SOVEREIGN BOND RESTRUCTURING: COMMITMENT VS. FLEXIBILITY*

Jason Roderick Donaldson† Lukas Kremens‡ Giorgia Piacentino§

April 11, 2023

Abstract

Policies facilitating sovereign bond restructuring are said to help countries emerge from distress and even reduce negative spillovers across classes of bonds. We use natural experiments to assess such policies. Our findings suggest they can backfire: Committing *not* to restructure one class of bonds is valuable for not only that class, but for others too. These positive spillovers contrast with received theory, so we develop a model to rationalize them. The model points to a system of cross-bond equations that, we show, can be exploited to estimate unobservable elasticities in terms of a few sufficient statistics. We use this method to quantify how much commitment decreases default costs.

*For valuable comments, thanks to Viral Acharya, Tobias Berg, Patrick Bolton, Charlie Calomiris, Christopher Clayton, Andras Danis, Olivier Darmouni, Wenxin Du, Brent Glover, Todd Gormley, Ben Hébert, Ritt Keerati, Doron Levit, Yiming Ma, Enrico Mallucci, Stefan Nagel, Chris Parsons, Carolin Pflueger, Bernardo Ricca, Jesse Schreger, Janis Skrastins, Amir Sufi, Paul Tetlock, Fabrice Tourre, Pierre Yared, and audiences at the 2021 FIRS Conference, the 2022 AFA Meeting, the 11th Advances in Macro Finance Tepper-LAEF Conference, the 2022 SFS Cavalcade, the IU Junior Finance Conference, the 2023 ASSA Meeting LSE, MIT Sloan, BI Oslo, and the University of Washington.

†USC, WashU, and CEPR.

‡University of Washington.

§USC, Columbia, CEPR, and NBER.

1 Introduction

As of February 2023, more than half of all low-income countries are in debt distress.[1] Many have already defaulted in the last few years, including Argentina, Belize, Ecuador, Lebanon, Mozambique, Russia, Ukraine, Sri Lanka, Suriname, and Zambia.[2] Others are pleading for debt relief.[3] Some defaults are hard defaults in which bondholders get little or nothing. But most defaults are soft/partial defaults in which they take a haircut but still get something in a debt restructuring.

The policy world is demanding action to address the crisis. The World Bank and the IMF are advocating making restructuring easier, focusing on "resolving the impasse in debt restructurings" and "working hard...to avoid a disorderly debt restructuring."[4] The G20 is also on board, having agreed to the Common Framework for Debt Treatments, which facilitates sovereign debt restructuring by, inter alia, helping dispersed bondholders coordinate.[5]

But support for such policies is not universal. Detractors argue that they encourage opportunistic restructuring, decreasing the value of bonds and even choking off credit ex ante.[6] Bonds, they say, are made hard to restructure by design, via covenants and governing law, as a way to commit not to do it.[7] They conclude that the commitment not to restructure must be valuable.

But the case for policy intervention turns on whether the value of this commitment accrues to bonds collectively or, as Bolton and Jeanne (2007, 2009) argue, to some

[1]See, e.g., "Why a Plan to Help Low-income Nations Rework their Debt Isn't Working" (*Bloomberg*, February 14, 2024).

[2]See, e.g., "Argentina clinches near-unanimous backing for debt restructuring" (*Financial Times*, August 31, 2020), "How to restructure sovereign debt" (*Financial Times*, September 16, 2022).

[3]"Emerging economies plead for more ambitious debt relief programmes" (*Financial Times*, October 11, 2020).

[4]See blogs.worldbank.org/voices/action-needed-debt-crisis and imf.org/en/News/Articles/-current-sovereign-debt-challenges-and-priorities-in-the-period-ahead.

[5]See imf.org/en/About/FAQ/sovereign-debt; see, e.g., Krueger (2002) and White (2002) for earlier proposals.

[6]See, e.g., Dooley (2000) and Shleifer (2003).

[7]See, e.g., Dooley and Verma (2003).

individually at the expense of others. They say that committing not to restructure a bond, e.g., via governing law and covenants, could be like installing a burglar alarm: It could provide protection for an individual only by diverting costs onto others—other bonds are selectively restructured/other homes targeted by burglars—making it self-defeating collectively.

The policy debate needs empirical answers to the following questions: (i) Is the commitment not to restructure one class of bonds valuable for that class? (ii) Does that commitment have negative spillovers on other classes, like a burglar alarm? But the answers remain elusive. The evidence on the first is mixed, perhaps due to inherent identification challenges—bonds made hard to restructure (e.g., via foreign law or cross-default clauses) could differ in any number of ways from those made easy (e.g., currency, maturity, and other covenants).[8] There is no evidence on the second question as yet (to our knowledge).

In this paper, we bring evidence from natural experiments to bear on these questions. Our results suggest that the commitment not to restructure one class of bonds is valuable not only (i) for that class but (ii) for other classes as well: Commitment has positive cross-bond spillovers—it seems more like a successful neighborhood crime watch than a self-defeating burglar alarm. This result contrasts with the predictions of existing models, which stress the burglar-alarm view that the commitment not to restructure one class of bonds leads to selective default on other bonds.

Lacking a model to explain our findings, we develop one. Rather than selective default, we focus on cross-default, assuming a sovereign restructures all its bonds if any. The assumption is consistent with recent practice—e.g., cross-default and pari passu clauses have made selective default harder—and with standard strategic default incentives (cf. Appendix C). It helps generate the crime-watch spillovers we document, since committing not to restructure one class of bonds can benefit all of them by

[8]Chamon, Schumacher, and Trebesch (2018) and Clare and Schmidlin (2014) exploit heterogeneity in bonds' governing law; Becker, Richards, and Thaicharoen (2003), Bradley and Gulati (2014), Carletti et al. (2020), Eichengreen and Mody (2004), and Picarelli, Erce, and Jiang (2019) exploit heterogeneity in their covenants.

making it unattractive to start a restructuring at all. It also points to heretofore overlooked cross-bond relationships that, we show, can be used to quantify the impact of experiments and to estimate otherwise unobservable elasticities in terms of a few sufficient statistics. Applying the methodology delivers sensible numbers, which could be useful for quantitative (i.e. less highly stylized) models of sovereign debt. It can also be useful for counterfactual policy analysis, as we illustrate by applying it to estimate how much an increase in commitment decreases default costs.

Our starting point is the collective action problem inherent in bond restructuring. As we summarize in Section 2.1, dispersed bondholders could reject a restructuring that would make them better off collectively (the "hold-out problem") or accept one that would make them worse off (the "hold-in problem"). Sovereigns mitigate/exploit the problems using threats and bribes of questionable legality to pit bondholders against one another.

Our baseline experiments are court rulings determining whether or not certain bribes and threats are legal. The first experiment, which is new to the economics literature (to our knowledge), is a case heard in the English High Court against a firm that had used an effective bribe called "consent payments" to cut every coupon that came due for two years. A bondholder sued in court, but the judge ruled against him, deeming the bribes legal. Thus the ruling made all outstanding bonds under English law easier to restructure and, being the Court's first on consent payments, thus had "worldwide ramifications" (White & Case (2013)). The second, which is taken from Hébert and Schreger (2017), is a sequence of rulings culminating in a case heard in the Southern District of New York about Argentina's threat to punish hold-out bondholders, paying them nothing while servicing bonds that agreed to a haircut. The judge deemed the threat illegal. Thus the ruling made all outstanding bonds issued under NY law harder to restructure and precipitated another Argentine default. (We drop Argentine bonds everywhere as, contrary to Hébert and Schreger (2017), we use the experiment as a shock to NY law, not to Argentina.)

The experiments provide a laboratory to address our questions on sovereign debt restructuring, as we can compare the same bonds on the trading days before and after the rulings and thereby avoid relying on comparisons of different bonds. Both court decisions mattered for hundreds of sovereign bonds, as many countries issue under English or NY law. Other countries do not. Their bonds provide us with a control group for difference-in-differences (DiD) analyses.

We use a different treatment group to address each of our two motivating questions, on how making on class of bonds harder to restructure affects (i) that class—the direct effect, hereinafter—and (ii) other classes—the indirect/spillover effect. For the direct effect, we consider "directly treated" bonds, defined as those issued under the affected law, be it English or NY. For the spillover effect, we consider "indirectly treated" bonds, defined as the other bonds issued by countries with some outstanding directly treated ones.

We then compare the differences in credit spreads of each treatment to those of the control group before and after the rulings. For the English-law experiment, we find that the direct effect of an increase in commitment is about 5 bps and the indirect effect is about the same size. For the NY-law experiment, which includes multiple events, we aggregate them by weighting them by the (signed) change in the risk-neutral default probability P^{AR}, as Hébert and Schreger do. We find that for every 1 percentage point increase in P^{AR} the direct effect is about 0.35 bps and the indirect effect is about 0.15. While these effects are large enough to be statistically and economically significant, they reflect only marginal changes in the law, not an overhaul of market practice. That proves useful below, where, using a linear approximation around the equilibrium, our model helps us interpret their magnitudes. Besides, such a marginal effect is what we want to understand policy interventions given the prevailing equilibrium, and thus to address our motivating questions: An increase in commitment not to restructure one class of debt helps both (i) that class and (ii) other classes as well—commitment is not a burglar alarm.

We interpret our results as a causal effect of the ruling on treated bonds. To do so, we rely on two assumptions. The first is that the experiments do not affect control bonds. Our research design ensures that holds by focusing on experiments specific to restructuring English- or NY-law bonds and choosing control bonds issued by countries without any treated bonds. The second assumption is that, conditional on controls variables, treated and control bond outcomes would have moved the same way absent the experiments. Our research design ensures that holds by including region- and currency-time fixed effects as control variables, focusing on narrow event windows (just two days), and matching treated and control bonds on risk and other characteristics. Thus, to threaten our identification, a confounding event must be nearly simultaneous with our experiments and must affect treatment and matched control bonds differently within the same region and currency. That sounds unlikely prima facie and we do a data-driven placebo-type test to confirm as much; our results almost never arise spuriously in our sample (see Section 3.4).

Our interpretation also relies on the assumption of market efficiency, namely that short-term price changes reflect new information, not merely short-term inefficiencies due to, e.g., overreaction or fire sales. The assumption is relatively weak for the bond markets we study, which are frequented by professional investors and are particularly informative empirically (e.g., Philippon (2009)). Moreover, those investors have been attuned to changes in governing law, like those in our experiments, since at least the Greek restructuring, when newspapers said things like, "We also really, really want to know if the new bonds under this new offer will be governed under English law or will remain under Greek law."[9] Nonetheless, where possible, we show that our results do not rely on it: We repeat our analysis with a longer event window (four trading days). Our results are robust.

Our baseline results encompass both bribes and threats under both of English and NY law, effectively the only foreign laws countries borrow under. Yet we still provide

[9]"Who Wants to Be a Greek Bond Holdout" (*Financial Times*, October 26, 2011).

further assurance of external validity, repeating our analysis using two similar (though arguably more imperfect) experiments, on yet another type of restructuring device: exit consents (Section 5). We find analogous results.

Next, we develop a simple model of a sovereign restructuring, which serves to investigate the mechanism behind our results, to interpret their magnitudes, and to understand policy counterfactuals. The model is based on two key ingredients. The first is the willingness-to-pay problem, which is common in the literature: The only thing that deters restructuring, a form of partial default, is costs of default such as exclusion from financial markets (Eaton and Gersovitz (1981)) or trade sanctions (Bulow and Rogoff (1989)). We assume the costs are independent of the scale of default, so, as touched on above, the sovereign defaults on all bonds if any (except in extensions in Appendix C, where we show that our findings are robust to including selective default). The second ingredient is heterogeneity in the ease of restructuring, something that appears in only a few papers:[10] Bonds differ in the haircuts suffered in default, an assumption that, as shown formally in Section 2.1, captures differences in what bribes and threats are legal. Specifically, there are hard-to-restructure "rigid bonds" associated with a low haircut h_r and easier-to-restructure "flexible bonds" associated with a higher (expected) haircut $h_f > h_r$. For example, flexible bonds could correspond to local-law bonds, which a sovereign can restructure unilaterally by changing the law as Greece did in 2012 (Zettelmeyer, Trebesch, and Gulati (2014)).

The model has close counterparts in our baseline experiment: Rigid bonds correspond to foreign-law, flexible to local, and, per our analysis in Section 2.1, our experiments to changes in the haircut on rigid bonds h_r.

Our main results are comparative statics with respect to h_r that mirror the empirical effects of our experiments. A decrease in h_r has two effects on yields. (i) There is an effect that works via recovery values: It increases the payoff to rigid bonds in default, decreasing their spreads. (ii) There is an effect that works via default

[10]See Bolton and Jeanne (2007, 2009) and Carletti et al. (2020).

probabilities: It discourages strategic default, decreasing the spread on *both* types of bonds.

The model captures our main empirical findings on what happens when commitment increases: Per our first question, rigid spreads decrease (by both (i) and (ii)). Per our second, flexible spreads in countries with outstanding rigid debt decrease too (by (ii)). It also explains how the indirect effect can be just as large as the direct effect: Bonds with the most to lose in default have the most to gain from deterring it (by $h_f > h_r$ and (ii)).

The model can also help us to interpret our experiments, which do not immediately lend themselves to quantification—some rulings made restructuring easier, others harder, but by how much did they affect the haircut h_r? We use the model to back out this number. We derive a formula for the implied change in h_r in terms of our estimated yield changes and a few other sufficient statistics, notably estimates of risk-neutral default probabilities (from CDS) and expected haircuts. We estimate that average haircuts increased by about 37 bps from High Court ruling and, coincidentally, decreased about the same amount from each event in the Argentina saga on average. These estimates accord with our intuition that the rulings are meaningful, reflecting landmark legal decisions, but their immediate economic effects are not large, pertaining to specific restructuring devices.

We show how to use the same sufficient statistics to translate our baseline estimates into an elasticity that could be useful for policy counterfactuals: We estimate that a 1% increase in haircuts increases the (risk-neutral) default probability by 14.16% for English law and 4.74% for NY. A back-of-the-envelope application of this number, in combination with Hébert and Schreger's (2017) estimates of the costs Argentine default, suggests that a one percentage point decrease in haircuts on Argentine debt could save 1.7 billion USD in deadweight costs.

Finally, a proviso: Although our results suggest the commitment *not* to restructure debt is good for sovereign creditors, and could even decrease deadweight costs of default, they say little about whether it is good for sovereign borrowers. It could be good, e.g., facilitating access to credit, or bad, e.g., limiting the self-insurance role of default (Zame (1993)). One way to get at that question in future research could be to look at what sovereigns do following a shock to creditor commitment, something our experiments are unfortunately not well suited to do: Their strength in identifying price responses—a narrow event window—is a weakness in identifying behavioral responses, which do not happen fast enough for us to pick up.

2 Institutional Background and Data

Here we lay out the ingredients behind our empirical analysis, starting with a primer on the problems of debt restructurings. We describe the threats and bribes used to impose haircuts on creditors and the legal remedies creditors use to resist them (Section 2.1). We also derive a relationship between equilibrium haircuts and the strength of such remedies that is useful to map our theory to our experiments, which are summarized in the next section (Section 2.2). Finally, we describe the data we use (Section 2.3).

2.1 Threats and Bribes in Bond Restructuring

Sovereign default is not uncommon. As Reinhart and Rogoff (2009) catalog, "[m]ost countries in all regions have gone through a prolonged phase as serial defaulters" (p. 49). Most defaults are negotiated partial defaults in which creditors take a haircut but still get something (Arellano et al. (2019); Cruces and Trebesch (2013)). For bonds, such partial defaults typically start with a restructuring offer from the sovereign.

Given bondholders are often dispersed, bond restructurings are plagued by collective action problems. On the one hand, there is what is known as the "hold-out problem": An individual bondholder could reject an offer that could make bondholders collectively better off. On the other hand, there is what is sometimes called the "hold-in problem": An individual bondholder could accept an offer that makes them collectively worse off.[11] (See Buchheit and Gulati (2000), Donaldson et al. (2022), Gertner and Scharfstein (1991), Kahan and Tuckman (1993), and Roe (1987).)

To see how the hold-out problem works, consider a sovereign that has so much debt that it is tempted to default outright. Collectively, bondholders might be better off taking a haircut to reduce the sovereign's debt burden and thereby increase the likelihood it repays them at all. But an individual bondholder still might not accept the haircut. By holding out, it can free ride on others' accepting, benefiting from the overall debt reduction without taking a haircut itself. As all bondholders have an incentive to act this way, the whole restructuring can fall apart even though it would have made everyone better off.

Sovereigns tackle the hold-out problem with threats and bribes. One threat is to pay nothing to "hold-out" bondholders that do not accept a haircut, as Argentina did in 2001 (see Section 2.2). But the legality of such threats is dubious and the hold-outs might litigate and end up getting paid in full. In equations, if a bondholder with face value D accepts, its face value is decreased by the haircut h. If it holds out, it risks getting paid nothing, but could litigate and get paid in full with probability π. Thus it accepts if

$$(1-h)D \geq \pi D, \tag{1}$$

or $h \leq 1-\pi$. I.e. haircuts are higher when hold-outs are less likely to win in court, a relationship we employ off the shelf in the model below (see Section 4).

Such threats can be especially effective if invoked using a collective action clause (CAC), a contractual term, now commonplace in sovereign bonds, that addresses free rider problems by allowing all bonds in a series to be modified by a specified majority, so minority hold-outs effectively cannot litigate (π is close to zero in equation (1)).[12]

[11] These are not the only ways that collective action problems can lead to inefficiencies in sovereign-debt restructuring, something Pitchford and Wright (2012) demonstrate in a model of bargaining delays.

[12] Most countries enshrine some bondholder rights in legal statutes for corporate bonds (e.g., the Trust Indenture Act in the US). Not so for sovereign bonds, in which they are specified, by contract, in the indenture; see example prospectuses: under NY, English, and Italian laws, respectively: Suriname, Poland, and Italy.

But, while helpful (Fang, Schumacher, and Trebesch (2021)), CACs are not a panacea. Free rider problems persist when CACs cannot be aggregated across series (Gelpern, Heller, and Setser (2016)), and they "might be an invitation to litigation" otherwise.[13]

Bribes are the alternative to threats. One contractual device, so-called "consent payments," constitutes a payment to bondholders when they agree ("consent") to changes in terms (of all bonds in a series). However, legality is again uncertain and hold-outs could likewise litigate for full payment.[14] In equations, if a bondholder with face value D accepts, it gets the (proportional) bribe b but decreases its face value by a haircut h. If it holds out, it forgoes the bribe but can litigate and be paid in full with probability π. Thus it accepts if

$$bD + (1-h)D \geq \pi D + (1-\pi)(1-h)D, \tag{2}$$

or $h \leq b/\pi$. The take-way of the analysis of bribes is the same as of threats: Haircuts are higher when hold-outs are less likely to win in court, so we can apply the relationship in the model later regardless of the reason for litigation.

Overall, such threats and bribes are a double-edged sword. By making it easier to restructure, they mitigate the hold-out problem. But they can exacerbate its counterpart, the hold-in problem (especially when combined with other anti-hold-out provisions like CACs), because they can force bondholders into a prisoner's dilemma in which they tender even if restructuring is against their collective interest. Either way, what haircuts are feasible depends on how likely a legal jurisdiction is to protect bondholders—on how big π is.

[13] See creditslips.org/creditslips/2020/06/the-argentine-re-designation-drama-notes-from-two-frustrated-readers.html

[14] Panizza, Sturzenegger, and Zettelmeyer (2009) enumerate cases of such litigation; hold-outs received full payment or close to it in some cases but little or nothing in others. See Schumacher, Trebesch, and Enderlein (2021) for an overview of sovereign-debt litigation in recent decades.

2.2 Baseline Rulings: Azevedo v Imcopa and Argentina v NML Capital

At the center of our analysis are two legal disputes over bond restructurings, one under English law, about bribes, and another under NY law, about threats.

2.2.1 Azevedo v Imcopa: Bribes under English Law

Our first baseline experiment is the May 2012 ruling by the English High Court that opposed a challenge to the legality of an effective bribe called consent payments, which has proved useful in sovereign debt restructuring lately.[15] The ruling thus made it easier to restructure bonds in exchange offers.[16] It applied to all English-law bonds, including foreign sovereign bonds, but was otherwise unrelated to the sovereign bond market.

The case, brought by two individual bondholders, Sergio Barreiros Azevedo and Vera Cintia Alvarez, against Imcopa Group, a Brazilian company in the soybeans business, constitutes a "landmark decision," according to the *Financial Times*.[17] The newspaper also stresses that consent payments, while not uncommon, had never been considered before by English courts, making the case "hugely important." However, the case seems not to have been studied in the finance and economics literature. Hence, we offer a précis now.

Things started in 2008, when soybean prices plummeted from more than $16 a bushel in July to less than $8 in December.[18] Imcopa embarked on a plan to reduce its interest payments via a restructuring of its bonds, which were issued under English law. Over the next two years, it restructured its bonds every time an interest payment

[15]E.g., Suriname deployed it repeatedly in 2020 and 2021 to defer interest payments on two bond series; see reporting by *Fitch* (fitchratings.com/research/sovereigns/fitch-upgrades-suriname-lt-fc-idr-to-cc-16-07-2020) and *LatinFinance* (latinfinance.com/daily-briefs/2020/11/17/suriname-seeks-breathing-room-for-debt-payments). Primary sources are available at bourse.lu/security/US86886PAB85/299502.

[16]See Azevedo v Imcopa Importacao, Exportacao e Industria de Oleos Ltda [2012] EWHC 1849 (Comm) (30 May 2012)).

[17]See "The consent of the (bondholder) governed" (April 22, 2013).

[18]See, e.g., macrotrends.net.

came due. Each time, it offered bondholders consent payments if they agreed to put off the interest payment. As predicted by the prisoner's dilemma described above (Section 2.1), nearly all bondholders accepted.

But the last time, Azevedo and Alvarez held out. Equation (2) suggests they found the bribe b insufficient compensation for the chance π of getting paid in full following litigation. Indeed, they sued Imcopa in England, claiming, inter alia, that the consent payments constituted an illegal bribe.

The High Court ruled in favor of the defendant on May 30, 2012, on the grounds that the consent payments were offered to all bondholders, and therefore were not an illegal bribe. The decision was not a foregone conclusion (see, e.g., Jones Day (2012)). Azevedo and Alvarez appealed. But the original ruling was upheld on April 22, 2013.

2.2.2 Argentina v NML Capital: Threats under NY Law

Our second baseline experiment is a series of court decisions throughout the Argentine restructuring saga, culminating in a 2014 Manhattan federal district court ruling that made it harder to punish hold-out bondholders.

Things started in 2001, when Argentina missed payments on NY-law bonds with face value of about \$82bn. It proceeded with exchange offers, offering bondholders a choice: accept a 70% haircut or hold out and get nothing. Per equation (1), 90% accepted. It made good on its threat toward the hold-outs, servicing only its restructured debt.

Litigation ensued. Most notably, a hedge fund, NML Capital, sued Argentina in New York, demanding full repayment. It argued that defaulting on hold-outs while repaying other bonds violated a covenant in the bond indenture: the pari passu clause. The presiding judge, Thomas P. Griesa, ruled with the plaintiff. The ruling prevented Argentina (via the trustee that distributed payments on its behalf) from servicing any debt until the hold-out bonds were paid in full.

To avoid another hard default, Argentina appealed, counter-sued, and even tried to service its debt outside of US jurisdiction. Ultimately, the US judiciary affirmed the decision and blocked the attempts to circumvent its implementation. Argentina defaulted on its bonds serviced via US-based payment systems in July, 2014, and only resumed payments in April, 2016, after settling with NML Capital and other hold-outs.

2.3 Data Sources and Sample Construction

Bond data come from Dealogic DCM and Bloomberg's records of sovereign bond issues from 1980 onwards. We collect data on bonds outstanding between May 2012 and October 2014, a sample period that includes both baseline experiments, as well as two others used for external validity (Section 5). The sample starts after the Greek restructuring in early 2012, an event which made it salient that sovereign debt haircuts can depend on governing law.

We focus on bonds issued by countries with significant English- and NY-law borrowing and those issued by relevant control countries, e.g., countries in Europe and Latin America. We include only bonds with some data available on governing law, our key variable of interest. Dealogic provides it for most issues; Bloomberg for fewer.

We must make a few adjustments when we merge the governing law data from the two sources, as sometimes multiple laws are listed. Occasionally, (i) Bloomberg lists two laws. In this case, if Dealogic provides a single governing law for the bond, we use that law; if not, we use the one that coincides with the currency denomination of the bond. Other times, (ii) Bloomberg and Dealogic list different laws. In this case, if one coincides with the currency denomination, we use that; if not, we use the one provided by Dealogic. (See also Kropp, Gulati, and Weidemaier (2018) for a discussion of data sources and quality for governing law.)

We exclude bonds with missing ISIN and those issued by publicly-owned companies, state/local governments, or Argentina, which, as of 2012, was still in litigation following its 2001 default and subsequent restructuring. The sample of bonds with available governing law and credit spread information that are outstanding during the sample period contains 2,024 bonds issued by 76 countries, denominated in 23 currencies, and 907,415 bond-day observations.

We collect daily yields and swap rates from Bloomberg, which we use to construct credit spreads and thereby difference out risk-free rate changes.

For our English-law experiment, our event window is two days: from May 29 to May 31, 2012. Since bonds are traded OTC in different time zones, this is the narrowest window that ensures that bonds traded in European, Middle Eastern, and Asian markets reflect the new information. (We also consider a wider event window as a robustness check, to show our results are not an artifact of short-term market overreaction; see 3.4.) On May 31, the sample contains 1,281 bonds, issued by 67 countries in 22 currencies. We winsorize credit spread changes at 1% and 99%.

For our NY-law experiment, we supplement these data with data on changes in the estimated risk-neutral probability of Argentine default on fifteen event-days throughout the saga—including not only NML's filing and the court's ruling, but also Argentina's appeals and counter-suits—taken from Hébert and Schreger (2017). On November 23, 2012, the sample contains 1,325 bonds, issued by 68 countries in 23 currencies. Figure 1 plots treated and control countries.

Table 1 shows the summary statistics, aggregated by country, for bonds outstanding on the day before (the start of) each experiment. The median number of bond issues a country has outstanding is 12 in May 2012 and 14 in November 2012; the median face value is around $1bn on both dates; the median maturity is about 8 years on both dates; the median credit spread is 2.87% in May and 2.11% in November. The average amount of debt issued under English law is about a third, under NY-law about a quarter, and under local about 39%. Table SA.6 in the Supplementary Appendix shows statistics by country on the date of the English-law experiment.

3 Empirical Framework and Results

Here we set up our empirical specifications (Section 3.1), we argue they have a causal interpretation (Section 3.2), we report our baseline empirical results (Section 3.3), and we argue for their robustness (Section 3.4).

3.1 Estimation Strategy

To evaluate the bond market response to a change in the ease of bond restructuring, we use a high-frequency difference-in-differences (DiD) approach. Our baseline model, which we label (R1), nests two DiD specifications, each addressing one of our motivating questions.

In each of the two DiD specifications nested in regression (R1), the first difference is an increase in commitment not to restructure over an event window, which we label ΔComm_t. The second difference is between our control bonds and one of two treatment groups. The first is "directly treated" bonds, namely those issued under English and NY law, for each respective experiment. The second is "indirectly treated" bonds, namely other bonds issued by countries with some outstanding directly treated bonds. Nesting these specifications in a single regression has the advantage of estimating the coefficients on control variables on the full sample (absent controls the nested specification is equivalent to two separate DiD specifications). Formally, we estimate:

$$\begin{aligned} s_{i,t} = \beta_1 \Delta\text{Comm}_t \times \mathbb{1}_i^{\text{Direct}} + \beta_2 \Delta\text{Comm}_t \times \mathbb{1}_{i,c}^{\text{Indirect}} \\ + \beta_3 \mathbb{1}_i^{\text{Direct}} + \beta_4 \mathbb{1}_{i,c}^{\text{Indirect}} + \alpha^\top \text{FE}_{i,t} + \varepsilon_{i,t}. \end{aligned} \tag{R1}$$

In our baseline specifications, the dependent variable s_{it} is the change in the credit spread. Relative to other measures of default risk (viz. yields), the credit spread has the advantage of isolating risk by differencing out the risk-free rate. We construct the credit spread by subtracting the maturity- and currency-matched swap rate from the yield, as detailed in Appendix A.2.[19] (These baseline specifications compare changes, the analog of "abnormal returns," in and out of the event window, as is standard in event studies; however, as a robustness check, we also compare levels immediately before and after the event (see Table 4), as is more common in DiDs.)

The key dependent variable, ΔComm_t, is defined as follows. For the English-law experiment, which is a single event that made restructuring easier (Section 2.2.1), the variable is just the indicator of the event with a minus sign: $\Delta\text{Comm}_t = -\mathbb{1}_t^{\text{Event}}$, where $\mathbb{1}_t^{\text{Event}} = 1$ if date t is the first trading day after the High Court ruling; for the NY-law experiment, which is a sequence of events, some making restructuring easier and others harder, we weight the events by their (signed) intensity by multiplying the indicator by the change in the estimated risk-neutral probability of default: $\Delta\text{Comm}_t = \Delta P_t^{\text{AR}} \mathbb{1}_t^{\text{Event}}$, where, analogously, $\mathbb{1}_t^{\text{Event}} = 1$ if date t is the end of one of the two-day event windows catalogued in Hébert and Schreger (2017).

The remaining dummy variables are as follows: $\mathbb{1}_i^{\text{Direct}} = 1$ if bond i is in the direct treatment group; $\mathbb{1}_{i,c}^{\text{Indirect}} = 1$ if bond i is in our indirect treatment group. The terms $\text{FE}_{i,t}$ denote vectors of fixed effects, including region- and currency-time for panel specifications, and pair-time for matching specifications (detailed in Table 2).[20]

We also control for bond risk, either by including granular credit rating FEs or by running the specification on matched bond pairs, comparing English- to NY-law bonds for the direct effect and local-law bonds of English-law issuers to local-law bonds of non-English-law issuers for the indirect effect. We match bonds by currency and maturity. We also match on credit rating, as tightly as we can given the sample size; for the English-law experiment that is within two notches (so, e.g., a bond rated AA− may be matched to one rated as high as AA+ or as low as A) and, for the NY-law experiment, that is within one notch for the direct effect and within the same rating (so it is equivalent to a notch-level fixed effect) for the indirect. The details of

[19] Swap rates avoid measurement problems arising from credit risk and convenience yields in seemingly risk-free sovereign bonds (Du and Schreger, 2016; Kremens, 2023).

[20] The regions are Europe, South America, Central America, Middle East/North Africa, Africa, Central Asia, East Asia. A list of sample countries and the regions we assign them to is in Appendix A.

the matching procedure and a complete list of matched bond pairs are in Appendix B and, respectively, in the Supplementary Appendix (Table SA.4 and Table SA.5).

The key coefficients of interest are β_1 and β_2. They measure the changes in s_{it} around the event dates for, respectively, directly treated and indirectly treated bonds relative to our control groups. They represent averages across heterogeneous bonds conditional on the market conditions—countries' debt levels, investors' risk aversion, etc.—on the event date; i.e. they are local average treatment effects (LATE).

3.2 Identifying Assumptions

Our empirical strategy is a difference-in-differences (DiD) over a narrow event window. We can interpret the estimates as causal results of the rulings if the following assumption is satisfied: The expected change in the dependent variable over the window, conditional on our controls (but not on the ruling), must be equal for treated and control bonds. The assumption could be violated if there were another event that both (i) happened at the same time as our events and (ii) affected our treatment and control bonds differently.

The assumption is inherently untestable (and tests that use observations prior to events to proxy for counterfactual behavior contemporaneous with them inappropriate in our setting; here the dependent variable, being derived from an asset price, is inherently forward looking, not backward). But we can mitigate concerns about it. We do so in five principle ways.

(a) We include region-time, currency-time, and risk (credit rating)-time fixed effects in our panel regressions. These absorb variation on any given date common to a given region, currency, or risk category.

(b) We include a specification run on only matched pairs of closely comparable bonds which we list entirely in the Supplementary Appendix (Table SA.4 and Table SA.5). These bonds are in the same currency and are closely matched on maturity and credit rating. To be a threat to identification, confounding shocks would have to hit matched bonds in systematically different ways.

(c) We choose a narrow event window of just two days. To be a threat to our identification, an event must both affect treatment and control bonds differently and also happen at nearly exactly the same time as the ruling. We have searched the news and not have not found any such relevant event released around the ruling.

(d) We control for movements in risk-free rates non-parametrically by using credit spreads as our dependent variable (subtracting currency- and maturity-matched swap rates from yields).

(e) We include multiple experiments; the English-law experiment uses different bonds, at a different date, from the New York-law one.

For the NY-law experiment, the events did not only make NY-law bonds harder to restructure, but also made Argentina more likely to default. Thus, although the conditions above suffice to say our findings are caused by the rulings, to interpret them as a result of bonds becoming harder to restructure requires an additional assumption, albeit an apparently weak one: Countries' exposure to Argentine default is uncorrelated with their propensity to borrow under New York law *within region/after matching*. (Recall that we drop all Argentine bonds from our sample (Section 2.3) and note we can also drop all South American bonds (Section 3.4.3).)

3.3 Do Sovereign Creditors Benefit from Committing Not to Restructure?

Table 2 reports the results from estimating equation (R1), our DiD regression of spread reactions on directly and indirectly treated bonds for both the English-law and NY-law experiments.

3.3.1 The Effect on Directly Treated Bonds

The direct treatment effect for the English-law experiment ranges from about 5–10 bps, with larger effects for the specifications without risk controls. The change in coefficient across specifications could reflect heterogeneous treatment effects: As treated countries have both higher risk and larger effects, the specification without risk controls (which compares high- and low-risk bonds) yields a larger coefficient. With risk controls, the coefficients are similar. But the standard errors, which are clustered at the country level, are not. The reason is that there are too few countries with the same bond ratings for us to maintain statistical power with the FE specification. Thus we favor the matching specification from now on.

The direct treatment effect for the NY-law experiment is about 0.29–0.39 bps for every 1 percentage point increase in the risk-neutral probability of Argentine default, with coefficients and standard errors relatively stable across specifications.

The results suggest increased commitment benefited bondholders. In the language of Section 2.1, that means bondholders care more about protection against the hold-in problem than the risk of the hold-out problem.

3.3.2 The Spillover Effect on the Indirectly Treated Bonds

The indirect treatment effect for the English-law experiment ranges from about 4–13 bps. The coefficients and standard errors demonstrate the same patterns as for the direct effect (and we offer the same explanation for them).

The indirect treatment effect for the NY-law experiment is about 0.13–0.15 bps for every 1 percentage point increase in the risk-neutral probability of Argentine default. Again both the coefficients and standard errors are similar across specifications.

This result suggests a positive spillover from increased commitment, in line with easy-to-restructure bonds failing to discipline a sovereign, harming holders not only of those bonds but of others as well. That channel also explains why the indirect effect is about the same size as the direct, a result verified in a within-country event study in Table 6, shown to be consistent with the disciplining role of rigid/foreign-law debt in Section 4.2.

3.4 Empirical Concerns and Robustness Tests

We now address residual concerns about our baseline tests with a number of robustness exercises.

3.4.1 Concern: Spurious Results (Omitted Variables and/or "Noise")

One concern is that our results are spurious, caused not by the experiments we study, but by something else, unknown to us, that happened in the same event window(s). As discussed above (Section 3.2), such omitted variables seem hard to imagine; they would have to coincide with our experiments and affect our treatment and control groups differently, net of our controls/FEs/matching. That seems especially far fetched for the NY-law experiment, since it comprises a sequence of fifteen events, but perhaps not impossible to imagine for the English-law experiment, since it is just one. Thus, for that experiment, we do two data-driven placebo exercises to provide further assurance that what we estimate is due to the ruling.

The first, a "time-series placebo," shows, by brute force, that events that could generate spurious results are rare in our sample. We assign the event dummy to each of the 644 days in our sample period, excluding dates $t-1$ to $t+1$ for the baseline event and the two other rulings described in Section 5. We re-estimate regression (R1) and construct an empirical distribution of the estimated t-statistics of interest for these placebo events. We then calculate "placebo p-values" for the results, each counting the sample days on which we observe differences in credit spread changes between treatment and control groups that produce larger effects (as measured by their t-statistics) than those obtained for the event date. The results are in Table 3.

These placebo p-values should be interpreted as (an estimate of) an upper bound on how likely our results are to occur by chance. The bound need not be tight as the results could reflect not only "noise," but also systematic events on other days in the sample, especially given the treatment and control bonds are not identical in some dimensions, such as risk. Indeed, absent risk controls, effects like ours, while uncommon, are not exceedingly rare: The placebo p-value is 10.7% for the direct effect, 14.4% for the indirect effect, and 4.8% for the two jointly. But using the matching specification tightens the bounds: The placebo p-value is 7.9% for the direct effect, 3.0% for the indirect effect, and 0.2% for the two jointly—a set of direct and indirect effects like ours happens on only one other day in the sample.

The second test, a "cross-sectional placebo," is based on the idea that, in the unlikely circumstance that a confounding event did occur in the event window, it would probably have been something that affected local-law bonds differently from all foreign-law bonds, not just English-law ones. Hence we test whether (i) bonds governed by New York law (direct placebo) or (ii) bonds issued by countries which have such bonds outstanding (indirect placebo), react to the English-law ruling (controlling for the true treatment effect by including all regressors from (R1)). We use the specification without risk controls because, first, it stacks the deck in favor of our finding an effect and, second, the matching specification, in which NY-law bonds are direct controls for EN-law bonds, is infeasible. As reported in Table 3, we find "nothing," i.e. the estimates are small, both statistically and relative to the analogous baseline estimates in columns 1 and 4 of Table 2, and some of them have the "wrong" sign.

3.4.2 Concern: Market Inefficiency/Measurement Error

We use a narrow event window in our baseline because it makes it hard to imagine violations of our identifying assumptions (Section 3.2). But it could also make it easier to imagine that (i) some bond prices reflect short-term market inefficiencies, like overreaction or fire sales, or (ii) some bonds do not trade over the window at all, so quoted prices might not be trading prices (as most of the bonds trade OTC, we cannot observe volumes). To address concerns about such market inefficiencies or measurement error, we double the event window to four trading days (plus an intermittent weekend) for the English-law experiment. This mitigates (i) by allowing for inefficiencies to work themselves out and (ii) by making it unlikely that a bond does not trade at all over the event window. (Such a test is, unfortunately, impossible for Argentina, as many of the events in the experiment are within a few days of each other or of potentially confounding events excluded for other reasons (Hébert and Schreger, 2017).) We report the results in Table 4. Our findings are robust.

3.4.3 Concern: Exclusion restriction

A general concern is that our experiments cause our results through a channel other than increasing commitment not to restructure and its associated spillover effect.

One possibility is that the experiments affected treated bonds other than via governing law. That is unlikely for the English-law experiment, in which the defendant was a firm, not a country. But it could be conceivable for the NY-law experiment, which changed not only how hard it was to restructure NY-law bonds, but also how likely Argentina was to default. Thus, as mentioned in Section 3.2, to say our results follow from a change in commitment not to restructure, we must make an additional assumption (the exclusion restriction): Countries' propensity to borrow under NY law is not conditionally correlated with their exposure to Argentina's default. This seems like a weak assumption, in light of our matching and our region and risk controls, but we provide an additional test by removing bonds issued by South American countries from the sample. We report the results in Table 4. Our findings are robust.

Another, albeit more remote, concern is that the rulings under foreign law spill over onto local law by setting international precedents. Courts' use of foreign precedents is limited (e.g., Simon (2013)) and, to the extent that it happens, it would pose a problem only to the extent that countries that issue under one foreign (English or NY) law are systematically more likely to follow precedents from that law. Either way, we can use our cross-sectional placebo analysis to rule out such spillovers, at least in the case of English and NY law. As discussed in Section 3.4.1, the English-law experiment had no effect on NY-law bonds. An analogous placebo test shows that the NY-law experiment had no effect on English-law bonds (also in Table 3). We think such potential precedent-based spillovers are most likely between English and New York law, as they are the main foreign laws countries issue under. Their absence reassures us that such spillovers to other laws are unlikely.

4 Model and Quantification

Here we develop a model of sovereign default/restructuring (Section 4.1). We use it to investigate the mechanism behind the empirical findings above (Section 4.2), to interpret their magnitudes (Section 4.3), and to calculate elasticities for a (back-of-the-envelope) policy counterfactual (Section 4.4).

4.1 Model Set-up

We consider a one-period model of a sovereign debtor. At the end of the period, it generates random output and its outstanding bonds come due. We denote the distribution function of output under the risk-neutral ($\mathbb{Q}$) measure by $F(z)$ and its density by $f(= F')$.

The model has two key ingredients.

1. The sovereign has a willingness-to-pay problem.[21] It has the option to default strategically, but default destroys a fraction c of the output, i.e. the default cost is cz.

[21] On the importance of the problem, see Reinhart and Rogoff (2009), who say

> If the reader has any doubt that willingness to pay rather than ability to pay is typically the main determinant of country default, he or she need only [observe] that more than half of defaults by middle-income countries occur at levels of external debt relative to GDP below 60 percent, when, under normal circumstances, real interest payments of only a few percent of income would be required to maintain a constant level of debt relative to GDP (p. 54).

2. The sovereign has two different types of bonds with different haircuts in the event of default (viz. restructuring). It has "rigid bonds" D_r with haircut h_r in default and it has "flexible bonds" D_f with haircut h_f in default. Rigid bonds are harder to restructure, corresponding to a lower haircut: $h_r < h_f$.

In our empirical environment, foreign-law bonds correspond to rigid and local-to flexible. The dependent variable in our experiments, ΔComm_t, corresponds to a decrease in the haircut h_r on rigid bonds since, in the language of Section 2.1, it increases the probability π with which hold-outs win in litigation. (The set-up in Section 2.1 also provides a micro-foundation for the heterogeneous haircuts ($h_f \neq h_r$) as b and π vary with governing law.)

The ingredients above generate the following trade-off between default and repayment. If it defaults, the sovereign suffers a deadweight loss in terms of destroyed output cz (per the first ingredient) but enjoys a lower repayment, repaying $(1-h_i)D_i$ instead of D_i on each type of bond $i \in \{r, f\}$ (per the second ingredient). This can be seen from the following expression for the sovereign's payoff:

$$\text{payoff} = \begin{cases} z - D_r - D_f & \text{if repay} \\ \\ z - cz - (1-h_r)D_r - (1-h_f)D_f & \text{if default.} \end{cases} \tag{3}$$

Observe that the sovereign defaults on either all debt or none—we focus on "cross default," abstracting from selective default for now.[22] However, we incorporate it explicitly in two extensions in Appendix C and show that our results are robust.

We normalize the face value of each of the sovereign's bonds to one and assume

[22]Such selective default has become harder in recent decades due to the rise in bond terms, such as cross-default clauses, which say a default on one class constitutes a default on another, and pari passu clauses, which preclude payments to one class without payments to another (see Choi, Gulati, and Posner (2012))—that is what the Argentina experiment is all about: Argentina's selective default was deemed a violation of pari passu.

that they are priced competitively: The price p_i of a bond of type $i \in \{r, f\}$ is

$$p_i = \mathbb{Q}[\text{repay}] + \mathbb{Q}[\text{default}](1 - h_i), \tag{4}$$

where, for notational simplicity, the price is inflated by the risk-free rate.

4.2 Model Results

The sovereign defaults if its payoff from defaulting exceeds its payoff from repayment, or, from equation (3), if

$$z - cz - (1 - h_r)D_r - (1 - h_f)D_f > z - D_r - D_f. \tag{5}$$

Re-writing, we see that it defaults whenever its assets z are below a threshold, which we denote by z^*:

$$z^* := \frac{h_r D_r + h_f D_f}{c}. \tag{6}$$

Thus, the $\mathbb{Q}$-probability of default is $F(z^*)$. We have the following immediate comparative statics:

Lemma 1. *The $\mathbb{Q}$-probability of default $F(z^*)$ is*

1. *decreasing in the default cost c,*
2. *increasing in the amount of outstanding debt of each type D_i, and*
3. *increasing in the haircut h_i on each type of debt.*

The first two comparative statics are typical of models with strategic default. The third points to something that is more specific to our environment: Increasing the haircut on either type of debt increases the probability of default on both types.

Our main results follow from comparative statics of the bond prices p_r and p_f with respect to h_r, capturing how the prices of each type of bond respond to our experiments. To derive them, we can use the default threshold in equation (6) to re-write the bond price in equation (4) as

$$p_i = 1 - F(z^*)h_i. \tag{7}$$

We first compute the sensitivity of the price of rigid bonds to its own haircut:

$$\frac{\partial p_r}{\partial h_r} = -F'(z^*)\frac{\partial z^*}{\partial h_r}h_r - F(z^*) \tag{8}$$

$$= -f(z^*)\frac{D_r}{c}h_r - F(z^*) < 0. \tag{9}$$

This expression captures how an increase in the haircut h_r harms r-bondholders in two ways. First, it decreases what they get in default (per the second term) and, second, it makes default more likely, because a higher haircut in default is attractive to the sovereign (per the first term).

We also compute the sensitivity of the price of flexible bonds to the haircut on rigid bonds:

$$\frac{\partial p_f}{\partial h_r} = -F'(z^*)\frac{\partial z^*}{\partial h_r}h_f \tag{10}$$

$$= -f(z^*)\frac{D_r}{c}h_f < 0. \tag{11}$$

This expression captures that an increase in the haircut h_r harms f-bondholders. This is the spillover effect of the haircut on rigid bonds on the price of flexible: It increases the default probability, making f-bonds less likely to be repaid.

Equations (9) and (11) correspond to our first two main empirical findings (Section 3.3.1 and Section 3.3.2), translated in the language of prices and haircuts instead of spreads and commitment. The first describes how the price of rigid bonds responds to a change in its haircut:

Proposition 1. *The price of rigid bonds decreases if their haircut increases, i.e.*

$$\left.\frac{\partial p_r}{\partial h_r}\right|_{D_r>0} < 0. \tag{12}$$

The second main result describes how the price of flexible bonds responds to a change in the haircut on rigid:

Proposition 2. *The price of flexible bonds decreases if the haircut on rigid bonds increases, i.e.*

$$\left.\frac{\partial p_f}{\partial h_r}\right|_{D_r>0} < 0. \tag{13}$$

4.3 Implied Haircut Changes

The model can also help us to interpret our experiments, which do not immediately lend themselves to quantification—the Imcopa ruling made restructuring easier and, on net, the Argentina rulings made it harder, but by how much did they affect the haircuts h_r? Here we show how to use the model to back this number out from our empirical estimates and a few other sufficient statistics.

Using a first-order approximation, we derive a system of equations for the changes in yields on r- and f-bonds, denoted by Δs_r and Δs_f, in terms of the change in haircut on rigid debt induced by the experiments, denoted by Δh_r:

Proposition 3. *A small change in the haircut on rigid bonds Δh_r induces changes in the spreads of r- and f-bonds, Δs_r and Δs_f, approximately as follows:*

$$\Delta s_r \approx \frac{1}{p_r T_r}\Big(F_{h_r} h_r + F(z^*)\Big)\Delta h_r, \tag{14}$$

$$\Delta s_f \approx \frac{1}{p_f T_f} F_{h_r} h_f \Delta h_r, \tag{15}$$

where T_i is the maturity of bonds of type $i \in \{r, f\}$ and F_{h_r} is short-hand for $\frac{\partial}{\partial h_r} F\big(z^*(h_r)\big)$.

We solve the system for countries with both directly and indirectly treated bonds outstanding. As we detail in Appendix E, we use our estimates of Δs_r and Δs_f from our matching specifications, the five-year CDS premiums from Markit to calculate $F(z^*)$, five-year bond maturities T_r and T_f to match the CDS, the model-implied haircuts from equation (7) for h_r and h_f, and yields adjusted for risk-free rates and currency risk from our data. These quantities are sufficient statistics to solve for Δh_r; we do not need to know the costs of default c or the sensitivity of the default probability to the r-haircut, F_{h_r} (in fact we back that out too; see the next section).

Solving (14) and (15), we estimate that haircuts increased by about 37 bps from the High Court ruling and decreased by, coincidentally, the same amount of 37 for each event in the Argentina saga on average. These magnitudes accord with our perception of the events as mattering for specific restructuring instruments but not revolutionizing market practice. We think that affords some validation that our model, despite being highly stylized, could be quantitatively relevant, or at least that the framework could be a useful building block for future (less highly-stylized) quantitative work.

4.4 The Default-to-haircut Elasticity and a Policy Counterfactual

We can also use the model to estimate by how much an increase in haircuts increases the default probability, a potentially relevant input for sovereign debt modeling and policy analysis. As many models and policies do not distinguish between r- and f-bonds, we work with the average haircut here: $\bar{h} := \varphi h_r + (1-\varphi)h_f$, where $\varphi := D_r/(D_r + D_f)$. The next result describes how to use observable quantities to translate F_{h_r}, which we can back out from the system in Proposition 3, into the desired elasticity:

Proposition 4. *The elasticity of the default probability $F(z^*)$ with respect to the average haircut $\bar{h}$ is*

$$\mathscr{E} \equiv \frac{\mathrm{d}F(z^*)/F(z^*)}{\mathrm{d}\bar{h}/\bar{h}} = \frac{\bar{h}}{\varphi F(z^*)} F_{h_r}, \tag{16}$$

where F_{h_r} solves the system in Proposition 3.

Applying formula (16) to averages from the raw data and the solution of the system

(14) and (15), we find that $\mathscr{E} \approx 14.16$ for English-law issuers and $\mathscr{E} \approx 4.74$ for NY-law issuers. I.e. a 1% increase in haircuts increases the risk-neutral default probabilities by 14.16% for English law and 4.74% for NY. That commitment benefits English- more than NY-law issuers likely reflects that English-law issuers have higher debt and lower default costs, i.e. that the term $F_{h_r}/\varphi = f(z^*)(D_r + D_f)/c$ is higher in the expression for $\mathscr{E}$.

To illustrate how $\mathscr{E}$ can be used for policy counterfactuals, we use the NY-law number to approximate by how much a change in haircuts changes the expected deadweight costs of default for Argentina, for which we have an estimate of these costs from Hébert and Schreger (2017). To do so, we use the Euler approximation to write

$$\Delta \text{ } \mathbb{Q}\text{-expected cost} = \Delta \int_0^{z^*} cz \mathrm{d}F(z) \approx \frac{cz^* \mathscr{E}}{\bar{h}/F(z^*)} \Delta \bar{h}, \tag{17}$$

having used the fundamental theorem of calculus and substituted for $\mathscr{E}$ from above. Replacing cz^* with Hébert and Schreger's estimate, $\mathscr{E}$ with ours, and the other parameters with data on Argentina, we find that a one percentage point increase in haircuts saves about 1.7 billion USD in deadweight costs of default. (See Appendix E for omitted details from this section.)

5 External Validity: Additional Events

We interpret our results broadly—bondholders' commitment not to restructure disciplines a sovereign debtor. They are based on two distinct experiments, covering bribes and threats under English and NY laws. But external validity can always be a question. To address it further, we consider two additional experiments, both of which made it harder to restructure English-law debt (Section 5.1). Although each presents greater identification challenges than our baselines (Section 5.2), together they affirm its message: The commitment not to restructure one class of bonds benefits not only that class but others as well.

5.1 The Ruling and (Dropped) Appeal: Assenagon Asset Management v Irish Bank Resolution Corp

Here we exploit a ruling and its (dropped) appeal pertaining to another controversial restructuring device: exit consents, which are the stick to consent payments' carrot: When tendering creditors part with their bonds ("exit"), they agree ("consent") to changes in the terms of other bonds in the same class. Like the other threats and bribes, they counter the hold-out problem, but can create a hold-in problem, and therefore be deemed coercive.[23]

Just two months after its ruling on consent payments in Azevedo v Imcopa (Section 2.2), the High Court ruled on exit consents.[24] The case was brought against Anglo Irish by a bondholder, Assenagon Asset Management, who held out from an exchange offer. After having received just €170 for Anglo-Irish bonds with face value of €17 million(!), Assenagon sued, submitting, inter alia, that the exit consents in the exchange offer were abusive. The presiding judge, Michael Briggs, agreed, ruling on July 27, 2012 that the restructuring was illegal. It was the Court's first decision on exit consents and thus, as noted in the decision, "of wide importance within the bond market." It made restructuring all English-law bonds harder, thus increasing debtor commitment.

Anglo Irish appealed. The appeal was withdrawn on April 22, 2013, affirming that exit consents could not be used and thus increasing commitment further.

5.2 Estimation and Identification Challenges

We run the same regressions as in our baseline specifications (R1), setting the treatment variable ΔComm_t equal to one on the trading day after the respective event, that is, on July 30, 2012 (the Assenagon ruling) and April 23, 2013 (the withdrawal

[23] They are nonetheless alive and well in the sovereign debt market broadly, as evidenced by a recent proposal to restructure Lebanese debt, namely Luo, Smith, and Xiao (2020).

[24] See Assenagon Asset Management SA v Irish Bank Resolution Corp Ltd (formerly Anglo Irish Bank Corp Ltd) [2012] EWHC 2090 (Ch) (27 July 2012).

of the appeal).

We mitigate identification concerns in the same way as in the baseline (Section 3.2). Thus, as above, to be a threat, an event must happen at nearly the same time as the experiments and affect both treatment and control bonds differently within the same currency and region/risk category.

Unlike above, such a threat is more than a remote possibility. Indeed, each experiment is at risk of being confounded by another event. The ruling could be confounded by a speech that ECB President Mario Draghi gave the previous day, in which he promised to do "whatever it takes to preserve the euro," a statement that could have affected treated and control bonds differently. The withdrawal of the appeal could be confounded by the Appellate Court upholding the baseline ruling in Azevedo v Imcopa, a decision that, by construction, affects treated and control bonds differently and goes in the opposite direction of the withdrawn appeal in Assenagon v Anglo Irish.

Although these threats make us more tentative in our interpretation, neither undermines our identification. The first experiment remains valid in an efficient market in which the information in Draghi's speech would be fully incorporated in prices at the start of our event window. The second remains valid as long as the effect is interpreted as the net effect of the two events: the withdrawn appeal in the Assenagon case and the upheld ruling in Azevedo. There is little doubt that the net effect of the appeals is to make restructuring harder: Assenagon outweighs Azevedo. It overturned rather than upheld market practice, and it was surprising to market participants (e.g., Skadden (2012)).

5.3 Do Sovereign Creditors Benefit from Committing Not to Restructure?

Table 5 reports the estimated coefficients for the matching specification. The signs of the estimates on all four coefficients of interest line up with the baseline experiments:

Increasing commitment leads spreads on both directly and indirectly treated bonds to decrease for both the ruling and withdrawal of the appeal. However, given Draghi's confounding speech, we are cautious in our interpretation of the estimates for the ruling, especially for the indirect effect, which relies heavily on comparisons of local-currency Eurozone bonds. Indeed, that estimate is not statistically significant.

6 Conclusion

We bring evidence from natural experiments to bear on the value of commitment not to restructure sovereign bonds. Our research design allows us to distinguish between the value for bondholders individually and collectively, and thereby assess cross-bond spillovers, like those implied by the "burglar alarm" view that to commit not to restructure one bond is only to encourage the restructuring of others (Bolton and Jeanne (2007, 2009)).

Our empirical results suggest a perspective new to the literature. We find that committing not to restructure one bond is more like a successful neighborhood crime watch than a self-defeating burglar alarm, viz. it benefits both that bond and others as well. We develop a model of the "crime watch" perspective that not only formalizes the mechanism behind it but also proves useful to quantify natural experiments and analyze policy counterfactuals.

Figure 1: **Treated and control countries.** Countries in the treatment group for the shocks to English law and New York law are in blue/red. Countries in gray are in the control group. Countries in white are not part of the sample.

Panel A. English-law experiment

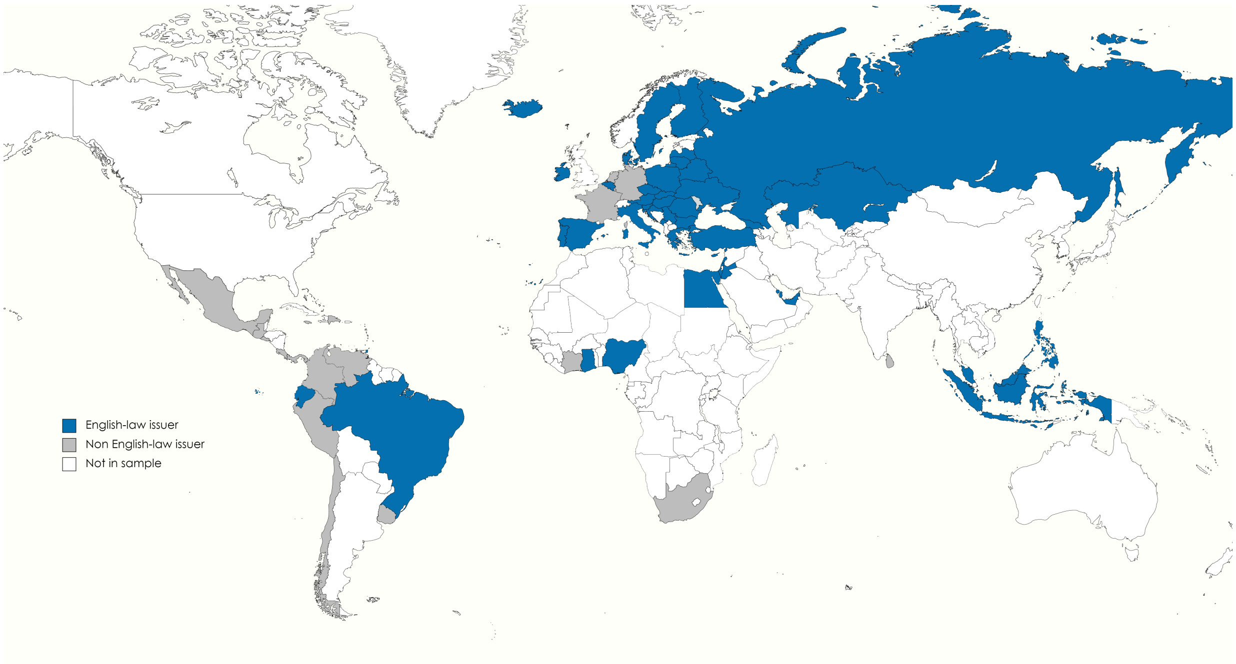

Panel B. New York-law experiment

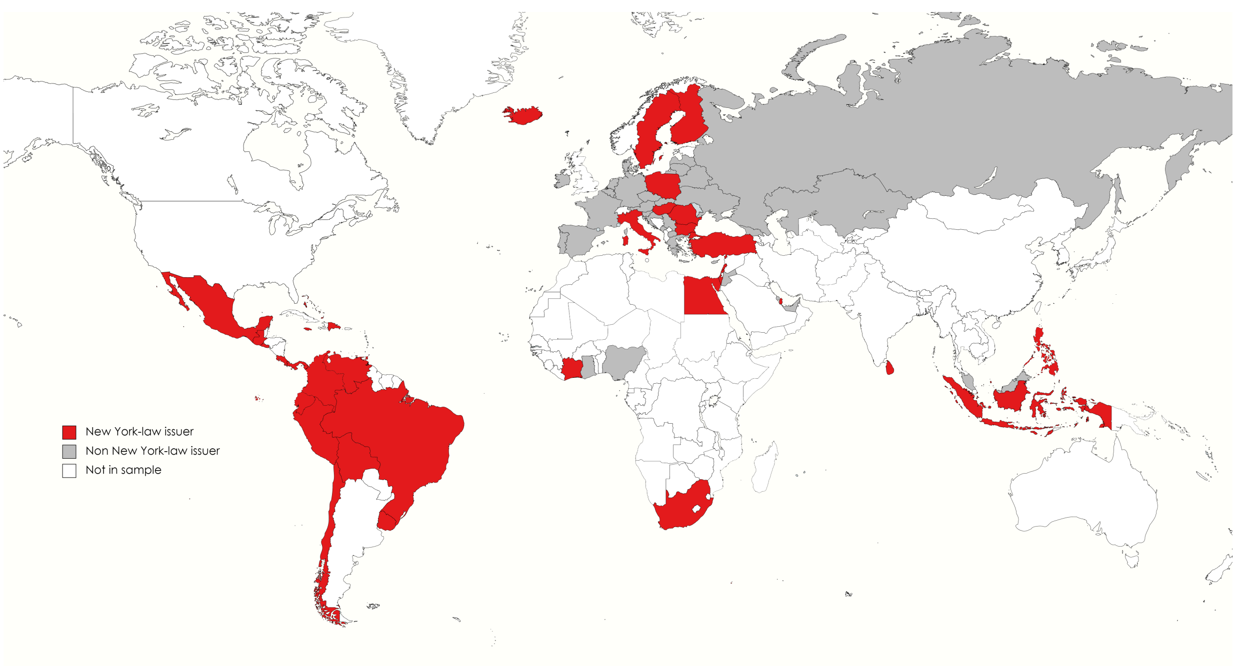

Table 1: **Summary statistics on the event date.** We report statistics for the 67 countries with bonds in our sample on the eve of the (first) event, For the Imcopa experiment, this is Tuesday, May 29, 2012, ahead of the announcement of the ruling on May 30. For the Argentina experiment, this is Friday November, 23, 2012, ahead of the first ruling identified by Hébert and Schreger (2017) (Judge Griesa denying a request for stay by exchange bondholders, immediately appealed by Argentina and exchange bondholders), on Monday, November 26 (3:43pm EST). In the first three columns, we report statistics by bond group (direct treatment, indirect treatment, and control), with the total number of bonds in the first row and group-means in the remaining rows. These columns provide bond-level comparisons between the two treatment groups and the control group. The other columns report the distributions across countries of the same variables averaged by country (e.g., the average number of bonds per country is 20.48). These columns describe the country-level heterogeneity in the respective variables.

	Bond-level statistics			Distribution of country means				
Variable	Direct	Indirect	Control	Mean	St. Dev.	1%	Median	99%
				English-law experiment				
Number of bonds	235	715	331	19.12	21.74	1	12	123
Face value per bond ($bn)	1.45	2.20	3.29	1.52	1.42	0.18	0.96	6.87
Yield (%)	5.20	4.29	3.01	5.19	3.08	0.58	4.52	13.85
Credit spread (%)	3.91	2.11	1.36	3.39	3.08	-1.38	2.87	11.51
Maturity	8.17	8.22	9.81	9.15	4.63	2.20	8.05	22.22
Fraction English law	1.00	0.00	0.00	0.35	0.39	0.00	0.14	1.00
Fraction New York law	0.00	0.23	0.40	0.24	0.35	0.00	0.00	1.00
Fraction local law	0.00	0.73	0.59	0.39	0.37	0.00	0.29	1.00
				New York-law experiment				
Number of bonds	311	521	668	19.49	21.19	1	14	116
Face value per bond ($bn)	1.13	3.89	2.62	1.55	1.58	0.18	1.00	8.37
Yield (%)	3.61	3.62	3.06	4.23	2.84	0.53	3.84	13.85
Credit spread (%)	2.47	1.01	1.68	2.80	2.82	-1.55	2.11	11.51
Maturity	10.74	7.93	7.56	9.16	4.49	1.71	8.27	21.74
Fraction English law	0.00	0.15	0.26	0.34	0.39	0.00	0.14	1.00
Fraction New York law	1.00	0.00	0.00	0.25	0.36	0.00	0.00	1.00
Fraction local law	0.00	0.81	0.73	0.39	0.37	0.00	0.28	1.00

Table 2: **Baseline estimates: The effect of the ruling on directly treated and indirectly treated bonds.** We estimate the DiD regression (R1) (Section 3.3). We report standard errors in parentheses, clustered at the country level, and at the country-match level for matching specifications to account for sampling with replacement of control bonds. For the English-law specifications, $\Delta\text{Comm}_t = -1$ on May 31st, 2012, and zero otherwise. For the NY-law specifications, $\Delta\text{Comm}_t = \Delta P_t^{\text{AR}} \mathbb{1}_t^{\text{Event}}$, that is, the change in the CDS-implied, risk-neutral default probability for Argentina, multiplied by an event indicator equal to one for 15 dates corresponding to the end of the two-trading day event window around rulings in the Argentine restructuring saga catalogued by Hébert and Schreger (2017). The dependent variable—the bond-level credit spread—is expressed in percentage points.

	English Law			New York Law		
Risk control	None	Rating FE	Matching	None	Rating FE	Matching
Direct$\times\Delta\text{Comm}_t$	-0.1101	-0.0496	-0.0580	-0.3545	-0.2917	-0.3869
	(0.0413)	(0.0289)	(0.0182)	(0.1728)	(0.0965)	(0.1790)
Indirect$\times\Delta\text{Comm}_t$	-0.1324	-0.0630	-0.0425	-0.1472	-0.1272	-0.1523
	(0.0503)	(0.0286)	(0.0088)	(0.0666)	(0.0476)	(0.0759)
Direct	-0.0065	0.0001	-0.0028	-0.0010	0.0027	-0.0000
	(0.0020)	(0.0010)	(0.0011)	(0.0013)	(0.0014)	(0.0010)
Indirect	-0.0051	0.0004	-0.0015	-0.0011	0.0017	0.0019
	(0.0013)	(0.0009)	(0.0003)	(0.0015)	(0.0011)	(0.0006)
Region-Time	Yes	Yes	No	Yes	Yes	No
Currency-Time	Yes	Yes	matched	Yes	Yes	matched
Rating-Time	No	Yes	matched	No	Yes	matched
Bond pair-time	No	No	Yes	No	No	Yes
Bonds / Pairs	1281	1281	207	1325	1325	251

Table 3: **Placebo tests.** We run two placebo tests. *Time-series:* We re-estimate the DiD regression (R1) for the ruling in *Azevedo v Imcopa* assigning the event indicator to 644 placebo dates. We then report in brackets the fraction of these dates for which the t-statistic of the estimated placebo-coefficients is larger than that obtained for the event window. The row labelled "Joint" reports this statistic for the joint results (β_1, β_2). *Cross-sectional:* We estimate the placebo-treatment effect by switching the treatment law: for the English-law ruling in *Azevedo v Imcopa*, we estimate the effect on NY-law bonds (direct) and bonds of NY-law issuers (indirect) on the event date against the remaining control group. For the New York-law rulings in *Argentina v NML Capital* we estimate the effect on English-law bonds (direct) and bonds of Egnlish-law issuers (indirect) relative to the remaining control group.

	Time-series: placebo dates English law		Cross-sectional: placebo treatment English law	New York law
Risk control	None	Matching	None	None
Direct $\times\Delta\text{Comm}_t$	[0.107]	[0.079]	−0.0101	0.0561
			(0.0523)	(0.0642)
Indirect $\times\Delta\text{Comm}_t$	[0.144]	[0.030]	−0.0670	−0.0432
			(0.0673)	(0.0404)
Joint	[0.048]	[0.002]		
Region-Time	Yes	No	Yes	Yes
Currency-Time	Yes	matched	Yes	Yes
Bond pair-Time	No	Yes	No	No

Table 4: **Robustness tests** We estimate the DiD regression (R1) for the English-law experiment (i) for a longer event window of four trading days (credit spread changes from $t-1$ to $t+3$), and (ii) in levels of credit spreads before and after the ruling with a bond fixed effect rather than two-day changes. The last two columns report a robustness test on the NY-law experiment excluding all South American issuers from the sample. We report all estimates for the baseline specifications without risk controls and with matched bond pairs.

	English Law				New York Law	
Robustness	Four-day window		Levels		Excl. South America	
Risk control	None	Matching	None	Matching	None	Matching
Direct$\times\Delta\text{Comm}_t$	-0.1244	-0.1365	-0.0567	-0.0550	-0.3545	-0.3869
	(0.0360)	(0.0314)	(0.0303)	(0.0281)	(0.1728)	(0.1790)
Indirect$\times\Delta\text{Comm}_t$	-0.0879	-0.0348	-0.1243	-0.0410	-0.1472	-0.1523
	(0.0321)	(0.0147)	(0.0521)	(0.0086)	(0.0666)	(0.0759)
Region-Time	Yes	No	Yes	No	Yes	No
Currency-Time	Yes	matched	Yes	matched	Yes	matched
Rating-Time	No	matched	No	matched	No	matched
Bond pair-time	No	Yes	No	Yes	No	Yes
Bond	No	No	Yes	No	No	No

Table 5: **Estimates from a landmark ruling on exit consents and its appeal.** We estimate the DiD regressions (R1), using two-day event windows around the ruling (July 27, 2012) and the withdrawal of its appeal (April 22, 2013). We report clustered standard errors in parentheses. These are clustered at the country level for fixed-effects specifications and at the country-control bond level for matching specifications.

	Ruling (July 27,2012)	Appeal (April 22, 2013)
Direct$\times\Delta$Comm$_t$	-0.0514	-0.0555
	(0.0238)	(0.0194)
Indirect$\times\Delta$Comm$_t$	-0.0019	-0.0426
	(0.0070)	(0.0171)
Direct	-0.0026	-0.0026
	(0.0012)	(0.0012)
Indirect	-0.0014	-0.0013
	(0.0003)	(0.0003)
Currency-Time	matched	matched
Rating-Time	matched	matched
Bond pair-time	Yes	Yes
Bond pairs	207	207

Table 6: **Baseline estimates: The relative effect of the ruling on directly treated versus indirectly treated bonds within country.** We estimate the DiD regression (R1) (Section 3.3). We report standard errors in parentheses, clustered at the country level.

	English Law		New York Law	
Risk control	None	Rating FE	None	Rating FE
Direct $\times \Delta \text{Comm}_t$	0.0062	0.0062	-0.0206	-0.0076
	(0.0204)	(0.0206)	(0.0569)	(0.0577)
Direct	-0.0008	-0.0008	0.0014	0.0013
	(0.0004)	(0.0004)	(0.0007)	(0.0007)
Region-Time	Yes	Yes	Yes	Yes
Currency-Time	Yes	Yes	Yes	Yes
Rating-Time	No	Yes	No	Yes
Bonds	1281	1281	1325	1325

A Data

A.1 Geography

Our sample includes the following issuers, assigned to regions as listed below:

- **Europe**: Austria, Belgium, Bosnia and Herzegovina, Bulgaria, Croatia, Cyprus, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Iceland, Ireland, Isle of Man, Italy, Latvia, Lithuania, Luxembourg, Montenegro, the Netherlands, North Macedonia, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden.
- **South America**: Bolivia, Brazil, Chile, Colombia, Ecuador, Paraguay, Peru, Uruguay, Venezuela.
- **Central America**: Barbados, Bahamas, Belize, Costa Rica, Dominican Republic, El Salvador, Grenada, Guatemala, Honduras, Jamaica, Mexico, Panama, Trinidad and Tobago.
- **Middle East**: Bahrain, Egypt, Israel, Jordan, Lebanon, Qatar, UAE.
- **Africa**: Côte d'Ivoire, Ghana, Nigeria, South Africa.
- **Central Asia**: Azerbaijan, Belarus, Georgia, Kazakhstan, Russia, Turkey, Ukraine.
- **East Asia**: Indonesia, Malaysia, Philippines, Sri Lanka.

A.2 Risk-Free Interest Rate Proxies

The spread at time t of bond i with yield $y_{i,t}$ maturity T, currency denomination f is defined as $s_{i,t} := y_{i,t} - r_{t,T,f}$, where the yield and $r_{t,T,f}$ is the currency-specific risk-free rate from t to T. As $r_{t,T,f}$ is unobservable, we proxy for it with the rate on the appropriate (i.e. currency and maturity matched) swap rate.[25]

[25]Both rates are (basically) risk-neutral expectations of the average spot rate over the maturity. The difference is the time weighting: The swap rate discounts longer maturities. If the discount rate is about

We match bonds to swap rate maturities of 1, 2, 3, 4, 5, 10, 15, 20, and 30 years. Where swap rates for the longer maturities are not available, we use the longest available maturity. Where no swap rate is available, we drop all bonds denominated in that currency from the sample. In Table SA.1 in the Supplementary Appendix, we list the different currencies in our sample along with the floating rate underlying the swap used to construct credit spreads.

B Matching procedure

The matching procedure is as follows. For each (directly or indirectly) treated bond, we consider all untreated bonds that have (i) the same currency denomination as, (ii) maturity within one year of, and (iii) a credit rating within two notations of the treated bond's. If the set is empty, we discard the treated bond. If it is not, we choose the bond with the closest pre-event credit spread to the treated bond as its control bond. (Note that we do not exclude matching multiple treated bond with the same control bonds, i.e. we "match with replacement.") In Table SA.4 and Table SA.5 of the Supplementary Appendix, we list the matched pairs for directly and indirectly treated bonds.

C Selective Default

In our baseline model, we capture the "crime watch" effect of the ruling by assuming that making it unattractive to restructure one bond makes it unattractive to start a restructuring altogether. That abstracts from the "burglar alarm" effect stressed in the literature, which focuses on selective default.

Although that abstraction explains our empirical results, it is not a perfect description of reality: Selective defaults do happen in practice (Erce and Mallucci (2018)).

constant (e.g., because maturities are short or spot rates low) then the two rates are about equal. Even if it is not, swap rate is still an effective proxy for the risk-free rate for our purposes, as it captures variation in currency-specific risk-free term-structures.

To show that our results do not turn on an unrealistic assumption, we show that our mechanism does not turn on there being no selective default whatsoever. We do that in two ways. First, we show that if the haircuts in Section 4.2 are interpreted as expected, not realized, haircuts that come out of a risky version of the restructuring game in Section 2.1, then sometimes ostensible selective defaults do occur. The rest of our analysis, however, does not change. Second, we show that our main mechanism is robust to including strategic selective default with bond-specific default costs: Decreasing the haircut on one class makes default on the other more attractive/likely. The extension also provides a microfoundation for our baseline assumption of all-or-nothing default (albeit under a parametric assumption).

C.1 Expected Haircuts in Risky Restructurings

So far, we have assumed that restructuring is riskless. If the sovereign chooses to default it restructures its debt for sure. Combined with the assumption that it chooses to default on either all debt or none, this precludes selective default. In reality, however, (i) restructuring is risky and (ii) selective defaults do occur. Here we show how including uncertainty in the restructuring game in Section 2.1 captures both of these things but does not affect the main analysis. The only difference is one of interpretation. The haircuts in the baseline model should be interpreted as expected, not realized, haircuts.

In equation (2), the sovereign offers each class of bondholders the maximum haircut that would make them break even for sure. In practice, however, the sovereign would not know bondholders' belief π that litigation would be successful. To accommodate such uncertainty, we assume the sovereign treats π as a random variable and maximizes the expected haircut (for a fixed b): Using h to denote the $\mathbb{Q}$-expected haircut[26] and

[26]With this interpretation, the assumption that $h_f > h_r$ says that f- are more likely than r-bonds to be restructured selectively. That lines up with Erce, Mallucci, and Picarelli's (2021) finding that sovereigns are more likely to default selectively on local- than foreign-law bonds.

$\hat{h}$ the offered haircut, we can write:

$$h := \mathbb{Q}[\text{accept}]\hat{h} = \mathbb{Q}[\hat{h} \leq b/\pi]\hat{h} = G\big(b/\hat{h}\big)\hat{h}, \tag{18}$$

where G is the distribution function of π.

Supposing the first-order approach is valid here, the maximizer solves

$$G\big(b/\hat{h}\big) = g\big(b/\hat{h}\big)b/\hat{h}, \tag{19}$$

where $g := G'$ is the density of π. In this case, the restructuring fails, so there is no observed haircut, with probability $1 - G(b/\hat{h})$, i.e. if $\pi > b/\hat{h}$ at the optimum. A "selective default" is observed if the condition holds for one class of debt and not the other.

We conclude with an example, which illustrates a case in which the first-order approach assumed above is indeed valid: Suppose $G(\pi) = \frac{(A+B\pi)^{1+\alpha}}{A+B}$ for $\alpha > 0$ and $0 \leq \pi \leq 1$. Thus equation (19) for the offered haircut reads

$$(A + Bb/\hat{h})^{1+\alpha} = (1+\alpha)B(A + Bb/\hat{h})^{\alpha}b/\hat{h}. \tag{20}$$

In this case, $\hat{h} = \alpha bB/A$, the probability of acceptance $G(b/\hat{h}) = A^{1+\alpha}(1+1/\alpha)^{1+\alpha}/(A+B)$, and the expected haircut $h = G(b/\hat{h})\hat{h} = \alpha bA^{\alpha}(1+1/\alpha)^{1+\alpha}/(A+B)$. So increasing the bribe b increases the expected haircut, keeping the probability of acceptance constant.

C.2 Bond-specific Default Costs

So far, we have assumed that any choice to default/restructure results in a loss of the fraction of output c. But one reason selective default could arise in practice is that defaulting on one class of bonds brings separate costs from defaulting on another.

Here we suppose that defaulting on one class i causes a loss of a fraction c_i of output z and defaulting on the other i' causes a loss of $c_{i'}$ of the remaining output. That means that defaulting on one class is relatively less costly to the sovereign if it defaults on the other as well, because it destroys the same fraction of a smaller amount: It destroys $c_{i'}(1-c_i)z$ instead of $c_{i'}z$. The sovereign's payoff in this case is:

$$\text{payoff} = \begin{cases} z - D_r - D_f & \text{if repay,} \\ (1-c_r)z - (1-h_r)D_r - D_f & \text{if default on } r\text{-bonds,} \\ (1-c_f)z - D_r - (1-h_f)D_f & \text{if default on } f\text{-bonds,} \\ (1-c_r)(1-c_f)z - (1-h_r)D_r - (1-h_f)D_f & \text{if default on both.} \end{cases} \tag{21}$$

Observe that even if defaulting on class i' bonds selectively is unattractive, defaulting on them in concert with i bonds can be attractive, i.e. it can be that $c_{i'}z > h_{i'}D_{i'} > c_{i'}(1-c_i)z$.

Now we define the threshold z^{**} to make a parametric assumption. Define z^{**} as the output at which the sovereign is indifferent between defaulting on all debt or none, so

$$z^{**} - D_r - D_f = (1-c_r)(1-c_f)z^{**} - (1-h_r)D_r - (1-h_f)D_f \tag{22}$$

or

$$z^{**} = \frac{h_r D_r + h_f D_f}{1-(1-c_r)(1-c_f)}. \tag{23}$$

The parametric assumption is

$$z^{**} - D_r - D_f \geq (1-c_i)z^{**} - (1-h_i)D_i - D_{i'} \tag{24}$$

for $i \in \{r, f\}$. I.e., near z^{**}, the sovereign would prefer not to default selectively. The assumption above implies that the sovereign defaults on all debt or none, and z^{**} is the default threshold.

From equation (23) we see that default on *each* type of bond—indeed on all bonds—becomes more attractive as the haircut on *either* increases.

Thus the "crime watch" effect by which low haircuts on one bond benefit others arises even when selective default is allowed, whenever inequality (24) is satisfied. (As default costs and haircuts are likely to be random, depending on market and political conditions, it is enough for it to be satisfied with positive probability.)

The parametric assumption in equation (24) above says that it is attractive to default on both classes together, but neither individually. That might seem unintuitive or, even, pathological. But a little algebra suggests it is not. Substituting for z^{**} and re-writing, it reads

$$\frac{c_i}{c_{i'}} > (1-c_i)\frac{h_i D_i}{h_{i'} D_{i'}} \tag{25}$$

for $i \in \{r, f\}$. That is satisfied, e.g., whenever the bonds are symmetric, $c_r = c_f, h_r = h_f$, and $D_r = D_r$. That provides some support for our assumption of all-or-nothing default in the baseline model.

D Proofs

D.1 Proof of Proposition 1

The result is immediate from equation (9). □

D.2 Proof of Proposition 2

The result is immediate from equation (11). □

D.3 Proof of Proposition 3

To derive the result, we use the definition of yield, a linear approximation, and the chain rule:

1. By definition, the spread on a (zero-coupon) bond of type $i \in \{r, f\}$ is $s_i = -\frac{1}{T_i} \log p_i$ (keeping in mind our notational convention that the price is inflated by the risk-free rate).

2. Differentiating the price, we get

$$\mathrm{d}p_i = \mathrm{d}e^{-s_i T_i} = -e^{-s_i T_i} T_i \mathrm{d}s_i = -p_i T_i \mathrm{d}s_i. \tag{26}$$

3. Rearranging, approximating, and applying the chain rule, we get

$$\Delta s_i \approx -\frac{\Delta p_i}{p_i T_i} \approx -\frac{1}{p_i T_i} \frac{\partial p_i}{\partial h_r} \Delta h_r. \tag{27}$$

   Using equations (8) and (10) for $\frac{\partial p_i}{\partial h_r}$, we have the system in the statement of the proposition. □

D.4 Proof of Proposition 4

We write $F_{h_r} \equiv \frac{\partial}{\partial h_r} F(z^*)$ and $F_{\bar{h}} \equiv \frac{\partial}{\partial \bar{h}} F(z^*)$ using $cz^* \equiv h_r D_r + h_f D_f \equiv \bar{h} D$, for $D := D_r + D_f$:

$$F_{h_r} = F'(z^*) \frac{D_r}{c}, \tag{28}$$

$$F_{\bar{h}} = F'(z^*) \frac{D}{c}. \tag{29}$$

Dividing one by the other gives the result, recalling that $\varphi = D_r / D$.

E Details of Quantitative Analysis in Section 4.3 and Section 4.4

Here we describe omitted details from our quantitative analysis.

E.1 Data Selection

We begin with the full cross section of our sample of bonds the days before our baseline experiments (May 29, 2012 and November 23, 2012, respectively) and we retain only bonds that are (i) issued under English, New York, or local law, (ii) not missing yield or credit spread data, and (iii) issued in US dollars, euros, or local currency.[27] For the English-law experiment, we are left with sixteen countries for which we observe both local- and English-law bonds that satisfy all criteria; for the New York-law experiment, we are left with twelve countries.[28]

E.2 Yield Adjustments

We need to adjust the raw yields to account for our modeling assumptions, namely that risk-free rates are normalized to zero and that all bonds are denominated in the same currency:

- *Risk-free rate.* To undo the rate normalization, we can just replace yields in the model with credit spreads in the data (this follows immediately from multiplying both sides of equation (7) by e^{-r_f}).
- *Currency risk.* The model is written from the point of view of a US investor, that is, under the dollar risk-neutral measure.[29] We therefore need to adjust observed credit spreads on non-dollar bonds (see the "quanto adjustment" in Du and Schreger (2016)). The intuition is as follows: if the dollar is expected to appreciate against the currency of denomination (say, the euro) in the event of default, the dollar-denominated credit spread will exceed the observed euro-denominated credit spread even if the default probability and percentage haircut

[27] We discard the few bonds issued in Japanese yen, Swiss francs, or British pounds to simplify the credit spread currency adjustment below.

[28] Those countries are AT, BE, CZ, DK, ES, FI, HR, HU, MY, PH, PL, RO, RU, SE, SK, and TR for English law. For New York law, they are BG, CL, CO, CR, FI, HU, IT, PH, PL, RO, TR, ZA.

[29] I.e. the $\mathbb{Q}$ measure in the model should be interpreted as the one that makes asset prices discounted by the US dollar-denominated money market account martingales.

are equal for dollar- and euro-bonds. This is because, expressed in a common numéraire, the dollar haircut exceeds the euro haircut. As a result, dollar-denominated CDS on eurozone sovereigns tend to have higher premiums than euro-denominated ones, even though they share a common trigger event and recovery. Since euro-denominated CDS quotes tend to be reliable for most sovereigns, we adjust the credit spread on euro-denominated bonds by adding the difference between US dollar and euro denominated CDS premiums. For emerging markets sovereigns, however, local-currency CDS quotes are often unreliable (Du and Schreger, 2016). We therefore use the dollar-denominated five-year CDS premium as a proxy for the currency-adjusted credit spread in local-currency, local-law bonds.[30]

For each country, we define the spread variable s_i as the average currency-adjusted credit spread on its English-/local-law bonds for $i = r/i = f$.

E.3 Variables Definitions and Values

We use the following variables to solve the system in Proposition 3 and to compute the elasticity $\mathscr{E}$ in Proposition 4.

- Credit spreads: $s_r = 2.52\%$ and $s_f = 2.47\%$ from equal-weighted average of (currency-adjusted) English- and local-law spreads across countries. For the New York-law experiment, these numbers are $s_r = 1.88\%$ and $s_f = 1.76\%$ for (currency-adjusted) New York- and local-law spreads.
- Maturity: $T_r = T_f = 5$ from assuming a flat term structure to match the maturity of the most reliable CDS quotes. (For most countries in the English-law experiment, bonds in this sample have a similar average maturity of about six years anyway).

[30]While sovereign CDS in emerging markets are not triggered by default on local-currency, local-law bonds, using the CDS premium as a proxy for currency adjusted credit spreads is consistent with our model, in which the default event coincides between local- and foreign-law bonds.

- $\mathbb{Q}$-default probability: $F(z^*) = 1 - e^{-\frac{\xi T}{1-R}} = 17.5\%$, per the "triangle method" (Hébert and Schreger (2017, Appendix A.1). Here ξ is T-year CDS premium and R the recovery value from the Markit CDS data. We use $T = 5$ years because the five-year premiums are deemed most reliable (Hébert and Schreger (2017)). We compute the five-year default probability for each of the countries involved in this exercise and take an equal-weighted average.
- Prices p_i: $p_i = e^{-s_i T_i}$ by definition.
- Haircuts: $h_i = (1 - p_i)/F(z^*)$ from the model (equation (7)).
- Fraction of English-law bonds: $\varphi = 20.5\%$, the equal-weighted cross-country average.
- Fraction of New York-law bonds: $\varphi = 14.6\%$, the equal-weighted cross-country average.
- Average haircut: $\bar{h} = \varphi h_r + (1 - \varphi) h_f$ by definition.
- Spread changes: $\Delta s_r = 5.8$ bps and $\Delta s_f = 4.25$ bps, and, respectively, $\Delta s_r = 38.69$ bps and $\Delta s_f = 15.23$ bps, from Table 2 (these are the estimates from the matched regression, which is more representative of the restricted sample of bonds issued by countries both directly and indirectly treated bonds outstanding used here).
- The sensitivity of default to h_r: $F_{h_r} = 0.76$, from solving the system in Proposition 3 for the English-law experiment, and $F_{h_r} = 0.09$ for the New York-law experiment.

E.4 Variable Values for Argentina Policy Counterfactual

For the policy counterfactual in Section 4.4, we use data from Argentina the day before our ruling to line up be consistent with the other variables. The default costs, taken from Hébert and Schreger (2017), are estimated over a longer period that includes that date.

- $\mathbb{Q}$-default probability: $F(z^*) = 87\%$ applying the triangle method described above for five-year dollar-denominated CDS on Argentina.
- Haircut: $\bar{h} = 66\%$ (based on $F(z^*)$, $D_r = 31.6\%$—calculated as the fraction of NY-law and English-law debt relative to total debt—and Argentine credit spreads on dollar-denominated five-year bonds under local and New York law. For this exercise, we are careful to derive credit spreads from bonds that are not in default at the time of our baseline ruling).
- Total default cost: $cz^* = 27.8$ billion USD from Hébert and Schreger (2017, Table 3). They describe this number as the cost of the $\mathbb{Q}$-default probability going from zero to one, which arguably corresponds to the cost at our threshold z^*.

References

Arellano, C., X. Mateos-Planas, and J.-V. Ríos-Rull (2019). Partial Default. NBER Working Papers 26076, National Bureau of Economic Research, Inc.

Becker, T., A. Richards, and Y. Thaicharoen (2003). Bond restructuring and moral hazard: Are collective action clauses costly? *Journal of International Economics 61*(1), 127–61.

Bolton, P. and O. Jeanne (2007). Structuring and restructuring sovereign debt: The role of a bankruptcy regime. *Journal of Political Economy 115*(6), 901–924.

Bolton, P. and O. Jeanne (2009). Structuring and restructuring sovereign debt: The role of seniority. *The Review of Economic Studies 76*(3), 879–902.

Bradley, M. and M. Gulati (2014). Collective action clauses for the Eurozone. *Review of Finance 18*, 2045–2102.

Buchheit, L. C. and M. Gulati (2000). Exit consents in sovereign bond exchanges. *UCLA Law Review 48*, 59–80.

Bulow, J. and K. Rogoff (1989). A constant recontracting model of sovereign debt. *Journal of Political Economy 97*(1), 155–178.

Carletti, E., P. Colla, M. Gulati, and S. Ongena (2020). The price of law: The case of the Eurozone collective action clauses. *Review of Financial Studies* forthcoming.

Chamon, M., J. Schumacher, and C. Trebesch (2018). Foreign-law bonds: Can they reduce sovereign borrowing costs? *Journal of International Economics 114*(C), 164–179.

Choi, S. J., M. Gulati, and E. Posner (2012). The evolution of contractual terms in sovereign bonds. *Journal of Legal Analysis 4*(1), 131–179.

Clare, A. and N. Schmidlin (2014). The impact of foreign governing law on European government bond yields. Working paper, City University London, London, Mimeo.

Cruces, J. and C. Trebesch (2013). Sovereign defaults: The price of haircuts. *American Economic Journal: Macroeconomics 5*, 85–117.

Donaldson, J. R., E. Morrison, G. Piacentino, and X. Yu (2022). Restructuring vs. bankruptcy. Working paper, USC.

Dooley, M. (2000). Can output losses following international financial crises be avoided? NBER Working Papers 7531, National Bureau of Economic Research, Inc.

Dooley, M. and S. Verma (2003). Rescue packages and output losses following crises. In *Managing Currency Crises in Emerging Markets*, NBER Chapters, pp. 125–186. National Bureau of Economic Research, Inc.

Du, W. and J. Schreger (2016). Local currency sovereign risk. *Journal of Finance 71*(3), 1027–70.

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

Eichengreen, B. and A. Mody (2004). Do collective action clauses raise borrowing costs? *Economic Journal 114*, 247–64.

Erce, A. and E. Mallucci (2018). Selective sovereign defaults. International finance discussion paper 1239, Board of Governors of the Federal Reserve System.

Erce, A., E. Mallucci, and M. Picarelli (2021). A journey in the history of sovereign defaults on domestic-law public debt. Working paper D.T. 2106, Departamento de Economía, Universidad Pública de Navarra.

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.

Gelpern, A., B. Heller, and B. Setser (2016). Count the limbs: Designing robust aggregate clauses in sovereign bonds. In G. Martin, J. A. Ocampo, and J. Stiglitz (Eds.), *Too Little, Too Late: The Quest to Resolve Sovereign Debt Crises*, pp. 109–143. Columbia University Press.

Gertner, R. and D. Scharfstein (1991). A theory of workouts and the effects of reorganization law. *The Journal of Finance 46*(4), 1189–1222.

Hébert, B. and J. Schreger (2017). The costs of sovereign default: Evidence from Argentina. *American Economic Review 107*(10), 3119–45.

Jones Day (2012). Two recent cases test legality of consent payments and exit consents under English law. Insights.

Kahan, M. and B. Tuckman (1993). Do bondholders lose from junk bond covenant changes. *Journal of Business 66*(4), 499–516.

Kremens, L. (2023). Currency redenomination risk. Working paper, University of Washington.

Kropp, A., M. Gulati, and M. Weidemaier (2018). Sovereign bond contracts: Flaws in the public data? Working paper, Duke Law School Public Law & Legal Theory Paper No. 2018-42.

Krueger, A. (2002). A new approach to sovereign debt restructuring. Working paper, IMF.

Luo, B., C. Smith, and A. Xiao (2020). Restructuring Lebanese sovereign debt: Tackling the holdout problem. Working paper, Duke.

Panizza, U., F. Sturzenegger, and J. Zettelmeyer (2009). The economics and law of sovereign debt and default. *Journal of Economic Literature 47*(3), 651–98.

Philippon, T. (2009). The bond market's *q*. *The Quarterly Journal of Economics 124*(3), 1011–56.

Picarelli, M. O., A. Erce, and X. Jiang (2019). The benefits of reducing holdout risk: Evidence from the Euro CAC experiment, 2013–2018. *Capital Markets Law Journal 14*(2), 155–177.

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

Reinhart, C. and K. Rogoff (2009). *This Time Is Different: Eight Centuries of Financial Folly.* Princeton University Press.

Roe, M. J. (1987). The voting prohibition in bond workouts. *Yale Law Journal 97*, 232–279.

Schumacher, J., C. Trebesch, and H. Enderlein (2021). Sovereign defaults in court. *Journal of International Economics forthcoming.*

Shleifer, A. (2003). Will the sovereign debt market survive? *American Economic Review 93*(2), 85–90.

Simon, S. (2013). The supreme court's use of foreign law in constitutional rights cases: An empirical study. *Journal of Law and Courts 1*(2), 279–301.

Skadden (2012). Rewards and penalties in bond covenant consent solicitations under English law. Skadden, Arps, Slate, Meagher & Flom LLP & Affiliates Memorandum.

White & Case (2013). Bond consent solicitation payments. Insight: Financial restructuring and insolvency.

White, M. (2002). Sovereigns in distress: Do they need bankruptcy? *Brookings Papers on Economic Activity 1*, 287–319.

Zame, W. R. (1993). Efficiency and the role of default when security markets are incomplete. *American Economic Review 83*(5), 1142–1164.

Zettelmeyer, J., C. Trebesch, and M. Gulati (2014). The Greek debt restructuring: An autopsy. *Economic Policy 28*, 513–64.