

Transparency in Debt Crises ^{*}

Lucas Belmudes[†] Angelo Mendes[‡]

May 28, 2026

Abstract

We study a sovereign's choice of whether to disclose information about its repayment capacity to international lenders in a 3-period [Eaton and Gersovitz \(1981\)](#) environment. Under full disclosure, debt is priced at actuarially fair rates for each type; under non-disclosure, a pooling equilibrium generates cross-subsidization across types. We show that non-disclosure is preferred when the probability of being the low type is small: the adverse selection discount the high type bears under pooling is then small, making opacity relatively cheap. Moreover, non-disclosure is preferred when dead-weight losses from default are small: the insurance value of pooling then dominates the pricing gains from transparency. Two pieces of evidence support the framework. Mexico's pre-1995 reserve-disclosure regime illustrates a durable, pre-committed disclosure rule consistent with the model. Using IMF Data Standards Initiatives, we document that no country has ever moved to a lower transparency tier, that countries facing higher recession risk are more often observed to adopt transparency tiers, an association consistent with the model's threshold comparative static, and that early adopters exhibit, for two of the three tier margins, more dispersed residualized issuances than eligible non-adopters, in a direction consistent with the model.

^{*}We thank Manuel Amador, Marco Bassetto, and Tim Kehoe for their guidance during this project. We also thank Ariel Akerman, V.V. Chari, Hasan Çetin, Alessandro Dovis, Jason Hall, Illenin Kondo, Fernando Lopes, Juanpa Nicolini, Guillermo Ordoñez, Chris Phelan, David Rahman, and seminar participants at the University of Minnesota Trade Workshop, the Minnesota-Wisconsin International Macro Workshop, the 23rd annual SAET Conference, and the Spring 2026 Midwest Macro.

[†]belmu002@umn.edu: University of Minnesota.

[‡]mende449@umn.edu: University of Minnesota.

1 Introduction

Lender-borrower relationships are inherently shaped by informational asymmetries. Sovereign debt is no exception: governments typically possess better information about their repayment capacity and policy intentions than their creditors. This information gap creates scope for strategic behavior, as sovereigns decide how much to reveal about their fundamentals to influence market perceptions.

A large body of evidence documents how countries benefit from transparency ([Goldstein et al., 1998](#); [Alesina et al., 1999](#); [Alt and Lassen, 2006](#)). In particular, [Choi and Hashimoto \(2018\)](#) and [Gonzalez-Garcia \(2022\)](#), using event-study and local-projection identification designs, find that countries that adopted IMF data-dissemination standards experience lower sovereign spreads. These pricing gains are consistent with the hypothesis that better information reduces the adverse-selection premium that creditors demand.

Yet opacity persists.¹ Bolivia has suspended publication of international reserve statistics. Greece systematically misreported fiscal data in the years preceding its debt crisis. Several sovereigns have accumulated undisclosed liabilities through bilateral arrangements with China. Mexico withheld detailed information about central bank reserves throughout the early 1990s, even as large external obligations and the peso peg made reserves the key indicator of repayment capacity. If transparency reduces spreads, this persistence of opacity is puzzling. This motivates our question: when should a sovereign that cannot commit to repay disclose information to creditors?

We study this question in a two-period sovereign debt model following [Eaton and Gersovitz \(1981\)](#) with an endogenous choice of disclosure. Following the information-design literature ([Kamenica and Gentzkow, 2011](#)), we assume the government can commit to a disclosure regime at the outset: either full disclosure (FD), under which lenders observe the sovereign’s type before pricing debt, or no disclosure (ND), under which the type remains private. It then chooses how much debt to issue, and whether to default or repay. The sovereign faces a fundamental trade-off between type-specific pricing and insurance. Under full disclosure, debt is sold at actuarially fair prices for each type: the H-type (high-repayment-capacity type) obtains favorable spreads, while the L-type (low-repayment-capacity type) may be constrained in the quantities it can issue. Under no disclosure, a pooling equilibrium forms in which lenders set prices based on the average probability of default. This pool creates an insurance arrangement: the H-type cross-subsidizes the L-type, allowing the L-type to borrow on terms better than full informa-

¹Throughout the paper, we use “opacity,” “no disclosure,” and “non-disclosure” interchangeably.

tion would permit. The cost of opacity is that the L-type may borrow more than it can repay, incurring deadweight losses in default, while the H-type gets an adverse selection discount on its debt price, since it is pooled with L-types. Endogenous disclosure creates a pool over which insurance operates; the extent of that insurance is limited by the deadweight losses incurred when default occurs.

We establish three results. First, under a mild condition on the L-type’s repayment capacity, no separating equilibrium exists under no disclosure: the H-type cannot credibly signal its quality through debt issuance alone (Nachman and Noe, 1994). Second, when the prior probability of the L-type is sufficiently small, the unique high-type-preferred perfect Bayesian equilibrium under no disclosure is a pooling equilibrium in which both types issue the same face value but only the H-type repays. Third, and centrally, no disclosure is preferred when the prior probability of the L-type is sufficiently small: the adverse selection discount the H-type bears under pooling is then limited, making opacity relatively cheap. Moreover, the range of priors over which no disclosure is preferred expands as deadweight losses from default shrink, that is, when default costs are low or the default endowment is high (we use these terms interchangeably throughout). When the cost of default is low, the insurance value of pooling dominates the pricing gains from transparency.

The insurance role of opacity is related to Hirshleifer (1971)’s observation that releasing public information can destroy risk-sharing opportunities. Our contribution is to show how this force operates through defaultable sovereign debt, where the information structure is endogenously chosen, and default is possible. When deadweight losses are small, the L-type’s default is close to what Grossman and Van Huyck (1985) call excusable default, the bad-state realization of an implicit risk-sharing arrangement. In our framework, this arrangement is implemented through the sovereign’s no disclosure policy.

Two robustness exercises show that the main result is not driven by the specific equilibrium selection or information structure. First, we characterize the full set of pooling equilibria under no disclosure and derive a complete welfare ordering: the ex-ante optimal ND equilibrium dominates FD under a strengthened condition, while FD in turn dominates the worst ND equilibrium, safe pooling (Appendix B.1). Second, we re-derive the main result under an alternative information structure in which private information concerns the cost of default rather than the repayment endowment (Appendix B.2). The insurance channel operates analogously in both extensions.

We then provide suggestive evidence in two complementary ways. First, we use Mex-

ico’s 1982–1995 reserve-disclosure history as a case study to illustrate two features consistent with the model’s environment: the durability of a pre-crisis disclosure rule, and its informational consequences. Banco de México released reserve data three times per year throughout the post-1982 modernization period; this infrequent disclosure regime predated the 1994 crisis by more than a decade and was not revised in response to deteriorating fundamentals. The sparse release schedule left investors without timely information about reserve levels during a period of rapidly changing conditions, consistent with the interpretation that investors lacked timely official information about a repayment-relevant state variable. Second, we use the IMF Data Standards Initiatives (SDDS, GDDS, eGDDS+NSDP) to provide cross-country evidence on three dimensions. We document that disclosure regimes are durable and that sovereigns do not reverse stricter standards, supporting the commitment assumption. We construct a WEO-implied proxy for the prior probability of low growth (the probability of being a L-type), from historical IMF forecast errors, and find that countries with a high prior are more likely to adopt a transparency tier, a finding that is statistically significant for SDDS adoption and in the pooled sample, consistent with the model’s comparative statics. Specifically, the model predicts that disclosure is preferred when it is likely for a government to face low repayment capacity in the coming period. Among countries at the frontier of an available transparency tier, early adopters of two out of three margins exhibit more dispersed residualized issuances than countries that remain opaque, consistent with the model’s prediction that transparency reveals type and introduces dispersion in borrowing outcomes.

Finally, we characterize the sovereign’s optimal policy from a constrained disclosure problem (Section 4). Rather than restricting attention to the binary choice between full and no disclosure, we allow the sovereign to commit to any one-sided signal structure indexed by the probability with which the low type is revealed. This is analogous to the sender’s problem in [Kamenica and Gentzkow \(2011\)](#), but adds a constraint specific to our environment: the signal must be informative enough to sustain the risky-pooling equilibrium. We show that the optimal disclosure probability is decreasing in the default endowment. When the default endowment is high, the pooling equilibrium can be sustained under a less informative signal, and the marginal benefit of revealing the low type falls because the revealed type’s safe borrowing capacity shrinks. Both forces push the sovereign toward opacity as default costs decline, extending the baseline result to the full range of disclosure policies.

Related Literature. Our paper is related to three strands of literature. The first of these is the theoretical literature on sovereign debt initiated by [Eaton and Gersovitz \(1981\)](#) and

Cole and Kehoe (2000), among others. Unlike those papers, our model accounts for incomplete information as in Sandleris (2008), Perez (2017), Amador and Phelan (2021), and others. Zabai (2014) and Szkup (2022) also work in incomplete information environments, but differently from the previous papers, they rely on dispersed information to eliminate the equilibrium multiplicity commonly observed under the Cole-Kehoe timing assumption. Bassetto and Galli (2019) study sovereign default in an environment in which the currency choice of debt affects debt prices as it shifts information held by marginal investors. We differ from this literature by introducing an endogenous choice of disclosure in the sovereign’s problem, making lenders’ information (or the lack thereof) an endogenous object in our model.

We are related to the sovereign-debt literature on *hidden debt* and transparency, including Gu and Stangebye (2023), Gamboa (2023), Guler et al. (2022), and Horn et al. (2022). We differ by focusing on a distinct source of information asymmetry: limited transparency about the Mexican central bank’s U.S. dollar reserves (rather than undisclosed liabilities). Among the closest antecedents, Guler et al. (2022) quantifies the macro effects of disclosing non-Paris Club debt in a quantitative model with *exogenous* information regimes, finding modest welfare losses from transparency; Horn et al. (2022) studies hidden debt and transparency with *lender monitoring* that endogenously triggers revelations and affects prices and default. Our framework departs in two key ways. First, the object of opacity is reserves, which is conceptually parallel in our model because hidden debt due imminently maps into our type (repayment capacity). Second, and more importantly, we assume the sovereign can *commit ex ante to a disclosure policy*, giving it a credible channel to communicate with lenders; this contrasts with Horn et al. (2022)’s environment where lenders choose whether to monitor. This shift in the institutional locus of transparency makes the information regime itself a policy choice.

The closest antecedent is Croitorov (2017), who also develops a sovereign-debt model with private information and an ex-ante choice of disclosure. We differ in four ways. First, our comparative static with respect to the default endowment is novel. It connects the idea from Hirshleifer (1971), that transparency can hurt risk-sharing, to Grossman and Van Huyck (1985)’s notion of excusable default: as default costs vanish, the L-type’s default approaches the bad-state realization of an implicit risk-sharing arrangement implemented through no disclosure. Second, we characterize the full set of pooling equilibria and derive a complete welfare ordering (best ND dominates FD, which dominates worst ND, under a strengthened condition), and we derive the mechanism when private information concerns the cost of default rather than the repayment endowment. Third, we

exploit the discrete adoption of IMF Data Standards tiers (SDDS, GDDS, eGDDS+NSDP) as a direct analog of the binary disclosure choice, instead of a continuous transparency index. We document the durability of these regimes across 186 countries in support of the commitment assumption, examine whether the model’s threshold comparative statics are consistent with observed adoption decisions, and compare issuance outcomes across transparent and opaque regimes. Moreover, we focus on emerging-market and developing economies, excluded from [Croitorov \(2017\)](#)’s OECD-only analysis. Fourth, we characterize the optimal policy from a constrained disclosure problem that allows for partial disclosure, showing that the optimal disclosure probability is decreasing in the default endowment; [Croitorov \(2017\)](#) does not consider intermediate disclosure rules.

Second, we contribute to the debate on how information disclosure affects crises. [Gorton and Ordoñez \(2020\)](#) studies how to reveal information *during* an ongoing crisis to manage it in equilibrium. In contrast, in our model, the government chooses the information rule *ex ante*, before a crisis materializes. Crucially, in our framework, the disclosure choice itself can *generate or avert* a crisis. This places us with the pre-crisis design literature that treats disclosure as a commitment device: stress test design in [Goldstein and Leitner \(2018\)](#) and disclosure to prevent runs in [Faria-e Castro et al. \(2017\)](#).² Despite the banking-sovereign differences—and our departure from multiplicity toward a unique pooling outcome under no disclosure—the insurance logic is closely related to [Goldstein and Leitner \(2018\)](#).

Finally, our paper is related to the information design literature: [Kamenica and Gentzkow \(2011\)](#), [Bergemann et al. \(2015\)](#), [Inostroza and Pavan \(2023\)](#).³ More specifically, we focus on a choice of disclosure in environments with adverse selection, as in [Garcia and Tsur \(2021\)](#), [Immorlica et al. \(2022\)](#), and [Dovis and Martellini \(2024\)](#). Among our applications, we use the Mexican Peso Crisis as a case study of a pre-committed sparse disclosure rule.

Outline. The rest of the paper is organized as follows. Section 2 presents the environment and main results, including two robustness extensions (equilibrium set, alternative private information) in Appendix B. Section 3 provides empirical evidence: Section 3.1 uses Mexico’s pre-1995 reserve-disclosure regime as a case study, and Section 3.2 uses IMF Data Standards Initiatives to document the durability of disclosure regimes, examine whether recession-risk priors are consistent with transparency adoption decisions in a manner predicted by the model’s threshold comparative static, and compare issuance

²For a survey, see [Goldstein and Yang \(2017\)](#).

³For a detailed survey, see [Bergemann and Morris \(2019\)](#).

outcomes across transparent and opaque regimes. Section 4 characterizes the optimal disclosure policy in the constrained information-design problem. Section 5 concludes.

2 Model

We study a two-period sovereign-debt environment in which the government privately learns its repayment capacity before issuing debt. Before this information is realized, the government commits to a disclosure regime. Under full disclosure, lenders observe the government's repayment capacity and price debt type by type. Under no disclosure, lenders do not observe the government's type and can only update their beliefs from the debt issuance they observe. We show that no disclosure allows low-repayment-capacity governments to borrow on pooled terms, while high-repayment-capacity governments bear the associated pricing discount. The model therefore isolates the trade-off between transparency, which improves pricing in good states, and opacity, which can provide insurance in bad states through pooling in defaultable debt markets.

Environment. Time is discrete and indexed by $t \in \{0, 1, 2\}$. The economy is populated by a risk-averse sovereign and a continuum of risk-neutral competitive investors. The sovereign derives utility from consumption at $t = 1, 2$ and discounts the future at rate $\beta \in (0, 1)$. The risk-free gross interest rate is constant and equal to R . At $t = 0$, the sovereign's expected payoff is

$$\mathbb{E} [u(c_1) + \beta u(c_2)].$$

Throughout this section, we assume $u(c) = \log(c)$. The expectation is taken with respect to the sovereign's period-2 repayment capacity θ , which can take two values:

$$\theta = \begin{cases} \theta_L, & \text{with probability } \pi, \\ \theta_H, & \text{with probability } 1 - \pi, \end{cases} \quad \theta_H > \theta_L.$$

We refer to the sovereign with repayment capacity θ_H as the H-type (high-repayment-capacity type) and to the sovereign with repayment capacity θ_L as the L-type (low-repayment-capacity type).

At $t = 1$, the sovereign learns its type θ , receives a constant endowment y , and issues one-period debt b at price q . For simplicity, the sovereign enters $t = 1$ with no inherited

debt. First-period consumption is therefore

$$c_1 = y + qb.$$

The sovereign cannot commit at $t = 1$ to repay at $t = 2$. If it repays, it uses its repayment capacity θ to repay debt and consume $\theta - b$. If it defaults, it consumes the default endowment y_d . Hence,

$$c_2 = \delta y_d + (1 - \delta)(\theta - b),$$

where $\delta \in \{0, 1\}$ equals one if the sovereign defaults and zero otherwise.

The object y_d summarizes the value of default. For a given repayment capacity θ , a higher y_d means that default is less costly. Thus, throughout the paper, we use “low default costs,” “small deadweight losses from default,” and “high default endowment” interchangeably.

Information and disclosure. At $t = 0$, before θ is realized, the sovereign commits to a disclosure policy $\phi \in \{0, 1\}$. Throughout the paper, we refer to $\phi = 1$ as full disclosure (FD) and to $\phi = 0$ as no disclosure (ND). At $t = 1$, the sovereign learns θ , the disclosure rule is implemented, and the sovereign chooses debt. Lenders know the prior distribution over types. Under FD, lenders observe θ directly. Under ND, lenders do not observe θ ; they observe only the debt choice b and update beliefs accordingly. Figure 1 displays the timing of the model.

Figure 1: Timing of events

$t = 0$	$t = 1$	$t = 2$
1) Choice of disclosure policy $\phi \in \{0, 1\}$	1) Government learns its type (θ_i) 2) Govt type is revealed ($\phi = 1$) or not ($\phi = 0$) 3) Choice of debt 4) Price is determined q 5) Consumption takes place	1) Default decision 2) Consumption

Modeling assumptions. Two assumptions deserve emphasis. First, the sovereign can commit at $t = 0$ to the disclosure rule that governs information at $t = 1$. This commitment assumption is based on [Kamenica and Gentzkow \(2011\)](#): the information structure

is chosen ex ante and then taken as fixed by lenders.⁴ Second, under FD, the signal truthfully reveals θ . Truthfulness is a property of the disclosure technology, not an equilibrium outcome of a reporting game. The empirical section therefore treats observed disclosure regimes as motivating objects, not as tests of these assumptions.

2.1 Full disclosure

We first characterize the equilibrium under full disclosure, $\phi = 1$. In this regime, the sovereign's type is publicly observed before debt is issued. Lenders therefore price debt separately for each type.

Given type θ , the sovereign solves

$$v_{FD}(\theta; y_d) = \max_{b(\theta)} \{ \log(y + q_{FD}(b, \theta)b) + \beta \log(c_2^*(b, \theta)) \}, \quad (1)$$

where

$$c_2^*(b, \theta) = \max_{\delta \in \{0,1\}} \{ \delta y_d + (1 - \delta)(\theta - b) \}.$$

Definition 1. *An equilibrium under full disclosure is a strategy for the sovereign $(\delta^*(b, \theta), b_{FD}^*(\theta))$ and a price schedule $q_{FD}(b, \theta)$ such that:*

1. *The default policy $\delta^*(b, \theta)$ maximizes second-period consumption.*
2. *Given $q_{FD}(b, \theta)$, the debt choice $b_{FD}^*(\theta)$ maximizes the sovereign's lifetime payoff.*
3. *Investors make zero profits:*

$$\frac{1}{R} [\delta^*(b, \theta) \times 0 + (1 - \delta^*(b, \theta))] - q_{FD}(b, \theta) = 0.$$

The default decision is simple. For a type- θ sovereign, repayment is preferred whenever $\theta - b \geq y_d$. Define the repayment threshold

$$\bar{b}(\theta) \equiv \theta - y_d.$$

Debt is risk-free for type θ up to $\bar{b}(\theta)$ and worthless above that threshold. Hence, the

⁴Appendix A.5 discusses the role of commitment: without it, both types would always send the high-signal message in a cheap-talk equilibrium, eliminating any information transmission. The alternative timing is depicted in Figure 6.

full-disclosure price schedule is

$$q_{FD}(b, \theta) = \begin{cases} 1/R, & \text{if } b \leq \bar{b}(\theta), \\ 0, & \text{if } b > \bar{b}(\theta). \end{cases}$$

We impose the following condition on the L-type.

Assumption 1 (Constrained L-type).

$$y_d > \frac{\beta R}{1 + \beta} \left(\frac{\theta_L}{R} + y \right).$$

Figure 2: Debt and prices in the FD equilibrium under Assumption 1




Figure 2 illustrates the equilibrium debt and price schedule under Assumption 1. Assumption 1 implies that the unconstrained Euler debt choice of the L-type exceeds the maximum debt level that can be repaid without default. Therefore, under FD, the L-type borrows at the repayment threshold $\bar{b}(\theta_L) = \theta_L - y_d$. At this threshold, it is indifferent between repayment and default. We refer to this case by saying that the L-type is constrained.

Lemma 1. Define $\bar{b}(\theta) \equiv \theta - y_d$. Under FD, the equilibrium default policy, price schedule, and debt choice are

$$\delta^*(b, \theta) = \begin{cases} 1, & \text{if } b > \bar{b}(\theta), \\ 0, & \text{if } b \leq \bar{b}(\theta), \end{cases}$$

$$q_{\text{FD}}(b, \theta) = \begin{cases} 0, & \text{if } b > \bar{b}(\theta), \\ 1/R, & \text{if } b \leq \bar{b}(\theta), \end{cases}$$

and

$$b_{\text{FD}}^*(\theta) = \min \left\{ \frac{\theta - \beta R y}{1 + \beta}, \bar{b}(\theta) \right\}.$$

Proof. See Appendix A.1. □

Full disclosure therefore gives the sovereign an ex ante lottery over type-specific borrowing opportunities. When the H-type is realized, the sovereign faces the H-type's repayment threshold and can borrow on risk-free terms up to $\bar{b}(\theta_H)$. When the L-type is realized, the sovereign faces the lower repayment threshold $\bar{b}(\theta_L)$, and Assumption 1 implies that this constraint binds.

The ex ante value of full disclosure is

$$V_{\text{FD}}(\pi; y_d) \equiv \pi v_{\text{FD}}(\theta_L; y_d) + (1 - \pi) v_{\text{FD}}(\theta_H; y_d).$$

Thus, FD generates type-contingent prices and type-contingent borrowing, but it also exposes the L-type to its limited repayment capacity.

2.2 No disclosure

We now consider the no-disclosure regime, $\phi = 0$. In this regime, the sovereign's type is not directly observed by lenders. Lenders observe the debt choice b and form posterior beliefs

$$\mu(b) \equiv \mathbb{P}(\theta = \theta_L \mid b).$$

Given beliefs and prices, a type- θ sovereign solves

$$v_{\text{ND}}(\theta; y_d) = \max_{b(\theta)} \{ \log(y + q_{\text{ND}}(b)b) + \beta \log(c_2^*(b, \theta)) \}, \quad (2)$$

where $c_2^*(b, \theta)$ is defined as above.

Definition 2. A *Perfect Bayesian Equilibrium* under ND is a strategy for the sovereign $(b_{\text{ND}}^*(\theta), \delta^*(b, \theta))$, a price schedule $q_{\text{ND}}(b)$, and beliefs $\mu(b)$ such that:

1. The default policy $\delta^*(b, \theta)$ maximizes second-period consumption.
2. Given $q_{\text{ND}}(b)$, the debt choice maximizes the sovereign's lifetime payoff.

3. Beliefs follow Bayes' rule whenever possible.

4. Investors make zero profits:

$$\frac{1}{R} \mathbb{E}_{\mu(b)} [\delta^*(b, \theta) \times 0 + (1 - \delta^*(b, \theta))] - q_{ND}(b) = 0.$$

The default decision is unchanged relative to the FD case. Information matters at $t = 1$, through debt prices and debt choices. At $t = 2$, after debt has been issued, the sovereign defaults whenever repayment leaves it with less than y_d . Hence the default threshold remains $\bar{b}(\theta) = \theta - y_d$.

Debt issuance can potentially reveal information about the sovereign's type.

Definition 3. A *separating equilibrium* is a PBE in which the two types choose different debt levels:

$$b_{ND}^*(\theta_L) \neq b_{ND}^*(\theta_H).$$

Proposition 1. Under Assumption 1, there is no separating equilibrium under ND.

Proof. See Appendix A.2. □

The logic is that Assumption 1 places the L-type's desired borrowing in the risky region: debt levels that are attractive to the L-type are precisely debt levels at which the L-type defaults, while the H-type can repay. This eliminates the single-crossing logic that would allow debt issuance to separate types. The result is related to the classic pooling result in Nachman and Noe (1994). If Assumption 1 were dropped, a separating equilibrium could arise in which the L-type borrows safely while the H-type separates by choosing a debt level that is costly for itself. Since full disclosure is available as a costless information policy, such costly signaling outcomes are not the central object of the paper.

Definition 4. A *pooling equilibrium* is a PBE in which both types choose the same debt level:

$$b_{ND}^*(\theta_L) = b_{ND}^*(\theta_H) = \tilde{b}.$$

There are two natural pooling candidates. In a safe pooling equilibrium, both types choose the L-type's repayment threshold, $\theta_L - y_d$, and both repay. In a risky pooling equilibrium, both types choose a higher common debt level, $\tilde{b} > \theta_L - y_d$. At this higher debt level, the L-type defaults while the H-type repays. The risky pooling outcome is the one that generates the paper's insurance mechanism: the L-type borrows on pooled

terms and defaults, while the H-type repays and bears the pooling discount. Safe pooling avoids default, but it also limits borrowing for both types.

Define the candidate risky-pooling debt level

$$\tilde{b}(\pi) \equiv \max \left\{ \theta_L - y_d, \min \left\{ \frac{\theta_H(1-\pi) - \beta R y}{(1+\beta)(1-\pi)}, \theta_H - y_d \right\} \right\}.$$

Consider the following belief system, default policy, and price schedule:

$$\mu(b; \pi) = \begin{cases} \pi, & \text{if } b \leq \tilde{b}(\pi), \\ 1, & \text{if } b > \tilde{b}(\pi), \end{cases}$$

$$\delta^*(b, \theta) = \begin{cases} 1, & \text{if } b > \bar{b}(\theta), \\ 0, & \text{if } b \leq \bar{b}(\theta), \end{cases}$$

and

$$q_{ND}(b) = \begin{cases} 0, & \text{if } b > \theta_H - y_d, \\ (1 - \mu(b))/R, & \text{if } b \in (\theta_L - y_d, \theta_H - y_d], \\ 1/R, & \text{if } b \leq \theta_L - y_d. \end{cases}$$

In the risky region, lenders are repaid only when the borrower is the H-type. Hence the pooling price is discounted by the posterior probability of facing the L-type.

Let

$$Y_d \equiv \left(\frac{R\beta}{1+\beta} \left(y + \frac{\theta_L}{R} \right), \frac{R\beta}{1+\beta} \left(y + \frac{\theta_H}{R} \right) \right)$$

denote the set of default endowments for which the L-type is constrained, but the H-type is not, when facing the risk-free price. Assumption 1 is equivalent to y_d exceeding the lower bound of Y_d ; the upper bound of Y_d ensures that the H-type is unconstrained at the risk-free price. Both bounds are imposed in Propositions 2 and 3.

Proposition 2. *Let $y_d \in Y_d$. There exists a cutoff $\pi^*(y_d)$ such that, for $\pi < \pi^*(y_d)$, the H-type-preferred pooling PBE is the risky pooling equilibrium. Moreover,*

$$\frac{\partial \pi^*(y_d)}{\partial y_d} > 0.$$

Proof. See Appendix A.4. □

Appendix A.3 provides an alternative derivation of this result via a mechanism-design

problem in which the H-type is the designer.

The intuition is as follows. In a pooling equilibrium, lenders price debt using the average repayment probability. The L-type defaults in the risky pooling equilibrium, so the H-type repays while receiving a price below the risk-free price. The H-type therefore bears the cost of being pooled with a defaulter. When π is small, however, lenders put little weight on the L-type, and the pooling price is close to $1/R$. The H-type's pooling cost is then small. At the same time, risky pooling permits a higher common debt level than safe pooling. For sufficiently small π , the H-type is willing to pool with the L-type in order to borrow more at a price that is only mildly discounted.

A higher y_d means lower default costs. This lowers the safe debt threshold $\theta_L - y_d$, making the safe pooling outcome more restrictive, reducing its payoff. Thus, the range of priors for which risky pooling is the H-type-preferred pooling equilibrium expands with y_d .

In the risky pooling region, the ex ante value of no disclosure is

$$V_{ND}(\pi; y_d) \equiv u\left(y + \frac{1-\pi}{R}\tilde{b}(\pi)\right) + \beta [\pi u(y_d) + (1-\pi)u(\theta_H - \tilde{b}(\pi))].$$

This expression highlights the insurance embedded in ND. Both types obtain the same first-period resources from pooled borrowing. In the second period, the L-type defaults and consumes y_d , while the H-type repays and consumes $\theta_H - \tilde{b}(\pi)$. The arrangement therefore transfers borrowing capacity toward the low-repayment-capacity state by pooling it with the H-type.

The main text focuses on this risky pooling equilibrium because it is the outcome that generates the paper's mechanism. Appendix B.1 characterizes the broader set of pooling equilibria, including safe pooling, and derives welfare bounds across the equilibrium set.

2.3 Disclosure choice

The sovereign chooses the disclosure rule ex ante, before learning its type. The relevant comparison is therefore between the ex ante value of ND and the ex ante value of FD. Define

$$\Delta(\pi; y_d) \equiv V_{ND}(\pi; y_d) - V_{FD}(\pi; y_d).$$

No disclosure is preferred whenever $\Delta(\pi; y_d) > 0$, while full disclosure is preferred whenever $\Delta(\pi; y_d) < 0$.

Before stating the main result, impose the following condition:

$$\log \left(\frac{Ry + \theta_H}{Ry + \theta_L} \right) + \frac{R\beta y - \theta_H}{Ry + \theta_H} > 0. \quad (\star)$$

Condition $(\star)$ is a sufficient condition ensuring that ND is valuable when the probability of the low state is small. At $\pi = 0$, the L-type never occurs, so FD and ND coincide. The relevant object is therefore the derivative of $\Delta(\pi)$ at zero. Condition $(\star)$ guarantees that introducing a small probability of the low-repayment-capacity state raises the value of ND relative to FD.

Economically, the condition requires the bad state to be sufficiently severe. One simple way for it to hold is for the endowment gap $\theta_H - \theta_L$ to be large. When the L-type's repayment capacity is much lower than the H-type's, FD exposes the L-type to a tight borrowing constraint. Pooling then provides valuable insurance by allowing the L-type to borrow against the average repayment probability rather than only against its own repayment capacity.

Proposition 3. *Suppose condition $(\star)$ holds. Then, for any $y_d \in Y_d$, there exists a cutoff $\bar{\pi}(y_d) > 0$ such that*

$$\Delta(\pi; y_d) > 0 \quad \text{for all } \pi \in (0, \bar{\pi}(y_d)).$$

Moreover, $\Delta(\pi; y_d)$ is strictly concave in π . Hence the set of priors for which no disclosure dominates full disclosure is an interval. If $\bar{\pi}(y_d)$ is an interior cutoff, then

$$\frac{d\bar{\pi}(y_d)}{dy_d} > 0.$$

Proof. Fix $y_d \in Y_d$; write $\Delta(\pi) \equiv \Delta(\pi; y_d)$ for brevity throughout this proof. We start by rewriting $\Delta(\pi)$:

$$\begin{aligned} \Delta(\pi) = & \log \left(\frac{y + \frac{1-\pi}{R}\theta_H}{1 + \beta} \right) + \beta(1 - \pi) \log \left[\frac{R\beta}{1 + \beta} \left(\frac{y}{1 - \pi} + \frac{\theta_H}{R} \right) \right] \\ & - \pi \log \left(y + \frac{\theta_L - y_d}{R} \right) \\ & - (1 - \pi) \left[\log \left(\frac{y + \frac{\theta_H}{R}}{1 + \beta} \right) + \beta \log \left(\frac{R\beta}{1 + \beta} \left(y + \frac{\theta_H}{R} \right) \right) \right]. \end{aligned}$$

Differentiating term by term,

$$\frac{\partial}{\partial \pi} \log \left(\frac{y + \frac{1-\pi}{R} \theta_H}{1 + \beta} \right) = \frac{-\frac{\theta_H}{R}}{y + \frac{1-\pi}{R} \theta_H} = -\frac{\theta_H}{Ry + (1-\pi)\theta_H}.$$

For the second term,

$$\begin{aligned} & \frac{\partial}{\partial \pi} \left\{ \beta(1-\pi) \log \left[\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right) \right] \right\} \\ &= -\beta \log \left[\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right) \right] \\ & \quad + \beta(1-\pi) \frac{\frac{R\beta}{1+\beta} \frac{y}{(1-\pi)^2}}{\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right)}. \end{aligned}$$

Canceling the common factor $\frac{R\beta}{1+\beta}$,

$$\begin{aligned} &= -\beta \log \left[\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right) \right] + \beta(1-\pi) \frac{\frac{y}{(1-\pi)^2}}{\frac{y}{1-\pi} + \frac{\theta_H}{R}} \\ &= -\beta \log \left[\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right) \right] + \beta \frac{\frac{y}{1-\pi}}{\frac{y}{1-\pi} + \frac{\theta_H}{R}} \\ &= -\beta \log \left[\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right) \right] + \frac{\beta Ry}{Ry + (1-\pi)\theta_H}. \end{aligned}$$

For the third term,

$$\frac{\partial}{\partial \pi} \left[-\pi \log \left(y + \frac{\theta_L - y_d}{R} \right) \right] = -\log \left(y + \frac{\theta_L - y_d}{R} \right).$$

For the fourth term,

$$\begin{aligned} & \frac{\partial}{\partial \pi} \left\{ -(1-\pi) \left[\log \left(\frac{y + \frac{\theta_H}{R}}{1 + \beta} \right) + \beta \log \left(\frac{R\beta}{1+\beta} \left(y + \frac{\theta_H}{R} \right) \right) \right] \right\} \\ &= \log \left(\frac{y + \frac{\theta_H}{R}}{1 + \beta} \right) + \beta \log \left(\frac{R\beta}{1+\beta} \left(y + \frac{\theta_H}{R} \right) \right). \end{aligned}$$

Therefore,

$$\begin{aligned}
\Delta_\pi(\pi) = & -\frac{\theta_H}{Ry + (1-\pi)\theta_H} + \frac{\beta Ry}{Ry + (1-\pi)\theta_H} \\
& - \beta \log \left[\frac{R\beta}{1+\beta} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right) \right] \\
& - \log \left(y + \frac{\theta_L - y_d}{R} \right) \\
& + \log \left(\frac{y + \frac{\theta_H}{R}}{1+\beta} \right) + \beta \log \left(\frac{R\beta}{1+\beta} \left(y + \frac{\theta_H}{R} \right) \right).
\end{aligned}$$

Combining terms,

$$\begin{aligned}
\Delta_\pi(\pi) = & \frac{\beta Ry - \theta_H}{Ry + (1-\pi)\theta_H} \\
& + \log \left(\frac{y + \frac{\theta_H}{R}}{(1+\beta) \left(y + \frac{\theta_L - y_d}{R} \right)} \right) \\
& + \beta \log \left(\frac{y + \frac{\theta_H}{R}}{\frac{y}{1-\pi} + \frac{\theta_H}{R}} \right).
\end{aligned}$$

Evaluating at $\pi = 0$,

$$\Delta_\pi(0) = \frac{\beta Ry - \theta_H}{Ry + \theta_H} + \log \left(\frac{y + \frac{\theta_H}{R}}{(1+\beta) \left(y + \frac{\theta_L - y_d}{R} \right)} \right).$$

Multiplying numerator and denominator inside the logarithm by R ,

$$\boxed{\Delta_\pi(0) = \frac{\beta Ry - \theta_H}{Ry + \theta_H} + \log \left(\frac{Ry + \theta_H}{(1+\beta)(Ry + \theta_L - y_d)} \right).}$$

Notice that $\Delta(0) = 0$, and that $\Delta'(0)$ depends on y_d only through $-\log(Ry + \theta_L - Ry_d)$, which makes $\Delta_\pi(0, y_d)$ increasing in y_d . This means that if $\Delta_\pi(0) > 0$ for the minimum admissible value of $y_d \in Y_d$ it must be greater than zero for all admissible values of y_d . Recall that the minimum admissible value of y_d , denoted by $y_d^{\min}$ is:

$$y_d^{\min} = \frac{R\beta}{1+\beta} \left(y + \frac{\theta_L}{R} \right)$$

$$(1 + \beta)(Ry + \theta_L - y_d^{\min}) = (1 + \beta) \left(Ry + \theta_L - \frac{\beta(Ry + \theta_L)}{1 + \beta} \right) = Ry + \theta_L$$

We then get condition ($\star$):

$$\Delta_\pi(0; y_d^{\min}) = \log \left(\frac{Ry + \theta_H}{Ry + \theta_L} \right) + \frac{R\beta y - \theta_H}{Ry + \theta_H} > 0$$

It remains to show that $\Delta(\pi; y_d)$ is globally concave in π . Differentiate Δ_π term by term.

For the first term,

$$\frac{\partial}{\partial \pi} \left[\frac{\beta Ry - \theta_H}{Ry + (1 - \pi)\theta_H} \right] = \frac{\theta_H(\beta Ry - \theta_H)}{[Ry + (1 - \pi)\theta_H]^2}.$$

For the last term,

$$\begin{aligned} \frac{\partial}{\partial \pi} \left[\beta \log \left(\frac{y + \frac{\theta_H}{R}}{\frac{y}{1-\pi} + \frac{\theta_H}{R}} \right) \right] &= -\beta \frac{\frac{\partial}{\partial \pi} \left(\frac{y}{1-\pi} + \frac{\theta_H}{R} \right)}{\frac{y}{1-\pi} + \frac{\theta_H}{R}} \\ &= -\beta \frac{\frac{y}{(1-\pi)^2}}{\frac{y}{1-\pi} + \frac{\theta_H}{R}}. \end{aligned}$$

Since

$$\frac{y}{1-\pi} + \frac{\theta_H}{R} = \frac{Ry + (1-\pi)\theta_H}{R(1-\pi)},$$

we get

$$-\beta \frac{\frac{y}{(1-\pi)^2}}{\frac{y}{1-\pi} + \frac{\theta_H}{R}} = -\frac{\beta Ry}{(1-\pi)[Ry + (1-\pi)\theta_H]}.$$

Therefore,

$$\Delta_{\pi\pi}(\pi) = \frac{\theta_H(\beta Ry - \theta_H)}{[Ry + (1-\pi)\theta_H]^2} - \frac{\beta Ry}{(1-\pi)[Ry + (1-\pi)\theta_H]}.$$

Putting both terms over the common denominator,

$$\Delta_{\pi\pi}(\pi) = \frac{(1-\pi)\theta_H(\beta Ry - \theta_H) - \beta Ry [Ry + (1-\pi)\theta_H]}{(1-\pi)[Ry + (1-\pi)\theta_H]^2}.$$

Expanding the numerator,

$$\begin{aligned} &(1-\pi)\theta_H(\beta Ry - \theta_H) - \beta Ry [Ry + (1-\pi)\theta_H] \\ &= \beta Ry\theta_H(1-\pi) - (1-\pi)\theta_H^2 - \beta(Ry)^2 - \beta Ry\theta_H(1-\pi). \end{aligned}$$

The cross terms cancel:

$$\beta R y \theta_H (1 - \pi) - \beta R y \theta_H (1 - \pi) = 0.$$

Hence,

$$\Delta_{\pi\pi}(\pi) = \frac{-(1 - \pi)\theta_H^2 - \beta(Ry)^2}{(1 - \pi)[Ry + (1 - \pi)\theta_H]^2}.$$

Equivalently,

$$\Delta_{\pi\pi}(\pi) = -\frac{(1 - \pi)\theta_H^2 + \beta(Ry)^2}{(1 - \pi)[Ry + (1 - \pi)\theta_H]^2} < 0.$$

Thus $\Delta(\pi; y_d)$ is strictly concave in π . Since the previous part of the proof establishes

$$\Delta(0; y_d) = 0 \quad \text{and} \quad \Delta_{\pi}(0; y_d) > 0,$$

Δ initially rises above zero and, by strict concavity, can cross zero at most once. Hence the set of priors for which no disclosure dominates full disclosure is an interval of the form

$$(0, \bar{\pi}(y_d)).$$

Finally, using the expression for $\Delta(\pi; y_d)$,

$$\Delta_{y_d}(\pi; y_d) = \frac{\pi}{Ry + \theta_L - y_d} > 0 \quad \text{for every } \pi > 0.$$

If $\bar{\pi}(y_d)$ is an interior cutoff, then

$$\Delta(\bar{\pi}(y_d); y_d) = 0.$$

By the implicit function theorem,

$$\frac{d\bar{\pi}(y_d)}{dy_d} = -\frac{\Delta_{y_d}(\bar{\pi}(y_d); y_d)}{\Delta_{\pi}(\bar{\pi}(y_d); y_d)}.$$

At the positive root of a strictly concave function that starts at zero with positive slope,

$$\Delta_{\pi}(\bar{\pi}(y_d); y_d) < 0.$$

Since

$$\Delta_{y_d}(\bar{\pi}(y_d); y_d) > 0,$$

we obtain

$$\frac{d\bar{\pi}(y_d)}{dy_d} > 0.$$

Therefore, a higher default endowment y_d , equivalently lower default costs, expands the interval of priors for which no disclosure is preferred. $\square$


The result formalizes the insurance value of opacity. Under FD, the H-type borrows on favorable terms, but the L-type is constrained by its own repayment capacity. Under ND, both types borrow on pooled terms. The L-type obtains more first-period resources than it could under FD and then defaults in the second period. The H-type repays, but faces a price discount because lenders account for the possibility of default by the L-type.

When π is small, the H-type's cost from pooling is limited because the pooling price remains close to the risk-free price. But the L-type's gain can be large because opacity relaxes a severe borrowing constraint. If default costs are low, so that y_d is high, the second-period cost of default in the low state is limited. In this case, the insurance value of ND can dominate the pricing gains from transparency.

The mechanism is therefore not that less information is always better. Full disclosure improves pricing conditional on being the H-type. The point is instead that, in a sovereign debt environment with limited commitment, opacity can act as a second-best insurance device. It reallocates market access toward the low-repayment-capacity state through adverse-selection pricing and endogenous default. Figure 3 illustrates the welfare differential $\Delta(\pi; y_d) = V_{ND}(\pi; y_d) - V_{FD}(\pi; y_d)$ as a function of π : the differential is positive for small π and eventually reverses sign, consistent with the local result in Proposition 3.

Robustness and extensions. Two robustness exercises in Appendix B support the main result. First, we characterize the full set of pooling equilibria under ND and derive a complete welfare ordering: the ex-ante optimal ND equilibrium dominates FD under a strengthened version of condition ($\star$), while FD in turn dominates the worst ND equilibrium (safe pooling) (Appendix B.1). Second, we re-derive the mechanism under an alternative information structure in which private information concerns default costs rather than repayment endowments (Appendix B.2). In both cases, the insurance channel operates through the same sovereign-debt-specific forces: pooling prices, endogenous borrowing, and default incentives. Third, Appendix A.6 derives a knife-edge parametric case

Figure 3: Difference between expected payoffs under ND and FD



Note: Parameters used: $\beta = 0.92$, $R = 1$, $y = 1$, $\theta_L = 1.25$, and $\theta_H = 5$.

in which the main result admits closed-form analytical verification.

3 Empirical Evidence

We now present descriptive evidence consistent with the model's environment. Section 3.1 uses Mexico's pre-1995 reserve-disclosure regime to document a pre-existing, sparse disclosure rule and its information consequences during the 1994 crisis. Section 3.2 uses IMF Data Standards Initiatives to document that distinct, durable disclosure regimes exist across countries, that sovereigns facing higher recession risk are more likely to adopt transparency in a manner consistent with the model's threshold prediction, and that the dispersion of issuances differs between transparent and opaque regimes in a direction consistent with the model.

3.1 Mexico as a case study

We use Mexico's pre-1995 reserve-disclosure regime as a case study of a predetermined and infrequent disclosure rule. The purpose of the case is to show that, before the 1994 crisis, Mexico followed a long-standing schedule for releasing reserve information, and that this schedule left investors without current information about a repayment-relevant state variable during a period of rapidly changing fundamentals.

Our model assumes a disclosure rule chosen before the repayment state is realized. Mexico's pre-crisis reserve-disclosure regime is consistent with this interpretation. Banco de

México released reserve data three times per year: the March Annual Report, the October Bankers' Convention speech, and the November President's Report. This rule was in place throughout the post-1982 modernization period, not adopted in response to the 1994 deterioration. After the 1982 debt crisis, Mexico undertook a profound economic transformation: fiscal consolidation reduced public expenditure from 45% of GDP in 1983 to 27% in 1993; more than 75% of state-owned enterprises were privatized; and trade liberalization eliminated tariffs on over 85% of imports by 1985. These reforms attracted large foreign capital inflows in the early 1990s. Through this entire period, fiscal consolidation, trade liberalization, privatization, and large capital inflows, the disclosure rule was not revisited.

Reserves in this period were a repayment-relevant state variable. They served both to service dollar debt and to defend the peso, so their level directly signaled repayment capacity. Across multiple vintages of the World Bank's *Global Development Finance* and *The World Debt Tables*, reserve levels were reported consistently before, during, and after the crisis.⁵

In 1994, a sequence of political and social shocks undermined investor confidence: the EZLN uprising, the assassination of presidential candidate Luis Donaldo Colosio, the resignation of the interior minister, and the later assassination of José Francisco Ruiz Massieu. Despite this turbulent backdrop, the government did not revise its reserve-disclosure schedule. The left panel of Figure 4 shows reserves relative to short-term dollar debt throughout 1994, with vertical dashed lines at the three official release dates, illustrating how much information accumulated between scheduled disclosures. The right panel compares contemporaneous press statements with the values later disclosed: points referring to past reserve levels lie close to the 45-degree line, whereas current statements, or nowcasts, lie below it, revealing systematic overestimation. Taken together, these patterns illustrate a sparse and delayed reporting regime, suggesting that official reserve releases left a substantial information gap; whether investors relied primarily on these releases cannot be established from this evidence alone.

3.2 IMF Transparency Programs

In direct response to the Mexican crisis, the IMF established the Special Data Dissemination Standard (SDDS) in 1996 for countries with capital-market access, followed by

⁵This vintage stability is only a descriptive fact about the reported reserve series. It does not establish that truthful reporting was incentive compatible. We treat truthfulness as a property of the disclosure technology in the model.




Figure 4: Panel A (left) shows the ratio of foreign exchange reserves to short-term dollar debt in 1994. Vertical dashed lines mark the dates of official reserve releases: the Annual Report (March), the Bankers’ Convention (October), and the President’s Report (November). Panel B (right) compares contemporaneous statements about reserve levels with the values later disclosed; deviations from the 45-degree line indicate systematic overestimation before 1995.

the General Data Dissemination System (GDDS) in 2000 for countries developing their statistical systems. The GDDS was later replaced by the enhanced GDDS (eGDDS) in 2015, which introduced National Summary Data Pages (NSDPs) as a dissemination vehicle. An additional tier, SDDS Plus, was introduced in 2012 for systemically important financial sectors, but is used mostly by developed countries and is excluded from the analysis below. These standards define codified disclosure regimes with documented adoption dates, and can be interpreted as the cross-country analogue of the model’s disclosure choice $\phi \in \{0, 1\}$. [Choi and Hashimoto \(2018\)](#) and [Gonzalez-Garcia \(2022\)](#), using event-study and local-projection identification, establish the causal pricing benefit of transparency; our cross-country analysis is descriptive and complementary to those designs. Figure 5 shows the history of adoptions across tiers, with the initial wave of SDDS subscriptions in the late 1990s following in the wake of the Mexican crisis.

Disclosure regimes are durable. For the commitment assumption to be empirically meaningful, adopted disclosure regimes should be stable over time. We assess this using the IMF’s Dissemination Standards Bulletin Board (DSBB), which records exact adoption dates for each tier. For each transition margin, we define an eligibility clock that starts at

$$t_{ir}^{\text{elig}} = \max\{\text{date country } i \text{ enters the origin tier}, T_r\},$$




Figure 5: Countries participating in the IMF Data Standards Initiatives (GDDS, e-GDDS, SDDS, SDDS Plus) by year. The initial wave of SDDS subscriptions followed in the wake of the 1994–95 Mexican crisis.

where T_r is the date the destination tier became available. A country-spell contributes to margin r only once both conditions hold simultaneously: the country is in the origin tier, *and* the destination tier exists. The adoption probability is the fraction of eligible countries that adopted the specific destination tier; spells still ongoing at the time of writing are included in waiting-time averages at their current duration.

Table 1 summarises the results for all six margins. Among 186 countries in the sample (listed in Appendix C.5), *no country has ever moved to a lower tier*. Four of the six margins exhibit positive adoption rates (43.0%, 66.5%, 69.6%, and 40.0% for No prior recorded tier $\rightarrow$ SDDS, No regime $\rightarrow$ GDDS/e-GDDS, e-GDDS $\rightarrow$ e-GDDS+NSDP, and SDDS $\rightarrow$ SDDS Plus, respectively), while the remaining two (No regime $\rightarrow$ e-GDDS+NSDP and e-GDDS+NSDP $\rightarrow$ SDDS) show zero adoptions. Median waiting times before upgrading are long, ranging from 4.6 to 8.9 years across the three upward margins to standard transparency tiers (excluding SDDS Plus, where the median is 12.8 years). We interpret this durability as consistent with the commitment assumption in Section 2: sovereigns treat disclosure regimes as rule-like choices rather than period-by-period decisions. Crucially,

this evidence rests on the positive fact of non-reversal across many countries, rather than on the absence of disclosure during a single deterioration episode.

Table 1: Adoption probabilities and waiting times by tier frontier

Transition margin	N_{elig}	N_{adopted}	$P(\text{adopt})$	Wait (years)		
				Mean	Median	Std. dev.
No prior recorded tier $\rightarrow$ SDDS [†]	186	80	0.430	11.8	8.9	8.0
No regime $\rightarrow$ GDDS/e-GDDS	164	109	0.665	9.3	6.3	7.3
No regime $\rightarrow$ e-GDDS+NSDP	48	0	0.000	4.1	4.0	3.2
e-GDDS $\rightarrow$ e-GDDS+NSDP	102	71	0.696	6.4	4.6	3.9
e-GDDS+NSDP $\rightarrow$ SDDS	75	0	0.000	7.1	7.4	2.2
SDDS $\rightarrow$ SDDS Plus	80	32	0.400	10.4	12.8	4.5

Notes: Data from the DSBB master file. N_{elig} counts distinct countries whose eligibility clock was active at some point after the destination tier became available. N_{adopted} counts countries that moved to the destination tier. $P(\text{adopt}) = N_{\text{adopted}}/N_{\text{elig}}$. Waiting years measures time in the origin tier after the destination tier became available; ongoing spells are included at current duration. [†]The DSBB source is a cross-sectional snapshot in which each country appears only under its current highest tier. Former e-GDDS members that subsequently subscribed to SDDS are no longer listed in the e-GDDS history table, so the e-GDDS $\rightarrow$ SDDS path is unobservable from this source; those cases are absorbed into the first row.

3.2.1 WEO-Implied Recession Prior and Transparency Adoption

The model's key comparative static is that when being an L-type is unlikely ($\pi < \bar{\pi}$), the sovereign prefers to remain opaque. On the other hand, for high probabilities π , the sovereign prefers to be transparent. We test whether this prediction is consistent with observed transparency choices by constructing an empirical proxy for π from IMF World Economic Outlook (WEO) forecasts. Let $\hat{g}_{i,t,t+1}^{\text{WEO}}$ denote the one-year-ahead WEO growth forecast for country i , constructed from the fall (October) vintage of year t and targeting year $t + 1$. We retrieve these forecasts from the IMF's historical WEO archive (series NGDP_RPCH) and compute realized forecast errors $\varepsilon_{i,t+1} = g_{i,t+1}^{\text{realized}} - \hat{g}_{i,t,t+1}^{\text{WEO}}$. The proxy is

$$\pi_{it}^{\text{WEO}} = \Phi\left(\frac{\bar{g} - \hat{g}_{i,t,t+1}^{\text{WEO}} - \hat{\mu}_{\varepsilon}}{\hat{\sigma}_{\varepsilon}}\right), \quad (3)$$

where $\hat{\mu}_{\varepsilon}$ is the pooled mean, and $\hat{\sigma}_{\varepsilon}$ the standard deviation across all emerging market and developing countries in the panel. Φ is the standard normal CDF and $\bar{g} = 1$ pp is the low-growth threshold below which we define the bad state θ_L . This proxy is analogous to the model prior $\pi = \Pr(\theta = \theta_L)$: a higher π_{it}^{WEO} means the sovereign is more likely to face low growth in the coming year.

To avoid using information realized after the transparency decision, the regressor entering adoption equations in year t is the predetermined prior

$$\pi_{it}^{\text{pre}} = \Phi\left(\frac{\bar{g} - \hat{g}_{i,t-1,t}^{\text{WEO}} - \hat{\mu}_{\varepsilon}}{\hat{\sigma}_{\varepsilon}}\right), \quad (4)$$

that is, the WEO probability built in fall $t - 1$ for growth in year t . The WEO proxy captures publicly available growth expectations; whether this aligns with the sovereign's private repayment-capacity risk depends on the extent to which growth shortfalls translate into reduced debt-service capacity.

The model predicts that adoption of a transparency tier is more likely when π exceeds a threshold $\bar{\pi}$: agents who face a high probability of being an L-type tomorrow prefer disclosure because the adverse selection discount is too costly. We operationalize this with the indicator $\mathbf{1}\{\pi_{it}^{\text{pre}} > \hat{\pi}\}$, setting $\hat{\pi} = 0.15$. The event variable Adopt_{it} equals one in the year a country first subscribes to a given tier. The feasible choice set for each margin consists of country-years in which the destination tier was available and the country had not yet adopted; adoption events are appended to this set. Country-years during default or restructuring spells (identified from [Asonuma and Trebesch 2016](#)) are excluded throughout. We estimate a linear probability model (LPM),

$$\text{Adopt}_{it} = \lambda_t + \gamma \mathbf{1}\{\pi_{it}^{\text{pre}} > \hat{\pi}\} + \mathbf{X}_{it}' \Gamma + u_{it}, \quad (5)$$

with year fixed effects λ_t , and controls $\mathbf{X}_{it}$ comprising log GDP, real GDP growth, external debt to GDP, the ratio of reserves to external debt, and tax revenue to GDP, and a logit specification,

$$\Pr(\text{Adopt}_{it} = 1) = \Lambda(\lambda_t + \gamma \mathbf{1}\{\pi_{it}^{\text{pre}} > \hat{\pi}\} + \mathbf{X}_{it}' \Gamma), \quad (6)$$

where Λ is the logistic CDF; we report log-odds coefficients. The model prediction is $\gamma > 0$ in both equations.

Table 2 reports results for each tier frontier separately and pooled. The feasible choice set for each margin is restricted to the IDS panel of 92 low- and middle-income economies; high-income SDDS-Plus subscribers are excluded at the data stage. Appendix C.6 lists all countries in each margin's sample. The coefficient on the threshold indicator is positive across all four samples in both specifications. The per-margin estimates are imprecisely estimated given the small number of events in each tier, but are statistically significant for SDDS and the pooled sample with all events. The threshold $\hat{\pi} = 0.15$ is not derived from the model's structural parameters; Appendix C.1 replicates the regressions

for $\hat{\pi} \in \{0.15, 0.17, 0.20, 0.35\}$ and shows that the positive and significant results for SDDS and the pooled sample hold at the benchmark $\hat{\pi} = 0.15$, with evidence weakening at higher thresholds where the indicator variable loses variation. Appendix C.2 replaces the indicator with the continuous probability π_{it}^{pre} ; the resulting pattern positive but imprecise for SDDS, near-zero or negative for the smaller per-margin samples is consistent with a threshold-shaped rather than monotonically increasing relationship, in line with the model’s comparative static.

Table 2: Adoption regressions: $\mathbf{1}\{\pi_{it}^{pre} > \hat{\pi}\}$ and transparency tier adoption

Sample	LPM (5)	Logit (6)	N	Events
SDDS	0.027** (0.013)	1.272 (1.126)	929	24
GDDS	0.028 (0.075)	0.713 (1.132)	207	24
eGDDS+NSDP	0.001 (0.087)	0.607 (0.875)	163	27
Pooled	0.040** (0.018)	1.087* (0.557)	1,299	75

Notes: Dependent variable: $\text{Adopt}_{it} = 1$ in the year of first subscription to the relevant tier. Regressor: $\mathbf{1}\{\pi_{it}^{pre} > \hat{\pi}\}$, with $\hat{\pi} = 0.15$, where π_{it}^{pre} is the WEO-implied probability of real GDP growth below $\bar{g} = 1$ pp, constructed from the fall WEO vintage in year $t - 1$. LPM includes year fixed effects with HC1 standard errors (parentheses); Logit reports log-odds coefficients with MLE standard errors. Both include controls for log GDP, real GDP growth, external debt-to-GDP, reserves-to-external-debt, and tax revenue-to-GDP. Default and restructuring spells excluded (Asonuma and Trebesch 2016). The pooled sample stacks all three tier frontiers; the per-margin samples are restricted to the feasible choice set for that tier. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

3.2.2 Residualized Issuances and the Transparent-Opaque Split

The model predicts that, at the moment a new transparency tier becomes available, early adopters reveal their type and price accordingly, so their post-adoption issuance distribution should be more dispersed than in countries that remain pooled. We use the World Bank International Debt Statistics (IDS), series PPG_Bonds_DIS: annual disbursements of public and publicly-guaranteed bond debt, scaled by contemporaneous GDP. The panel covers 93 low- and middle-income economies. We retain country-years with at least one positive issuance anywhere in the panel (country-years with zero issuance are kept; countries that *never* issued are dropped). We use the dataset from Asonuma and Trebesch (2016) to exclude countries in default or debt restructuring. The outcome is

$y_{it} = \log(1 + \text{Issuance}_{it}/\text{GDP}_{it})$. We project it on country fixed effects, year fixed effects, and macro controls:

$$y_{it} = \alpha_i + \lambda_t + \beta_1 \log \text{GDP}_{it} + \beta_2 \Delta \ln \text{GDP}_{it} + \beta_3 \frac{\text{ExtDebt}_{it}}{\text{GDP}_{it}} + \beta_4 \frac{\text{Reserves}_{it}}{\text{GDP}_{it}} + u_{it}. \quad (7)$$

The residual $\hat{u}_{it}$ captures issuance variation not explained by size, growth, leverage, or reserve buffers. For each frontier tier r with launch year T_r (SDDS: 1996, GDDS: 2000, eGDDS+NSDP: 2015), let $\mathcal{C}_r$ be the cohort of countries whose eligibility clock is active at T_r . For each tier we define a five-year observation window $[T_r, T_r + 5]$ and classify every country-year (i, t) with $i \in \mathcal{C}_r$ and $t \in [T_r, T_r + 5]$ as:

$$\begin{aligned} \text{Transparent}_{it}^r &= \mathbf{1}\{i \in \mathcal{C}_r, t \in [T_r, T_r + 5], \text{ and } i \text{ has adopted tier } r \text{ by year } t\}, \\ \text{Opaque}_{it}^r &= \mathbf{1}\{i \in \mathcal{C}_r, t \in [T_r, T_r + 5], \text{ and } i \text{ has not yet adopted tier } r \text{ by year } t\}. \end{aligned}$$

A country that adopts mid-window therefore contributes pre-adoption years to the Opaque distribution and post-adoption years to the Transparent distribution. SDDS Plus is excluded as it is used almost entirely by high-income countries. The headline statistic is $\text{Var}(\hat{u} \mid \text{Transparent})/\text{Var}(\hat{u} \mid \text{Opaque})$; the model predicts this ratio exceeds one.

Example. Consider the GDDS tier, launched in 2000 ($T_r = 2000$, window $[2000, 2005]$). A country eligible at 2000 that subscribes to GDDS in 2003 contributes its years 2000–2002 to the Opaque distribution and its years 2003–2005 to the Transparent distribution (3 years each). A country that never adopts within the window contributes all six years to Opaque. The Transparent distribution therefore consists exclusively of post-adoption observations, while the Opaque distribution pools never-adopters with pre-adoption years of eventual adopters.

Table 3 reports the variance ratio for each tier and the pooled sample. Two of the three tier margins show ratios consistent with the model’s prediction: GDDS (3.75) and eGDDS+NSDP (1.24). The SDDS margin is an exception (ratio 0.54): a possible explanation is that SDDS early adopters in 1996–2000 were predominantly larger, more creditworthy economies whose repayment capacity was already well known to markets, so formal SDDS adoption may have conveyed little new information. The GDDS estimate rests on only $N_T = 14$ transparent country-years; small cell sizes limit inference for that tier in isolation. Pooling all three frontier episodes, the variance ratio is 1.28. These variance ratios are descriptive and consistent with the model’s type-revelation prediction, but higher variance among transparent country-years could also reflect pre-existing heterogeneity, regime-change-

induced market scrutiny, or selection of more volatile issuers into early adoption; residualization partially mitigates, but does not eliminate, these concerns. Appendix C.3 reports the kernel density plots of the residual distributions. Appendix C.4 reports a robustness check using Dealogic quarterly bond issuance data (81 countries); the main pattern is preserved.

Table 3: Variance ratio of residualized issuances: Transparent vs. Opaque country-years

Transparency tier	N_T	N_O	$\text{Var}(\hat{u} T)/\text{Var}(\hat{u} O)$	Consistent with prediction?
SDDS (1996)	60	72	0.54	No
GDDS (2000)	14	61	3.75	Yes
eGDDS+NSDP (2015)	57	75	1.24	Yes
Pooled	181	343	1.28	Yes

Notes: IDS annual panel, 93 low- and middle-income countries. N_T (N_O) counts country-year observations classified as Transparent (Opaque) in the regression sample. Residuals $\hat{u}_{it}$ from equation (7). SDDS Plus excluded. Countries with no positive issuance in the relevant window are dropped before residualizing; because each tier row applies this filter within its own five-year window $[T_r, T_r + 5]$, a country that did not issue during the GDDS window but issued during the eGDDS+NSDP window is excluded from the GDDS row but included in the pooled row. The pooled row stacks all three frontier windows and drops countries with no positive issuance across *any* of the three windows, so its observation counts ($N_T = 181$, $N_O = 343$) exceed the sum of the three tier-specific rows ($N_T = 131$, $N_O = 208$).

Taken together, Sections 3.1 and 3.2 establish that distinct disclosure regimes exist, that they are durable enough to be treated as commitments, that adoption decisions are consistent with the model's prediction that sovereigns with higher recession risk are more likely to choose transparency, and that issuance outcomes differ between transparent and opaque regimes in a direction the model can speak to. None of this evidence identifies the structural mechanism. On the Mexico side, the evidence supports a narrow claim: the sparse reserve-disclosure rule was not introduced during the 1994 deterioration, but rather persisted through the post-1982 modernization period, when reforms were generally well received and Mexico was attracting large capital inflows. This pre-existence gives the regime its rule-like interpretation and explains why investors lacked current information about reserves in 1994. At the same time, we do not interpret the post-1995 move to weekly reporting as clean evidence of a borrower-induced switch to transparency: the shift coincided with U.S. and IMF support packages and has been characterized as partly creditor-imposed.⁶ On the cross-country side, adoption of disclosure

⁶See U.S. Government Accountability Office, *Mexico's Financial Crisis: Origins, Awareness, Assistance, and Initial Efforts to Recover*, GAO/GGD-96-56 (1996). Available at <https://www.govinfo.gov/content/pkg/GAOREPORTS-GGD-96-56/html/GAOREPORTS-GGD-96-56.htm>.

tiers is not random, and selection into higher transparency may correlate with unobserved fundamentals; the residualization mitigates but does not eliminate this concern, and the WEO-prior regressions establish consistency with the comparative static rather than a causal effect of recession risk on disclosure choices. [Choi and Hashimoto \(2018\)](#) and [Gonzalez-Garcia \(2022\)](#), using event-study and local-projection identification, establish the spread-reduction result more rigorously than the variance-comparison exercise here.

4 Optimal Disclosure

The baseline model compares two extreme information regimes: full disclosure and no disclosure. We now allow the sovereign to choose an intermediate disclosure rule and characterize the optimal policy from a constrained disclosure problem in the spirit of [Kamenica and Gentzkow \(2011\)](#). Our environment adds one feature not present in the standard information-design problem: the constraint that ensures the pooling equilibrium is sustained, which limits how opaque the sovereign can credibly be.

Environment. The timing is as in the baseline. Before the repayment-capacity type is realized, the sovereign commits to a signal structure. The type is

$$\theta \in \{\theta_L, \theta_H\}, \quad \Pr(\theta = \theta_L) = \pi, \quad \Pr(\theta = \theta_H) = 1 - \pi.$$

The signal takes values $m \in \{h, \ell\}$. We focus on a one-sided disclosure rule: the high type is always reported truthfully, while the low type is revealed only with probability ϕ :

$$\Pr(h|\theta_H) = 1, \quad \Pr(h|\theta_L) = 1 - \phi, \quad \Pr(\ell|\theta_L) = \phi.$$

Thus $\phi \in [0, 1]$ indexes the degree of disclosure. When $\phi = 0$ the signal is completely uninformative and both types send message h , corresponding to no disclosure. When $\phi = 1$ the signal fully reveals the type. Intermediate values $\phi \in (0, 1)$ correspond to partial disclosure: the low type is sometimes revealed and sometimes pooled with the high type.

The probability of the pooling message h is $\Pr(h) = 1 - \pi\phi$, and after message h lenders form the posterior

$$\mu(\phi) = \Pr(\theta_L|h) = \frac{\pi(1 - \phi)}{1 - \pi\phi}, \quad 1 - \mu(\phi) = \frac{1 - \pi}{1 - \pi\phi}.$$

In the risky-pooling equilibrium following message h , the low type defaults and the high type repays, so the bond price is

$$q(\phi) = \frac{1 - \mu(\phi)}{R} = \frac{1 - \pi}{R(1 - \pi\phi)}.$$

In the interior risky-pooling region the pooled debt choice is

$$b(\phi) = \frac{1}{1 + \beta} \left[\theta_H - \frac{\beta R y (1 - \pi\phi)}{1 - \pi} \right],$$

so that current consumption after the pooling message is

$$y + q(\phi)b(\phi) = \frac{1}{1 + \beta} \left[y + \frac{(1 - \pi)\theta_H}{R(1 - \pi\phi)} \right],$$

and the high type's repayment consumption is

$$\theta_H - b(\phi) = \frac{\beta}{1 + \beta} \left[\theta_H + \frac{R y (1 - \pi\phi)}{1 - \pi} \right].$$

If the low type is revealed (message ℓ), it borrows safely up to $b_L = \theta_L - y_d$ at price $1/R$.

Using log utility, the ex-ante value of a disclosure policy ϕ is

$$\begin{aligned} G(\phi, y_d) = & (1 - \pi\phi) \log \left\{ \frac{1}{1 + \beta} \left[y + \frac{(1 - \pi)\theta_H}{R(1 - \pi\phi)} \right] \right\} \\ & + \pi\phi \log \left(y + \frac{\theta_L - y_d}{R} \right) \\ & + \beta(1 - \pi) \log \left\{ \frac{\beta}{1 + \beta} \left[\theta_H + \frac{R y (1 - \pi\phi)}{1 - \pi} \right] \right\} \\ & + \beta\pi \log y_d. \end{aligned} \tag{8}$$

The first term is the period-1 payoff when the pooling message is sent (both types borrow $b(\phi)$); the second is the period-1 payoff of the revealed low type; the third is the period-2 payoff of the high type who repays; and the fourth is the period-2 payoff of the low type who defaults whether pooled or revealed.

Constrained disclosure problem. The risky-pooling equilibrium following message h is feasible only if the posterior probability of the low type does not exceed the threshold

$\pi^*(y_d)$ identified in Section 2:

$$\mu(\phi) = \frac{\pi(1-\phi)}{1-\pi\phi} \leq \pi^*(y_d).$$

As established there, $\pi'^*(y_d) > 0$: a higher default endowment expands the set of beliefs under which risky pooling can be sustained. The sovereign's problem is

$$\max_{\phi \in [0,1]} G(\phi, y_d) \quad \text{subject to} \quad \frac{\pi(1-\phi)}{1-\pi\phi} \leq \pi^*(y_d). \quad (9)$$

Proposition 4 (Partial disclosure and the default endowment). *Let $\phi^*(y_d)$ solve problem (9). Then*

$$\frac{d\phi^*(y_d)}{dy_d} \leq 0.$$

Proof. See Appendix A.7. □

A higher default endowment y_d exerts two forces on the optimal disclosure probability, both pointing toward less disclosure. First, it reduces the safe borrowing capacity of the revealed low type: $b_L = \theta_L - y_d$ is decreasing in y_d , so the period-1 consumption of a disclosed low type falls, reducing the marginal benefit of revelation. Second, it relaxes the feasibility constraint: a higher y_d raises $\pi^*(y_d)$, so the pooling equilibrium can be sustained under a less informative signal. When the constraint binds, this directly lowers the minimum disclosure level the sovereign must commit to. Both effects push $\phi^*(y_d)$ downward. This comparative static is the continuous analogue of the baseline result: a higher default endowment corresponds to lower deadweight losses from default, making the insurance value of pooling more attractive relative to the pricing gains from transparency.

5 Conclusion

We develop an analytic sovereign-default model in the Eaton-Gersovitz tradition with incomplete information, where the government chooses, *ex ante*, whether to disclose its repayment capacity to lenders. Under no disclosure, there are no separating equilibria; instead, a *risky pooling* equilibrium emerges in which both types issue the same face value, the L-type receives a positive price yet defaults in period 2, and the H-type repays but bears a price discount due to adverse selection. The mechanism relies on a strong assumption: the sovereign can commit *ex ante* to a disclosure policy. Under this assumption, by committing not to disclose, the sovereign secures insurance across states via this

cross-subsidization. We show that when deadweight losses from default are small, the set of lender priors for which no disclosure is preferred is larger (equivalently, as default becomes more costly, the region favoring no disclosure shrinks and disclosure becomes relatively more attractive).

We use Mexico's 1994–95 episode as a case study consistent with this mechanism, without claiming it tests the models assumptions. Before the crisis, Banco de Mexico released reserve data three times per year under a rule that had been in place throughout the post-1982 modernization period and was not revised during the 1994 deterioration. After the crisis, Mexico shifted durably to weekly reserve reporting. In the cross-country sample, we construct a proxy for the prior probability of low growth from historical IMF forecast errors, and find that countries whose priors are high are significantly more likely to adopt a new transparency tier in the year it becomes available, consistent with the models prediction that disclosure is preferred when the probability of being the L-type is sufficiently high. Early adopters of two of the three transparency tier margins also exhibit more dispersed issuance residuals than eligible non-adopters, consistent with the models implication that transparency makes issuances more dispersed across countries.

The commitment assumption is strong. An extension that microfound commitment through reputation, institutional design, or costly signaling would strengthen the theoretical foundation. We leave this to future work.

References

- Alesina, Alberto, Ricardo Hausmann, Rudolf Hommes, and Ernesto Stein**, "Budget institutions and fiscal performance in Latin America," *Journal of development Economics*, 1999, 59 (2), 253–273.
- Alt, James E and David Dreyer Lassen**, "Fiscal transparency, political parties, and debt in OECD countries," *European economic review*, 2006, 50 (6), 1403–1439.
- Amador, Manuel and Christopher Phelan**, "Reputation and sovereign default," *Econometrica*, 2021, 89 (4), 1979–2010.
- Asonuma, Tamon and Christoph Trebesch**, "Sovereign debt restructurings: preemptive or post-default," *Journal of the European Economic Association*, 2016, 14 (1), 175–214.
- Bassetto, Marco and Carlo Galli**, "Is inflation default? The role of information in debt crises," *American Economic Review*, 2019, 109 (10), 3556–3584.

- Bergemann, Dirk and Stephen Morris**, “Information design: A unified perspective,” *Journal of Economic Literature*, 2019, 57 (1), 44–95.
- , **Benjamin Brooks, and Stephen Morris**, “The limits of price discrimination,” *American Economic Review*, 2015, 105 (3), 921–957.
- Choi, Sangyup and Yuko Hashimoto**, “Does transparency pay? Evidence from IMF data transparency policy reforms and emerging market sovereign bond spreads,” *Journal of International Money and Finance*, 2018, 88, 171–190.
- Cole, Harold L and Timothy J Kehoe**, “Self-fulfilling debt crises,” *The Review of Economic Studies*, 2000, 67 (1), 91–116.
- Crawford, Vincent P and Joel Sobel**, “Strategic information transmission,” *Econometrica: Journal of the Econometric Society*, 1982, pp. 1431–1451.
- Croitorov, Olga**, “Optimal government transparency,” Technical Report DPS17.04, KU Leuven, Department of Economics, Discussion Paper Series February 2017.
- Dovis, Alessandro and Paolo Martellini**, “Long-Term Contracts, Commitment, and Optimal Information Disclosure,” 2024.
- e Castro, Miguel Faria, Joseba Martinez, and Thomas Philippon**, “Runs versus lemons: information disclosure and fiscal capacity,” *The Review of Economic Studies*, 2017, 84 (4), 1683–1707.
- Eaton, Jonathan and Mark Gersovitz**, “Debt with potential repudiation: Theoretical and empirical analysis,” *The Review of Economic Studies*, 1981, 48 (2), 289–309.
- Gamboa, Juliana**, “Hidden chinese lending,” *Available at SSRN* 4617798, 2023.
- Garcia, Daniel and Matan Tsur**, “Information design in competitive insurance markets,” *Journal of Economic Theory*, 2021, 191, 105160.
- Goldstein, Itay and Liyan Yang**, “Information disclosure in financial markets,” *Annual Review of Financial Economics*, 2017, 9, 101–125.
- **and Yaron Leitner**, “Stress tests and information disclosure,” *Journal of Economic Theory*, 2018, 177, 34–69.
- Goldstein, Morris et al.**, “The Asian financial crisis,” 1998.
- Gonzalez-Garcia, Jesus**, *Improving sovereign financing conditions through data transparency*, International Monetary Fund, 2022.

- Gorton, Gary and Guillermo Ordoñez**, "Fighting crises with secrecy," *American Economic Journal: Macroeconomics*, 2020, 12 (4), 218–245.
- Grossman, Herschel I and John B Van Huyck**, "Sovereign debt as a contingent claim: Excusable default, repudiation, and reputation," Technical Report, National Bureau of Economic Research 1985.
- Gu, Grace Weishi and Zachary R Stangebye**, "Costly information and sovereign risk," *International Economic Review*, 2023, 64 (4), 1397–1429.
- Guler, Bulent, Yasin Kürşat Önder, and Temel Taskin**, "Asymmetric Information and Sovereign Debt Disclosure," *Available at SSRN 4048018*, 2022.
- Hirshleifer, Jack**, "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 1971, 61 (4), 561–574.
- Horn, Sebastian, David Mihalyi, Philipp Nickol, and César Sosa-Padilla**, "Hidden debt revelations," *Unpublished. World Bank, Washington, DC*, 2022.
- Immorlica, Nicole, Andre M Sztutman, and Robert M Townsend**, "Optimal Credit Scores Under Adverse Selection," in "Proceedings of the 23rd ACM Conference on Economics and Computation" 2022, pp. 737–738.
- Inostroza, Nicolas and Alessandro Pavan**, "Adversarial coordination and public information design," *Available at SSRN 4531654*, 2023.
- Kamenica, Emir and Matthew Gentzkow**, "Bayesian persuasion," *American Economic Review*, 2011, 101 (6), 2590–2615.
- Nachman, David C and Thomas H Noe**, "Optimal design of securities under asymmetric information," *The Review of Financial Studies*, 1994, 7 (1), 1–44.
- Perez, Diego J**, "Sovereign debt maturity structure under asymmetric information," *Journal of International Economics*, 2017, 108, 243–259.
- Sandleris, Guido**, "Sovereign defaults: Information, investment and credit," *Journal of international Economics*, 2008, 76 (2), 267–275.
- Szkup, Michal**, "Preventing self-fulfilling debt crises: A global games approach," *Review of Economic Dynamics*, 2022, 43, 22–55.
- Zabai, Anna**, "Managing default risk," 2014.

A Appendix

A.1 Proof of Lemma 1

Default decision. At $t = 2$ the sovereign decides whether to default or repay on its debt b to maximize consumption at $t = 2$. We denote the default decision by $\delta^*(\theta, b) \in \{0, 1\}$:

$$\delta^*(b, \theta) = \arg \max_{\delta \in \{0, 1\}} \delta \times y_d + (1 - \delta) \times (\theta - b) \quad (10)$$

The default policy at $t = 2$ depends solely on the amount of debt picked at $t = 1$, and the sovereign's type:

$$\delta^*(b, \theta) = \begin{cases} 1, & \text{if } b > \bar{b}(\theta) \\ 0, & \text{if } b \leq \bar{b}(\theta) \end{cases}$$

$$y_d = \theta - \bar{b}(\theta)$$

where $\bar{b}(\theta) \equiv \theta - y_d$ is the upper repayment threshold for debt levels that make each type indifferent between repaying and defaulting. The default policy determines consumption at $t = 2$:

$$c_2^*(b, \theta) = \delta^*(b, \theta) \times y_d + [1 - \delta^*(b, \theta)] \times (\theta - b)$$

Lenders' problem. The strategy for the lenders is a price schedule $q_{FD}(b, \theta)$:

$$q_{FD}(b, \theta) = \begin{cases} 0, & \text{if } b > \bar{b}(\theta) \\ 1/R, & \text{if } b \leq \bar{b}(\theta) \end{cases}$$

The following expression gives the interior solution for debt:

$$\frac{\theta - \beta R y}{1 + \beta}$$

However, conditional on repayment the sovereign's choice of debt can be a corner solution:

$$b(\theta) = \min \left\{ \frac{\theta - \beta R y}{1 + \beta}, \theta - y_d \right\}$$

□

A.2 Proof of Proposition 1

Let $b_{ND}^*(\theta_H)$ and $b_{ND}^*(\theta_L)$ be the separating choices of debt. By definition, in a separating equilibrium, lenders can distinguish between types, that is, $\mu(b_{ND}^*(\theta_L)) = 1$ and $\mu(b_{ND}^*(\theta_H)) = 0$. By the zero profit condition, debt is priced at $q_{ND}(b) \in \{0, 1/R\}$ for $b \in \{b_{ND}^*(\theta_L), b_{ND}^*(\theta_H)\}$.

Claim 1. *For any type θ , we must have that $q_{ND}(b_{ND}^*(\theta)) > 0$*

Proof. Suppose not, that is, there exists $\theta^* \in \{\theta_H, \theta_L\}$ such that $q_{ND}(b_{ND}^*(\theta^*)) = 0$. Then $b_{ND}^*(\theta^*) > \theta^* - y_d$, which gives the utility:

$$\begin{aligned} u(y) + \beta u(y_d) &< \underbrace{u\left(y + \frac{\theta_L - y_d}{R}\right) + \beta u(y_d)}_{\text{L-type's payoff by picking } b_{ND}^*(\theta_L) = \theta_L - y_d} \\ u(y) + \beta u(y_d) &< \underbrace{u\left(y + \frac{\theta_L - y_d}{R}\right) + \beta u(\theta_H - \theta_L + y_d)}_{\text{H-type's payoff by picking } b_{ND}^*(\theta_H) = \theta_L - y_d} \end{aligned}$$

We thus conclude that any choice of debt that induces zero price (and default) cannot be our separating equilibrium. $\square$

By Assumption 1 the L-type sovereign $\theta_L - y_d$ is preferred over any $b < \theta_L - y_d$. Therefore $b_{ND}^*(\theta_L) = \theta_L - y_d$, and the L-type's utility at $t = 2$ is given by $u(y_d)$. If the L-type is constrained at $\theta_L - y_d$, the H-type prefers $\theta_L - y_d$ to any debt level smaller than this.

Our candidate for separating equilibrium must be such that $b_{ND}^*(\theta_H) \in (\theta_L - y_d, \theta_H - y_d)$, that is, the H-type's choice of debt is in the default region for the L-type, a region where the L-type's payoff is strictly increasing in its debt level. Moreover by the above claim: $q_{ND}(b_{ND}^*(\theta_H)) = 1/R$. Therefore:

$$u\left(y + \frac{\theta_L - y_d}{R}\right) + \beta u(y_d) < u\left(y + \frac{b_{ND}^*(\theta_H)}{R}\right) + \beta u(y_d)$$

We thus conclude that $b_{ND}^*(\theta_H)$ is a profitable deviation for the L-type regardless of the belief system. $\square$

A.3 Mech. Design Approach: H-type preferred pooling PBE

This section considers a mechanism design problem in which the designer is the H-type sovereign. We start by defining some useful objects:

$$v_L^a = u(y) + \beta u(\theta_L) \quad (11)$$

where superscript a denotes autarky (repaying without borrowing: $b = 0$, so $c_1 = y$ and $c_2 = \theta_L$). We split the problem in two: (i) investors break even for each of the types; (ii) investors pick a pooling contract in which the L-type defaults.

A.3.1 The Best Outcome with Default

We now consider a problem in which the designer offers a pooling contract and the L-type defaults. In this case, ICR-L (Incentive Compatibility Repayment for the L-type) is not required further and ICD-L (Incentive Compatibility of Default for L-type) is reduced to $\tilde{b} > \theta_L - y_d$. Since investors are offering a pooling contract, its participation constraint (PC-I) takes into account the probability of facing each of the types.

$$\max_{\{\tilde{b}, \tilde{q}\}} u(y + \tilde{q}\tilde{b}) + \beta u(\theta_H - \tilde{b}) \quad (P)$$

subject to

$$\tilde{b} \leq \theta_H - y_d \quad (\text{SUST-H}, \nu)$$

$$\tilde{b} \geq \theta_L - y_d \quad (\text{ICD-L}, \chi)$$

$$u(y + \tilde{q}\tilde{b}) + \beta u(y_d) \geq v_L^a \quad (\text{PC-L}, \eta)$$

$$\underbrace{\pi \times [-\tilde{q}\tilde{b}] + (1 - \pi) \times \left[\frac{\tilde{b}}{R} - \tilde{q}\tilde{b} \right]}_{\frac{(1-\pi)}{R} \tilde{b} - \tilde{q}\tilde{b}} \geq 0 \quad (\text{PC-I}, \xi)$$

Lemma 2. Suppose $y_d \leq \theta_L$, and the L-type defaults, that is, $\tilde{b} > \theta_L - y_d$ ($\chi = 0$). Therefore, investors make zero profits ($\xi > 0$), that is,

$$\tilde{q} = \frac{1 - \pi}{R}$$

Proof. Let $(\tilde{q}, \tilde{b})$ be a solution to (P). Assume, by contradiction, that investors make positive profits, that is, $\tilde{q} < \frac{1-\pi}{R}$. Since the H-type's payoff is strictly increasing in $\tilde{q}$, there

exists $q' > \tilde{q}$ that increases the H-type's payoff. Since the payoff of type L is also increasing in the price of the debt, and when $\tilde{q}$ PC-L is satisfied, we must have that q' also satisfies PC-L, which means that $\tilde{q}$ is not a solution to (P). $\square$

Proposition 5. *The risky pooling contract maximizes the payoff of the H-type among all contracts in which the L-type defaults.*

Proof.

$$(1 + \eta)u' \left(y + \frac{1 - \pi}{R} \times \tilde{b} \right) (1 - \pi) - R\beta u'(\theta_H - \tilde{b}) = R(v - \chi)$$

By looking for a solution in which the L-type, defaults, that is, $\tilde{b} \in (\theta_L - y_d, \theta_H - y_d)$, we have that $v = \chi = 0$. If the L-type's PC-I binds and PC-L is slack ($\eta = 0$), we have:

$$\tilde{b} = \frac{\theta_H - \frac{R\beta}{1 - \pi} \times y}{1 + \beta}$$

We now look for a parametrization in which $\eta = 0$. If $y_d = \theta_L$, PC-L is slack. By continuity of $u \left(y + \frac{1 - \pi}{R} \tilde{b} \right) + \beta u(y_d)$ in y_d , there must exist $\epsilon > 0$ such that PC-L holds with strict inequality for $y_d \in (\theta_L - \epsilon, \theta_L)$. Define:

$$\bar{\epsilon} \equiv \sup\{\epsilon : \epsilon > 0, \text{ such that PC-L holds with strict inequality for } y_d = \theta_L - \epsilon\}$$

By taking $y_d \equiv \theta_L - \bar{\epsilon}$ we have that $\eta = 0$. $\square$

A.3.2 Best Repayment Outcome

Proposition 6. *The safe pooling outcome maximizes the H-type's payoff subject to: (i) lenders make zero profits; (ii) the sustainability constraint for the L-type binds; (iii) the L-type prefers to default rather than repay the H-type's contract.*

Proof. We start considering the problem in which both types repay and investors break even for each type. This case asks for two ICs for the L-type: (i) repays H-type's contract; (ii) defaults on the H-type's contract. Since both types repay, we also need a sustainability constraint for the H-type.

$$\max_{\{b_H, b_L, q_H, q_L\}} u(y + q_H b_H) + \beta u(\theta_H - b_H) \tag{S}$$

subject to

$$\begin{aligned}
b_H &\leq \theta_H - y_d & (\text{SUST-H}, \nu_H) \\
b_L &\leq \theta_L - y_d & (\text{SUST-L}, \nu_L) \\
u(y + q_L b_L) + \beta u(\theta_L - b_L) &\geq u(y + q_H b_H) + \beta u(y_d) & (\text{ICD-L}, \chi) \\
u(y + q_L b_L) + \beta u(\theta_L - b_L) &\geq u(y + q_H b_H) + \beta u(\theta_L - b_H) & (\text{ICR-L}, \lambda) \\
u(y + q_L b_L) + \beta u(\theta_L - b_L) &\geq v_L^a & (\text{PC-L}, \eta) \\
b_L \times \left[\frac{1}{R} - q_L \right] &\geq 0 & (\text{PCI-L}, \xi_L) \\
b_H \times \left[\frac{1}{R} - q_H \right] &\geq 0 & (\text{PCI-H}, \xi_H)
\end{aligned}$$

We look for a solution with $\xi_H, \xi_L > 0$, which implies $q_H = q_L = 1/R$. By the FOC with respect to b_L :

$$u' \left(y + \frac{b_L}{R} \right) - R\beta u'(\theta_L - b_L) = \frac{R\nu_L}{\chi + \lambda + \eta}$$

Consider first that case in which the sustainability constraint binds for the L-type: $\nu_L > 0$, that is, $b_L = \theta_L - y_d$. We can then rewrite the ICD-L as:

$$u \left(y + \frac{\theta_L - y_d}{R} \right) \geq u \left(y + \frac{b_H}{R} \right) \implies b_H \leq \theta_L - y_d$$

Since SUST-L binds we know that the L-type is indifferent between default and repayment of b_L . By the FOC with respect to b_H :

$$u'(y + b_H/R) - R\beta u'(\theta_H - b_H) = \chi u'(y + b_H/R) + \lambda[u'(y + b_H/R) - R\beta u'(\theta_L - b_H)]$$

If the relevant constraint is ICD-L (picking the H-type's contract and defaulting is better than picking the H-type's contract and repaying):

$$u(y_d) \geq u(\theta_L - b_H)$$

That is:

$$b_H \geq \theta_L - y_d$$

□

A.4 Proof of Proposition 2

Lemma 3. *There exists $\pi_L(y_d)$ such that, for $\pi < \pi_L(y_d)$, the L-type sovereign prefers the risky pooling debt level.*

Proof. It is enough to check the L-type's deviation to the *safe pooling*

$$\begin{aligned} y + \frac{1-\pi}{R} \left[\frac{\theta_H - \frac{\beta R y}{1-\pi}}{1+\beta} \right] &> y + \frac{\theta_L - y_d}{R} \\ \underbrace{1 - \left[\frac{(1+\beta)(\theta_L - y_d) + \beta R y}{\theta_H} \right]}_{\equiv \pi_L(y_d)} &> \pi \end{aligned}$$

□

We now check whether the H-type prefers the *risky pooling* over issuing $\theta_L - y_d$ at a risk-free price. Let the payoff of the *risky pooling* be defined as⁷:

$$\mathcal{F}_1(\pi) \equiv (1+\beta) \log \left(\frac{y + \frac{1-\pi}{R} \times \theta_H}{1+\beta} \right) + \beta \log \left(\frac{\beta R}{1-\pi} \right)$$

Lemma 4. $\mathcal{F}_1(\pi)$ has a minimum ρ' . Moreover, at its minimum:

$$\mathcal{F}_1(\rho') = \log(y) + \beta \log(\theta_H)$$

Proof. The first order condition of $\mathcal{F}_1$ is:

$$\mathcal{F}'_1(\rho') = (1+\beta) \times \frac{-\theta_H}{R y + (1-\rho')\theta_H} + \frac{\beta}{1-\rho'} = 0$$

Denote by $\rho' \in (0, 1)$ the unique root of $\mathcal{F}'_1(\pi) = 0$. Solving gives $\rho' = 1 - \frac{\beta R y}{\theta_H}$. The second derivative of $\mathcal{F}_1$:

$$\mathcal{F}''_1(\pi) = -\frac{(1+\beta)\theta_H^2}{[R y + (1-\pi)\theta_H]^2} + \frac{\beta}{(1-\pi)^2}$$

⁷Function $\mathcal{F}_1$ is obtained by plugging the equilibrium debt and prices of the *risky-pooling* equilibrium into the payoff of the high-type under repayment.

Evaluating the second derivative at the critical point ρ' we get:

$$\begin{aligned}\mathcal{F}_1''(\rho') &= \frac{-(1+\beta)\theta_H^2}{[Ry + (1-\rho')\theta_H]^2} + \frac{\beta}{(1-\rho')^2} \\ &= \frac{-1}{1+\beta} \times \left(\frac{\theta_H}{Ry}\right)^2 + \frac{1}{\beta} \times \left(\frac{\theta_H}{Ry}\right)^2 > 0\end{aligned}$$

□

Lemma 5. *The safe pooling payoff for the H-type is decreasing in y_d .*

Proof. We have already shown that when evaluated at ρ' , $\mathcal{F}$ attains its minimum and $\mathcal{F}(\rho') = \log(y) + \beta \log(\theta_H)$. We want to know under which conditions the *safe pooling* payoff of the H-type is decreasing in y_d , that is:

$$\begin{aligned}\frac{-1}{Ry + \theta_L - y_d} + \frac{\beta}{\theta_H - \theta_L + y_d} &< 0 \\ \frac{\beta}{\theta_H - \theta_L + y_d} &< \frac{1}{Ry + \theta_L - y_d} \\ \beta Ry + \beta \theta_L - \beta y_d &< \theta_H - \theta_L + y_d \\ \beta Ry + (1+\beta)\theta_L &< (1+\beta)y_d + \theta_H \\ \underbrace{\theta_L - y_d}_{b_{FD}^*(\theta_L)} &< \underbrace{\frac{\theta_H - \beta Ry}{1+\beta}}_{b_{FD}^*(\theta_H)}\end{aligned}$$

□

Lemma 6. *There exists $\pi_H(y_d)$ such that for $\pi < \pi_H(y_d)$, the H-type prefers the risky pooling debt level.*

Proof. $\mathcal{F}_1$ is equal to the safe pooling payoff at its minimum, attained when $y_d = \theta_L$. Also, the safe pooling payoff is decreasing in y_d , which means that for $y_d < \theta_L$, the safe pooling payoff lies above $\mathcal{F}_1$ at its minimum. By continuity of $\mathcal{F}_1$ we get the existence of $\pi_H(y_d)$ that makes the H-type indifferent between the safe and risky pooling outcomes. □

Define $\pi^*(y_d) \equiv \min\{\pi_L(y_d), \pi_H(y_d)\}$. Both cutoffs are positive for $y_d \in Y_d$, which establishes the single cutoff asserted in Proposition 2. □

A.5 Comparison to Hidden Debt Models

A.5.1 Commitment vs. Monitoring

To make our model relatable to [Horn et al. \(2022\)](#), we provide a different interpretation of the disclosure policy. In the main text we have:

$$\phi(\pi) = \mathbb{1}_{\{V_{FD}(\pi; y_d) > V_{ND}(\pi; y_d)\}}$$

Suppose now that the sovereign has a set of available *messages* to be sent to the lenders: $M \equiv \{h, \ell\}$. We can express the disclosure policy as a conditional distribution over messages:

$$\begin{aligned}\phi &= \mathbb{P}(m = \ell | \theta = \theta_L) \\ 1 - \phi &= \mathbb{P}(m = h | \theta = \theta_L)\end{aligned}$$

Suppose that $\mathbb{P}(m = h | \theta = \theta_H) = 1$. The posterior belief of the lenders is then:

$$\mathbb{P}(\theta = \theta_L | m = h) = \frac{(1 - \phi)\pi}{(1 - \phi)\pi + 1 - \pi}$$

When $\phi = 0$, the sovereign *always sends the h-message*, making the posterior of the lender to be the same as the prior π (which is equivalent to no-disclosure in the main text). If $\phi = 1$, on the other hand, lenders are always informed (equivalent to FD in the main text). Suppose now we remove the commitment assumption, that is, the equivalent of Figure 1 is now:

Figure 6: Timing without the commitment assumption

t = 0	t = 1	t = 2
1) Government learns its type (θ_i)	1) Message is chosen $m \in \{h, \ell\}$ 2) Choice of debt 3) Price is determined 4) Consumption takes place	1) Default decision 2) Consumption


Under this timing, our model would feature a cheap talk equilibrium where both types would always send an h-message ([Crawford and Sobel, 1982](#)). This assumption micro-

finds the assumption from [Horn et al. \(2022\)](#) that governments cannot credibly communicate their state, transferring to the lenders the ability to monitor the sovereign and trigger revelations.

A.5.2 Comparative Statics (y).

[Horn et al. \(2022\)](#) shows that hidden sovereign debt tends to accumulate during boom years and is typically revealed during downturns. To relate this empirical pattern to our framework, we perform a comparative statics exercise with respect to the endowment parameter y . We interpret a low y as representing bad times, when resources are scarce. In [Horn et al. \(2022\)](#), lenders have stronger incentives to pay monitoring costs and uncover information during such periods. In contrast, in our setting, information revelation is a policy choice made by the government.

Figure 7: Difference between expected payoffs under ND and FD, for different beliefs.



Note: Parameters used: $\beta = 0.92$, $R = 1$, $y_d = (y_d^{\max} + y_d^{\min})/2$, $\theta_L = 1.25$, $\theta_H = 5$.

Figure 7 shows that for lower values of y , the range of parameters for which no disclosure is preferred expands. The reason is that opacity enables low-type governments to borrow more; when income is low, this insurance motive dominates. Although this pattern may appear opposite to [Horn et al. \(2022\)](#), the difference arises because, in our model, the disclosure is endogenously chosen by the sovereign rather than driven by lenders incentives to monitor⁸.

⁸In a world where the government can credibly communicate its state to the market *and* lenders can monitor the sovereign, we would expect to see the monitoring fees being paid exactly in moments when the government wants to hide information from lenders.

A.6 Simple knife edge case

To simplify the analysis, we describe a set of conditions that express a limiting knife-edge case. Throughout this subsection, $\tilde{b}$ denotes $\tilde{b}(\pi)$ for a given prior π ; the dependence on π is suppressed for brevity.

Assumption 2 (Limiting Case). *Suppose the following parametric restrictions hold:*

1. *The H-type is constrained (chooses the maximum amount of debt it could repay)*

$$y_d > \frac{\beta R}{1 + \beta} \left(\frac{\theta_H}{R} + \frac{y}{1 - \pi} \right)$$

2. *Low deadweight loss from default*

$$y_d \rightarrow \theta_L$$


3. *Positive borrowing by H-type.*

$$\pi < 1 - \beta R$$

In the next proposition, we characterize the pooling equilibrium under Assumption (2).

Lemma 7. *Under Assumption (2), $\tilde{b} = \theta_H - y_d$ is the unique debt level sustained in a pooling equilibrium.*

Figure 8: Price schedule for the pooling eq.



Proof. Under the suggested belief system, the price schedule is represented in Figure 8.

The optimal choice of debt for a H-type sovereign that faces a price $\frac{1-\pi}{R}$ is:

$$\tilde{b} = \min \left\{ \frac{\theta_H(1-\pi) - R\beta y}{(1+\beta)(1-\pi)}, \theta_H - y_d \right\}$$

In the limiting case $y_d \rightarrow \theta_L$ we thus have:

$$\tilde{b} = \min \left\{ \frac{\theta_H - R\beta \frac{y}{1-\pi}}{1+\beta}, \theta_H - \theta_L \right\}$$

The H-type sovereign has positive marginal utility for an additional unit of debt, even when debt is at its upper limit. We can thus conclude that there is no profitable deviation for the H-type that would yield the same price, $\frac{1-\pi}{R}$. Since $y_d \rightarrow \theta_L$, no debt level can be issued at a risk-free price, which means that autarky is the only candidate for a profitable deviation for the H-type sovereign. For autarky to be preferred to the pooling equilibrium the following inequality must hold:

$$u \left(y + \frac{1-\pi}{R} \times (\theta_H - \theta_L) \right) - u(y) < \beta [u(\theta_H) - u(\theta_L)]$$

This holds if $\pi > 1 - R\beta$, a contradiction. On the other hand, L-types' payoff is strictly increasing in debt, conditional on getting a positive price, which means they also will not deviate from $\theta_H - \theta_L$.

□

The above result is important for us to evaluate the sovereigns' disclosure policy as it ensures there is a *unique* payoff at $t = 0$ associated with not disclosing the aggregate state θ . We can now define such payoff as follows:

$$\begin{aligned} V_{ND}(\pi; y_d) \equiv & \pi \underbrace{\left[\log(y + q_{ND}(\tilde{b})\tilde{b}) + \beta \log(y_d) \right]}_{v_{ND}(\theta_L; y_d)} \\ & + (1-\pi) \underbrace{\left[\log(y + q_{ND}(\tilde{b})\tilde{b}) + \beta \log(\theta_H - \tilde{b}) \right]}_{v_{ND}(\theta_H; y_d)} \end{aligned}$$

$$V_{ND}(\pi; y_d) = \log(y + q_{ND}(\tilde{b})\tilde{b}) + \beta[(1-\pi) \log(\theta_H - \tilde{b}) + \pi \log(y_d)]$$

The disclosure policy picks the *sovereign-preferred equilibrium*, that is:

$$\phi^* = \arg \max_{\phi \in \{0,1\}} \phi \times V_{FD}(\pi; y_d) + (1 - \phi) \times V_{ND}(\pi; y_d)$$

Proposition 7. Under Assumption (2) $V_{ND}(\pi; y_d) > V_{FD}(\pi; y_d)$, that is, $\phi^* = 0$

Proof. Using the equilibrium outcomes derived, the expression for the payoffs under FD and ND reduces to:

$$\begin{aligned} V_{FD}(\pi; y_d) &= \pi \times u(y) + (1 - \pi) \times u\left(y + \frac{\theta_H - \theta_L}{R}\right) + \beta u(\theta_L) \\ V_{ND}(\pi; y_d) &= u\left(y + \frac{1 - \pi}{R} \times (\theta_H - \theta_L)\right) + \beta u(\theta_L) \end{aligned}$$

The result follows from the strict concavity of u . □

The above result shows that if there are no deadweight losses from defaulting at the L-state, not disclosing information to international lenders is preferred to disclosing it fully. By committing not to reveal θ , the sovereign can transfer resources from a realization of the H-state to the L-state by playing a pooling equilibrium. In this equilibrium, the L-type takes as much debt as the H-type, at a positive price, and defaults, while the H-type receives a price below $1/R$ even though it would optimally repay. Observe that our results do not critically depend on assuming a logarithmic utility function. Because the high-type sovereign is constrained, that is, it issues the maximum amount of debt that lenders are willing to purchase at a positive price, the sovereign's payoff is equalized across all realizations of the state variable θ . Due to the strict concavity of the utility function, the sovereign prefers this equalized payoff to the risky, lottery-like payoff that would result from a full disclosure policy.

Although restrictive, the limiting assumptions imposed above are useful to illustrate the mechanism that makes not disclosing preferred to the sovereign. In equilibrium, given that both the L and H-types are constrained, debt choices are $b \in \{\theta_L - y_d, \theta_H - y_d\}$. Considering the limiting case $y_d \rightarrow \theta_L$ ensures that one of such choices is zero, providing a sharp characterization of the mechanism in our model by ruling out the issuance of risk-free debt. Bounding the probability of being an L-type by $1 - \beta R$ is slightly stronger than the $\beta R < 1$ assumption, and ensures that the H-type issues positive debt in equilibrium.

A.7 Proof of Proposition 4

The proof considers two cases depending on whether the risky-pooling feasibility constraint binds.

Case 1: Slack feasibility constraint. Suppose the constraint is slack and $\phi^*(y_d) \in (0, 1)$ satisfies the unconstrained first-order condition $G_\phi(\phi^*(y_d), y_d) = 0$. An interior optimum is assumed to exist; sufficient conditions are $G_\phi(0, y_d) > 0$ and $G_\phi(1, y_d) < 0$. Differentiating G from (8) with respect to ϕ gives

$$G_\phi(\phi, y_d) = -\pi \log \left\{ \frac{1}{1+\beta} \left[y + \frac{(1-\pi)\theta_H}{R(1-\pi\phi)} \right] \right\} + \frac{(1-\pi\phi) \frac{\pi(1-\pi)\theta_H}{R(1-\pi\phi)^2}}{y + \frac{(1-\pi)\theta_H}{R(1-\pi\phi)}} \\ - \beta(1-\pi) \frac{\frac{\pi R y}{1-\pi}}{\theta_H + \frac{R y (1-\pi\phi)}{1-\pi}} + \pi \log \left(y + \frac{\theta_L - y_d}{R} \right).$$

The only term in G_ϕ that depends on y_d is the last one. Differentiating,

$$G_{\phi y_d}(\phi, y_d) = \pi \frac{\partial}{\partial y_d} \log \left(y + \frac{\theta_L - y_d}{R} \right) = -\frac{\pi}{R \left(y + \frac{\theta_L - y_d}{R} \right)} < 0.$$

To sign $G_{\phi\phi}$, compute the second derivative of each term in G . The pooling current-consumption term $(1-\pi\phi) \log\{\dots\}$ contributes

$$-(1-\pi\phi) \left[\frac{\frac{\pi(1-\pi)\theta_H}{R(1-\pi\phi)^2}}{y + \frac{(1-\pi)\theta_H}{R(1-\pi\phi)}} \right]^2,$$

because the two terms from the product rule cancel (see the supplemental computation). The high-type repayment term $\beta(1-\pi) \log\{\dots\}$ contributes

$$-\beta(1-\pi) \left[\frac{\frac{\pi R y}{1-\pi}}{\theta_H + \frac{R y (1-\pi\phi)}{1-\pi}} \right]^2.$$

The revealed-low-type term is linear in ϕ and contributes zero. Hence

$$G_{\phi\phi}(\phi, y_d) = -(1 - \pi\phi) \left[\frac{\frac{\pi(1 - \pi)\theta_H}{R(1 - \pi\phi)^2}}{y + \frac{(1 - \pi)\theta_H}{R(1 - \pi\phi)}} \right]^2 - \beta(1 - \pi) \left[\frac{\frac{\pi Ry}{1 - \pi}}{\theta_H + \frac{Ry(1 - \pi\phi)}{1 - \pi}} \right]^2 < 0, \quad (12)$$

so G is strictly concave in ϕ under log utility. Applying the implicit function theorem to $G_\phi(\phi^*(y_d), y_d) = 0$:

$$\frac{d\phi^*(y_d)}{dy_d} = -\frac{G_{\phi y_d}}{G_{\phi\phi}} = \frac{\pi}{R\left(y + \frac{\theta_L - y_d}{R}\right) G_{\phi\phi}(\phi^*(y_d), y_d)} < 0.$$

Case 2: Binding feasibility constraint. Since $\partial\mu/\partial\phi = -\pi(1 - \pi)/(1 - \pi\phi)^2 < 0$, a higher ϕ lowers the posterior after h . When the constraint binds,

$$\phi^*(y_d) = \frac{\pi - \pi^*(y_d)}{\pi[1 - \pi^*(y_d)]},$$

provided $\pi^*(y_d) < \pi$ (if $\pi^*(y_d) \geq \pi$ the constraint is non-binding at $\phi = 0$). Differentiating,

$$\frac{d\phi^*(y_d)}{dy_d} = -\frac{1 - \pi}{\pi[1 - \pi^*(y_d)]^2} \pi'^*(y_d) < 0,$$

since $\pi'^*(y_d) > 0$ by Proposition 2.

Combining both cases, $\phi^*(y_d)$ is weakly decreasing in y_d , with the inequality strict in the interior of each case. $\square$

B Robustness and extensions

B.1 Welfare bounds across pooling equilibria

In this section we compare full-disclosure against the entire set of pooling PBEs under no disclosure. We define the full set of pooling PBEs, identify ex-ante welfare bounds, and establish two complementary results: (i) full disclosure strictly dominates the *worst* no-disclosure equilibrium; (ii) the ex-ante welfare-maximizing no-disclosure equilibrium is itself a valid PBE for sufficiently small π and, under condition ($\star$), dominates full disclosure locally. Together these results show that the main welfare result holds at the ex-ante

constrained optimum not merely at a favorable selection.

Equilibrium set and welfare bounds. Define the set of all pooling PBEs under no disclosure:

$$\mathcal{E}_{\text{ND}}(\pi, y_d) \equiv \{e = (b, \delta(b, \theta), q_{\text{ND}}(b), \mu(b)) : e \text{ is a pooling PBE as in Section 2}\}.$$

For each $e \in \mathcal{E}_{\text{ND}}$ with the L-type defaulting, the ex-ante payoff is:

$$V_0(e) = u(y + q_{\text{ND}}(b)b) + \beta [\pi u(y_d) + (1 - \pi) u(\theta_H - b)].$$

Define the ex-ante welfare bounds:

$$\bar{V}_{\text{ND}}(\pi, y_d) \equiv \sup_{e \in \mathcal{E}_{\text{ND}}} V_0(e), \quad \underline{V}_{\text{ND}}(\pi, y_d) \equiv \inf_{e \in \mathcal{E}_{\text{ND}}} V_0(e).$$

Safe and risky pooling payoffs. Recall $\bar{b}(\theta) \equiv \theta - y_d$ from Lemma 1. The *safe-pooling* debt is $b^S \equiv \theta_L - y_d$, at which both types repay at price $1/R$:

$$V^S(\pi, y_d) \equiv u\left(y + \frac{\theta_L - y_d}{R}\right) + \beta [\pi u(y_d) + (1 - \pi) u(\theta_H - \theta_L + y_d)]. \quad (13)$$

Risky pooling uses $b \in (b^S, \theta_H - y_d]$, the L-type defaults, and the price is $(1 - \pi)/R$:

$$V^R(b; \pi, y_d) \equiv u\left(y + \frac{1 - \pi}{R} b\right) + \beta [\pi u(y_d) + (1 - \pi) u(\theta_H - b)]. \quad (14)$$

Under the off-path belief $\mu(b') = 1$ for all $b' \neq b$ in the risky region, the relevant IC comparison for each type is between candidate debt b and the safe-pooling deviation b^S . The L-type IC reduces to:

$$b \geq \frac{b^S}{1 - \pi} = \frac{\theta_L - y_d}{1 - \pi}.$$

Denoting $V_H^R(b; \pi) \equiv u(y + \frac{1 - \pi}{R} b) + \beta u(\theta_H - b)$ and $V_H^S \equiv u(y + \frac{b^S}{R}) + \beta u(\theta_H - b^S)$, the *feasible risky-pooling set* is:

$$\mathcal{B}^R(\pi, y_d) \equiv \left\{ b \in (b^S, \theta_H - y_d] : b \geq \frac{b^S}{1 - \pi} \text{ and } V_H^R(b; \pi) \geq V_H^S \right\}.$$

Lemma 8 (Worst ND equilibrium). *Restricting to equilibria with $b \geq b^S$, $\underline{V}_{\text{ND}}(\pi, y_d) =$*

$V^S(\pi, y_d)$.

Proof. Every $b \in \mathcal{B}^R(\pi, y_d)$ satisfies both IC constraints relative to b^S , so in particular $V^R(b; \pi, y_d) \geq V^S(\pi, y_d)$ for all $b \in \mathcal{B}^R$. Hence safe pooling is the lower bound within this set. $\square$

Remark. The infimum statement is restricted to the pooling region $b \geq b^S$. Under Assumption 1, the L-type's constrained optimum lies at b^S , so no type would deviate to lower debt under the maintained off-path belief system; pooling equilibria with $b < b^S$ therefore do not arise.

Lemma 9 (FD dominates safe pooling). $V^S(\pi, y_d) < V_{FD}(\pi, y_d)$ for all $y_d \in Y_d$ and $\pi \in (0, 1)$.

Proof. Under Assumption 1, $b_{FD}^*(\theta_L) = \theta_L - y_d = b^S$, so the L-type is indifferent between safe pooling and full disclosure. The H-type's full-disclosure debt is $b_{FD}^*(\theta_H) = \frac{\theta_H - \beta R y}{1 + \beta}$, which strictly exceeds b^S for $y_d \in Y_d$. The H-type is therefore strictly better off under full disclosure, and ex-ante welfare under FD strictly exceeds V^S . $\square$

Lemmas 8 and 9 together give $V_{FD}(\pi, y_d) > \underline{V}_{ND}(\pi, y_d)$: full disclosure dominates the worst no-disclosure equilibrium.

Ex-ante optimal risky-pooling debt. To characterize the upper bound $\bar{V}_{ND}$, maximize $V^R(b; \pi, y_d)$ over $b \in (b^S, \theta_H - y_d]$. With $u = \log$, the unconstrained interior FOC is:

$$\frac{\frac{1-\pi}{R}}{y + \frac{1-\pi}{R}b} = \frac{\beta(1-\pi)}{\theta_H - b},$$

yielding the ex-ante optimal risky-pooling debt:

$$\hat{b}(\pi) \equiv \frac{\theta_H - \beta R y}{1 + \beta(1-\pi)}. \quad (15)$$

For $\pi > 0$, $\hat{b}(\pi) > \tilde{b}(\pi)$ (the H-type preferred debt in the main text), with equality at $\pi = 0$; the ex-ante perspective places weight on L-type states and pushes the welfare-maximizing debt level higher.

Proposition 8. Let $y_d \in Y_d$. Define:

$$\tilde{\pi}_C(y_d) \equiv \frac{\frac{\beta R}{1+\beta} \left(y + \frac{\theta_H}{R} \right) - y_d}{\frac{\beta}{1+\beta} (\theta_H - y_d)},$$

$$\tilde{\pi}_L(y_d) \equiv 1 - \frac{\theta_L - y_d}{\theta_H - \beta R y - \beta (\theta_L - y_d)},$$

$$\tilde{\pi}_H(y_d) \equiv \text{the unique root of } \mathcal{F}_2(\pi) = K(y_d),$$

$$\text{where } \mathcal{F}_2(\pi) \equiv \frac{Ry + (1-\pi)\theta_H}{1+\beta(1-\pi)}, \text{ and } K(y_d) \equiv \left[\frac{(Ry + \theta_L - y_d)(\theta_H - \theta_L + y_d)^\beta}{\beta^\beta} \right]^{\frac{1}{1+\beta}}.$$

Set $\tilde{\pi}(y_d) \equiv \min\{\tilde{\pi}_C(y_d), \tilde{\pi}_L(y_d), \tilde{\pi}_H(y_d)\}$. Then for every $\pi < \tilde{\pi}(y_d)$:

$$\hat{b}(\pi) \in \mathcal{B}^R(\pi, y_d), \quad \bar{V}_{ND}(\pi, y_d) = V^R(\hat{b}(\pi); \pi, y_d).$$

Proof. We verify the three conditions for $\hat{b}(\pi) \in \mathcal{B}^R(\pi, y_d)$.

Region condition ($\tilde{\pi}_C$). $\hat{b}(\pi) \leq \theta_H - y_d$ iff $\pi \leq \tilde{\pi}_C$. The lower bound $\hat{b}(\pi) > b^S$ holds for all $\pi \in [0, 1)$ since $b_{FD}^*(\theta_H) = \hat{b}(0) > b^S$ when $y_d \in Y_d$, and $\hat{b}'(\pi) = \frac{\beta(\theta_H - \beta Ry)}{[1+\beta(1-\pi)]^2} > 0$.

L-type IC ($\tilde{\pi}_L$). $(1-\pi)\hat{b}(\pi) \geq b^S$ iff $\pi \leq \tilde{\pi}_L$; positivity of $\tilde{\pi}_L$ follows from $y_d \in Y_d$.

H-type IC ($\tilde{\pi}_H$). The H-type IC is $\mathcal{F}_2(\pi) \geq K(y_d)$. At $\pi = 0$: $\mathcal{F}_2(0) = \frac{Ry + \theta_H}{1+\beta}$, which strictly exceeds $K(y_d)$ because $\hat{b}(0) = b_{FD}^*(\theta_H) > b^S$ implies the H-type strictly prefers its unconstrained FD point to safe pooling. Since $\mathcal{F}_2'(\pi) = \frac{\beta Ry - \theta_H}{[1+\beta(1-\pi)]^2} < 0$ (as $\theta_H > \beta Ry$ in the relevant region), $\mathcal{F}_2$ is strictly decreasing, so a unique $\tilde{\pi}_H > 0$ with $\mathcal{F}_2(\tilde{\pi}_H) = K(y_d)$ exists. $\square$

Combined with the main welfare result of Section 2, Proposition 8 establishes the following complete ordering for $\pi < \min\{\tilde{\pi}(y_d), \pi^*(y_d), \varepsilon(y_d)\}$:

$$\underbrace{V^S(\pi, y_d)}_{= \underline{V}_{ND}} < V_{FD}(\pi; y_d) < \underbrace{V^R(\hat{b}(\pi); \pi, y_d)}_{= \bar{V}_{ND}}.$$

Full disclosure is neither the best nor the worst outcome: it dominates the worst pooling equilibrium (safe pooling) but is dominated by the ex-ante optimal pooling equilibrium.

B.2 Alternative source of private information

In this section, we consider an alternative information structure where private information concerns the sovereign's *cost of default* rather than its period-2 endowment. We show that the same pooling logic operates under this alternative information structure: a pooling equilibrium relaxes the constrained type's borrowing limit, and for sufficiently small π the ex-ante best no-disclosure equilibrium dominates full disclosure under an analogue of condition $(\star)$.

Environment. The timing and disclosure structure are unchanged. The modification is that the sovereign's hidden type $\theta \in \{\theta_L, \theta_H\}$, $0 < \theta_L < \theta_H$, now represents the *cost of defaulting*: a higher θ means default is more costly, so the L-type is more willing to default. Period endowments are y_1 (period 1) and y_2 (period 2, publicly known). Budget equations are:

$$c_1 = y_1 + qb, \quad c_2^R = y_2 - b, \quad c_2^D(\theta) = y_2 - \theta.$$

The sovereign defaults whenever $c_2^D(\theta) > c_2^R$, i.e. $b > \theta$. Under full disclosure, the price is $q_{FD}(b, \theta) = 1/R$ if $b \leq \theta$ and 0 otherwise, giving:

$$b_{FD}^*(\theta) = \min \left\{ \frac{y_2 - \beta R y_1}{1 + \beta}, \theta \right\}.$$

Assumption 3 (Interesting region).

$$\theta_L < \frac{y_2 - \beta R y_1}{1 + \beta} \leq \theta_H.$$

This parallels Assumption 1: the L-type is constrained under full disclosure while the H-type is unconstrained. Under Assumption 3, a separating equilibrium does not exist by the same argument as Proposition 1: the L-type's optimal constrained debt is such that $b_{FD}^*(\theta_L) > \theta_L$, which lies in the region where the L-type defaults, making it a profitable deviation from any separating candidate.

Pooling benchmarks. *Safe pooling* sets $b^S = \theta_L$, at which both types repay:

$$\underline{V}_{ND} = \log \left(y_1 + \frac{\theta_L}{R} \right) + \beta \log(y_2 - \theta_L).$$

This is independent of π since both types face the same allocation.

Lemma 10. Under Assumption 3, $V_{\text{FD}}(\pi; y_d) > \underline{V}_{\text{ND}}$ for all $\pi \in (0, 1)$.

Proof. The L-type is indifferent between safe pooling and FD since $b_{\text{FD}}^*(\theta_L) = \theta_L = b^S$. The H-type's FD debt is $\frac{y_2 - \beta R y_1}{1 + \beta} > \theta_L$ by Assumption 3, so the H-type strictly prefers FD. Hence $V_{\text{FD}}(\pi) > \underline{V}_{\text{ND}}$. $\square$

Risky pooling uses $b \in (\theta_L, \theta_H]$, the L-type defaults, and the price is $(1 - \pi)/R$. The ex-ante payoff is:

$$\hat{V}(b; \pi) = \log\left(y_1 + \frac{1 - \pi}{R}b\right) + \beta[\pi \log(y_2 - \theta_L) + (1 - \pi) \log(y_2 - b)].$$

Denoting $\hat{v}_H(b; \pi) \equiv \log(y_1 + \frac{1 - \pi}{R}b) + \beta \log(y_2 - b)$ and $v_H^{\text{SP}} \equiv \log(y_1 + \frac{\theta_L}{R}) + \beta \log(y_2 - \theta_L)$, the L-type IC reduces to $b \geq \frac{\theta_L}{1 - \pi}$ and the feasible risky-pooling set is:

$$\mathcal{B}^R(\pi) \equiv \left\{ b \in (\theta_L, \theta_H] : b \geq \frac{\theta_L}{1 - \pi}, \hat{v}_H(b; \pi) \geq v_H^{\text{SP}} \right\}.$$

Ex-ante optimal risky-pooling debt. Maximizing $\hat{V}(b; \pi)$ over $b \in (\theta_L, \theta_H]$, the interior FOC gives:

$$\hat{b}(\pi) = \frac{y_2 - \beta R y_1}{1 + \beta(1 - \pi)}. \quad (16)$$

Define $\bar{V}_{\text{ND}}(\pi) \equiv \hat{V}(\hat{b}(\pi); \pi)$.

Proposition 9. Under Assumption 3, define:

$$\begin{aligned} \tilde{\pi}_C &\equiv \frac{(1 + \beta)\theta_H - y_2 + \beta R y_1}{\beta \theta_H}, \\ \tilde{\pi}_L &\equiv \frac{y_2 - \beta R y_1 - (1 + \beta)\theta_L}{y_2 - \beta R y_1 - \beta \theta_L}, \\ \tilde{\pi}_H &\equiv \text{the unique root of } \mathcal{F}_3(\pi) = \left[\frac{(R y_1 + \theta_L)(y_2 - \theta_L)^\beta}{\beta^\beta} \right]^{\frac{1}{1 + \beta}}, \end{aligned}$$

where $\mathcal{F}_3(\pi) \equiv \frac{R y_1 + (1 - \pi)y_2}{1 + \beta(1 - \pi)}$. Set $\tilde{\pi} \equiv \min\{\tilde{\pi}_C, \tilde{\pi}_L, \tilde{\pi}_H\}$. Then for every $\pi < \tilde{\pi}$, $\hat{b}(\pi) \in \mathcal{B}^R(\pi)$.

Proof. Region $(\tilde{\pi}_C)$. $\hat{b}(\pi) \leq \theta_H$ iff $\pi \leq \tilde{\pi}_C$. The lower bound $\hat{b}(\pi) > \theta_L$ follows from Assumption 3 at $\pi = 0$ and $\hat{b}'(\pi) = \frac{\beta(y_2 - \beta R y_1)}{[1 + \beta(1 - \pi)]^2} > 0$.

L-type IC ($\tilde{\pi}_L$). $(1 - \pi)\hat{b}(\pi) \geq \theta_L$ iff $\pi \leq \tilde{\pi}_L$; positivity follows from Assumption 3 since $y_2 - \beta Ry_1 > (1 + \beta)\theta_L$.

H-type IC ($\tilde{\pi}_H$). At $\pi = 0$, $\hat{b}(0) = \frac{y_2 - \beta Ry_1}{1 + \beta}$ is the H-type's unconstrained optimum, so the H-type strictly prefers $\hat{b}(0)$ to safe pooling θ_L . Since $\mathcal{F}'_3(\pi) = \frac{\beta Ry_1 - y_2}{[1 + \beta(1 - \pi)]^2} < 0$ whenever $y_2 > \beta Ry_1$, a unique $\tilde{\pi}_H > 0$ exists by continuity. $\square$

Welfare comparison. Define $\Delta(\pi) \equiv \bar{V}_{ND}(\pi) - V_{FD}(\pi)$. Since $\hat{b}(0) = b_{FD}^*(\theta_H)$, we have $\Delta(0) = 0$, and the sign near zero is governed by $\Delta'(0)$.

Proposition 10. *Suppose:*

$$\log\left(\frac{Ry_1 + y_2}{(1 + \beta)(Ry_1 + \theta_L)}\right) > \frac{y_2 - \beta Ry_1}{Ry_1 + y_2}. \quad (**)$$

Then there exists $\varepsilon > 0$ such that $\bar{V}_{ND}(\pi) > V_{FD}(\pi)$ for all $\pi \in (0, \varepsilon)$.

Proof. Using the envelope theorem at $\pi = 0$:

$$\Delta'(0) = \log\left(\frac{Ry_1 + y_2}{(1 + \beta)(Ry_1 + \theta_L)}\right) - \frac{y_2 - \beta Ry_1}{Ry_1 + y_2},$$

which is positive under (**). The conclusion follows by continuity of Δ . $\square$

The economic mechanism is identical to the baseline: near $\pi = 0$, lenders assign almost no weight to the low-default-cost type, the pooling price approaches $1/R$, and the insurance gain from relaxing the L-type's borrowing constraint dominates the pricing distortion borne by the H-type. Condition (**) is the analogue of condition (*) in this environment and is satisfied when the endowment gap $y_2 - \theta_L$ is large relative to y_1 .

C Empirical supplement

C.1 Robustness: adoption regressions under alternative thresholds $\hat{\pi}$

Table 4 replicates the adoption regressions of Table 2 for four values of the recession-probability threshold $\hat{\pi} \in \{0.15, 0.17, 0.20, 0.35\}$. The benchmark uses $\hat{\pi} = 0.15$. Each column pair reports the OLS coefficient (LPM) with HC1 standard errors and the log-odds coefficient (Logit) on $\mathbf{1}\{\pi_{it}^{pre} > \hat{\pi}\}$; all other controls and fixed effects are unchanged. Two patterns emerge. First, the SDDS and pooled results are statistically significant at $\hat{\pi} = 0.15$: the SDDS LPM coefficient is positive and significant (**, $p < 0.05$) and the pooled

LPM coefficient is positive and significant (**, $p < 0.05$). Second, at $\hat{\pi} \geq 0.17$ the evidence is weaker and sign reversals appear in the per-margin GDDS and eGDDS+NSDP samples, which have few events (24 and 27, respectively). We interpret the benchmark threshold $\hat{\pi} = 0.15$ as the informative specification: at higher thresholds, the indicator variable covers progressively fewer observations, the regressor loses variation, and the per-margin estimates are very imprecise. The sign reversals at high thresholds in the small per-margin samples should not be over-interpreted.

Table 4: Adoption regressions: robustness to alternative thresholds $\hat{\pi}$

Sample	$\hat{\pi} = 0.15$		$\hat{\pi} = 0.17$		$\hat{\pi} = 0.20$		$\hat{\pi} = 0.35$	
	LPM	Logit	LPM	Logit	LPM	Logit	LPM	Logit
SDDS	0.027** (0.013)	1.272 (1.126)	0.016 (0.012)	1.022 (0.803)	0.008 (0.012)	0.624 (0.616)	0.001 (0.013)	0.411 (0.526)
N = 929, events = 24								
GDDS	0.028 (0.075)	0.713 (1.132)	-0.061 (0.077)	0.014 (0.728)	-0.040 (0.072)	0.089 (0.642)	-0.086 (0.054)	-0.706 (0.600)
N = 207, events = 24								
eGDDS+NSDP	0.001 (0.087)	0.607 (0.875)	0.029 (0.082)	1.023 (0.782)	-0.051 (0.068)	0.252 (0.641)	-0.149** (0.062)	-0.400 (0.525)
N = 163, events = 27								
Pooled	0.040** (0.018)	1.087* (0.557)	0.026 (0.017)	0.684* (0.401)	0.011 (0.017)	0.404 (0.319)	-0.018 (0.015)	-0.095 (0.289)
N = 1299, events = 75								

Notes: Each column pair corresponds to a different threshold $\hat{\pi}$ applied to the WEO-implied recession probability π_{it}^{pre} . LPM is the OLS coefficient on $1\{\pi_{it}^{pre} > \hat{\pi}\}$ with HC1 standard errors (parentheses); Logit reports the log-odds coefficient from a logit specification. All columns include year fixed effects and controls for log GDP, real GDP growth, external debt-to-GDP, reserves-to-external-debt, and tax revenue-to-GDP. Default and restructuring spells excluded (Asonuma and Trebesch 2016). N and event counts are identical across threshold columns because sample restrictions are held fixed. The benchmark specification in Table 2 uses $\hat{\pi} = 0.15$. † denotes quasi-complete separation in the logit (standard error > 100); the LPM estimate for that cell remains valid. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.2 Adoption regressions: continuous specification

Table 5 replaces the threshold indicator $1\{\pi_{it}^{pre} > \hat{\pi}\}$ with the continuous recession probability π_{it}^{pre} itself. The model's comparative static predicts $\gamma > 0$ under both specifications, but the two need not agree: a threshold crossing matters theoretically (the sovereign switches from ND to FD at π^*), so the relationship may be non-linear and the continuous coefficient can be near zero even when the threshold effect is strong.

The results in Table 5 are consistent with this interpretation. For SDDS, the LPM coefficient is positive (0.018) but imprecisely estimated. The GDDS and eGDDS+NSDP per-margin samples show negative and in some cases significant coefficients, which combined with the positive and significant threshold coefficients in Table 2 suggests the relationship is concentrated around the threshold rather than monotonically increasing across the full support of π^{pre} . The pooled coefficient is also negative but small and insignificant. We therefore prefer the threshold specification as the primary test, with the continuous specification serving to diagnose functional form.

Table 5: Adoption regressions: continuous recession-probability regressor π_{it}^{pre}

Sample	LPM (5)	Logit (6)	N	Events
SDDS	0.018 (0.044)	1.658 (2.022)	929	24
GDDS	−0.383* (0.228)	−3.255 (2.402)	207	24
eGDDS+NSDP	−0.660*** (0.246)	−2.550 (2.087)	163	27
Pooled	−0.053 (0.057)	−0.345 (1.062)	1,299	75

Notes: Dependent variable: $\text{Adopt}_{it} = 1$ in the year of first subscription to the relevant tier. Regressor: π_{it}^{pre} , the continuous WEO-implied probability of real GDP growth below $\bar{g} = 1$ pp, constructed from the fall WEO vintage in year $t - 1$. This specification replaces the threshold indicator $1\{\pi_{it}^{\text{pre}} > \hat{\pi}\}$ used in Table 2 with the underlying probability directly. LPM includes year fixed effects with HC1 standard errors (parentheses); Logit reports log-odds coefficients. Both include controls for log GDP, real GDP growth, external debt-to-GDP, reserves-to-external-debt, and tax revenue-to-GDP. Default and restructuring spells excluded (Asonuma and Trebesch 2016). The pooled sample stacks all three tier frontiers. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

C.3 Residualized-issuance distribution plots

Figures 9 and 10 show the kernel density of residualized log issuance/GDP for transparent and opaque country-years, estimated using the World Bank International Debt Statistics (IDS) annual panel (93 low- and middle-income countries). These figures are the distributional counterpart of the variance ratios reported in Table 3 of Section 3.2. The GDDS and eGDDS+NSDP margins show the transparent distribution shifted toward the tails relative to the opaque distribution, consistent with the model’s prediction.




Figure 9: Kernel density of residualized log issuance/GDP for transparent (early adopters) and opaque (eligible non-adopters) country-years, by transparency frontier (SDDS 1996, GDDS 2000, eGDDS+NSDP 2015). IDS annual data, 93 countries. Countries in default or restructuring are excluded.




Figure 10: Pooled kernel density of residualized log issuance/GDP across all three frontier episodes. IDS annual data. Transparent obs: 181; opaque obs: 343. Variance ratio (T/O) = 1.28.

C.4 Dealogic quarterly-panel robustness

This section replicates the variance-ratio analysis of Section 3.2 using Dealogic data on USD-denominated sovereign bond tranches, aggregated to country-quarter frequency, covering 81 countries from 1990 to 2025. The outcome is $w_{iq} = \log(1 + \text{Issuance}_{iq}/\text{GDP}_{iq})$ (we use w to avoid a clash with the model endowment y). We project it on country fixed effects, quarter fixed effects, and macro controls available at quarterly frequency:

$$w_{iq} = \alpha_i + \lambda_q + \beta_1 \log \text{GDP}_{iq} + \beta_2 \Delta \ln \text{GDP}_{iq} + u_{iq}. \quad (17)$$

The frontier-cohort classification follows Section 3.2 exactly, with country-quarter (rather than country-year) observations. Table 6 reports the variance ratios. The GDDS margin again has the highest ratio (2.67) and the pooled ratio (1.11) exceeds one, consistent with the IDS results.

Table 6: Variance ratio of residualized issuances: Transparent vs. Opaque country-quarters (Dealogic)

Transparency tier	N_T	N_O	$\text{Var}(\hat{u} T)/\text{Var}(\hat{u} O)$	Predicted >1?
SDDS (1996)	302	538	0.20	No
GDDS (2000)	140	436	2.67	Yes
eGDDS+NSDP (2015)	228	492	0.84	No
Pooled	877	2,435	1.11	Yes

Notes: Dealogic data, USD-denominated sovereign bond tranches aggregated to country-quarter, 81 countries, 1990–2025. N_T (N_O) counts country-quarter observations classified as Transparent (Opaque) in the regression sample. Residuals $\hat{u}_{i,q}$ from equation (17). SDDS Plus excluded. Countries with no positive issuance in the relevant window are dropped before residualizing; each tier row applies this filter within its own five-year window $[T_r, T_r + 5]$, while the pooled row stacks all three frontier windows and drops countries with no positive issuance across *any* of the three windows. As a result, the pooled observation counts ($N_T = 877$, $N_O = 2,435$) exceed the sum of the three tier-specific rows ($N_T = 670$, $N_O = 1,466$).

C.5 Durability sample: all 186 countries

Table 1 covers all 186 countries with records on the IMF Dissemination Standards Bulletin Board (DSBB). Their current tier as of the data extraction date is shown: GDDS/eGDDS (109 countries), SDDS (48 countries), or SDDS Plus (29 countries). Countries without a subscription appear under GDDS/eGDDS because they are enrolled in the enhanced GDDS by default.

Table 7: DSBB sample: countries A–K (1 of 2)

ISO	Country	Tier	ISO	Country	Tier	ISO	Country	Tier
ABW	Aruba	GDDS/eGDDS	CHL	Chile	SDDS+	GIN	Guinea	GDDS/eGDDS
AGO	Angola	GDDS/eGDDS	CHN	China	SDDS	GMB	Gambia	GDDS/eGDDS
ALB	Albania	GDDS/eGDDS	CIV	Côte d'Ivoire	GDDS/eGDDS	GNB	Guinea-Bissau	GDDS/eGDDS
AND	Andorra	GDDS/eGDDS	CMR	Cameroon	GDDS/eGDDS	GNQ	Equatorial Guinea	GDDS/eGDDS
ARE	UAE	GDDS/eGDDS	COG	Congo	GDDS/eGDDS	GRC	Greece	SDDS
ARG	Argentina	SDDS	COK	Cook Islands	GDDS/eGDDS	GRD	Grenada	GDDS/eGDDS
ARM	Armenia	SDDS	COL	Colombia	SDDS	GTM	Guatemala	GDDS/eGDDS
ATG	Antigua and Barbuda	GDDS/eGDDS	COM	Comoros	GDDS/eGDDS	GUY	Guyana	GDDS/eGDDS
AUS	Australia	SDDS	CPV	Cabo Verde	GDDS/eGDDS	HKG	Hong Kong SAR	SDDS
AUT	Austria	SDDS+	CRI	Costa Rica	SDDS	HND	Honduras	GDDS/eGDDS
AZE	Azerbaijan	GDDS/eGDDS	CYP	Cyprus	SDDS	HRV	Croatia	SDDS
BDI	Burundi	GDDS/eGDDS	CZE	Czech Republic	SDDS+	HTI	Haiti	GDDS/eGDDS
BEL	Belgium	SDDS+	DEU	Germany	SDDS+	HUN	Hungary	SDDS+
BEN	Benin	GDDS/eGDDS	DJI	Djibouti	GDDS/eGDDS	IDN	Indonesia	SDDS
BFA	Burkina Faso	GDDS/eGDDS	DMA	Dominica	GDDS/eGDDS	IND	India	SDDS
BGD	Bangladesh	GDDS/eGDDS	DNK	Denmark	SDDS+	IRL	Ireland	SDDS
BGR	Bulgaria	SDDS+	DOM	Dominican Rep.	GDDS/eGDDS	IRN	Iran	GDDS/eGDDS
BHR	Bahrain	GDDS/eGDDS	DZA	Algeria	GDDS/eGDDS	IRQ	Iraq	GDDS/eGDDS
BHS	Bahamas	GDDS/eGDDS	ECU	Ecuador	SDDS	ISL	Iceland	SDDS
BIH	Bosnia & Herz.	SDDS	EGY	Egypt	SDDS	ISR	Israel	SDDS+
BLR	Belarus	SDDS	ESP	Spain	SDDS+	ITA	Italy	SDDS+
BLZ	Belize	GDDS/eGDDS	EST	Estonia	SDDS+	JAM	Jamaica	GDDS/eGDDS
BOL	Bolivia	GDDS/eGDDS	ETH	Ethiopia	GDDS/eGDDS	JOR	Jordan	SDDS
BRA	Brazil	SDDS+	FIN	Finland	SDDS+	JPN	Japan	SDDS+
BRB	Barbados	GDDS/eGDDS	FJI	Fiji	GDDS/eGDDS	KAZ	Kazakhstan	SDDS
BRN	Brunei Darussalam	GDDS/eGDDS	FRA	France	SDDS+	KEN	Kenya	GDDS/eGDDS
BTN	Bhutan	GDDS/eGDDS	FSM	Micronesia	GDDS/eGDDS	KGZ	Kyrgyzstan	SDDS
BWA	Botswana	GDDS/eGDDS	GAB	Gabon	GDDS/eGDDS	KHM	Cambodia	GDDS/eGDDS
CAF	Cent. African Rep.	GDDS/eGDDS	GBR	United Kingdom	SDDS+	KIR	Kiribati	GDDS/eGDDS
CAN	Canada	SDDS+	GEO	Georgia	SDDS	KNA	St. Kitts & Nevis	GDDS/eGDDS
CHE	Switzerland	SDDS+	GHA	Ghana	GDDS/eGDDS	KOR	Korea	SDDS

Table 8: DSBB sample: countries K–Z (2 of 2)

ISO	Country	Tier	ISO	Country	Tier	ISO	Country	Tier
KWT	Kuwait	GDDS/eGDDS	NIC	Nicaragua	GDDS/eGDDS	SVK	Slovak Republic	SDDS+
LAO	Lao PDR	GDDS/eGDDS	NLD	Netherlands	SDDS+	SVN	Slovenia	SDDS+
LBN	Lebanon	GDDS/eGDDS	NOR	Norway	SDDS	SWE	Sweden	SDDS+
LBR	Liberia	GDDS/eGDDS	NPL	Nepal	GDDS/eGDDS	SWZ	Eswatini	GDDS/eGDDS
LBY	Libya	GDDS/eGDDS	OMN	Oman	GDDS/eGDDS	SYC	Seychelles	SDDS
LCA	St. Lucia	GDDS/eGDDS	PAK	Pakistan	GDDS/eGDDS	SYR	Syria	GDDS/eGDDS
LKA	Sri Lanka	SDDS	PAN	Panama	GDDS/eGDDS	TCD	Chad	GDDS/eGDDS
LSO	Lesotho	GDDS/eGDDS	PER	Peru	SDDS	TGO	Togo	GDDS/eGDDS
LTU	Lithuania	SDDS+	PHL	Philippines	SDDS	THA	Thailand	SDDS
LUX	Luxembourg	SDDS+	PLW	Palau	GDDS/eGDDS	TJK	Tajikistan	GDDS/eGDDS
LVA	Latvia	SDDS+	PNG	Papua New Guinea	GDDS/eGDDS	TLS	Timor-Leste	GDDS/eGDDS
MAR	Morocco	SDDS	POL	Poland	SDDS	TON	Tonga	GDDS/eGDDS
MDA	Moldova	SDDS	PRT	Portugal	SDDS+	TTO	Trinidad & Tobago	GDDS/eGDDS
MDG	Madagascar	GDDS/eGDDS	PRY	Paraguay	SDDS	TUN	Tunisia	SDDS
MDV	Maldives	GDDS/eGDDS	PSE	West Bank & Gaza	SDDS	TUR	Türkiye	SDDS
MEX	Mexico	SDDS	QAT	Qatar	GDDS/eGDDS	TUV	Tuvalu	GDDS/eGDDS
MHL	Marshall Islands	GDDS/eGDDS	ROU	Romania	SDDS+	TZA	Tanzania	GDDS/eGDDS
MKD	North Macedonia	SDDS+	RUS	Russia	SDDS	UGA	Uganda	GDDS/eGDDS
MLI	Mali	GDDS/eGDDS	RWA	Rwanda	GDDS/eGDDS	UKR	Ukraine	SDDS
MLT	Malta	SDDS+	SAU	Saudi Arabia	SDDS	URY	Uruguay	SDDS
MMR	Myanmar	GDDS/eGDDS	SDN	Sudan	GDDS/eGDDS	USA	United States	SDDS+
MNE	Montenegro	GDDS/eGDDS	SEN	Senegal	SDDS	UZB	Uzbekistan	GDDS/eGDDS
MNG	Mongolia	SDDS	SGP	Singapore	SDDS	VCT	St. Vincent	GDDS/eGDDS
MOZ	Mozambique	GDDS/eGDDS	SLB	Solomon Islands	GDDS/eGDDS	VEN	Venezuela	GDDS/eGDDS
MRT	Mauritania	GDDS/eGDDS	SLE	Sierra Leone	GDDS/eGDDS	VNM	Vietnam	GDDS/eGDDS
MUS	Mauritius	SDDS+	SLV	El Salvador	SDDS	VUT	Vanuatu	GDDS/eGDDS
MWI	Malawi	GDDS/eGDDS	SMR	San Marino	GDDS/eGDDS	WSM	Samoa	GDDS/eGDDS
MYS	Malaysia	SDDS	SOM	Somalia	GDDS/eGDDS	XKX	Kosovo	GDDS/eGDDS
NAM	Namibia	SDDS	SRB	Serbia	GDDS/eGDDS	YEM	Yemen	GDDS/eGDDS
NER	Niger	GDDS/eGDDS	STP	São Tomé	GDDS/eGDDS	ZAF	South Africa	SDDS
NGA	Nigeria	GDDS/eGDDS	SUR	Suriname	GDDS/eGDDS	ZMB	Zambia	GDDS/eGDDS

C.6 Regression sample countries

The adoption regressions in Table 2 draw on three separate feasible choice sets, one per transparency frontier. Each set consists of IDS panel countries (92 low- and middle-income economies) whose eligibility clock was active at the relevant launch year and that had not pre-adopted. High-income SDDS-Plus subscribers (e.g., Australia, Germany, Japan) are excluded at the data stage. The regression sample includes, for each country, all years in which the country had not yet adopted (the pre-adoption spell) plus the first year of adoption; observations during default or restructuring spells ([Asonuma and Trebesch 2016](#)) and observations with missing controls are dropped. Countries that adopted in the launch year itself appear only as the adoption event. Adoption years are shown in parentheses; countries without a year did not adopt the relevant tier within the sample period.

SDDS sample (92 countries, launch year 1996).

ISO	Country	Year	ISO	Country	Year
AFG	Afghanistan	—	LBN	Lebanon	—
AGO	Angola	—	LBR	Liberia	—
ALB	Albania	—	LKA	Sri Lanka	2015
ARG	Argentina	1996	LSO	Lesotho	—
ARM	Armenia	2003	MAR	Morocco	2005
AZE	Azerbaijan	—	MDA	Moldova	2006
BDI	Burundi	—	MDG	Madagascar	—
BEN	Benin	—	MEX	Mexico	1996
BFA	Burkina Faso	—	MLI	Mali	—
BGD	Bangladesh	—	MMR	Myanmar	—
BIH	Bosnia & Herz.	2024	MNG	Mongolia	2019
BLR	Belarus	2004	MOZ	Mozambique	—
BOL	Bolivia	—	MRT	Mauritania	—
BRA	Brazil	—	MWI	Malawi	—
BWA	Botswana	—	NER	Niger	—
CAF	Cent. African Rep.	—	NGA	Nigeria	—
CHN	China	2015	NIC	Nicaragua	—
CIV	Côte d’Ivoire	—	NPL	Nepal	—
CMR	Cameroon	—	PAK	Pakistan	—
COD	Congo, Dem. Rep.	—	PER	Peru	1996
COG	Congo	—	PHL	Philippines	1996
COL	Colombia	1996	PNG	Papua New Guinea	—
DOM	Dominican Republic	—	PRY	Paraguay	2024
DZA	Algeria	—	RWA	Rwanda	—
ECU	Ecuador	1998	SDN	Sudan	—
EGY	Egypt	2005	SEN	Senegal	2017
ETH	Ethiopia	—	SLE	Sierra Leone	—
GAB	Gabon	—	SLV	El Salvador	1998
GEO	Georgia	2010	SOM	Somalia	—
GHA	Ghana	—	SRB	Serbia	—
GIN	Guinea	—	TCD	Chad	—
GMB	Gambia	—	TGO	Togo	—
GTM	Guatemala	—	THA	Thailand	1996
HND	Honduras	—	TJK	Tajikistan	—
HTI	Haiti	—	TKM	Turkmenistan	—
IDN	Indonesia	1996	TUN	Tunisia	2001
IND	India	1996	TUR	Türkiye	1996
IRN	Iran	—	TZA	Tanzania	—
IRQ	Iraq	—	UGA	Uganda	—
JAM	Jamaica	—	UKR	Ukraine	2003
JOR	Jordan	2010	UZB	Uzbekistan	—
KAZ	Kazakhstan	2003	VNM	Viet Nam	—
KEN	Kenya	—	YEM	Yemen	—
KGZ	Kyrgyzstan	2004 ⁶²	ZAF	South Africa	1996
KHM	Cambodia	—	ZMB	Zambia	—
LAO	Lao PDR	—	ZWE	Zimbabwe	—

GDDS sample (63 countries, launch year 2000). Eligible countries are IDS panel countries that had not adopted SDDS as of 2000 and had not pre-adopted GDDS.

ISO	Country	Year	ISO	Country	Year
AFG	Afghanistan	—	LAO	Lao PDR	2018
AGO	Angola	2004	LBN	Lebanon	2003
ALB	Albania	2000	LBR	Liberia	2005
AZE	Azerbaijan	2001	LSO	Lesotho	2003
BDI	Burundi	2011	MDG	Madagascar	2004
BEN	Benin	2001	MLI	Mali	2001
BFA	Burkina Faso	2001	MMR	Myanmar	2013
BGD	Bangladesh	2001	MOZ	Mozambique	2003
BOL	Bolivia	2000	MRT	Mauritania	2004
BRA	Brazil	—	MWI	Malawi	2002
BWA	Botswana	2002	NER	Niger	2002
CAF	Cent. African Rep.	2004	NGA	Nigeria	2003
CIV	Côte d'Ivoire	2000	NIC	Nicaragua	2005
CMR	Cameroon	2000	NPL	Nepal	2001
COD	Congo, Dem. Rep.	—	PAK	Pakistan	2003
COG	Congo	2004	PNG	Papua New Guinea	2012
DOM	Dominican Republic	2005	RWA	Rwanda	2003
DZA	Algeria	2009	SDN	Sudan	2003
ETH	Ethiopia	2002	SLE	Sierra Leone	2003
GAB	Gabon	2002	SOM	Somalia	2022
GHA	Ghana	2005	SRB	Serbia	2009
GIN	Guinea	2003	TCD	Chad	2002
GMB	Gambia	2000	TGO	Togo	2001
GTM	Guatemala	2004	TJK	Tajikistan	2004
HND	Honduras	2005	TKM	Turkmenistan	—
HTI	Haiti	2009	TZA	Tanzania	2001
IRN	Iran	2012	UGA	Uganda	2000
IRQ	Iraq	2009	UZB	Uzbekistan	2018
JAM	Jamaica	2003	VNM	Viet Nam	2003
KEN	Kenya	2002	YEM	Yemen	2001
KHM	Cambodia	2002	ZMB	Zambia	2002
ZWE	Zimbabwe	—			

eGDDS+NSDP sample (59 countries, launch year 2015). Eligible countries are IDS panel countries that had adopted GDDS/eGDDS but not SDDS as of 2015 and had not pre-adopted the eGDDS+NSDP tier.

ISO	Country	Year	ISO	Country	Year
AGO	Angola	2018	LBN	Lebanon	—
ALB	Albania	2017	LBR	Liberia	—
AZE	Azerbaijan	2019	LSO	Lesotho	2016
BDI	Burundi	2023	MDG	Madagascar	2019
BEN	Benin	2017	MLI	Mali	—
BFA	Burkina Faso	2018	MMR	Myanmar	2019
BGD	Bangladesh	2017	MOZ	Mozambique	2019
BOL	Bolivia	—	MRT	Mauritania	2019
BWA	Botswana	2016	MWI	Malawi	2016
CAF	Cent. African Rep.	2025	NER	Niger	—
CIV	Côte d'Ivoire	2017	NGA	Nigeria	2016
CMR	Cameroon	2017	NIC	Nicaragua	—
COD	Congo, Dem. Rep.	—	NPL	Nepal	2017
COG	Congo	2025	PAK	Pakistan	2019
DOM	Dominican Republic	2019	PNG	Papua New Guinea	—
DZA	Algeria	—	RWA	Rwanda	2017
ETH	Ethiopia	2019	SDN	Sudan	—
GAB	Gabon	2019	SLE	Sierra Leone	2016
GHA	Ghana	2018	SOM	Somalia	2022
GIN	Guinea	2019	SRB	Serbia	2018
GMB	Gambia	2018	TCD	Chad	2022
GTM	Guatemala	2020	TGO	Togo	2018
HND	Honduras	2016	TJK	Tajikistan	2020
HTI	Haiti	—	TZA	Tanzania	—
IRN	Iran	—	UGA	Uganda	2016
IRQ	Iraq	—	UZB	Uzbekistan	2018
JAM	Jamaica	2017	VNM	Viet Nam	2019
KEN	Kenya	2018	YEM	Yemen	—
KHM	Cambodia	2018	ZMB	Zambia	2016
LAO	Lao PDR	2018			