CONTENTS

1

Introduction

1.1 Overview

Sovereign debt markets are prone to frequent and serial defaults involving prolonged and costly periods of restructuring, interspersed with tranquil periods in which countries accumulate more debt as prelude to the next debt crisis. This pattern is most prominent in low-income and emerging markets, but, as the euro-area's crisis in the 2010s demonstrated, extends to developed countries as well.

Conceptually, the distinguishing feature of sovereign debt markets is the limited ability for creditors to enforce contracts. Seminal papers such as Eaton and Gersovitz (1981) and Bulow and Rogoff (1989a) are concerned with why governments repay external debt at all. However, lack of commitment to contracts is only the first of many frictions, a partial list of which includes: (i) no or limited state contingency in debt contracts; (ii) an incentive to "dilute" legacy creditors arising from the inability to contract/commit to future debt policies; (iii) large deadweight costs to default and lengthy renegotiations; (iv) vulnerability to self-fulfilling crises; (v) currency mismatch; and (vi) political economy distortions that drive a wedge between the welfare of the citizenry and the objectives of the political incumbents. In the subsequent chapters, we explore the connection between these underlying frictions and the observed behavior of sovereign debt markets. In particular, we emphasize how these frictions influence the government's *equilibrium choices*, including whether to save or borrow, default or repay, dilute, repurchase, or restructure existing bonds, and whether to issue short-term or long-term bonds. We also show how these choices mitigate or exacerbate the inefficiencies stemming from the underlying frictions.

The approach we take emphasizes analytical clarity as a prelude to quantitative implementation. For each topic, we begin with tractable analytical models that isolate the core friction of interest. To the extent possible, we characterize the equilibria of a model by showing an equivalence to an optimization problem, despite the fact that the equilibria may not be constrained efficient. These pseudo-planning problems avoid the need to analyze a complicated fixed point, and instead allow us to study the much more transparent problem of choosing an allocation subject to constraints. After analyzing the simple analytical environments, we turn to the richer quantitative models used in the literature, leveraging the insights of the analytical models to illuminate the "black box" that encloses the quantitative models.

Chapter 2 begins with the core friction of limited commitment. In particular, we study the canonical model of a risk-neutral representative lender insuring a risk-averse government, subject to the constraint that the government can always renege on promised payments. To isolate the role of limited commitment, the contract can be state contingent. The punishment for deviation is allowed to be quite general, and we show how to nest both the canonical punishment of financial autarky, as in Eaton and Gersovitz (1981), as well as the ability to save but not borrow, as in Bulow and Rogoff (1989a). Lack of commitment limits the extent of insurance, but also provides an incentive for the government to save (or reduce debt). In terms of allocations, lack of commitment tilts the path of consumption in favor of delaying consumption into the future, the well-known "backloading" of incentives.

While Chapter 2 focuses on using financial markets for risk sharing and inter-temporal substitution, another important role for international financial flows is to fund investment. This is the focus of Chapter 3. That chapter is motivated by a fact documented in Gourinchas and Jeanne (2013) and Aguiar and Amador (2011) that high-growth emerging markets have net *outflows* of capital, a pattern generated by private net inflows that are more than offset by public outflows. Chapter 3 uses a production economy overseen by a government that lacks commitment to both debt payments and capital taxation/expropriation. This economy is vulnerable to "debt overhang." Specifically, a heavily indebted government has a greater incentive to default and expropriate capital. This, in turn, depresses private investment. Thus public debt and private investment move in opposite directions, generating the empirical pattern described above.

Government policy is set by political incumbents that are present biased and rotate in and out of power stochastically. The political economy frictions are shown to be the key determinant of both public capital flows and the speed at which the economy converges to the frontier. The deeper the political frictions, the less political incumbents are willing to reduce debt, and the slower the pace of private capital accumulation. Through the lens of the model, the heterogeneity in rates of economic convergence seen in the data is driven by the heterogeneous severity of political myopia.

Chapters 2 and 3 focus on the core friction of limited commitment, abstracting from other restrictions on financial markets. This baseline provides a useful starting point, as the remaining frictions have clear implications for whether the incentive to save is enhanced or diminished. The introduction of market incompleteness is a good case in point.

In Chapter 4, we show that market incompleteness provides an additional incentive to save. We initially do this in an environment in which consumption is deterministic (absent default), highlighting that more than precautionary savings is at work. The source of uninsurable risk is the severity of punishment in the case of default, with the costs of default following a simple *iid* process. As markets are incomplete, the government may decide to default in equilibrium. In doing so, there is a deadweight loss, as ideally the government and creditor would like to contract on a state-contingent payment rather than default. This loss succinctly proxies for disruption of economic activity, both real and financial, that often coincides with default status. We show that the larger is this deadweight cost, the greater the incentive for the government to reduce debt.

Why this is the case provides a useful guide to some additional results. The government can always opt not to default at a given point in time—that is, default is *strategic* rather than due to an inability to pay. At the moment of default, despite the deadweight costs, the government strictly prefers default to repayment. So why avoid a situation in which it gets to choose the best of two options? The answer is in how bonds are priced. Competitive bond markets ensure the lenders break even in expectation. Thus, contracted payments include compensation for the entirety of lenders' expected losses; however, if there are deadweight costs of default, the government captures only a fraction of lenders' losses when it defaults. The expected wedge between what the government pays to the lenders absent default and what it gains in default varies with the probability of default, and thus, as we show, varies with the level of debt. The deadweight costs of default, via the equilibrium bond price schedule, provide the government with the incentive to save above and beyond that provided by limited commitment in a complete-markets environment.

This raises the question of whether the equilibrium incentive to save is "strong enough." That is, would a planner (or the lenders) prefer an alternative fiscal policy. We show that short-term bonds ensure the government and lenders agree on the government's debt policy. That is, it does not matter whether debt issuance (conditional on strategic default) is chosen by the government, the lenders, or a planner maximizing the weighted sum of the respective utilities.

In Chapter 5 we show how to extend this result to the richer environments with consumption risk used in the early one-period-bond quantitative models of Aguiar and Gopinath (2006) and Arellano (2008). Along the way, we establish that the competitive equilibria in the one-period bond model can be computed using a pseudo-planning problem. Importantly, the planning problem includes an "implementability constraint" that ensures the planner's chosen allocation respects the absence of state-contingent contracts in the competitive equilibrium. In fact, the planning problem can be computed by iterating on a Bellman operator that is a contraction, establishing existence and uniqueness of equilibrium in the one-period bond model.

While the canonical one-period model is a useful starting point, the efficiency properties described above and in Chapters 4 and 5 make it ill suited to address two other frictions that appear to be empirically relevant; namely, self-fulfilling debt crises and debt dilution. We take up the first of these in Chapter 6. Following Cole and Kehoe (2000), we show that relaxing an implicit assumption in the literature building on Eaton and Gersovitz (1981)—that the government commits to repayment of existing bonds before auctioning new debt—opens the door to self-fulfilling debt crises. Chapter 6 discusses how and when the canonical Cole-Kehoe model can generate a failed auction and induce a self-fulfilling default. Building on Aguiar, Chatterjee, Cole, and Stangebye (2019b), the chapter also discusses that the environment admits several other plausible self-fulfilling scenarios, including sudden stops and overborrowing during a crisis.

Debt dilution is the focus of Chapter 7, which explores debt dynamics in an Eaton-Gersovitz model with long-term bonds. Again, the anlaysis begins with a stripped-down analytical model before moving on to the richer environments used for quantitative analysis. The chapter sheds light on why equilibrium prices fail to provide strong incentives to avoid the costs of default when bonds are rolled over infrequently. In equilibrium, this manifests itself as inefficiently high levels of new borrowing when the government has long-term bonds outstanding.

We also show that the same force can lead to multiplicity, what we term "self-fulfilling debt dilution." This multiplicity is distinct from the rollover crises analyzed in Chapter 6. The latter is static in nature; that is, multiplicity arises even when holding future equilibrium behavior fixed. Self-fulfilling debt dilution is inherently dynamic. Creditor beliefs about future borrowing generate a long-term-bond price schedule that induces a fiscal policy that validates those beliefs. There may be a "borrowing" equilibrium that features high debt and likely default as well as a "saving" equilibrium in which the government deleverages over time to reduce the probability of default. Which equilibrium is realized depends on self-fulfilling beliefs.

The inefficiency generated by the presence of legacy debt raises the question of whether the government should repurchase long-term bonds in equilibrium. We show that this is never optimal for the government to do via arms-length market transactions. However, a coordinated restructuring that avoids holdouts can be Pareto improving.

With these analytical results in hand, we turn to the quantitative models with long-term debt popular in the literature, such as those proposed by Chatterjee and Eyigungor (2012b) and Hatchondo and Martinez (2009). The combination of stochastic deadweight costs of default and the incentive for debt dilution combine to generate patterns that match key empirical moments.

1.2 Empirical Landscape

Before delving into analytical and quantitative models, it is useful to survey the empirical landscape. This will establish what the theory is trying to explain, as well as clarify what is included and what is left out of the models.

Incomplete Contracts. The vast majority of sovereign debt contracts promise fixed, non-contingent payments. There are some examples of payments linked to GDP or commodity prices. For example, Argentina's 2005 restructuring included warrants linked to GDP growth. The Brady plan included bonds with payments linked to oil revenues or GDP, as well. Nevertheless, these are the exceptions rather than the rule.

However, the process of default and renegotiation makes sovereign debt implicitly state contingent. Grossman and Van Huyck (1988) model de facto state-contingent contracts by distinguishing "excusable default" from debt repudiation. Bulow and Rogoff (1989b) introduce a post-default bargaining protocol to generate ex post state-contingent payments.

Sovereign debt contracts also do not typically make promises regarding future fiscal policy or have rates that vary with changes in default probabilities. This leaves bondholders vulnerable to risk induced by future government decisions. Hall and Sargent (2020) discuss US Liberty Bonds that allowed holders to exchange their initial bonds for those with a higher coupon if the Treasury auctioned additional debt in the future at more favorable terms. The same authors also characterize various put and call options that have been included in US bonds to insure bondholders or the US Treasury against future price changes. Such clauses are not typical in sovereign debt markets.

Sovereign bonds are also issued on a “pari passu” (Latin for equal footing) basis. While open to interpretation, this means there is no explicit seniority across bond issuances (Wright, 2014). This makes selective default—that is, defaulting on bond series A while repaying series B—difficult if not impossible for bonds issued in the same jurisdiction (Schumacher, Trebesch, and Enderlein 2012).

Sovereigns rely on a variety of instruments and intermediaries to tap global financial markets. Using World Bank data, Schlegl, Trebesch, and Wright (2019) document the mix of official and private lending to emerging and developing economies since 1980. A large fraction of credit is extended by governments or multilateral institutions, such as the World Bank and International Monetary Fund. Private sector lending was primarily bank loans in the 1970s and 1980s, but the Brady Plan, and the associated “Brady bonds” introduced in 1989 to resolve the Latin American debt crisis of the 1980s, helped spur a switch toward bond finance. More precisely, it was a return to bond finance, which had been relatively common in the nineteenth and the first half of the twentieth centuries (see Meyer, Reinhart, and Trebesch 2019).

For the most part, the models of this book involve bonds traded in competitive markets. To the extent banking relationships are competitive and non-exclusive, the insights carry over to bank lending as well. We do not consider models of non-competitive lending, such as government-to-government or multilateral agency loans, or long-term exclusive relationships.

Legal Enforcement and Reputation. Sovereign debt can be governed by domestic law or foreign law. For example, Argentine bonds issued in New York are subject to US courts, while those issued in Buenos Aires are governed by Argentine law. If the consequences of default were purely reputational, the location of issuance should not be a major factor. However, it is well documented that it matters for pricing and payment under whose jurisdiction a bond is issued. For example, at the time of its restructuring in 2012, Greece had bonds issued under local law as well as English law.[1] The behavior of yields leading up to the restructuring as well as the treatment of creditors ex post differed depending on whether foreign or domestic law was operative (Choi, Gulati, and Posner 2011; Zettelmeyer, Trebesch, and Gulati 2013). Schumacher et al. (2012) review how foreign courts enforce sovereign debt contracts. Sovereign immunity limits the scope for punishment or restitution available to creditors in the event of default. In some cases, foreign courts have frozen export revenue or assets on behalf of creditors. However, the main threat is to disrupt issuance or payments of new bonds until an agreement is reached with existing bondholders.

While legal considerations play a role in sovereign debt markets, there is evidence that reputation (as captured, for example, by the previous history of default) matters as well. Cruces and Trebesch (2011) document how the history of default and repayment is correlated with future spreads.[2] Related to this, George Hall and Thomas Sargent have a body of work that documents how the United States built up its reputation as a reliable debtor in order to establish deep and reliable markets for government bonds.[3]

In the models described in the remaining chapters, any notion of "reputation" is captured by a default state variable, which potentially involves direct costs as well as financial autarky. But richer notions of reputation exist in the literature. Such notions rely on the financial market's belief about a hidden government type. Naturally, this market belief (which corresponds to the country's reputation) evolves over time and is affected by the country's history of default and repayment. The papers by Cole, Dow, and English (1995), Alfaro and Kanczuk (2005), D'Erasmo (2011), Egorov and Fabinger (2016), Perez (2017), and Amador and Phelan (2020) studied models where the market has imperfect information about the government's type and learns from the government's equilibrium actions. There is a related work that focuses on the market's asymmetric information and learning about a hidden productivity state of the country. For examples of these, see Sandleris (2008), Phan (2017), and Dovis (2019).

1. For more on the data and modeling of selective default between domestic and foreign-issued bonds, see Erce and Mallucci (2018).

2. See also Meyer et al. (2019) and Asonuma (2016).

3. See, for example, Hall and Sargent (2014, 2019, 2020).

Currency. As a rule of thumb, sovereign bonds issued in financial centers like New York or London are typically issued in global currencies, such as dollars, pounds, or euros. Conversely, bonds issued in domestic financial markets are typically denominated in domestic currency, although there are notable exceptions of domestically issued bonds paying out in foreign currency.[4] The fact that external borrowing by developing countries is denominated in foreign currency was famously explored by Eichengreen and Hausmann (1999). Recently, there is a growing trend toward more foreign participation in domestic sovereign bond markets in emerging markets.[5]

Debt Intolerance. In a world in which many developed economies have public debt-to-GDP ratios that often exceed 100%, it is striking how low external debt-to-income ratios are for emerging markets at the time of default. For example, Argentina's 2001 default occurred when public debt was roughly 50% of annual GDP. Reinhart, Rogoff, and Savastano (2003) document that many debt crises occur at relatively low debt levels, coining the term "debt intolerence" to describe this phenomenon.

Maturity. Broner, Lorenzoni, and Schmukler (2013) and Arellano and Ramanarayanan (2012) document maturity choice for emerging market governments. In the Broner et al. (2013) sample, which covers the majority of external, US-dollar-denominated debt issuances for a sample of eleven emerging markets between 1993 and 2009, debt issuances ranged between one-year and thirty-year maturities, with the 75th percentile centered around ten years. For the same period, they report the 75th percentile of US Treasury issuances of five years. Based on Broner et al. (2013), Chatterjee and Eyigungor (2012a) calibrate their model to a maturity of five years to match the median issuance of Argentina. Arellano and Ramanarayanan (2012) document that the average maturity of issuances for Argentina, Brazil, Mexico, and Russia ranged between nine and twelve years for roughly the same years.[6]

Both Broner et al. (2013) and Arellano and Ramanarayanan (2012) document that as spreads increase, governments shorten the maturity of issuances relative to issuance in relatively tranquil times. A similar pattern of issuances

4. For example, both Argentina and Russia issued domestic-law US dollar bonds as part of their restructurings. See Chamon, Schumacher, and Trebesch (2018).

5. See Burger and Warnock (2007); Arslanalp and Tsudo (2014), and Du and Schreger (2016a).

6. See the respective papers for details, including how defaults and restructuring are handled.

can be observed for the affected countries during the 2011–2012 European sovereign debt crisis (Bocola, Bornstein, and Dovis 2019). Broner et al. (2013) attribute this behavior to an increase in the risk premium charged to longer-maturity bonds during crises, while Arellano and Ramanarayanan (2012) appeal to the balance between hedging and incentives related to the discussion in Chapter 7.

Default. Tomz and Wright (2007, 2013) construct a database of 176 sovereign entities spanning 1820 to 2012. They document that the unconditional probability of default is 1.8% per year. Excluding 1945–1980, a period with little external borrowing, the probability rises to 2.2% per year.

Meyer et al. (2019) collect data on ninety-one countries with debt issued in London or New York over the period 1815–2016. They identify 313 defaults or distressed restructurings, defined as exchanges of debt at a loss. Reinhart and Rogoff (2004) document that countries tend to default more than once (so-called serial defaulters). Moreover, the nineteenth and twentieth centuries contain a number of defaults by rich countries, indicating that default has not been exclusively a phenomenon of poor or developing economies.

The canonical models of default predict countries' default during recessions. Tomz and Wright (2007) find modest evidence in this regard. Roughly two-thirds of defaults begin when output is below trend. Average deviation of output from trend at the start of default is a modest -1.6%. The cross-sectional correlation between output and default status is only -0.08. While output is typically above trend in periods of non-default, there are many instances of sharp recessions without default.

Reinhart and Rogoff (2004) document that many debt crises are preceded by or coincide with banking crises. Bailouts as well as the drop in revenue associated with financial crises strain fiscal authorities. A notable example of this mechanism is Ireland's debt crisis after the global financial crisis of 2008. Ireland's debt-to-annual GDP was less than 25% in 2007. By the end of 2010, the crisis accompanied by a large banking system bailout had forced Ireland to seek assistance from the EU. Debt-to-GDP peaked at over 130% by 2013. More broadly, Reinhart and Rogoff (2009) document that government debt increases 86% in the three years following systemic financial crisis.

Restructuring and Haircuts. Once a country has defaulted on external creditors, they remain in default status until a settlement is reached with creditors. Tomz and Wright (2007) document in their sample that the mean length of default is 9.9 years and the median length is 6.5 years. There is substantial heterogeneity in the duration between initial default and settlement; the standard deviation of duration of default is 10.5 years. Debtors spend roughly 19% of the sample period in default status.

In most cases, the defaulting government and creditors come to an agreement in which creditors accept restructured bonds with different payment terms and maturities than the original bonds. That is, from the creditors' perspective, a default implies a loss due to delay, negotiation costs, and a potential write-down of debt obligations—the so-called haircut.

There are different methods of computing haircuts. The task is complicated by the need to assess the impact of maturity changes as well as changes in face value. The canonical approach was pioneered in Sturzenegger and Zettelmeyer (2008). Their method takes the implied yield, r, using the secondary market prices of the restructured bonds at the time of settlement. The haircut is then the present discounted value of the promised sequence of payments in the restructured bonds divided by the present value of the original bonds, both discounted at the rate r. By construction, the numerator is the secondary market price of a restructured bond. The denominator is the counterfactual market value of an original bond priced using the same yield as the new bonds. Thus, this ratio captures the loss (if any) due to the bond exchange relative to hypothetically keeping the original bonds to sell in the post-default market.

Cruces and Trebesch (2013) document in a sample of 180 restructurings spanning sixty-eight countries and the period 1978–2010 a mean haircut of 37% with a standard deviation of 27%. They find larger haircuts (mean 87%) in the subsample of highly indebted poor countries. They also find that longer default durations are associated with larger haircuts. Interestingly, Benjamin and Wright (2008) find that the size of haircut and length of default are uncorrelated with initial debt levels. However, Benjamin and Wright (2008) and Trebesch and Zabel (2017) document that a more severe decline in output post-default is associated with a larger haircut.

While default typically involves haircuts for creditors, it is not the case that default necessarily lowers the face value of debt. This reflects that restructured bonds tend to be longer maturity, and the losses are imposed by delaying payment into the future rather than reducing the face value of payments. Benjamin and Wright (2008) document that the median and average post-default debt-to-GDP ratio is 5% and 25% higher, respectively, than at the time of default. In the Cruces and Trebesch sample, only 57 of 180 cases involved reduction in face value.

As noted above, while in default status, governments are excluded from debt markets. Even after coming to a resolution with creditors, governments do not immediately access primary debt markets. Cruces and Trebesch (2013) look at the duration between default resolution and first loan or bond issuance (or new credit to public sector). They find that the average duration is 2.3 years when haircuts were less than 30%, but 6.1 years on average when haircuts were greater than 30%. Moreover, the spreads on newly issued bonds after re-entry are increasing in the size of the haircut imposed on the defaulted bonds.

Interest Rate Spreads. Sovereign bonds from emerging economies as well as some developed economies have yields greater than comparable bonds from the United States or Germany. This "spread" can be quite large; for example, at the height of the euro-area debt crisis in the summer of 2012, Spanish 10-year bonds had a yield of close to 7%, compared to German 10-year bonds, which had a yield of less than 2%. The spreads for emerging market bonds like those of Argentina can be even greater; based on data from the Emerging Market Bond Index (EMBI) for Argentina in the 1990s, Chatterjee and Eyigungor (2012b) target an average spread over three-month US T-bills of roughly 8%.

The high spreads compensate creditors for the risk of default. However, even including the losses due to the default, sovereign bonds return more on average than comparable US bonds (a "risk premium" or "excess return").[7] Meyer et al. (2019) document that between 1815 and 2016, a sample of sovereign bonds from sixty "serial defaulters"[8] had an annualized realized (including defaults and restructurings) mean return of 7%, or 4% more than comparable US or UK bonds.[9] For the modern period 1995–2016, investors in the bonds of serial defaulters earned a mean annual return of 9.3%, or 5.2 percentage points more than investors in US bonds.

Returns on sovereign bonds are also volatile. Meyer et al. (2019) document that the annual standard deviation of real returns in their 1815–2016 sample is 14%, more than double the volatility of US equities. The ratio of average spreads to standard deviation (the Sharpe ratio) is 0.37 for this

7. For sovereign debt models featuring risk premia, see Borri and Verdelhan (2011), Lizarazo (2013), Pouzo and Presno (2016), and Tourre (2017).

8. Serial defaulters are countries that defaulted at least twice in the sample period or that spent 20% of the years since 1815 (or independence) in default status.

9. Spreads prior to 1918 use 10-year UK bonds as a benchmark, and 10-year US bonds for the subsequent 100 years.

sample, compared to 0.32 for US equities, suggesting that as an asset class, the risk-return tradeoff is relatively favorable for sovereign bonds, despite the frequent defaults.

Sovereign bond spreads are correlated with major international global risk factors. For example, Longstaff et al. (2011) look at spreads on credit default swaps (CDS) for twenty-six countries spanning October 2000 to January 2010.[10] The spreads are correlated across countries: The first principal component explains 64% of the variation. Moreover, sovereign CDS spreads are more related to the US stock and high-yield markets than they are to local economic measures. The first principal component of spreads has a correlation with the US Stock Market of 0.61 and the VIX -0.75. Longstaff, Pan, Pedersen, and Singleton (2011) regress spreads on local variables as well as a number of global indicators. Local variables explain roughly 40% of total explained variation (R^2) in CDS spreads, with the remaining 60% accounted for by global factors.[11]

1.3 Goals

The literature on sovereign debt has exploded since Eaton and Gersovitz (1981). The repeated waves of debt crises since the early 1980s have driven much of this interest, as each crisis allows us to evaluate existing models in a new light as well as point out directions for future research. Coinciding with the theoretical and quantitative research has been a boom in empirical research, both on modern crises as well as historical episodes, only a fraction of which is outlined in the previous section.[12] The interplay of empirics and theory has been a tremendous strength of this literature.

10. Credit default swaps are insurance contracts that pay out if and when a debtor defaults or restructures at unfavorable terms. Thus, the premiums (or "CDS spreads") reflect the probability of default plus any risk premium. CDS spreads are correlated with, but not identical to, the spreads on the underlying bonds. Although we do not discuss them in this book, see Salomao (2017) for an analysis of the role CDS contracts played during the Greek debt crisis.

11. After documenting the negative covariance between spreads and output, Neumeyer and Perri (2005) build a model in which exogenous shocks to a country's borrowing rate partially drive the business cycle. Uribe and Yue (2006) provide a methodology to measure the relative impact of foreign versus domestic shocks in driving emerging market interest rates, as well as the subsequent feedback between foreign interest rate shocks and domestic output.

12. For surveys, see the handbook chapters Eaton and Fernandez (1995); Aguiar and Amador (2014); Aguiar et al. (2016). See also the textbook treatment in Chapter 13 of Uribe and Schmitt-Grohé (2017).

A number of books on sovereign debt have been published in recent years.[13] The bulk of this work has been focused on documenting empirical patterns and analyzing key historical episodes. There has been less focus on the theoretical modeling of debt and default, and we hope this work at least partially fills this gap. The book has elements of both a textbook and a monograph. The focus is what we have learned in our own research on sovereign debt. But, we do not work in a vacuum. We build on the large body of work that precedes us, and, perhaps more important, have benefited from a strong cohort of contemporaneous researchers. Our own research includes work on debt overhang, growth under limited commitment and political economy frictions, constrained (in)efficiency of debt markets, debt dilution, debt buybacks, indeterminacy of equilibrium and self-fulfilling crises, and maturity choice. These contributions constitute the bulk of the book, but, to the extent possible, we include results from other researchers that speak directly to or have inspired our own contributions.

A unifying thread throughout is a tractable analytical framework that isolates the theoretical mechanisms at work, after which, when relevant, we discuss how the analytical models help us understand the richer, quantitative models that have become popular in the literature. These latter models are useful to understand to what extent our models can quantitatively match key empirical moments, but often operate as a "black box." A primary goal is to bridge the gap between simple models and the quantitative literature.

The book can thus serve as an introduction to these themes, a companion work to those interested in quantitative sovereign debt research, and a synthesis of our own work. While the analysis is geared to the level of a graduate student, each chapter begins and ends with a non-technical overview of the key questions and core results. When appropriate, the concluding sections also summarize the policy implications of the respective chapter.

13. A partial list includes Sturzenegger and Zettelmeyer (2007); Tomz (2007); Reinhart and Rogoff (2009); Drelichman and Voth (2011) and Abbas, Pienkowski, and Rogoff (2019).

INDEX