

THE PRICING OF SOVEREIGN RISK UNDER COSTLY INFORMATION*

January 11th, 2017

Grace Weishi Gu[†] and Zachary R. Stangebye[‡]

The consequences of investors' endogenous costly information acquisition behavior for the pricing of sovereign risk are explored in a quantitative sovereign default model. We develop an approach to identify information costs empirically using sovereign credit rating publication data. The calibrated model delivers a number of interesting results: First, it endogenously generates levels of time-varying volatility in the country risk spread that approach the data; second, it suggests that risk premia fluctuate with information acquisition behavior and thus with country-specific states, which has econometric consequences; and third, it highlights a non-trivial welfare trade-off in the promotion of sovereign transparency.

Keywords: costly information; sovereign default; time-varying volatility; risk premia; transparency

JEL code: F34, D83, G12

*The authors would like to thank Manuel Amador, Mark Aguiar, Sebnem Kalemli-Ozcan, Boragan Aruoba, Luminita Stevens, Felipe Saffie, John Shea, Juan Carlos Conesa, Marina Azzimonti, Ting Liu, Raoul Minetti, Charles Engel, Kamran Bilir, Kenneth West, Leonardo Martinez, Sandra Lizarazo, Ruediger Bachmann, Olivier Jeanne, Michal Szkup, Mark Wright, Marco Bassetto, Linda Tesar, Pablo Ottonello, Gita Gopinath and seminar participants at the University of Maryland, Michigan State University, Stony Brook University, the University of Pittsburgh, the University of Wisconsin, the International Monetary Fund, the Federal Reserve Bank of Chicago, the University of Michigan, and the NBER Summer Institute for useful comments and suggestions. All errors and mistakes are our own.

[†]E-mail: grace.gu@ucsc.edu, Phone: (831) 459-4791, Address: 463 Engineering 2, 1156 High Street, Santa Cruz, CA 95064

[‡]E-mail: zstangeb@nd.edu, Phone: (574) 631-1054, Fax: (574) 631-1054, Address: 3015 Nanovic Hall, Notre Dame, IN 46556 (Corresponding Author)

Research highlights:

1. Costly information acquisition is introduced into a sovereign default model.
2. Lenders' information acquisition and the sovereign's defaults are jointly endogenous.
3. Information acquisition produces significant time-variation in spread volatility.
4. Information acquisition shifts the composition of spreads during crises.
5. Sovereign welfare as a function of information costs features a U-shape.

1. Introduction

Yields in sovereign bond markets in emerging economies largely reflect the risk that a domestic borrower may default on foreign creditors since it does not in principle have any reason to care about the well-being of these creditors. But a sovereign borrower's lack of welfare concern for its foreign lenders is not the only relevant friction that arises from the international nature of these markets: Information frictions also play a large role in cross-border financial transactions: Investors are likely to be less informed about payoff-relevant shocks in other developing countries ([Hatchondo \[2004\]](#), [Van Nieuwerburgh and Veldkamp \[2009\]](#), or [Bacchetta and van Wincoop \[2010\]](#)).

This is true even for institutional investors, as few developing countries fully adopt International Public Sector Accounting Standards (IPSAS) or other international standards for government finance reporting. Further, their sovereign bonds are usually traded over-the-counter with limited information and thin trading activities relative to assets traded in exchanges, such as stocks. Nevertheless, information about emerging markets' payoff-relevant shocks can often be acquired at a cost.

In this paper we seek to understand how costly information acquisition affects the equilibrium pricing of sovereign default risk. In particular, we construct a model in which the sovereign's default and borrowing decisions as well as the lenders' acquisition of payoff-relevant information are jointly endogenous: Lenders always observe some public states, such as output growth and debt levels, but cannot directly observe other potentially payoff-relevant states when making investment decisions, such as the severity of a recession implied by a potential future default or populist sentiment in that country. Information regarding these sorts of shocks can only be acquired at a cost.

In any market-based model of costly information acquisition, one encounters the dilemma highlighted by [Grossman \(1976\)](#), which is that market prices tend to reveal too much information and thus kill incentives to acquire information in the first place. We circumnavigate this issue by separating the information acquisition decision from market participation, much

like [Angeletos and Werning \(2006\)](#). An independent contractor/forecaster interested only in the integrity of its forecast such as the credit rating agencies in [Holden et al. \(2018\)](#) or [Manso \(2013\)](#) conducts the information acquisition. As market participants, lenders pay a fixed small fee to gain access to the forecasts provided by the forecaster, the usefulness of which will depend on the acquisition efforts of the forecaster in that state. Real world analogues of the forecaster might be financial software and analytics firms, such as Bloomberg or Reuters; alternatively they could be public, supernational entities such as the IMF or the World Bank, or credit rating agencies such as Moody's or S&P. Our calibration strategy reflects this market framework.

Costly information acquisition generates a dependence of lender knowledge regarding unobserved states on those that are publicly observed. For instance, when debt levels are low and output growth is high, there is little to gain in terms of accurately inferring unobserved shocks since default risk is negligible for most of their realizations. However, for moderately high debt levels and low growth, information is more valuable since unobserved shocks may substantially affect default risk.

Intuitively, the forecaster will start acquiring more information during crises, more carefully studying the borrower and its associated default risk. This will be reflected in either more or more precise information being transmitted to lenders. To quantify the impact of this mechanism, in the benchmark model we select an empirical counterpart for the forecaster, namely credit rating agencies. We then employ total number of additional rate *publishings* in a quarter as a proxy for information acquisition activity.¹ For instance, Moody's standard procedure is to publish a rating twice a year, but occasionally they will publish more if they deem it necessary. It is precisely when they deem it necessary that we interpret them to be undertaking greater information acquisition efforts.

To calibrate information costs, we target the fraction of quarters in which intense attention is paid to the borrower country. We use sovereign credit rating publishing as an empirical

¹Publishings include reports in which the country's rating did not change. This is an important distinction, since greater information acquisition occasionally reaffirms prior beliefs, but is costly nevertheless.

proxy for attention in the data, and intense attention is defined as the number of rating publications during a quarter being at the peak volume of the entire sample period. If information is infinitely costly this intense-attention fraction will be zero; as it becomes free, it goes to one. The attention fraction uniformly increases as information costs decrease; thus, the latter cleanly identifies the former. We could perform our quantitative exercise for any country, but we calibrate our benchmark model to Ukraine from 2004-2014 for two reasons: First, in contrast to many Latin American economies, it experienced substantial macroeconomic volatility during the sample period;² second, it does not suffer from the added complications of Eurozone or E.U. membership.

The model generates three new sets of results. First, it serves as a source for time-varying volatility in the country risk spread and can generate this feature endogenously without assuming that any fundamental processes exhibit it. During normal times, lenders receive little to no accurate information regarding payoff-relevant unobserved shocks. Consequently, they consider these shocks to be at their mean for the purpose of inferring and thus pricing default risk. This implies that bond yields do not respond to realizations of these unobserved shocks during normal times and so spread volatility is lower. As crises approach, however, lenders acquire additional information about these shocks, which then become priced. This implies that bond yields *do respond* to realizations of unobserved shocks during crisis times, which increases spread volatility.

This time-variation in the macroeconomic volatility is a well-documented empirical fact (Justiniano and Primiceri [2008] or Bloom [2009]) but little has been done as of yet to understand its causes,³ despite the fact that Fernández-Villaverde et al. (2011) show that second-moment fluctuations in the risk-spread can have substantial first-moment effects on investment and output.

To quantify the effect, we develop a model-free metric of time-varying volatility, which

²Latin American economies experience substantial volatility in the preceding decades, but the Moody's data begins in 1999.

³Some notable recent exceptions are Seoane (2015) and Johri et al. (2015), but these papers explain time-varying volatility in the country risk spread by assuming exogenous time-varying volatility in fundamentals. Sedlacek (2016) provides an alternative solution, but it generates time-varying cross-sectional dispersion among firms. Chari and Kehoe (2003) also show that investor herding behavior can increase uncertainty during crises.

we call the *Crisis Volatility Ratio* (CVR). We find that in the data the CVR is 3.67. Our calibrated benchmark model, without targeting this metric, can explain roughly 79% of it⁴ by generating a CVR of 2.89. And it does so without assuming any time-variation in the volatility of any underlying shocks. This is a substantial improvement over, for example a standard sovereign default model in the absence of costly information, which generates a CVR of 1.27.

Our second set of results is that the state-contingent nature of information acquisition translates into bond risk premia. In standard sovereign default models, the spread on a short-term sovereign bond can be decomposed into two components: Default risk and a risk premium for that default risk. The latter is affected not only by investors' degree of risk aversion, but also by the conditional distribution of payoff-relevant shocks. From the investors' point of view, information acquisition reduces the variance of this conditional distribution. Consequently, it exerts downward pressure on the risk premium. This novel channel alters the relative contributions of default risk and the risk premia to the overall spread across publicly observed states.

The calibrated model suggests that during crisis times default risk constitutes about 33.2% of crisis-level spreads,⁵ whereas a standard model without costly information acquisition would put this figure substantially nearly 20 percentage points lower, at about 13.6%. During non-crisis times, the average difference across these models is a mere 0.84 percentage points, so the effect is strongly state-contingent. This has substantial implications for the inference of default risk from spread data, which is a common practice in the literature ([Bi and Traum \[2012\]](#), [Bocola \[2016\]](#), [Bocola and DAVIS \[2016\]](#), and [Stangebye \[2015\]](#)). In particular, it implies that a standard sovereign default model may *understate* default risk during crises by assuming a risk premium that is too large.⁶

Third, we find a non-monotonicity in sovereign welfare as a function of information costs.

⁴Note our model suggests there is still room for other sources of time-varying volatility, such as uncertainty shocks.

⁵We define a 'crisis' directly from spreads following [Aguiar et al. \(2016a\)](#).

⁶It is worth noting that while this channel offsets the surge in risk premia associated with crises, it does *not* reverse it. Risk premia in our model remain countercyclical.

The intuition is simple and operates through debt prices. If there is no transparency, i.e., large information costs, then lenders demand a greater information-driven risk premium, making it more expensive to borrow and service debt. This risk premium falls as transparency increases, which is consistent with findings in the empirical literature (Kopits and Craig [1998], Poterba and Rueben [1999], Bernoth and Wolff [2008], and Iara and Wolff [2014]). This is tantamount to a reduction in borrowing costs for the sovereign, which improves its welfare.

But there is a countervailing force: At high debt levels, the sovereign will be fully exposed to the volatility that results from the pricing of normally ignored unobserved shocks, which reduces welfare. The calibrated model suggests that sovereign welfare as a function of information costs features a U-shape. Initially, increasing transparency hurts the sovereign since price volatility rises. Eventually, though, the welfare benefits of reduced borrowing costs begin to dominate and welfare attains its maximum level in the full-information model. This may explain why some countries, such as Argentina and Greece, tend to resist incremental changes in transparency policy despite the apparent benefits of implementation.

It is important to note that we are focusing on information frictions that are *country-specific*, i.e., those that arise between a single country and its lenders. This has substantial implications for our results' application. For instance, one cannot account for all the time-varying volatility or the state-contingent composition of the risk-spread by controlling for global metrics such as the CBOE VIX or the P/E ratio, as has been done in the literature (Bocola and Dovis [2016] and Aguiar et al. [2016a]). With regard to the time-varying volatility, our finding corroborates the careful empirical work of Fernández-Villaverde et al. (2011), who find that the bulk of the time-varying interest rate volatility in emerging markets is country-specific rather than global.

Our focus on relations between a single borrower and its lenders over time and the implications for the ergodic price distribution distinguishes our analysis from the related work of Cole et al. (2016). These authors also explore a model of costly information acquisition in sovereign debt markets, but their focus is static. They highlight the potential for this channel

to cause contagion effects across many countries and generate multiplicity. [Angeletos and Werning \(2006\)](#) and [Carlson and Hale \(2006\)](#) also explore how market-based information acquisition or rating agencies affect equilibrium multiplicity or uniqueness in variations on the canonical model of [Morris and Shin \(1998\)](#). [Durdu et al. \(2013\)](#) explore the impact of news shocks in a similar model. Recently, [Bassetto and Galli \(2017\)](#) also study the role of information on bond pricing in a two-period Bayesian trading game, focusing on implications for inflation risk. [Cole and Kehoe \(1998\)](#), [Sandleris \(2008\)](#), [Catao et al. \(2009\)](#), and [Pouzo and Presno \(2015\)](#) have all shown that information asymmetries are key to explaining various features of these markets, though none have considered the consequences of allowing information to be gathered at a cost.

The remainder of this paper is divided as follows: Section 2 describes the model; Section 3 discusses the data and quantitative implementation of the model, as well as counterfactual analyses and the model’s novel implications; and Section 4 concludes.

2. Model

We consider a small open economy model of endogenous sovereign default in the vein of [Eaton and Gersovitz \(1981\)](#). This is in part for tractability, but also to demonstrate our model’s applicability and compare results to the recent, expanding quantitative literature e.g. [Aguiar and Gopinath \(2006\)](#), [Arellano \(2008\)](#), [Hatchondo and Martinez \(2009\)](#), or [Mendoza and Yue \(2012\)](#).⁷ There is a sovereign borrower who issues one-period non-state-contingent debt to a unit mass of foreign lenders. This borrower lacks the ability to commit to repay this debt in subsequent periods and will default if it is optimal to do so ex post.

For clarity, we distinguish a random variable from its realization by placing a tilde over the former.

⁷For simplicity of exposition, as the above papers we focus on Markov Perfect Equilibria that can be expressed recursively, though the set of equilibria is potentially much larger ([Passadore and Xandri \[2015\]](#)).

2.1. Shocks

There are two shocks in this model, both to the sovereign. The first is a growth shock to output that the sovereign receives each period. More specifically, the endowment Y_t can be expressed in terms of a sequence of growth rates g_s as follows:

$$Y_t = Y_0 \times \prod_{s=1}^t e^{g_s}$$

where Y_0 is given. We assume g_t follows an AR(1) process:

$$g_t = (1 - \rho)\mu_g + \rho g_{t-1} + \sigma_g \epsilon_t$$

where μ_g is average growth rate, ϵ_t is a standard normal, and σ_g is the standard deviation of the growth innovation. The endowment and its growth processes are publicly observed by everyone. We collect the publicly observed exogenous states into a vector $s_t = \{Y_t, g_t\}$. This vector follows a Markov process with transition density $f_t(s_{t+1}|s_t)$.

The second shock in the model is an iid default cost shock, m_t , which applies only in the first period of a default.⁸ It is the only shock that is unobserved by foreign lenders and the forecaster when they make investment and attention decisions (we will detail the timing below); information regarding it can only be acquired at a cost. As a default cost shock, it could represent the magnitude of private capital outflows, the impact of likely fiscal consolidation, or the severity of the ensuing international sanctions, among other default-induced costs about which lenders may not be perfectly informed.

For the benchmark, we assume that m_t is normally distributed around zero with an unconditional standard deviation σ_m . In [Appendix A](#), we also show that when m_t is normally distributed it can alternatively be interpreted as a default preference shock, such as populist sentiment or political change. In this sense, m_t can stand in for a wide variety of payoff-relevant factors that are not immediately observed by the lenders when they make investment

⁸Allowing the shock to be applied to more than the first period of a default does not change our mechanism. This assumption is merely for tractability.

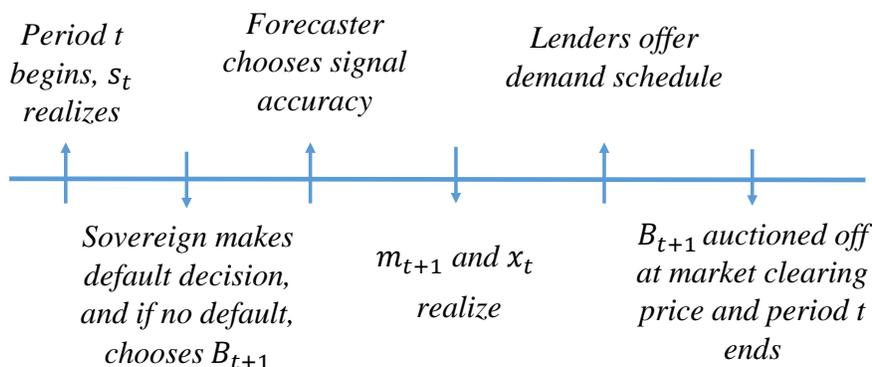
decisions.

For clarity, we treat σ_m as a fixed parameter that we will calibrate. We are not interested in the exact interpretation or source of these hidden shocks, so much as how their presence causes lenders to respond to those that are obviously observed.⁹

2.2. Timing

The timing of events can be found in Figure 1 and is as follows: Period t begins with the realization of s_t , following which the sovereign makes a default decision. Conditional on repayment, it then chooses a level of debt issuance B_{t+1} to maximize its expected utility prior to the realization of m_{t+1} .

Figure 1: Timing of Events in a Period



Next, a professional forecaster, who observes the public states s_t and B_{t+1} , chooses the accuracy with which he acquires information about m_{t+1} given some information cost. He designs a signal of the unobserved shock, x_t , and can pay a cost to increase its accuracy.

Following the information acquisition decision, m_{t+1} and x_t are jointly realized in the middle of period t . The market coordinates on the signal: Competitive lenders know both the signal and its accuracy, and then decide bond demand. The sovereign then determines an issuance price that clears the bond market and period t ends.

Notice that we assume that the sovereign cannot change its bond supply following the

⁹Further, even if these shocks featured time-varying volatility, lenders would not respond differently unless there were some public state that signaled this.

realization of m_{t+1} .¹⁰ This allows us to focus on the role of information acquisition and avoid the complicated and, for our purposes, unnecessary signaling game that would ensue.

2.3. Sovereign Borrower

As is standard in the literature, we use a recursive, Markov-Perfect specification with limited commitment on the part of the sovereign. At the beginning of each period, it compares the value of repaying debt, $V_{R,t}$, with that of default, $V_{D,t}$, and chooses the option that provides a greater value:

$$V_t(s_t, B_t, m_t) = \max\{V_{R,t}(s_t, B_t), V_{D,t}(s_t, m_t)\}$$

Given the timing assumption, we can express the value of repayment at the beginning of period t as follows:

$$V_{R,t}(s_t, B_t) = \max_{B_{t+1} \in \mathcal{B}_t} E_{\tilde{m}_{t+1}, \tilde{x}_t} \left[\log(C_t(\tilde{x}_t)) + \beta E_{\tilde{s}_{t+1}|s_t} V_{t+1}(\tilde{s}_{t+1}, B_{t+1}, \tilde{m}_{t+1}) \right]$$

subject to $C_t(\tilde{x}_t) = Y_t - B_t + q_t(B_{t+1}|s_t, \tilde{x}_t)B_{t+1}$ (1)

The sovereign has time-separable log-preferences over consumption¹¹ and is a monopolist in his own debt market. The determination of the issuance price schedule, $q_t(B_{t+1}|s_t, x_t)$, will be discussed in the market clearing section below. If the budget set is ever empty, either before or after the realization of x_t , we follow the literature standard and assume that the sovereign defaults.¹²

We assume that when a default happens in period t all debt is wiped out. In the first period of the default regime, the sovereign is subject to the shock m_t realized in period $t - 1$. During the entire default regime, the sovereign is excluded from capital markets for a

¹⁰We do not specify whether sovereign gets to know m_{t+1} before or after the lenders do. Since the sovereign decides its bond supply before m_{t+1} realizes and cannot change its issuance decision in either case, it does not matter.

¹¹None of the intuition behind our results relies on the assumption of log-utility. Any concave function will work. The benefit of using log-utility is that the unobserved shock can be interpreted either as an endowment/supply shock or as a preference/demand shock. We demonstrate this in [Appendix A](#).

¹²In the calibrated model, we can easily construct the set \mathcal{B}_t such that the budget set is never empty at any point in the state space and the upper bound on B_t never binds.

random number of periods and faces persistent output losses, though it is subject to the m_t shock only in the first period of the default.

The costs during the period of exclusion are output losses that consist of two components, known and unknown. We will denote the former with $\psi > 0$; the latter will be the shock m_t . These costs could be interpreted as the usual consequences of tightening credit conditions, a disruption of trade credit, or a banking slump (Mendoza and Yue [2012] or Sosa-Padilla [2012]). m_t is intended to stand in for the components of those costs that are harder to predict, such as the magnitude of private capital outflows, the impact of likely fiscal consolidation, or the severity of the ensuing international sanctions, about which lenders may not be perfectly informed. Under these assumptions, the value of default can be expressed recursively as follows:

$$V_{D,t}(s_t, m_t) = \log(C_t) + \beta E_{\tilde{s}_{t+1}|s_t} [\phi V_{t+1}(\tilde{s}_{t+1}, 0, 0) + (1 - \phi)V_{D,t+1}(\tilde{s}_{t+1}, 0)]$$

subject to $C_t = Y_t \times e^{-\psi+m_t}$ (2)

We assume that m_t is normally distributed around zero.¹³ ϕ is the Poisson rate at which the sovereign regains access to international capital markets.

We define the sovereign's default decision with a binary operator:

$$d_t(m_t, s_t, B_t) = \mathbf{1}\{V_{R,t}(s_t, B_t) < V_{D,t}(s_t, m_t)\}$$

2.4. Forecaster

There is an inherent difficulty associated with market-based information acquisition problems: The price tends to convey too much information. Perfect Bayesian investors can infer all relevant information from market prices, which gives them no incentive to acquire information in the first place (Grossman [1976] or Dow and Gorton [2006]).

Rather than bounding rationality by including noise traders (Kyle [1985]), we circumvent

¹³The calibrated level of σ_m will be sufficiently small that 'efficient' defaults in which $m_t > \psi$ will have near zero probability and will never materialize on any simulated path.

this problem by designing the market to operate with complete rationality and transparency on a signal of the true hidden information rather than the hidden information itself.¹⁴ Thus, no additional information regarding payoff-relevant states can be gleaned from the price besides what participants already have.

To do so, we separate the information acquisition decision from market participation, much like [Angeletos and Werning \(2006\)](#). We assume that all lenders are fully rational, but that each acquires information by employing the same contractor, whom we call the forecaster. In a sense, the forecaster serves as a market maker by coordinating market information. [Veldkamp \(2011\)](#) argues that information is a commodity that is often difficult to obtain but essentially costless to disseminate. Our market set-up reflects this feature of information production by allocating the task to a single specialist. Once the information is produced, it is costlessly transmitted to market participants. Real-world analogues of the forecaster might be financial software and information firms such as Bloomberg or Reuters; alternatively, they could be or financial news or media outlets such as the Wall Street Journal of the Financial Times; credit rating agencies such as Moody's or S&P; or multinational public institutions such as the International Monetary Fund.

The forecaster has a technology capable of gathering information regarding the unobserved shock, m_{t+1} , at a per-unit cost κ that we specify later. He sells his services to lenders and sovereign as an independent contractor for an exogenous, fixed, non-state-contingent fee. The fee for the lenders is $l_t \geq 0$. For simplicity, we assume that the fee is cointegrated with the sovereign's output.¹⁵

A fixed, non-state-contingent fee is not unrealistic. For instance, one pays a fixed subscription fee to access IMF's International Financial Statistics data; the fee is not contingent on the data updating frequency. For rating agencies, once a credit rating has been published,

¹⁴The other popular option is to add unseen noise to the aggregate supply e.g. [Grossman and Stiglitz \(1980\)](#). In our model, however, supply is determined by the optimal behavior of the sovereign. Thus, there is little plausible room for such noise.

¹⁵Credit rating agencies often charge fees to the borrowers rather than the investors. We could allow for a fee for the sovereign as well e.g. some non-state-contingent $L_{s,t} \geq 0$ that was cointegrated with output. Output in the current set-up would then be interpreted at $Y_t - L_{s,t}$ and the model and results would otherwise not change. In practice, such fees are small relative to the size of the country.

they monitor that credit rating on an ongoing basis, rectify it as necessary in response to changes, and notify their subscribers.¹⁶ We also assume for simplicity that lenders must pay this subscription fee to have access to the sovereign bond market, such that all lenders in this market have access to a forecaster-provided signal of the unobserved shock.

The forecaster is interested in the integrity of its forecasts, much like the rating agencies in [Holden et al. \(2018\)](#) or [Manso \(2013\)](#). It actively weights this objective against information acquisition costs. In each period it produces a signal, x_t , of the next period's unobserved shock, m_{t+1} . The signal and the unobserved state are jointly normal, and the information contained in this signal is reflected in $\rho_{mx,t} = \text{corr}(x_t, m_{t+1}) \in [0, 1]$, which is the forecaster's choice.¹⁷ That is, the forecaster can modify the signal to make it more or less informative about m_{t+1} : More informative signals will feature a larger $\rho_{mx,t}$. For this reason we will call $\rho_{mx,t}$ the *accuracy* or *precision* of the signal. This information structure is isomorphic to a fully Gaussian signal-noise model in which the forecaster chooses the variance in the noise. We prefer our specification to this, though, since it features a closed domain i.e. an infinitely noisy signal would arise with some regularity along the equilibrium path.

The forecaster uses this information to publish a forecast (distribution) over all future states, observed and unobserved: $h_t(s_{t+1}, m_{t+1} | x_{t+1}, s_t) = f_t(s_{t+1} | s_t) g_{\rho_{mx,t}}(m_{t+1} | x_t)$. The goal is to minimize the mean-square-error of the default-risk forecast under this distribution.

In our benchmark, we assume m_{t+1} and x_t to be orthogonal to observed states s_t , but our framework is flexible enough to allow for some correlation with no change in the mechanism. The forecaster would simply acquire residual information that is not conveyed through observed states.

The information required to obtain a signal is given by a time-invariant function, $I(\rho_{mx,t})$, which is increasing in signal accuracy. The per-unit cost of information is a constant κ . In the benchmark, we assume that $I(\cdot)$ is the reduction in entropy in m_{t+1} that comes from

¹⁶Admittedly, sovereign ratings are sometimes solicited by the issuer countries themselves, but this typically does not happen during crises, which we will see are the periods in which the most information is acquired.

¹⁷Our restriction to signals with positive correlation is without loss of generality, since negatively correlated signals have the same information content.

knowledge of x_t , but our results do not hinge on this functional form.¹⁸ Any increasing function would work.

We formulate the forecaster's information acquisition problem as below, given s_t and B_{t+1} :

$$\begin{aligned} \min_{\rho_{mx,t} \in [0,1]} \quad & E_{\tilde{x}_t} E_{\tilde{m}_{t+1}, \tilde{s}_{t+1} | \tilde{x}_t, s_t} [d_t(\tilde{m}_{t+1}, \tilde{s}_{t+1}, B_{t+1}) - \bar{d}_t]^2 + \kappa I(\rho_{mx,t}) \quad (3) \\ \text{subject to} \quad & \bar{d}_t = E_{\tilde{m}_{t+1}, \tilde{s}_{t+1} | s_t} [d_t(\tilde{m}_{t+1}, \tilde{s}_{t+1}, B_{t+1})] \end{aligned}$$

where $d_t(\cdot)$ is the binary default identifier.¹⁹ To see the benefit of information acquisition, notice that the variance of the interior expectation is decreasing in $\rho_{mx,t}$. Consequently, the variance in the forecaster's default forecast can be reduced if he is willing to undergo costly information acquisition. When s_t and B_{t+1} indicate a greater risk of default, acquisition of more accurate information will be optimal, since the unobserved shock matters for the variance of the forecast. When these publicly known states indicate instead that there is little to no default risk, the forecaster can save on information costs and provide imprecise or even orthogonal signals (i.e., $\rho_{mx,t} = 0$), because m_{t+1} is adding little to no variance to the forecast.

2.5. Foreign Lenders

There is a unit mass of risk-averse foreign lenders who invest in risky sovereign debt. These lenders act competitively, similarly to [Lizarazo \(2013\)](#) or [Aguiar et al. \(2016a\)](#). Lenders arrive in overlapping generations and each lives for two periods. Each lender is endowed with wealth, w_t , pays the fixed contractor fee, l_t , which is assumed to be relatively small in relation to B_t or w_t , to gain access to the forecaster's signals. Then they solve a portfolio allocation problem, deciding how much to invest in risky sovereign debt and how much to invest in a risk-free asset yielding a return, r_t . For simplicity, we follow [Aguiar et al. \(2016b\)](#)

¹⁸This notion of information was developed primarily by [Shannon \(1958\)](#) and applied to economics by [Sims \(2003, 2006\)](#). Here we have $I(\rho_{mx,t}) = \frac{1}{2} \log_2 \left(\frac{1}{1 - \rho_{mx,t}^2} \right)$.

¹⁹We prefer to think of the objective as a utility rather than a resource cost, so we do not incorporate their fees into the objective. Since the fee is non-state-contingent, doing so would not change the model.

and assume that their wealth is cointegrated with sovereign output.

When making investment decisions in period t , lenders observe s_t and B_{t+1} . Further, they can access the forecaster's signal x_t . The lenders also know the signal's informativeness since they know the policy function of the forecaster: $\rho_{mx,t} = \rho_t(s_t, B_{t+1})$.

Given the above setup, when choosing to include x_t in the information set,²⁰ investor i takes the bond issuance price, q , as given and solves the following problem:

$$\max_{b_{i,t+1}} E_{\tilde{s}_{t+1}, \tilde{m}_{t+1} | s_t, x_t} \left[\frac{c_{i,t+1}^{1-\gamma_L}}{1-\gamma_L} \right] \quad (4)$$

$$\text{subject to } c_{i,t+1} = (w_t - l_t - b_{i,t+1}q)(1+r) + b_{i,t+1} [1 - d_t(\tilde{m}_{t+1}, \tilde{s}_{t+1}, B_{t+1})]$$

The equilibrium pricing function, as well as all other equilibrium objects, are also contained in the lenders' information set, but since it contains no hidden information there is no Bayesian extraction problem to be undertaken as in [Robert E. Lucas \(1972\)](#), [Grossman \(1976\)](#), or [Bassetto and Galli \(2017\)](#).

We denote aggregate bond demand in any period by:

$$B_{D,t+1}(s_t, x_t, B_{t+1}, q) = \int_0^1 b_{i,t+1}^*(s_t, x_t, B_{t+1}, q) di.$$

Note that there is no heterogeneity amongst lenders, and so the i -index will be irrelevant in the benchmark.

2.6. Market Clearing

Once the correlation is chosen by the forecaster, the signal is realized and distributed to all lenders, who then enter a competitive market with a common information set. The sovereign issues its predetermined debt stock, B_{t+1} at the highest possible price. Market clearing requires that $B_{D,t+1} = B_{t+1}$. This yields a pricing schedule identical in structure to that in [Aguiar et al. \(2016b\)](#) but with the inclusion of the signal realization as an additional

²⁰For the sake of brevity we do not write out explicitly the problem without x_t in the information set.

state.

2.7. Equilibrium Definition

Having described the model, we can now define our equilibrium:

Definition 1. *A Markov Perfect Equilibrium is a set of functions,*

$\{V_t(s_t, B_t, m_t), V_{R,t}(s_t, B_t), A_t(s_t, B_t), V_{D,t}(s_t, m_t), q_t(B_{t+1}|s_t, x_t), \rho_t(s_t, B_{t+1})\}_{t=0}^\infty$ *such that*

1. $V_{R,t}(s_t, B_t)$ and $V_{D,t}(s_t, m_t)$ *solve Recursions 1 and 2 and imply the policy*
 $B_{t+1} = A_t(s_t, B_t)$. *Further, $V_t(s_t, B_t, m_t) = \max\{V_{R,t}(s_t, B_t), V_{D,t}(s_t, m_t)\}$.*
2. $\rho_t(s_t, B_{t+1})$ *solves Problem 3.*
3. $q_t(B_{t+1}|s_t, x_t)$ *ensures that $B_{t+1} = B_{D,t+1}(s_t, x_t, B_{t+1}, q_t(B_{t+1}|s_t, x_t))$ where bond demand is derived from Problem 4.*

In [Appendix B](#), we demonstrate how this model can be stationarized for solution purposes. This model will behave in most respects as a standard quantitative sovereign default model in the vein of [Eaton and Gersovitz \(1981\)](#). The novel feature is how information is collected and transmitted and what that implies for bond price.

The model endogenously generates acquisition of information that is contingent on observable states. During non-crisis times, information regarding unobserved shocks is not particularly valuable for improving the default-risk forecast. Thus, such information is not acquired and sovereign debt is priced assuming these shocks to be at their average.

During crises, however, such information becomes valuable. Consequently, it is both acquired and priced, which increases the volatility of the sovereign spreads. This generates a number of interesting results including time-varying spread volatility, information-driven risk premia that are relevant for econometric inference, and novel welfare results on transparency for the sovereign borrower, which we demonstrate below.

3. Quantitative Analysis

To determine the impact that costly information acquisition has on the pricing of sovereign risk, we calibrate the model to match a set of empirical moments from Ukraine from 2004-2014. We choose Ukraine since the volatility of its real growth rate process during this time is similar to Argentina during the 1990’s, which is the canonical calibration choice for models in this vein (Aguiar and Gopinath [2006] or Arellano [2008]). Aguiar et al. (2016a, 2016b) show that growth volatility is particularly important for generating realistic spread dynamics. Further, Ukraine was at the heart of several news cycles over the course of this period, including political upheavals during the Russo-Georgian War in August 2008 and the annexation of Crimea by the Russian Federation in early 2014. We solve the model using an iterative procedure over a discrete grid. Details regarding the solution method can be found in [Appendix B](#).

3.1. Data and Calibration

We take data from three primary sources: First is the JP Morgan Emerging Market Bond Index (EMBI) database taken from Datastream; second is the World Bank database; and third are rating publications data from S&P, Moody’s, and Fitch.

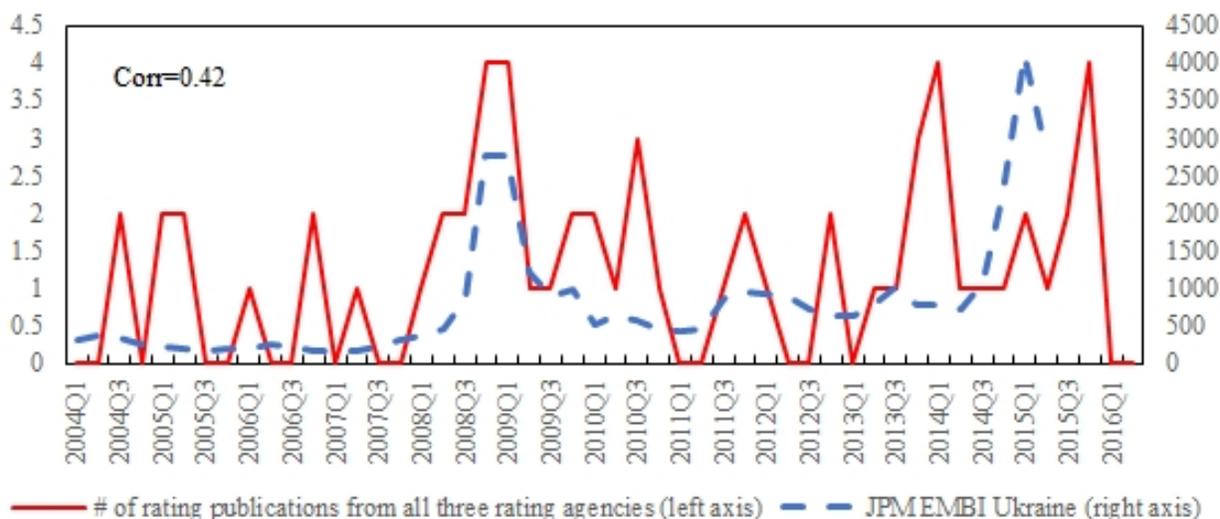
3.1.1. Information Cost Identification

First and most importantly, to find a proper cost value per unit information, i.e., κ , we match the variability of information acquisition in the model and in the data. The model has a direct measure of information acquisition, $I(\rho_{mx})$. We now describe how we compute this metric in the data.

The literature has proposed many different measures of information acquisition or ‘attention’ ranging from news to trading volume to extreme returns to Google search trends (Barber and Odean [2008], Gervais et al. [2001], Seasholes and Wu [2007], or Da et al. [2011]). Given our market set-up, however, we instead select an empirical entity to serve as the forecaster and use information on the products they offer to investors.

In the benchmark model we impute this role to credit rating agencies. Credit rating agencies have been thought of before as serving a coordinating role in the market for sovereign debt, e.g., [Holden et al. \(2018\)](#). In particular, as a proxy for information output, we consider the total number of sovereign credit rating publications on Ukraine by the three largest credit rating agencies (Moody’s, S&P, and Fitch) in a given quarter. This series can be seen in Figure 2 alongside the EMBI return for Ukraine. The correlation between these series is quite high at 0.42. In [Appendix D](#), we also show that the total number of credit rating publications is highly correlated with other indicators for lenders’ or forecasters’ attention, such as extreme returns on sovereign bonds (with correlation 54%).

Figure 2: EMBI and Total Number of Credit Rating Publications, Ukraine



We define an information threshold to be having strictly greater than three publications total issued by the three rating agencies in a given quarter. We define such quarters to be periods of ‘intense attention.’ The fraction of such periods with intense attention is our calibration target. In this case, it is 8.0%. It captures the frequency of large, positive, and relatively discrete jumps in information acquisition. More specifically, the intense-attention periods include the Russo-Georgian conflict around 2008 and Ukraine’s conflict with Russia and the annexation of the Crimean peninsula in early 2014. This is the sort of binary attention behavior that our model predicts and so it is a natural target for the information

cost identification.

3.1.2. Calibration

To obtain the output process, we estimate via MLE an AR(1) process ($\rho_g = 0.5058, \mu_g = 0.0126, \sigma_g = 0.0846$) on dollar-valued GDP growth for Ukraine from 2004-2014 at a quarterly frequency.

Table 1: Calibration by Simulated Method of Moments

Parameter	Value	Target Description	Data	Model
Discount factor	$\beta=0.8110$	Annual Default Frequency	1.5%	1.2%
Output cost (known)	$\psi=0.0227$	Average Debt-Service-to-GDP Ratio	12.6%	11.3%
Lender wealth	$w=2.5000$	Average Spread	6.5%	5.7%
Unobs shock std dev	$\sigma_m=0.0153$	Spread Volatility	5.5%	5.6%
Unit info cost	$\kappa = 5.272 \times 10^{-4}$	Fraction of Quarters with $IA > \zeta$	8.0%	7.5%

We assume that the risk-free rate is fixed at 1% quarterly; that the lenders exhibit constant relative risk-aversion preferences with CRRA $\gamma_L = 2$, which is standard; and that $\theta = 0.083$, which is an estimate used by [Mendoza and Yue \(2012\)](#) for an average duration of 6 years before returning to international bond market. For simplicity, we also assume that lender wealth is constant over time barring its cointegration with output i.e. $w_t = wY_t$ and $l_t = 0$, since these rating agencies do not receive funds directly from investors.

We calibrate the remaining five parameters, $\{\beta, \psi, w, \sigma_m, \kappa\}$, using simulated method of moments (SMM) to target the simulated results from our model at five moments from the corresponding data:²¹ Annual default frequency, average debt-service-to-GDP ratio, annual spread volatility, average annual spread, and fraction of time in which information acquisition (IA) is above its midpoint, $\zeta = I(0.5)$.

This last target is the most novel in our approach. Here, we define periods of ‘intense attention’ and match the fraction of such periods with its counterpart in the data. This fraction is monotone in the cost of information, as we will see in the comparative static exercises. Model results are not sensitive to changing the threshold ζ from the midpoint. In

²¹Our model and its results are at quarterly frequency. The model results are adjusted to annual statistics to match the listed annual targets.

the model, information acquisition tends to highly binary even though it is allowed to be continuous. The forecaster tends to pay either no attention at all or nearly full attention. Consequently, the exact choice of ζ is relatively inconsequential so long as it divides these two types of behavior.

These parameters are given in Table 1. Each parameter is primarily identified by its corresponding target moment, though there are significant cross-partial effects. The resulting parameterization is fairly standard²² and the match is quite close for a 5-dimensional non-linear matching problem.

In the following sections, we explore the model’s quantitative implications.

3.2. Model Behavior

Before we exposit the properties that are unique to this model, we show that along many dimensions it preserves key features of a standard quantitative sovereign default model. For instance, Figure 3a gives the bond demand functions in equilibrium. It exhibits a simple, downward sloping feature that looks remarkably similar to [Aguiar and Gopinath \(2006\)](#) or [Arellano \(2008\)](#) despite the added complexity of endogenous information acquisition. We can see that better growth shocks lead to higher price schedules.

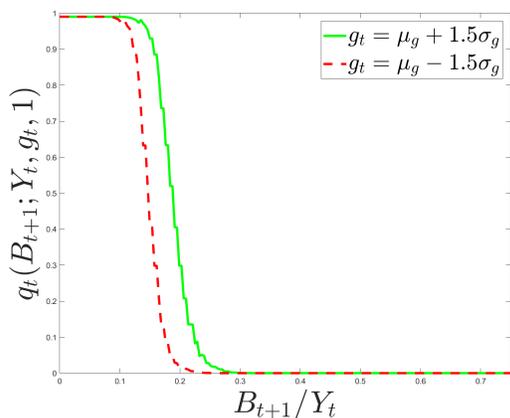
The equilibrium policy functions are given in Figure 3b. We can see that better growth shocks lead to higher debt issuance. This is a standard feature of the quantitative models discussed in [Aguiar et al. \(2016a\)](#).

Figure 4 provides an event study surrounding a default. Again, the model behaves as a standard model would. We can see that spreads increase prior to a default. This is because growth follows the boom-bust cycle documented in [Aguiar et al. \(2016a\)](#): Leading up to a default is a series of benevolent growth shocks that induce excessive borrowing, which raises spreads. A default occurs when the sovereign accumulates a large amount of debt in this fashion and then experiences a severe and unanticipated drop in growth (not shown in the

²²While β may at first glance appear to be very low, it is in the neighborhood of estimates from similar models e.g. [Aguiar and Gopinath \(2006\)](#) or [Aguiar et al. \(2016a\)](#).

Figure 3: Benchmark Behavior: Policy and Pricing Functions

(a) Equilibrium Pricing Function



(b) Equilibrium Policy Function

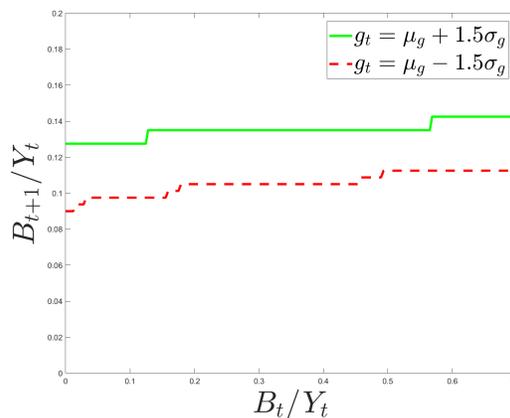
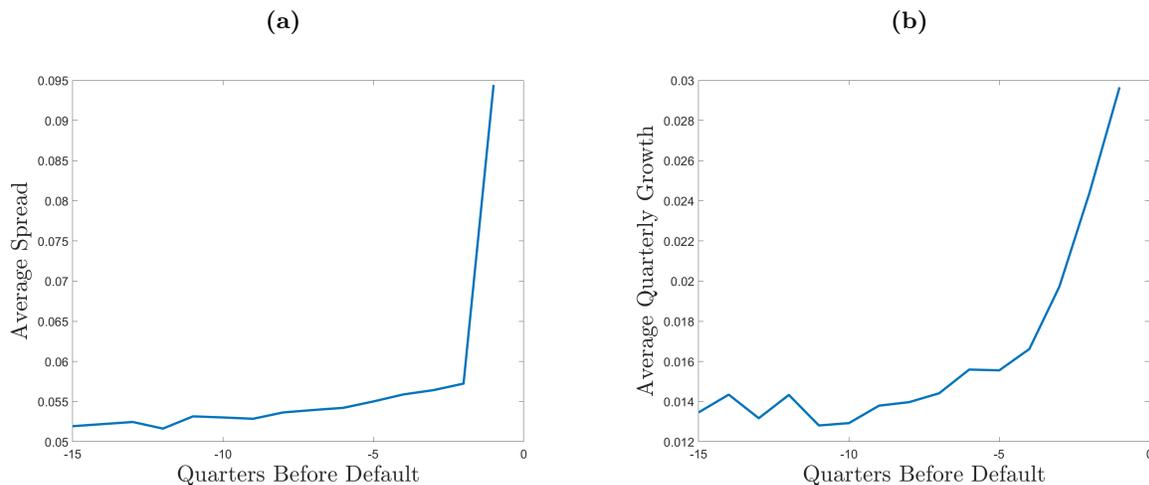


figure).

What is novel in this model is the state-contingent acquisition of information: More information is endogenously acquired during times of crises i.e. near default. Recall from the calibration that the ergodic mean of the debt-to-GDP ratio is about 11.3%. We can see in the policy functions in Figure 5a that for debt levels above this, information precision depends on the underlying growth shock i.e. the observed shock. If growth is high, no useful information is acquired unless debt levels become very large. However, when growth is low, the forecaster acquires useful information and the lenders pay attention to it for lower debt levels as well. Notice further that for very high debt levels information acquisition drops off. This is because default is near certain in these regions, regardless of the realization of the unobserved shocks. Consequently, there is no point for the forecaster to pay a cost to learn about those shocks.

The bond demand schedule across x_t can be seen Figure 5b. For low debt levels, the forecaster does not acquire useful information about unobserved shocks; this implies that lenders do not pay attention to these uninformative signals and thus prices do not react. Hence, there is no difference across the two price schedules for debt-to-GDP levels lower than about 10% or so. As debt levels rise, though, the forecaster begins to acquire useful

Figure 4: Benchmark Behavior Around Default



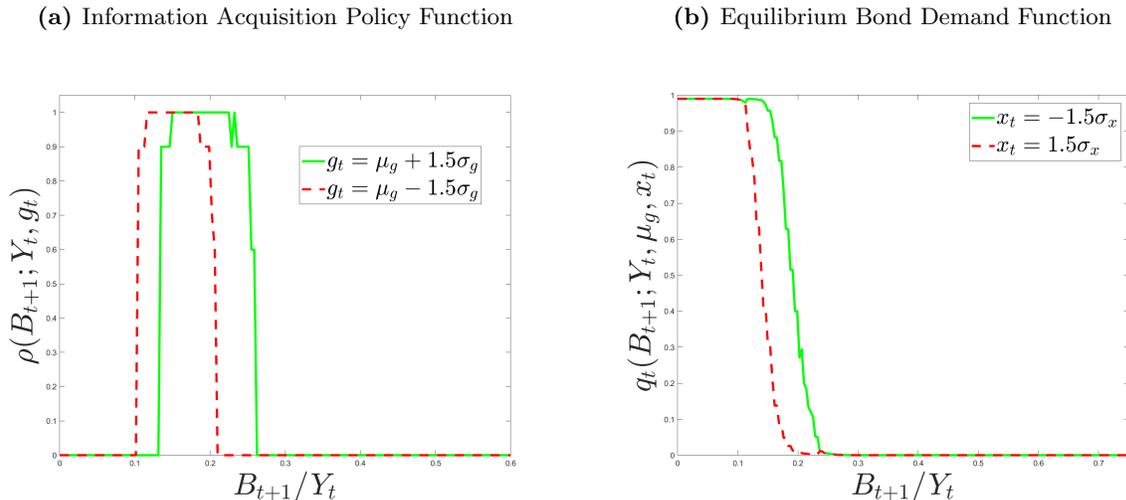
information, the lenders start to pay attention to it, and the bond demand schedules begin to diverge according to the different signals. This divergence lines up exactly with the non-zero information-acquisition regions from Figure 5a.

What do these policy functions imply for the behavior of endogenous objects? Figure 6 provides a graph of optimal signal correlation leading up to a default event. We can see that attention increases before such crises. Signal precision is relatively low in the periods prior to the event i.e. during non-crisis times. As observed states indicate that default risk is rising, however, the forecaster becomes more aggressive in its acquisition of information, increasing it from an average of about $\rho_{mx} = 0.055$ to $\rho_{mx} = 0.086$ on the cusp of a default event. The acquisition of more information near defaults or crises has a number of interesting implications and results that we will explore over the next few subsections.

3.3. Comparative Statics

The state-contingent acquisition of information and information cost have large effects on some model moments and minimal effects on others. Among those that are relatively invariant to information costs are average and median debt levels as well as default frequency, all of which hover around their benchmark levels for the complete relevant set of information costs.

Figure 5: Information Acquisition: Model Objects



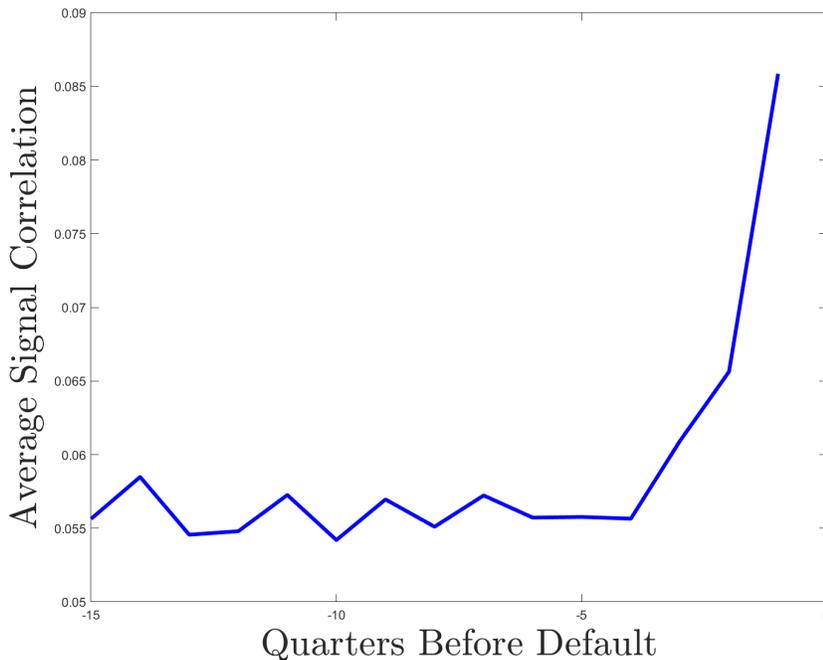
The most relevant affected moments are shown in Figure 7, which provides a comparative static of how the model responds to an array of information costs.²³ These moments consist of the average spread, the median spread, the spread volatility, and the attention fraction. In the figure, black lines correspond to the left y-axis, while blue lines with intermittent crosses correspond to the right y-axis.

As information costs decrease, both the fraction of time spent paying attention to the sovereign and the spread volatility increase significantly. The intuition for the former is trivial and was in fact used for the calibration; the intuition for the latter is that cheaper information means that unobserved shocks are priced more often through the signals instead of being considered at their average. This naturally increases overall spread volatility.

The impact across measures of central tendency as information costs fall is quite interesting. The average spread rises modestly as information costs fall near zero, but the median spread falls substantially. The reason for the former is the increased spread levels during crises. Crises are no more frequent as information costs fall, but during them imminent default risk is more accurately priced so spreads are higher. These equally rare but increasingly large spread spikes bring the average up as information costs fall.

²³The benchmark model is very near the center of these plots since $\kappa = 5.272e - 4$.

Figure 6: Information Acquisition Before Default

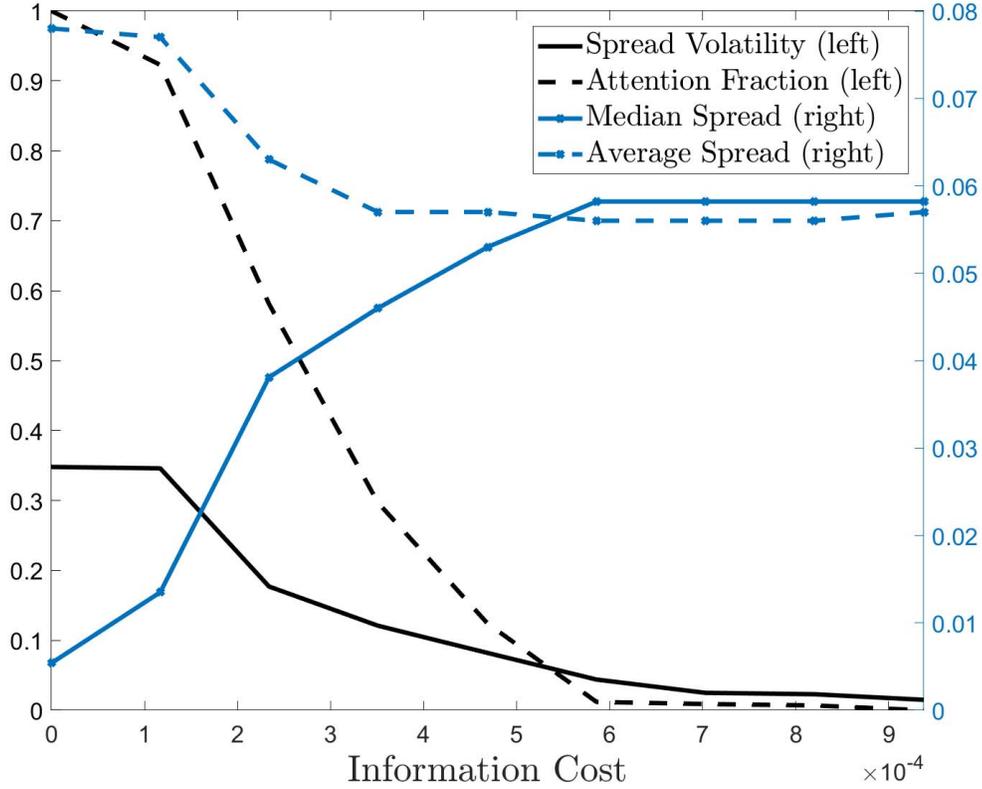


On the other hand, the median spread falls with information costs. This is because the information-driven risk premium demanded by the lenders falls as the information uncertainty generated by m_{t+1} is more frequently tamed by more accurate information regarding it. And this reduction in the risk premium happens in most states of the world i.e. non-crisis times, which brings the median spread down. This result accords with the findings of [Bernoth and Wolff \(2008\)](#) and [Iara and Wolff \(2014\)](#) that increased transparency tends to reduce borrowing costs.

The comparative statics here also provide us a robustness check on the information cost calibration. Even if we recalibrate our attention fraction target to as large as 11% or as small as 4%, our calibrated information cost changes only between 5.0×10^{-4} to 5.5×10^{-4} . The consequent impact on the average spread, the median spread, and the spread volatility are about 0.25%, 0.1%, and 1.5%, respectively.

It is also worth noting that once information cost moves above 6.0×10^{-4} , the changes to those simulated moments become minimal. That is, once information cost is high enough,

Figure 7: Relevant Moments Across κ

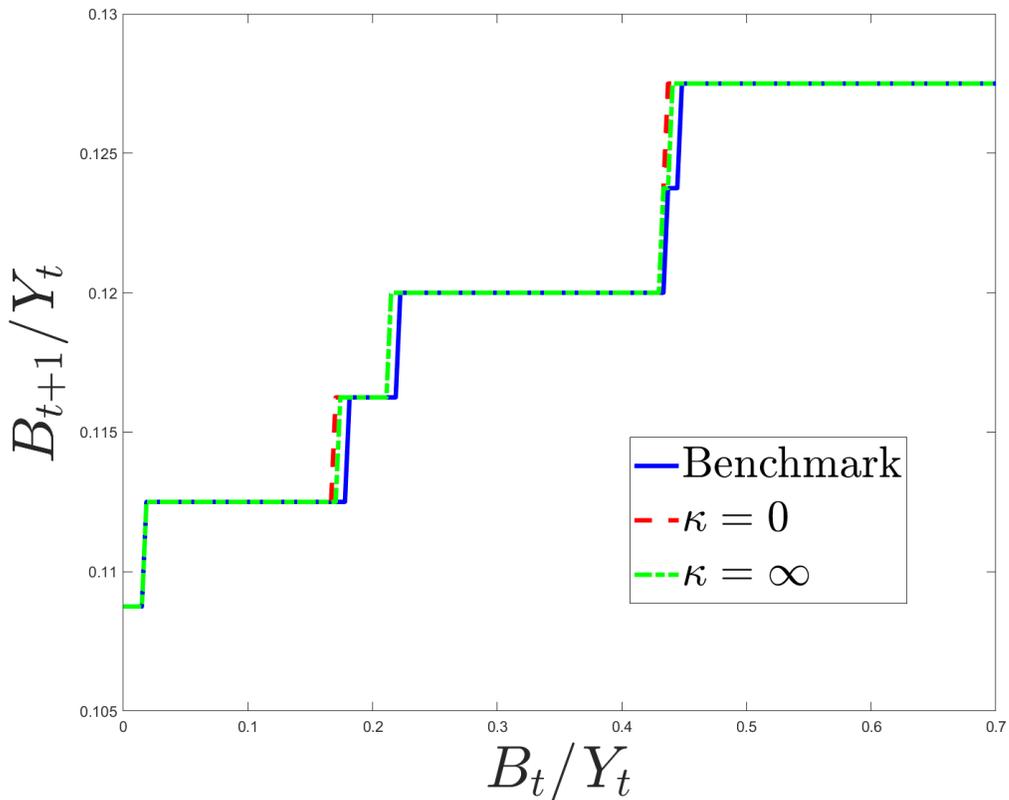


the investors/forecasters stop acquiring information as the attention fraction is near zero, as if the cost is infinite, and thus the model moments do not change much.

Finally, as mentioned above, the impact on average and median debt levels as well as default frequency is minute. This follows from the more general quantitative result that general equilibrium effects, for which we do allow, appear to be relatively small. This can be seen in the sovereign’s debt issuance policy function in Figure 8.

Figure 8 provides the policy function for the benchmark case as well as the full-information and no-information cases evaluated at steady state growth. We can see that the differences, while greater than convergence tolerance are small. They also appear to be non-monotone. More borrowing takes place in the full-information version than the no-information version, but the benchmark, which falls between them, implies the least amount of borrowing ceteris paribus. Consequently, our key results are derived from the direct effect of information costs

Figure 8: Policy Functions Across Information Costs



rather than general equilibrium repercussions.

3.4. Implications and Results

This section details three important bond price implications from our sovereign default model with costly information acquisition.

3.4.1. Time-Varying Volatility

The first key result is that our model endogenously generates time-variation in sovereign spread volatility. In the model, spreads exhibit increased volatility during crises when lenders price the unobserved shocks more accurately, rather than considering them to be at their mean, which they do during non-crisis times. Further, the time-varying volatility is strongly countercyclical and positively correlated with spread levels. This implies that one could

interpret our framework as a micro-foundation for such models as [Melino and Turnbull \(1990\)](#) or [Fernández-Villaverde et al. \(2011\)](#).

To assess our model’s ability to generate time-variation in the spread volatility, we propose a model-free metric that we call the **Crisis Volatility Ratio** or CVR.²⁴ It is defined as follows: In a series of data, either simulated or empirical, let \hat{T} denote the set of all periods in which the change in the spread from the prior period is above the 97.5 percentile of its distribution. We follow [Aguiar et al. \(2016a\)](#) in calling such events “crises,” and by construction they are about $8.3\times$ more likely to happen than default events and can be observed in the data even if a default cannot. With this notation, we define the CVR as

$$CVR = \frac{1}{|\hat{T}|} \sum_{t \in \hat{T}} \frac{\hat{\sigma}_{t:t+w}}{\hat{\sigma}_{t-w-1:t-1}}$$

where $\hat{\sigma}_{x:y}$ is the sample standard deviation calculated using the periods from x to y . This ratio compares the volatility in a window of w periods immediately prior to a crisis to the volatility in a window of w periods after. Neither window includes the crisis itself. In the benchmark, we set $w = 5$. If the CVR is larger than one, then crisis periods tend to be more volatile than non-crisis periods. To give a sense of the metric, any AR(1) process would feature an average CVR equal to one since the volatility of the innovation is independent of the any prior observation.

We first compute the CVR for the data (the first column in [Table 2](#)). We can see that time-varying volatility is a strong feature of the data by this metric: The CVR in the data is more than 3.5 times what than an AR(1) would suggest. This is no surprise, as strong time-variation in the volatility is documented by [Fernández-Villaverde et al. \(2011\)](#).

What is surprising is the capacity of our model to explain this. Our benchmark model (the second column in [Table 2](#)) can explain roughly 79% of this non-targeted moment: 2.89

²⁴Typically, when measuring time-varying volatility, the literature imposes quite a bit of structure. For instance, [Melino and Turnbull \(1990\)](#) and [Fernández-Villaverde et al. \(2011\)](#) can measure the impact of time-varying volatility in the context of a stochastic volatility model. Imposing this or another similar structure to measure the quantitative efficacy of the model would be inappropriate in our case, since we know that the data-generating process for the simulated data is not a stochastic volatility model.

Table 2: Crisis Volatility Ratios (First Row) and Counterfactual Moments

Data (Ukraine)	Benchmark Model	$\kappa = \infty$ Counterfactual	$\sigma_m = 0$ Counterfactual
3.67	2.89	1.32	1.27
Def. Freq.	1.2%	1.1%	0.7%
E[Debt Service/GDP]	11.3%	11.3%	13.3%
E[Spread]	5.7%	5.7%	3.7%
Std(Spread)	5.6%	1.5%	2.0%
Fraction with $IA > \zeta$	7.5%	0.0	0.0

out of 3.67. This is worth noting, since all time-variation in the volatility is generated endogenously.

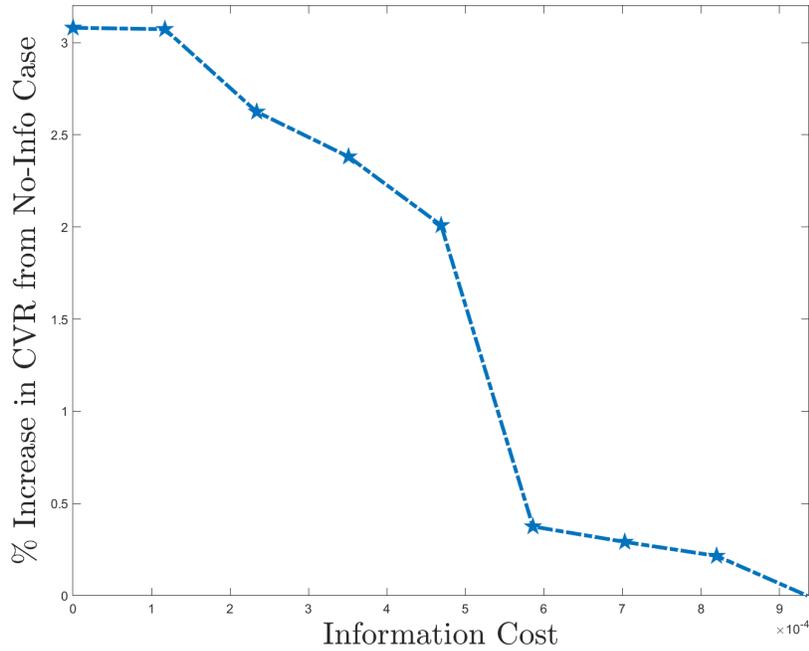
The vast bulk of this explanatory power is due to costly information acquisition. To see this, we compare our model to two different counterfactuals and report the CVRs in the last two columns in Table 2. First, we re-solve the model for the same parameters but with infinite information costs, such that unobserved shocks are never priced. Our benchmark model generates a 119% increase in the CVR (from 1.32 to 2.89) relative to the this counterfactual. Second, we re-solve the model with the same parameters but no hidden information i.e. $\sigma_m = 0$. Our model generates a 128% increase in the CVR (from 1.27 to 2.89) relative to this counterfactual, which is equivalent to the model of [Aguiar et al. \(2016b\)](#) with permanent shocks, short-term debt, and a modestly different calibration.

It is worth noting that there is some time-variation in volatility in the model even without information frictions. This is because even those publicly known shocks' fluctuations are not always relevant for default risk and consequently are not always priced. This implies lower price volatility in normal times than crises. This is true in any model of endogenous sovereign default and accords with the findings of [Bocola and DAVIS \(2016\)](#). Nevertheless, our model's ability to more than double this underlying the time-variation in spread volatility and thus bring it to empirically relevant levels is substantial and noteworthy.

While our benchmark model is calibrated to match the data and these counterfactuals are not, Table 2 also gives the relevant model moments for these counterfactual models, which are within the empirically relevant neighborhood for a typical emerging market.

Further evidence of the power of costly information acquisition to generate time-variation

Figure 9: Crisis Volatility Ratios Across κ



in the spread volatility can be found in Figure 9, which gives the percent increase in the CVR from the infinite-cost counterfactual for a wide range of different information costs. It makes clear that the CVR is monotonically decreasing in the information cost. Further, this relationship is quite steep, with the full-information model generating a more than 300% increase over the no-information model.

3.4.2. State-Contingent Risk Premia

The second result we highlight is that the composition of spreads is not the same during crisis times and non-crisis times. To understand this, first note that we can intuitively break the spread on sovereign debt into two categories: default risk and a risk premium for that

default risk.²⁵

$$\begin{aligned}
 Sprd_t \approx & \overbrace{E_{\tilde{s}_{t+1}, m_t | s_t, x_t} [1 - R_{t+1}(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1})]}^{\text{Risk-Neutral Spread/Default Probability}} + \\
 & \overbrace{\frac{-cov_{\tilde{s}_{t+1}, m_t | s_t, x_t}(R_{t+1}(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1}), u'(c_L(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1})))}{E_{\tilde{s}_{t+1}, m_t | s_t, x_t} [R_{t+1}(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1})] E_{\tilde{s}_{t+1}, m_t | s_t, x_t} [u'(c_L(\tilde{s}_{t+1}, \tilde{m}_{t+1}, B_{t+1}))]}}^{\text{Risk Premium}}
 \end{aligned} \tag{5}$$

where R is the binary repayment function and $u'(\cdot)$ is the lenders' marginal utility. Notice that the risk premium will always be positive, since $u'(\cdot)$ is a decreasing function and c_L is an increasing function of debt repayment. Thus, the covariance is negative and the overall term is positive.

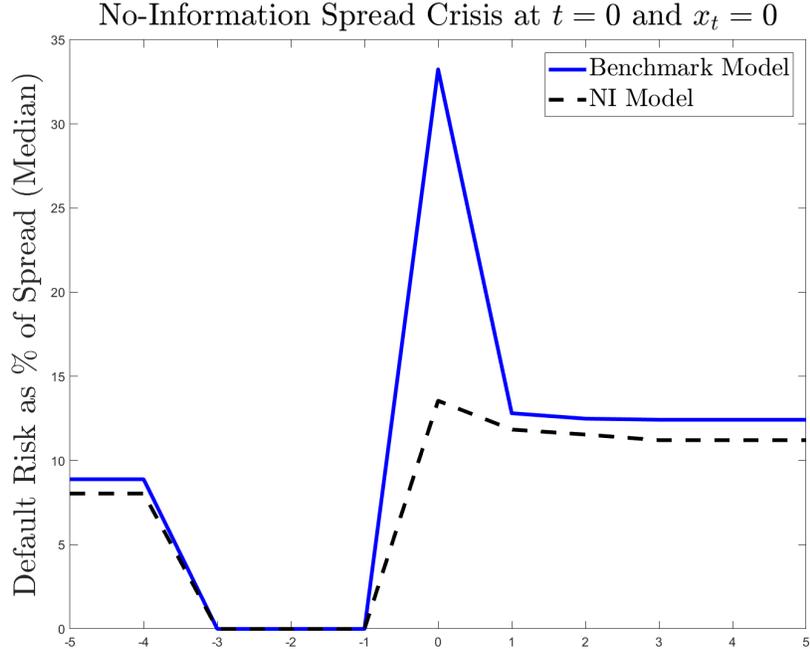
When lenders pay the cost to observe normally unobserved states, they learn more about the realization of those shocks. In particular, the conditional volatility $\sigma_{m|x}$ shrinks. This pushes down the risk premium term in Equation 5 by reducing the absolute value of the covariance. This occurs even in the case when $x_t = 0$ and the signal merely supports the prior of the lenders, which is that $m_{t+1} = 0$.

What does this imply for default risk inference? While default risk is high during a crisis, so too is the lenders' ability to learn and contain that risk since they are acquiring more precise signals about unobserved shocks. This implies that their effective risk-aversion is lower and thus that the risk premium comprises a relatively smaller share of the spread during a crisis than during normal times. Consequently, if an econometrician were not to take this into account, instead employing a more standard sovereign default model with constant investor attention to infer default risk from spread data, she would *underestimate* default risk during crises: She would assume the risk premium to be higher than it actually was.

Our model allows us to quantify this compositional shift in the risk spread. To do so, we construct an artificial, non-equilibrium no-information price schedule, which prices the exact same default risk as the benchmark but under the assumption that $\kappa = \infty$. We

²⁵This decomposition follows from a first-order approximation. Derivations can be found in [Appendix C](#).

Figure 10: Spread Decomposition During Crises



then compare simulated spreads from the benchmark model to this alternative spread series around a typical crisis.

In particular, we consider the following event study: We isolate crises *not* driven unobserved shocks by isolating the top 2.5%ile of the spread change distribution of the no-information counterfactual spread series, which does not price m_{t+1} in any capacity. We further condition on such crises in which $x_t = 0$, and thus the benchmark spread is not unusually high or low as a result of learning, since in these cases informed investors had their prior beliefs confirmed. The only difference between the two spread series during such events should be the risk premium: In the benchmark model it will be lower than in the no-information counterfactual, since more information is acquired during a crisis in the benchmark model.

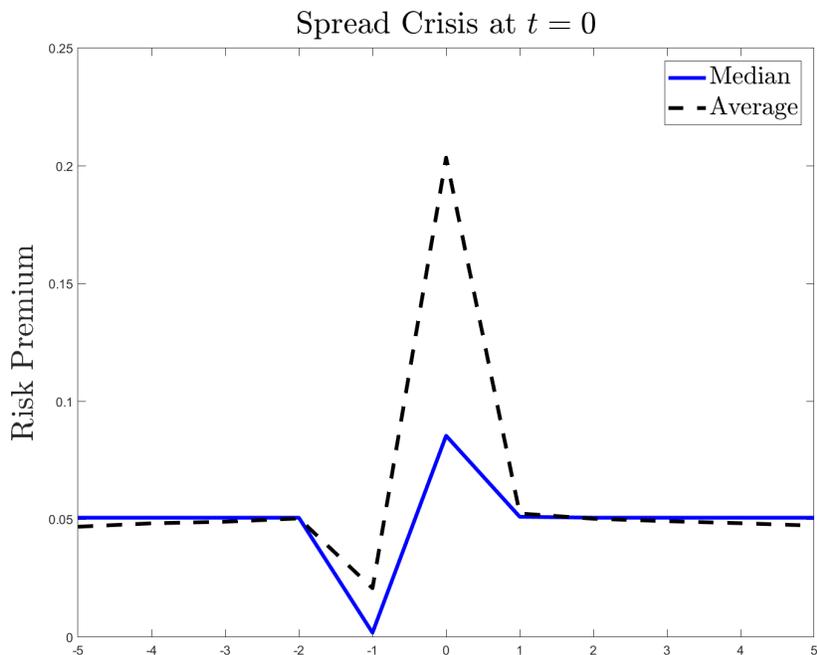
We isolate all such events in a simulation with length 1.5 million periods, and we compute the median share of default risk as a fraction of the total spread. The stochastic impulse-response function can be found in Figure 10. During a crisis, the benchmark model puts this figure at 33.2%, while the no-information counterfactual puts this same figure at 13.6%,

a difference of nearly 20 percentage points. The average difference between the series in non-crisis times, on the other hand, is a mere .85 percentage points, an order of magnitude smaller. Thus, default risk *as a share of the total spread* increases by substantially more during crises than a model without costly information acquisition would suggest.

This default-risk composition is particularly interesting also since it follows from the fact that risk-premia depend on *country-specific* states. Thus, it cannot be controlled for using global metrics, such as the CBOE VIX or the P/E ratio, as is often done (Aguiar et al. [2016a] or Bocola and DAVIS [2016]). Rather, our theory suggests that in order to accurately assess default risk, some metric of forecaster/investor information acquisition must be controlled for.

It is worth noting that this offsetting effect is not so strong as to reverse the cyclicity of risk premia. Risk premia remain countercyclical in our model. After all, default risk is high during crises, and thus the expected repayment is low, which drives up the risk premia. This can also be seen in Figure 11, which gives the average and median risk premium during a typical crisis. We can see that it is lower before the crisis, i.e., when times are good and growth is high, and jumps substantially during the crisis itself.

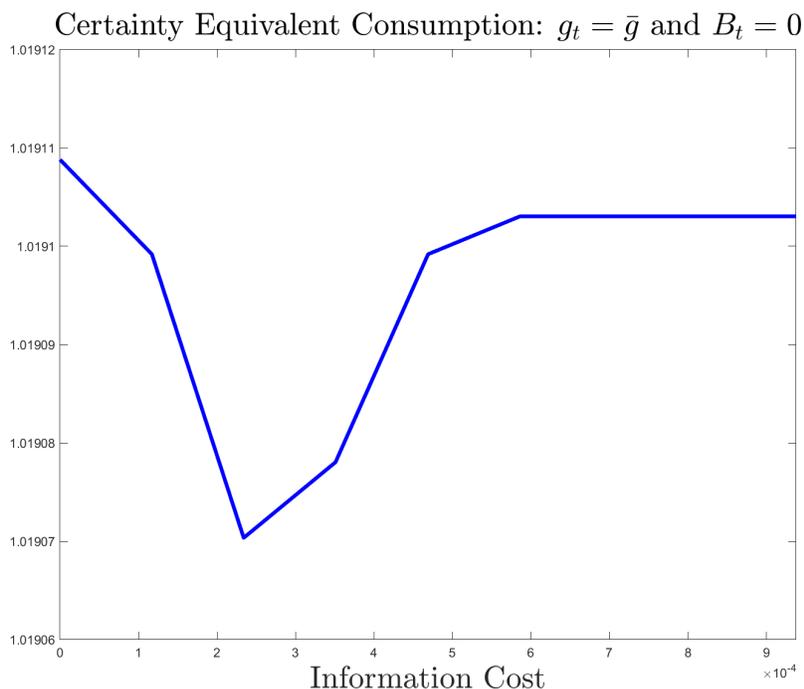
Figure 11: Risk Premium During Crises



3.4.3. Transparency

The last result we highlight regards the benefits and costs of transparency in light of our model mechanism. Costly information acquisition can be interpreted in many ways in the context of our model. One way the unit information cost κ can be understood as is the level of transparency that a sovereign has about its domestic affairs and finances. Interpreted this way, our model can also offer some interesting insights for sovereign transparency.

Figure 12: Sovereign Welfare Across Information Costs



Our model suggests that transparency is a double-edged sword. When there is zero transparency i.e. information is infinitely costly, lenders always demand a risk premium for the unobserved shocks, especially during crisis times. This makes it more expensive for the sovereign to borrow and service debt. Having more transparency will benefit the sovereign by lowering the risk premium. On the other hand, when the sovereign is fully transparent, although risk premium for the unobserved shocks disappears, it is replaced by substantial spread volatility, since now what used to be unobserved shocks are always reflected in bond prices. This will hurt the risk-averse sovereign. This welfare cost of public information

precision differs from that of [Morris and Shin \(2002\)](#), which falls instead on investors and results from an overweighting of public information in coordination games.

Therefore, transparency brings about two forces: There is the benefit of lower risk premia, but also the cost of higher price volatility. To illustrate this insight, we show the sovereign's welfare levels along different information costs in [Figure 12](#), which is evaluated at the steady state growth and at zero debt.²⁶ We can see that as the information cost decreases from the highest end, the sovereign's welfare first decreases in response to the higher volatility; but once the cost gets low enough, its welfare eventually reverses course and increases sharply as we approach full information. In this region, the enjoyment afforded by lowered borrowing costs overtakes the excessive volatility and the sovereign is better off.

The model thus suggests that there may be a period of pain associated with a country undertaking greater transparency measures before it reaches the welfare-maximizing level of full information.

4. Conclusion

Costly information acquisition plays an important role in the pricing of sovereign risk. We constructed and calibrated a structural model of endogenous default and information acquisition, using credit rating publication data to identify information costs.

We demonstrated that costly information acquisition generates country-specific time-varying volatility in sovereign bond spread; implies a compositional shift in that spread during crises; and highlights both the benefits and costs for emerging markets' welfare with regard to increasing transparency.

Possible extensions to our framework could include rollover crises in the vein of [Cole and Kehoe \(1996\)](#), long-maturity debt ([Hatchondo and Martinez \[2009\]](#) or [Chatterjee and Eyigungor \[2012\]](#)), or persistent unobserved shock processes. The intuition of our results

²⁶This is the welfare metric used by [Chatterjee and Eyigungor \(2012\)](#). The valley-shape pattern remains the same if we evaluate welfare at the ergodic mean of the debt-to-GDP ratio instead of zero debt.

would not change with any of these extensions, though the quantitative results may be affected.

References

- Aguiar, Mark and Gita Gopinath**, “Defaultable Debt, Interest Rates, and the Current Account,” *Journal of International Economics*, 2006, 69 (1), 64–83.
- , **Satyajit Chatterjee, Harold Cole, and Zachary R. Stangebye**, “Quantitative Models of Sovereign Debt Crises,” *Handbook of Macroeconomics, Volume II*, 2016.
- , —, —, and —, “Quantitative Models of Sovereign Debt Crises,” *Handbook of Macroeconomics, Volume II (Forthcoming)*, 2016.
- Angeletos, George-Marios and Ivan Werning**, “Crises and Prices: Information Aggregation, Multiplicity, and Volatility,” *American Economic Review*, 2006, 96 (5), 1720–1736.
- Arellano, Cristina**, “Default Risk and Income Fluctuations in Emerging Economies,” *American Economic Review*, 2008, 98 (3), 690–712.
- Bacchetta, Philippe and Eric van Wincoop**, “Infrequent Portfolio Decisions: A Solution to the Forward Discount Puzzle,” *American Economic Review*, 2010, 100 (3), 870–904.
- Barber, Brad M. and Terrance Odean**, “All That Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors,” *Review of Financial Studies*, 2008, 21 (2), 785–818.
- Bassetto, Marco and Carlo Galli**, “Is Inflation Default? The Role of Information in Debt Crises,” *Mimeo*, 2017.
- Bernoth, Kerstin and Guntram B. Wolff**, “Fool the Markets? Creative Accounting, Fiscal Transparency, and Sovereign Risk Premia,” *Scottish Journal of Political Economy*, 2008, 55 (4), 465–487.
- Bi, Huixin and Nora Traum**, “Estimating Sovereign Default Risk,” *American Economic Review*, 2012, 102 (3), 161–166.
- Bloom, Nicholas**, “The Impact of Uncertainty Shocks,” *Econometrica*, 2009, 77 (3), 623–685.
- Bocola, Luigi**, “The Pass-Through of Sovereign Risk,” *Journal of Political Economy*, 2016, 124 (4), 879–926.
- and **Alessandro Dovis**, “Self-Fulfilling Debt Crises: A Quantitative Assessment,” *Mimeo*, 2016.
- Carlson, Mark and Galina B Hale**, “Rating Agencies and Sovereign Debt Rollover,” *Topics in Macroeconomics*, 2006, 6 (2).
- Catao, Luis A.V., Ana Fostel, and Sandeep Kapur**, “Persistent Gaps and Default Traps,” *Journal of Development Economics*, 2009, 89 (2), 271–284.
- Chari, V.V. and Patrick J. Kehoe**, “Hot Money,” *Journal of Political Economy*, 2003, 111 (6), 1262–1292.
- Chatterjee, Satyajit and Burcu Eyigungor**, “Maturity, Indebtedness, and Default Risk,” *American Economic Review*, 2012, 102 (6), 2674–2699.
- Cole, Harold L. and Patrick J. Kehoe**, “Models of Sovereign Debt: Partial Versus General Reputations,” *International Economic Review*, 1998, 39 (1), 55–70.
- and **Timothy J. Kehoe**, “A Self-Fulfilling Model of Mexico’s 1994–1995 Debt Crisis,” *Journal of International Economics*, 1996, 41 (3–4), 309–330.
- , **Daniel Neuhann, and Guillermo Ordonez**, “Debt Crises: For Whom the Bell Tolls,” *NBER Working Paper No. 22330*, 2016.

- Da, Zhi, Joseph Engelberg, and Pengjie Gao**, “In Search of Attention,” *Journal of Finance*, 2011, 66 (5), 1461–1499.
- Dow, James and Gary Gorton**, “Noise Traders,” *NBER Working Paper No. 12256*, 2006.
- Durdu, C. Bora, Ricardo Nunes, and Horacio Sapriza**, “News and Sovereign Default Risk in Small Open Economies,” *Journal of International Economics*, 2013, 91 (1), 1–17.
- Eaton, Jonathon and Mark Gersovitz**, “Debt with Potential Repudiation: Theoretical and Empirical Analysis,” *Review of Economic Studies*, 1981, 48 (2), 289–309.
- Fernández-Villaverde, Jesús, Pablo A. Guerron-Quintana, Juan Rubio-Ramirez, and Martin Uribe**, “Risk Matters: The Real Effect of Volatility Shocks,” *American Economic Review*, 2011, 101 (6), 2530–2561.
- Gervais, Simon, Ron Kaniel, and Dan H. Mingelgrin**, “The High-Volume Return Premium,” *The Journal of Finance*, 2001, 56 (3), 877–919.
- Grossman, Sanford**, “On the efficiency of competitive stock markets where traders have diverse information,” *Journal of Finance*, 1976, 31 (2), 573–585.
- and **Joseph Stiglitz**, “On the impossibility of informationally efficient markets,” *American Economic Review*, 1980, 70 (3), 393–408.
- Hatchondo, Juan Carlos**, “Asymmetric Information and the Lack of International Portfolio Diversification,” *Working Paper University of Rochester*, 2004.
- and **Leonardo Martinez**, “Long-Duration Bonds and Sovereign Default,” *Journal of International Economics*, 2009, 79 (1), 117–125.
- Holden, Steinar, Gisle James Natvik, and Adrien Vigier**, “Equilibrium Rating and Debt Crises,” *International Economic Review*, 2018, *Forthcoming*.
- Iara, Anna and Guntram B. Wolff**, “Rules and Risk in the Euro Area,” *European Journal of Political Economy*, 2014, 34, 222–236.
- Johri, Alok, Shahed K. Kahn, and Cesar Sosa-Padilla**, “Interest Rate Uncertainty and Sovereign Default Risk,” *Mimeo*, 2015.
- Justiniano, Alejandro and Giorgio E. Primiceri**, “The Time-Varying Volatility of Macroeconomic Fluctuations,” *American Economic Review*, 2008, 98 (3), 604–641.
- Kopits, George and Jon Craig**, “Transparency in Government Operations,” *IMF Occasional Paper No. 158*, 1998.
- Kyle, Albert S.**, “Continuous auctions and insider trading,” *Econometrica*, 1985, 53 (6), 1315–1335.
- Lizarazo, Sandra V.**, “Default Risk and Risk Averse International Investors,” *Journal of International Economics*, 2013, 89, 317–330.
- Lucas, Jr. Robert E.**, “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, 1972, 4 (2), 103–124.
- Manso, Gustavo**, “Feedback Effects of Credit Ratings,” *Journal of Financial Economics*, 2013, 109 (2), 535–548.
- Melino, Angelo and Stuart M. Turnbull**, “Pricing Foreign Currency Options with Stochastic Volatility,” *Journal of Econometrics*, 1990, 45 (1-2), 239–265.
- Mendoza, Enrique G. and Vivian C. Yue**, “A General Equilibrium Model of Sovereign Default and Business Cycles,” *Quarterly Journal of Economics*, 2012, 127 (2), 889–946.
- Morris, Stephen and Hyun Song Shin**, “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 1998, 88 (3), 587–597.
- and — , “Social Value of Public Information,” *American Economic Review*, 2002, 92 (5), 1521–1534.

- Passadore, Juan and Juan Pablo Xandri**, “Robust Conditional Predictions in Dynamic Games: An Application to Sovereign Debt,” *Job Market Paper*, 2015.
- Poterba, James M. and Kim Rueben**, “State Fiscal Institutions and the US Municipal Bond Market,” *Chapter in NBER Book: Fiscal Institutions and Fiscal Performance*, 1999.
- Pouzo, Demian and Ignacio Presno**, “Sovereign Default Risk and Uncertainty Premia,” *SSRN Working Paper No. 2710469*, 2015.
- Sandleris, Guido**, “Sovereign Defaults: Information, Investment and Credit,” *Journal of International Economics*, 2008, *76* (2), 267–275.
- Seasholes, Mark S. and Guojun Wu**, “Predictable Behavior, Profits, and Attention,” *Journal of Empirical Finance*, 2007, *14* (5), 590–610.
- Sedlacek, Petr**, “Growth and Uncertainty over the Business Cycle,” *CEPR Working Paper No. 11296*, 2016.
- Seoane, Hernan D.**, “Time-Varying Volatility, Default and the Sovereign Risk Premium,” *Mimeo*, 2015.
- Shannon, Claude E.**, “Channels with Side Information at the Transmitter,” *IBM Journal of Research and Development*, 1958, *2* (4), 289–293.
- Sims, Christopher A.**, “The Implications of Rational Inattention,” *Journal of Monetary Economics*, 2003, *50* (3), 665–690.
- , “Rational Inattention: Beyond the Linear-Quadratic Case,” *American Economic Review*, 2006, *96* (2), 158–163.
- Sosa-Padilla, Cesar**, “Sovereign Defaults and Banking Crises,” *MPRA Paper No. 41074*, 2012.
- Stangebye, Zachary R.**, “Dynamic Panics: Theory and Application to the Eurozone,” *MPRA Working Paper No. 69967*, 2015.
- Van Nieuwerburgh, Stijn and Laura Veldkamp**, “Information Immobility and the Home Bias Puzzle,” *Journal of Finance*, 2009, *64* (3), 1187–1215.
- Veldkamp, Laura**, *Information Choice in Macroeconomics and Finance*, Princeton University Press, 2011.

Appendix A. Proof of Isomorphism Between Preference and Cost Shocks

Notice that under log-preferences, we can express the value of default as

$$V_{D,t}(s_t, m_t) = \log(\hat{Y}_t) + m_t + \beta E_{\tilde{s}_{t+1}|s_t} [\phi V_t(\tilde{s}_{t+1}, 0, 1) + (1 - \phi)V_{D,t}(\tilde{s}_{t+1}, 1)]$$

where $\hat{Y}_t = Y_t e^{-\psi}$ is output net of default costs. Now we can define $\hat{V}_{D,t}(s_t) = V_{D,t}(s_t, m_t) - m_t$, where $\hat{V}_{D,t}(s_t)$ is the value of default when the only cost of defaulting is $\phi(g_t)$.

With this notation, the default decision can be written as

$$d_t(m_t, s_t, B_t) = \mathbf{1}\{V_{R,t}(s_t, B_t) < \hat{V}_{D,t}(s_t) + m_t\}$$

In this isomorphic environment, m_t becomes a preference shock to the utility of defaulting.

Appendix B. Solving the Model

We consider the case of CRRA preferences for generality. We stationarize this model by dividing the sovereign resource constraint in every period by Y_t , and the value functions by $Y_t^{1-\gamma}$. This delivers a convenient recursive structure which is independent of both time and the level of output. We denote $b' = B_{t+1}/Y_t$ and $c = C_t/Y_t$.

$$v_R(g, b) = \max_{b' \geq 0} E_{\tilde{m}', \tilde{x}} \left[\frac{c(\tilde{x})^{1-\gamma}}{1-\gamma} + \beta E_{\tilde{g}'|g} e^{\tilde{g}'(1-\gamma)} v(\tilde{g}', b', \tilde{m}') \right]$$

s.t. $c(\tilde{x}) = 1 - b e^{-g} + q(g, b', \tilde{x}) b'$

The value of default is scaled similarly, yielding

$$v_D(g, m) = \frac{e^{(-\psi+m)(1-\gamma)}}{1-\gamma} + \beta E_{\tilde{g}', \tilde{m}'|g} \left[\phi e^{\tilde{g}'(1-\gamma)} v(\tilde{g}', 0, 0) + (1 - \phi) e^{\tilde{g}'(1-\gamma)} v_D(\tilde{g}', 0) \right]$$

This stationarization implies that we can express the default policy function using only stationarized model objects, since Y_t does not influence the default decision once g_t is known.

$$d(m, g, b) = \mathbf{1}\{e^{-g(1-\gamma)} v_R(g, b) < e^{-g(1-\gamma)} v_D(g, m)\}$$

The benchmark model with sovereign's log-preferences will simply be the limiting case as $\gamma \rightarrow 1$. This will imply that the stationarized model will feature log flow utility and that the impact of g on the effective discount factor vanishes.

The forecaster's problem is already stationarized, since it deals only with the distributions. We can stationarize the lenders' problem as well under the assumption that $w_t = wY_t$.

$$\begin{aligned} \max_{b'_i} E_{\tilde{m}', \tilde{g}' | x_i, g} \left[\frac{c_i^{1-\gamma_L}}{1-\gamma_L} \right] \\ \text{s.t. } c'_i = (w - b'_i q)(1+r) + b'_i [1 - d(\tilde{m}', \tilde{g}', b')] \end{aligned}$$

We solve the model using a discretized grid over the state space: We approximate the growth process using a Tauchenized Markov process with 25 grid points; we use a uniform grid over debt levels from $[0, .75]$ of 201 points; and we discretize m_t , x_t , and $\rho_{m_x, t+1}$ across 11 points each. Relevant model moments remain virtually unchanged when we increase the size of this grid along any dimension.

Appendix C. Spread Decomposition

For simplicity of notation, we drop the time subscripts and work with the recursive notation. We first consider the risk-neutral case. Here, the lenders' first-order condition together with the market clearing condition implies the following pricing expression:

$$q(B' | s, x) = \frac{E_{\tilde{s}', \tilde{m}' | s, x} [R(\tilde{s}', \tilde{m}', B')]}{1+r}$$

where $R(\tilde{s}', \tilde{m}', B')$ is the binary repayment function of the sovereign. The effective interest rate is $\hat{r}_{RN} = 1/q$. Taking a log of the previous expression yields:

$$\begin{aligned} \log(1 + \hat{r}_{RN}) &= \log(1+r) - \log(E_{\tilde{s}', \tilde{m}' | s, x} [R(\tilde{s}', \tilde{m}', B')]) \\ \implies \hat{r}_{RN} - r &\approx -E_{\tilde{s}', \tilde{m}' | s, x} [\log(R(\tilde{s}', \tilde{m}', B'))] \\ \implies \text{sprd}_{RN} &\approx E_{\tilde{s}', \tilde{m}' | s, x} [1 - R(\tilde{s}', \tilde{m}', B')] \end{aligned}$$

The second line follows from the expectation of a first-order approximation around the mean. The third follows from the first-order approximation $x \approx \log(1+x)$ applied to $R(s, m, B) = 1 - D(s, m, B)$, assuming default risk to be relatively small.

To compute the overall spread, we use a similar strategy. Note the first-order necessary condition is

$$q(B' | s, x) = \frac{E_{\tilde{s}', m' | s, x} [R(\tilde{s}', \tilde{m}', B') u'(c_L(\tilde{s}', \tilde{m}', B'))]}{(1+r) E_{\tilde{s}', m' | s, x} [u'(c_L(\tilde{s}', \tilde{m}', B'))]}$$

i.e., the ratio of the expected marginal utilities in repayment states over the expected marginal utilities in all states. Following a similar procedure as before, we arrive at

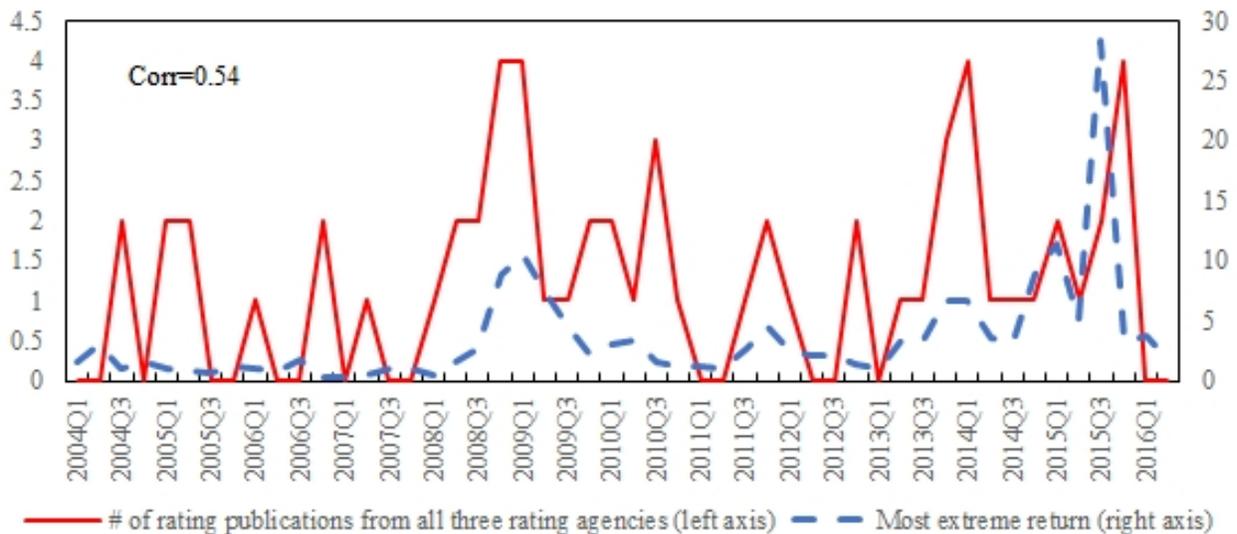
$$\log(1 + \hat{r}) - \log(1 + r) = -\log(E_{\tilde{s}', \tilde{m}|s, x}[R(\tilde{s}', \tilde{m}, B')]) - \log\left(1 + \frac{\text{cov}(R(\tilde{s}, \tilde{m}, B'), u'(c_L(\tilde{s}, \tilde{m}, B')))}{E[R(\tilde{s}, \tilde{m}, B')]E[u'(c_L(\tilde{s}, \tilde{m}, B'))]}\right)$$

If both the default risk and the covariance expression are relatively small (close to zero), then a first-order approximation yields Equation 5.

Appendix D. Attention Measure Robustness

We explore the correlation of our chosen attention metric, total number of sovereign credit rating publications, with another common metric of investor attention in the literature, extreme daily returns (Barber and Odean [2008]). In particular, we use daily return data on the stripped sovereign spread from JP Morgan’s EMBI database and compute the largest quarterly return in absolute value. We then compare this series to the rating publication series in Figure D.13. While they do not line up perfectly, there is high co-movement ($\rho = 0.54$) and thus our metric lines up reasonably well with this other proxy. Most of this co-movement is driven by extreme negative returns, which is consistent with our theory since investors pay attention more during crises.

Figure D.13: Extreme Returns and Total Number of Credit Rating Publications, Ukraine



Another alternative considered in the literature are daily trade volumes. However, such data is difficult to acquire and may not exist for sovereign bonds. This is the case for

our benchmark choice of Ukraine, for which almost all trades are executed over-the-counter instead of in an exchange.

Finally, given that the relevant literature (e.g., [Da et al. \[2011\]](#)) has also used Google Trend’s search volume index, we also check the correlation between our total rating publication metric and the search volume index on search term “Ukraine” or “Ukraine IMF” over our sample period. Both correlations are above 40%.