Sovereign Borrowing from Private External Creditors during Default *

Hasan Cetin†

August 2026

Abstract

Emerging-market sovereigns are not fully excluded from private credit during default, contrary to what the literature assumes. Using World Bank data on 66 default episodes, I show that private creditors keep lending, although about two-thirds less than in normal times. Most of it comes early, and the first half of an episode receives roughly 3.7 times as much as the second. This access can help the government smooth its expenditure through stressful times (an *insurance effect*). But it can also weaken the sovereign's commitment to repay and raise spreads in normal times (a *commitment effect*). To weigh the two, I develop a model with two sequentially accessed private markets, calibrated to Colombia. Removing the second market makes the sovereign better off by 0.28% in consumption-equivalent terms. The commitment effect dominates unless re-entry to the regular market is fast enough. The equilibrium can be read as a coordination failure among private lenders that hurts the borrower itself.

Keywords: sovereign default, private external creditors, borrowing during default, long-term debt, limited commitment, *pari passu*, emerging markets.
JEL codes: E44, F32, F34, F41, G15, H63.

*I thank Manuel Amador, Marco Bassetto, Tim Kehoe and Illenin Kondo for their guidance during this project. I also thank Lucas Belmudes, Jason Hall, Angelo Mendes, Zach Swaziek, and seminar participants at the University of Minnesota Trade Workshop.

†cetin019@umn.edu: University of Minnesota.

1 Introduction

Standard sovereign-debt models assume full exclusion from private markets while in default. That is, no new borrowing until re-entry, whether re-entry is modeled as a random event or as the close of the negotiations over the defaulted debt. However, this is not the case in reality. Figure 1 shows disbursements and net flows from private external creditors for Argentina and Brazil, with shaded regions marking default episodes. Even during default episodes, both governments were still able to borrow new funds from private external creditors. And this fact is not specific to these two countries.

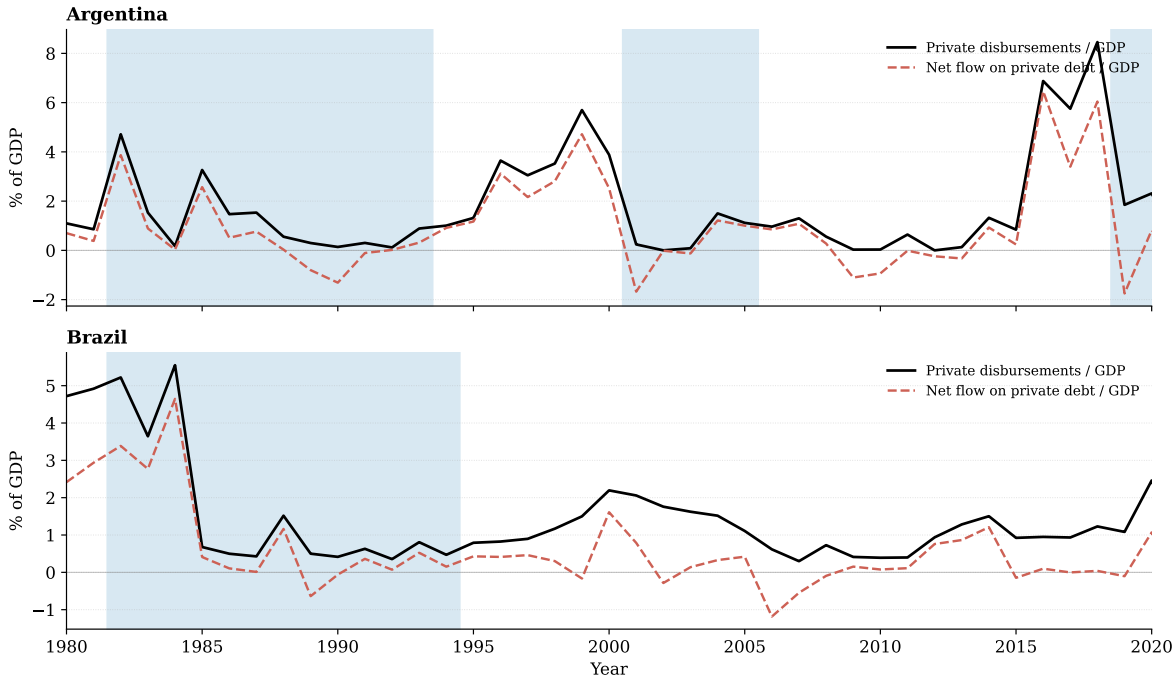


Figure 1: Private disbursements / GDP (%, solid) and net flows on private external debt / GDP (%, dashed), Argentina and Brazil, 1980–2020. Shaded regions mark Asonuma and Trebesch (2016) default episodes with missed payments (the AT-PDW classification, defined in Section 2). In every episode in the sample, both countries continue to receive private disbursements, concentrated in the early years of each episode. *Source:* World Bank International Debt Statistics, default dating from Asonuma and Trebesch (2016), and GDP from IMF World Economic Outlook.

Across 66 default episodes from the Asonuma and Trebesch (2016) sovereign-debt-restructuring database (henceforth AT), covering 49 emerging-market and developing economies over 1980–2020, sovereigns have real but limited access to new borrowing from private external creditors during default. It is about two-thirds below the non-default level once country and year fixed effects are absorbed. This limited access is also front-loaded: in

the first half of an episode the government receives roughly 3.7 times the disbursements of the second half, and its net flow on private debt is still positive.

On the one hand, access to new funds during a default episode might help governments smooth their expenditure through stressful times (an *insurance effect*). But on the other hand, having an outside option for funding might weaken their commitment to repay, resulting in higher spreads outside default episodes (a *commitment effect*). The main question this paper tries to answer is: does having access to new borrowing during default actually make the sovereign better off?

To answer this question, I develop a sovereign-default model with two sequentially accessed private markets. Market 1 is the regular market where the sovereign borrows in normal times. Once the sovereign defaults on Market 1, it can still borrow from Market 2 lenders. Market 2 lenders can be thought of as "loan sharks": they lend to the sovereign while it is in default. The sovereign has limited commitment to these new lenders as well. If it defaults again, it goes to autarky for a random period of time. Both default states carry an output cost. Income is reduced in Market 2, and reduced further in autarky. While in Market 2, a forgiveness shock arrives from time to time and allows the sovereign to settle its Market 2 debt at face value and re-enter Market 1. With long-term debt, what the sovereign does when this shock arrives depends on its state. It can settle in full, keep borrowing in Market 2 under the output penalty, or default again and go to autarky. In both markets, lenders are competitive and risk-neutral, and bond prices reflect the sovereign's default risk. Market 2 lenders also price the risk of yet another default, so borrowing terms in Market 2 get worse as Market 2 debt accumulates. As a result, the model produces the front-loaded borrowing as shown in the data.

The model is calibrated to Colombia at annual frequency. The main counterfactual removes Market 2. A Market 1 default then sends the sovereign straight to autarky, which brings back the standard one-market exclusion structure of the sovereign-default literature. In effect, it imposes full commitment across lenders: no one lends to the sovereign while it is in default. In the calibrated model, removing Market 2 makes the sovereign better off by 0.28% in consumption-equivalent terms at the ergodic center. So at the calibrated parameters, the commitment effect dominates the insurance effect. Which effect dominates is not obvious a priori, though. The gain is robust to the default cost in the second market. However, it is contingent on the market re-entry probability. If the re-entry to Market 1 is fast enough, the sovereign stays only a short period of time excluded from Market 1. Then Market 2 works as cheap insurance against bad times, and the insurance effect wins. The calibrated economy sits well inside the region where removing Market 2

helps.

These findings suggest that the coordination failure among private lenders is actually bad for the sovereign itself. While the sovereign is in default against one group of lenders, the other group still lends to it. Each lender prices its own loans fairly, but none internalizes what its lending does to the sovereign's borrowing terms in normal times. Contractual or institutional arrangements that help private lenders act as a single cohort might therefore serve the sovereign itself, not only the creditors such arrangements were designed to protect. The most-favored-creditor reading of the *pari passu* ("equal footing") clause is one existing example (Olivares-Caminal, 2013).[1]

Related literature. The quantitative framework follows the strategic-default tradition of Eaton and Gersovitz (1981) and Arellano (2008), with the random-maturity long-term-bond convention of Hatchondo and Martínez (2009). The difference in debt maturity between the two markets has a counterpart in the short-versus-long bond distinction of Arellano and Ramanarayanan (2012), and the output penalty in Market 2 plays the same role as the endogenous output cost of Mendoza and Yue (2012). What all of these models have in common is that a sovereign in default cannot borrow from any private lender.

Another strand of the literature models default as a renegotiation instead of an exclusion state. Yue (2010) has the sovereign and its lenders bargain over a haircut, and in Benjamin and Wright (2009) the bargaining itself is the source of delay. Dvorkin et al. (2021) (henceforth DSSY) calibrate a restructuring model with maturity choice to Colombia, and their income process and spread moments motivate my calibration choices in Section 4. Still, in all of these papers the default episode is a bargaining period over the legacy debt, and the sovereign does not borrow anything new until it settles and repays the negotiated recovery. Positive disbursements during the episode, which is what I find in Section 2, cannot happen in these models. To get them, a model needs a borrowing option that survives default, and in my model this is the second market.

Bocola and Dovis (2019) come closer to my two-market structure, with a sovereign that borrows from foreign and domestic creditor pools *simultaneously* and a focus on rollover risk. My sovereign accesses the two pools *sequentially* instead, and it can still borrow in Market 2 while it is in default on Market 1. Erce and Mallucci (2018) also study selective default across creditor types, but as a cross-sectional choice made within an episode. The coordination problem across lender cohorts is related, too, to the bargaining frictions

[1] In practice, *pari passu* has rarely bound in this form. The main exception is the *NML Capital v. Argentina* litigation of 2011–2014 (Schumacher et al., 2021).

caused by holdout creditors in Pitchford and Wright (2012), although there the friction arises within a single restructuring.

That the exclusion punishment is rarely enforced is a puzzle documented and rationalized by Tomz and Wright (2007) and Bulow and Rogoff (1989). Gelos et al. (2011) show that defaulters often regain market access within a few years, long before the formal restructuring is over. I find the same thing happening inside the default itself. Seeing it requires separating post-default restructurings from preemptive ones, and the AT database draws this line more sharply. The closest paper to mine is Arellano et al. (2023). They document that sovereigns partially repay their existing debt during default, and they model a continuous repayment intensity on that stock. Even though both papers analyze borrowing under default, the empirical findings are still different. Their repayment intensity is roughly symmetric over the episode, while the disbursements I find are front-loaded. To explain this front-loading empirical fact, I introduce sequential markets, which is different from the one-market case in Arellano et al. (2023).

Outline. The rest of the paper goes as follows. Section 2 presents the facts on private lending during default, Section 3 builds the model, and Section 4 calibrates it to Colombia. Section 5 runs the counterfactuals, and Section 6 concludes. The robustness checks and the solution algorithm are in the appendices.

2 Empirical Facts

This section documents the two empirical facts that the model is built to match. Sovereigns have real but limited access to external private credit during default (Fact 1). This flow is also front-loaded, and the first half of an episode receives roughly 3.7 times the disbursements of the second half (Fact 2). Along the way, I also find that disbursements stay low after the episodes end, that the private debt-to-GDP ratio goes up and then down over the episode, and that the mix of creditors changes during default.

Data. The private-creditor flows I use come from the World Bank's International Debt Statistics (IDS). When a sovereign signs a long-term contract with a private external creditor, IDS records its face value in the year of signing as a *new commitment to private creditors*. The cash actually drawn on these commitments during a year is recorded as *disbursements to private creditors*. Disbursements received minus the amortization paid to the external creditors is the *net flow on private external debt*. It is positive if the sovereign is a net receiver of funds, and negative if the sovereign is a net repayer. All three series cover public and publicly guaranteed (PPG) debt owed to external private creditors. External private creditors cover bondholders, commercial banks, and other private creditors (excluding bilateral and multilateral creditors).[2] I divide each series by nominal GDP, taken from the IMF World Economic Outlook. The IDS series codes and the details of the country-year panel are in Appendix B.

Default-episode dataset. I identify default episodes using the AT sovereign-debt-restructuring database (Asonuma and Trebesch, 2016), which classifies restructurings by whether and when the government missed a payment. A restructuring is *post-default* if a scheduled payment is missed before negotiations begin. It is *weakly preemptive* if the government first publicly announces that it wants to restructure and then misses a payment while negotiating, and *strictly preemptive* if no payment is missed at all. I keep episodes that start between 1980 and 2020. My baseline sample (AT-PDW) covers the post-default and weakly preemptive default episodes, which excludes the strictly preemptive episodes. In other words, the baseline includes the episodes where there are actual missed payments. Appendix A shows that the results are robust to other default-episode sample selections. The baseline AT-PDW sample has 66 episodes in 49 emerging-market and developing economies, and Appendix B lists all of them.

[2] Restricting attention to private creditors fits the model of Section 3, where lenders are numerous, risk-neutral, and competitive. Bilateral and multilateral official lending does not fit this description well, since political considerations and conditionality dominate it.

Disbursements versus new commitments. Since disbursements are drawings on commitments, the two series should be comoving. Table 1 compares both series. The mean ratio of disbursements to commitments is 1.02 in default years and 0.97 in non-default years, and the correlation is above 0.94. In the analysis that follows I use disbursements as the dependent variable. IDS reports them on 1,272 additional country-years where the commitment series is missing, and provides a finer creditor-type decomposition (bonds, commercial banks, and other private) that is not available for commitments. The same regressions on new commitments yield qualitatively the same results (Appendix C).

Table 1: Disbursements versus commitments: equivalence on the common sample, true-zero pattern on the gap sample. Means in real 2015 US$ millions per country-year.

	N	mean commit	mean DIS	DIS/commit or zero-share	corr.
Panel A. Common sample (commit and DIS both observed)					
default years	264	757	769	1.016	0.961
non-default years	1,838	2,678	2,589	0.967	0.945
Panel B. Gap sample (commit NaN, DIS observed)					
default years	119	n.a.	12.1	61% zero, 25% >$10m	n.a.
non-default years	1,153	n.a.	11.2	77% zero, 10% >$10m	n.a.

Notes: The sample is AT-PDW interior years, 1980–2020. Panel A uses the country-years where IDS reports both commitments and disbursements. Panel B uses the country-years where only disbursements are reported. Most of these have no private credit activity, and 76% of their disbursement values are zero. The fifth column is the mean ratio of disbursements to commitments in Panel A, and the share of country-years with zero disbursements (or above $10m) in Panel B. The last column is the correlation between the two series. *Source:* World Bank International Debt Statistics.

Episode selection and halves. While AT reports episodes by month and year, the IDS dataset is annual. Thus, I drop the beginning and end years of each default episode because they might mix borrowing activity inside and outside the default episode. The remaining episode years are split into a *first half* and a *second half*.[3] I consider default episodes that last at least 4 years.[4] There are 46 such AT-PDW episodes, in 43 countries, where disbursement and GDP data are available. Table 2 covers these episodes. The event study in Figure 2 also needs non-missing observations both before and after the default episode. Only 35 episodes satisfy this, and Figure 2 depicts the event study on these episodes.

Event study. For each episode, I average five outcomes over four bins of years. The two middle bins are the first and second halves defined above. The bins before and after the

[3]When the number of years is odd, the middle year goes to the first half.

[4]Since I remove the beginning and end years, 4 years is the minimum episode length that can have a proper first and second half.

episode match the widths of the two halves. The *Before* bin has as many years as the first half and ends just before the episode. The *After* bin has as many years as the second half and starts right after it. For example, an 8-year episode gives bin widths of $(3, 3, 3, 3)$, and a 7-year episode gives $(3, 3, 2, 2)$. The five outcomes are real new disbursements, real GDP, the ratio of disbursements to GDP, the net flow on private external debt (disbursements minus amortization) to GDP, and the private external debt stock to GDP.

Panel (a) of Figure 2 shows that real new private disbursements fall from 100 in the Before bin to about 24 in the first half and 10 in the second half. They recover only partly in the After bin, to about 12. So access is limited, but some of it remains. Panel (b) shows that real GDP falls only modestly over the episode, to 81 and 79 in the two halves, and it recovers to 106 in the After bin. Since income falls much less, the fall in disbursements reflects a fall in the credit flow itself. Panel (c) shows the same fall as a ratio to GDP. Disbursements are 1.43% of GDP in the first half and 0.67% in the second half, so the front-loading in the event-study sample is roughly two to one. In the broader within-episode sample, which I use for the Fact 2 regression below, the raw ratio of the first half to the second half is larger, at about 3.7 to 1.

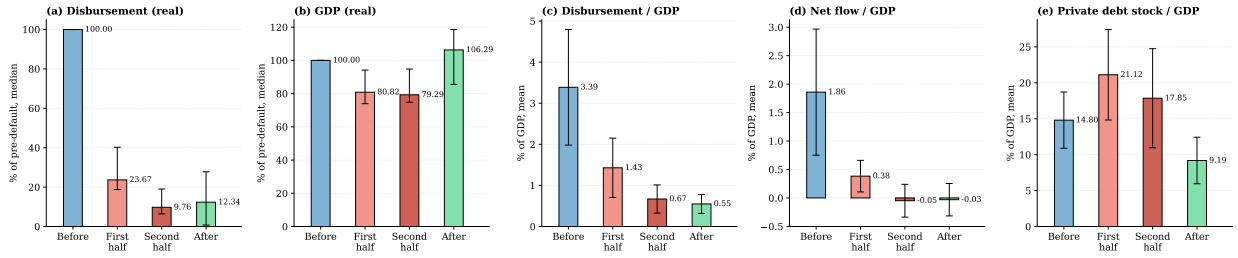


Figure 2: Event study around default onset, AT-PDW sample (35 balanced episodes). Panels (a) and (b): real private disbursements and real GDP, indexed to Before = 100 (cross-episode medians, bootstrap 95% CIs). Panel (c): disbursement / GDP (cross-episode means, asymptotic 95% CIs). Panel (d): net flow on private external debt / GDP, defined as disbursement minus amortization on PPG private external debt (cross-episode means, asymptotic 95% CIs, World Bank IDS series `DT.NFL.PRVT.CD`). Panel (e): private external debt stock / GDP (cross-episode means, asymptotic 95% CIs, World Bank IDS series `DT.DOD.PRVT.CD`). *Source:* World Bank International Debt Statistics, IMF World Economic Outlook, and default dating from Asonuma and Trebesch (2016).

Within-episode debt hump. Panel (e) of Figure 2 shows the path of the private debt stock to GDP across the four bins. It rises from 14.8% in the Before bin to a peak of 21.1% in the First half, falls to 17.9% in the Second half, then drops sharply to 9.2% in the After bin. This is similar to the partial-repayment hump that Arellano et al. (2023) document. Panel (d) shows where the hump comes from on the flow side. Net flows on private external

debt (disbursement minus amortization, divided by GDP) average 1.86% in the Before bin, 0.38% in the First half, and turn close to zero in the Second half (−0.05%) and After bin (−0.03%). Before the default episode and in the first half of the episode, sovereigns are net borrowers and the debt stock increases. In the second half of the episode the net flow is almost zero. At the same time, GDP falls with default, which makes private debt to GDP peak in the first half of the episode. After the episode, the sharp fall in the debt-to-GDP ratio comes mostly from the recovery of GDP (Panel b, 79 → 106) and from the debt restructuring at the end of the episode. It is not driven by net amortization within the episode. To a first approximation, the hump is debt built up before the default meeting a falling flow, and borrowing during default that is later repaid plays little role in it.

Fact 1 regression: real but limited access. I run a panel regression of private disbursements to GDP on a default dummy and a post-episode dummy, with country and year fixed effects:

$$\frac{\text{Disbursements}_{i,t}}{\text{GDP}_{i,t}} \times 100 \;=\; \beta\, D_{i,t} \;+\; \gamma\, P_{i,t} \;+\; \alpha_i \;+\; \tau_t \;+\; \varepsilon_{i,t}, \tag{1}$$

where $D_{i,t}$ is a default-year indicator equal to one in interior years of any AT-PDW episode in country i, $P_{i,t}$ is a post-episode indicator equal to one in the five years following each episode's end, α_i is a country fixed effect, τ_t is a year fixed effect, and $\varepsilon_{i,t}$ is the country-year residual.[5] The estimating sample is the panel of 92 emerging-market and developing economies that (i) appear in the IDS country list with at least one non-missing observation over 1980–2020 and (ii) have at least two million population. Of these countries, 49 have an AT-PDW episode in the window and 43 do not. Country fixed effects absorb level differences across countries, and year fixed effects absorb global shocks. The default dummy is non-zero only in the interior episode years of defaulting countries. So β measures how much lower private disbursements are during default, conditional on the fixed effects. Standard errors are clustered at the country level.

Table 2 reports the estimates. Columns 1 and 4 use total private disbursements, my main series. Columns 2 and 5 keep only bonds and commercial banks, the part of private external credit closest to the arm's-length lenders of the model in Section 3. Columns 3 and 6 use new commitments to private creditors, the contract side of the disbursements discussed above. The first three columns use the full 92-country panel, and the last three

[5] The coefficients β and γ in equations (1) and (2) are different objects from the discount factor β and the CRRA risk aversion γ of the model in Section 3. When the post-episode window overlaps with a later episode in the same country, I move the overlapping years from the post indicator to the default indicator, so the two indicators never overlap.

drop the HIPC Initiative completion-point countries.[6]

Table 2: Regression estimates on private-credit flows, AT-PDW sample, 1980–2020.

	All 92 countries			HIPC excluded		
	Private DIS	Bond+Bank DIS	Commitments	Private DIS	Bond+Bank DIS	Commitments
Default + post-episode panel regression (country + year FE)						
Default-year dummy	-0.710***	-0.500**	-1.334***	-0.852**	-0.642**	-1.653***
	(0.257)	(0.214)	(0.427)	(0.333)	(0.306)	(0.482)
Post-episode dummy, 5 yr	-0.779***	-0.537***	-1.394***	-0.975***	-0.767***	-1.293***
	(0.165)	(0.124)	(0.306)	(0.206)	(0.160)	(0.268)
Observations	3,351	3,351	2,087	2,086	2,086	1,538
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Within-episode regression (episode FE, log GDP control)						
First-half dummy	+0.961**	+0.526	+0.417**	+0.531***	+0.295**	+0.560**
	(0.487)	(0.326)	(0.185)	(0.132)	(0.119)	(0.231)
Observations	373	373	256	148	148	109
Episodes	46	46	41	22	22	21
Episode FE	Yes	Yes	Yes	Yes	Yes	Yes
log(GDP) control	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The sample is AT-PDW interior years, 1980–2020. All dependent variables are in percent of GDP. Private DIS is total disbursements from private creditors, Bond+Bank DIS keeps only bonds and commercial banks, and Commitments is new commitments from private creditors. The HIPC excluded columns drop the 33 HIPC completion-point countries. Standard errors are in parentheses, clustered by country in the top panel and by episode in the bottom panel. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* World Bank International Debt Statistics and IMF World Economic Outlook.

There is still private credit during default episodes. Private disbursements still average 1.00% of GDP in default years. However, new borrowing is limited in default years: the default-dummy coefficient is -0.71 percentage points of GDP. I find the same fall with all three dependent variables and in both samples. It is smaller but still significant for bonds and commercial banks alone (-0.50 pp in the full panel), and larger for new commitments (-1.33 pp, on the smaller commitment sample). Without the HIPC countries, the Private DIS coefficient is larger in magnitude, at -0.85 pp.

Country fixed effects take care of the time-invariant factors behind how much a country can borrow during default, but the sovereigns that keep any positive flow during default are still a selected group, because countries that lose all market access drop out of the regression after the within transformation. The size of the default-year coefficient also depends on where AT place the episode boundaries, since these dates decide which country-years count as default years.

[6] I drop all 33 HIPC completion-point countries in the panel together. For these countries, official rescheduling (Paris Club, IMF, World Bank) may matter more for private flows during default than competitive market forces. The list comes from the IMF and World Bank HIPC Initiative.

Post-episode dynamics. Conditional on the fixed effects, private disbursements stay low for at least five years after an episode ends. The post-episode coefficient γ in equation (1) is -0.78 pp in the full panel and -0.98 pp without HIPC countries, both significant at the 1% level and larger in magnitude than the default-year coefficients (Table 2). The strict AT-PD subsample gives near-identical values (Appendix A). The model of Section 3 is silent on this margin, and I return to it in Section 6.

Fact 2 regression: front-loading within the episode. Conditional on being inside a default episode, sovereigns receive substantially more private disbursements early in the episode than late. To formalize the asymmetry I estimate

$$\frac{\text{Disbursements}_{i,t}}{\text{GDP}_{i,t}} \times 100 \;=\; \beta\, \mathbf{1}\{t \in \text{first half of episode}_e\} \;+\; \gamma\, \log(\text{GDP}_{i,t}) \;+\; \mu_e \;+\; \varepsilon_{i,t}, \tag{2}$$

on the within-episode sample (interior years of length-$\geqslant$ 4 episodes), where μ_e is an episode (country $\times$ episode-start) fixed effect and $\varepsilon_{i,t}$ the country-year residual. The coefficient β compares the first and second halves of the same episode. I control for $\log(\text{GDP}_{i,t})$ because country size can change during an episode, and the episode fixed effect alone would not pick that up. Standard errors are clustered at the episode level.

Table 2 (third row) shows that the first-half coefficient for Private DIS is $+0.96$ pp of GDP, significant at the 5% level. That is, private disbursements are about 1 pp of GDP higher in the first half of an episode than in the second half of the same episode, once country size and episode fixed effects are controlled for. As with Fact 1, this is only an association. The raw numbers give a sense of the size. Disbursements average 1.53% of GDP in the first half and 0.41% in the second, with a pooled mean of 1.00%, a ratio of about 3.7 to 1. Without the HIPC countries the coefficient drops to $+0.53$ pp, but it is significant at the 1% level on this tighter sample, so front-loading looks more homogeneous among the non-HIPC episodes. I get the same result with new commitments ($+0.42$ pp on the smaller commitment sample, $p < 0.05$) and on the strict AT-PD subsample (Appendix A). With bonds and commercial banks alone, the coefficient is $+0.53$ pp in the full panel but not significant, and $+0.30$ pp ($p < 0.05$) once the HIPC countries are out. For the calibration I also need this contrast in the six peer countries, measured with commitments so that it matches the targets in Section 4. In that sample, the first half gets a share of 0.69 of the within-episode new commitments to GDP. This share is the Fact 2 target in the calibration, since it is the summary that can be compared directly with the model.

Creditor-type composition. I split total private disbursements by creditor type to see which creditors lend to a sovereign during default. For each creditor type, Table 3 gives the mean disbursement per country-year in percent of the country's own GDP, which puts countries on the same footing, and the dollar-weighted share of total disbursements, separately for default and non-default years. Outside default, bonds are the largest source of private external credit (58% of credit), while credit from commercial banks (27%) and other creditors (15%) is limited. In default years, bonds nearly vanish (7%), and the limited new credit in these years mainly comes from commercial banks (59%) and other private creditors (34%). Other private credit is a residual category that contains trade credit and bank loans guaranteed by export credit agencies, and the notes to Table 3 give the full IDS definition. Appendix E repeats the Fact 1 and Fact 2 regressions with net flows by creditor type, and Appendix F runs the Fact 1 regression on disbursements separately for each creditor type.

Table 3 shows that looking only at bond or commercial bank credit would leave most of the private lending in default years unexplained. Thus, the findings use sovereign borrowing from all kinds of external private creditors. The radical shift in borrowing composition from bondholders to commercial banks and other private creditors during default suggests that the competitive lenders described in Section 3 can be read as a pool of bondholders and commercial banks in non-default years, and of commercial banks and other private creditors in default years.

Table 3: Creditor-type composition of private external disbursements, default vs non-default years. Per-country-year mean disbursement as percent of own GDP (size-normalized) and dollar-weighted within-period share (aggregate composition).

	Default years		Non-default years	
Creditor type	mean DIS/GDP (%)	$-weight share	mean DIS/GDP (%)	$-weight share
Bonds	0.02	7%	0.45	58%
Commercial banks	0.43	59%	0.34	27%
Other private	0.55	34%	0.27	15%
Total private DIS	1.00	100%	1.06	100%

Notes: AT-PDW interior years, 1980–2020, on the 92-country regression sample. Each row is one creditor-type component of total private disbursements (`DT.DIS.PRVT.CD`). *Mean DIS/GDP* is the average per country-year within the year-status sub-sample, in percent of own GDP. *$-weight share* is the dollar-weighted aggregate share within the year-status sub-sample. The default-year sample has 246 country-years in 43 episodes, and the non-default sample has 3,105 country-years in 92 countries. *Other private*, per the World Bank IDS definition, is "other private credits from manufacturers, exporters, and other suppliers of goods, and bank credits covered by a guarantee of an export credit agency." *Source:* World Bank International Debt Statistics and IMF World Economic Outlook.

3 Model

This section presents a small open economy model of sovereign default. The sovereign borrows in two private bond markets that it accesses one after the other, and its debt is long-term and zero-coupon. A forgiveness option allows the sovereign to settle its second-market debt and return to the first market. The timing within a period and the competitive risk-neutral pricing follow Eaton and Gersovitz (1981) and Chatterjee and Eyigungor (2012).

The timing of lending of these two groups of creditors is sequential. The lenders in the first market are the usual bondholders, who lend to the government while the government has good credit standing. Lenders in the second market are a separate and competitive pool that the sovereign can borrow from only after defaulting on the lenders in the first market. The government also has limited commitment to these second lenders. This sequential access to two separate sets of lenders lets the model produce the two empirical facts shown in Section 2: (i) limited new borrowing in the default episode compared to non-default years, and (ii) front-loaded borrowing within the default episode.

Environment. There is a sovereign with discount factor $\beta \in (0,1)$ and a period utility function $u(c)$ that is continuous, strictly increasing, and strictly concave. There are two groups of competitive lenders, first-market lenders (L_1) and second-market lenders (L_2). The two groups are *distinct sets of agents*, but otherwise they are identical. They are risk-neutral and competitive, discount at the constant world risk-free rate $r^f > 0$, and price bonds at zero expected profit.

Total per-period income is $x_t = y_t + m_t$, where y_t is a persistent log-AR(1) endowment shock,

$$\log y_{t+1} = \rho \log y_t + \varepsilon_{t+1}, \qquad \varepsilon \sim N(0, \sigma_y^2),$$

and m_t is a small i.i.d. transitory shock, independent of y, with a continuous CDF $G(m)$ on $[-\bar{m}, \bar{m}]$. The transitory shock is only there to help the numerical solution (Appendix G).

The sovereign maximizes

$$\sum_{t=0}^{\infty} \beta^t u(c_t).$$

Sequential market access. The sovereign begins in *Market 1* (M_1) with a debt stock b_1 owed to L_1. If it defaults on L_1, it loses access to M_1 and moves to *Market 2* (M_2), where it

faces an output penalty

$$y_2^{\text{def}}(y) \leqslant y,$$

and it starts with a new debt account $b_2 = 0$ owed to L_2. Its M_1 stock is extinguished, with no recovery across markets in either direction. A second default, this time on L_2, sends the sovereign to *autarky*, where the output penalty is strictly more severe,

$$y_{\text{aut}}^{\text{def}}(y) \leqslant y_2^{\text{def}}(y) \leqslant y,$$

and the M_2 stock is extinguished.

In each period that the sovereign spends in M_2 or in autarky, a forgiveness shock arrives with probability $\theta \in (0,1)$. From autarky, the shock returns the sovereign to Market 1 with $b_1 = 0$ in the period in which it arrives, while in Market 2 it triggers the three-way choice described below. The re-entry probabilities from M_2 and from autarky could in principle differ, but the AT data identify them separately only weakly, so I use the same θ for both.

Long-term zero-coupon bonds. Debt in each market $i \in \{1,2\}$ is long-term and zero-coupon, with random maturity. Each period, every unit of outstanding b_i matures with probability $\lambda_i \in (0,1]$ and then pays one unit of consumption. If $\lambda_i = 1$, then the bond is a one-period bond. Let b_i be the existing stock of debt at the start of the period, and b_i' the new stock of debt chosen for the next period. The sovereign repays $\lambda_i b_i$ of its existing debt and borrows $b_i' - (1-\lambda_i)\, b_i$ in new debt at the price $q_i(y, b_i')$. When the sovereign repays in market i, its budget constraint is

$$c \;=\; y_{\text{state}} \;-\; \lambda_i\, b_i \;+\; q_i(y, b_i')\left(b_i' - (1-\lambda_i)\, b_i\right), \qquad c \geqslant 0, \tag{3}$$

where y_{state} is y in Market 1 and $y_2^{\text{def}}(y)$ in Market 2. Because debt is long-term, the price q_i also depends on how much the sovereign is expected to borrow next period. More expected issuance means a capital loss for the existing bondholders. Competitive lenders price this standard debt-dilution force in up front.

State variables and within-period timing. The state at the start of period t has two components: the persistent endowment shock y, and the regime-and-stock pair, which is one of:

- *In Market 1* with debt stock b_1 owed to L_1.

- *In Market 2 with the forgiveness shock not yet drawn this period*, with debt stock b_2 owed to L_2. I refer to this regime as *NF* ("no forgiveness this period").
- *In Market 2 with the forgiveness shock drawn this period*, debt stock b_2 owed to L_2. I refer to this regime as *F* ("forgiveness offered").
- *Autarky*, with no debt.

The timing within a period follows Eaton and Gersovitz (1981). Shocks are realized at the start of the period, and then the sovereign chooses its action. If it repays, it *first* pays the matured fraction of its existing debt and *then* issues new debt at the equilibrium price. Consumption is the residual. Finally, the regime and the stock for $t+1$ follow from the action and from next period's forgiveness draws. Figure 3 summarizes the within-period decision tree.

The action choice depends on the regime:

- *Market 1.* Repay or default. Defaulting consumes the M_2-penalty endowment $y_2^{\text{def}}(y)$ this period (no same-period M_2 borrowing) and moves the sovereign to NF next period with $b_2 = 0$.
- *Market 2 / NF.* Repay (under the M_2 output penalty $y_2^{\text{def}}(y)$) or default, sending the sovereign to autarky.
- *Market 2 / F.* A *three-way choice* $\chi^F \in \{\text{accept}, \text{reject-repay}, \text{default}\}$. *(i) Accept F*: settle b_2 to L_2 in full, re-enter M_1 next period with b_2 extinguished and a freshly chosen b_1'. *(ii) Reject F (stay in M_2)*: decline this period's F draw, repay $\lambda_2 b_2$, choose b_2', bear the M_2 penalty, and face an independent F draw next period. *(iii) Default*: repudiate the M_2 debt and enter autarky next period.
- *Autarky.* Consume penalized output. With probability θ, return to M_1 with $b_1 = 0$ next period.

The F-regime trade-off. With one-period bonds ($\lambda_2 = 1$), accepting F is mechanically equivalent to repaying in M_2 at full face value: both require the sovereign to settle the entire current b_2 this period. With long-term bonds ($\lambda_2 < 1$), accepting F costs an extra lump sum of $(1-\lambda_2)b_2$ compared with rejecting it and staying in M_2. If returning is chosen, the sovereign escapes the M_2 output penalty from then on. The choice is not trivial, and it is state dependent: if b_2 is high, the repayment cost may dominate the benefit of getting rid of the output penalty, and the sovereign may choose to reject F.

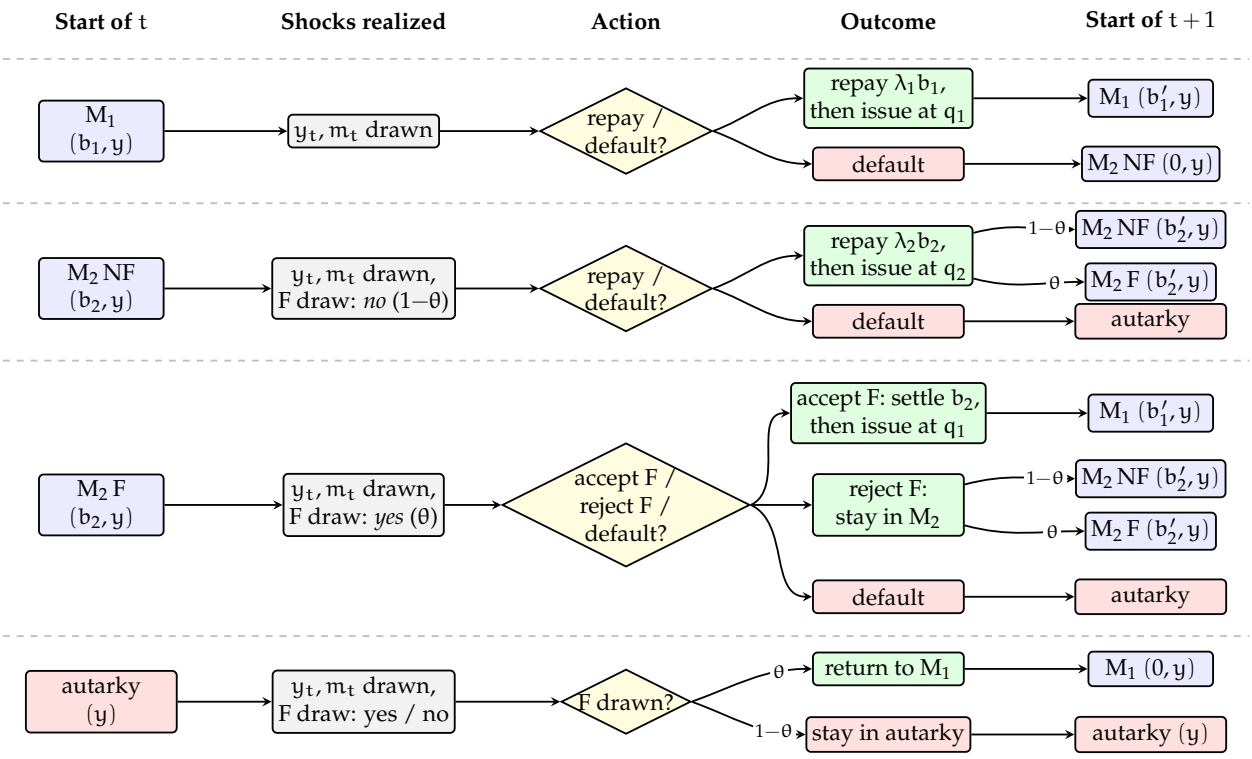


Figure 3: Within-period timing.

Value functions. Value functions take (y, m, b) as arguments. Because m is i.i.d., the equilibrium bond prices do not depend on it.

Market 1. Define the unconditional value W_1 as the maximum of repaying and defaulting:

$$W_1(y, m, b_1) \;=\; \max\Big\{\, V_1(y, m, b_1),\;\; V_d^{M_1}(y) \,\Big\}. \tag{4}$$

The repayment value V_1 solves

$$V_1(y, m, b_1) \;=\; \max_{b_1'} \left\{\, u(c) + \beta\, Z_1(y, b_1') \,\right\} \tag{5}$$

subject to

$$c \;=\; y + m - \lambda_1\, b_1 + q_1(y, b_1')\,(b_1' - (1-\lambda_1)\, b_1), \qquad c \geqslant 0, \tag{6}$$

and $Z_1(y, b_1') = \mathbb{E}_{(y', m')|y}[W_1(y', m', b_1')]$. In the period of an M_1 default, the sovereign does not borrow and consumes its M_2-penalty endowment, with the transitory shock set at its lowest value $-\bar{m}$, so the default value does not depend on the realized m. It then

starts the next period in NF with $b_2 = 0$:

$$V_d^{M_1}(y) \;=\; u\big(y_2^{\text{def}}(y) - \bar{m}\big) \;+\; \beta\, \mathbb{E}_{(y',m')|y}\big[\, W_2^{\text{NF}}(y', m', 0)\,\big]. \tag{7}$$

Market 2 / NF. Symmetrically,

$$W_2^{\text{NF}}(y, m, b_2) \;=\; \max\Big\{\, V_2(y, m, b_2),\;\; V_d^{M_2}(y)\,\Big\}, \tag{8}$$

where the M_2-default value $V_d^{M_2}(y)$ works the same way, with the autarky-penalty endowment this period (again at $m = -\bar{m}$) and autarky with no debt from the next period on,

$$V_d^{M_2}(y) \;=\; u\big(y_{\text{aut}}^{\text{def}}(y) - \bar{m}\big) \;+\; \beta\, Z^{\text{Aut}}(y), \tag{9}$$

with $Z^{\text{Aut}}(y)$ defined in equation (17) below. The repayment value V_2 solves

$$V_2(y, m, b_2) \;=\; \max_{b_2'} \,\big\{\, u(c) + \beta\, Z_2(y, b_2')\,\big\} \tag{10}$$

subject to

$$c \;=\; y_2^{\text{def}}(y) + m - \lambda_2\, b_2 + q_2(y, b_2')\,(b_2' - (1-\lambda_2)\, b_2), \qquad c \geqslant 0. \tag{11}$$

The M_2 continuation expectation is

$$Z_2(y, b_2') \;=\; \mathbb{E}_{(y',m')|y}\Big[\, \theta\, W_2^{\text{F}}(y', m', b_2') \;+\; (1-\theta)\, W_2^{\text{NF}}(y', m', b_2')\,\Big], \tag{12}$$

where the F shock is drawn next period with probability θ.

Market 2 / F. The F-regime envelope is

$$W_2^{\text{F}}(y, m, b_2) \;=\; \max\Big\{\, V_1^{\text{F}}(y, m, b_2),\;\; W_2^{\text{NF}}(y, m, b_2)\,\Big\}. \tag{13}$$

The second argument is the value of rejecting F, which already includes the option to default to autarky through (8), so the three-way choice reduces to this two-way max. The new piece is V_1^{F}, the value of accepting F:

$$V_1^{\text{F}}(y, m, b_2) \;=\; \max_{b_1'} \,\big\{\, u(c) + \beta\, Z_1(y, b_1')\,\big\} \tag{14}$$

subject to

$$c \;=\; y + m - b_2 + q_1(y, b_1')\, b_1', \qquad c \geqslant 0. \tag{15}$$

A sovereign that accepts F pays off all of b_2 this period, as a lump sum, and is back in M_1 right away. So it already consumes the unpenalized endowment y instead of $y_2^{\text{def}}(y)$, and it issues new M_1 debt b_1' at the M_1 price $q_1(y, b_1')$. From the next period on it is in M_1 with no M_2 debt, which is why the continuation value is Z_1.

Autarky. Consumption equals penalized output plus the transitory shock. With probability θ the sovereign returns to M_1 with zero debt:

$$W^{\text{Aut}}(y, m) \;=\; u\big(y_{\text{aut}}^{\text{def}}(y) + m\big) \;+\; \beta\, Z^{\text{Aut}}(y), \tag{16}$$

$$Z^{\text{Aut}}(y) \;=\; \mathbb{E}_{(y', m')|y}\Big[\, \theta\, W_1(y', m', 0) \;+\; (1-\theta)\, W^{\text{Aut}}(y', m') \,\Big]. \tag{17}$$

Bond pricing. Both creditor groups are competitive and risk-neutral. The price that each group charges depends only on the persistent income state y and on next period's stock in its own market, b_i'. This depends on two simplifications in the model. Since m is i.i.d., it does not enter the distribution of next period's state, and so it does not show up in either price. Also, L_1 and L_2 have separate balance sheets, and the markets are accessed in sequence. When L_1 prices an issuance of b_1', the only part of the sovereign's future behavior that matters for L_1's expected payoff is its behavior in M_1. By (5), this behavior depends on (y, b_1') and not on b_2. A sovereign in M_1 has no current M_2 liability, and any later default on M_1 extinguishes the whole M_1 stock. In the same way, when L_2 prices an issuance of b_2', only the M_2 continuation matters, and (10) depends on (y, b_2') and not on b_1. The sovereign in M_2 has no current M_1 liability, having extinguished it on M_1 default. Bonds are priced at zero expected profit.

Market 1 prices. Let $d_1(y', m', b_1') \in \{0, 1\}$ be the M_1 default decision and $a_1(y', m', b_1')$ the M_1 borrowing policy when not defaulting next period. Here a_1 is the optimal b_1' chosen at state (y', m', b_1'), i.e., the argmax of equation (5). A unit of M_1 debt that survives next period pays λ_1 today (matured fraction) plus a continuation value $q_1(y', a_1)$ on the unmatured fraction. So

$$q_1(y, b_1') \;=\; \frac{1}{1 + r^f}\, \mathbb{E}_{(y', m')|y}\Big[\, (1 - d_1(y', m', b_1'))\, \big(\lambda_1 + (1 - \lambda_1)\, q_1(y', a_1(y', m', b_1'))\big) \,\Big]. \tag{18}$$

The expectation is over both shocks, y' and m'. The dependence of the integrand on a_1 is the debt-dilution force.

Market 2 prices. The M_2 price is more involved because the per-unit payoff to L_2 depends on what the sovereign does next period. Analogously to a_1, let $a_2^{NF}(y', m', b_2')$ denote the M_2 borrowing policy under the no-forgiveness regime, i.e., the optimal b_2' chosen at state (y', m', b_2') as the argmax of equation (10).

Conditional on (y', m', b_2'), next period either the F shock arrives (probability θ) or it does not (probability $1-\theta$). If the F shock happens, the sovereign can accept F and go back to M_1, or reject it and stay in M_2, or reject it and also default on M_2 and go to autarky. If the F shock does not happen, the sovereign repays or defaults. The table gives the per-unit payoff to L_2 in each case.

F shock	**Sovereign action**	**Payoff per unit to L_2**
arrives (θ)	Accept F (settle b_2 in full)	1
arrives (θ)	Reject F (repay in M_2)	$\lambda_2 + (1-\lambda_2)\, q_2(y',\, a_2^{NF})$
arrives (θ)	Default ($M_2 \to$ autarky)	0
does not ($1-\theta$)	Repay in NF	$\lambda_2 + (1-\lambda_2)\, q_2(y',\, a_2^{NF})$
does not ($1-\theta$)	Default in NF (to autarky)	0

The accept-F branch pays 1 because the entire stock is settled in cash this period.

The price equates the expected discounted payoff to the unit cost:

$$q_2(y, b_2') \;=\; \frac{1}{1+r^f}\, \mathbb{E}_{(y',m')|y}\Big[\, \theta \cdot \Big(\mathbf{1}_{\text{accept F}} \cdot 1 \;+\; \mathbf{1}_{\text{reject F,repay}} \cdot \big(\lambda_2 + (1-\lambda_2) q_2(y', a_2^{NF})\big)\Big) \\ + (1-\theta) \cdot \big(1 - d_2^{NF}(y', m', b_2')\big)\, \big(\lambda_2 + (1-\lambda_2) q_2(y', a_2^{NF})\big)\, \Big], \tag{19}$$

where $\mathbf{1}_{\text{accept F}}$ and $\mathbf{1}_{\text{reject F,repay}}$ are indicators of the corresponding sub-events on the F-arrives branch (the default-from-F branch pays 0 and so contributes nothing), and d_2^{NF} is the NF default indicator on the no-F-arrives branch. Compared with a counterfactual without F, the accept-F payoff of 1 raises q_2. Full settlement under F is better for L_2 than the ongoing long-term stream, which is worth $\lambda_2 + (1-\lambda_2)\, q_2 < 1$ per unit. It is also much better than the zero that L_2 recovers in autarky.

Recursive equilibrium. A recursive equilibrium is a set of value functions $\{V_1, V_2, V_1^F, W_1, W_2^{NF}, W_2^F,$ policy functions $\{a_1, a_2^{NF}, a_1^F, d_1, d_2^{NF}, \chi^F\}$, and price functions $\{q_1, q_2\}$ such that:

1. Given the prices q_1 and q_2, the value functions solve the sovereign's problems in equations (4) to (17), and the policies are the sovereign's optimal choices in these problems. The borrowing policies a_1, a_2^{NF}, and a_1^F solve (5), (10), and (14). The default decisions d_1 and d_2^{NF} come from (4) and (8). The F-regime decision χ^F is the choice between accepting F, rejecting F, and defaulting in (13).

2. Given the sovereign's policies, the prices q_1 and q_2 satisfy the zero-profit conditions (18) and (19).

3. The budget constraints (3) and (15) hold in every state.

I solve for the equilibrium numerically by value function iteration. Appendix G gives the details.

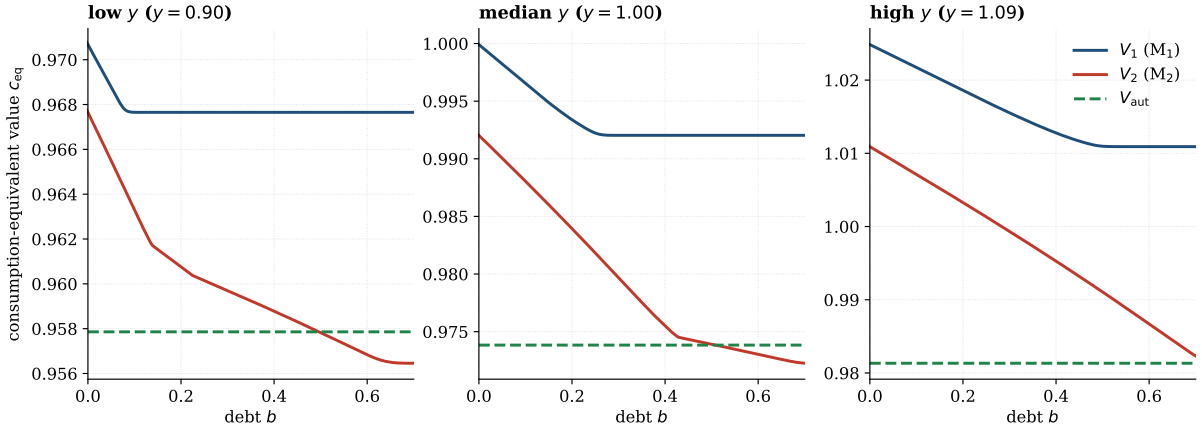


Figure 4: Value functions at the Colombia calibration of Section 4, plotted across three income states (y at the 10th, 50th, and 90th percentiles of the income grid) and over the full debt grid in units of mean endowment. Blue line: $V_1(y, b_1)$ in Market 1. Red line: $V_2(y, b_2)$ in Market 2. Green dashed line: $V_{aut}(y)$ in autarky (flat in debt). V_1, V_2, and V_{aut} stand for W_1, W_2^{NF}, and W^{Aut} in (4), (8), and (16). Values are computed in consumption-equivalent units. All curves are evaluated at $m = -\bar{m}$. In the M_1-default region, V_1 lies weakly below V_2 at $b = 0$, and the two cannot be told apart in the figure.

Figure 4 plots the solved value functions at the calibrated parameters of Section 4. Outside the M_1-default region, $V_1 \geqslant V_2 \geqslant V_{aut}$ at every (y, b), since the option to repay M_1 is worth at least as much as being in M_2 already, and being in M_2 is worth at least as much as autarky. Inside that region, V_1 lies weakly below V_2 at $b = 0$, because the sovereign does not borrow in the period of an M_1 default (equation (7)). At $b = 0$, the gap $V_1 - V_2$ measures the option value of M_1 over a new M_2 account. This gap is the M_1 default

cost that the model generates endogenously, and it shapes the M_1 price schedule. At low income, the sovereign defaults on M_1 whatever the realized b, so V_1 is flat at the M_1 default value $V_d^{M_1}(y)$.

The gap $V_2 - V_{aut}$ measures the option value of M_2 over autarky. This gap is small at low income, where autarky costs only a little more output than M_2 under the calibrated penalties of equation (20), and it grows as income rises. So the flat M_1-default region of V_1 gets smaller at higher income. At higher y, the sovereign defaults on M_1 only at some debt levels, instead of at all of them.

4 Calibration and Quantitative Fit

The model in Section 3 has 12 parameters. Seven are fixed externally from standard literature values or Colombia data, and the remaining five are chosen jointly by simulated method of moments to match the six targets in Table 4. The model is run at annual frequency, with Colombia as the target country. I use Colombia because it is a major emerging-market borrower that earlier papers have calibrated models to, Dvorkin et al. (2021) among them. My parameters and untargeted moments can then be compared with theirs directly. Appendix H redoes the welfare exercise for Mexico.

Table 4: Externally fixed parameters, internally estimated parameters, and targeted moments in the Colombia calibration of the two-market sovereign-default model.

Parameter / Moment	Target / Value	Model
Externally fixed parameters		
Relative risk aversion, γ	2.00	
Risk-free rate, r^f	3.32%	
Income persistence, ρ	0.87	
Income innovation s.d., σ_y	0.020	
M_1 maturity hazard, λ_1	0.084	
M_2 maturity hazard, λ_2	0.107	
Transitory income shock s.d., σ_m	0.005	
Internally estimated parameters (SMM)		
Discount factor, β	0.9511	
Re-entry probability, θ	0.0933	
M_2 output-penalty cap, d_2	0.9481	
M_2 output-penalty elasticity, d_2^c	0.6411	
Autarky multiplier, d_{aut}	0.9810	
Targeted moments (SMM loss)		
Mean M_1 debt to GDP	9.92%	10.41%
Mean episode length (years)	10.4	10.7
Commit-asymmetry share	0.690	0.732
In-default to non-default new-commitment ratio	0.590	0.606
Mean spread on Market-1 borrowing	2.61%	1.06%
Default frequency (per country-year)	2.60%	0.93%

Notes: SMM weights: $w_k = 1$ on the first four targeted moments, $w_k = 0.001$ on the mean spread and the default frequency. *Source:* World Bank International Debt Statistics, IMF World Economic Outlook, Bloomberg, and FRED.

Output penalties. The output-penalty parameters (d_2, d_2^c, d_{aut}) enter the model through the kinked-cap form

$$y_2^{def}(y) = \min\{ y,\ \ d_2\,\bar{y} + d_2^c\,(y - d_2\,\bar{y})\ \}, \qquad y_{aut}^{def}(y) = d_{aut} \cdot \min\{ y,\ \ d_2\,\bar{y}\ \}. \tag{20}$$

Here $\bar{y}$ is the long-run mean of income, so the kink sits at d_2 times mean income. Figure 5 plots the three regime-conditional output curves $\{y, y_2^{def}(y), y_{aut}^{def}(y)\}$ at the calibrated d_2, d_2^c, d_{aut}. In each market, the penalty is the gap between y and the corresponding curve. In Market 2, it is zero at low y, below the kink at $y = d_2\,\bar{y}$, and it is strictly positive and increasing in y above the kink. In autarky, the penalty is positive at every y, because $d_{aut} < 1$, and it also increases with y. Because the absolute output loss increases with income, default is countercyclical and the model is well behaved, as with the kinked cap of Arellano (2008) and the asymmetric quadratic form of Chatterjee and Eyigungor (2012).

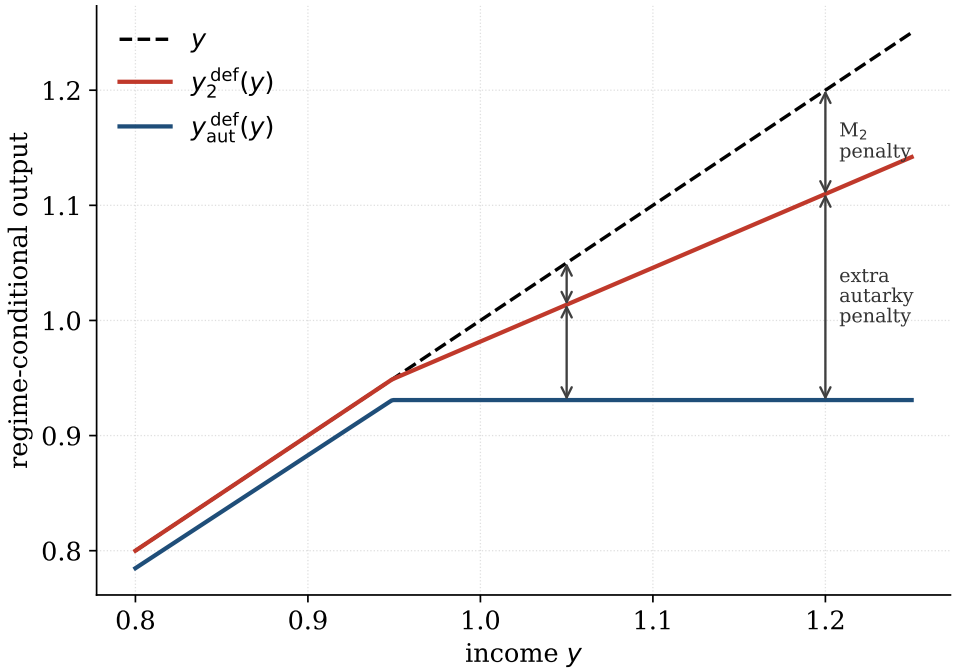


Figure 5: Regime-conditional output at the Colombia calibration. Horizontal axis: persistent income state y, in units of mean endowment. Vertical axis: post-penalty output, same units. Dashed: y (the unpenalized 45-degree line). Red: M_2 regime $y_2^{def}(y) = \min\{y,\ d_2\,\bar{y} + d_2^c\,(y - d_2\,\bar{y})\}$. Blue: autarky regime $y_{aut}^{def}(y) = d_{aut}\ \min\{y,\ d_2\,\bar{y}\}$. The gap between the 45-degree line and the red curve is the M_2 output penalty from defaulting in Market 1. The gap between the red and blue curves is the additional autarky penalty from defaulting in Market 2. The arrows at $y = 1.05$ and $y = 1.20$ mark both gaps. Both penalties grow with income, which makes default countercyclical.

Externally calibrated parameters. The parameters at the top of Table 4 come from the standard sovereign-default literature or directly from Colombia data. I set the CRRA curvature to $\gamma = 2$. For the income process $\log y_{t+1} = \rho \log y_t + \varepsilon_{t+1}$, with $\varepsilon \sim N(0, \sigma_y^2)$, an AR(1) on linearly detrended log real GDP for Colombia (1980–2019, leaving out 2020 because of the COVID disturbance) gives $\rho = 0.87$ and $\sigma_y = 0.020$. The transitory shock's standard deviation, $\sigma_m = 0.005$, does not come from the data. As in Chatterjee and Eyigungor (2012), it only keeps the numerical solution smooth (Appendix G). For the risk-free rate I take the annualized mean of the 10-year US Treasury yield over 2002–2017, which gives $r^f = 3.32\%$. Colombia's mean USD bond spread is measured over the same years, so the spread and the risk-free rate line up in both window and maturity. The Market 1 maturity hazard $\lambda_1 = 0.084$ is the inverse of the simple mean maturity of new private external commitments to Colombia in IDS, which is 11.9 years.

Peer-set construction and λ_2. Colombia has not defaulted in the sample window. I pick the defaulters in the AT dataset that have similar borrowing patterns to Colombia in the years outside default episodes. A country is in if its mean private external debt to GDP in non-default years is within ± 5 percentage points of Colombia's and its mean maturity of new private commitments in non-default years is within ± 3 years of Colombia's. Six AT-PDW defaulters pass the filter: Argentina, Brazil, the Dominican Republic, Morocco, Peru, and the Philippines (Table 5). There are six countries within the boundaries around Colombia, mostly Latin American. Figure 6 visualizes them.

Table 5: Colombia and its six peers.

	Debt/GDP, non-def (%)	Maturity, non-def (yr)	Maturity, in default (yr)
Colombia	9.92	11.94	–
Morocco	10.63	13.19	10.09
Peru	8.57	13.18	8.30
Philippines	10.31	12.76	11.32
Dominican Republic	8.91	11.73	7.96
Argentina	13.11	11.54	9.46
Brazil	8.40	10.75	8.99
Peer mean	9.99	12.19	9.35

Notes: 1980–2020, AT-PDW labeling. Peers are chosen so that their mean private external debt to GDP is within ± 5 pp, and their mean maturity of new private commitments within ± 3 years, of Colombia's in the non-default years. In the last column, each country's mean maturity in default years is the empirical analog of $1/\lambda_2$. Colombia has no AT-PDW episode in the window, so it has no in-default maturity ("–").
Source: World Bank International Debt Statistics and IMF World Economic Outlook.

The seven length-$\geqslant 4$ default episodes from this peer subset (Argentina 1982–1993, Ar-

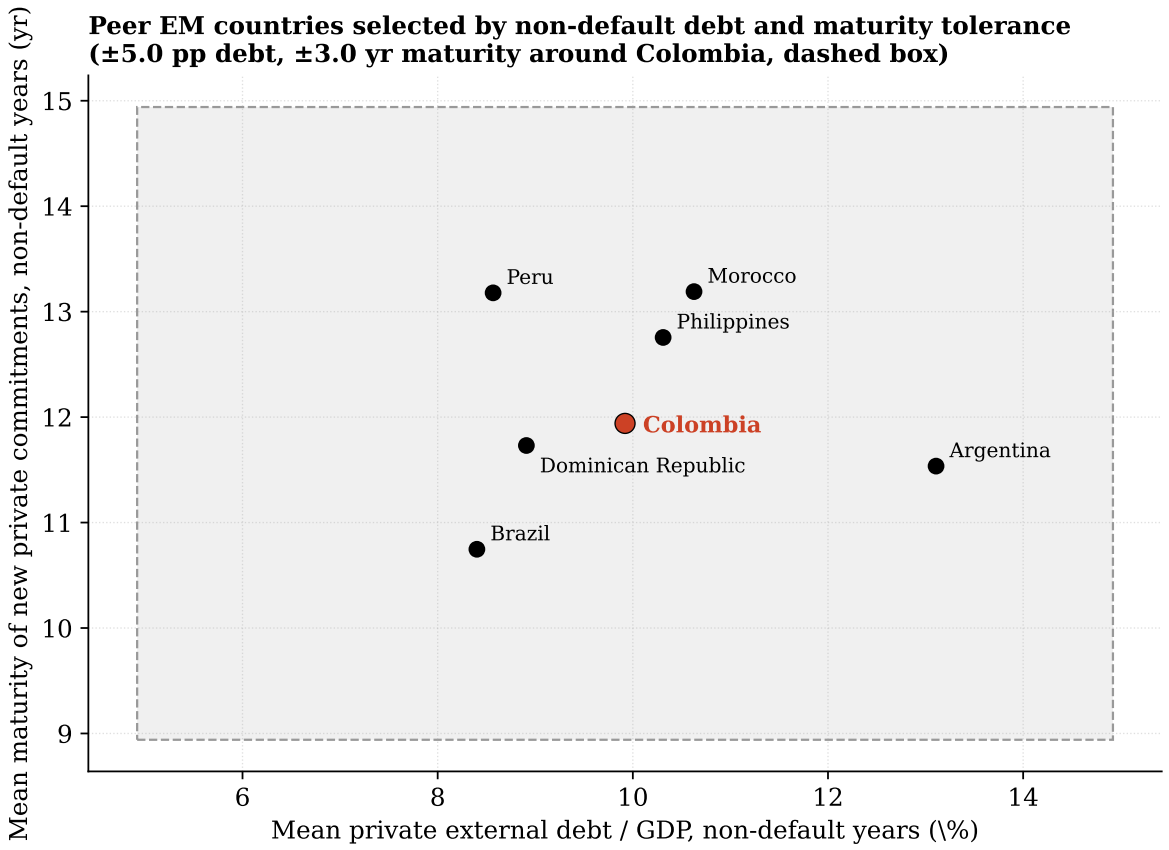


Figure 6: Comparability check: Colombia (gray dashed line) versus the broader pool of AT-PDW defaulters, 1980–2020. Left: mean private external debt / GDP. Right: mean maturity of new private commitments in non-default years. The six countries inside the dotted box on both dimensions form the peer subset.

gentina 2001–2005, Brazil 1982–1994, Dominican Republic 1982–1994, Morocco 1983–1987, Peru 1983–1997, Philippines 1983–1992) are the empirical sample used to construct the default-related target moments below. The event-study patterns of Section 2 also hold in this smaller sample (Appendix D).

The Market 2 maturity hazard $\lambda_2 = 0.107$ is the inverse of the mean maturity of new private commitments *during* default years in the six peers, which is 9.35 years. The mean maturity of new commitments is 12.19 years in non-default years. For all peers, the maturity is shorter in default years.[7]

[7]The same pattern holds for all 92 countries: mean maturity is 9 years in default years and 11.2 years in non-default years. One caveat is the timing of the defaults. Most sovereign default episodes occurred in the 1980s. The maturity of sovereign debt has a positive trend over time.

Internally calibrated parameters (SMM). The five SMM-determined parameters $(\beta, \theta, d_2, d_2^c, d_{aut})$ are jointly identified by the six targeted moments in Table 4. *Mean Market-1 debt to GDP* (target 9.92%, Colombia's private external debt over GDP, 1980–2020) is most informative of the discount factor β. *Mean episode length* (target 10.4 yr, mean length across the 7 length-$\geqslant 4$ peer episodes) is paired with the re-entry probability θ, the common value imposed in Section 3 on the $M_2 \to M_1$ and autarky $\to M_1$ transitions. The two within-episode moments pool the country-years of the 7 peer episodes and use new commitments to GDP, the closest data counterpart to new borrowing in the model. The *commit-asymmetry share* (target 0.690) is the share of these commitments that arrive in the first half of the episode, and it is paired with the M_2 output-penalty elasticity d_2^c. The *in-default to non-default new-commitment ratio* (target 0.590) is the pooled mean of commitments to GDP in the episode interiors, divided by the pooled mean in non-default years, again across the 7 peer episodes. It is paired with the autarky multiplier d_{aut}. The *mean spread* target (2.61%, Colombia's 10-year USD bond spread over 2002–2017) is most informative of the M_2 output-penalty cap d_2. The reason is that d_2 sets the level of the M_1 default boundary, and so the level of the equilibrium default-risk premium. The *default frequency* target (2.6% per country-year, the rate at which post-default episodes start across 119 low- and middle-income countries over 1980–2020) is not paired with any one parameter, since all of them affect it. Appendix B explains how I construct the spread and default-frequency targets. These pairings are only a guide, since all five parameters are determined jointly.

SMM loss. I choose the five parameters to minimize a weighted sum of squared percentage deviations of the model moments from their targets, with much less weight on the spread and the default frequency than on the other four moments (Appendix G gives the loss function and the weights). At the optimum, mean debt to GDP, episode length, the commit-asymmetry share, and the in-default to non-default ratio, which are the four flow and stock targets, are all within 7% of their data values (Table 4). The mean spread and the default frequency are undershot, for a structural reason.

Sources of the spread and default-frequency undershoot. In the model, the mean spread is 1.06% and the default frequency is 0.93% per year. The targets are a bond spread of 2.61% and the IDS-119 default hazard of 2.6%. A better optimizer would not close these gaps, because they come from a tension inside the model. Matching the empirical Fact 1 ratio of 0.59 requires Market 2 to be a moderately unattractive lender, so that borrowing during default stays strictly below borrowing in normal times. But when Market 2 is unattractive, defaulting in Market 1 is also less appealing, which limits the default risk

that Market 1 lenders have to price, and the equilibrium default frequency and spread both go down. Getting the spread and the default frequency up to their targets would need a more attractive Market 2, and that would push the ratio of in-default to non-default borrowing above one, against Fact 1. The model cannot deliver the Fact 1 ratio together with a high spread or a high default frequency, so the calibration gives priority to the within-episode moments, including the Fact 1 ratio. This is a standard limitation of one-asset sovereign-default models. Arellano (2008) also undershoots the empirical level of sovereign spreads, and so do Aguiar and Gopinath (2006), who point this out themselves. Allowing partial debt carryover when the sovereign moves from M_2 back to M_1 would ease the tension. With carryover, the sovereign would re-enter M_1 with $b_1 = \kappa\, b_2$ for some $\kappa \in [0, 1]$ instead of $b_1 = 0$, which makes re-entry itself costly and gives the depression in commitments after the episode a structural channel. I leave this extension for future work.

Untargeted moments. Beyond the calibration targets in Table 4, Table 6 reports additional moments produced by the model at the calibrated parameters, alongside data values for the 6-peer set and Colombia. The data moments use annual series, log-linearly detrended in the same way as in the income-process estimation (see the notes to Table 6). The magnitudes reported here do not match the canonical emerging-market business-cycle moments of Neumeyer and Perri (2005). Their quantification uses quarterly HP-filtered series. This paper reports the comparison on annual linearly detrended data, for internal consistency with the SMM calibration. The model matches Colombia's income volatility closely ($\sigma(\log y) = 4.03$ pp in the model versus 4.05 pp in the data) and the consumption-income comovement closely ($\rho(\log c, \log y) = 0.99$ versus 0.96), while understating relative consumption volatility ($\sigma(\log c)/\sigma(\log y) = 0.99$ versus 1.34). The spread–income correlation in the model (-0.72) is more negative than the typical empirical values across the peers and Colombia (which range from about -0.17 to -0.43), but of the right sign and order of magnitude. The NX to y correlation is closer to zero in the model than in the peers, a known weakness of this class of models, pointed out by Aguiar and Gopinath (2007).

Model-vs-data event study. Figure 7 compares the model's simulated event study around default episodes with the 6-peer data event study of Figure 12 in Appendix D, bin by bin. Inside the episode, the two flow margins match. The net flow on private debt is positive in the first half and close to zero in the second, in both the model and the data ($+0.60$ and $+0.06$ percent of GDP in the model against $+0.87$ and -0.08 in the data). This profile

Table 6: Model fit on untargeted business-cycle moments at the Colombia calibration, alongside the corresponding data values for Colombia and the six peer defaulters.

Moment	Arg	DR	Mor	Bra	Per	Phi	Col	Model
$\sigma(\log y)$ (pp)	8.22	5.72	3.56	5.11	14.41	10.10	4.05	4.03
$\sigma(\log c)/\sigma(\log y)$	1.20	0.88	1.29	1.04	0.57	0.36	1.34	0.99
$\rho(\log c, \log y)$	0.80	0.39	0.71	0.83	0.33	0.49	0.96	0.99
$\sigma(NX/y)$ (pp)	3.68	3.17	3.78	2.30	4.03	3.48	3.24	0.53
$\rho(NX/y, \log y)$	-0.56	-0.33	$+0.02$	-0.11	$+0.02$	-0.43	-0.18	$+0.14$
$\rho(\text{spread}, \log y)$	n.a.	n.a.	n.a.	-0.38	-0.32	-0.17	-0.43	-0.72
spread coverage	n.a.	n.a.	n.a.	1998–2016	2003–2016	2005–2016	2002–2017	

Notes: Annual real GDP and consumption 1980–2019, log-linearly detrended (2020 dropped to remove the COVID disturbance, matching the Colombia income-process estimation). Spread–income correlation: each country's Bloomberg USD 10-year sovereign yield minus DGS10, against linearly detrended log real GDP, over the country's spread-coverage window. Real GDP is from World Bank WDI in constant USD. Real consumption is indexed from quarterly real private-consumption growth (S&P Capital IQ Pro), chained to annual frequency. NX to y is exports minus imports as a share of GDP, from GMD. Argentina, the Dominican Republic, and Morocco lack pre-Colombia-window Bloomberg coverage (n.a.). The *Model* column uses a $T = 200{,}000$ annual simulation at the calibrated parameters, with a linear detrend on the log series. *Source:* World Bank WDI, S&P Capital IQ Pro, Bloomberg, FRED (DGS10), and IMF World Economic Outlook.

is not targeted by the calibration. The gross-flow contrast across halves is 2.1 to 1 in the model against 2.6 to 1 in the data. Here the calibration does target the related first-half commit share, so part of this match is built in. The two outer bins are outside the scope of the model. In the Before bin, the simulated sovereign is already reducing its debt on the way to default, while the peer economies were still borrowing heavily. In the After bin, the model produces a borrowing boom at re-entry, because the sovereign returns to Market 1 with zero debt and borrows again. In the data, there is at most a partial recovery. This failure in the After bin is the other side of the post-episode depression that I document in Section 2. Both come from the sovereign re-entering M_1 with no debt, and the debt-carryover extension discussed above is the most direct way to address both.

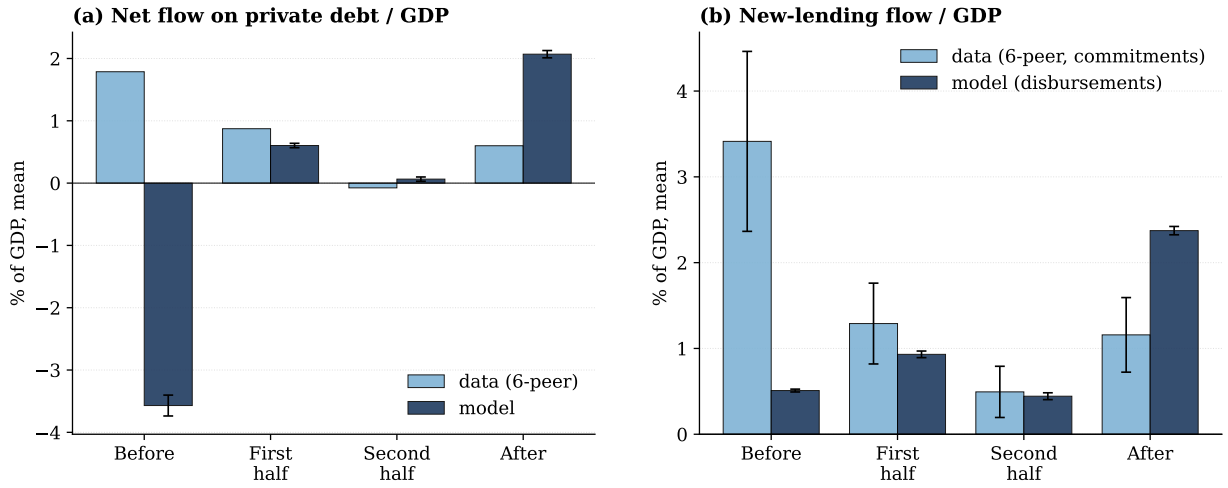


Figure 7: Model versus data around default episodes. Data: pooled 6-peer bins as in Figure 12. Panel (b) uses new commitments to GDP, and the data net-flow bins have no confidence bands. Model: simulated bins at the calibrated parameters with 95% confidence intervals. The model flow is new lending $b'_j - (1 - \lambda_j)b_j$ over output.

5 Counterfactuals

The main experiment removes the second private market from the model calibrated in Section 4. That shuts down lending by a second cohort of creditors, which is what my model adds to the standard setup with one creditor group and an exclusion state. The experiment is also close in spirit to the most-favored-creditor reading of *pari passu*. In two robustness sweeps, I then vary the M_2 output-penalty elasticity d_2^c and the re-entry probability θ one at a time, to find where the welfare gain changes sign. Removing M_2 gives a consumption-equivalent gain of $\Delta c = +0.28\%$ at the ergodic center $(y = 1, b = 0)$. This gain is robust to the harshness of M_2 and changes sign only when re-entry is fast enough, and the calibrated re-entry rate, $\theta = 0.0933$, is well inside the region where removing M_2 improves welfare.

Removing the second market: *pari passu*. I ask how welfare and equilibrium quantities change when the second private market is removed entirely. A default in Market 1 then sends the sovereign directly to autarky, which gives back the standard long-term-debt sovereign-default structure of Chatterjee and Eyigungor (2012). All other parameters stay at their calibrated values in Table 4.

To compare welfare, I use the consumption-equivalent gain Δc defined by the Lucas compensation

$$V^{cf}(y, b = 0) \;=\; \mathbb{E}\sum_{t=0}^{\infty} \beta^t \, u\big(c_t^{base}(\cdot \mid \Phi^{base}) \cdot (1 + \Delta c) \big), \tag{21}$$

evaluated at the ergodic center $(y = 1, b = 0)$, where $y = 1$ is the median of the persistent income process. Φ^{base} stands for the policy and price functions of the calibrated baseline economy, $c_t^{base}(\cdot \mid \Phi^{base})$ is date-t consumption along the baseline policy path from (y, b), and $V^{cf}(y, b)$ is the Market-1 value function in the counterfactual equilibrium without M_2. So Δc is how much baseline consumption would have to go up, by the same percentage in every period and state, for the baseline sovereign to be as well off as the counterfactual one. With CRRA utility $u(c(1 + \Delta c)) = (1 + \Delta c)^{1-\gamma} u(c)$, so the constant comes out of the discounted sum. For $\gamma = 2$ we have $u(c) = -1/c$, both value functions are negative, and equation (21) gives $1 + \Delta c = V^{base}/V^{cf}$. A positive Δc therefore means the sovereign is better off without M_2.

Removing Market 2 improves the sovereign's welfare by $\Delta c = +0.28\%$ at the ergodic center (Table 7). At zero M_1 debt, the gain is positive at every income state. Figure 8 shows it going from $+0.15\%$ at the highest y to $+0.56\%$ near the lowest.

Table 7: Equilibrium moments and welfare-equivalent consumption gain in the baseline two-market economy versus the counterfactual with M_2 removed, at the Colombia calibration.

	Baseline (M_2 active)	M_2 removed	Δ
default frequency (annual)	0.93%	0.39%	−0.54 pp
mean M_1 debt to GDP (in M_1)	10.4%	45.3%	+34.9 pp
mean spread (annualized)	1.06%	0.42%	−0.64 pp
sd of spread	0.75%	0.24%	−0.51 pp
corr(spread, log y)	−0.72	−0.78	−0.06
mean episode length	10.74 yr	10.56 yr	−0.19 yr
welfare gain Δc at $(y = 1, b = 0)$			+0.28%

Notes: Simulation at the parameters of Table 4, $T_{sim} = 200{,}000$ annual periods. Model spread: $\lambda_1/q_1 - \lambda_1 - r^f$, averaged over the simulated M_1 periods. This steady-state approximation ignores how q_1 depends on next-period borrowing through debt dilution. Δc is from equation (21) at $(y = 1, b_1 = 0)$.

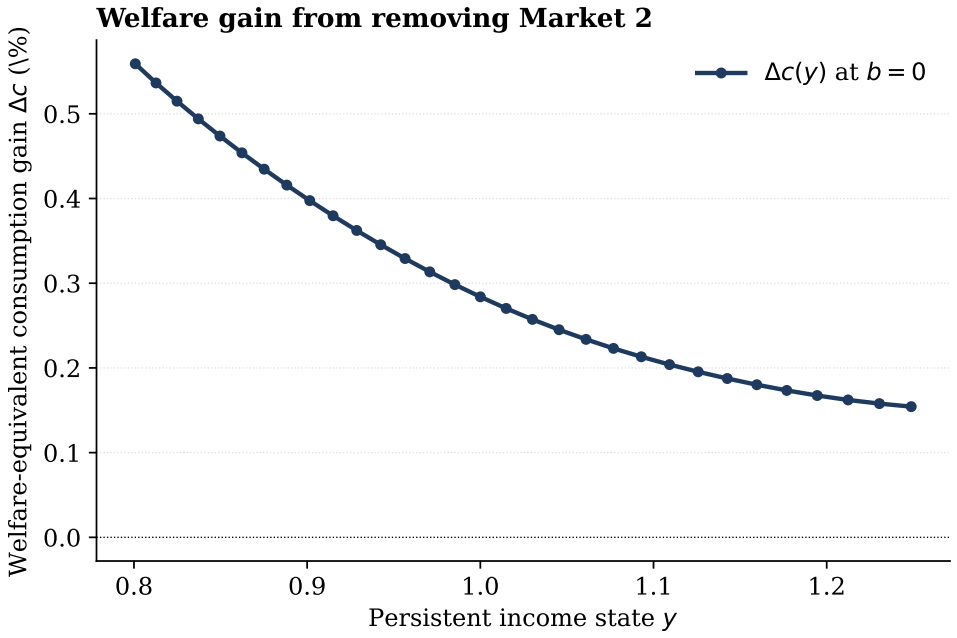


Figure 8: Welfare-equivalent consumption gain Δc from removing Market 2, by persistent income state y, evaluated at zero M_1 debt ($b = 0$). It is positive at every income state and broadly falls as y rises. Low-income states suffer more from the M_2 option in the baseline because below the kink at $y = d_2 \bar{y}$ (equation (20)) the M_2 output penalty is zero, so a default on M_1 costs them no output.

The gain comes from commitment, in a form that only exists with two markets. In the baseline, defaulting on L_1 closes off L_1 but not private credit in general, since Market 2 is still there. That puts a cap on what a Market 1 default can cost. So the sovereign defaults on M_1 more easily, and L_1 creditors, being competitive, charge for it from the start. The M_1 spread ends up high, at 1.06%, and borrowing capacity ends up small, at 10.4% of GDP. Take M_2 away and a default in M_1 leads straight to autarky, which is a strictly worse outside option. Defaults drop from 0.93% to 0.39% per year, the spread from 1.06% to 0.42%, and borrowing capacity goes up about four times, from 10.4% to 45.3% of GDP. In the model, the expansion of debt smooths consumption against income shocks more efficiently. The welfare gain is positive at every y, even though the sovereign now spends most of its time in autarky in the stationary distribution (90.5% versus close to 0%). The time in autarky rises mechanically, since $\theta = 0.0933$ implies a mean autarky spell of about 10.7 years.

Removing M_2 is the structural counterpart of strict across-cohort separation, which is suggestive of the most-favored-creditor reading of *pari passu*. In the baseline, L_1 lenders take a haircut at default and L_2 lenders then lend fresh money on clean terms. Read as in Olivares-Caminal (2013), *pari passu* is designed to deter exactly this kind of difference between cohorts. Without M_2 there is no second cohort, so no new creditors come in on better terms while the sovereign is in default on M_1, and the most-favored-creditor concern goes away. Note that the welfare gain from removing M_2 measures the consequence of strictly separating cohorts of private creditors. It is not a calculation of what enforcing *pari passu* would do, because as a contractual provision *pari passu* does not categorically prohibit new lending. The mapping defined here is structural only, and it should not be read in legal terms. Arellano et al. (2023) run the same experiment in their single-creditor partial-default model, where the across-vintage differential takes the form of newer bonds carrying smaller implicit haircuts than original issues. They likewise interpret eliminating new issuance during default as "adding more stringent *pari passu* clauses to the bond contracts."

M_2 harshness dimension. I vary the M_2 output-penalty elasticity d_2^c over the grid and keep the other SMM parameters at their calibrated values. At each grid point, I solve the model again with M_2 active and report the welfare-equivalent gain Δc from removing M_2.

From $d_2^c = 0.05$ up to about $d_2^c = 0.85$ the welfare gain stays strictly positive, peaking at around $+0.31\%$, and the value of $+0.28\%$ at the calibrated $d_2^c = 0.6411$ is just below this

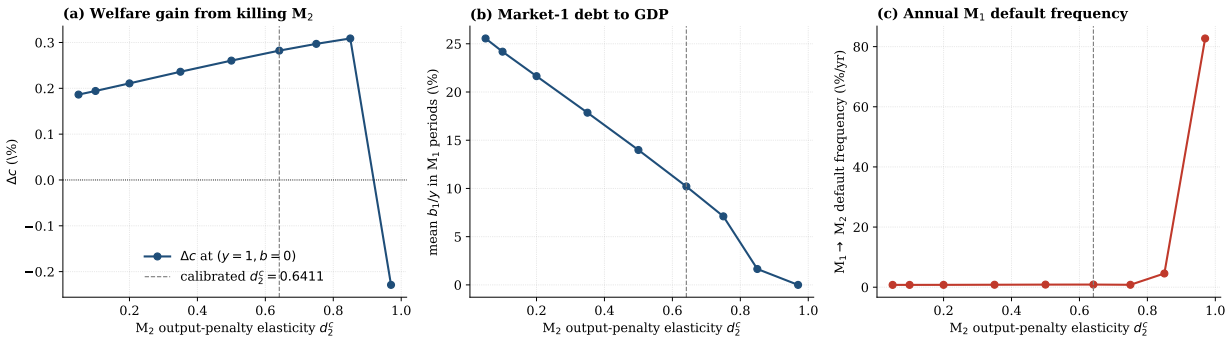


Figure 9: Regime sweep over the M_2 output-penalty elasticity $d_2^c \in [0.05, 0.97]$, holding all other SMM parameters at their values in Table 4. Panel (a) shows the welfare-equivalent gain Δc (%) from removing M_2 at the ergodic center $(y = 1, b = 0)$. Panel (b) shows mean Market-1 debt to GDP (%) in Market-1 periods, in the baseline equilibrium. Panel (c) shows the annual default frequency from M_1 to M_2 (%). The vertical dotted line in each panel marks the calibrated $d_2^c = 0.6411$.

peak. A higher d_2^c means a less harsh M_2, which makes default in M_1 cheaper.

Re-entry dimension. The second sweep varies the re-entry probability θ over the grid, with all other parameters at their calibrated values. Here θ governs two different transitions: in M_2 it is the per-period probability that the forgiveness offer arrives (which triggers the three-way decision of Section 3), and in autarky it is the per-period probability of automatic re-entry into M_1 with $b_1 = 0$.

Here the gain falls as θ rises and turns negative just below $\theta = 0.17$. Below that cutoff, re-entry is slow and the sovereign expects a long stay in M_2 or autarky, long enough for the M_2 output penalty to matter, so removing M_2 helps. At $\theta = 0.04$, the slowest re-entry rate I try, the gain is $+0.62\%$. Above the cutoff, re-entry is fast and both default states are short. M_2 then works as cheap insurance against bad M_1 realizations, and removing it hurts, with $\Delta c = -0.41\%$ at $\theta = 0.20$ and -4.41% at $\theta = 0.35$. At $\theta = 0.35$, the M_1 market has all but collapsed, because default in M_1 becomes close to costless under rapid re-entry and competitive L_1 creditors lend almost nothing. The calibrated $\theta = 0.0933$ is in the region where removing M_2 improves welfare, roughly 0.08 below the cutoff. The trade-off between the insurance effect and the commitment effect shows up most clearly in this sweep. In the d_2^c sweep, the sign changes only far into the corner where M_2 is soft.

One might think that this welfare result contradicts the borrowing documented in Section 2, but it does not, because the comparison is *ex ante*. *Ex post*, a sovereign already in default is better off using the second market than refusing it. What the *ex ante* result says

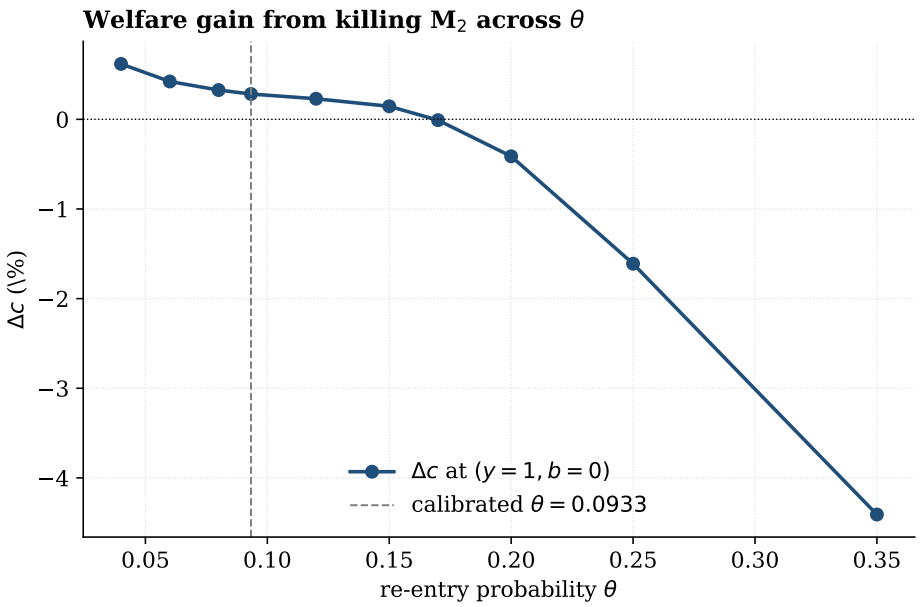


Figure 10: Re-entry sensitivity: welfare-equivalent gain Δc (%) from removing M_2 at the ergodic center $(y = 1, b = 0)$, as a function of the re-entry probability θ swept over the grid $[0.04, 0.35]$. All other SMM parameters are held at their values in Table 4. The vertical dotted line marks the calibrated $\theta = 0.0933$.

is that the equilibrium played by the sovereign and L_1 creditors when M_2 is available is worse for the sovereign than the one they would play if M_2 did not exist. In the calibrated regime the commitment effect dominates. The sovereign knows that M_2 will be there *ex post*, so it cannot credibly promise to repay L_1, and L_1 prices that in. An obvious next question is which policies could close this commitment gap. Candidates include creditor coordination treaties that bind second-market lenders to refuse credit during default, or an international institution doing the same job. The model gives one structural way to evaluate such proposals, but I leave that for future work.

6 Conclusion

Emerging-market sovereigns keep borrowing from private external creditors while they are in default. The access is limited, and most of it comes in the first half of the episode. I build a sovereign default model with two private markets that the sovereign reaches one after the other, and I ask whether this borrowing option is good for the sovereign.

In the calibrated model, it is not. The sovereign knows that a default on Market 1 does not cut it off from private credit. So it defaults more easily, and Market 1 lenders charge for this in normal times. In the Colombia calibration, the Market 1 spread is 1.06% and debt capacity is 10.4% of GDP. Without Market 2, the spread falls to 0.42% and debt capacity is about four times larger. The welfare gain from removing Market 2 is $\Delta c = +0.28\%$ of consumption at the ergodic center, and it is positive at every income level. The commitment effect is stronger than the insurance effect.

I read this as a coordination failure among private lenders. Each Market 2 lender prices its own loans fairly. But none of them cares about what its lending does to the terms the sovereign gets in normal times. The sovereign pays for this. Any arrangement that makes private creditors act as one group would remove the loss. The most-favored-creditor reading of *pari passu* is one such arrangement. Still, I do not want to push this too far. When re-entry is fast, Market 2 is cheap insurance, and removing it hurts. The answer depends on how fast the sovereign gets back to Market 1.

The empirical facts come from countries that still have some private access during default, so they are a selected group of defaulters. The welfare result comes from one calibration, to Colombia. The Mexico calibration in the appendix is only a partial check. And the *pari passu* reading is only my interpretation. Others are possible.

The model has one prediction that I have not tested yet. Episodes with more Market 2 borrowing should last longer, because a larger Market 2 debt makes the forgiveness offer more expensive to accept. Splitting the AT episodes by how much the sovereign borrows during default would be a direct test. On the model side, one extension is to let part of the Market 2 debt carry over to Market 1 at re-entry. The removal experiment then becomes a continuous policy lever. Carryover would also give the model a reason for the drop in borrowing after the episode, which Section 2 documents and the current model cannot produce.

References

Aguiar, Mark and Gita Gopinath, "Defaultable debt, interest rates and the current account," *Journal of International Economics*, 2006, *69* (1), 64–83.

_ and _, "Emerging market business cycles: The cycle is the trend," *Journal of Political Economy*, 2007, *115* (1), 69–102.

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

_ and Ananth Ramanarayanan, "Default and the maturity structure in sovereign bonds," *Journal of Political Economy*, 2012, *120* (2), 187–232.

_, Xavier Mateos-Planas, and José-Víctor Ríos-Rull, "Partial default," *Journal of Political Economy*, 2023, *131* (6), 1385–1439.

Asonuma, Tamon and Christoph Trebesch, "Sovereign debt restructurings: preemptive or post-default," *Journal of the European Economic Association*, 2016, *14* (1), 175–214.

Benjamin, David and Mark L. J. Wright, "Recovery Before Redemption: A Theory of Delays in Sovereign Debt Renegotiations," 2009. Working paper.

Bocola, Luigi and Alessandro Dovis, "Self-fulfilling debt crises: A quantitative analysis," *American Economic Review*, 2019, *109* (12), 4343–4377.

Bulow, Jeremy and Kenneth Rogoff, "Sovereign debt: Is to forgive to forget?," *American Economic Review*, 1989, *79* (1), 43–50.

Chatterjee, Satyajit and Burcu Eyigungor, "Maturity, indebtedness, and default risk," *American Economic Review*, 2012, *102* (6), 2674–2699.

Dvorkin, Maximiliano, Juan M Sánchez, Horacio Sapriza, and Emircan Yurdagul, "Sovereign debt restructurings," *American Economic Journal: Macroeconomics*, 2021, *13* (2), 26–77.

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

Erce, Aitor and Enrico Mallucci, "Selective sovereign defaults," *FRB International Finance Discussion Paper*, 2018, (1239).

Gelos, R Gaston, Ratna Sahay, and Guido Sandleris, "Sovereign borrowing by developing countries: What determines market access?," *Journal of International Economics*, 2011, *83* (2), 243–254.

Hatchondo, Juan Carlos and Leonardo Martínez, "Long-duration bonds and sovereign defaults," *Journal of International Economics*, 2009, *79* (1), 117–125.

Mendoza, Enrique G and Vivian Z Yue, "A general equilibrium model of sovereign default and business cycles," *The Quarterly Journal of Economics*, 2012, *127* (2), 889–946.

Neumeyer, Pablo A and Fabrizio Perri, "Business cycles in emerging economies: the role of interest rates," *Journal of Monetary Economics*, 2005, *52* (2), 345–380.

Olivares-Caminal, Rodrigo, "The Pari Passu Clause in Sovereign Debt Instruments: Developments in Recent Litigation," Technical Report 72, Bank for International Settlements 2013.

Pitchford, Rohan and Mark LJ Wright, "Holdouts in sovereign debt restructuring: A theory of negotiation in a weak contractual environment," *Review of Economic Studies*, 2012, *79* (2), 812–837.

Rouwenhorst, K. Geert, "Asset pricing implications of equilibrium business cycle models," in Thomas F. Cooley, ed., *Frontiers of Business Cycle Research*, Princeton University Press, 1995, pp. 294–330.

Schumacher, Julian, Christoph Trebesch, and Henrik Enderlein, "Sovereign defaults in court," *Journal of International Economics*, 2021, *131*, 103388.

Tomz, Michael and Mark LJ Wright, "Do countries default in "bad times"?," *Journal of the European Economic Association*, 2007, *5* (2–3), 352–360.

Yue, Vivian Z, "Sovereign default and debt renegotiation," *Journal of International Economics*, 2010, *80* (2), 176–187.

A AT-PD strict-subsample robustness

The main empirical section restricts to AT-PDW episodes: cases where missed payments occurred, whether the restructuring was post-default or weakly preemptive. As a robustness check this appendix re-runs the regressions of Table 2 on the strict AT post-default subsample (AT-PD), which drops the weakly preemptive episodes. Dropping weakly preemptive years can split a single AT-PDW episode into multiple shorter AT-PD episodes (Appendix B).

Table 8: Facts 1 and 2 on the AT-PD strict subsample (post-default only, weakly preemptive episodes dropped).

	All 92 countries			HIPC excluded		
	Private DIS	Bond+Bank DIS	Commitments	Private DIS	Bond+Bank DIS	Commitments
Fact 1: default + post-episode panel regression (country + year FE)						
Default-year dummy (Fact 1)	-0.663***	-0.464**	-1.281***	-0.865***	-0.673**	-1.715***
	(0.243)	(0.207)	(0.435)	(0.321)	(0.293)	(0.460)
Post-episode dummy, 5 yr	-0.769***	-0.516***	-1.417***	-1.007***	-0.767***	-1.365***
	(0.163)	(0.125)	(0.306)	(0.219)	(0.172)	(0.298)
Observations	3,360	3,360	2,096	2,092	2,092	1,543
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Fact 2: within-episode regression (episode FE, log GDP control)						
First-half dummy (Fact 2)	+0.998*	+0.492	+0.396*	+0.445**	+0.103	+0.393**
	(0.559)	(0.374)	(0.201)	(0.183)	(0.118)	(0.193)
Observations	341	341	226	120	120	81
Episodes	46	46	41	22	22	21
Episode FE	Yes	Yes	Yes	Yes	Yes	Yes
log(GDP) control	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The sample is AT-PD interior years, 1980–2020. All dependent variables are in percent of GDP. Private DIS is total disbursements from private creditors, Bond+Bank DIS keeps only bonds and commercial banks, and Commitments is new commitments from private creditors. The HIPC excluded columns drop the 33 HIPC completion-point countries. Standard errors are in parentheses, clustered by country in the top panel and by episode in the bottom panel. *p $<$ 0.10, **p $<$ 0.05, ***p $<$ 0.01. *Source:* World Bank International Debt Statistics and IMF World Economic Outlook.

The qualitative facts hold under the strict classification. Default years still show a lower level for all dependent variables, in both the full panel and the sample without HIPC countries. The default-dummy coefficients range from -0.46 to -1.72 pp across dependent variables, and all are significant at the 5% level or below. The first-half premium within the episode also stays positive for all dependent variables. The Private DIS first-half coefficient stays significant both in the full panel ($+1.00$, $p < 0.10$) and without HIPC countries ($+0.45$, $p < 0.05$). The coefficient for bonds and banks only falls to an insignificant value in the strictest combination of AT-PD, bonds and banks, and no HIPC countries ($+0.10$, $p = 0.39$). This weaker result comes from the AT-PD classification, which breaks some AT-PDW episodes into shorter intervals and so dilutes the contrast between the first

and the second half. So the Fact 2 result in the main text is not driven by the inclusion of weakly preemptive episodes.

B Data Sources

This appendix describes the variables, the datasets, and the full list of AT-PDW episodes used in the paper.

Table 9 lists every variable in the empirical analysis, together with its World Bank indicator code where there is one, and its source dataset.

Table 9: Data sources used throughout the paper.

Object	Code	Provider
New private-creditor commitments (USD)	DT.COM.PRVT.CD	World Bank IDS
New private-creditor disbursements (USD)	DT.DIS.PRVT.CD	World Bank IDS
Private external net flow (USD)	DT.NFL.PRVT.CD	World Bank IDS
PPG external debt to private creditors (USD)	DT.DOD.PRVT.CD	World Bank IDS
Avg. maturity, new priv. commitments (yr)	DT.MAT.DPPG.PR	World Bank IDS
GDP, current USD (primary)	NGDPD	IMF WEO
GDP, current USD (fallback)	NY.GDP.MKTP.CD	World Bank WDI
Real GDP, constant 2015 USD	NY.GDP.MKTP.KD	World Bank WDI
Real GDP growth (annual %)	NGDP_RPCH	IMF WEO
US implicit GDP deflator (base 2015)	–	US BEA / FRED
10-year US Treasury constant-maturity yield	DGS10	FRED
Colombia 10Y USD synthetic sovereign yield	I248 (COLOM 8 11/14/35 ref.)	Bloomberg
Brazil 10Y USD synthetic sovereign yield	BRAZIL sovereign curve	Bloomberg
Peru 10Y USD synthetic sovereign yield	PERU sovereign curve	Bloomberg
Philippines 10Y USD synthetic sovereign yield	PHILIPPINES sovereign curve	Bloomberg
Mexico 10Y USD synthetic sovereign yield	MEXICO sovereign curve	Bloomberg
Real private consumption growth (quarterly)	–	S&P Capital IQ Pro
Exports / GDP, Imports / GDP	–	GMD (IMF)
Default-episode dating (baseline)	AT_PDW	Asonuma and Trebesch (2016)

The 92-country regression panel. The estimation sample is the panel of 92 emerging-market and developing economies that satisfy the two selection rules in Section 2. Every country is in the World Bank IDS country list (119 sovereign entries), has at least one non-missing IDS observation over 1980–2020, and has a population of at least two million. Below, I list first the 49 countries with at least one AT-PDW episode starting in 1980–2020, and then the 43 without one.

Defaulters (49): Albania, Algeria, Argentina, Bolivia, Bosnia and Herzegovina, Brazil, Cameroon, Democratic Republic of the Congo, Dominican Republic, Ecuador, Ethiopia, Gabon, Gambia, Guinea, Honduras, Iraq, Ivory Coast, Jamaica, Jordan, Kenya, Liberia, Madagascar, Malawi, Mauritania, Mexico, Moldova, Morocco, Mozambique, Nicaragua, Niger, Nigeria, Pakistan, Paraguay, Peru, Philippines, Republic of the Congo, Senegal, Serbia, Sierra Leone, South Africa, Sudan, Tanzania, Togo, Turkey, Uganda, Ukraine, Vietnam, Yemen, Zambia.

Non-defaulters (43): Afghanistan, Angola, Armenia, Azerbaijan, Bangladesh, Belarus, Benin, Botswana, Burkina Faso, Burundi, Cambodia, Central African Republic, Chad, China, Colombia, Egypt, El Salvador, Georgia, Ghana, Guatemala, Haiti, India, Indonesia, Iran, Kazakhstan, Kyrgyzstan, Laos, Lebanon, Lesotho, Mali, Mongolia, Myanmar, Nepal, Papua New Guinea, Rwanda, Somalia, Sri Lanka, Tajikistan, Thailand, Tunisia, Turkmenistan, Uzbekistan, Zimbabwe.

AT sub-classifications. Section 2 describes the three categories into which Asonuma and Trebesch (2016) classify each sovereign-debt restructuring. The baseline sample for the empirical analysis combines the post-default and weakly preemptive episodes, and I denote it `AT_PDW`. The strict post-default subsample, `AT_PD`, drops the weakly preemptive episodes. Table 2 in the main text reports the results for AT-PDW, and Table 8 in Appendix A reports the same results for AT-PD.

The two labelings give different episode lists, and one is not simply a subset of the other. Under AT-PDW, consecutive years that are either post-default or weakly preemptive are joined into a single episode. Under AT-PD, the weakly preemptive years are excluded, and this can split a single AT-PDW episode into several shorter AT-PD episodes. For example, Argentina 1982–1993 is one AT-PDW episode of length 12 (post-default 1982–1985, weakly preemptive 1986–1987, post-default 1988–1993), but it is two separate AT-PD episodes (1982–1985 with length 4, and 1988–1993 with length 6). The merging under AT-PDW and the splitting under AT-PD roughly cancel out in the counts. At length $\geqslant 4$, both labelings give 48 episodes, in 43 and 42 countries, respectively. At any length, AT-PDW

has 66 episodes in 49 countries, and AT-PD has 60 episodes in 45 countries.

Borrowing variables. Section 2 defines the three IDS borrowing series, `DT.DIS.PRVT.CD`, `DT.COM.PRVT.CD`, and `DT.NFL.PRVT.CD`. To compute real values, I use the US implicit GDP deflator (base 2015).

Sovereign spread. The spread target in the calibration (Section 4) is Colombia's 10-year USD synthetic sovereign yield minus the 10-year US Treasury constant-maturity yield. For Colombia I take Bloomberg's USD sovereign curve (identifier I248), anchored on the on-the-run reference bond COLOM 8 11/14/35 (mid yield to maturity). The US yield is FRED's DGS10. I take monthly averages of both series and subtract. For Colombia the window runs from 2002-04 to 2017-01. The mean spread is 2.61% (Table 4). The spreads for the other peers in the untargeted-moments table (Brazil, Peru, and the Philippines) are built the same way, each from its own Bloomberg sovereign curve.

Default frequency. The default-frequency target is the post-default (AT-PDW) hazard rate among the 119 low- and middle-income sovereigns in the World Bank IDS universe, 1980–2020. I count 110 post-default starts in roughly 4,096 at-risk country-years, about 2.69%, and truncate this to 2.6% for the target.

GDP. Most of the nominal GDP data in current USD come from the IMF WEO series NGDPD. Before 1980, and wherever WEO is missing, I use the World Bank WDI series NY.GDP.MKTP.CD instead, and in the years where both exist the two series are close. Real GDP in constant 2015 USD is the World Bank WDI series NY.GDP.MKTP.KD.

Maturity and the external private-creditor debt stock. Average maturity of new private commitments (`DT.MAT.DPPG.PR`, in years) is the empirical analog of the model's $1/\lambda_1$ when averaged over a country's non-default years (and of $1/\lambda_2$ when averaged over its default years). The stock of sovereign-backed external debt owed to private creditors (`DT.DOD.PRVT.CD`, USD current) is the empirical analog of the model's debt stock b.

The 48 length-$\geqslant 4$ AT-PDW episodes (1980–2020). Table 10 lists every AT-PDW episode of length at least four years in the 1980–2020 window. Of these 48 episodes, 42 are post-default (PD) and 6 are weakly preemptive (WP).

Table 10: Length-$\geqslant$ 4 AT-PDW default episodes, 1980–2020.

Country	Episode	Length	AT
Albania	1991–1995	5	PD
Algeria	1993–1996	4	PD
Argentina	1982–1993	12	PD
Argentina	2001–2005	5	PD
Bolivia	1980–1993	14	PD
Bosnia and Herzegovina	1992–1997	6	PD
Brazil	1982–1994	13	WP
Cameroon	1985–2003	19	PD
Democratic Republic of the Congo	1982–1989	8	PD
Dominican Republic	1982–1994	13	PD
Ecuador	1982–1995	14	WP
Ethiopia	1990–1996	7	PD
Gabon	1989–1994	6	PD
Gambia	1984–1988	5	PD
Guinea	1985–1988	4	PD
Guinea	1991–1998	8	PD
Honduras	1981–2001	21	PD
Iraq	1986–2006	21	PD
Ivory Coast	1983–1998	16	PD
Ivory Coast	2000–2012	13	PD
Jamaica	1986–1990	5	WP
Jordan	1989–1993	5	PD
Kenya	1992–1998	7	PD
Madagascar	1981–1990	10	PD
Mauritania	1992–1996	5	PD
Mexico	1982–1990	9	WP
Moldova	2001–2004	4	PD
Morocco	1983–1987	5	PD
Mozambique	1983–1991	9	PD
Mozambique	2016–2019	4	PD
Nicaragua	1983–1995	13	PD
Niger	1983–1991	9	WP
Nigeria	1986–1991	6	PD

Table 10 continued

Country	Episode	Length	AT
Paraguay	1986–1993	8	PD
Peru	1983–1997	15	WP
Philippines	1983–1992	10	PD
Republic of the Congo	1983–2007	25	PD
Senegal	1981–1985	5	PD
Senegal	1992–1996	5	PD
Serbia	1992–2004	13	PD
Sierra Leone	1980–1995	16	PD
Sudan	1980–1985	6	PD
Tanzania	1981–2004	24	PD
Togo	1991–1997	7	PD
Uganda	1980–1993	14	PD
Vietnam	1982–1997	16	PD
Yemen	1983–2001	19	PD
Zambia	1983–1994	12	PD

Notes: AT sub-classifications: *PD* = post-default, *WP* = weakly preemptive. *Source:* Asonuma and Trebesch (2016) and World Bank International Debt Statistics.

C Commitment-based robustness for Section 2

Table 1 in the main text shows that disbursements and new commitments track each other closely on the country-years where IDS reports both.

Event-study robustness on commitments. Figure 11 reports the event study with new commitments (DT.COM.PRVT.CD) as the level series, on the 22 AT-PDW episodes with commitment data in all four bins. The commitment sample is smaller than the disbursement sample, because IDS does not report commitments in country-years with zero private creditor activity. These are the 1,272 country-years of the gap sample discussed in the main text. The qualitative pattern is the same as in Figure 2, which uses disbursements. Real commitments collapse from the Before bin to the first half and stay low through the second half and the After bin. The commit-to-GDP ratio is front-loaded within the episode, with a first-half to second-half ratio of 1.80 here against 2.13 on the disbursement-based 35-episode sample. The debt-stock panel shows the same hump shape.

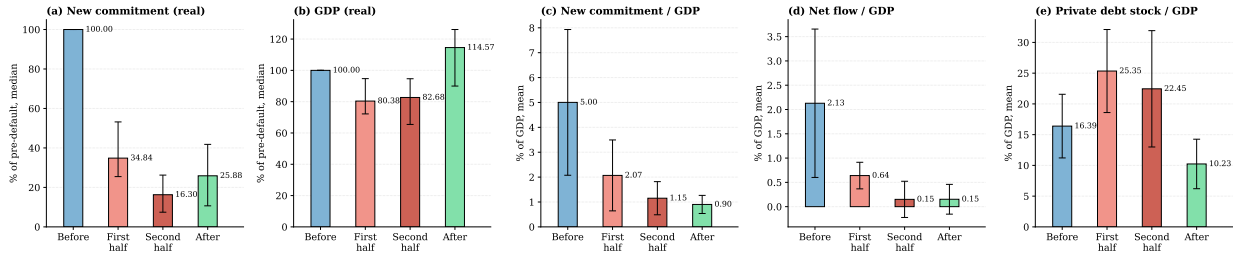


Figure 11: Event-study robustness on new commitments, AT-PDW sample (22 length-$\geqslant 4$ episodes with non-zero, non-missing new commitments in all four bins). Panels (a) and (b): real new commitments and real GDP, indexed to Before $= 100$ (cross-episode medians, bootstrap 95% CIs). Panel (c): new commitment / GDP (cross-episode means, asymptotic 95% CIs). Panel (d): net flow on private external debt / GDP (cross-episode means, asymptotic 95% CIs, World Bank IDS series DT.NFL.PRVT.CD). Panel (e): private external debt stock / GDP (cross-episode means, asymptotic 95% CIs, World Bank IDS series DT.DOD.PRVT.CD).

D Event study on the peer sample

Before committing to the calibration, I check that the empirical patterns documented in Section 2 on the broader AT-PDW sample also hold on the 7-episode peer subset that I use for the default-related targets in Section 4. Figure 12 reports the result, replicating Figure 2 on the 6-peer pooled sample. The same qualitative shape carries through. Real private disbursements decline steeply from Before to First half (from 100 to 37 on the indexed series, and from 3.41% to 1.29% of GDP), decline further in the Second half (15 indexed, 0.49% of GDP), and recover only partially in the After bin (1.16% of GDP). The debt stock has the same hump as in the broad sample, going from 14.5% of GDP in the Before bin to 25.0% in the first half, then down to 17.3% in the second half and 9.6% in the After bin. Net borrowing is positive in the Before bin (1.79% of GDP) and in the first half (0.87%), and close to zero in the second half (−0.08%), as in the broad sample of Section 2.

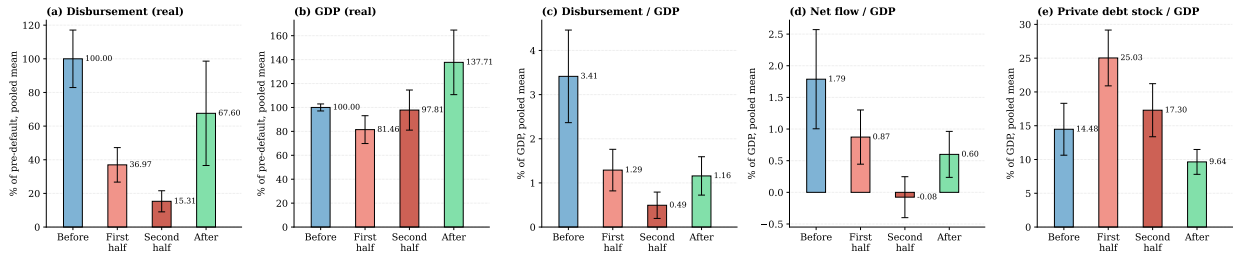


Figure 12: Event study around default onset, 6-peer subsample (7 episodes of length $\geqslant 4$). Pooled means across episode-year observations, 95% asymptotic CIs. Panels (a) and (b): real private disbursements and real GDP, indexed to each episode's Before bin $= 100$. Panel (c): disbursement / GDP. Panel (d): net flow on private external debt / GDP. Panel (e): private external debt stock / GDP.

E Net-flow decomposition

Disbursements are a gross flow. The *net* flow, $DIS_t - AMT_t$, also subtracts AMT_t, the amortization on the existing stock in year t, so it is the closest data counterpart to the net new private debt taken on in a year. Table 11 reports Facts 1 and 2 with net flows decomposed by creditor type.

Table 11: Facts 1 and 2 on net flows (DIS − AMT) by creditor type, AT-PDW.

	Total NET	Bonds NET	Banks NET	Other-priv. NET
F1 default-year	-0.417**	-0.171	-0.253*	-0.082
	(0.201)	(0.104)	(0.140)	(0.085)
F1 post-episode	-0.582***	-0.181**	-0.271***	-0.166***
	(0.133)	(0.073)	(0.085)	(0.058)
F2 first-half dummy	+0.449***	-0.044**	+0.220***	+0.283***
	(0.110)	(0.022)	(0.068)	(0.078)
Observations (Fact 1)	3,526	3,351	3,351	3,351
Observations (Fact 2)	373	373	373	373
Country FE / Year FE (Fact 1)	Yes	Yes	Yes	Yes
Episode FE / log(GDP) (Fact 2)	Yes	Yes	Yes	Yes

Notes: The sample is AT-PDW interior years, 1980–2020. The dependent variables are net flows by creditor type, disbursements minus amortization, in percent of GDP. Standard errors are in parentheses, clustered by country in the Fact 1 rows and by episode in the Fact 2 row. *p < 0.10, **p < 0.05, ***p < 0.01. *Source:* World Bank International Debt Statistics and IMF World Economic Outlook.

Total net flows fall in default years. The default-year coefficient is negative for every creditor type, but it is significant only for banks. Raw averages in default years look different. Bond net flows are negative because existing bonds are amortized and no new bonds are issued. Bank and other private net flows stay positive. Within the episode, the first-half dummy is positive and significant for total, bank, and other private net flows. So the front-loading in Fact 2 comes from bank credit and other private credit. Bond net flows do not front-load. Their first-half coefficient is small and slightly negative, so the bond stock runs down evenly over the episode.

F Per-creditor disbursement regression

Table 3 in the main paper reports the dollar-weighted shares of bond, commercial-bank, and other private disbursements in default and non-default years. Here I run the same Fact 1 specification once for each disbursement sub-series.

The creditor groups are bondholders, commercial banks, and other private creditors, where the last group covers credits from manufacturers, exporters, and other suppliers of goods, and bank credits guaranteed by an export credit agency (the IDS definition in the notes to Table 3). The results in Table 12 use a common sample of 2,087 observations. For the bank series and for total private disbursements, the default-year coefficient is significant at the 5% level. For the "other private" residual it is significant at the 10% level, and for bonds it is statistically indistinguishable from zero. Bond disbursements are already close to zero during default, though. The bond regression is estimating a null within a narrow range around zero, and it does not rule out a sizable fall. That matches the main text, where commercial banks are the main lenders during default and non-bank private creditors form a large second group.

Table 12: Creditor-type decomposition of the Fact 1 regression.

	Total (DIS)	Bonds (DIS)	Banks (DIS)	Other priv. (DIS)	Commitments
Default dummy	-0.812**	-0.149	-0.403**	-0.260*	-1.334***
	(0.317)	(0.198)	(0.183)	(0.134)	(0.427)
Post-episode dummy (5 yr)	-1.100***	-0.377***	-0.467***	-0.256**	-1.394***
	(0.239)	(0.135)	(0.133)	(0.121)	(0.306)
Mean y, non-default yrs	1.76	0.77	0.55	0.44	1.98
Mean y, default yrs	1.40	0.03	0.62	0.76	1.44
Mean y, post-episode yrs	0.78	0.33	0.23	0.22	0.90
Country FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	2087	2087	2087	2087	2087
R^2 (within)	0.016	0.012	0.002	-0.014	0.017

Notes: AT-PDW interior years, 1980–2020. The dependent variables are in percent of GDP. The FE specification is the same as in Table 2, with country and year FE and SE clustered by country. Columns 1–4 are IDS disbursement sub-series, and column 5 is new commitments. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. *Source:* World Bank International Debt Statistics and IMF World Economic Outlook.

G Solution and calibration details

I solve the model of Section 3 numerically by value-function iteration. The iteration is adapted to the three-envelope structure created by the three-way decision in the F regime.

Transitory shock. The transitory shock m is there for numerical reasons, because it smooths the bond-pricing operator and lets the price iteration converge. Its standard deviation, $\sigma_m = 0.005$, is the smallest round number at which the value-function iteration converges robustly on the calibration grid, and it is small enough that the simulated moments stay close to those in the limit with $m \equiv 0$. The endpoint $\bar{m}$ of its support is unrelated to $\bar{y}$, the long-run mean of income. In the period of a default, m is fixed at $-\bar{m}$ (equations (7) and (9)), so the default decision does not depend on the realized m, while in every later autarky period W^{Aut} uses the actual draw of m through that period's budget.

Discretization. I discretize the persistent shock y with the Rouwenhorst (1995) method, which matches the AR(1) persistence exactly. The transitory shock m is put on a grid spanning $\pm 2\sigma_m$, and the expectations over m' use midpoint quadrature with truncated-normal weights. The debt grid allows small positive asset holdings, and it goes high enough for the debt that builds up in the counterfactuals, where defaults are rare. The two sensitivity sweeps in Section 5 use a coarser debt grid and shorter simulations for speed. At the calibrated parameters, the coarse grid gives the same Δc of $+0.28\%$. The sweep figures leave out the grid points at which the iteration does not converge ($\theta = 0.02$ and 0.50, and $d_2^c = 0.92$).

Three envelopes. Given the persistent state y, the within-period Bellman problem comes down to choosing b_i' to maximize $u(c) + \beta Z_i(y, b_i')$, with c linear in m. As $u'(c) > 0$ and m enters consumption linearly, every candidate b_i' traces a line in m-space with slope $u'(c_0)$ and intercept $u(c_0) + \beta Z_i(y, b_i')$, and the optimal $b_i'(m)$ is a piecewise constant step function of m whose switch points come from pairwise indifference between adjacent candidates (Chatterjee and Eyigungor 2012, Lemma 2(ii)). In each iteration I keep track of three such envelopes:

- `env_M1`: $c_0 = y - \lambda_1 b_1 + q_1(y, b_1')(b_1' - (1-\lambda_1)b_1)$. Used to construct V_1 and the M_1 default threshold against $V_d^{M_1} = W_2^{NF}(y, -\bar{m}, 0)$.
- `env_M2`: $c_0 = y_2^{def}(y) - \lambda_2 b_2 + q_2(y, b_2')(b_2' - (1-\lambda_2)b_2)$. Used to construct V_2 and the NF default threshold against $V_d^{NF} = W^{Aut}(y, -\bar{m})$.

- `env_M2_F`: $c_0 = y - b_2 + q_1(y, b_1')b_1'$, the M_1 budget under full settlement. Used to construct V_1^F and the F-acceptance threshold against W_2^{NF}.

F-acceptance threshold. The F-acceptance threshold $m_d^F(y, b_2)$ is the m at which the sovereign is indifferent between accepting F and staying in NF, $V_1^F(y, m, b_2) = W_2^{NF}(y, m, b_2)$. Both sides are piecewise linear in m, being upper envelopes of finitely many b_i' branches (on the right side, together with the autarky branch below the NF default threshold), so their difference $V_1^F(m) - W_2^{NF}(m)$ is piecewise linear too and can change sign more than once in $[-\bar{m}, \bar{m}]$. I follow a *rightmost-crossing convention*, setting m_d^F to the highest m at which the difference changes sign. F is then accepted in the upper tail $[m_d^F, +\bar{m}]$, where the sovereign has enough resources to settle b_2 in one payment. This crossing moves continuously with the prices, so the price iteration stays continuous and a fixed point exists.

Iteration loop. Each iteration starts from the current prices (q_1, q_2). Given these prices, I build the three envelopes and find the thresholds for M_1 default, NF default, and F acceptance. Then I update the continuation values Z_1, Z_2, and Z^{Aut}. The integral over m' is computed analytically within each piece of the envelope, and the sum over y' uses the Markov chain. The pricing equations (18) and (19) then give new prices. Write $\mathcal{H}[q_i^{(k)}]$ for the right-hand side of the pricing equation at iterate k. As in Chatterjee and Eyigungor (2012), I damp the price update,

$$q_i^{(k+1)}(y, b_i') \;=\; (1-\zeta)\cdot\big(\mathcal{H}[q_i^{(k)}]\big)(y, b_i') \;+\; \zeta\cdot q_i^{(k)}(y, b_i'), \tag{22}$$

with $\zeta = 0.5$, and stop when the prices in both markets change by less than 10^{-5}.

Simulation and moments. I compute the model moments from one long simulation that starts with zero debt in Market 1, after a burn-in. Each period draws y' and m', applies the sovereign's decision, and records the outcome. Both the calibration and the counterfactuals use the value functions integrated over m. The commit-asymmetry share is built the same way as in the data, with the episode halves of Section 2, and new commitments are zero in autarky years. The share is $\bar{c}_{FH}/(\bar{c}_{FH} + \bar{c}_{SH})$, where $\bar{c}_{FH}$ and $\bar{c}_{SH}$ are the pooled first-half and second-half means of commitments to GDP.

SMM loss. The SMM loss is a weighted sum of squared relative deviations over the six targeted moments,

$$\mathcal{L}(\Theta) \;=\; \sum_{k=1}^{6} w_k \big((\hat{m}_k - m_k(\Theta))/|\hat{m}_k| \big)^2, \tag{23}$$

with weight $w_k = 1$ on mean debt to GDP, episode length, the commit-asymmetry share, and the in-default to non-default ratio, and $w_k = 0.001$ on the mean spread and the default frequency. I minimize it by Nelder–Mead with sigmoid box transforms on each parameter.

Counterfactual without $\mathbf{M}_2$. In the counterfactual, W_2^{NF} and W_2^{F} are replaced by W^{Aut}. There is no b_2 to settle, so the forgiveness branch is inactive and the three-way decision in the F regime reduces to the autarky branch.

H Calibration robustness: Mexico

The calibration in Section 4 is fit to Colombia, with default-related moments from the six-peer set because Colombia has no AT-PDW episode of its own. To check the welfare conclusion of Section 5, I recalibrate the model to Mexico. Mexico is not one of the six peers, and it had an AT-PDW episode of its own, the post-default restructuring of 1982–1990. Its episode length, within-episode moments, and λ_2 can therefore come from Mexico's own data. Mexico also fails the peer box on debt, with a mean ratio of private external debt to GDP of 16.7% over 1980–2020 against Colombia's 9.92% and the ± 5 pp box.

Country-specific exogenous parameters. For the risk-free rate I take the mean DGS10 yield over 2000–2017, $r^f = 3.53\%$. This window starts in 2000 because that is when Mexico's Bloomberg 10-year USD sovereign yield begins, and, as for Colombia, it matches the spread window. Re-estimating the income process on Mexico's WEO real GDP with the same linear detrend gives $\rho_y = 0.724$ and $\sigma_y = 0.031$. Mexico's non-default years give $\lambda_1 = 0.064$, and its own default years give $\lambda_2 = 0.114$, a mean maturity of 8.8 years. Risk aversion and the dispersion of the transitory shock are left at $\gamma = 2$ and $\sigma_m = 0.005$.

Country-specific SMM targets. I build the six targets as in Section 4, but with Mexico data. Mean debt to GDP is 16.7%. The commit-asymmetry share on the 1982–1990 episode is 0.61, the in-default to non-default commitment ratio is 0.71, and the episode lasts 9 years. The mean spread is 1.73%. The spread is the Mexico 10-year USD sovereign yield minus DGS10 over 2000–2017. The sixth moment is the panel default frequency of 2.6%, which is not country-specific.

NM optimum. Nelder–Mead converges to the optimum in Table 13. The four primary moments fit about as well as for Colombia, with mean debt to GDP of 18.6% in the model against 16.7% in the data, a commit-asymmetry share of 0.716 against 0.610, an in-default to non-default ratio of 0.733 against 0.710, and an episode length of 9.17 against 9 years. As for Colombia (Section 4), the spread and the default frequency are undershot, at 0.55% and 0.49%, for the same structural reason. With the optimum in hand, I re-run the removing-M_2 counterfactual.

Removing M_2 at the Mexico calibration. For Mexico, removing M_2 gives $\Delta c_{MEX} = +0.175\%$ at the ergodic center. The sign is the same as for Colombia, and the gain is

Table 13: Mexico calibration robustness: externally fixed parameters, internally estimated parameters, and targeted moments for the Mexico recalibration of the two-market sovereign-default model.

Parameter / Moment	Target / Value	Model
Externally fixed parameters		
Relative risk aversion, γ	2.00	
Risk-free rate, r^f	3.53%	
Income persistence, ρ	0.724	
Income innovation s.d., σ_y	0.031	
M_1 maturity hazard, λ_1	0.064	
M_2 maturity hazard, λ_2	0.114	
Transitory income shock s.d., σ_m	0.005	
Internally estimated parameters (SMM)		
Discount factor, β	0.9506	
Re-entry probability, θ	0.1118	
M_2 output-penalty cap, d_2	0.9228	
M_2 output-penalty elasticity, d_2^c	0.5790	
Autarky multiplier, d_{aut}	0.9977	
Targeted moments (SMM loss)		
Mean M_1 debt to GDP	16.70%	18.56%
Mean episode length (years)	9.0	9.17
Commit-asymmetry share	0.610	0.716
In-default to non-default new-commitment ratio	0.710	0.733
Mean spread on Market-1 borrowing	1.73%	0.55%
Default frequency (per country-year)	2.60%	0.49%

Notes: SMM weights are the same as for Colombia: 1 on the first four targeted moments and 0.001 on the mean spread and the default frequency. The income process is an AR(1) on linearly detrended log real GDP for Mexico, 1980–2019. The spread target is Mexico's 10-year USD sovereign yield minus the 10-year US Treasury yield, 2000–2017. The default-frequency target is the same all-country hazard rate as for Colombia. *Source:* World Bank International Debt Statistics, IMF World Economic Outlook, Bloomberg, and FRED.

smaller. It is positive at every income state (Figure 13), from +0.12% at the top of the grid to +0.37% near the bottom. The channel is also the same as before. Without M_2, the M_1 default rate goes from 0.49% to 0.37% per year and borrowing capacity from 18.6% to 45.4% of GDP. Why is the gain smaller? Mexico's calibrated $\theta = 0.112$ sits closer to the sign-flip threshold of about 0.17 in Figure 10 than Colombia's 0.0933 does.

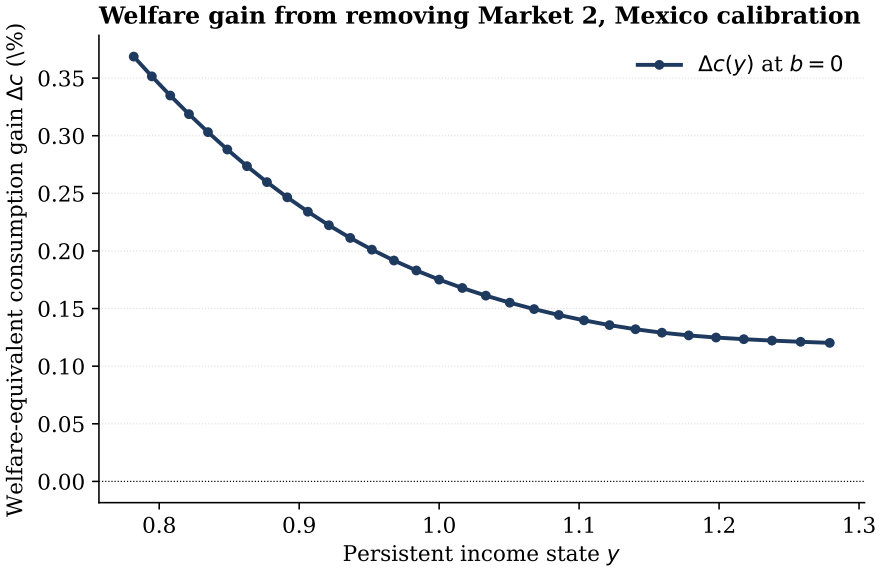


Figure 13: Welfare-equivalent consumption gain $\Delta c(y)$ from removing M_2 at the Mexico calibration, evaluated at $b = 0$ for each of the 31 persistent income states. As for Colombia (Figure 8), the gain is positive at every income state and falls with income.