The Economics of Comparability of Treatment in Sovereign Debt Restructurings *

Cristián Cuevas
Universidad de los Andes (Chile)

Martín Guzmán
Columbia University

Joseph E. Stiglitz
Columbia University

October 2026

Abstract

A central source of tension in sovereign debt restructurings concerns the allocation of debt relief across creditors. Each creditor group seeks to minimize its own contribution to debt reduction while insisting that the losses incurred are at least "comparable" to those borne by others. This challenge is compounded by the heterogeneous nature of the creditor base (in instruments, seniorities, and objectives). Despite its policy relevance, the concept of comparable treatment (CoT) lacks a clear foundation, often contributing to protracted and inefficient restructurings. This paper develops an economic framework to define comparability, focusing on creditor heterogeneity, and the allocation criteria that underlie alternative loss-allocation rules. We distinguish two creditor-facing notions of comparability (equal treatment of claims and origination-adjusted treatment) from a debtor-welfare criterion that allocates debt relief so as to minimize the welfare cost to the debtor of obtaining it. We also show that the metric used to assess comparability matters in itself: valuing restructured claims at endogenous exit yields turns the equal-treatment assessment into a fixed-point problem that, in our environment, can admit two self-fulfilling assessments or none, whereas risk-free valuation always delivers a unique one. Finally, we assess how prevailing CoT practices employed by policy institutions can be rationalized -if at all- within this framework. The analysis is primarily positive: we clarify the conceptual and analytical foundations of CoT rather than provide prescriptive reforms.

Keywords: debt restructuring, debt relief, creditor heterogeneity

JEL classification: F34, G15, H63.

*For useful comments and suggestions, we thank participants at Debtcon 8, the 2026 Spring Midwest Macro Conference, and the European Economic Association Annual Meeting.

Corresponding author: ccuevas@uandes.cl

1 Introduction

Sovereign debt restructurings involve negotiations with multiple heterogeneous creditors that hold different instruments, lending models and motives, among other dimensions. There are two key conflicts associated with sovereign debt restructurings: first, the definition of the level of debt that is sustainable, with both distributional (between debtors and creditors) as well as efficiency consequences (Krugman, 1988; Guzman and Stiglitz, 2020); second, the distribution of relief among creditors—a matter known as "comparability of treatment" (CoT) among creditors. In practice, the issue of CoT arises from each creditor's demand that the sovereign debtor does not treat her worse than other creditors—collectively, leading to a need for equitable treatment, which does not necessarily mean equal treatment. No creditor wants the others to receive better (i.e., more lenient) debt restructuring terms, and each fears contributing a disproportionately large share of the needed debt relief. In this context, the lack of a clear definition of comparability often results in stalemates and protracted negotiations.

This paper focuses on the second conflict. Comparability is poorly defined, in theory and in practice, leading to inequitable burden-sharing or creating ambiguities that delay restructuring processes, entailing efficiency losses. The concept is largely absent from economics: to our knowledge, no paper defines comparability in economic terms or derives, in equilibrium, the allocation that a comparability criterion implies. We provide a relatively simple framework that allows us to characterize the CoT problem and obtain a set of analytical results. This contrasts with much of the recent literature, which tends to rely on numerical simulations. We do three things.

First, we build a three-period model in which restructuring is triggered by a debt sustainability event, not by strategic default: total relief is fixed by a debt sustainability analysis, and the CoT question is who provides it. Creditors differ in the terms on which they lent and in what their relationship is worth to the debtor. Private investors price restructuring risk into their claims. An official lender, resembling a multilateral development bank, lends on concessional terms that do not respond to that risk and stands ready to lend in a later liquidity crisis, in which the government cannot borrow from private investors, with a capacity that falls with the haircut on its claims.[1]

Second, we characterize the allocations implied by three criteria that differ in what they compare and for whose benefit. They rest on three objects disciplined by the model, each requiring the assessor to observe more: contracts, then prices, then the crisis-lending technology. Two criteria are creditor-facing. Equal treatment of claims (Criterion 1) compares haircuts: the proportional rule, which imposes equal haircuts, is the only continuous rule that sees claims

[1] We add bilateral official creditors in the Supplemental Appendix.

alone, treats equal claims equally, and cannot be gamed by repackaging claims across accounts. Origination-adjusted treatment (Criterion 2) nets out the compensation each class received, or the concession it granted, when it lent. The private spread pre-pays the expected loss, while concessional official terms are a subsidy, so under equal haircuts private creditors exit the restructuring better off than official ones, by exactly the compensation differential per unit of face value. Because the spread prices the rule's own anticipated treatment, the no-windfall allocation that nets this differential is a fixed point, which we show is unique and lies between equal treatment and preferred creditor status.[2]

The debtor-cost criterion (Criterion 3) is of a different kind: it measures no creditor's position and sources relief where it costs the debtor least. With total relief fixed, the allocation moves only the price at which private claims are issued and the crisis lending that survives. With commitment, it breaches officials fully when post-restructuring crises are unlikely or the overhang is too deep to preserve any crisis lending, gives them priority when such crises are likely enough, and may impose graduated participation in between. Without commitment, as long as protecting officials preserves some crisis lending, the debtor itself gives them priority at the restructuring date: preferred creditor status is the ex-post choice of the debtor. Criterion 2 thus favors a lender for compensation it forwent at origination, Criterion 3 for financing whose supply depends on its treatment, and the two can rank the same lenders in opposite orders.

Confronted with the practice and implementation of CoT, the framework interprets the Paris Club's equal present-value-haircuts standard as Criterion 1 in principle, the unconditional preferred creditor status of multilateral lenders as the allocation chosen to minimize the debtor's welfare cost of relief, and the proposal to rank creditors by the concessionality of their terms as Criterion 2.

Third, we show that the metric used to assess comparability matters in itself. The Paris Club discounts all claims at a common rate; market-value assessments value the new claims of each class at its own exit yield, which prices the risk that they are not honored. At the risk-free rate, or at a rate with a common premium not based on expected repayment, the equal-treatment allocation exists, is unique, and does not depend on beliefs about exit yields. At class-specific exit yields, which depend on the restructuring terms, the same rule becomes a fixed point between expected repayment, measured haircuts and the allocation of relief. On some open sets of parameters it has two solutions, the pessimistic one self-fulfilling and certifying equal treatment while officials grant more relief than equal treatment requires; on others it has none. The market-value haircut remains the right measure of expected investor losses; what fails is its use as the input to a rule that determines the claims it values. These are possibility results for the rule and assessment class we analyze, not an impossibility theorem for market-value

[2]The latter refers to the practice of excluding certain creditors from the restructuring.

metrics.

Related literature. A workable definition of CoT is a policy need (Georgieva and Pazarbasioglu, 2021; Gill, 2022), epecially when seen as a source of protracted negotiations.[3] The G20's Common Framework requires the debtor to seek from its other creditors a treatment at least as favorable as the one agreed with participating official bilateral creditors (G20, 2020), yet "comparable" is vague enough to accommodate very different treatments (Bolton, Gulati, and Panizza, 2023). Proposals to operationalize comparability are accounting rules that take valuations as given (Rivetti, 2022; Diwan, Harnoys-Vannier, and Kessler, 2023; Amegashie, 2023).

Work on the 1980s debt crisis estimated and modeled how official and private creditors shared its burden, but proposed no criterion for when their shares are comparable (Demirgüç-Kunt and Fernández-Arias, 1992; Spiegel, 1996). An empirical literature measures the haircuts, computed at exit yields, private creditors take in sovereign restructurings (Sturzenegger and Zettelmeyer, 2008; Cruces and Trebesch, 2013; Fang, Schumacher, and Trebesch, 2021; Graf von Luckner et al., 2025). We ask instead how the relief granted by all creditors, official ones included, can be compared. Moreover, our object of interest is debt relief, understood as debt reduction from the debtor's perspective, a concept distinct from the losses incurred by creditors during a restructuring. The restructurings studied in the papers cited above often affect only a fraction of the total outstanding debts, and they cover private creditors only.

This paper is also related to the literature on creditor heterogeneity. Creditors differ in seniority (Schlegl, Trebesch, and Wright, 2019) and legal jurisdiction (Erce and Mallucci, 2018). In Bolton and Jeanne (2009), selective defaults make hard-to-restructure debt de facto senior, and an enforced priority rule can improve on equal treatment. Powell (2026) formalizes strict comparability as *pari passu* treatment: two creditor types receive equal restructuring haircuts, and examines how this restructuring rule affects debt dilution. In the default models of Dellas and Niepelt (2016) and Wicht (2025), official lenders are costlier to default on, and Wicht compares seniority regimes by borrower welfare. Arellano and Barreto (2024) find that official debt flows in during default episodes and model it as more concessional in default. In Cordella and Powell (2021), a lender that commits to limited lending at the risk-free rate and does not lend into arrears is always repaid. These papers ask how creditor heterogeneity shapes borrowing and default; we take the relief requirement as given and ask how a comparability criterion allocates it.

A separate literature studies collective action clauses and creditor litigation (Fang, Schumacher, and Trebesch, 2021; Schumacher, Trebesch, and Enderlein, 2021) and holdout delay, which equal-repayment rules can lengthen (Pitchford and Wright, 2012); legal proposals would

[3]Longer default episodes entail twice the output losses of quickly resolved defaults (Farah-Yacoub, Graf von Luckner, and Reinhart, 2024).

enforce comparability across creditor groups by contract or statute (Buchheit and Gulati, 2023; Reichert-Facilides, 2023). Our model abstracts from several important real-world features. The aim is to understand the fundamentals of the CoT problem. The framework has no voting, holdout, bargaining, or litigation margin, so it evaluates none of these devices; its commitment problem works through prices, not delay. Rules that span creditor groups presuppose a definition of the treatment to be equalized; our three criteria are such definitions.

The paper proceeds as follows. Section 2 presents the model. Section 3 defines the CoT problem and characterizes the allocation implied by each of the three criteria. Section 4 shows how using exit yields as the metric of a comparability rule creates a fixed-point problem for the assessment. Section 5 asks which criterion, if any, supports the rules used or proposed in practice, and Section 6 concludes. Appendix A contains the proofs of the results stated in the text, and the Supplemental Appendix collects additional results.

2 Framework

The model is deliberately minimal: each ingredient delivers one economic object. There are three periods and two creditor classes, private investors who price default risk and an official lender, resembling a multilateral development bank, who provides crisis support and concessional terms; Supplemental Appendix S.2 adds bilateral official creditors. Of the two shocks, a restructuring event at date 1 creates the allocation problem and, because it is anticipated, the ex-ante pricing margin; a liquidity crisis at date 2 gives the official relationship its value.

2.1 Dates, preferences, and endowments

Time is discrete. There are three payoff dates $t = 0, 1, 2$ and a settlement date $t = 3$. A small open economy ("the government") values consumption streams according to

$$\mathbb{E}_0\big[\, u(c_0) + \beta\, u(c_1) + \beta^2 u(c_2) + \beta^3 c_3 \,\big]\,. \tag{1}$$

The endowment is y_0 at date 0. At date 1 the endowment is $y_1 \in \{y^n, y^c\}$ with $y^c < y^n$: the *crisis* state y^c occurs with exogenous probability $F \in (0,1)$ and triggers a debt restructuring (Assumption 3). At date 2 the endowment is $y_2 \in \{y^h, y^\ell\}$ with $y^\ell < y^h$: the *liquidity-crisis* state y^ℓ occurs with exogenous probability $\pi \in (0,1)$, independently of y_1. The settlement endowment $\bar{y}_3$ is large enough that date-3 consumption is always positive.

Assumption 1 (Preferences and discounting)**.** u is twice continuously differentiable, with $u' > 0$ and $u'' < 0$, and satisfies $\lim_{c\to 0} u'(c) = \infty$; the government's discount factor equals the inverse of the gross world risk-free rate, $\beta = \beta^* \equiv 1/R^*$; and date-3 utility is linear.

Three comments on Assumption 1. First, $\beta = \beta^*$ removes the impatience wedge between the government and its creditors, as in Cordella and Powell (2021). In Dellas and Niepelt (2016) and Wicht (2025) that wedge generates borrowing; here, issue quantities are given (Section 2.2).[4] Second, the linear settlement date is a transferable-utility device in the tradition of Holmström and Tirole (1998): it gives the date-2 emergency loan a repayment date without a fourth round of consumption risk, so that the value of crisis liquidity is governed by date-2 curvature alone. Third, the Inada condition guarantees that crisis liquidity is valuable when crisis consumption is low, the empirically relevant case; no result depends on the specific curvature of u, and the figures use $u = \log$ for simplicity.

2.2 Creditors and predetermined contracts

There are two creditor classes, indexed by $i \in \{P, O\}$: competitive private investors, P, and an official policy lender, O. At date 0 the government issues one-period zero-coupon claims with face values $d_P > 0$ and $d_O > 0$ due at date 1, with $D \equiv d_P + d_O$. The composition (d_P, d_O) is exogenous: the CoT question is how to allocate a *given* relief requirement across an *inherited* creditor structure (the object negotiators face); endogenizing composition is the subject of Dellas and Niepelt (2016), and holding it fixed isolates burden-sharing at date 1. Predetermined quantities do not imply predetermined debt values: the claims are placed at date 0 at prices (q_P from (2), $1/R_O$ from the official terms), and the private price responds to the anticipated allocation. There is thus no issuance margin, but the date-0 price margin, which the allocation rule moves, remains active: it is the pricing channel of Lemma 1 below.

Assumption 2 (Creditor technologies and market access)**.** *(i)* Private investors are deep-pocketed, risk-neutral, and competitive, with outside return R^*; the date-0 price per unit of face value is therefore

$$q_P \;=\; \beta^*\,\mathbb{E}\big[\,\text{repayment per unit of face value}\,\big] \;=\; \beta^*\big(1 - F\,h_P^e\big), \tag{2}$$

where h_P^e is the private haircut investors anticipate in the restructuring state at $t = 1$. *(ii)* The official lender disburses d_O/R_O at date 0 against the claim d_O, at a contractual gross rate R_O that is independent of anticipated treatment and concessional, $R_O \le R^*$. *(iii)* The government's liabilities are the claims above, the restructured claims that replace them at date 1 (Assumption 3 below), and the date-2 emergency loan (Section 2.4); it issues no other claim and holds no asset. Private investors price and may trade the claims it has issued, but advance

[4] The normalization does not, by itself, remove borrowing motives: under uncertainty the government would still wish to trade consumption across dates and states; the absence of a quantity margin comes instead from holding (d_P, d_O) fixed and from the asset menu of Assumption 2(iii).

no new funds after date 0.

Part (i) is the standard small-open-economy pricing condition (Eaton and Gersovitz, 1981). Part (ii) has two components. Treatment-insensitivity is an institutional feature of official lending: multilateral institutions price uniformly across members at rates tied to their own funding costs plus a fixed margin, and their exposures expand rather than reprice when sovereign risk rises.[5] The second, $R_O \leq R^*$, does not follow from treatment-insensitivity: a cost-plus rate, R_O, can lie below the rate at which the sovereign borrows from the market without lying below the risk-free rate, so $R_O \leq R^*$ refers to the concessional terms offered by the official class.

We summarize the official terms by the *concessionality margin* $g_O \equiv R^*/R_O - 1 \geq 0$, a strictly increasing transform of the loan's grant element computed at the risk-free rate, $1 - R_O/R^* = g_O/(1+g_O)$, the standard measure of concessionality (IMF/World Bank; OECD DAC).

Part (iii) is the market-access assumption behind the crisis-lending channel: private investors price and trade the claims the government has issued, but in the date-2 liquidity crisis the government cannot borrow from them against its settlement resources, however certain those are. On the restructuring branch, this resembles a lack of market access, of the kind that follows a restructuring (Gelos, Sahay, and Sandleris, 2011; Cruces and Trebesch, 2013), in which the official sector is the lender that still operates (Arellano and Barreto, 2024); on the repayment branch it only fixes a constant. The model takes this asymmetry as given and studies its implications for burden-sharing; it does not explain its origin.[6] Part (iii) also closes the budgets: from date 1 on, consumption is the endowment net of contractual payments, plus the emergency loan in the date-2 liquidity-crisis state and its settlement at date 3, so the settlement endowment is pledged to that loan alone. We lift assumptions (ii) and (iii) in Supplemental Appendix S.1, where we let private creditors lend in crisis and officials price anticipated treatment.

The date-0 budget constraint is

$$c_0 \;=\; y_0 + q_P\, d_P + d_O/R_O. \tag{3}$$

The gap between private lending (Part (i)) and official lending (Part (ii)) is what makes the allocation problem non-trivial: private losses are priced ex ante, official losses are not.

2.3 The restructuring event and the relief requirement

Assumption 3 (Restructuring trigger and sustainability envelope)**.** In the date-1 restructuring state y^c, all claims are exchanged for one-period claims (m_P, m_O) due at date 2, with no cash

[5] Wicht (2025) documents that multilateral debt *increases* during default episodes.

[6] This would require a model of why sovereign income is pledgeable to an official creditor, with its conditionality and non-toleration of arrears, but not to private lenders in a crisis.

payment at date 1, subject to the sustainability envelope $m_P + m_O \leq \bar{M}$ with $0 < \bar{M} < y^\ell$ and $\beta^* \bar{M} < D$. In the no-crisis state $y^n > D$, claims are repaid in full.

The envelope $\bar{M}$ is the model's one-period analogue of a maximum politically feasible repayment capacity, and the trigger is a debt sustainability event, not a strategic default. This is a design choice, made for two reasons. First, it matches the premise of the paper: CoT is the problem of allocating relief whose *total* is determined by a debt sustainability analysis (DSA), so the total should be a constraint, not a choice. Second, it separates the burden-sharing margin from the default-incentive margin studied by Eaton and Gersovitz (1981) and, in the composition context, Dellas and Niepelt (2016). The condition $\bar{M} < y^\ell$ ensures that restructured claims are repaid in every date-2 state: the exchange restores sustainability *by construction*; Section 4 discusses how re-activating the default margin at date 2 endogenizes exit yields.

Definition 1 (Relief requirement and haircuts)**.** *In the date-1 restructuring state, with present values taken at the restructuring date and at the risk-free rate, the aggregate* relief requirement *is the reduction in the value of all claims that the envelope $\bar{M}$ imposes, and the* haircut *on class $i \in \{P, O\}$ is the proportional reduction in the value of its claim:*

$$\mathcal{R} \;\equiv\; D - \beta^* \bar{M} \;\in\; (0, D), \tag{4}$$

$$h_i \;\equiv\; 1 - \frac{\beta^* m_i}{d_i} \;\in\; [0, 1]. \tag{5}$$

Because claims fall due at the restructuring date, their pre-restructuring value is their face value and (5) is the standard present-value haircut discounted at the risk-free rate. Feasible allocations satisfy the relief constraint

$$h_P\, d_P + h_O\, d_O \;=\; \mathcal{R}, \qquad h_P, h_O \in [0, 1]. \tag{6}$$

Because the allocation is one equation in two unknowns, it is fully summarized by the official haircut h_O, with $h_P(h_O) = (\mathcal{R} - h_O d_O)/d_P$. Feasibility restricts h_O to $[\underline{h}, \overline{h}]$, where

$$\underline{h} \equiv \max\left\{0, \tfrac{\mathcal{R} - d_P}{d_O}\right\} \qquad \text{and} \qquad \overline{h} \equiv \min\left\{1, \tfrac{\mathcal{R}}{d_O}\right\}. \tag{7}$$

The corner $\underline{h}$ is *preferred creditor status* (PCS): private claims absorb losses first, and the official creditor is touched only if private capacity is exhausted ($\mathcal{R} > d_P$). The corner $\overline{h}$ is *full breach*.

2.4 Crisis lending and the access technology

In the date-2 liquidity-crisis state the government may borrow an amount $\ell \geq 0$ from the official sector's emergency facility at the concessional gross rate $R_E \leq R^*$, repaying $R_E \ell$ at

the settlement date. Access is capped at $\bar{\ell}$, which on the restructuring branch depends on how official claims were treated at date 1.

Assumption 4 (Access technology and scarce crisis liquidity)**.** On the restructuring branch, emergency-lending capacity is

$$\bar{\ell}(h_O) \;=\; \max\{\,0,\; \bar{\ell}_0 - \lambda\, h_O d_O\,\}, \qquad \lambda > 0, \quad \bar{\ell}_0 > 0, \tag{8}$$

and $\bar{\ell}_0$ on the repayment branch. Crisis liquidity is scarce: $u'\big(y^\ell - \bar{M} + \bar{\ell}_0\big) > \beta R_E$.

The coefficient λ is the response of this government's own crisis access to the haircut on its official claims, and (8) is a reduced form with two possible interpretations. One is a balance-sheet channel. Multilateral lending capacity is a gearing multiple of usable capital under capital-adequacy frameworks (G20, 2022), so an uncompensated haircut that destroys $h_O d_O$ of claims removes $L\, h_O d_O$ of headroom across the institution's whole portfolio, of which the share $s \in [0,1]$ that is passed-through to this borrower's facility enters (8): $\lambda = sL$. L is a marginal response, which depends on the capital constraint that binds and on the headroom available, and $s = 1$ would mean that the MDB concentrates the lost headroom on the borrower that caused it. Another interpretation is discrete: the non-toleration-of-arrears rule of Cordella and Powell (2021) makes support available if and only if official claims are performing (the limit of (8) in which capacity collapses at a threshold). We adopt the former interpretation because it delivers interior solutions from a measurable institutional primitive rather than from an assumed curvature; every corner-solution result below survives the discrete version unchanged. The scarcity condition says that even unimpaired access does not satiate crisis liquidity needs, as official crisis finance is rationed by access limits and quotas and, by Assumption 2(iii), no private lender replaces it; it is what gives access strictly positive marginal value, and hence the debtor a possible motive to protect the official creditor.

Define crisis consumption before emergency borrowing as $c_2^c \equiv y^\ell - \bar{M} > 0$ and the *access surplus* from capacity $x \geq 0$ as

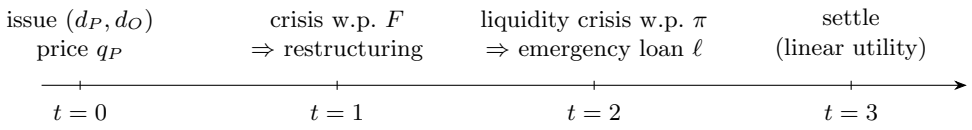
$$A(x) \;\equiv\; \max_{0 \leq \ell \leq x} \; u\big(c_2^c + \ell\big) - \beta R_E\, \ell \;-\; u\big(c_2^c\big) \;\geq\; 0. \tag{9}$$

Under scarcity the cap binds at every feasible capacity level, so A is strictly increasing and strictly concave on $[0, \bar{\ell}_0]$ with derivative equal to the *crisis liquidity wedge*

$$\Delta_\ell(h_O) \;\equiv\; u'\big(c_2^c + \bar{\ell}(h_O)\big) - \beta R_E \;>\; 0 \qquad \text{for } h_O < h_{\mathrm{a}} \equiv \bar{\ell}_0/(\lambda d_O), \tag{10}$$

where h_{a} is the haircut at which access is exhausted.

Figure 1: Model: Timing of Events



Writing $\beta R_E = 1 - (1 - R_E/R^*)$, the wedge has two parts:

$$\Delta_\ell(h_O) \;=\; \underbrace{u'\big(c_2^c + \bar{\ell}(h_O)\big) - 1}_{\text{scarcity premium}} \;+\; \underbrace{1 - \beta R_E}_{\text{grant element of the facility}}, \tag{11}$$

the excess of crisis marginal utility over its settlement value, which is what liquidity is worth, plus the facility's grant element at the risk-free rate. The premium is positive if and only if $u'(c_2^c + \bar{\ell}(h_O)) > 1$, which can fail at low official haircuts when unimpaired capacity is large. The two parts resemble different market structures. Were private investors to lend any amount at R^* in the crisis, contrary to Assumption 2(iii), the government would borrow from them whenever $u' > 1$, and the marginal value of official capacity would be $\min\{\Delta_\ell(h_O),\, 1 - \beta R_E\}$: at most the facility's grant element, and zero at $R_E = R^*$ (Supplemental Appendix, Proposition S.2). Conversely, with private markets closed, access at the market rate $R_E = R^*$ is still worth $u' - 1 > 0$ at the margin. Assumption 2(iii) is therefore what lets the official lender's value rest on liquidity provision and not only on subsidized funding, a distinction Section 3.3 returns to.

Figure 1 summarizes the timing: claims are issued and priced at date 0 given the anticipated allocation and, at date 1, either repaid or exchanged; restructured claims are repaid at date 2, when the liquidity-crisis state triggers the emergency loan settled at date 3.

2.5 Equilibrium and two channels

Definition 2 (Allocation rule and competitive equilibrium)**.** *An* allocation rule *is a value $h_O \in [\underline{h}, \overline{h}]$ for the official haircut in the date-1 restructuring state, with $h_P(h_O)$ given by (6). Given a rule, a* competitive equilibrium *is a price q_P and an emergency-borrowing policy such that (i) q_P satisfies (2) with $h_P^e = h_P(h_O)$; (ii) in the liquidity-crisis state ℓ is chosen optimally subject to the cap, as in (9); and (iii) the government's budget constraints hold at every date and state.*

The first result establishes that the model has two allocation-relevant channels, formalizing the sense in which the CoT problem is a problem about effects on prices and access rather than about total relief resources needed.

Lemma 1 (Debt-service invariance and the two channels)**.** *Fix the relief requirement $\mathcal{R}$. In any*

competitive equilibrium, contractual debt service and the government's resources before emergency borrowing are invariant to the allocation (h_P, h_O) at every date and state from date 1 on. The allocation moves consumption only at date 0, through q_P, and in the restructuring-branch liquidity-crisis state, through emergency borrowing and its settlement, where the cap binds, $\ell = \bar{\ell}(h_O)$. Up to an additive constant independent of the allocation, date-0 expected utility is

$$\mathcal{W}(h_O) \;=\; u\big(c_0(h_O)\big) \;+\; \beta^2 F\pi\, A\big(\bar{\ell}(h_O)\big), \qquad c_0(h_O) = y_0 + \beta^*\big(d_P - F\mathcal{R} + Fh_O d_O\big) + \frac{d_O}{R_O}, \tag{12}$$

with c_0 strictly increasing in h_O.

Proof in Appendix A.1. Once a DSA fixes $\mathcal{R}$, the government's resources at dates 1 and 2 are pinned down: it pays nothing in the crisis, $\bar{M}$ at date 2, and D on the repayment branch, regardless of who bears the losses. What the split changes is, first, the price at which private claims were sold at date 0 (every unit of relief shifted onto the official creditor is a unit private investors no longer expect to lose, and competition passes it to the government as $\beta^* F$ per unit of face value) and, second, the emergency-lending capacity that survives into the date-2 crisis. CoT is therefore not a zero-sum division of a pie; it is a choice between an *ex-ante price* and an *ex-post option*.[7]

3 Comparability of Treatment: three implementations

We now define the CoT problem, and characterize its solution under three criteria. The model of Section 2 admits, at the restructuring node, three disciplined objects on which an allocation criterion can be based. They support two creditor-facing rules, one axiomatic (Criterion 1) and one accounting (Criterion 2), which the welfare function of Lemma 1 prices, and one debtor-cost criterion (Criterion 3), obtained by maximizing that function.

Definition 3 (Three objects at the restructuring node)**.** *For each class $i \in \{P, O\}$, the three objects in the restructuring state are:*

(i) its haircut h_i of Definition 1;

(ii) its excess payoff per unit of face value, in date-1 present value, $x_i \equiv (1 - h_i) - R^* q_i$: the payoff net of the capitalized date-0 price, q_P from (2) and $q_O = 1/R_O$ (Assumption 2);

[7] Lemma 1 does not depend on the three-date horizon. What repetition in independent blocks could add is a reputational equilibrium of the infinitely repeated economy, in which a past private haircut moves later prices, as it does empirically (Cruces and Trebesch, 2013); such prices depend on history rather than on the state, and we do not pursue them.

(iii) the marginal cost of relief through class i, σ_i: the cost, in the debtor's date-0 welfare, of a marginal unit of relief sourced from class i, obtained by differentiating the welfare function of Lemma 1 (Eq. (16)).[8]

Criterion 1 (Section 3.1) equalizes (i) and Criterion 2 (Section 3.2) equalizes (ii), creditor-facing measurements computable from contracts alone and from contracts and observed date-0 prices, respectively; Criterion 3 (Section 3.3) takes the relief requirement as given and sources it across classes at minimum welfare cost to the debtor, and its interior first-order condition equalizes (iii), a debtor-welfare object that requires the technology (u, π, λ, A).[9]

3.1 Claims parity: equal treatment (Criterion 1)

First, we show that equalizing object (i) means equal haircuts. At the restructuring node, restructured claims are default-free (Assumption 3), so the estate to be divided in date-1 present value is $E \equiv \beta^* \bar{M} \in (0, D)$, awards are $a_i = \beta^* m_i$ with $\sum_i a_i = E$ and $0 \leq a_i \leq d_i$, and $h_i = 1 - a_i/d_i$: a canonical rationing problem $(E; d)$.[10]

Domain and axioms. Let $\mathcal{D}$ be the set of problems (N, c, E): N a finite nonempty set of *accounts*, $c \in \mathbb{R}^N_{++}$ claims, $0 < E \leq \sum_i c_i$. A rule f assigns awards with $0 \leq f_i \leq c_i$ and $\sum_i f_i = E$. Claims-measurability (the rule sees (N, c, E) only) is the normative stance: entitlement is the claim, and the technologies driving Lemma 1's two channels are invisible to the rule by construction. The axioms:

(A1) *Equal treatment of equals:* $c_i = c_j$ implies $f_i = f_j$.

(A2) *Reconfiguration invariance:* if (N', c', E) replaces a subset $S \subseteq N$ of accounts by accounts S' with the same total claim, other accounts unchanged, then the total award to S' equals the total award to S.

(A3) *Continuity* of $f(N, \cdot, E)$ in claims.

[8] Because the feasible set (6) is one-dimensional, any allocation can be rationalized ex post by *some* equalization statement. The contribution is that all three objects are equilibrium objects: h_i is contractual, x_i is an arbitrage identity in observed prices, and σ_i is a derivative of the welfare function. The single free scalar, introduced below, the netting share $\kappa \in [0, 1]$, has principled endpoints: object (i) at $\kappa = 0$ and object (ii) at $\kappa = 1$; representing these objects through empirical proxies is a separate exercise, deliberately not undertaken here.

[9] In practice, CoT is implemented through negotiations between the debtor and its creditors. In this context, the criteria we develop here should be seen as a fair benchmark for those negotiations, which we abstract from.

[10] The rationing ("bankruptcy") problem of O'Neill (1982) and Aumann and Maschler (1985). Axiom (A2) combines this literature's no advantageous reallocation (Chun, 1988) with its non-manipulability via merging and splitting of claims (de Frutos, 1999; Ju, Miyagawa, and Sakai, 2007). Proposition 1 is a self-contained variant stated for group reconfigurations, with (A3) used only to pass from rational to real claims; the contribution here is the mapping of the restructuring node onto the problem and the identification of the doctrines with its canonical rules (Moulin, 2000).

Proposition 1 (Equal treatment characterized)**.** *A rule on $\mathcal{D}$ satisfies (A1)–(A3) if and only if it is the proportional rule $f_i = (c_i/\sum_j c_j)E$. At the restructuring node the proportional rule is equal treatment: $h_P = h_O = \mathcal{R}/D \equiv \theta$.*

Proof in Appendix A.2; it is valid for any number of classes. (A2) is the main ingredient: claims are repackageable in practice (participations, SPVs, on-lending), so a class-based rule that is not reconfiguration-invariant invites boundary gerrymandering.[11] Proportionality is the unique claims-measurable, gerrymander-proof rule. Preferred creditor status, by contrast, is the *priority* rule with the official class last in the loss order: the polar canonical rationing method, which violates (A1) by construction.

3.2 Origination-adjusted treatment: the no-windfall rule (Criterion 2)

Criterion 2 equalizes the excess payoff x_i (object (ii) of Definition 3), which can be read off the date-0 contracts.

3.2.1 What origination terms pre-pay

Lemma 2 (Origination concessions and realized excess payoffs)**.** *Define the* pre-paid compensation *of class i as the discount, per unit of face value and in date-1 value, at which the class acquired its claim relative to the risk-free price β^* of a safe unit-face claim:*

$$\chi_i \equiv \frac{\beta^* - q_i}{\beta^*} = 1 - R^* q_i.$$

Under Assumption 2,

$$\chi_P = F h_P^e, \qquad \chi_O = 1 - \frac{R^*}{R_O} = -g_O.$$

Moreover, per unit of face value:

(i) χ_i is class i's realized excess payoff in the no-restructuring *state, and $x_i = \chi_i - h_i$ its realized excess payoff in the* restructuring *state;*

(ii) expected excess payoffs, $(1-F)\chi_i + F x_i = \chi_i - F h_i$, are zero for P under any anticipated rule and equal to $-(g_O + F h_O) \le 0$ for O, with equality if and only if $g_O = 0$ and $h_O = 0$.

Proof in Appendix A.6. The private spread is not profit but the actuarially fair premium for anticipated treatment: it exactly pre-funds the expected loss ($\chi_P = F h_P$). No anticipated

[11] The same conclusion holds in equilibrium. If accounts could re-register across classes at no cost before the exchange, any anticipated $h_P \neq h_O$ would make migration into the class with the lower haircut strictly profitable, and the relief constraint (6) would then force θ on every account; differential treatment presupposes segmentation. In practice, IMF credit is non-transferable by legal construction (IMF, *Financial Organization and Operations of the IMF*, Pamphlet 45), IBRD loans have no secondary market (IBRD Information Statement), and Paris Club agreed minutes cap the sale of non-ODA claims to investors, in the Club's own words to preserve CoT.

rule is therefore unfair to the market ex ante; a rule shapes only the state-contingent profile of private payoffs, which is the margin on which Criterion 2 acts. The official terms run the other way: the lender concedes g_O in every state and, because its rate does not respond to default risk (Assumption 2(ii)), collects nothing for the expected haircut, so it breaks even in expectation only on market-rate lending that the rule fully exempts.[12] Hence, under equal treatment of claims ($h_P = h_O = \theta$), restructuring-state excess payoffs are *not* equal:

$$x_P - x_O \;=\; \chi_P - \chi_O \;=\; F\theta + g_O \;>\; 0. \tag{13}$$

The private class outperforms the official class by exactly the compensation differential.[13] Equal treatment of claims is unequal origination-adjusted treatment.

3.2.2 The rule and its equilibrium

As a free-standing norm, the no-windfall allocation is fair when the accepted entitlement metric is the claim net of priced origination compensation, not the gross claim: creditors already paid, in price or in rate, for restructuring exposure have pre-funded that share of their loss.

Definition 4 (No-windfall rule with netting share κ)**.** *Fix $\kappa \in [0,1]$. Given the measured compensations (χ_P, χ_O) of Lemma 2, the* no-windfall allocation *makes the netted losses $h_i - \kappa\chi_i$ as equal as the relief constraint (6) and the bounds $h_i \in [0,1]$ permit: it minimizes $|(h_P - \kappa\chi_P) - (h_O - \kappa\chi_O)|$ subject to them. When the bounds do not bind, it satisfies*

$$h_P - h_O \;=\; \kappa\,(\chi_P - \chi_O). \tag{14}$$

An equilibrium *of the rule is a rational-expectations fixed point: the treatment anticipated in the date-0 price is the treatment the rule implements, $h_P^e = h_P$.*

Under (14), treatment differentials are set proportional to compensation differentials: $\kappa = 0$ is claims accounting (equal haircuts, Section 3.1); $\kappa = 1$ is origination-adjusted accounting, which by Lemma 2(i) equalizes realized excess payoffs x_i, so that no class exits the restructuring having done better than another; and interior κ is partial recognition of origination terms, which is a not-model-derived policy stance. The fixed point in Definition 4 captures the key intuition. The compensation the rule nets out for the private class is the market's price of the rule's *own*

[12] The sign $g_O \geq 0$ captures only the concessionality in Assumption 2(ii). Treatment-insensitivity priced above the risk-free rate ($g_O < 0$) collects a premium its treatment does not justify and can reverse the sign in part (ii). Criterion 3's pricing/access tradeoff does not require inherited $R_O \leq R^*$, provided its other conditions hold.

[13] x_i is an excess payoff per unit of face value, not a rate of return on the funds advanced ($x_i/(R^* q_i)$); origination prices differ across classes, so equalizing one does not equalize the other. We normalize by face-value to state Criterion 2 in the units of haircuts and of the relief constraint (6), keeping condition (14) linear; the signs in Lemma 2(ii) hold for rates as well, since q_i is predetermined (Suppl. Appendix, Remark S.1 and Proposition S.3).

anticipated treatment: a rule that claws back more induces a larger spread, which enlarges the measured compensation, which justifies a larger differential. The equilibrium resolves this self-reference, and the resolution is unique and in closed form.

Proposition 2 (The no-windfall equilibrium)**.** *Maintain Assumption 2(ii), so that* $g_O \geq 0$*, together with* $\mathcal{R} \leq d_P$ *(so* $\underline{h} = 0$*),* $F \in (0,1)$*, and* $\kappa F < 1$*, and let* $\kappa^* \equiv \mathcal{R}/(F\mathcal{R} + g_O d_P)$*. For each* $\kappa \in [0,1]$ *the no-windfall rule has a unique equilibrium* $(h_P^{\mathrm{nw}}, h_O^{\mathrm{nw}})$*:*

(i) (Allocation.) *For* $\kappa \leq \kappa^*$ *the equilibrium solves (14) and (6):*

$$h_P^{\mathrm{nw}} = \frac{\mathcal{R} + \kappa g_O d_O}{d_P + (1-\kappa F) d_O}, \qquad h_O^{\mathrm{nw}} = (1-\kappa F)\, h_P^{\mathrm{nw}} - \kappa g_O,$$

so that $h_P^{\mathrm{nw}} - h_O^{\mathrm{nw}} = \kappa(F h_P^{\mathrm{nw}} + g_O)$*;* h_P^{nw} *is strictly increasing and* h_O^{nw} *strictly decreasing in* κ*, from equal treatment,* $h_P^{\mathrm{nw}} = h_O^{\mathrm{nw}} = \theta$*, at* $\kappa = 0$*. For* $\kappa \in [\kappa^*, 1]$ *(nonempty iff* $\kappa^* \leq 1$*) the bound* $h_O \geq 0$ *binds and the equilibrium is PCS:* $h_O^{\mathrm{nw}} = 0$*,* $h_P^{\mathrm{nw}} = \mathcal{R}/d_P$*.*

(ii) (Windfall gap.) *In equilibrium,*

$$x_P - x_O = \begin{cases} (1-\kappa)\big(F h_P^{\mathrm{nw}} + g_O\big) = (1-\kappa)\,\dfrac{F\mathcal{R} + g_O D}{d_P + (1-\kappa F) d_O} & \text{if } \kappa \leq \kappa^*, \\[2ex] g_O - (1-F)\,\dfrac{\mathcal{R}}{d_P} & \text{if } \kappa \geq \kappa^*. \end{cases}$$

The gap is continuous and weakly decreasing in κ *(strictly for* $\kappa \leq \kappa^*$*), from the equal-treatment windfall* $F\theta + g_O$ *of (13) at* $\kappa = 0$*. Full netting (*$\kappa = 1$*) eliminates it if and only if* $\kappa^* \geq 1$*, that is,* $g_O \leq (1-F)\mathcal{R}/d_P$*.*

Proof in Appendix A.7. As κ rises, the recognized private compensation $\kappa F h_P^{\mathrm{nw}}$ and the recognized official concession κg_O both grow, and the allocation moves monotonically from equal treatment toward preferred creditor status. If the concession is shallow ($\kappa^* \geq 1$), full netting closes the gap. If it is deep ($\kappa^* < 1$), the official haircut reaches its floor first: the concession is sunk, and once $h_O = 0$ there is nothing left to concede. Past κ^* the rule delivers preferred creditor status and a residual windfall $g_O - (1-F)\mathcal{R}/d_P$, the excess of the official class's sunk concession over the private class's unpriced loss in the restructuring state. No netting share can remove it: that would require a transfer to the official class, and haircuts cannot implement transfers. Full netting then justifies complete exemption and still leaves the official class with the lower realized excess payoff.

The rule sees only priced margins: it can favor the official class only up to the compensation differential embedded at origination, $F h_P^{\mathrm{nw}} + g_O$, whereas the crisis access that motivates protection under Criterion 3 is never priced (Assumption 2(ii)) and never enters χ_O. Supple-

mental Appendix S.1 characterizes when a netting share nonetheless implements the allocation of Criterion 3, and when Criterion 2 reduces to equal haircuts.

3.3 The debtor-cost criterion: relief at minimum welfare cost (Criterion 3)

Criterion 3 takes the relief requirement fixed by the DSA as given and sources it across classes at minimum welfare cost to the debtor.

Assumption 5 (Commitment)**.** The rule $h_O \in [\underline{h}, \overline{h}]$ is chosen at date 0 to maximize (1) subject to Definition 2, internalizing the dependence of q_P on the rule through (2).

Commitment is what keeps the pricing channel of Lemma 1 live: the split is announced before private claims are priced, so the criterion is a design object; chosen at date 1 instead, with q_P sunk, the same objective delivers preferred creditor status (see Corollary 2). By Lemma 1, the planner's problem is $\max_{h_O \in [\underline{h}, \overline{h}]} \mathcal{W}(h_O)$.

Proposition 3 (The debtor-cost criterion)**.** *Maintain Assumptions 1–4, and let $h_O^*(\pi)$ denote the Ramsey allocation, a maximizer of $\mathcal{W}$ on $[\underline{h}, \overline{h}]$.*
(i) Degenerate access. *If $\underline{h} \geq h_{\mathrm{a}}$, equivalently $\bar{\ell}_0 \leq \lambda(\mathcal{R} - d_P)$, then $\mathcal{W}$ is strictly increasing on $[\underline{h}, \overline{h}]$ and the allocation is full breach, $h_O^* = \overline{h}$, for every π.*
(ii) Nondegenerate access. *If $\underline{h} < h_{\mathrm{a}}$, then $\mathcal{W}$ is strictly concave on $[\underline{h}, h_{\mathrm{a}} \wedge \overline{h}]$, with a unique maximizer $h_O^\circ(\pi)$ there, and, if $h_{\mathrm{a}} < \overline{h}$, strictly increasing on $[h_{\mathrm{a}}, \overline{h}]$. There are thresholds $\hat{\pi}$ and π_p, with $\hat{\pi} \in \{\pi_a, \pi_b\}$ and $0 < \pi_a < \pi_p < \infty$ (all defined below), such that the allocation is full breach, $h_O^* = \overline{h}$, for $\pi < \hat{\pi}$; interior burden-sharing, $h_O^* = h_O^\circ(\pi) \in (\underline{h}, h_{\mathrm{a}} \wedge \overline{h})$, the unique solution of the access margin*

$$\underbrace{u'\big(c_0(h_O^\circ)\big)\,\beta^*}_{\substack{\text{marginal date-0 value of the} \\ \text{pricing gain per unit of relief}}} \;=\; \underbrace{\beta^2 \pi\, \lambda\, \Delta_\ell\big(h_O^\circ\big)}_{\substack{\text{expected marginal value of the} \\ \text{crisis access destroyed per unit}}}, \tag{15}$$

strictly decreasing in π, for $\pi \in (\hat{\pi}, \pi_p)$; and preferred creditor status, $h_O^* = \underline{h}$, for $\pi \geq \max\{\hat{\pi}, \pi_p\}$. If $\overline{h} \leq h_{\mathrm{a}}$, then $\hat{\pi} = \pi_a < \pi_p$ and h_O^* is single-valued and continuous in π. If $\overline{h} > h_{\mathrm{a}}$, then $\hat{\pi} = \pi_b > \pi_a$ is the unique zero of Δ; h_O^* is single-valued except at π_b, where full breach and $h_O^\circ(\pi_b)$ both attain the maximum, so the allocation jumps down at π_b; and the interior band is empty if and only if $\pi_b \geq \pi_p$, i.e., $\Delta(\pi_p) \leq 0$.

Here π_p and π_a are the crisis probabilities at which (15) holds at the two ends of the access-preserving branch $[\underline{h}, h_{\mathrm{a}} \wedge \overline{h}]$, and Δ (for $h_{\mathrm{a}} < \overline{h}$) is the value gap between the best access-preserving allocation and full breach:

$$\pi_p \equiv \frac{u'(c_0(\underline{h}))}{\beta\lambda\, \Delta_\ell(\underline{h})}, \qquad \pi_a \equiv \frac{u'(c_0(h_{\mathrm{a}} \wedge \overline{h}))}{\beta\lambda\, \Delta_\ell((h_{\mathrm{a}} \wedge \overline{h})^-)}, \qquad \Delta(\pi) \equiv \max_{h_O \in [\underline{h}, h_{\mathrm{a}}]} \mathcal{W}(h_O; \pi) - \mathcal{W}(\overline{h}; \pi).$$

Regimes are stated on the ray $\pi \in (0, \infty)$ and read on the admissible domain $\pi \in (0, 1)$, in which any of them may be empty.

Proof in Appendix A.3. Equation (15) prices the marginal value of access to countercyclical crisis lending.[14] The left side is what one more unit of official relief buys today: resources front-loaded through a higher bond price, valued at date-0 marginal utility. The right side is what it costs: with probability π a crisis arrives, λ turns the haircut into lost emergency capacity, and each lost unit is worth the wedge between crisis marginal utility and the concessional cost of funds.[15] Crisis risk raises the protection the criterion affords to official creditors: on the access-preserving branch h_O° falls with π, and the allocation passes from full breach to graduated participation to preferred creditor status as π crosses $\hat{\pi}$ and π_p.[16] Neither threshold need lie in $(0, 1)$: if $\pi_p > 1$ the official haircut stays above $\underline{h}$ at every admissible crisis probability, and if $\hat{\pi} > 1$ full breach is optimal throughout, however likely the crisis; at $\hat{\pi} = \pi_b$ the maximizer is not unique. Preferred creditor status is exemption, $h_O = 0$, only if $\mathcal{R} \leq d_P$; otherwise it is private-first relief with official participation $\underline{h} > 0$. Part (i) is the deep-overhang configuration: relief so large relative to private capacity that no feasible allocation preserves the facility, so protecting officials buys nothing. Such an outcome allows an analogy with the Heavily Indebted Poor Countries Initiative (HIPC) and the Multilateral Debt Relief Initiative (MDRI).[17]

Write $r_O \equiv h_O d_O$ for the relief borne by the official class, and let

$$\sigma_P(h_O) \equiv u'\big(c_0(h_O)\big)\,\beta^* F, \qquad \sigma_O(h_O) \equiv \beta^2 F \pi \lambda\, \Delta_\ell(h_O)\, \mathbf{1}\{h_O < h_{\mathrm{a}}\}, \tag{16}$$

the marginal cost, in the debtor's date-0 welfare, of a unit of relief sourced from each class: $d\mathcal{W}/dr_O = \sigma_P - \sigma_O$ for $h_O \neq h_{\mathrm{a}}$ (Appendix A.3), so (15) follows $\sigma_P = \sigma_O$. Relief is sourced where it is cheapest and, at an interior optimum, marginal costs are equalized across classes; beyond h_{a} official relief is free at the margin, so full breach is settled by the global comparison Δ rather than by a marginal condition. This is the structure of Ramsey's problem, a fixed levy ($\mathcal{R}$) raised from several bases (the classes) at minimum welfare cost, with (15) as its equal-marginal-

[14] The normalization $\beta = \beta^*$ (Assumption 1) puts the two sides on a common footing. Holding the rule and R^* fixed, changing β leaves issue prices and c_0 unchanged, but changes both the weight $\beta^2 F\pi$ on access surplus and the borrowing-cost wedge $u' - \beta R_E$ (via the repayment cost βR_E). These are allocation-relevant continuation-value effects because access depends on the official haircut, and they can change prices through a newly optimal allocation. The direction of that allocation change is not generally determined by the stated scarcity assumption.

[15] Criterion 3 does not deliver global efficiency. The objective omits the headroom that an uncompensated impairment removes from the official lender's other borrowers, the share $1 - s$ of Section 2.4, and the cost to the lender's shareholders of compensating it.

[16] This is the allocation-side mirror of Proposition 1 in Cordella and Powell (2021), which characterizes when a sovereign honors preferred creditor status; here a probability threshold governs when a planner designing the restructuring should honor it.

[17] HIPC and the subsequent MDRI used dedicated financing and, in important cases, donor compensation to sustain multilateral lending after relief. Those arrangements are not represented in the access technology.

sacrifice condition.[18] Unlike in taxation, the sacrifice equalized is the debtor's, not the creditors': the private class's expected excess payoff is zero under any anticipated rule (Lemma 2(ii)), and σ_i is the sacrifice the debtor bears through class i.

The protection is conditional on the access technology and on the market structure behind the wedge, not on official status. $\sigma_O > 0$ only because Assumption 4 makes this government's crisis capacity fall with the official haircut and Assumption 2(iii) makes that capacity scarce: with $\lambda = 0$ the same criterion would breach the official class fully for every π, and so would it with private creditors lendig during a crisis at R^* and a facility priced at R^*, where the wedge of (11) is replaced by $\min\{\Delta_\ell, 1 - \beta R_E\} = 0$ (Supplemental Appendix, Proposition S.2). Treatment-insensitive pricing, Assumption 2(ii), plays the opposite role: it is what keeps $\sigma_P > 0$. Were official claims priced like private ones, c_0 would be invariant to the allocation and Criterion 3 would be preferred creditor status for every π whenever $\underline{h} < h_{\mathrm{a}}$ (Supplemental Appendix, Corollary S.1). Two official attributes thus play distinct roles: pricing rigidity opens the date-0 transfer margin, and treatment-sensitive, scarce future access earns protection. Criterion 3 protects whichever class controls such access; the official class earns protection as the empirically motivated provider of emergency lending.

Figure 2 displays the regimes at the illustrative parameters of Table S.1. The illustration is an existence claim, not a proper calibration. The parameters are chosen so that all three regimes occur within the probability domain, with full breach below $\pi_b = 0.25$, an interior band in which the official haircut falls from 29.5 percent to zero, and full protection above $\pi_p = 0.64$, and no frequency is attached to them. At $\pi = 0.5$ the criterion assigns officials a haircut of 6.9 percent against 70.3 percent for private creditors, where equal treatment would impose $\theta = 43.9$ percent on both, and the issue price is 0.826 against 0.877 under committed equal treatment. The protection in the example is bought by scarce liquidity: the wedge at $h_O = 0$ is 0.353, of which the facility's grant element (the rate subsidy) is 0.019, and with private lender's backup liquidity at R^* full breach would be optimal at every admissible π (Supplemental Appendix, Proposition S.2). The regimes also move with the own-access response: Table S.2 in the Supplemental Appendix varies λ from 1 to 6; at $\lambda \leq 1.5$ the protection threshold π_p exceeds one, so officials are never fully protected, and at $\lambda \geq 4$ they are exempt already at $\pi = 0.5$.

Corollary 1 (Comparative statics of interior burden-sharing)**.** *On the interior regime, the optimal official haircut h_O° satisfies*

$$\frac{\partial h_O^\circ}{\partial \pi} < 0, \qquad \frac{\partial h_O^\circ}{\partial \lambda} < 0, \qquad \frac{\partial h_O^\circ}{\partial R_E} > 0, \qquad \frac{\partial h_O^\circ}{\partial R_O} > 0, \qquad \frac{\partial h_O^\circ}{\partial \mathcal{R}} > 0,$$

where the last derivative varies crisis depth holding y^ℓ fixed (so that $\bar{M} = (D - \mathcal{R})/\beta^*$ falls as

[18] Pigou (1928) identifies equal *marginal* sacrifice with the utilitarian optimum.

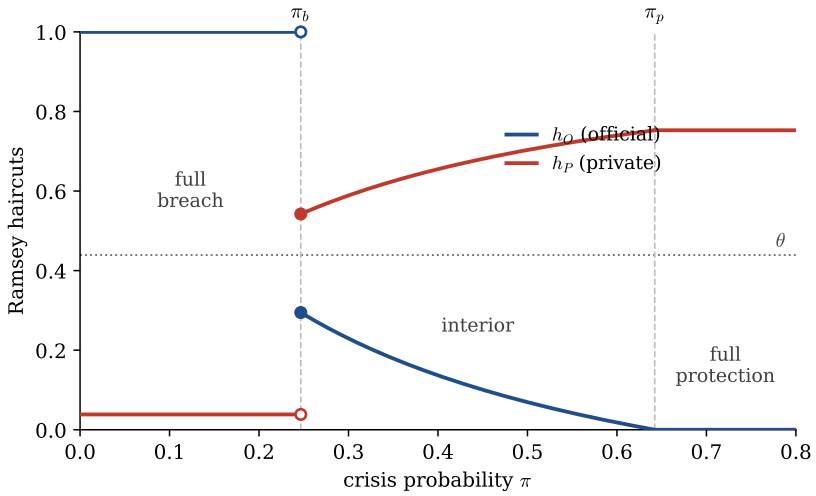


Figure 2: Ramsey haircuts as a function of the post-restructuring crisis probability π. Illustrative parameters of Table S.1 (Supplemental Appendix): $\pi_b = 0.25$, $\pi_p = 0.64$. The dotted line is the equal-treatment haircut θ. At π_b the criterion is indifferent between full breach and the interior allocation (open and filled circles); the figure plots the concave-branch selection for $\pi \geq \pi_b$.

$\mathcal{R}$ rises). In particular, the official haircut is decreasing in the concessionality of crisis support ($R_E \downarrow$) and in the concessionality of ex-ante lending terms ($R_O \downarrow$), and increasing in the depth of the crisis.

Proof in Appendix A.4. Each sign is a statement about the two sides of (15). A higher crisis probability or a larger access response raises the cost of impairing access, so officials are protected more; a less concessional facility (higher R_E) shrinks the wedge, so they are haircut more: haircuts fall with the concessionality of the ongoing relationship, a future-terms channel. The R_O result is a more modest channel: cheaper ex-ante lending makes the government richer at date 0 and lowers the value of the pricing gain, an income effect rather than a punishment effect. Deeper crises raise the official haircut through both channels (a poorer date 0, a lighter $\bar{M}$), which is why the sign is unambiguous; this matches Wicht's (2025) fact that official creditors take losses only in the deepest crises, alongside the largest private haircuts.

Corollary 2 (Discretion delivers preferred creditor status)**.** *Suppose $\underline{h} < h_{\mathrm{a}}$ and the allocation is chosen at date 1, with q_P predetermined and expectations rational. Then:* (i) *the unique equilibrium allocation is preferred creditor status, $h_O = \underline{h}$ and $h_P = \min\{1, \mathcal{R}/d_P\}$;* (ii) *no rule with $h_O > \underline{h}$ is sequentially rational; and* (iii) *the equilibrium issue price, $q_P = \beta^*(1 - F\min\{1, \mathcal{R}/d_P\})$, is the lowest attainable under any feasible rule, strictly so whenever the rule assigns $h_O > \underline{h}$. If $\underline{h} \geq h_{\mathrm{a}}$, the date-1 planner is indifferent among all feasible allocations.*

Proof in Appendix A.5. At the restructuring date the price is sunk and access is not, so the

only live margin is access (Lemma 1) and every dollar of relief goes to private creditors first. The model thus derives the exemption observed in virtually every restructuring: PCS is not a privilege extracted by officials but what the debtor itself chooses ex post. It would like to promise official burden-sharing, which raises q_P and front-loads resources, but investors price the outcome they know will occur, and the economy lives with the lowest issue price. The indifference under $\underline{h} \geq h_{\mathrm{a}}$ is the discretionary face of Proposition 3(i): once access is gone, protecting officials buys nothing. Nor does it buy anything when access is not scarce: if private lenders provide liquidity at R^* the date-1 planner protects the official class only for the facility's grant element, and at $R_E = R^*$ is indifferent among all feasible allocations (Supplemental Appendix, Proposition S.2). In the model, preferred creditor status is earned by liquidity that private markets do not supply, or by a rate that they do not match.

Discussion. We note four avenues for possible extensions. First, the framework characterizes the three criteria but does not rank them. A welfare comparison of the criteria could include creditors' payoffs. Private investors break even under any anticipated rule, while the official lender's expected loss rises with its haircut, so such a comparison would turn on how official losses are valued. Second, contracts are one-period, so claims differ in present value and, across classes only, in grant element, but not in timing. With payments at several dates, instruments equal on both counts can place debt service at dates with different marginal utilities, so timing would have allocative content of its own. Third, the likelihood F and the size of the restructuring are exogenous, so a rule is priced into private claims at issuance but, by design (Assumption 3), changes neither. Fourth, all claims are in one currency, so the framework is silent on comparability between local- and foreign-currency debt. Addressing these limitations is left for future research.

Three creditor classes. Supplemental Appendix S.2 adds bilateral official creditors, who lend for development goals or geopolitical reasons, and sometimes at concessional terms. This development lending is sometimes conditioned on the treatment received on their claims. A lender that does not condition is the residual claimant: its haircut destroys nothing the debtor values, so the date-1 planner is indifferent between it and bondholders and Criterion 3 sources relief from it first.[19] A lender that conditions inherits the multilateral lender's protection (Proposition S.4). The no-windfall rule ranks official classes inversely by the concessionality of the claims they delivered (Corollary S.2).

[19] Schlegl, Trebesch, and Wright (2019) document this juniority of bilateral official debt.

4 Endogenous exit yields and the measurement of comparability

Under Assumption 3 restructured claims are default-free ($\bar{M} < y^\ell$), so the risk-free rate is the unique internally consistent discount rate for measuring haircuts. In practice they are not, and the rate is a key source of disagreement among creditors: market-value assessments discount the new instruments at their *exit yields*—the secondary-market yields prevailing at, or right after, the exchange, which are still forecasts when terms are set.[20] Exit yields, however, price the risk that the restructured claims will themselves not be honored, and that risk depends on the restructuring's terms. We therefore re-open the date-2 repayment decision (Lemma 3), characterize the market-value haircut that exit yields imply (Lemma 4), and show that the equal-treatment rule always selects a unique allocation under the risk-free metric (Proposition 4) but need not under the market-value metric (Proposition 5).

4.1 A default margin at date 2

Assumption 6 (Date-2 default option)**.** Modify the performance clause of Assumption 3 as follows. At date 2, after y_2 is realized and before emergency borrowing, the government chooses a repayment action $\delta = (\delta_P, \delta_O) \in \{0,1\}^2$, where $\delta_i = 1$ denotes default on the restructured claim m_i. A defaulted claim is extinguished: nothing is recovered on it, at date 2 or at date 3, and restructured claims are neither renegotiated nor rescheduled. Any default ($\delta \neq (0,0)$) entails a resource cost φy_2, $\varphi \in (0,1)$, incurred once per default event. Assumption 2(iii) holds under every action, so the only access that depends on the action is the official window, which closes on official arrears:

$$\ell \;\leq\; \bar{\ell}(h_O)\,(1-\delta_O) \qquad \text{in state } y^\ell. \tag{17}$$

The envelope is serviceable in the recovery state: $\varphi \geq \bar{M}/y^h$. Consumption must be strictly positive: u is extended by $u(c) = -\infty$ for $c \leq 0$, so an action is feasible only if it implies positive consumption, inclusive of emergency borrowing where the window is open.

The default cost is proportional to output, as in Aguiar and Gopinath (2006):[21] with two date-2 states only the levels φy^h and φy^ℓ enter, so the shape of the cost in y_2 is immaterial, while proportionality delivers the serviceability floor and the closed-form windows below in a single restriction. Enforcement is *symmetric*: the same φ applies to arrears on either class. In Dellas and Niepelt (2016) official creditors levy harsher sanctions by assumption; here any de facto date-2 seniority must be *derived* from the access technology (Lemma 3). The window condition (17)

[20] This is the convention of Sturzenegger and Zettelmeyer (2008), which underlies the haircut series of Cruces and Trebesch (2013); it was designed to estimate investor losses in past restructurings, not to assess comparability.

[21] Alternatives are the market-exclusion penalty of Eaton and Gersovitz (1981), asymmetric costs (Arellano, 2008), and quadratic costs (Chatterjee and Eyigungor, 2012).

is the non-toleration-of-arrears reading of Assumption 4 (Cordella and Powell, 2021): crisis support flows only to a borrower current on its official claims, the rule to which lending-into-arrears policies are the codified exception. Zero recovery follows from the extinction of defaulted claims and the absence of a second renegotiation, which are restrictions on the action set.[22] The cost is incurred once per default event, the standard reading of φ as the output disruption of default status.[23] The floor $\varphi \geq \bar{M}/y^h$ makes the envelope serviceable in the recovery state: a restructuring that fails even when output recovers has not restored sustainability. It also places the live default margin in renewed crises, where Wicht (2025) finds re-defaults, and the baseline is nested as the strong-enforcement limit $\varphi \geq \bar{M}/y^\ell$ (Lemma 3).

Let $\Gamma(x; c)$ denote the value in state y^ℓ of a consumption base c with emergency capacity x,

$$\Gamma(x; c) \equiv \max_{0 \leq \ell \leq x} \{u(c + \ell) - \beta R_E \ell\}, \tag{18}$$

so that $\Gamma(x; c_2^c) = u(c_2^c) + A(x)$ for the access surplus A of Eq. (9). At allocation h_O, capacity is $\bar{\ell}(h_O)$ and the payments due are $(m_P, m_O) = \big(R^*(d_P - \mathcal{R} + h_O d_O),\ R^*(1 - h_O)d_O\big)$, which sum to $\bar{M}$ by (4). Define the *official-service surplus*

$$S(h_O) \equiv \Gamma\big(\bar{\ell}(h_O); b\big) - u(b + m_O), \qquad b \equiv c_2^c + m_P - \varphi y^\ell = (1 - \varphi)y^\ell - m_O, \tag{19}$$

the value of staying current on the official class, conditional on having defaulted on the private class.

Lemma 3 (Date-2 repayment policy)**.** *Under Assumptions 1–4 and 6, with indifference resolved toward repayment, the government repays both classes in full in state y^h at every feasible allocation. In state y^ℓ, at allocation h_O, it never defaults on the official class alone; it repays both classes, $\delta = (0, 0)$, iff $m_P \leq \varphi y^\ell$ and $\Gamma\big(\bar{\ell}(h_O); c_2^c\big) \geq u\big((1 - \varphi)y^\ell\big)$; it defaults on the private class alone, $\delta = (1, 0)$, iff $m_P > \varphi y^\ell$ and $S(h_O) \geq 0$; and it defaults on both classes, $\delta = (1, 1)$, otherwise. If $\varphi \geq \bar{M}/y^\ell$, then $\delta = (0, 0)$ in every state at every feasible allocation, and the model of Sections 2 and 3 obtains unchanged.*

Proof in Appendix A.8. Ex post, the classes differ only in that arrears to officials close the emergency window. The government therefore defaults selectively on the marketable claim—pocketing m_P against the cost φy^ℓ—and services the official claim exactly when the relationship

[22] Date 3 exists and settles the emergency loan, so restructured claims could in principle be rescheduled into it. A rescheduled private claim would be a new private claim on $\bar{y}_3$, which Assumption 2(iii) excludes under every action: the exclusion follows the date-1 restructuring, not the date-2 default, so the four actions are compared on the same asset menu. A second renegotiation would re-open at date 2 the allocation problem the paper poses once, at date 1. Lemma 3 and the assessment equilibria of Definition 6 are conditional on both restrictions.

[23] With a recovery rate $\rho \in [0, 1)$, the pricing formulas below hold with π replaced by $\pi(1 - \rho)$. Were the cost levied once per defaulted class, defaulting on the official class alone could be optimal when $\varphi y^\ell > m_P$, so the selective-default ranking of Lemma 3 is specific to the event-level cost (Supplemental Appendix, Remark S.4).

still buys something, $S(h_O) \geq 0$: the date-2 policy re-derives preferred creditor status ex post. Paying m_O to keep capacity $\bar{\ell}(h_O)$ is borrowing to repay the official sector, which cannot pay for itself unless capacity strictly exceeds the payment due, $\bar{\ell}(h_O) > m_O$ (the *servicing bound*). When $\lambda > R^*$, $\bar{\ell}$ falls faster than m_O as h_O rises, and a deep official haircut can gut the relationship until arrears extend to the official class—the deep-overhang configuration, now at date 2.

4.2 Exit prices, metrics, and assessment equilibria

After the exchange, the new claims are one-period instruments, repaid in y^h (Lemma 3). Rational date-1 valuation gives the *exit price* and *gross exit yield* of class i's claim per unit of face value,

$$q_i^x \;=\; \beta^*\big(1 - \pi\,\delta_i^\ell\big), \qquad Y_i^x = 1/q_i^x, \tag{20}$$

where $\delta_i^\ell \in \{0,1\}$ is the anticipated default action in state y^ℓ; the exit spread $Y_i^x - R^*$ is zero for a performing claim and $R^*\pi/(1-\pi)$ otherwise.

Definition 5 (Market-value haircut)**.** *The market-value haircut on class i is the proportional loss in the value of its claim at the restructuring date, with the new instruments valued at exit prices:*

$$h_i^m \;\equiv\; 1 - \frac{q_i^x\, m_i}{d_i} \;=\; 1 - \big(1 - \pi\delta_i^\ell\big)\big(1 - h_i\big). \tag{21}$$

Because old claims are due at the restructuring date, (21) is the Sturzenegger–Zettelmeyer (2008) haircut; it differs from h_i only in the rate that discounts the *new* instruments.

Lemma 4 (Measured haircuts and date-0 pricing)**.** *(i) $h_i^m - h_i = \pi\delta_i^\ell(1-h_i) \geq 0$, and for a common anticipated action $\delta_P^\ell = \delta_O^\ell = \delta$ measured gaps compress: $h_P^m - h_O^m = (1-\pi\delta)(h_P - h_O)$. (ii) In any competitive equilibrium in which investors anticipate the crisis allocation and the date-2 policy of Lemma 3, the date-0 price generalizes Eq. (2) to*

$$q_P = \beta^*\big(1 - F\,h_P^m\big). \tag{22}$$

(iii) h_i^m is the expected present-value loss per unit of the old claim, anticipated re-default included.

Proof in Appendix A.9. Part (iii) is a deliberate concession: as *accounting*, the market-value haircut is not wrong—it is the actuarially correct expected loss, and part (ii) says the primary market prices precisely this object, so the ex-ante channel of Lemma 1 runs through h_P^m, not h_P. The defect is conceptual, and it matters once the haircut is the metric of a rule: h_i^m conflates the relief a class has *granted* (the terms h_i) with losses it is merely *anticipated to absorb later* (the *forward re-default premium* of part (i))—losses that may never materialize and whose probability is created by the allocation itself, that is, by how the *other* class was treated.

Definition 6 (Haircut metric and assessment equilibrium)**.** *A haircut metric assigns to each feasible allocation and anticipated action δ^ℓ measured haircuts (η_P, η_O): the* risk-free metric *sets $\eta_i = h_i$ (Definition 1); the* market-value metric *sets $\eta_i = h_i^m$ at rational exit prices (20). The equal-treatment rule under metric η selects the allocation satisfying $\eta_P = \eta_O$ and the relief constraint (6). An* assessment equilibrium *under η is a triple—allocation h_O, anticipated action $\delta^\ell \in \{0,1\}^2$ (a degenerate, or pure, assessment), exit prices (q_P^x, q_O^x)—such that (i) exit prices are rational given δ^ℓ; (ii) the allocation satisfies the rule at those prices; (iii) δ^ℓ is the date-2 policy of Lemma 3 at that allocation. Date-2 indifference is resolved toward repayment.*

The rule is committed; what no contract can bind is date-2 repayment—the Eaton–Gersovitz margin that Assumption 6 re-opens—so the failures below are not credibility failures. Equal treatment is the natural rule to study: it is the usual benchmark ($\theta = \mathcal{R}/D$) and, under the market-value metric, it renders "comparable effort" assessments, which value private claims at exit yields and official claims, untraded and serviced in the relevant assessment, at par.

4.3 Determinacy and its failure

Proposition 4 (Determinacy of the risk-free metric)**.** *Under the risk-free metric, for every parameter vector satisfying Assumptions 1–4 and 6 an assessment equilibrium exists, is unique, and has allocation $h_P = h_O = \theta$, for all exit-price beliefs. The date-2 action is then uniquely determined by fundamentals at θ.*

Proof in Appendix A.10. The risk-free metric measures only what the exchange grants, so beliefs about exit yields never enter the rule and the allocation is a function of fundamentals. Default may still occur (for $\varphi y^\ell < m_P(\theta)$ the government defaults on P in the liquidity-crisis state), but it is a *fundamental* event: the metric transmits no beliefs for it to validate.

For the market-value metric, conjecturing $\delta^\ell = (1,0)$ and solving the rule $1-(1-\pi)(1-h_P) = h_O$ jointly with (6) yields the unique candidate allocation

$$h_P^d = \frac{\mathcal{R} - \pi d_O}{d_P + (1-\pi)d_O}, \qquad h_O^d = \pi + (1-\pi)h_P^d, \tag{23}$$

with $m_P^d = R^*(1-h_P^d)d_P$ and $m_O^d = R^*(1-h_O^d)d_O$.

Proposition 5 (Self-fulfilling exit yields under the market-value metric)**.** *Suppose $\pi d_O < \mathcal{R}$, so that the allocation (23) is feasible. Under the market-value metric:*

*(i) (*Repayment assessment.*) $\big(\theta, \delta^\ell = (0,0)\big)$ is an assessment equilibrium iff $m_P(\theta) \le \varphi y^\ell$ and $\Gamma\big(\bar{\ell}(\theta); c_2^c\big) \ge u\big((1-\varphi)y^\ell\big)$.*

(ii) (Default assessment.) $\big((h_P^d, h_O^d), \delta^\ell = (1,0)\big)$ *is an assessment equilibrium iff* $m_P^d > \varphi y^\ell$ *and* $S(h_O^d) \geq 0$. *In it,* $h_P^d < \theta < h_O^d$, *the private exit spread is* $R^*\pi/(1-\pi)$, *and measured haircuts satisfy* $h_P^m = h_O^m = h_O^d$.

(iii) (Multiplicity.) *The conditions in (i) and (ii) hold simultaneously on a parameter set with nonempty interior, on which the two assessments are the only assessment equilibria.*

(iv) (Ranking.) *Whenever both exist, the default assessment carries a strictly lower issue price,* $q_P^d < q_P^{rep}$, *and a strictly lower official recovery,* $m_O^d < m_O(\theta)$.

(v) (Emptiness.) *There exist open parameter sets on which no assessment equilibrium exists.*

Proof in Appendix A.11. *The loop.* The default assessment closes on itself in three steps. Pessimism inflates the measured private haircut by the forward re-default premium; the rule, seeing private "effort" overstated, shifts the *contractual* burden onto the official class ($h_O^d > \theta > h_P^d$); the reallocation enlarges the residual private claim ($m_P^d > m_P(\theta)$) until defaulting on it is in fact optimal, $m_P^d > \varphi y^\ell$—validating the pessimism. The deeper official haircut also erodes emergency access ($\bar{\ell}(h_O^d) < \bar{\ell}(\theta)$), so the assessment survives only while the official claim stays worth servicing, $S(h_O^d) \geq 0$. The optimistic assessment closes symmetrically at θ.

The mechanism should be distinguished from self-fulfilling debt crises à la Calvo (1988), Ayres et al. (2018), and Cole and Kehoe (2000): beliefs there act through the government's budget or rollover constraint; here fundamentals, the envelope $\bar{M}$, and even total face-value relief are identical across assessments, and the identical environment under the risk-free metric has a unique outcome (Proposition 4). The loop runs entirely through the measurement convention.

What the metric reports. Part (ii) contains the section's sharpest indictment. In the default assessment the metric certifies deep and exactly equal treatment—$h_P^m = h_O^m = h_O^d$—while the relief actually granted, in the determinate risk-free sense of Definition 1, is $h_P^d < \theta$ from private creditors against $h_O^d > \theta$ from officials. The entire gap $h_O^d - h_P^d = \pi(1 - h_P^d)$ is anticipated re-default premium booked as if it were granted relief. A comparability regime run on this metric would report its most lopsided outcomes as its most exemplary ones.

Emptiness. Part (v) is the complementary failure: the metric's own reallocation can undermine every belief that would justify it. On the empty region the government defaults selectively at θ, which defeats repayment and total-default beliefs alike (both select θ), while the allocation mandated by selective-default beliefs destroys enough access that servicing officials no longer pays, $S(h_O^d) < 0$. No pure assessment is internally consistent: the negotiation cannot close under the metric without abandoning rational exit pricing or degenerate beliefs—a theorem-shaped rendering of the practitioners' complaint that exit-yield comparability chases its own tail.

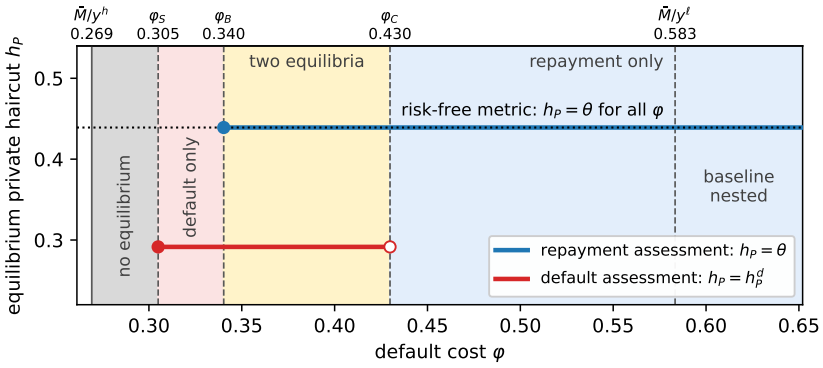


Figure 3: Self-fulfilling assessments under the market-value metric (Table S.3): equilibrium private haircut under the equal-treatment rule, by default cost φ. The dotted line is the risk-free metric's selection θ; φ_S solves $S(h_O^d) = 0$, $\varphi_B = m_P(\theta)/y^\ell$, and $\varphi_C = m_P^d/y^\ell$.

Welfare. The two rankings in (iv) hold wherever the assessments coexist: pessimism lowers the issue price (by Lemma 4(ii), since h_P^m rises from θ to h_O^d) and official recovery. The debtor's comparison is two-sided—the default assessment embeds a state-contingent transfer ($m_P^d - \varphi y^\ell > 0$ arrives in the crisis-upon-crisis state, the excusable-default direction of Grossman and Van Huyck, 1988) but buys it with the deadweight cost φy^ℓ, degraded access, and the lower issue price—and at the parameters of Table S.3 the purchase is a bad one throughout the coexistence window (Supplemental Appendix, Remark S.6).

Illustration. Figure 3 plots the private haircut selected by the market-value rule against the default cost φ, at the baseline parameters with $\bar{\ell}_0 = 0.65$ and $y^h = 1.3$ (Supplemental Appendix, Table S.3; $\pi d_O = 0.125 < \mathcal{R} = 0.263$). Above the serviceability floor $\bar{M}/y^h = 0.269$, four regimes appear in order: no assessment equilibrium up to 0.305; a unique default assessment on $[0.305, 0.340)$, where the metric *forces* the pessimistic allocation although the risk-free rule at θ yields a determinate (if defaulting) outcome; two assessments on $[0.340, 0.430)$; and a unique repayment assessment thereafter, which is the baseline once $\varphi \geq \bar{M}/y^\ell = 0.583$. At $\varphi = 0.40$ the repayment assessment has $h_P = h_O = \theta = 0.439$ and issue price 0.877; the default assessment has contractual haircuts $(h_P^d, h_O^d) = (0.291, 0.646)$, measured haircuts 0.646 for *both* classes, issue price 0.837, and official recovery 0.092 against 0.146.

Scope. Propositions 4 and 5 compare the two metrics under one rule, equal treatment (Definition 6), and assert nothing about other comparability rules under the market-value metric. The domain is deliberately spare (binary re-default with zero recovery and no second renegotiation, on a two-state date-2 endowment, and degenerate (pure) assessments) so the failures are established on open parameter sets, not under continuous repayment risk. Whether mixed assessments (interior re-default probabilities supported at date-2 indifference points) restore ex-

istence on the region of part (v) is left open; what mixing cannot restore is the invariance of Proposition 4: under the market-value metric the allocation varies with the assessed re-default probability, degenerate or not.

5 Analysis of Current CoT Practice

This section asks which criterion of Section 3, if any, supports each rule used or proposed to assess CoT, separating the principle a rule embodies from the procedure that implements it. One caveat applies throughout: the model's one-period contracts do not span duration or debt-service profiles, so where a rule turns on the timing of payments we say what the framework characterizes and what lies outside it.

5.1 Proportional to Reduction in Present Value of Claims

The rule most participants have in mind, explicitly or implicitly, is that every creditor bears the same proportional reduction in the present value of its claims: $h_i = h_j$ for all i, j. This is Criterion 1 (Section 3.1) and the NPV-parity principle at the core of the Paris Club's standard, whose members seek to ensure that their claims are not subordinated, in present value, to those of other creditors. The Club's operative assessment is broader: to the change in NPV it adds the change in duration and the change in nominal debt service over the IMF program period, and it lets a larger effort on one indicator offset a smaller effort on another (Paris Club, n.d.).[24]

Where the other two indicators only describe how a comparable NPV loss is delivered, the operative rule is Criterion 1 and the correspondence is exact. Where they offset a shortfall in NPV effort, it is a multi-indicator standard that Criterion 1 does not characterize. The model cannot express that standard: every claim falls due at the restructuring date and is exchanged for a claim due one period later, so the change in duration is common to all creditors, the change in nominal debt service is affine in h_i, and the three indicators are collinear across creditors. The correspondence is therefore exact for the principle and partial for the procedure, whose further content requires contracts that differ in timing.

The Club applies the same discount rate to all creditors: 5 percent for low-income countries, as in the IMF–World Bank debt sustainability framework, and a range of rates, as a sensitivity analysis, for market-access countries. Section 4 rationalizes the commonality of the rate: at a common gross rate $\bar{R}$ the measured haircut is $1-(R^*/\bar{R})(1-h_i)$, so equal treatment selects the unique allocation of Proposition 4 whatever the rate, while class-specific exit yields can make

[24] In the Club's account, some private creditors prefer early repayment and accept a larger NPV loss through a nominal haircut, while official creditors often cut program-period debt service and extend maturities but avoid a nominal loss.

the assessment multiple or empty (Proposition 5).

A variant, prominent in discussions of multilateral lenders' participation (Zucker-Marques et al., 2024), allocates relief in proportion to outstanding face value; its motivation is that long-maturity concessional claims have a low present value relative to face value, so present-value rules assign them little relief. Exposure shares involve neither the access technology nor origination prices, so the variant is not in general the solution to Criterion 3 or 2. With B_i the face value and V_i the risk-free present value of class i's claims, the variant assigns relief $\mathcal{R}B_i/\sum_j B_j$, hence, while no haircut is capped, $h_i = (\mathcal{R}/\sum_j B_j)(B_i/V_i)$, proportional to the face-to-value ratio.[25] It coincides with Criterion 1 if and only if B_i/V_i is common across creditors, which holds trivially in the model: every claim falls due at the restructuring date, so $B_i = V_i = d_i$, $h_i = \theta$, and no cap binds. Its distinct content, larger haircuts on long-maturity concessional claims, therefore requires the maturity heterogeneity that the model excludes.

5.2 Seniority and Preferred-Creditor Status

Under current practice the IMF and the multilateral development banks (MDBs), such as the World Bank and the African Development Bank, have preferred creditor status (PCS): they stay outside the restructuring perimeter and their claims are not restructured.[26]

The usual rationale is that MDBs are non-profit institutions that lend at below-market terms and keep lending in times of stress, when other creditors withdraw. PCS also supports their credit ratings (G20, 2022), which let them fund and lend cheaply.

The model derives seniority from the access technology of Assumption 4, under which impairing official claims destroys crisis-lending capacity at rate $\lambda > 0$, rather than assuming it through an enforcement advantage, as in Dellas and Niepelt (2016). The derivation is therefore conditional on the technology and on the market structure behind it: with $\lambda = 0$, or with private crisis finance at R^* and a facility priced at R^*, Criterion 3 breaches the official class for every π and a date-1 planner is indifferent over the official haircut (Section 3.3), so exemption would be at most a convention.

The unconditional PCS observed in practice is supported by the model's discretionary outcome: chosen at the restructuring date, the allocation protects the official class whenever protection preserves some access, $\underline{h} < h_{\mathrm{a}}$, and it carries the lowest issue price (Corollary 2(i) and (iii)). Exemption, $h_O = 0$, also requires $\mathcal{R} \leq d_P$. The date-2 default policy also spares the official class as long as the relationship retains value, $S(h_O) \geq 0$ (Lemma 3). Under commit-

[25] The formula respects $h_i \leq 1$ for every class only if $\mathcal{R} \leq (\sum_j B_j) \min_i V_i/B_i$, a bound stricter than aggregate feasibility, $\mathcal{R} \leq \sum_j V_j$, unless the ratios are common; beyond it the class with the lowest value-to-face ratio is written off in full and the excess must be reallocated by a rule that proportionality does not contain.

[26] PCS is a market practice, in general not written into treaties or loan contracts (Cordella and Powell, 2021). It has nonetheless been respected historically, with other creditors agreeing to the exemption.

ment PCS is not categorical: Criterion 3 protects officials fully only when post-restructuring crisis risk is high, $\pi \geq \max\{\hat{\pi}, \pi_p\}$, and breaches them when that risk is low or the overhang is deep, with graduated participation in between (Proposition 3). The model also qualifies the argument that without PCS "no one would lend during crises": under the access technology (8), equal treatment destroys emergency capacity only if the common haircut reaches h_{a}, and at the baseline of Figure 2 it cuts capacity from $\bar{\ell}_0 = 0.50$ to $\bar{\ell}(\theta) = 0.17$.

5.3 Ranking Creditors by Grant Element and Sequential Participation

A recurring proposal ties each creditor's burden to the concessionality of its lending terms: creditors are ranked by grant element, and relief is sourced from the least concessional first. Can the model rationalize such an ordering? Criterion 1 cannot, as it sees claims alone. Criterion 2 can, and its haircuts are ordered inversely to the concession each class granted at origination: for any netting share $\kappa > 0$ the private class, whose spread pre-paid its anticipated loss, takes the larger haircut, and the official haircut falls with g_O until the official class is exempt (the closed form in Proposition 2(i)); with several official classes, the rule ranks them inversely by concessionality (Supplemental Appendix, Corollary S.2). The two rationales usually given for the ordering, that a grant element is relief granted in advance and that lenders on harder terms were paid for risk and should carry more of it, are the two terms of the differential $h_P^{\mathrm{nw}} - h_O^{\mathrm{nw}} = \kappa(F h_P^{\mathrm{nw}} + g_O)$.

Diwan, Harnoys-Vannier, and Kessler (2023) make the ordering an explicit rule. A creditor's grant element is one minus the present value of the debt service it is owed, at a common discount rate, relative to the principal it advanced: the share of the loan that is not repaid in present value. Creditors are ranked from the lowest grant element to the highest. The least concessional creditor bears relief alone, which raises its post-relief grant element, until that grant element reaches the next creditor's; the two then share further relief, keeping their grant elements equal, until they reach the third creditor's; and so on. All participating creditors end at a common post-relief grant element, and creditors whose grant element already exceeds it are exempt.

In the model this is Criterion 2 at full netting, with each class's netted loss measured per unit of the funds it advanced rather than per unit of the face value it is owed. Class i advanced $q_i d_i$ against the promise d_i and holds, after relief, a claim worth $(1 - h_i)d_i$, so its post-relief grant element at the risk-free rate is

$$1 - \frac{(1 - h_i)d_i}{R^* q_i d_i} \;=\; \frac{h_i - \chi_i}{R^* q_i} \;=\; -\frac{x_i}{R^* q_i}, \tag{24}$$

which is the netted loss of Definition 4 at $\kappa = 1$, divided by the funds advanced per unit of face

value, R^*q_i in date-1 value.[27] Before relief it is $g_O/(1+g_O)$ for the official class, the grant element of Section 2.2, and $-Fh_P^e/(1-Fh_P^e)$ for the private class, negative because the spread placed the claim below the risk-free price. The private class is therefore always the less concessional, and the rule assigns that private creditors bear relief alone until their grant element rises to $g_O/(1+g_O)$, and beyond that point both classes share it at a common grant element. Because the private grant element depends on the haircut investors anticipate, through q_P, the allocation is a fixed point, as in Definition 4. It is unique, and the official class is exempt if and only if $\mathcal{R} \leq g_O d_P/(1-F+g_O)$ (Supplemental Appendix, Proposition S.3).

We note that the denominator used is consequential. Netting per unit of face value (as in Proposition 2 at $\kappa = 1$), exempts the official class for all $\mathcal{R} \leq g_O d_P/(1-F)$, a weakly larger set, and assigns it a weakly smaller haircut, strictly unless both rules exempt it (more details in Suppl. Appendix, Proposition S.3). Dividing by funds advanced scales the private loss up, because the spread means investors advanced less than the risk-free value of their claim, and the official loss down, because the concession means the official lender advanced more; equalizing the two then takes more from the official class.[28] The correspondence is to the rule's sequencing and normalization, not to its measurement of concessionality. In the model a one-period zero-coupon claim carries its concessionality in the issue price alone; the DHK's proposal exercise instead computes grant elements from interest, maturity, and grace at a 5 percent discount rate, applies group averages to nominal stocks, and sets negative grant elements to zero before ranking, a floor that would remove the spread from the numerator of (24) but not from its denominator.

Can Criterion 3 also rank creditors by their concessionality? The marginal cost of official relief, σ_O in (16), is proportional to the wedge Δ_ℓ, the facility's grant element plus the scarcity premium of (11),[29] so on the interior regime the official haircut falls as the facility becomes more concessional (Corollary 1). However, the concessionality of the outstanding stock is sunk: it enters only as an income effect, the R_O comparative static of Corollary 1, and with $\lambda = 0$ Criterion 3 breaches the official class in full however concessional its claims (Section 3.3), where the no-windfall rule still protects it. Inherited grant elements implement the debtor-cost allocation only when both criteria exempt the official class, or at isolated parameter values; they never implement full breach (Supplemental Appendix, Proposition S.1); and with several official classes they can order lenders in the opposite way, since the lender with the most concessional claims is

[27] The DHK's proposal measures grant elements against the nominal loan stock, which is the principal, and reduces the present value of debt service by the relief granted; the model's face value d_i is the promised debt service and $q_i d_i$ the principal, so (24) is the proposal's post-relief grant element at the risk-free rate.

[28] At the baseline of Figure 2, full netting assigns the official class a haircut of 37.1 percent per unit of face value and 40.0 percent per unit of funds advanced, against $\theta = 43.9$ percent under equal treatment and 6.9 percent under Criterion 3: the denominator moves the allocation, and the contrast with Criterion 3 survives either one.

[29] The wedge on the conditioned development lending of bilateral creditors decomposes in the same way (Supplemental Appendix S.2).

the most protected under the no-windfall rule and the first source of relief under Criterion 3 if it does not condition its future flows (Corollary S.2 and Proposition S.4(b)). A prospective grant element, that of the facility, is not enough either: σ_O depends on it only through the wedge, together with the scarcity premium, the access response λ, the crisis probability, and the pricing margin σ_P against which it is weighed. Agreement in the ordering of creditors, finally, is weaker than agreement in the allocation: at the baseline of Figure 2 both criteria haircut the private class more, with official haircuts of 37.1 and 6.9 percent.

6 Conclusions

Within a deliberately minimal dynamic model, we have characterized three criteria for allocating debt relief. "Comparable" treatment does not specify what should be comparable. With heterogeneous creditors, there are at least three natural candidates: losses on claims, losses net of the compensation embedded at origination, and the welfare cost to the debtor of the relief obtained through each class. Basing the allocation on them generally leads to different outcomes.

A second lesson concerns measurement. Within the same environment, the risk-free metric always induces a unique equal-treatment assessment, while the market-value metric—exit-yield valuation used as the rule's input—can generate two self-fulfilling assessments or none. Market-value haircuts remain the correct measure of expected investor losses; the failure arises only when an endogenous valuation is asked to govern the allocation that determines it.

As explained in the main text, promising future extensions to this framework would study the welfare implications of defaults in debt denominated in local versus foreign currency, or characterize a welfare comparison of the three criteria, both from the debtor and the creditors' perspective.

A Proofs

A.1 Proof of Lemma 1

By Assumption 2(iii), the government's resources at every date and state after date 0 are its endowment net of contractual payments, plus emergency borrowing where the facility is open. On the repayment branch, $c_1 = y^n - D$ and the government enters date 2 debt-free with capacity $\bar{\ell}_0$, so the branch's contribution to expected utility is independent of the allocation. On the restructuring branch, $c_1 = y^c$; date-2 resources are $y_2 - \bar{M}$ plus $\ell \le \bar{\ell}(h_O)$ in state y^ℓ, with $m_P + m_O = \bar{M}$ by (4)–(6); scarcity implies the cap binds whenever capacity is positive, $u'(c_2^c + \bar{\ell}(h_O)) \ge u'(c_2^c + \bar{\ell}_0) > \beta R_E$, so $\ell = \bar{\ell}(h_O)$ and (c_2, c_3) in that state vary with the allocation exactly through $\bar{\ell}(h_O)$, as stated. Settlement is $\bar{y}_3 - R_E\ell$, whose cost, $\beta R_E \ell$ in date-2 utils, is netted out in the definition of A in (9). Hence, in date-1 utils and up to a constant, the branch contributes $u(y^c) + \beta\big[\pi\{u(c_2^c) + A(\bar{\ell}(h_O))\} + (1-\pi)u(y^h - \bar{M})\big]$, and the only allocation-dependent term is $A(\bar{\ell}(h_O))$. At date 0, substituting $h_P(h_O) = (\mathcal{R} - h_O d_O)/d_P$ into (2)–(3):

$$c_0 = y_0 + \beta^*\big(1 - F\tfrac{\mathcal{R} - h_O d_O}{d_P}\big)d_P + \tfrac{d_O}{R_O} = y_0 + \beta^*\big(d_P - F\mathcal{R} + F h_O d_O\big) + \tfrac{d_O}{R_O},$$

with $\partial c_0/\partial h_O = \beta^* F d_O > 0$. Collecting the allocation-dependent terms and discounting to date 0 yields (12). □

A.2 Proof of Proposition 1

(Sufficiency.) The proportional rule satisfies (A1) and (A3) trivially. For (A2): both problems have the same estate E and aggregate claim $D \equiv \sum_i c_i$, and the proportional award to any group is (E/D) times its total claim, which the reconfiguration preserves.

(Necessity.) If $E = \sum_i c_i$, the bounds and budget balance force $f_i = c_i$, which is proportional. So let $E < \sum_i c_i \equiv D$.

Step 1 (rational claims). Let $c \in \mathbb{Q}^N_{++}$ and choose $n \in \mathbb{N}$ with $nc_i \in \mathbb{N}$ for all i (so $nD \in \mathbb{N}$). (a) *Full atomization:* let $\mathcal{A}$ be the problem with nD accounts of claim $1/n$ each. By (A1) every atom receives $E/(nD)$, feasible since $E \le D$. (b) *The i-star problem:* let $\mathcal{A}_i$ keep account i intact plus $n(D - c_i)$ atoms. $\mathcal{A}$ is obtained from $\mathcal{A}_i$ by replacing $S = \{i\}$ with its nc_i atoms (total c_i preserved, complement untouched), so by (A2) $f_i(\mathcal{A}_i) = nc_i \cdot E/(nD) = c_i E/D$. (c) *The original problem:* $\mathcal{A}_i$ is obtained from (N, c, E) by replacing $S = N \setminus \{i\}$ with its atomization (total $D - c_i$ preserved, account i untouched), so by (A2) and budget balance $\sum_{j \ne i} f_j(N, c, E) = E - c_i E/D$, whence $f_i(N, c, E) = c_i E/D$. Since i was arbitrary, f is proportional on rational claims.

Step 2 (real claims). For $E < \sum_i c_i$, take rational $c^k \to c$ with $E < \sum_i c_i^k$; (A3) gives $f_i = \lim_k c_i^k E / \sum_j c_j^k = c_i E / \sum_j c_j$.

Application. At the restructuring node $(E;d) = (\beta^*\bar{M}; d_P, d_O)$ and $a_i = \beta^* m_i$, so $a_i = d_i E/D$ delivers $h_i = 1 - E/D = \mathcal{R}/D = \theta$ for both classes; the argument nowhere used $|N| = 2$. □

A.3 Proof of Proposition 3

Part (i). Since $h_{\mathrm{a}} > 0$, $\underline{h} \geq h_{\mathrm{a}}$ if and only if $\bar{\ell}_0 \leq \lambda(\mathcal{R} - d_P)$, as stated. In that case $\bar{\ell}(h) = 0$ for every feasible h, so $\mathcal{W}(h) = u(c_0(h))$ up to the constant of Lemma 1, with $\mathcal{W}'(h) = u'(c_0(h))\beta^* F d_O > 0$; the unique maximizer is $\overline{h}$ for every π.

Part (ii): shape of $\mathcal{W}$. Let $\underline{h} < h_{\mathrm{a}}$ and write $\bar{h}_c \equiv h_{\mathrm{a}} \wedge \overline{h}$ for the upper end of the access-preserving branch $[\underline{h}, \bar{h}_c]$, extending Δ_ℓ to $\bar{h}_c$ by continuity, $\Delta_\ell(\bar{h}_c) \equiv \Delta_\ell(\bar{h}_c^-) > 0$ by scarcity (Assumption 4). On the branch, $\bar{\ell}(h) = \bar{\ell}_0 - \lambda d_O h \in [0, \bar{\ell}_0]$ and scarcity implies the cap binds, so $A(\bar{\ell}(h)) = u(c_2^c + \bar{\ell}(h)) - \beta R_E \bar{\ell}(h) - u(c_2^c)$ is C^2 with $\frac{d}{dh} A(\bar{\ell}(h)) = -\lambda d_O \, \Delta_\ell(h)$, and

$$\mathcal{W}'(h) = u'(c_0(h)) \, \beta^* F d_O - \beta^2 F \pi \lambda d_O \, \Delta_\ell(h), \tag{A.1}$$

$$\mathcal{W}''(h) = u''(c_0(h)) \, (\beta^* F d_O)^2 + \beta^2 F \pi \lambda^2 d_O^2 \, u''\big(c_2^c + \bar{\ell}(h)\big) < 0, \tag{A.2}$$

using $\Delta_\ell'(h) = -u''(c_2^c + \bar{\ell}(h))\lambda d_O > 0$. Hence $\mathcal{W}$ is strictly concave on the branch and has a unique maximizer there, $h_O^\circ(\pi)$. If $h_{\mathrm{a}} < \overline{h}$, then on $[h_{\mathrm{a}}, \overline{h}]$, $A(\bar{\ell}(h)) \equiv 0$ and $\mathcal{W}'(h) = u'(c_0(h))\beta^* F d_O > 0$, so $\mathcal{W}$ is strictly increasing; $\mathcal{W}$ is continuous at h_{a} since $A(x) \to 0$ as $x \downarrow 0$, and its right derivative there exceeds its left by $\beta^2 F \pi \lambda d_O \Delta_\ell(h_{\mathrm{a}}^-) > 0$.

The branch maximizer. On the branch write $G(h;\pi) \equiv u'(c_0(h))\beta^* - \beta^2 \pi \lambda \Delta_\ell(h)$, so that $\mathcal{W}' = F d_O \, G$; by (A.2) $G_h < 0$, and $G_\pi = -\beta^2 \lambda \Delta_\ell < 0$. Using $\beta = \beta^*$, $G(h;\pi) = 0$ if and only if $\pi = u'(c_0(h))/(\beta\lambda\Delta_\ell(h))$, so $G(\underline{h};\pi_p) = 0$ and $G(\bar{h}_c;\pi_a) = 0$, and $0 < \pi_a$; then $G(\underline{h};\pi_a) > G(\bar{h}_c;\pi_a) = 0 = G(\underline{h};\pi_p)$ and $G_\pi < 0$ give $\pi_a < \pi_p$. By strict concavity, $h_O^\circ(\pi) = \underline{h}$ iff $G(\underline{h};\pi) \leq 0$ iff $\pi \geq \pi_p$; $h_O^\circ(\pi) = \bar{h}_c$ iff $G(\bar{h}_c;\pi) \geq 0$ iff $\pi \leq \pi_a$; and for $\pi \in (\pi_a, \pi_p)$, $h_O^\circ(\pi)$ is the unique interior root of $G(\cdot;\pi) = 0$, which is (15) after dividing $\mathcal{W}' = 0$ by $F d_O$, with $dh_O^\circ/d\pi = -G_\pi/G_h < 0$ by the implicit function theorem. Being the unique maximizer of a jointly continuous function on a fixed compact set, $h_O^\circ(\cdot)$ is continuous by the maximum theorem.

Case $\overline{h} \leq h_{\mathrm{a}}$. The branch is the whole feasible set, so $h_O^*(\pi) = h_O^\circ(\pi)$ is single-valued and continuous, and the corner characterization gives full breach for $\pi \leq \pi_a$, an interior and strictly decreasing allocation for $\pi \in (\pi_a, \pi_p)$, and preferred creditor status for $\pi \geq \pi_p$: the statement with $\hat{\pi} = \pi_a$.

Case $\overline{h} > h_{\mathrm{a}}$. Now $\bar{h}_c = h_{\mathrm{a}}$ and $\bar{\ell}(\overline{h}) = 0$, so $\mathcal{W}(\overline{h};\pi) = u(c_0(\overline{h}))$ for every π, and $\Delta(\pi) = V(\pi) - u(c_0(\overline{h}))$ with $V(\pi) \equiv \max_{h \in [\underline{h}, h_{\mathrm{a}}]} \mathcal{W}(h;\pi) = \mathcal{W}(h_O^\circ(\pi);\pi)$. For each h, $\mathcal{W}(h;\cdot)$ is affine in π with slope $\beta^2 F A(\bar{\ell}(h)) \geq 0$, so V is convex, finite, hence continuous, and nondecreasing, and for $\pi' > \pi$, $V(\pi') \geq \mathcal{W}(h_O^\circ(\pi);\pi') = V(\pi) + \beta^2 F A(\bar{\ell}(h_O^\circ(\pi)))(\pi' - \pi)$. On $(0, \pi_a]$, $h_O^\circ(\pi) = h_{\mathrm{a}}$, so $V(\pi) = u(c_0(h_{\mathrm{a}}))$ and $\Delta(\pi) = u(c_0(h_{\mathrm{a}})) - u(c_0(\overline{h})) < 0$, as c_0 is strictly increasing. For $\pi > \pi_a$,

$h_O^\circ(\pi) < h_a$, so $A(\bar{\ell}(h_O^\circ(\pi))) > 0$ and the inequality makes V strictly increasing on (π_a, ∞), hence on $[\pi_a, \infty)$ since V is nondecreasing. For $\pi \geq \pi_p$, $V(\pi) = u(c_0(\underline{h})) + \beta^2 F \pi A(\bar{\ell}(\underline{h}))$ with $A(\bar{\ell}(\underline{h})) > 0$, so $\Delta(\pi) \to \infty$. Therefore Δ is continuous, constant and strictly negative on $(0, \pi_a]$, strictly increasing on $[\pi_a, \infty)$, and unbounded above: it has a unique zero $\pi_b \in (\pi_a, \infty)$, with $\Delta < 0$ on $(0, \pi_b)$ and $\Delta > 0$ on (π_b, ∞). As $\mathcal{W}(\overline{h})$ is the maximum of $\mathcal{W}$ on $[h_a, \overline{h}]$ and $V(\pi)$ its maximum on $[\underline{h}, h_a]$, each attained uniquely, the global maximizer is $\overline{h}$ for $\pi < \pi_b$, $h_O^\circ(\pi)$ for $\pi > \pi_b$, and both at $\pi = \pi_b$; since $h_O^\circ(\pi_b) \leq h_a < \overline{h}$, the allocation jumps down at π_b. The regimes follow from the corner characterization: for $\pi \in (\pi_b, \pi_p)$, nonempty iff $\pi_b < \pi_p$, $h_O^\circ(\pi)$ is interior, since $\pi_a < \pi_b < \pi < \pi_p$, and strictly decreasing; for $\pi \geq \max\{\pi_b, \pi_p\}$, $h_O^\circ(\pi) = \underline{h}$; and if $\pi_b \geq \pi_p$ the allocation passes from $\overline{h}$ to $\underline{h}$ at π_b. This is the statement with $\hat{\pi} = \pi_b$. Finally, Δ is strictly increasing on $[\pi_a, \infty) \ni \pi_p$ with $\Delta(\pi_b) = 0$, so $\pi_b \geq \pi_p$ iff $\Delta(\pi_p) \leq 0$. □

A.4 Proof of Corollary 1

Regard G, defined in the proof of Proposition 3, as a function of h and of the parameters $z \in \{\pi, \lambda, R_E, R_O, \mathcal{R}\}$. By the implicit function theorem applied to $G(h_O^\circ; z) = 0$, and since $G_h < 0$, the sign of $\partial h_O^\circ / \partial z$ is the sign of G_z. Using $\partial \bar{\ell} / \partial \lambda = -h_O^\circ d_O$,

$$G_\pi = -\beta^2 \lambda \Delta_\ell < 0, \qquad G_\lambda = -\beta^2 \pi \left[\Delta_\ell + \lambda \left|u''(c_2^c + \bar{\ell})\right| h_O^\circ d_O\right] < 0,$$
$$G_{R_E} = \beta^3 \pi \lambda > 0, \qquad G_{R_O} = u''(c_0)\, \beta^* \cdot \left(-d_O / R_O^2\right) > 0.$$

For $\mathcal{R}$ with y^ℓ fixed and $\bar{M} = (D - \mathcal{R})/\beta^*$: $\partial c_0 / \partial \mathcal{R} = -\beta^* F < 0$ and $\partial c_2^c / \partial \mathcal{R} = -\partial \bar{M} / \partial \mathcal{R} = R^* > 0$, so

$$G_{\mathcal{R}} = \underbrace{-u''(c_0)\, \beta^{*2} F}_{>0} + \underbrace{\left(-\beta^2 \pi \lambda\right) u''\left(c_2^c + \bar{\ell}\right) R^*}_{>0} > 0$$

unambiguously. □

A.5 Proof of Corollary 2

At the date-1 crisis node q_P, hence c_0, is predetermined, and by the proof of Lemma 1 the only allocation-dependent term in continuation utility is $\beta \pi A(\bar{\ell}(h_O))$. Under scarcity the cap binds at every $x \in (0, \bar{\ell}_0]$, so $A(x) = \int_0^x \left[u'(c_2^c + s) - \beta R_E\right] ds$ has a strictly positive integrand; hence A is strictly increasing on $[0, \bar{\ell}_0]$, and $A(\bar{\ell}(\cdot))$ is strictly decreasing on $[\underline{h}, h_a)$ and equal to zero on $[h_a, \overline{h}]$. Let $\underline{h} < h_a$. Then $A(\bar{\ell}(\underline{h})) > A(\bar{\ell}(h_O))$ for every feasible $h_O > \underline{h}$, whether h_O lies below or above h_a, so $h_O = \underline{h}$ is the unique date-1 optimum. Hence any rule whose crisis allocation has $h_O > \underline{h}$ admits the strictly profitable one-shot deviation to $\underline{h}$, which is (ii); and sequential rationality forces $h_O = \underline{h}$, hence $h_P = \min\{1, \mathcal{R}/d_P\}$ by (6), in every equilibrium, which with

rational expectations in (2) gives the price in (iii) and proves (i). For (iii), any feasible rule with crisis allocation $h'_O \geq \underline{h}$ is priced at $\beta^*\big(1-Fh_P(h'_O)\big) \geq \beta^*(1-Fh_P(\underline{h}))$, strictly if $h'_O > \underline{h}$, since $h_P(h_O) = (\mathcal{R} - h_O d_O)/d_P$ is strictly decreasing by (6). If instead $\underline{h} \geq h_{\mathrm{a}}$, then $A(\bar{\ell}(h_O)) = 0$ at every feasible allocation and the date-1 planner is indifferent. □

A.6 Proof of Lemma 2

From Eq. (2), $q_P = \beta^*(1 - Fh_P^e)$, so $R^* q_P = 1 - Fh_P^e$ and $\chi_P = 1 - R^* q_P = Fh_P^e$. The official lender disburses d_O/R_O against face d_O (Assumption 2), i.e., $q_O = 1/R_O$, so $\chi_O = 1 - R^*/R_O = -g_O$.

(i) In the no-restructuring state the claim pays 1 per unit face at date 1, so the excess payoff over the capitalized cost of funds is $1 - R^* q_i = \chi_i$. In the restructuring state it pays $\beta^* m_i/d_i = 1 - h_i$ (Definition 1; restructured claims default-free by Assumption 3), so the excess payoff is $(1 - h_i) - R^* q_i = \chi_i - h_i = x_i$.

(ii) Under any anticipated rule with $h_P^e = h_P$, $\mathbb{E}$[date-1 value per unit face] $= (1 - F) \cdot 1 + F(1 - h_P) = 1 - Fh_P = R^* q_P$, hence the expected excess payoff of class P is $(1-F)\chi_P + F(\chi_P - h_P) = \chi_P - Fh_P = 0$. For the official class it is $\chi_O - Fh_O = -(g_O + Fh_O) \leq 0$, with equality iff $g_O = 0$ and $h_O = 0$ (as $F > 0$ and $g_O \geq 0$). □

A.7 Proof of Proposition 2

Existence and uniqueness. Since $\mathcal{R} \leq d_P$, feasible allocations are indexed by $h_O \in [0, \overline{h}]$, with $h_P = (\mathcal{R} - h_O d_O)/d_P$. Let investors anticipate such an allocation, so that the measured compensations are $\chi_P = Fh_P$ and $\chi_O = -g_O$ (Lemma 2). Given these compensations, the netted gap of Definition 4 at a feasible allocation (h'_P, h'_O), $(h'_P - \kappa\chi_P) - (h'_O - \kappa\chi_O)$, is strictly decreasing in h'_O along (6), and at the anticipated allocation itself it equals

$$\psi(h_O) \equiv (1 - \kappa F)\, h_P - h_O - \kappa g_O.$$

The rule minimizes the absolute value of the gap on $[0, \overline{h}]$: it selects the zero of the gap if there is one, $h'_O = 0$ if the gap is negative throughout, and $h'_O = \overline{h}$ if it is positive throughout. The anticipated allocation is an equilibrium exactly when the rule returns it, that is, if and only if $\psi(h_O) = 0$, or $h_O = 0$ and $\psi(0) \leq 0$, or $h_O = \overline{h}$ and $\psi(\overline{h}) \geq 0$. As a function of the anticipated h_O, ψ is affine and strictly decreasing, since $\kappa F < 1$; $\psi(\overline{h}) \leq h_P - \overline{h} < 0$, because $\theta < \overline{h}$ and hence $h_P < \theta$ at $h_O = \overline{h}$; and $\psi(0) = (1 - \kappa F)\mathcal{R}/d_P - \kappa g_O$ has the sign of $\mathcal{R} - \kappa(F\mathcal{R} + g_O d_P)$, that is, of $\kappa^* - \kappa$. Hence there is exactly one equilibrium: for $\kappa < \kappa^*$ it is the zero of ψ, which lies in $(0, \overline{h})$, and for $\kappa \geq \kappa^*$ it is $h_O = 0$ with $h_P = \mathcal{R}/d_P$, which at $\kappa = \kappa^*$ is also the zero of ψ.

(i) The equation $\psi(h_O) = 0$ is (14) at the equilibrium compensations, $h_P - h_O = \kappa(Fh_P + g_O)$.

Solved for $h_O = (1-\kappa F)h_P - \kappa g_O$ and substituted into (6), it gives $h_P\,[d_P + (1-\kappa F)d_O] = \mathcal{R} + \kappa g_O d_O$, hence the pair displayed in part (i) for $\kappa \le \kappa^*$, which is (θ, θ) at $\kappa = 0$. The numerator of h_P^{nw} is positive and nondecreasing in κ and its denominator is positive and strictly decreasing, as $F > 0$ and $\kappa F < 1$; so h_P^{nw} is strictly increasing, and $h_O^{\mathrm{nw}} = (\mathcal{R} - h_P^{\mathrm{nw}} d_P)/d_O$ strictly decreasing. For $\kappa \in [\kappa^*, 1]$ the equilibrium is the corner $h_O = 0$, where the bound binds.

(ii) By Lemma 2(i), $x_P - x_O = (\chi_P - \chi_O) - (h_P - h_O)$, which at an anticipated allocation is $g_O - (1-F)h_P + h_O$. For $\kappa \le \kappa^*$, (14) gives $x_P - x_O = (1-\kappa)(F h_P^{\mathrm{nw}} + g_O)$, and $F h_P^{\mathrm{nw}} + g_O = (F\mathcal{R} + g_O D)/[d_P + (1-\kappa F)d_O]$, because κ cancels in $F(\mathcal{R} + \kappa g_O d_O) + g_O[d_P + (1-\kappa F)d_O] = F\mathcal{R} + g_O D$. The windfall gap thus equals $F\theta + g_O$ at $\kappa = 0$ and is strictly decreasing on this range, since

$$\frac{d}{d\kappa}\,\frac{1-\kappa}{d_P + (1-\kappa F)d_O} = -\frac{d_P + (1-F)d_O}{[d_P + (1-\kappa F)d_O]^2} < 0.$$

For $\kappa \in [\kappa^*, 1]$, $h_P = \mathcal{R}/d_P$ and $h_O = 0$ give $x_P - x_O = g_O - (1-F)\mathcal{R}/d_P$, constant in κ. The two expressions agree at κ^*, where the two allocations coincide, so the windfall gap is continuous. At $\kappa = 1$ it is zero if $\kappa^* \ge 1$, and it is $g_O - (1-F)\mathcal{R}/d_P > 0$ if $\kappa^* < 1$, that is, if $g_O d_P > (1-F)\mathcal{R}$. Finally, along (6) the expression $g_O - (1-F)h_P + h_O$ is strictly increasing in h_O, so its minimum over feasible allocations that investors anticipate is $g_O - (1-F)\mathcal{R}/d_P$, at $h_O = 0$; when $\kappa^* < 1$ this is positive, and closing the gap would require $h_O < 0$. □

A.8 Proof of Lemma 3

Throughout, (m_P, m_O) satisfy $m_P + m_O = \bar{M}$ (by (4)–(6)), so $0 \le m_P \le \bar{M}$. All values are computed with the extended u of Assumption 6: an action whose implied consumption is nonpositive (after emergency borrowing, where available) has value $-\infty$, and every dominance claim below holds *a fortiori* when the dominated action is infeasible. Thus $\Gamma(x; c)$ in (18) equals $-\infty$ when no $\ell \in [0, x]$ yields positive consumption, and its maximum is otherwise attained, by the Inada condition. Repay-all is feasible in y^h ($y^h > \bar{M}$, since $\varphi \ge \bar{M}/y^h$ and $\varphi < 1$), and repay-all and default-on-both are feasible in y^ℓ ($c_2^c > 0$ and $(1-\varphi)y^\ell > 0$), so each state has a feasible action, and the convention only prices out deviations that could not be consumed.

State y^h. There is no emergency borrowing in y^h; the four actions yield $u(y^h - \bar{M})$, $u((1-\varphi)y^h - m_O)$, $u((1-\varphi)y^h - m_P)$, and $u((1-\varphi)y^h)$. Repay-all dominates default-on-P iff $\varphi y^h \ge m_P$ and dominates default-on-both iff $\varphi y^h \ge \bar{M}$; both follow from $\varphi \ge \bar{M}/y^h$ and $m_P \le \bar{M}$. Default-on-O-only is dominated by default-on-both (it pays $m_P \ge 0$ for no benefit, the default cost being incurred either way).

State y^ℓ. Write $\bar{x} = \bar{\ell}(h_O)$. Net of terms common to all actions, repay-all, default-on-P-only, and default-on-both are valued by $V_{RR} = \Gamma(\bar{x}; c_2^c)$, $V_{DP} = \Gamma(\bar{x}; b)$, and $V_{DD} = u(c_2^c + \bar{M} - \varphi y^\ell) = u((1-\varphi)y^\ell)$, with b as in (19); since $b + m_O = (1-\varphi)y^\ell$, $V_{DP} - V_{DD} = S(h_O)$. Two properties of Γ

are used. First, $\Gamma(x;\cdot)$ is strictly increasing where finite, and trivially nondecreasing otherwise: for $c' > c$ with $\Gamma(x;c) > -\infty$, with ℓ_c optimal at c, $\Gamma(x;c') \geq u(c'+\ell_c) - \beta R_E \ell_c > u(c+\ell_c) - \beta R_E \ell_c = \Gamma(x;c)$. Second, $\Gamma(x;b) < u(b+x)$ whenever $x > 0$ and $b + x > 0$: if $b > 0$, then either the optimal $\ell^* > 0$, so $\Gamma = u(b+\ell^*) - \beta R_E \ell^* < u(b+\ell^*) \leq u(b+x)$, or $\ell^* = 0$, so $\Gamma = u(b) < u(b+x)$; if $b \leq 0 < b+x$, every feasible ℓ exceeds $-b \geq 0$, so each value $u(b+\ell) - \beta R_E \ell$ is strictly below $u(b+x)$, and the attained maximum inherits the strict inequality; if $b + x \leq 0$, $\Gamma(x;b) = -\infty$. *(a)* Action $(0,1)$ closes the window and yields $u((1-\varphi)y^\ell - m_P) \leq V_{DD}$, strictly if $m_P > 0$; if $m_P = 0$ it entails the same payments as $(1,1)$ and is recorded as $(1,1)$. *(b)* V_{RR} and V_{DP} share the cap $\bar{x}$ and differ only in the base, by $m_P - \varphi y^\ell$; by the first property, $V_{DP} \gtrless V_{RR}$ as $m_P \gtrless \varphi y^\ell$. *(c)* $V_{DP} \geq V_{DD}$ iff $S(h_O) \geq 0$. With the tie convention, $\delta = (0,0)$ iff $V_{RR} \geq \max\{V_{DP}, V_{DD}\}$, $\delta = (1,0)$ iff $V_{DP} > V_{RR}$ and $V_{DP} \geq V_{DD}$, and $\delta = (1,1)$ otherwise; by (b), (c), and the definitions of V_{RR} and V_{DD}, these are the conditions of the lemma.

Servicing bound. $S(h_O) \geq 0$ only if $\bar{\ell}(h_O) > m_O$ or $\bar{\ell}(h_O) = m_O = 0$. Suppose $\bar{x} \leq m_O$ with $(\bar{x}, m_O) \neq (0,0)$; note $b + m_O = (1-\varphi)y^\ell > 0$. If $\bar{x} > 0$, then either $b + \bar{x} \leq 0$ and $\Gamma(\bar{x};b) = -\infty < u(b+m_O)$, or $\Gamma(\bar{x};b) < u(b+\bar{x}) \leq u(b+m_O)$ by the second property; if $\bar{x} = 0 < m_O$, then $\Gamma(0;b) = u(b) < u(b+m_O)$ (with $u(b) = -\infty$ if $b \leq 0$). In both cases $S(h_O) < 0$. The corner $\bar{x} = m_O = 0$ is reached only at $h_O = 1$ with $\bar{\ell}_0 \leq \lambda d_O$; there $b = (1-\varphi)y^\ell$, $V_{DP} = \Gamma(0;b) = u(b) = V_{DD}$, and $S(h_O) = 0$: default-on-P and default-on-both yield identical consumption.

Nesting. If $\varphi \geq \bar{M}/y^\ell$: $m_P - \varphi y^\ell \leq \bar{M} - \varphi y^\ell \leq 0$, so $V_{DP} \leq V_{RR}$ by the first property; and $V_{DD} = u(c_2^c + \bar{M} - \varphi y^\ell) \leq u(c_2^c) \leq \Gamma(\bar{x}; c_2^c) = V_{RR}$. With y^h covered above, $\delta = (0,0)$ everywhere: the performance clause of Assumption 3 is restored, the default cost is never incurred, and (17) reduces to the cap of Assumption 4, so the model of Sections 2 and 3 obtains unchanged. □

A.9 Proof of Lemma 4

(i) Immediate from (21): $h_i^m - h_i = \pi \delta_i^\ell (1 - h_i) \geq 0$, zero iff $\delta_i^\ell = 0$ or $h_i = 1$; subtracting the two expressions at common δ gives the compression identity. *(ii)* Consider one unit of old class-P face value at date 0. In the no-crisis state it pays 1 at date 1. In the crisis state it is exchanged for m_P/d_P units of new claims; by the law of one price each is worth $q_P^x = \beta^*(1 - \pi\delta_P^\ell)$ at date 1, the discounted expected date-2 payoff. Hence

$$q_P = \beta^* \Big[(1-F) + F \, \frac{q_P^x m_P}{d_P} \Big] = \beta^* \big[(1-F) + F(1 - \pi\delta_P^\ell)(1 - h_P) \big] = \beta^* \big(1 - F h_P^m \big),$$

using $m_P/d_P = (1-h_P)/\beta^*$; with $\delta_P^\ell = 0$ this is Eq. (2). *(iii)* The date-1 expected loss per unit of old face is $1 - q_P^x m_P/d_P = h_P^m$ by definition; the Sturzenegger–Zettelmeyer haircut discounts new flows at the exit yield, $m_i/Y_i^x = q_i^x m_i$, and old claims due at date 1 have value d_i at any

rate, giving (21). □

A.10 Proof of Proposition 4

Under the risk-free metric the rule reads $h_P = h_O$ regardless of (δ^ℓ, q^x); with (6) this has the unique solution $h_P = h_O = \theta$, feasible since $\theta = \mathcal{R}/D \in (0,1)$. Condition (iii) of Definition 6 then determines δ^ℓ uniquely (Lemma 3 with the tie convention), and (i) determines prices. Existence, uniqueness, and belief-invariance follow. □

A.11 Proof of Proposition 5

(i) Under $\delta^\ell = (0,0)$, $h_i^m = h_i$ and the rule selects θ. Consistency requires the policy at θ to be $(0,0)$, which by Lemma 3 holds iff $m_P(\theta) \le \varphi y^\ell$ and $\Gamma(\bar{\ell}(\theta); c_2^c) \ge u((1-\varphi)y^\ell)$.

(ii) Under $\delta^\ell = (1,0)$, $h_P^m = 1-(1-\pi)(1-h_P)$ and $h_O^m = h_O$; the rule plus (6) give $h_P\big[d_P+(1-\pi)d_O\big] = \mathcal{R}-\pi d_O$, i.e. (23), the unique candidate. Writing $h_P^d = (\mathcal{R}-\pi d_O)/(D-\pi d_O)$ and noting $x \mapsto (\mathcal{R}-x)/(D-x)$ is strictly decreasing on $[0,D)$ (derivative $(\mathcal{R}-D)/(D-x)^2 < 0$) with $0 < \pi d_O < \mathcal{R}$, we get $0 < h_P^d < \theta$; then (6) forces $h_O^d > \theta$, and $h_O^d = 1-(1-\pi)(1-h_P^d) < 1$. Both haircuts lie in $(0,1)$, so the candidate satisfies (6). Consistency requires the policy at h_O^d to be $(1,0)$: by Lemma 3, $m_P^d > \varphi y^\ell$ and $S(h_O^d) \ge 0$. Measured equality $h_P^m = h_O^m = h_O^d$ holds by construction of the rule, and the private exit spread is that of (20) at $\delta_P^\ell = 1$.

(iii) Let $u = \log$, $R^* = 1/\beta = 1.04$, $R_E = 1.02$, $\pi = 0.5$, $\lambda = 3$, $(d_P, d_O) = (0.35, 0.25)$, $\bar{M} = 0.35$, $(y^h, y^\ell) = (1.3, 0.6)$, and $\bar{\ell}_0 = 0.65$ (Table S.3). For every $\varphi \in (\bar{M}/y^h, 1)$ these values satisfy $\pi d_O < \mathcal{R}$ and, with strict inequalities, the restrictions that Section 2 and Assumption 6 place on them; no other parameter but φ enters (i)–(ii). With $\varphi = 0.40$ the four inequalities hold strictly ($0.204 < 0.240$; $-0.875 > -1.022$; $0.258 > 0.240$; $S(h_O^d) = 0.023 > 0$); every object is continuous in the parameter vector on the relevant domain, so the conditions hold on an open ball. On the coexistence set, exactly these two equilibria exist: a $(1,1)$ conjecture makes the metric distortion common to both classes, so the rule again selects θ (Lemma 4(i), compression identity), where the policy is $(0,0)$—consistency of the repayment assessment—not $(1,1)$; a $(0,1)$ conjecture induces (if feasible) an allocation with $m_P > 0$, at which the policy is never $(0,1)$ by Lemma 3.

(iv) $q_P^d = \beta^*(1-Fh_O^d) < \beta^*(1-F\theta) = q_P^{rep}$ by Lemma 4(ii) and $h_O^d > \theta$; $m_O^d = R^*(1-h_O^d)d_O < R^*(1-\theta)d_O = m_O(\theta)$.

(v) At the parameters of (iii) with $\varphi = 0.29 > \bar{M}/y^h = 0.269$: the $(0,0)$ conjecture selects θ, where $m_P(\theta) = 0.204 > 0.174 = \varphi y^\ell$, so the policy is not $(0,0)$; the $(1,0)$ conjecture selects (h_P^d, h_O^d), where $S(h_O^d) = -0.003 < 0$, so the policy is $(1,1)$, not $(1,0)$; the $(1,1)$ conjecture selects θ, where $S(\theta) = 0.029 > 0$, so the policy is $(1,0)$, not $(1,1)$; the $(0,1)$ conjecture is

vacuous by Lemma 3. All failures are strict; continuity delivers an open set. In general, by the argument of (iii) the only further candidate is θ with $\delta^\ell = (1,1)$, an assessment equilibrium iff the policy at θ is $(1,1)$. Hence none exists iff the policy at θ is $(1,0)$ and the policy at (h_P^d, h_O^d) is not, that is, since $m_P^d > m_P(\theta)$, iff $m_P(\theta) > \varphi y^\ell$ and $S(\theta) \geq 0 > S(h_O^d)$. □

References

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

Amegashie, J. A. (2023): "On the Preferred Creditor Status of Multilateral Development Banks," CESifo Working Paper 10521.

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

Arellano, C., and L. Barreto (2024): "Official Sovereign Debt," Working Paper, University of Minnesota and Federal Reserve Bank of Minneapolis, December.

Aumann, R. J., and M. Maschler (1985): "Game Theoretic Analysis of a Bankruptcy Problem from the Talmud," *Journal of Economic Theory*, 36(2), 195–213.

Ayres, J., G. Navarro, J. P. Nicolini, and P. Teles (2018): "Sovereign Default: The Role of Expectations," *Journal of Economic Theory*, 175, 803–812.

Bolton, P., M. Gulati, and U. Panizza (2023): "Sovereign Debt Puzzles," *Annual Review of Financial Economics*, 15, 239–263.

Bolton, P., and O. Jeanne (2009): "Structuring and Restructuring Sovereign Debt: The Role of Seniority," *Review of Economic Studies*, 76(3), 879–902.

Buchheit, L., and M. Gulati (2023): "Enforcing Comparable Treatment in Sovereign Debt Workouts," *Capital Markets Law Journal*, 18(1), 71–77.

Calvo, G. A. (1988): "Servicing the Public Debt: The Role of Expectations," *American Economic Review*, 78(4), 647–661.

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

Chun, Y. (1988): "The Proportional Solution for Rights Problems," *Mathematical Social Sciences*, 15(3), 231–246.

Cole, H. L., and T. J. Kehoe (2000): "Self-Fulfilling Debt Crises," *Review of Economic Studies*, 67(1), 91–116.

Cordella, T., and A. Powell (2021): "Preferred and Non-Preferred Creditors," *Journal of International Economics*, 132, 103491.

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

de Frutos, M. A. (1999): "Coalitional Manipulations in a Bankruptcy Problem," *Review of Economic Design*, 4(3), 255–272.

Dellas, H., and D. Niepelt (2016): "Sovereign Debt with Heterogeneous Creditors," *Journal of International Economics*, 99(S1), S16–S26.

Demirgüç-Kunt, A., and E. Fernández-Arias (1992): "Burden-Sharing among Official and Private Creditors," Policy Research Working Paper 943, World Bank, Washington, DC.

Diwan, I., B. Harnoys-Vannier, and M. Kessler (2023): "IDA in the Debt Crisis: Exploring Feasible Deals through Comparability of Treatments and New Loans," Finance for Development Lab, Policy Note 6, May.

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

Erce, A., and E. Mallucci (2018): "Selective Sovereign Defaults," International Finance Discussion Papers 1239, Board of Governors of the Federal Reserve System.

Fang, C., J. Schumacher, and C. Trebesch (2021): "Restructuring Sovereign Bonds: Holdouts, Haircuts and the Effectiveness of CACs," *IMF Economic Review*, 69(1), 155–196.

Farah-Yacoub, J., C. Graf von Luckner, and C. Reinhart (2024): "The Social Costs of Sovereign Default," NBER Working Paper 32600.

G20 (2020): "Common Framework for Debt Treatments beyond the DSSI," Annex to the Statement of the Extraordinary G20 Finance Ministers and Central Bank Governors' Meeting, November 13.

G20 (2022): *Boosting MDBs' Investing Capacity: An Independent Review of Multilateral Development Banks' Capital Adequacy Frameworks*.

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

Georgieva, K., and C. Pazarbasioglu (2021): "The G20 Common Framework for Debt Treatments Must Be Stepped Up," IMF Blog, December 2.

Gill, I. (2022): "It's Time to End the Slow-Motion Tragedy in Debt Restructurings," Brookings Institution, Commentary, February 25.

Graf von Luckner, C., J. Meyer, C. Reinhart, and C. Trebesch (2025): "Sovereign Haircuts: 200 Years of Creditor Losses," *IMF Economic Review*, 73(1), 150–195.

Grossman, H. I., and J. B. Van Huyck (1988): "Sovereign Debt as a Contingent Claim: Excusable Default, Repudiation, and Reputation," *American Economic Review*, 78(5), 1088–1097.

Guzman, M., and J. E. Stiglitz (2020): "Towards a Dynamic Disequilibrium Theory with Randomness," *Oxford Review of Economic Policy*, 36(3), 621–674.

Holmström, B., and J. Tirole (1998): "Private and Public Supply of Liquidity," *Journal of Political Economy*, 106(1), 1–40.

Ju, B.-G., E. Miyagawa, and T. Sakai (2007): "Non-Manipulable Division Rules in Claim Problems and Generalizations," *Journal of Economic Theory*, 132(1), 1–26.

Krugman, P. (1988): "Market-Based Debt-Reduction Schemes," NBER Working Paper 2587.

Moulin, H. (2000): "Priority Rules and Other Asymmetric Rationing Methods," *Econometrica*, 68(3), 643–684.

O'Neill, B. (1982): "A Problem of Rights Arbitration from the Talmud," *Mathematical Social Sciences*, 2(4), 345–371.

Paris Club (n.d.): "Methodology: The Comparability of Treatment," https://clubdeparis.org/en/sites/clubdeparis/accueil/club-de-paris/methodologie.html (accessed October 2, 2026).

Pigou, A. C. (1928): *A Study in Public Finance*. London: Macmillan.

Pitchford, R., and M. L. J. Wright (2012): "Holdouts in Sovereign Debt Restructuring: A Theory of Negotiation in a Weak Contractual Environment," *Review of Economic Studies*, 79(2), 812–837.

Powell, A. (2026): "Comparability of Treatment: Where Do We Stand?", CGD Working Paper 750. Washington, DC: Center for Global Development.

Reichert-Facilides, D. (2023): "Enforcing Comparability of Treatment – Why, What, How, Where? And Which Remedies?," available at SSRN: https://ssrn.com/abstract=4289082

Rivetti, D. (2022): "Achieving Comparability of Treatment under the G20's Common Framework," Equitable Growth, Finance and Institutions Note, World Bank, Washington, DC.

Schlegl, M., C. Trebesch, and M. L. J. Wright (2019): "The Seniority Structure of Sovereign Debt," NBER Working Paper 25793.

Schumacher, J., C. Trebesch, and H. Enderlein (2021): "Sovereign Defaults in Court," *Journal of International Economics*, 131, 103388.

Spiegel, M. M. (1996): "'Burden Sharing' in Sovereign Debt Reduction," *Journal of Development Economics*, 50(2), 337–351.

Sturzenegger, F., and J. Zettelmeyer (2008): "Haircuts: Estimating Investor Losses in Sovereign Debt Restructurings, 1998–2005," *Journal of International Money and Finance*, 27(5), 780–805.

Wicht, A. (2025): "Seniority and Sovereign Default: The Role of Official Multilateral Lenders," *Journal of International Economics*, 155, 104098.

Zucker-Marques, M., K. P. Gallagher, and U. Volz (2024): *Defaulting on Development and Climate: Debt Sustainability and the Race for the 2030 Agenda and Paris Agreement*. Boston University Global Development Policy Center, April.