



—  
INSTITUT DE HAUTES  
ÉTUDES INTERNATIONALES  
ET DU DÉVELOPPEMENT  
GRADUATE INSTITUTE  
OF INTERNATIONAL AND  
DEVELOPMENT STUDIES

---

Graduate Institute of International and Development Studies  
International Economics Department  
Working Paper Series

Working Paper No. HEIDWP05-2024

## **Obscure contract terms: an inadvertent pricing experiment**

**Stephen J. Choi**  
New York University

**Mitu Gulati**  
University of Virginia (Law)

**Ugo Panizza**  
Geneva Graduate Institute (International Economics)

**Robert E. Scott**  
Columbia University (Law)

**W. Mark C. Weidemaier**  
University of North Carolina at Chapel Hill (Law)

Chemin Eugène-Rigot 2  
P.O. Box 136  
CH - 1211 Geneva 21  
Switzerland

## OBSCURE CONTRACT TERMS: AN INADVERTENT PRICING EXPERIMENT

Stephen J. Choi, Mitu Gulati, Ugo Panizza, Robert E. Scott & W. Mark C. Weidemaier\*

### *Abstract*

*Contract terms that improve or reduce the likelihood of repayment of a debt should impact its price. That's basic economics. But what about a contract that is hundreds of pages long and has lengthy and complex terms that even the lawyers are unwilling to read? Believers in efficient markets might predict that variations that affect the likelihood of repayment in such obscure contract terms will be priced at the outset if there are profits to be made by exploiting these variations. An alternate view is that little attention is paid to the fine print in highly standardized contracts until the likelihood of default becomes sufficiently salient to make reading the fine print worthwhile. Using several inadvertent real-world experiments, we examine the question of how and when variations that are assumed to be standardized in obscure contract terms are priced.*

Keywords: Contract complexity; Market efficiency; Sovereign Debt

JEL Codes: G14; K12; G12

---

\* Respectively, faculty at New York University, the University of Virginia (Law), Geneva Graduate Institute (International Economics), Columbia University (Law) and the University of North Carolina at Chapel Hill (Law)

## Introduction

Bonds and other tradable securities are issued pursuant to detailed, lengthy contracts that govern investors' legal rights. These are largely standardized, form contracts, but the fine print can vary. From first principles, it seems that market prices should be sensitive to differences in the underlying contract, at least when those differences impact investors' legal rights. In markets that tend towards efficiency, the price of a security should incorporate public information about the security and its issuer. For example, if two securities are otherwise identical, but one confers contractual rights that might prove valuable in a default, one would expect investors to assign greater value to the more protective security. One should particularly expect that effect to manifest itself for riskier securities and as default becomes more likely. If this does not happen, it creates an arbitrage opportunity for sophisticated investors, whose trades should move the market towards efficiency.<sup>1</sup>

We focus on the market for sovereign bonds, where empirical research into the price impact of contract language has produced mixed results. Many pricing studies focus on highly-salient contract provisions such as Collective Action Clauses (“CACs”), which allow a supermajority of investors (e.g., holders of 75% of the outstanding principal amount) to approve a restructuring proposal and bind dissenters to the outcome.<sup>2</sup> The salience of CACs is highlighted by the fact that the prospectus or other sales document distributed to prospective investors typically describes the CAC in detail on the very first page.

A number of studies have examined the pricing implications of CACs with different voting thresholds (e.g., 100% versus 75%).<sup>3</sup> Some early studies found that CACs did not meaningfully impact bond prices, but more recent studies have found that, especially when the issuer is in

---

<sup>1</sup> For research on the pricing of legal covenants in debt instruments, see, e.g., Michael Bradley & Michael Roberts, *The Structure and Pricing of Corporate Debt Covenants*, 5 Q. J. FIN. 1550001 (2015); Umar Butt, *Debt Covenant Violation, Competition and the Cost of New Debt*, 44 AUSTRALIAN J. MGT. \_\_\_ (2018); Yini Liu & Ca Nguyen, *The Role of Financial Covenants in Pricing Private Investments in Public Equity*, 82 J. CORP. FIN. 102466 (2023); Adam B. Badawi, Vincent S.J. Buccola & Greg Nini & Vince Buccola, *Price Discipline for Non-Price Loan Terms*, 2024 Wharton School of Business Working Paper.

<sup>2</sup> E.g., Kay Chung & Michael G. Papaioannou, *Do Enhanced Collective Action Clauses Affect Sovereign Borrowing Costs?* 1 J. BANKING & FIN. ECON. 59 (2023); Mattia Picarelli, Aitor Erce & Xu Jiang, *The Benefits of Reducing Hold-Out Risk: Evidence From the Euro CAC Experiment, 2013-18*, 14 CAP. MKTS L. J. 155 (2019).

<sup>3</sup> E.g., Barry Eichengreen & Ashoka Mody, *Do Collective Action Clauses Raise Borrowing Costs?*, 114 ECON. J. 247 (2004). More generally on CACs, see Mark Weidemaier & Mitu Gulati, *A People's History of Collective Action Clauses*, 54 VA. J. INT'L L. 51 (2013).

financial distress, bond prices reflect differences in voting thresholds and mechanics.<sup>4</sup> Beyond CACs, a few studies have examined the pricing implications of other salient features of the bond contract, such as clauses designating the governing law (e.g., English versus New York versus the issuer’s domestic law) or the use of a trust indenture rather than a fiscal agency agreement.<sup>5</sup>

Pricing studies focused on CACs and other salient terms provide important insight. But the contract underlying a sovereign bond issuance is lengthy and detailed. Less salient provisions may have surprising legal implications. Disputes over the meaning of the *pari passu* clause, a seemingly innocuous boilerplate provision, led to years of litigation against Argentina and, ultimately, to a significant payout for a subset of creditors. At some key stages of the litigation, the precise language of the *pari passu* clause that Argentina used arguably mattered to the outcome. But discerning the pricing impact of small variations in *pari passu* language (which could arguably matter in litigation) has proved elusive.<sup>6</sup>

With some sovereigns, an investor’s legal rights may even depend on which of the sovereign’s many, seemingly identical bonds the investor holds. For instance, one bond may include a subtly different version of a clause than appears in other bonds issued by the same sovereign. If investors are unaware of the potential legal implications—whether of a clause that appears across a sovereign’s entire debt stock or in only a subset of bonds—these clauses operate as metaphorical “landmines,” generating litigation and affecting the dynamics of restructuring negotiations.<sup>7</sup> Creditors willing to play hardball can earn large payouts by identifying bonds with language that gives them leverage in a debt restructuring. This may require legal expertise, for it can be hard to tell whether a minor variation in contract language will benefit the creditor. But many investors in distressed debt specialize in just this activity.

---

<sup>4</sup> E.g., Elena Carletti, Paolo Colla, Mitu Gulati & Steven Ongena, *Pricing Contract Terms in a Crisis*, 11 CAP. MKTS. L. J. 540 (2016); Elena Carletti, Paolo Colla, Mitu Gulati & Steven Ongena, *The Price of Law: The Case of the Eurozone Collective Action Clauses* 34 REV. FIN. STUD. 5933 (2021); Torbjorn Becker, Anthony Richards & Yunyong Thaicharoen, *Bond Restructuring and Moral Hazard: Are Collective Action Clauses Costly?* 61 J. INT’L ECON. 127 (2003).

<sup>5</sup> E.g., Ugo Panizza & Mitu Gulati, *Make Wholes in Sovereign Bonds*, 16 CAP. MKTS L. 267 (2021); Marcos Chamon, Julian Schumacher & Christoph Trebesch, *Foreign Law Bonds: Can They Reduce Sovereign Borrowing Costs?* 114 J. INT’L ECON. 164 (2018); Michael Bradley, Elisabeth de Fontenay, Irving Arturo Salvatierra & Mitu Gulati, *Pricing Sovereign Debt: Foreign Versus Local Parameters*, 24 EUR. FIN. MGT. 261 (2018); Dilip Ratha, Supriyo De & Sergio Kurlat, *Does Governing Law Impact Bond Spreads?* 36 EMERGING MTS REV. 60 (2018); Stephen J. Choi & Mitu Gulati, *The Pricing of Non-Price Terms in Sovereign Bonds*, 1 J. L. FIN. & ACCT. 1 (2016).

<sup>6</sup> Paolo Colla, Anna Gelpern & Mitu Gulati, *The Puzzle of PDVSA Bond Prices*, 12 CAP. MKTS L. J. 66 (2017).

<sup>7</sup> See Robert E. Scott, Stephen J. Choi & Mitu Gulati, *Contractual Landmines*, 41 YALE J. REG. 307 (2024)

Despite this dynamic, we are not aware of studies meaningfully examining the pricing implications of potential contractual landmines. And this makes sense, given our focus on small variations in the language or particular clauses in bonds whose prices are impacted by myriad other factors. However, two inadvertent pricing experiments enable us to take a preliminary cut at answering the pricing question.

## **A Test of Pricing**

Ideally, one would examine the pricing question by comparing the issue price of two otherwise-identical bonds, one of which has the relevant contract language and the other does not. We do not have data to test this hypothesis for the universe of sovereign bonds, but we can make the following conjecture. Suppose the market price in fact fully incorporates all relevant information, including information about legally significant language in the sovereign's bond contracts. This should occur continuously from the date of issuance through default. As the risk of default rises, a contractual clause that increases investors' leverage in restructuring negotiations should increase in importance—if only because the probability that the clause will matter has increased—and the price should react accordingly. If the price is determined by arbitrageurs seeking to exploit a pro-creditor clause, it could potentially increase (although the risk that more aggressive investors will delay the restructuring might instead cause the price to decrease, but such negative pricing implication should affect all bonds that need to be restructured in the same way). If the clause increases the sovereign's bargaining power in a restructuring, the price will likely decrease, as the clause will reduce investor recoveries and may prompt litigation that will delay the restructuring. Ultimately, what matters is that, as the sovereign approaches default, the market price should continuously adjust to account for the growing importance of the clause. By contrast, if the market price does not fully incorporate information about the value of the clause, we would not expect this continuous price adjustment to occur.

To be sure, as a sovereign approaches default, many factors other than contract language will affect bond prices, including perceptions of default risk. So how can we discern whether market prices in fact incorporate information about the legal implications of clauses in the bond contract? It does not solve the problem to examine a cross section of bonds issued by different sovereigns, observing the price difference for bonds with and without particular clauses. Any

observed difference could result from a plethora of factors unrelated to contract language, including economic or political conditions that vary across sovereigns.

Instead, we compare prices across time for bonds with similar maturities, issued by the same sovereign but differing in one (but perhaps not the only) important respect. One bond has a clause that we view as a potential landmine (the treatment bond) and the other does not (the control bond). We focus on the price difference between the treatment and control bonds, but we are aware that these price differences could be due to the contractual language or to unrelated factors. To control for these factors unrelated to contractual language, we utilize a series of shock events to the market, described further below, that may have publicly announced to the market either that the treatment bond had a unique contractual clause or that the clause had increased in salience. We focus on a narrow window around these events and use the *relative* price change for the treatment and control bonds, described more fully below, to gauge the market's price response.

We posit that disclosure of the clause should not affect the relative price of the treatment bond if the market already has fully incorporated information about the value of the clause into bond prices. But if the market has not reacted, disclosure should move the price of the treatment bond relative to the control bond. This movement will be particularly acute for sovereigns nearing default. We further posit that, if market prices have not fully incorporated information about the clause as a sovereign nears default, this probably was also true at the time of the bond's issuance. If a clause has little or no impact on pricing at the time of issuance, one implication is that the players involved in the bond issuance, including attorneys catering to the short-term interests of their immediate clients, will have little reason to focus on the contract language. The relative price response of the treatment bond compared to the control bond on the date the landmine is made public thus provides a test of both (a) how the market views the contract clause *ex post*, when the sovereign is in distress and (b) whether prices fully incorporated information about the value of the clause prior to the public announcement, including at the time of issuance.

For example, assume a sovereign in deep crisis that is expected to restructure its debt. Imagine that one of its bonds includes contract language that a court might interpret to make the bond especially hard to restructure. Thus far, investors have paid no heed to this idiosyncratic feature of the bond. Now let us say that the Wall Street Journal or Financial Times publishes an article highlighting the potential importance of this overlooked contractual language. We should see a pricing impact on the bonds relative to similar bonds without the language. It is unfortunate,

from the perspective of those interested in studying the price implications of contract clauses, that experiments of this sort generally do not happen. Or so we thought.

## **Our Inadvertent Pricing Experiments**

Two of us co-host a podcast on sovereign debt and three of us sometimes write blog posts for the *Financial Times*. On occasion, our columns and podcasts discuss interesting contractual clauses that appear in sovereign bonds, including in bonds issued by countries experiencing financial distress. Typically, the clause that has captured our interest either appears in all of the sovereign's bonds or, if the clause appears in only a subset of bonds, captures our interest after the relevant bonds have stopped trading.

But in two cases, involving Sri Lanka and Ghana, the clause in question appeared in only one series of the sovereign's bonds. In both cases, the bond with the unusual clause happened to benefit from a guarantee.<sup>8</sup> This guaranteed bond interacted in unusual and, we thought, unexpected ways in the context of a restructuring that would include many of the sovereign's other "standard" bonds. In one case (Ghana), the unusual features of the guaranteed bond made it *more* vulnerable to a restructuring than many expected (or so we believed). In the other case (Sri Lanka), the guaranteed bond appeared to be *less* vulnerable to a restructuring than was commonly appreciated (again, so we believed). And both countries were in distress. We discussed both bonds in our podcast and also in posts for the *FT Alphaville* blog.

Although it did not occur to us at the time, we were inadvertently running two pricing experiments, publicly announcing the existence of unusual contract features that (we thought) contradicted widely-held views about the bonds. These announcements offered potential ways to assess the difference-in-difference between treatment and control bonds for Ghana and Sri Lanka. It was not obvious that these announcements would have an impact, even assuming we were right in believing that market participants had overlooked the unusual features of the bonds. For one thing, our podcast hardly commands a vast audience. A given episode might have just over 1,000 listeners, and while many are market actors, many others are students or academic colleagues. *FT*

---

<sup>8</sup> Bonds with sovereign guarantees tend to be unusual creatures, and sometimes more difficult to restructure than the sovereign's plain vanilla bonds. See Lee C. Buchheit & Mitu Gulati, *Restructuring a Sovereign Debtor's Contingent Liabilities*, in SOVEREIGN FINANCING AND INTERNATIONAL LAW: THE UNCTAD PRINCIPLES ON RESPONSIBLE SOVEREIGN LENDING AND BORROWING 287-294 (Carlos Esposito et al. eds., 2013).

*Alphaville* has more readers but is a blog, part of the Financial Times but not the main paper. Posts such as ours are unlikely to have more than 5,000 to 10,000 readers.<sup>9</sup> In addition to having a modest audience, we suspected that some market participants—especially at a few hedge funds—already knew about the unusual contractual features we had identified. However, even if this was true, the effect on bond prices was unclear. Investors with specialized knowledge have an incentive to hold that knowledge closely so as to allow time for transactions that exploit this knowledge.<sup>10</sup>

Readers interested in more detail about the unusual features of these bonds should consult the Appendix, which reproduces the key portions of our pieces from *FT Alphaville*.<sup>11</sup> Here, we provide only a brief description. The Ghanaian bond matured in 2030 and benefitted from a partial guarantee by the World Bank. In principle, the guarantee protected bondholders from a restructuring of Ghana’s debt relative to non-guaranteed bondholders. However, we explained how a clause in the bond contract allowed Ghana to give bondholders the same restructuring deal as holders of non-guaranteed bonds. We suspect this clause was copied from Ghana’s standard bonds without being modified to make it appropriate for a bond that benefitted from an official sector guarantee. Put simply, the clause was the CAC, and it let Ghana aggregate the guaranteed bond with other bonds, for purposes of holding a restructuring vote. If 75 percent of the group approved a restructuring proposal that offered every bondholder the same terms (which surely would not include a guarantee), all would be bound. Contrary to likely expectations, the effect would be to severely *undermine* the protection offered by the World Bank guarantee.

The second bond was issued by state-owned SriLankan Airlines and was guaranteed by Sri Lanka itself. The point of the guarantee was to improve the credit profile of the airline bond, effectively making it the same quality as a bond issued by the sovereign. Inexplicably—or so it seemed to us—the drafters of the SriLankan Airlines bond included an outdated CAC that was no longer standard for sovereign issuances at the time of issuance. The CAC did not allow the bond to be aggregated for voting purposes with other bonds. Instead, it would have to be restructured on its own, and the small size of the bond would make it easy for creditors to acquire a position large

---

<sup>9</sup> Our thanks to Alphaville’s editor, Robin Wigglesworth, for information on readership numbers.

<sup>10</sup> See Stephen J. Choi, Mitu Gulati & Robert E. Scott, *Anticipating Venezuela’s Debt Crisis: Hidden Holdouts and the Problem of Pricing Collective Action Clauses*, 100 BOSTON U. L. REV. 253 (2020).

<sup>11</sup> The full pieces are: Mark Weidemaier, Ugo Panizza & Mitu Gulati, *FUD and the Ghana 2030 Bond*, FT ALPHAVILLE (Nov. 29, 2022), <https://www.ft.com/content/fa3fddb7-72a7-475d-81f3-22bb68caaea6>; Mark Weidemaier, Ugo Panizza & Mitu Gulati, *About That Sri Lanka Airlines Guaranteed Bond*, FT ALPHAVILLE (June 9, 2022), <https://www.ft.com/content/024656d1-9f3b-4ba9-a38a-7d2f753cc84b>



enough to block a restructuring vote. Contrary to likely expectations, the airline bond was going to be harder to restructure than the sovereign’s own bonds (the opposite dynamic than is often assumed in the market), improving the credit profile of the airline bond *above* that of a bond issued by the sovereign.

## Data and Analysis

To conduct event studies around our podcasts and blog posts we use daily pricing data on the bonds in question sourced from Datastream, examining the relative price movements for the treatment bonds (with the guarantee and unusual contractual features) and control bonds (without these features) around the time of these events. To determine the relative price change, we develop a measure of the expected return for the treatment bond as a function of the control bond return using historical data on the relationship between the two returns. The expected return for the treatment bond allows us to control for price movements due to factors unrelated to the unusual contractual feature, which might affect the pricing of both the treatment and control bonds on the event date. For example, if the sovereign announces a windfall of tax revenue, presumably this will affect the returns for both the treatment and control bonds.

To determine the expected return, we start by regressing the daily return of the treatment bond over the daily return of the control bond issued by the same country and with similar maturity. Formally, we estimate the following model:

$$RT_t = \alpha + \beta RC_t + \varepsilon_t \quad (1)$$

$RT_t$  is the daily return of the treatment bond (defined as  $RT_t = \ln(PT_t/PT_{t-1})$ , where  $PT_t$  is the price of the treatment bond on day  $t$ ). Similarly, the daily return of the control bond is defined as  $RC_t = \ln(PC_t/PC_{t-1})$ . In our baseline model, we estimate Equation 1 over a window of 70 trading days (about 3 months), but we also experiment with a 90 trading days window. We always close the estimation window 4 trading days before the beginning of the event window (the event window is a five trading days window centered on the date of the event, i.e., the release of the podcast or the publication of the FT Blog). We use the estimated regression model to compute an expected daily return (ER) for the treatment bond based on the actual daily return for the control bond on any specific date. The ER is defined as:

$$ER_t = (\hat{\alpha} + \hat{\beta} RC_t) \quad (2)$$

We then compute an abnormal return (AR) equal to the actual daily return (RT) for the treatment bond minus the expected daily return (ER). The abnormal return gives us a measure of the relative price change due to the public disclosure of the contractual features of the treatment bond. Abnormal return (AR) is defined as:

$$AR_t = RT_t - ER_t \quad (3)$$

Finally, we compute average cumulated abnormal returns by computing the average abnormal return over the event window:

$$CAAR = \frac{1}{5} \sum_{t=E-2}^{E+2} AR_t \quad (4)$$

To test whether average abnormal returns are significantly different from zero, we use the standard deviation of abnormal returns computed over the estimation window ( $\sigma_{AR}$ ) to build the statistics:

$$t = \frac{\sum_{t=E-2}^{E+2} AR_t}{\sigma_{AR}\sqrt{5}} \quad (4)$$

### *Ghana*

In the case of Ghana, the treatment bond is Ghana's U.S. dollar denominated World Bank guaranteed bond series issued with a 10.75% coupon on October 14, 2005 and maturing in October 2030 (CUSIP QJ1388462). As a control we use two "plain vanilla" Ghana U.S. dollar denominated sovereign bonds. The first had a 8.625% coupon, was issued on April 7, 2021 and matured in April 2024 (CUSIP BO8436244). The second had a 8.125% coupon, was issued on March 26, 2023 and matured in March 2032 (37443GAB7).

Next, we consider the two events. The first occurs on November 29, 2022, when our *FT Alphaville* piece is published.<sup>12</sup> That piece reports our view that the Ghana 2030 bond, which

---

<sup>12</sup> See Mark Weidemaier, Ugo Panizza & Mitu Gulati, [FUD and the Ghana 2030 Bond](#), FT Alphaville (Nov. 9, 2022).

because of the guarantee was trading at a premium to the plain vanilla Ghana bonds, was less protected than some might think because of the idiosyncratic feature discussed above. The second event, roughly two weeks later, on December 12, 2022, was the release of a *Clauses and Controversies* podcast on the same topic.<sup>13</sup>

Each event announced that the Ghana 2030 bond was more vulnerable to restructuring than might have been assumed. If the market price did not already incorporate this information, we would expect the price to drop after one or both announcements—at least if market participants were aware of the announcements and thought it relevant. Less obviously, the value of Ghana’s other plain vanilla bonds should go up, since a restructuring of the guaranteed bond would make more funds available to pay the others. This second effect should be smaller, since there is only one guaranteed bond and numerous plain vanilla bonds. More importantly, for our purposes, is that the event should have affected the relative return of the two bonds and reduced their price difference.

Figure 1 plots the price of the treatment bond (the solid black line) together with the prices of the control bonds (the grey dashed lines) and their average (the solid grey line) around the two events (marked by the vertical lines).

### **Figure 1: Ghana Bonds**

This figure plots the price of the Ghana 2030 guaranteed bond (the solid black line, left axis) together with the price of two comparator bonds (the dashed grey lines, right axis) and their average. The vertical lines mark the November 29, 2022 and December 12, 2022 events.

---

<sup>13</sup> See *Ghana’s 2030 FUD Bond*, *Clauses and Controversies* Ep. 92 (Dec. 12, 2022), <https://podcasts.apple.com/us/podcast/ep-92-ft-mitu-mark/id1528208049?i=1000589924678>




Figure 1 shows that the prices of the treatment and control bonds converged after the first event. That is, they moved in the predicted directions, consistent with the idea that the market price had not fully incorporated information about the unusual features of the Ghana 2030 guaranteed bond. That bond went down in value and the ordinary Ghanaian bonds went up in value (however, they were already trending in this direction). After the second event, the release of the podcast, both bonds increased in value, but the price of the control bonds increased more than that of the treatment bonds.

Table 1 reports the results of the more formal event study analysis, which corroborates what we saw in Figure 1.

**Table 1: Ghana event study**

This table reports the results of the event studies for the Ghana guaranteed bond. The top panel focuses on the event of November 29, 2022 and the bottom panel on the event of December 12, 2022. The first column reports results for our baseline estimates (70 trading days estimation window and 5 trading days event window), column two uses a 3 trading days event window and column three a 90 trading days estimation window. In the baseline, the estimation window ends 5 trading days before November 21 for both events. The last column of Table 1 shows that the results for the December 12, 2022 event are robust to ending the estimation window on December 2<sup>nd</sup>.

	Baseline	3-trading days event window	90-trading days estimation window	Estimation window ends on December 2 <sup>nd</sup>
November 29, 2022 event				
AAR	-0.012*** (4.98)	-0.016*** (5.28)	-0.011*** (4.47)	
December 12, 2022 event				
AAR	-0.0042* (1.79)	-0.0036 (1.15)	-0.0037 (1.48)	-0022 (0.79)


\*\*\* statistically significant at the 1% confidence level, \*\* statistically significant at the 5% confidence level, \* statistically significant at the 10% confidence level.

We find a large (between 1.1% and 1.6% per day) and statistically significant effect of the first event (the Alphaville blog post). There is movement after the second event as well. But it is only statistically significant in the baseline model and the point estimates are much smaller (about one-third of what we obtained in the model that studies the first event). This is perhaps because the second event did not produce any additional news with respect to the first event. Or it could be that the podcast was too obscure.

As a final matter, we note that Figure 1 suggests that there might have been convergence before the event. To test for this possibility, we do a placebo test by estimating the model with a series of 3-trading days event window centered between 5 days before and 5 days after the actual event. Figure 2 shows that average abnormal returns are only significant when the 3-trading days event window is centered on the day of the event and the day before the event (which, by construction, includes the event day in event window).

## Figure 2: Ghana Bond, sensitivity analysis

This figure plots Average abnormal returns with 99% confidence intervals computed for three days windows centered between five days before and 5 days after the event.



Our first experiment is consistent with the market price's failure to incorporate information about the unique contractual features of the Ghana 2030 bond prior to public disclosure. After both events, there is price movement for the bond in the predicted direction (down). The first disclosure produces the bulk of the movement though.

### *Sri Lanka*

The Ghana 2030 bond was more vulnerable to restructuring than it seemed on its face. The SriLankan Airline bond involved the opposite dynamic. As noted, the airline bond had been guaranteed by Sri Lanka and had an unusual clause that made it difficult to restructure. In a podcast released April 25, 2022, we explained that this unusual clause made it more likely that the airline bond would escape a restructuring, making the holders of the bond much better off than other Sri Lankan bondholders.<sup>14</sup> On June 9, 2022, we reiterated this argument in an *Alphaville* post.<sup>15</sup> In

<sup>14</sup> See *Sri Lanka, SriLankan Airlines, and...*, *Clauses and Controversies* Ep. 71 (Apr. 25, 2022), <https://podcasts.apple.com/us/podcast/ep-71-ft-mitu-mark/id1528208049?i=1000558674659>


<sup>15</sup> See Mark Weidemaier, Ugo Panizza & Mitu Gulati, *About That Sri Lankan Airlines Guaranteed Bond . . .*, *FT Alphaville* (Jun 9, 2022), <https://www.ft.com/content/024656d1-9f3b-4ba9-a38a-7d2f753cc84b>.

other words, the order of the blog and the podcast were flipped from the Ghana case. Here, the podcast was first.

Figure 3 shows what happened. It plots the price of the SriLankan Airlines bond (the treatment bond), which is a U.S. dollar bond issued on June 25, 2019 with a 6.85% coupon and maturing on June 2024 (CUSIP AZ2259598). Our comparator bonds are three plain vanilla U.S. dollar bonds: (i) a Sri Lankan sovereign bond issued March 14, 2019, with a coupon of 6.85% and maturing in March 2024 (CUSIP AX5675811); (ii) a Sri Lankan sovereign bond issued in June 2019 with a coupon of 6.35% and maturing in June 2024 (CUSIP AZ2977041); and (iii) a sovereign bond issued in March 2015 with a coupon of 6.125% and maturing in March 2025 (CUSIP 85227AN6).

**Figure 3: Sri Lankan Airline and Government Bonds**

This figure plots the price of the SriLankan Airline Bond (the solid black line) together with the price of three comparator bonds (the dashed grey lines) and their average.



The black line in Figure 3 shows the treatment bond, the dashed grey lines the show the control bonds, and the solid grey line shows the average price of the control bonds. Here, we see divergences in prices at the time of both the podcast and the Alphaville blog post. Both the podcast and the blog post discussed how the Sri Lankan airline bond was less vulnerable to a restructuring

than might have been expected. If this information about the idiosyncratic features of the airline bond was new to the market, the price of the treatment bond should rise upon release of the information. Conversely, the price of the control bonds should go down, since paying the airline bond in full would mean less money available to pay other creditors. Again, the second effect should be smaller, since there are many sovereign bonds but only one guaranteed bond and the value of the guaranteed bond is relatively small compared to that of Sri Lanka’s external debt. This is what we observe. After the podcast, the airline bond increased in price and the control bonds go down. After the subsequent blog post, the airline bond again increases in price while the prices of the control bonds decrease.

Table 2 reports the formal analysis, and shows a large, positive, and statistically significant effect for the first event (top panel).

**Table 2: Sri Lankan Airline event study**

This table reports the results of the event studies for the Sri Lankan Airline guaranteed bond. The top panel focuses on the event of April 25, 2022 (the podcast) and the bottom panel on the event of June 9, 2022 (the blog post). The first column reports results for our baseline estimates (70 trading days estimation window and 5 trading days event window), column two uses a 3 trading days event window and column three a 90 days estimation window. In the baseline, the estimation window ends 5 trading days before April 25 for both events. The last column of Table 1 shows that the results for the June 9, 2022 event are robust to ending the estimation 5 trading days before June 9.

	Baseline	3 trading-days event window	90-trading days estimation window	Estimation window closes on June 3
April 25, 2022 event				
AAR	0.084* (1.85)	0.097*** (2.74)	0.084** (2.36)	
June 9, 2022 event				
AAR	0.045 (0.99)	0.012 (0.33)	0.044 (1.02)	0.027 (0.50)

\*\*\* statistically significant at the 1% confidence level, \*\* statistically significant at the 5% confidence level, \* statistically significant at the 10% confidence level.

The effect of the second event is also in the predicted directions, but not statistically significant and with point estimates that indicate a much smaller magnitude (bottom panel).



The two sets of experiments tell the same story. Disclosure of idiosyncratic contract clauses produces a significant relative price movement between the treatment and control bonds in the expected direction, assuming that the market price has not fully incorporated information about the impact of these clauses. And, if, prior to our public announcements, the market has not fully priced in the effect of the contract term as the sovereign nears default, it is also likely that the market did not fully price the contract term at the time of issuance.

## **Conclusions and Implications**

Theories of efficient markets and efficient contracts, taken to their extreme, would tell us the landmines in the fine print of contracts should not exist. In the case of efficient contracting theory, disciplinary pressures on high priced lawyers in a competitive legal market should ensure few instances of landmines of the magnitude that we have described. And even if there are drafting anomalies, the markets should fix them via the pricing mechanism. In either event, contractual landmines should not exist. Our inadvertent experiments suggest that they can exist. What we cannot show with but two experiment is how prevalent these are.

A puzzle with the results of the experiments is that at least some people in the market knew about these features of the bonds in question. They were not being discussed in any analyst reports or press articles. However, based on conversations with some of those in the market, we are confident that at least some had enough information for there to be trading. Yet, until our disclosures, there was no discernable impact on price.

A couple of possibilities occur to us. First, the smart money did trade, but it was able to conceal the information embedded in its trades. Second, the smart money was reluctant to trade until they knew that others were likely to recognize the value of what they were trading. In this case, our public discussion might have served as a kind of coordinating function. Again, we need more experiments to explore these dynamics further. But it strikes us as an exciting endeavor.

## Appendix

### The Alphaville Pieces

“[T]he Ghana 2030 Bond”<sup>16</sup>

Ghana is on the brink of default after a borrowing spree over the past decade, including a spate of recent kick-the-can issuances. It owes about \$50bn in total. Of this, roughly \$14bn is denominated in foreign currency, including a \$1bn bond due in 2030.

...

Issued in 2015, the Ghana 2030 bond has the latest next-generation aggregated ‘Collective Action Clause’, which have been gradually introduced since 2014 to ensure orderly debt restructuring. What distinguished these new provisions was an “aggregation” mechanism that was absent from prior generations of government bonds. Simplifying a little, the aggregated CACs let countries restructure multiple bond series in a single vote across all series as long as 75 per cent of creditors vote yes and everyone is offered the same deal (“uniformly applicable” treatment).

[ ]The fancy new CAC in the Ghana 2030 bond permits aggregation with other series of bonds that constitute “Debt Securities Capable of Aggregation.” That term is defined to mean: . . .

those debt securities which include or incorporate by reference this Condition 13 and Condition 14 (Aggregation Agent; Aggregation Procedures) or provisions substantially in these terms which provide for the debt securities which include such provisions to be capable of being aggregated for voting purposes with other series of debt securities.

---

<sup>16</sup> See Mark Weidemaier, Ugo Panizza & Mitu Gulati, *FUD and the Ghana 2030 Bond*, FT Alphaville (Nov. 9, 2022). See also . . .

In short, the Ghana 2030 can be pooled for voting purposes with other foreign currency bonds that have the same aggregated CAC, or “provisions substantially in these terms.” Then, if 75 per cent of the voting pool are in favour of a restructuring, the 2030 holders will be forced to take the same deal. And, because all bondholders must be offered the same deal, any restructuring proposal will almost certainly not benefit from a guarantee. Bye bye World Bank guarantee!

...

In principle, the bond documents do not have to be interpreted to allow this result. Ghana’s other foreign-law bonds with aggregated CACs — the ones without guarantees — do not have exactly the same aggregated CAC as the 2030 bond. The reason is that CACs list a dozen or so key matters (such as payment terms) that can be modified only by supermajority vote. These are called “Reserved Matters.” The Ghana 2030 has one more Reserved Matter than the others: the guarantee. But remember, the 2030 bond can be aggregated with any other bond that has an aggregated CAC that is “substantially in these terms” (i.e., the terms found in the Ghana 2030). And it is hardly obvious that this difference makes the other bonds not “substantially” the same as the 2030.

It might seem unfair to give holders of the Ghana 2030 bond the same restructuring deal as holders of non-guaranteed bonds. But the aggregated CAC does not require the bond issuer to respect differences in the economic value of different bonds. In fact, the “uniformly applicable” requirement seems to require the contrary. To simplify, it requires the issuer to offer each bondholder the same terms (or the same new instrument). The logic underlying this requirement was that investors holding different securities would find that their interests converged in a crisis. Except, well, they don’t.

Could all this happen? Ukraine recently concluded a ‘debt reprofiling’, with foreign creditors agreeing to a two-year freeze on payments. It has multiple bond series outstanding, some of which had aggregation provisions like those in Ghana’s post-2014 eurobonds. Ukraine’s reprofiling encompassed regular sovereign bonds and two corporate bonds that had a state guarantee. Ukraine did not hesitate to aggregate these bonds together, despite the fact that the guaranteed bonds were

arguably superior credits. We are willing to wager that the argument for aggregating guaranteed corporate bonds with regular sovereign bonds was that all counted as “Debt Securities Capable of Aggregation.”

---

“About That Sri Lankan Airline Guaranteed Bond . . .”<sup>17</sup>

Guaranteed bonds are a puzzle. One might assume they would carry a lower interest rate than the government’s own bonds. That’s because they are backed by both the issuing entity and the government itself. Even when the issuer is poorly-managed and hemorrhaging money, its guaranteed bonds would seem no worse than the sovereign’s own. Yet guaranteed bonds often carry a higher interest rate, possibly because they are less liquid than standard government bonds.

...

We know that deal documents often begin with a template used in a prior deal. But this usually means the deal from the prior week, not the prior century. Although the guarantee can be written so that it can be restructured along with the sovereign’s other obligations, this rarely happens. On the contrary, guarantee contracts often seem designed with little thought about what will happen if the government gets into distress and needs to reduce its obligations. So holders of guaranteed bonds sometimes escape the restructuring completely.

Is there a poster child for a badly-designed government guarantee slapped on bonds issued by a money-incinerating company? Before Sri Lanka, we would have pointed to the Greek railroad bonds issued before the Greek debt crisis from a decade ago.

Those bonds were such a hassle to restructure that the Greek government did little more than politely ask bondholders to take haircuts along with the government’s other creditors. Unsurprisingly, the answer was generally “no thanks, we’d rather get paid in full.” Holders of bonds issued by the loss-generating Greek

---

<sup>17</sup> See Mark Weidemaier, Ugo Panizza & Mitu Gulati, *About That Sri Lankan Airlines Guaranteed Bond . . .*, FT Alphaville (Jun 9, 2022).

railway not only got their regular interest payments, they got paid in full. Sovereign Greek bondholders . . . well, not so much.

But Sri Lankan Airlines may now take the cake. Not so long ago, the company was profitable and (as one of us recalls) pleasant to fly. But the ruling family decided that it could run the airline better on its own. Now, the airline loses about \$100mn a year; and it has incurred a net loss every year since 2008. Reports indicate the government wants to privatize it, but there are few takers. The government not only ran the airline into the ground, it supported its borrowing on the international debt markets with government guarantees.

And that brings us to the crux of the story. The SriLankan Airlines guaranteed bond is a veritable nightmare for a debt restructuring.

...

The modification term is the most important provision in a distressed sovereign bond. This term is often referred to as a Collective Action Clause, or CAC, because it enables collective decision making by dispersed creditors.

The CAC sets out the requirements for holding a vote to reduce the sovereign's payment obligations. What fraction of creditors must vote to approve the restructuring plan before it binds the entire group? If they hold bonds, can the government and affiliated entities cast votes? If the sovereign has multiple series of bonds outstanding, is the vote aggregated across the different series, or is it taken series by series? The CAC answers these and other questions.

That last question, the question of aggregation, is especially important for a country like Sri Lanka, which has a large debt stock represented by multiple series of bonds. If it must hold restructuring votes series by series — for example, persuading 75 per cent of each bond to vote in favour — then it becomes vulnerable to holdouts. This is especially true for smaller bonds, where a blocking position is easy to obtain.

When the vote is aggregated across series, it becomes significantly harder to buy a blocking position. The Greek restructuring of 2012 highlighted the importance of aggregated voting. Greece attempted series-by-series restructurings in thirty-some series of bonds, but holdouts blocked the restructuring in about half

of the bonds. These bonds were paid in full, although Greece's other bondholders took brutal haircuts. As a result, beginning in 2014, sovereign bonds began to include aggregation features.

For whatever reason, Sri Lanka came late to this party, but finally began to include aggregation features in its sovereign bonds starting in 2018. The guaranteed airline bond was issued in 2019 yet, inexplicably, it does not include the aggregation feature. Instead, it uses an antiquated CAC template that we've only seen a handful of times in all of our research on sovereign bonds.

Given the relatively small size of the airline bond —\$175mn, as compared to sovereign bonds that often exceed \$1bn — the absence of aggregation features is especially striking.

Typically, bonds without aggregation features can still be restructured in a fashion that deters holdouts. This involves a technique called the Exit Consent. Here, the sovereign takes advantage of the fact that, although the CAC often requires a creditor supermajority to modify payment terms, it can modify less important terms with the support of only a bare majority. But these less important terms can still be quite significant.

The idea is that, when a majority of creditors support the restructuring, they can threaten to leave prospective holdouts with bonds that have had key features removed or amended (e.g., a cross default clause). Holdouts do not always fall for the threat, but Exit Consents can be an effective restructuring tool.

Bizarrely, the Sri Lankan airline bond seems to require the same vote to change both important terms and less important terms (75 per cent in both cases). .

..

The more different this bond is from all the others, the harder it is to wrap it into the same restructuring. And because the government seems to have procrastinated in hiring restructuring advisers, there will be significant time pressure to design the restructuring. The airline bond may get paid in full simply for lack of time to design a clever method to restructure this odd little bond.

On balance, however, the goofy aspects of the airline bond . . . favour creditors. If some set of hedgies decide to target this bond and acquire a 25 per cent

stake (remember, it is just \$175mn and trading at less than half its face value), it is going to be awfully hard to cajole them into taking a haircut.

Time will tell whether Sri Lanka's restructurers figure out a special strategy to deal with the guarantee or whether it pays out like those Greek guarantee bonds — in full and on time.