minimizing a weighted sum of squared proportional deviations of model moments from data targets:

$$(\hat{\beta}, \hat{d}_0, \hat{d}_1, \hat{\sigma}_m) = \underset{\{\beta, d_0, d_1, \sigma_m\}}{\arg\min} \Bigg[w_1 \left(\frac{\bar{b}/\bar{y} - 0.1213}{0.1213} \right)^2 + w_2 \left(\frac{\bar{s} - 0.0223}{0.0223} \right)^2 + w_3 \left(\frac{\sigma_s - 0.0055}{0.0055} \right)^2 + w_4 \left(\frac{\pi_d - 0.0167}{0.0167} \right)^2 \Bigg],$$

where $\bar{b}/\bar{y}$ is the simulated mean debt-to-GDP, $\bar{s}$ the mean annualized spread, σ_s its standard deviation, and π_d the annual default probability, with weights $(w_1, w_2, w_3, w_4) = (10, 12, 2, 3)$; the heavy weights on the spread level and debt-to-GDP hold the model's spread and external debt near their targets. The default-probability target disciplines the transitory-shock scale σ_m, and the lenders' price of income risk is held fixed at $\chi = 8.75$ following Johri et al. (2022) rather than estimated, so it enters none of the moment conditions. The minimization uses a multi-start Nelder-Mead simplex algorithm initialized at values drawn from the related literature.

Stage 2. Given the Stage-1 parameters, the tilt intensity ω and the within-tercile shape ζ are calibrated jointly to two targets: the 24 bp spread increase and the change in the skewness of the cyclical component of Mexican real GDP between the pre- and post-pressure samples,

$$\bar{s}(\omega, \zeta) - \bar{s}(0) = 0.0024, \qquad \text{skew}^{\Pi^{(1)}}(\omega, \zeta) - \text{skew}^{\Pi^{(0)}} = -0.036,$$

where $\bar{s}(\omega, \zeta)$ is the simulated mean annualized spread under $\Pi^{(1)}(\omega, \zeta)$, 0.0024 (i.e., 24 bp in decimal) is the synthetic-control estimate of the causal spread increase, and the skewness

[39]The full composition and methodology are described in the SHCP's technical notes accompanying the publication, available here.

change (-0.036) is the difference between the skewness of the cyclical component of log GDP after the pressure episode (-1.097) and before it (-1.060).[40] The skewness change is a property of the income process alone and requires no model solution, so I exploit a nested procedure: for each candidate ω, I set ζ to match the skewness change exactly by a one-dimensional root-find on the ergodic distribution of the tilted chain (skewness is monotone in ζ), and I then solve and simulate the model to evaluate the spread increase. Because the spread increase is not globally monotone in ω (it rises, peaks, and then falls as the sovereign deleverages toward a lower debt level), I evaluate it on a fine grid over $\omega \in [0.0001, 0.004]$ and select the value on the rising branch closest to the 24 bp target, restricting attention to the ω at which the skewness target is attainable. This yields $\omega^* = 0.0003$ and $\zeta^* = 0.513$, which deliver a $+23$ bp spread increase while matching the skewness change exactly.

Short-term debt model

For comparison with the long-term debt model, this appendix presents the one-period bond ($\delta = 1$) obtained by shortening the maturity in the framework of Section 4. With $\delta = 1$ the continuation term $(1-\delta)q_{t+1}$ vanishes, so there is no dilution, and the price collapses to a single-period claim,

$$q_t = \mathbb{E}_t[\mathcal{M}_{t+1}\,(1 - d_{t+1})]\,, \tag{D.1}$$

which under risk-neutral pricing ($\mathcal{M}_{t+1} = 1/(1+r)$) is the standard Arellano (2008) bond price. I report two one-period variants: the standard risk-neutral benchmark, with the Arellano threshold output cost $h(y) = \min\{y, \hat{y}\}$ ($\hat{y}$ calibrated jointly with β), and a *risk-averse* version that uses the main text's stochastic discount factor and quadratic output cost $\phi(y)$, so it differs from the long-term model only in maturity.

Table D2 compares four specifications, each separately calibrated to the same Mexican targets: the one-period bond ($\delta = 1$) with risk-neutral lenders (the standard Arellano (2008) benchmark) and with the risk-averse lenders of the main text, and the long-term bond ($\delta = 0.0742$) under each lender type. Two lessons emerge. First, *maturity* sets the spread *level*: the short-term models generate mean spreads of only 0.41% (risk-neutral) and 0.62% (risk-averse), far below the 2.23% target, because a one-period bond has neither duration nor dilution and matching high spreads with it would require a counterfactually high default frequency; only the long-term models approach the target (1.65% and 2.04%). Second, *risk-averse lenders* govern the spread's *composition*. This is clearest in the long-term pair: moving from risk-neutral to risk-averse lenders *raises* the spread (1.65%→2.04%) while *lowering* the default rate (1.57%→1.16%), so the risk premium substitutes for realized default and the long-term risk-averse model matches the spread level at a realistic default frequency. All four produce countercyclical spreads. The welfare cost is larger under short-term debt (about 0.11%) than under long-term debt (0.06–0.09%): a one-period bond rolls its entire stock each quarter and is fully exposed to the geoeconomic repricing, whereas long-term debt is partly insulated. The long-term risk-averse model is the preferred specification.

[40] I target the change in skewness rather than its ratio because the model's pre-pressure income is a discretized Gaussian AR(1), whose cyclical component is essentially symmetric (skewness ≈ 0); a ratio would therefore be ill-posed against a near-zero denominator.

TABLE D2: Model comparison: short- vs. long-term debt, risk-neutral vs. risk-averse lenders

Moment	Data	ST (risk-neutral)	ST (risk-averse)	LT (risk-neutral)	LT (risk-averse)
Targeted moments					
Mean external debt/GDP	12.13%	10.7%	11.3%	13.2%	13.2%
Mean spread (ann.)	2.23%	0.41%	0.62%	1.65%	2.04%
Std(spread, ann.)	55 bp	57 bp	91 bp	98 bp	84 bp
Default prob. (ann.)	1.67%	0.42%	0.61%	1.57%	1.16%
Spread increase (ω)	24 bp	26 bp	50 bp	23 bp	23 bp
Untargeted moments					
$\text{Std}(C)/\text{Std}(Y)$	1.06	1.22	1.30	1.08	1.11
$\text{Corr}(\log C, \log Y)$	0.91	0.89	0.85	0.93	0.94
$\text{Corr}(\text{spread}, Y)$	-0.26	-0.46	-0.48	-0.65	-0.69
Model outcome					
Welfare cost (CEV)	—	-0.11%	-0.12%	-0.09%	-0.06%

Note: One-period bonds ($\delta = 1$) vs. long-term debt ($\delta = 0.0742$, $D \approx 12$q), *each separately calibrated* to the same Mexican targets. The short-term models use Stage-1 weights $(10, 4, 2, 3)$ and the long-term models $(10, 12, 2, 3)$ on (debt, spread, std, default rate); the two risk-averse columns use the main text's stochastic discount factor ($\chi = 8.75$) and quadratic output cost, so "LT (risk-averse)" is the model of the paper. Maturity sets the spread *level* (short-term $\leq 0.62\%$ vs. long-term up to 2.04%); within the long-term pair, risk-averse lenders raise the spread while lowering the default rate (1.57% $\to$ 1.16%), the premium substituting for realized default.

Simulation procedures

Long-run simulation (calibration targets)

To compute calibration targets in Stage 1, the model is simulated for $T_{\text{sim}} = 15{,}000$ quarters after a $T_{\text{burn}} = 2{,}000$-quarter burn-in starting from $b_0 = 0$ and the median income state. All moments are computed from the resulting time series. During default periods, the economy earns output $y - \phi(y)$ and the debt position remains at $b_0 = 0$ pending re-entry. Spread observations during default are excluded from the mean and standard deviation calculations, consistent with the empirical construction of EMBI spreads (which are not observed during default episodes).

Impulse response functions

The impulse responses reported in Figure 9 are constructed as follows.

Step 1: Burn-in (starting states). For each of $S = 10{,}000$ paths, I simulate a burn-in of $T_{\text{burn}} = 200$ quarters under $\Pi^{(0)}$ starting from $b_0 = 0$ and the median income state. The terminal state $(b_{\text{end}}, y_{\text{end}})$ of the burn-in is used as the starting state for the IRF path; the transitory shock m is redrawn i.i.d. each period. This ensures that starting states are drawn from a distribution close to the ergodic distribution of (b, y) under $\Pi^{(0)}$.

Step 2: Pre-shock period. Each path is simulated for a further $T_{\text{pre}} = 20$ quarters under $\Pi^{(0)}$, using the policy function and bond price schedule of the pre-shock equilibrium. These 20 quarters are used only to compute the pre-shock reference means for output, consumption, debt-to-GDP, and spreads; they are not displayed in the figures. Using $T_{\text{pre}} = 20$ quarters rather than a single observation mitigates finite-sample noise in the reference mean.

Step 3: Shock and post-shock path. At $t = 0$, the economy switches permanently to $\Pi^{(1)}(\omega^*, \zeta^*)$. From this point, the government uses the policy function and faces the bond price schedule of the distorted equilibrium. The path is simulated for $T_{\text{post}} = 7$ quarters ($t = 0$ to $t = 6$, corresponding to 2016Q3 through 2018Q1).

Step 4: Averaging and conditioning. All variables are averaged across paths at each date t over paths for which the economy is not in default at date t (the non-default conditioning used in the main text), or, alternatively, using the actual output-cost-adjusted values during default (the all-periods version, Figure D1). In the all-periods figure, the spread is still averaged over non-default periods only (it is undefined during default), while output, consumption, and debt-to-GDP record actual values during default episodes.

Step 5: Deviation and smoothing. Output and consumption deviations are expressed as percentage changes relative to their respective pre-shock non-default means (computed from the T_{pre} window). A two-quarter simple moving average is applied to both the model and the data series for output and consumption. External debt-to-GDP and spreads are shown in levels. The pre-shock model references are the T_{pre} ergodic means; the pre-shock data references are 2016Q3 values for debt-to-GDP and CDS, and the full-sample (2013Q1–2016Q2) means for the cyclical components of output and consumption.

Matched-pairs SCM premium

The model-equivalent SCM premium in Figure 10 is constructed using a matched-pairs design that eliminates idiosyncratic income variation from the estimated gap.

Step 1: Burn-in. Starting states (b_0, y_0) are drawn from the ergodic distribution of the pre-shock economy using the same burn-in procedure as above ($S = 10{,}000$ paths, $T_{\text{burn}} = 200$ quarters each).

Step 2: Matched pairs. For each path s, I simulate two trajectories from the same starting state (b_0^s, y_0^s) using the same sequence of random income draws:

- **Treated path:** T_{pre} quarters under $\Pi^{(0)}$, then permanent switch to $\Pi^{(1)}(\omega^*)$ at $t = 0$. The government uses policy $g^{\Pi^{(1)}}(b, y, m)$ and faces prices $q^{\Pi^{(1)}}(b', y)$ from $t = 0$ onward.
- **Counterfactual path:** $\Pi^{(0)}$ and $q^{\Pi^{(0)}}$ forever (the "synthetic control" within the model).

FIGURE D1. Impulse responses to a geoeconomic regime switch (all periods)

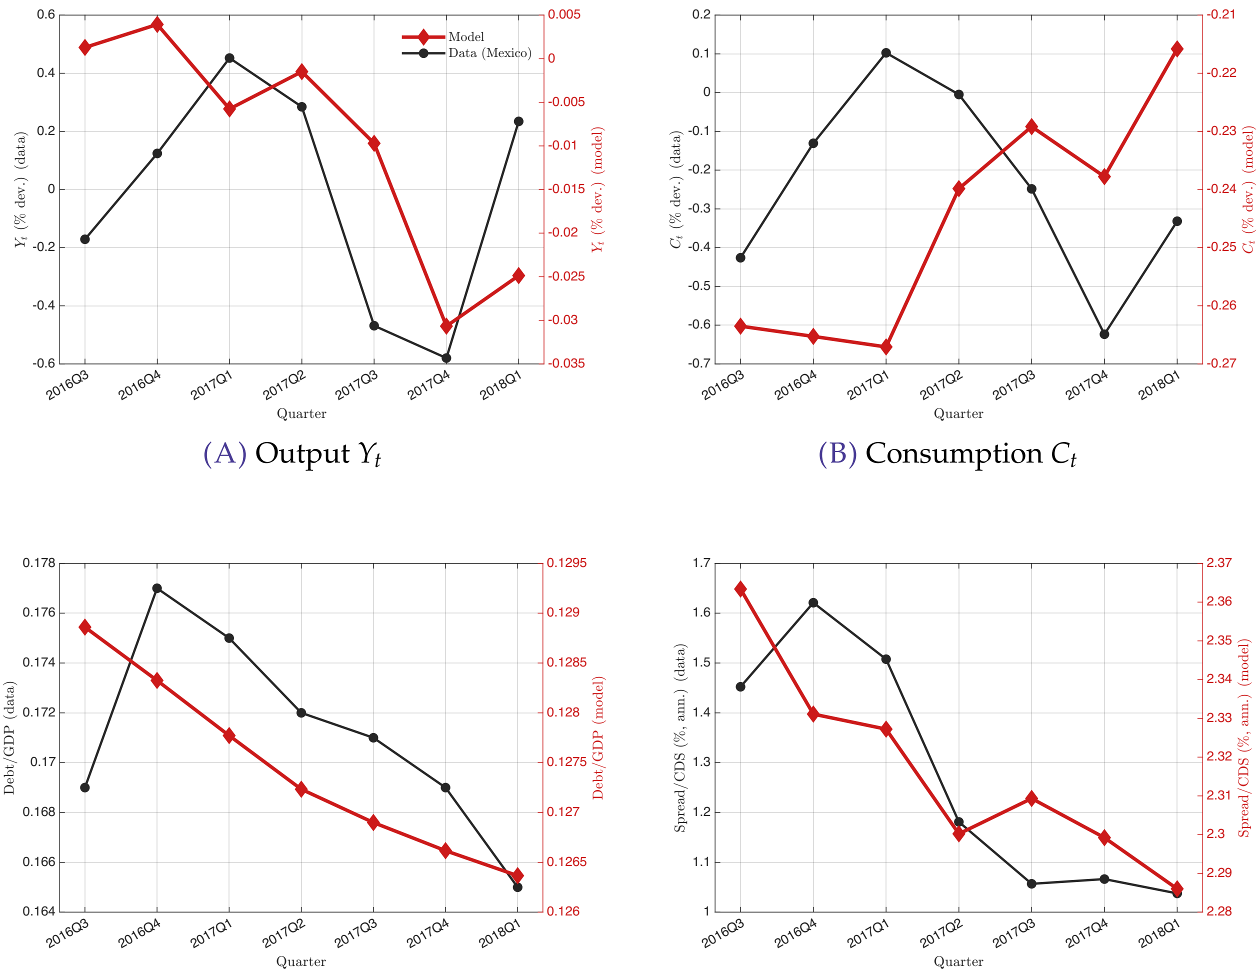

(A) Output Y_t

(B) Consumption C_t

(C) External debt/GDP

(D) Sovereign spread (CDS)

Note: All-periods version of Figure 9 (output, consumption, and debt/GDP record actual values during default). Red lines with diamond markers: model average over $S = 10{,}000$ paths (right axis). Black lines with circle markers: Mexican data (left axis). Output and consumption: % deviation from pre-shock mean, two-quarter moving average. External debt/GDP and CDS spread: levels.

Income transitions are drawn from the treated path's CDF at each period, and the same draw is used for both paths. This ensures that any difference in spreads between the two paths reflects only the regime switch, not differences in income realizations.

Step 3: Gap computation. The model-implied quarterly premium at date t is

$$\widehat{\Delta s}_t^{\text{model}} = \frac{1}{|\mathcal{S}_t^{\text{nd}}|} \sum_{s \in \mathcal{S}_t^{\text{nd}}} \left(z_t^{\Pi^{(1)},s} - z_t^{\Pi^{(0)},s} \right) \times 10{,}000 \text{ bp},$$

where $\mathcal{S}_t^{\text{nd}}$ is the set of paths for which neither the treated nor the counterfactual economy is in default at date t, and $z_t^{\Pi^{(1)},s}$ ($z_t^{\Pi^{(0)},s}$) is the annualized spread of the treated (counterfactual) path at date t. Conditioning on mutual non-default ensures that the gap reflects pricing differences rather than regime transitions.

Additional model details

Construction of the tilted transition matrix

The post-shock target distribution Π^{target} in equation (18) concentrates mass on the bottom tercile of the income grid, $\mathcal{L} = \{y_1, \ldots, y_7\}$ (states 1 through 7 of the 21-state grid):

$$\Pi^{\text{target}}(y'|y;\, \zeta) = \lambda_j(\zeta)\, \mathbb{1}\{y' \in \mathcal{L}\}, \qquad \lambda_j(\zeta) = \frac{\zeta^{j-1}}{\sum_{k=1}^{|\mathcal{L}|} \zeta^{k-1}}, \tag{D.2}$$

where j indexes the state within the bottom tercile. The parameter $\zeta > 0$ governs the *shape* of mass allocation within the tercile: $\zeta = 1$ places equal weight on all seven states (uniform); $\zeta < 1$ concentrates more mass on the worst states (e.g., y_1); and $\zeta > 1$ concentrates mass on the least adverse states within the tercile (e.g., y_7). The matrix $\Pi^{(1)}$ is row-stochastic for all $\omega \in [0, 1)$ and $\zeta > 0$. When $\omega = 0$, the post-shock process coincides with the baseline, and the model reduces to the standard Arellano framework with long-term debt. The special case $\zeta = 1$ takes the particularly transparent form

$$\Pi^{(1)}(y'|y;\, \omega) = \begin{cases} \Pi^{(0)}(y'|y) + \omega\left(\dfrac{1}{|\mathcal{L}|} - \Pi^{(0)}(y'|y)\right), & y' \in \mathcal{L}, \\ (1-\omega)\,\Pi^{(0)}(y'|y), & y' \notin \mathcal{L}, \end{cases} \tag{D.3}$$

which shows that the regime switch uniformly inflates the probability of each bad state by a fraction ω of the remaining probability mass.

Policy functions and the spread schedule

Figure D2 shows the associated debt policy and default region under the baseline chain $\Pi^{(0)}$. Panel A plots next-period debt $-b'(b, y)$ against current debt for three income levels: the government rolls over and gradually pays down debt in normal states, but at high debt and low income the policy drops to zero, marking default and the reset of the debt position. Panel B displays the default region as the default probability integrated over the transitory shock, $\mathbb{E}_m\, d(b, y, m)$, over the $(-b, z)$ plane: default is concentrated in the bottom-right corner (i.e., high debt and low persistent income), and the regime switch enlarges this region by raising the value of default relative to repayment at intermediate debt levels.

Figure D3 plots the equilibrium spread as a function of the debt chosen, $-b'$, for a low, median, and high value of the persistent income state, under the baseline chain $\Pi^{(0)}$ (solid) and the post-shock chain $\Pi^{(1)}$ (dashed). Spreads rise with debt and fall with income, as in the standard model. The regime switch shifts the entire schedule upward (i.e., for a given debt level the sovereign faces a higher spread) because lenders price in the elevated probability of low future income; the upward shift is largest at low income, where the additional mass is concentrated.

FIGURE D2. Debt policy and default region.

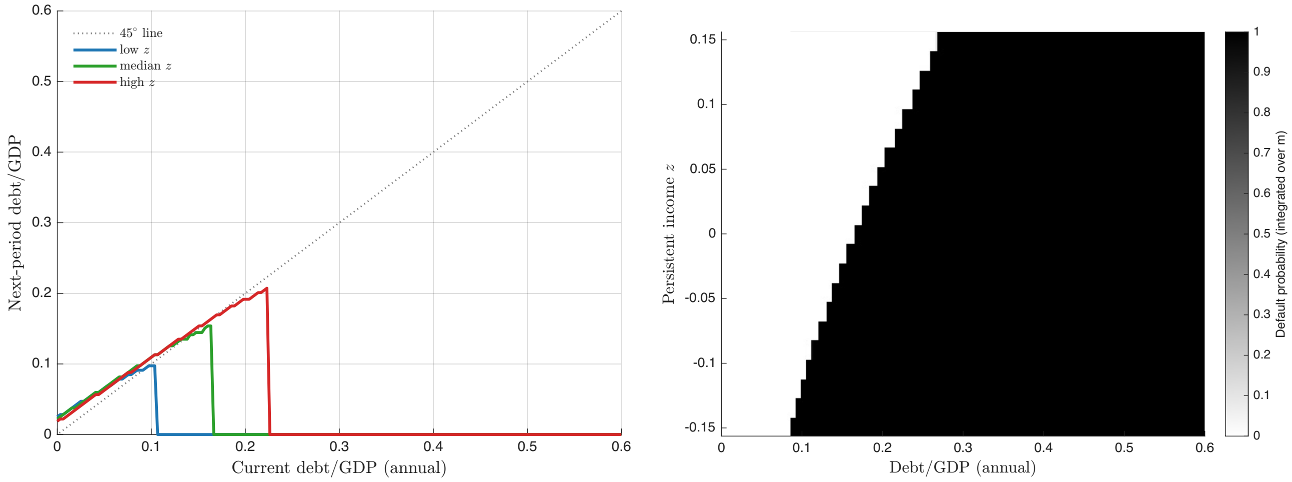

(A) Debt policy (B) Default region

Note: Panel A: next-period debt policy $-b'(b, z)$ at the median transitory shock, for low, median, and high persistent income z, together with the 45° line. Panel B: default region, measured as the default probability integrated over the transitory shock, $\mathbb{E}_m\, d(b, z, m)$, over debt $-b$ and persistent income z (darker = higher). Baseline chain $\Pi^{(0)}$.

FIGURE D3. Spread schedule under $\Pi^{(0)}$ and $\Pi^{(1)}$.

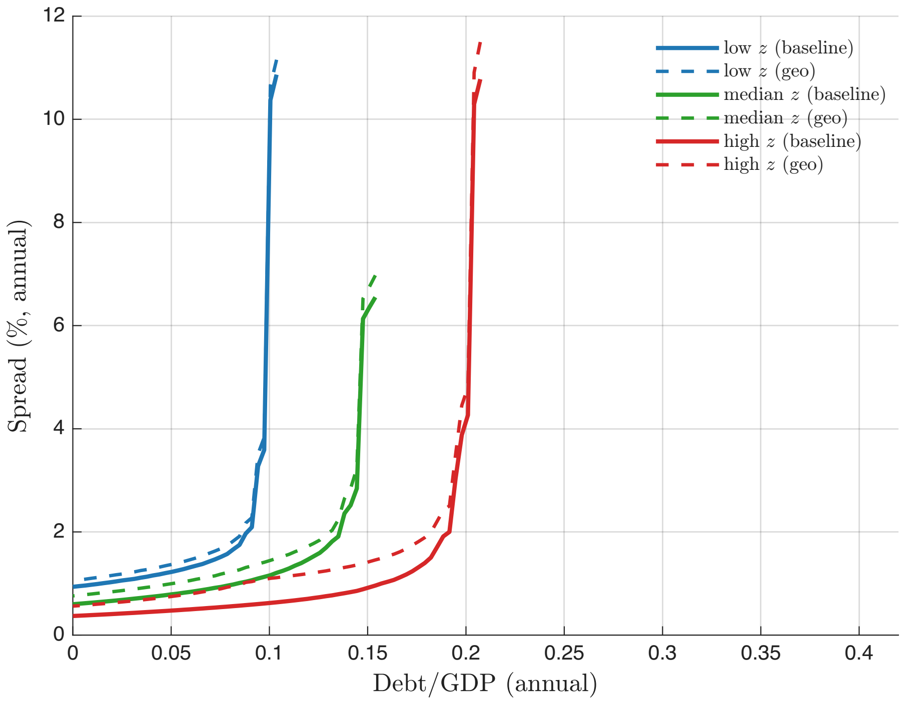

Note: Equilibrium annualized spread as a function of next-period debt $-b'$ for low, median, and high persistent income z, under the baseline transition matrix $\Pi^{(0)}$ (solid) and the post-shock matrix $\Pi^{(1)}$ (dashed, $\omega = \omega^* = 0.0003$, $\zeta = \zeta^* = 0.513$).

Welfare: derivation of the CEV and robustness

The permanent CEV is the proportional change in consumption that makes a household indifferent between living forever under the baseline regime ($\Pi^{(0)}$) and living forever under

TABLE D3: Welfare cost across within-tercile allocations ζ

ζ	CEV permanent (%)	CEV transitory (%)	Avg. cons. contraction (%)
0.25	0.077	0.019	0.10
0.513	0.063	0.016	0.10
1.00	0.050	0.017	0.08
2.00	0.033	0.008	0.07
5.00	0.027	0.015	0.07

Note: $\omega = \omega^* = 0.0003$ and $S = 20{,}000$ simulated paths throughout. "CEV permanent" is the consumption-equivalent variation of a permanent regime switch (aggregate closed form); "CEV transitory" is the CEV of a seven-quarter switch; "Avg. cons. contraction" is the average per-quarter percentage decline in consumption over the seven-quarter window. The calibrated within-tercile shape is $\zeta^* = 0.513$; $\zeta = 1$ corresponds to a uniform allocation.

the post-shock regime ($\Pi^{(1)}$). Define $\lambda(b, y, m)$ as the fraction of consumption that must be given to a household under $\Pi^{(0)}$ to make it indifferent between the two regimes; formally, λ solves

$$\mathbb{E}_0^{\Pi^{(0)}}\left[\sum_{t=0}^{\infty}\beta^t \frac{\left[(1+\lambda)\, c_t^{\Pi^{(0)}}\right]^{1-\gamma}}{1-\gamma}\right] = \mathbb{E}_0^{\Pi^{(1)}}\left[\sum_{t=0}^{\infty}\beta^t \frac{\left[c_t^{\Pi^{(1)}}\right]^{1-\gamma}}{1-\gamma}\right], \tag{D.4}$$

where $\{c_t^{\Pi^{(0)}}\}$ and $\{c_t^{\Pi^{(1)}}\}$ are the equilibrium consumption streams under the baseline and post-shock regimes and $\mathbb{E}_0^{\Pi}$ is the expectation under the transition matrix Π. Manipulating (D.4) yields the closed-form expression

$$\lambda(b, y, m) = \left[\frac{V^{\Pi^{(1)}}(b, y, m)}{V^{\Pi^{(0)}}(b, y, m)}\right]^{\frac{1}{1-\gamma}} - 1. \tag{D.5}$$

A negative λ indicates that the post-shock regime delivers lower welfare. I report the welfare cost as $|\lambda| \times 100$ percent of permanent consumption, averaged over non-default states (b, y, m) weighted by the stationary distribution under $\Pi^{(0)}$.

Although ζ is calibrated to the skewness change ($\zeta^* = 0.513$), it is instructive to trace how the welfare cost varies with the within-tercile shape. Table D3 reports both CEVs and the average consumption contraction across $\zeta \in \{0.25, 0.5, 1.0, 2.0, 5.0\}$, holding $\omega = \omega^*$ fixed so that the total additional mass on bad states is unchanged. Lower ζ concentrates the additional mass on the worst states (a fatter left tail) and raises all three quantities, while higher ζ shifts it toward the least-bad state and lowers them. The permanent CEV ranges from 0.077% ($\zeta = 0.25$) to 0.027% ($\zeta = 5$), the transitory CEV from about 0.008% to 0.019%, and the average consumption contraction from 0.07% to 0.10%; the calibrated $\zeta^* = 0.513$ is the second row of the table. Even at the most conservative allocation, the welfare loss remains economically meaningful.

Appendix E: Derivation of sufficient statistic

The sufficient-statistic approach of the main text is the household's intertemporal optimality condition from the government problem of Section 4, evaluated under three simplifications that render the calculation model-light and let the consumption loss be read directly from asset prices. In Section 4, a benevolent government maximizes household welfare subject to its flow budget constraint, issuing long-term bonds ($\delta = 0.0742$) that are priced by risk-averse foreign lenders, and choosing each period between repayment and default. Two features of that problem are relaxed here: the first because it would otherwise preclude a well-defined Euler equation, the second to keep the calculation tractable. First, endogenous default: the government's value function is the upper envelope $V = \max\{V^R, X\}$ of the repayment and autarky values, so it is kinked at the default frontier and marginal utility is discontinuous across a default event; a smooth intertemporal condition therefore does not hold globally. I take the default probability as exogenous, recovering it from the observed CDS, so default is no longer strategic and the household's optimality condition is only well-defined along the repayment path. Second, maturity: long-term debt does not preclude an Euler equation, but the resulting condition depends on the outstanding debt stock and the government's issuance policy (endogenous objects with no observable high-frequency counterpart), so their counterfactual paths cannot be recovered from the data. I therefore use a one-period bond ($\delta = 1$), for which the entire stock is refinanced each period and the intertemporal condition involves only the current price q_t, which I read from the CDS. Together, these two simplifications also make the government a price-taker, so no separate assumption is required. For a one-period bond the price depends on the government's own issuance b' only through the default probability, $\partial q / \partial b' = (\partial q / \partial \pi)(\partial \pi / \partial b')$, so once π is exogenous ($\partial \pi / \partial b' = 0$) the price-schedule term $\partial q / \partial b'$ vanishes from the government's first-order condition. Under long-term debt this term would not vanish, because the continuation price introduces a second, debt-dilution channel through which q depends on b'; the one-period bond shuts that channel down. The one-period simplification is not innocuous, however: a one-period bond exposes the whole stock to the repricing every period, whereas long-term debt reprices only a fraction δ of it, so the short-term formulation overstates the consumption contraction relative to the long-term model; Short-term debt model quantifies this maturity wedge.

Under these simplifications, and retaining the same CRRA preferences, discount factor β, and risk aversion γ as in Section 4, the household's problem reduces to

$$\max_{\{C_t, D_{t+1}\}_{t \geq 0}} \mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \right], \quad \text{with } \beta \in (0,1) \text{ and } \gamma > 0,$$

$$s.t. \quad C_t + q_t D_{t+1} = Y_t + D_t,$$

where C_t is consumption, $D_{t+1} = -b_{t+1}$ is new one-period borrowing, Y_t the stochastic income stream, and q_t the exogenous discount price of debt. The Lagrangian associated with this problem is given by:

$$\mathcal{L}(C_t, D_{t+1}; \lambda_t) = \mathbb{E}_t \left\{ \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} + \lambda_t [Y_t + D_t - C_t - q_t D_{t+1}] \right\},$$

which leads to the following first-order conditions (FOCs):

$$(C_t): \frac{\partial \mathcal{L}}{\partial C_t} = 0 \Leftrightarrow \beta^t C_t^{-\gamma} = \lambda_t,$$

$$(D_{t+1}) : \frac{\partial \mathcal{L}}{\partial D_{t+1}} = 0 \Leftrightarrow -q_t \lambda_t + \mathbb{E}_t \left[\lambda_{t+1} \right] = 0.$$

Combining the previous FOCs, we derive the following Euler equation:

$$q_t = \beta \mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right]. \tag{E.1}$$

The price q_t that appears in the Euler equation (E.1) is the observed market price of the one-period bond; I keep it as such throughout the derivation and defer to the end how it is recovered from the data. To pin down the effect on the level of consumption, I take logs of (E.1) under the assumption that consumption growth is conditionally log-normal:

$$\log(q_t) = \log(\beta) + \log \left(\mathbb{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \right] \right),$$

$$\Leftrightarrow \log(q_t) = \log(\beta) - \gamma \mathbb{E}_t \left[\Delta c_{t+1} \right] + \frac{\gamma^2}{2} \mathrm{Var}_t \left(\Delta c_{t+1} \right).$$

Since $\Delta c_{t+1} = \log(C_{t+1}) - \log(C_t)$, we can find an expression for $c_t = \log(C_t)$ by expanding the previous expression:

$$c_t = \frac{1}{\gamma} \left[\log(q_t) - \log(\beta) \right] + \mathbb{E}_t \left[c_{t+1} \right] - \frac{\gamma}{2} \mathrm{Var}_t \left(\Delta c_{t+1} \right),$$

which, after iterating forward 1 period ahead and using the Law of Iterated Expectations, we get:

$$\begin{aligned} c_t = \frac{1}{\gamma} \left[\log(q_t) - \log(\beta) \right] + \frac{1}{\gamma} \mathbb{E}_t \left[\log(q_{t+1}) - \log(\beta) \right] + \mathbb{E}_t \left[c_{t+2} \right] \\ - \frac{\gamma}{2} \mathbb{E}_t \left[\mathrm{Var}_{t+1}(\Delta c_{t+2}) \right] - \frac{\gamma}{2} \mathrm{Var}_t(\Delta c_{t+1}). \end{aligned}$$

Generalizing by iterating forward T periods ahead, we get:

$$c_t = \frac{1}{\gamma} \sum_{j=0}^{T} \mathbb{E}_t \left[\log(q_{t+j}) \right] - \frac{(T+1)}{\gamma} \log(\beta) - \frac{\gamma}{2} \sum_{j=0}^{T} \mathbb{E}_t \left[\mathrm{Var}_{t+j}(\Delta c_{t+j+1}) \right] + \mathbb{E}_t \left[c_{t+T+1} \right].$$

Now let c_t and $\tilde{c}_t$ denote log-consumption under the observed and counterfactual paths. Taking the difference between the two, the constant $\frac{(T+1)}{\gamma} \log(\beta)$ cancels, and we arrive at the expression displayed in the main text:

$$\begin{aligned} c_t - \tilde{c}_t = & \underbrace{\frac{1}{\gamma} \left(\sum_{j=0}^{T} \mathbb{E}_t \left[\log q_{t+j} \right] - \sum_{j=0}^{T} \mathbb{E}_t \left[\log \tilde{q}_{t+j} \right] \right)}_{\text{(1) Pricing component}} \\ & \underbrace{- \frac{\gamma}{2} \left(\sum_{j=0}^{T} \mathbb{E}_t \left[\mathrm{Var}_{t+j} \left(\Delta c_{t+j+1} \right) \right] - \sum_{j=0}^{T} \mathbb{E}_t \left[\mathrm{Var}_{t+j} \left(\Delta \tilde{c}_{t+j+1} \right) \right] \right)}_{\text{(2) Uncertainty component}} + \underbrace{\left(\mathbb{E}_t \left[c_{t+T+1} \right] - \mathbb{E}_t \left[\tilde{c}_{t+T+1} \right] \right)}_{\text{(3) Transitory vs. permanent component}}. \end{aligned} \tag{E.2}$$

The pricing component operates entirely through the market price q_t, which already embeds the risk premium lenders demand, including the risk aversion of the Section 4 lenders, so it requires no assumption on their price of income risk (i.e., χ). To evaluate it, I recover q_t from the observed and synthetic CDS. By no-arbitrage, the price of a zero-recovery one-period claim admits the risk-neutral representation

$$q_t = \frac{1-\pi_t}{1+r_f}, \tag{E.3}$$

where π_t is the risk-neutral (pricing) default probability, i.e., the default probability under the pricing measure, which is precisely what a CDS spread reports. Equation (E.3) is the change-of-measure form of the general pricing condition $q_t = \mathbb{E}_t[\mathcal{M}_{t+1}(1-d_{t+1})]$ and therefore holds for any lender preferences; it is not an assumption that lenders are risk-neutral. Under the risk-averse lenders of Section 4, π_t exceeds the physical default probability by the risk premium, so the observed spread, and hence q_t, already reflects that premium; only if lenders were risk-neutral would π_t coincide with the physical default probability. I recover π_t (and its counterfactual $\tilde{\pi}_t$) from the observed and synthetic 5-year CDS under a constant-hazard, zero-recovery convention. Under (E.3), the pricing component can equivalently be written as $\frac{1}{\gamma}\sum_{j=0}^{T}\mathbb{E}_t\left[\log(1-\pi_{t+j}) - \log(1-\tilde{\pi}_{t+j})\right]$, since $\log q_{t+j} = \log(1-\pi_{t+j}) - \log(1+r_f)$ and the risk-free rate is common to both paths.

The previous expression can be further simplified by using the fact that $\text{Var}_{t+j}(\Delta c_{t+j+1})$ represents the conditional volatility which is estimated via a GARCH(1,1):

$$\text{Var}_{t+j}(\Delta c_{t+j+1}) \equiv h_{t+j} = \varpi + a_1\varepsilon^2_{t+j-1} + b_1 h_{t+j-1}, \quad \text{with} \quad \varepsilon_{t+j-1} \overset{i.i.d.}{\sim} \mathcal{N}(0,1),$$

where (ϖ, a_1, b_1) denote the GARCH parameters.[41] Iterating backwards j periods, we can write h_{t+j} as a function of (a_1, b_1, ϖ) and h_t:

$$h_{t+j} = \varpi \sum_{i=0}^{j-1}\left[\prod_{k=1}^{i}(a_1\varepsilon^2_{t+j-k} + b_1)\right] + \left[\prod_{k=1}^{j}(a_1\varepsilon^2_{t+j-k} + b_1)\right] h_t, \tag{E.4}$$

now by taking the conditional expectation on information available at period t and setting $a_1 + b_1 = \phi$, we get:

$$\mathbb{E}_t[h_{t+j}] = \varpi \sum_{i=0}^{j-1}\phi^i + \phi^j h_t = \varpi\frac{1-\phi^j}{1-\phi} + \phi^j h_t.$$

Finally, by summing from $j = 0$ to T, we get that

$$\sum_{j=0}^{T}\mathbb{E}_t[\text{Var}_{t+j}(\Delta c_{t+j+1})] \equiv \sum_{j=0}^{T}\mathbb{E}_t[h_{t+j}] = \varpi\frac{(T+1)}{1-\phi} + \left(h_t - \frac{\varpi}{1-\phi}\right)\left(\frac{1-\phi^{T+1}}{1-\phi}\right).$$

[41] I use (ϖ, a_1, b_1) rather than the conventional (ω, α, β) to avoid conflict with the tilt intensity ω, the tilt shape α, and the discount factor β of the main text.

Appendix F: Filtering algorithm and the consumption cost of default

From the homoskedastic case of the LRR framework of Bansal and Yaron (2004), we have that consumption growth Δc_{t+1} and stock returns r_{t+1} are related via an unobservable and persistent component x_{t+1}:

$$\begin{aligned} \Delta c_{t+1} &= \mu_c + x_t + \sigma \eta_{t+1}, \\ x_{t+1} &= \rho x_t + \varphi_e \sigma e_{t+1}, \\ r_{t+1} &= \mu_r + \phi x_t + \varphi_r \sigma u_{t+1}, \end{aligned} \tag{F.1}$$

with η_{t+1}, e_{t+1}, and $u_{t+1} \overset{i.i.d}{\sim} \mathcal{N}(0,1)$. In most applications, this model is estimated using quarterly or monthly data on Δc_{t+1} and r_{t+1} with x_{t+1} being filtered out using the Kalman filter. Instead, in this paper, I use this structural relationship and introduce an aggregation constraint with the objective of measuring daily consumption growth at high-frequency by relying on daily stock returns and monthly consumption growth data. The rationale behind the introduction of this aggregation constraint is to enforce, up to some measurement error, that the filtered data matches the observed data at the highest available frequency (e.g. monthly).

In particular, I estimate the following state-space model with a monthly aggregation constraint:

$$\begin{aligned} \boldsymbol{z}_{t+1} &= \boldsymbol{\mu_z} + \boldsymbol{A}\boldsymbol{z}_t + \boldsymbol{B}\boldsymbol{u}_{t+1}, \quad \boldsymbol{u}_{t+1} \sim \mathcal{N}(0, \boldsymbol{I}), \\ r_{t+1} &= \mu_r + \boldsymbol{H}\boldsymbol{z}_t + \sigma_r \varepsilon_{t+1}, \quad \varepsilon_{t+1} \sim \mathcal{N}(0,1), \\ \Delta c_M &= \sum_{i \in M} \Delta c_i + \sigma_c \xi_M, \quad \xi_M \sim \mathcal{N}(0,1), \end{aligned} \tag{F.2}$$

with

$$\boldsymbol{z}_t = \begin{bmatrix} \Delta c_t & x_t \end{bmatrix}^T, \quad \boldsymbol{\mu_z} = \begin{bmatrix} \mu_c & 0 \end{bmatrix}^T, \quad \boldsymbol{A} = \begin{bmatrix} 0 & 1 \\ 0 & \rho \end{bmatrix}, \quad \boldsymbol{B} = \begin{bmatrix} \sigma_{\Delta c} & 0 \\ 0 & \sigma_x \end{bmatrix}, \quad \boldsymbol{H} = \begin{bmatrix} 0 & \phi \end{bmatrix},$$

where $\boldsymbol{z}_t$ corresponds to the vector of state variables, and r_{t+1} and Δc_M correspond to the measurement equations. I estimate the parameters of the state-space model using Maximum Likelihood within the Kalman filter.

Augmented-state formulation

The monthly aggregation constraint in (F.2) involves a sum of latent daily states across all trading days within a month, which cannot be handled as a standard Kalman measurement on the current state alone. To address this, I augment the state vector with a running within-month sum s_t that accumulates daily consumption growth and resets at the start of each month. The augmented state vector is

$$\tilde{\boldsymbol{z}}_t = \begin{bmatrix} \Delta c_t & x_t & s_t \end{bmatrix}^T,$$

with augmented drift, process-noise matrix, and measurement vectors:

$$\tilde{\boldsymbol{\mu}}_z = \begin{bmatrix} \mu_c \\ 0 \\ \mu_c \end{bmatrix}, \qquad \tilde{\boldsymbol{B}} = \begin{bmatrix} \sigma_{\Delta c} & 0 \\ 0 & \sigma_x \\ \sigma_{\Delta c} & 0 \end{bmatrix}, \qquad \tilde{\boldsymbol{H}}_r = \begin{bmatrix} 0 & \phi & 0 \end{bmatrix}, \qquad \tilde{\boldsymbol{H}}_c = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}.$$

The transition matrix $\tilde{\boldsymbol{A}}_{t+1}$ switches between two regimes depending on whether day $t+1$ falls within the same month or begins a new one:

$$\tilde{\boldsymbol{A}}_{\text{within}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & \rho & 0 \\ 0 & 1 & 1 \end{bmatrix}, \qquad \tilde{\boldsymbol{A}}_{\text{reset}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & \rho & 0 \\ 0 & 1 & 0 \end{bmatrix}.$$

Within a month, the third row of $\tilde{\boldsymbol{A}}_{\text{within}}$ yields $s_{t+1} = s_t + \mu_c + x_t + \sigma_{\Delta c} u_{1,t+1} = s_t + \Delta c_{t+1}$, so the running sum accumulates daily consumption growth. On the first trading day of a new month, $\tilde{\boldsymbol{A}}_{\text{reset}}$ zeros out the persistence of s_t, giving $s_{t+1} = \Delta c_{t+1}$. At the last trading day of month M, the third state component satisfies $s_{t_M^{\text{EOM}}} = \sum_{i \in M} \Delta c_i$, and the monthly aggregation constraint becomes a standard Kalman measurement:

$$\Delta c_M = \tilde{\boldsymbol{H}}_c \, \tilde{\boldsymbol{z}}_{t_M^{\text{EOM}}} + \sigma_c \xi_M.$$

This formulation allows the monthly observation to be incorporated via a proper Kalman state and covariance update, with the correct innovation variance $S_m^{(c)} = \tilde{\boldsymbol{H}}_c \, \tilde{P}_{t|t} \, \tilde{\boldsymbol{H}}_c' + \sigma_c^2$ that accounts for the full joint covariance of the accumulated daily states within the month.

Algorithm

The estimation proceeds in two stages. In the first stage, a forward Kalman filter produces filtered states $\tilde{\boldsymbol{z}}_{t|t}$ and the log-likelihood used for Maximum Likelihood estimation. In the second stage, a Rauch-Tung-Striebel (RTS) backward smoother refines these estimates to produce $\tilde{\boldsymbol{z}}_{t|T} = \mathbb{E}[\tilde{\boldsymbol{z}}_t \mid r_{1:T}, \Delta c_{M,1:M}]$, the optimal state estimates conditional on the full sample of daily returns and monthly consumption observations. The smoother redistributes the monthly information across all trading days within each month using the joint covariance structure implied by the augmented dynamics, yielding the smoothed daily consumption growth series $\Delta c_{t|T}$ used for constructing the log-consumption level and for the local projection analysis.

Algorithm 2 Augmented-state Kalman Filter (Forward Pass)

Require: Parameters $\Theta = (\mu_r, \mu_c, \rho, \phi, \sigma_{\Delta c}, \sigma_x, \sigma_r, \sigma_c)$; daily returns $\{r_t\}_{t=1}^{T}$; monthly observations $\{\Delta c_m\}_{m=1}^{M}$ (observed only at EOM); month-boundary indicator $\mathcal{R}(t) = \mathbb{1}\{\text{day } t \text{ starts a new month}\}$.

Ensure: Filtered states $\{\tilde{\boldsymbol{z}}_{t|t}\}$, covariances $\{\tilde{P}_{t|t}\}$, predicted states $\{\tilde{z}^{\text{pred}}_{t|t-1}\}$, predicted covariances $\{\tilde{P}^{\text{pred}}_{t|t-1}\}$, stored transition matrices $\{\tilde{\boldsymbol{A}}^{\text{store}}_t\}$, negative log-likelihood NLL.

Initialize: set $\tilde{z}_{0|0}$, $\tilde{P}_{0|0}$ (prior), $\text{logL} \leftarrow 0$.

Define augmented model matrices $\tilde{\boldsymbol{A}}_{\text{within}}$, $\tilde{\boldsymbol{A}}_{\text{reset}}$, $\tilde{\boldsymbol{\mu}}_z$, $\tilde{\boldsymbol{B}}$, $\tilde{\boldsymbol{H}}_r$, $\tilde{\boldsymbol{H}}_c$ as above. Set $\tilde{\boldsymbol{Q}} = \tilde{\boldsymbol{B}}\tilde{\boldsymbol{B}}'$.

for $t = 1$ **to** T **do**

  Select transition matrix: $\tilde{\boldsymbol{A}}_t \leftarrow \tilde{\boldsymbol{A}}_{\text{reset}}$ if $\mathcal{R}(t) = 1$ and $t > 1$; otherwise $\tilde{\boldsymbol{A}}_t \leftarrow \tilde{\boldsymbol{A}}_{\text{within}}$.

  Prediction:

$$\tilde{\boldsymbol{z}}_{t|t-1} \leftarrow \tilde{\boldsymbol{\mu}}_z + \tilde{\boldsymbol{A}}_t\, \tilde{\boldsymbol{z}}_{t-1|t-1}, \qquad \tilde{P}_{t|t-1} \leftarrow \tilde{\boldsymbol{A}}_t\, \tilde{P}_{t-1|t-1}\, \tilde{\boldsymbol{A}}_t' + \tilde{\boldsymbol{Q}}.$$

  Store: $\tilde{z}^{\text{pred}}_{t|t-1} \leftarrow \tilde{\boldsymbol{z}}_{t|t-1}$, $\tilde{P}^{\text{pred}}_{t|t-1} \leftarrow \tilde{P}_{t|t-1}$, $\tilde{\boldsymbol{A}}^{\text{store}}_t \leftarrow \tilde{\boldsymbol{A}}_t$.

  Daily return update:

$$\begin{aligned} \tilde{r}_t &\leftarrow r_t - \big(\mu_r + \tilde{\boldsymbol{H}}_r\, \tilde{\boldsymbol{z}}_{t|t-1}\big), \qquad S_t^{(r)} \leftarrow \tilde{\boldsymbol{H}}_r\, \tilde{P}_{t|t-1}\, \tilde{\boldsymbol{H}}_r' + \sigma_r^2, \\ \tilde{\boldsymbol{K}}_t^{(r)} &\leftarrow \tilde{P}_{t|t-1}\, \tilde{\boldsymbol{H}}_r' \,\big(S_t^{(r)}\big)^{-1}, \\ \tilde{\boldsymbol{z}}_{t|t} &\leftarrow \tilde{\boldsymbol{z}}_{t|t-1} + \tilde{\boldsymbol{K}}_t^{(r)}\, \tilde{r}_t, \qquad \tilde{P}_{t|t} \leftarrow \big(\boldsymbol{I} - \tilde{\boldsymbol{K}}_t^{(r)}\, \tilde{\boldsymbol{H}}_r\big)\, \tilde{P}_{t|t-1}. \end{aligned}$$

  Accumulate: $\text{logL} \leftarrow \text{logL} - \frac{1}{2}\Big(\log(2\pi) + \log S_t^{(r)} + \tilde{r}_t^2 / S_t^{(r)}\Big)$.

  if t is end-of-month for month m **then**

    Monthly update:

$$\begin{aligned} \tilde{\Delta c}_m &\leftarrow \Delta c_m - \tilde{\boldsymbol{H}}_c\, \tilde{\boldsymbol{z}}_{t|t}, \qquad S_m^{(c)} \leftarrow \tilde{\boldsymbol{H}}_c\, \tilde{P}_{t|t}\, \tilde{\boldsymbol{H}}_c' + \sigma_c^2, \\ \tilde{\boldsymbol{K}}_m^{(c)} &\leftarrow \tilde{P}_{t|t}\, \tilde{\boldsymbol{H}}_c' \,\big(S_m^{(c)}\big)^{-1}, \\ \tilde{\boldsymbol{z}}_{t|t} &\leftarrow \tilde{\boldsymbol{z}}_{t|t} + \tilde{\boldsymbol{K}}_m^{(c)}\, \tilde{\Delta c}_m, \qquad \tilde{P}_{t|t} \leftarrow \big(\boldsymbol{I} - \tilde{\boldsymbol{K}}_m^{(c)}\, \tilde{\boldsymbol{H}}_c\big)\, \tilde{P}_{t|t}. \end{aligned}$$

    Accumulate: $\text{logL} \leftarrow \text{logL} - \frac{1}{2}\Big(\log(2\pi) + \log S_m^{(c)} + \tilde{\Delta c}_m^2 / S_m^{(c)}\Big)$.

  end if

end for

Finalize: $\text{NLL} \leftarrow -\,\text{logL}$.

Algorithm 3 Rauch-Tung-Striebel Backward Smoother and Reconstruction

Require: Filtered states $\{\tilde{z}_{t|t}\}$, covariances $\{\tilde{P}_{t|t}\}$, predicted states $\{\tilde{z}^{\text{pred}}_{t|t-1}\}$, predicted covariances $\{\tilde{P}^{\text{pred}}_{t|t-1}\}$, stored transition matrices $\{\tilde{A}^{\text{store}}_t\}$ from Algorithm 2; initial log-consumption $c_0 = \log I_0$.

Ensure: Smoothed states $\{\tilde{z}_{t|T}\}$, smoothed daily growth $\{\Delta c_{t|T}\}$, log consumption $\{c_t\}$.

1: **Initialize:** $\tilde{z}^s_{T|T} \leftarrow \tilde{z}_{T|T}$, $\tilde{P}^s_{T|T} \leftarrow \tilde{P}_{T|T}$.

2: **for** $t = T-1$ **down to** 1 **do**

3: **Smoother gain:**

$$J_t \leftarrow \tilde{P}_{t|t} \left(\tilde{A}^{\text{store}}_{t+1}\right)' \left(\tilde{P}^{\text{pred}}_{t+1|t}\right)^{-1}.$$

4: **Smoothed state and covariance:**

$$\tilde{z}^s_{t|T} \leftarrow \tilde{z}_{t|t} + J_t\left(\tilde{z}^s_{t+1|T} - \tilde{z}^{\text{pred}}_{t+1|t}\right),$$
$$\tilde{P}^s_{t|T} \leftarrow \tilde{P}_{t|t} + J_t\left(\tilde{P}^s_{t+1|T} - \tilde{P}^{\text{pred}}_{t+1|t}\right)J_t'.$$

5: **end for**

6: **Extract smoothed daily growth:** $\Delta c_{t|T} \leftarrow \tilde{z}^{s,(1)}_{t|T}$ for $t = 1, \ldots, T$.

7: **Reconstruct log consumption level:**

$$c_t \leftarrow c_0 + \sum_{s=1}^{t} \Delta c_{s|T}, \quad t = 1, \ldots, T.$$

Smoothed daily consumption series

Figure F1 displays the result of the filtering and smoothing procedure. The black markers are the observed monthly *INEGI* private-consumption index; the blue line is the exponential of the smoothed daily series, $C_t = \exp\{c_t\}$, reconstructed from the smoothed daily growth rates as $c_t = c_0 + \sum_{s=1}^{t} \Delta c_s$ with the initial condition $c_0 = \log I_0$, where I_0 is the index value in the period immediately before the start of the sample. Maximum-likelihood estimation returns $\hat{\Theta} = (\mu_r, \mu_c, \rho, \phi, \sigma_{\Delta c}, \sigma_x, \sigma_r, \sigma_c) = (0.0002, 0.0001, 0.8968, 3.4968, 0.0021, 0.0001, 0.0026, 0.0001)$.

Consumption cost of a default

Given the estimated repricing of Mexico's sovereign risk, by how much would consumption fall if a sovereign default materializes as a consequence of this geoeconomic pressure? Answering this question is hard as Mexico has not defaulted in the last 40 years. In fact, the last time Mexico defaulted on its foreign debt was back in 1982, and it came quite close during the 1994 "Tequila Crisis." However, even though it is a *rare* event, I can use the implied risk-neutral default probability of Mexico to give a possible answer to this question. In particular, I would need to re-express the dynamic treatment effect $\hat{\alpha}_t$ in terms of the difference in implied probabilities of default for Mexico and the counterfactual instead of the difference in CDS spreads. I read this exercise primarily as an *external validation* of the filtering approach, since the implied consumption drop is built solely from the smoothed series and

FIGURE F1. Observed and smoothed consumption index

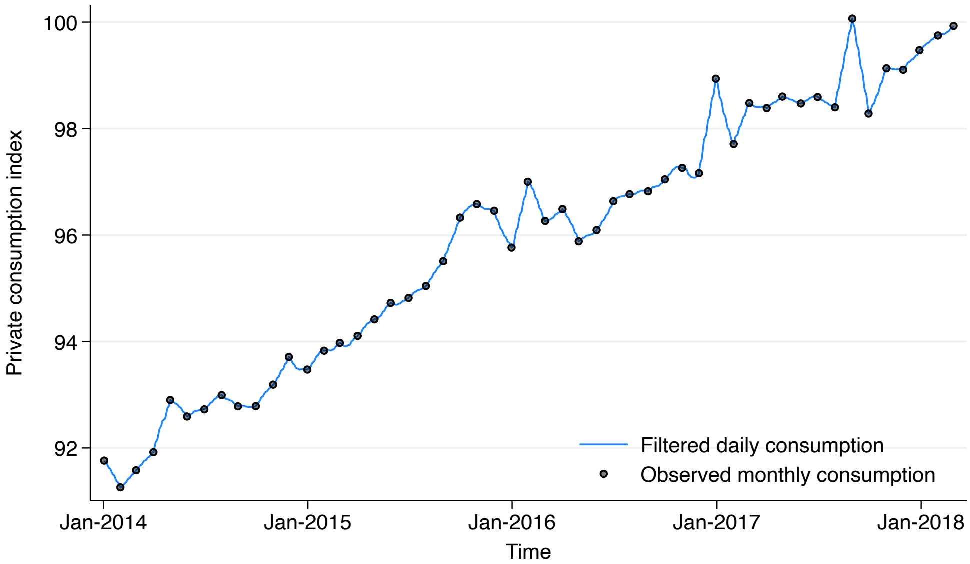

Note: The black markers represent the observed monthly observations of the private consumption index reported by *Instituto Nacional de Estadística, Geografía, e Información* (INEGI). The blue line corresponds to the smoothed series using the Kalman filter with daily stock returns and monthly consumption growth.

the synthetic-control premium, none of which uses the historical default episodes it is later compared against.

To that end, consider the following monotonic transformation of $\hat{\alpha}_t$:

$$\begin{aligned}\widetilde{\alpha}_t = \mathbb{P}_t^{\mathbb{Q}}\left[\text{Mexico defaults} \leq 5 \text{ years}\right] - \mathbb{P}_t^{\mathbb{Q}}\left[\text{Counterfactual defaults} \leq 5 \text{ years}\right], \\ = \left(1 - \exp\left(-5 \times \frac{CDS_t}{1-R}\right)\right) - \left(1 - \exp\left(-5 \times \frac{\widetilde{CDS}_t}{1-R}\right)\right), \\ = \exp\left(-5 \times \frac{\widetilde{CDS}_t}{1-R}\right) - \exp\left(-5 \times \frac{CDS_t}{1-R}\right), \end{aligned} \tag{F.3}$$

where $\mathbb{P}^{\mathbb{Q}}[\cdot]$ represents the cumulative five-year risk-neutral default probability, CDS_t and $\widetilde{CDS}_t$ correspond to the observed and counterfactual 5-year CDS for Mexico, respectively, and R is the recovery rate that I will set at 39.5% following Benjamin and Wright (2009). This transformation is particularly useful as $\widetilde{\alpha}_t \to 1$ can be interpreted as Mexico defaulting on its debt within the next 5 years relative to the counterfactual.[42]

By estimating a linear projection analogous to (29), extended to $h = 60$ and using this transformation of the treatment effect instead of $\hat{\alpha}_t$, I obtain an estimate of how much the daily smoothed consumption measure would change at $t + h$ if Mexico were to default at t. For comparability with the structural literature of strategic default, I use the estimated response at $h = 60$ (i.e., $\hat{\theta}^{60}$) as the corresponding quarterly response (a typical business-cycle frequency). This necessarily extrapolates a locally estimated coefficient to the boundary case $\widetilde{\alpha}_t \to 1$, so the resulting figure should be read as suggestive rather than as a precise point estimate. According to the estimates, consumption falls by 12.45% a quarter after default.

[42] This occurs when $\widetilde{CDS}_t \to 0$ and $CDS \to +\infty$.

TABLE F1: Fall in consumption after default or debt crisis

Estimated	Argentina (2002Q1)	Ecuador (1993Q3)	Mexico (1995Q1)	Russia (1999Q4)
-12.45%	-16.01%	-7.14%	-12.14%	-17.20%

Note: The numbers for Argentina, Ecuador, and Russia come from tables 1 and 2 of Arellano (2008), respectively. They represent the percent deviation of seasonally adjusted consumption from a linear trend. For the case of Mexico, I compute the percent deviation from the log-quarterly observations of the Private Consumption Index from a linear time trend.

In Table F1, I display the observed contraction in consumption documented in Arellano (2008) for Argentina, Ecuador, and Russia after a sovereign default episode. Also, I display the contraction in consumption for Mexico during the sovereign debt crisis of 1994–1995, even though this episode did not end up in a sovereign default.[43] The estimated contraction is roughly comparable to the contraction in quarterly consumption following the default episodes of Argentina, Ecuador, and Russia. Moreover, it is remarkably close to the documented contraction in consumption during the Tequila Crisis of 1994–1995.

Because the estimated contraction is constructed entirely from the smoothed consumption series and the synthetic-control premium, objects that never use these historical episodes, its proximity to the observed contractions during past default and near-default episodes provides an out-of-sample validation of the filtering approach.

[43] An ideal point of reference for Mexico would be the drop in consumption after the 1982 default; however, to my knowledge there is no quarterly consumption data available for that period. Using yearly data on real consumption (e.g., `ccon`) from Feenstra et al. (2015), I get a contraction -4.76% from 1982 to 1983.