Learning about Debt Crises

Radoslaw Paluszynski*

Department of Economics
University of Houston

First version: October 2015
This version: May 2021

Abstract

The European debt crisis presents a challenge to our understanding of the relationship between government bond yields and economic fundamentals. I argue that information frictions are an important missing element, and support that claim with evidence on the evolution of GDP forecast errors after 2008. I build a quantitative model of sovereign default where output features rare disasters and agents learn about their realizations. Debt crises coincide with economic depressions and develop gradually while markets update their expectations about future income. Calibrated to the Portuguese economy, the model replicates the comovement of bond spreads and output before and after 2008.

Keywords: Sovereign debt, disaster risk, Bayesian learning, long-term debt

JEL Classification Numbers: D83, F34, G15

*E-mail: rpaluszynski@uh.edu. This paper is based on Chapter 1 of my Ph.D. thesis at the University of Minnesota. I am grateful to Manuel Amador, Tim Kehoe and Motohiro Yogo for their continuous encouragement and support. I also thank Joao Ayres, Anmol Bhandari, Hyunju Lee, Ellen McGrattan, Bent Sørensen, George Stefanidis, Kei-Mu Yi, Pei Cheng Yu, three anonymous referees, and all participants of the Trade and Development workshop at the University of Minnesota for many helpful comments. I acknowledge the generous support of the Hutcheson-Lilly dissertation fellowship. The quantitative part of this paper was conducted using the resources of the Minnesota Supercomputing Institute. Consensus Economics Inc. owns the copyright to the Consensus Forecasts - G7 & Western Europe dataset which I use under a license agreement.

1 Introduction

The recent debt crisis in Europe has reopened the discussion about which factors put governments at risk of sovereign default. The weak correlation between interest rates on public debt and economic fundamentals of the southern European countries challenges the theoretical links established by a large body of research prior to 2008. This new evidence has led some researchers to revisit the hypothesis of self-fulfilling debt crises (Aguiar et al., 2020). Other economists, motivated by the same observations, have argued that the European episode was driven by external factors such as intra-EU politics (Brunnermeier, James and Landau, 2016). In this paper I propose a quantitative model of the European debt crisis based on idiosyncratic income fluctuations which feature disaster risk and information frictions.

Figure 1 reviews the comovement between government interest rates and external debt, and their relationship with economic fundamentals. The bond spreads[1] were negligible since the

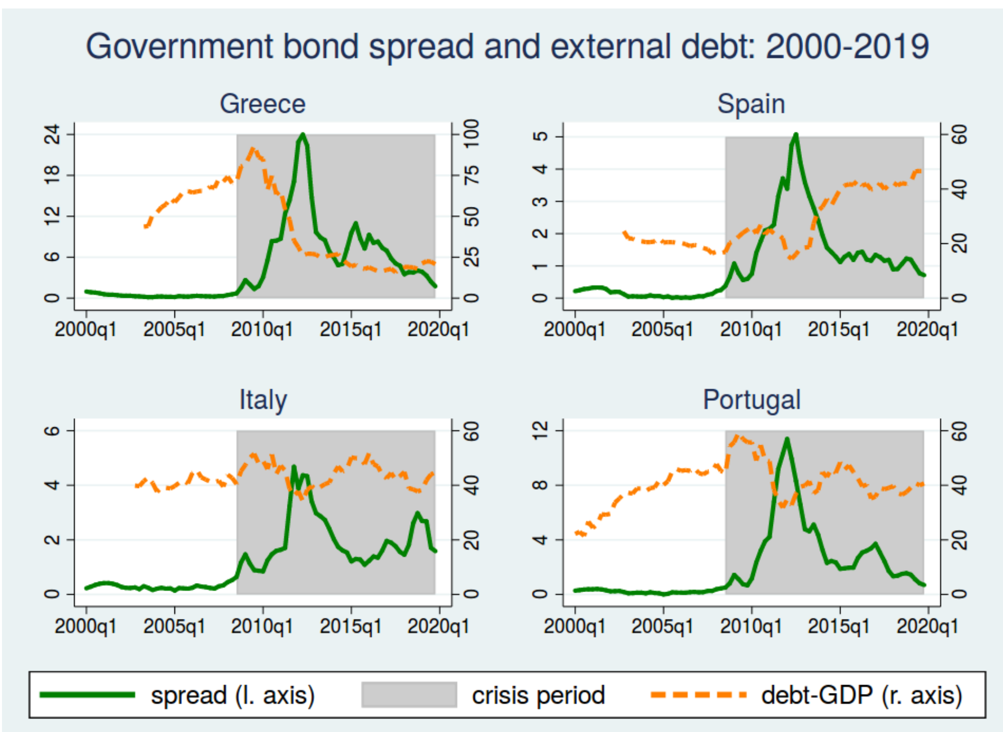

Note: The bond spreads are in percentage points and acquired from OECD. The debt series represent the external debt securities of the general government, and are expressed as a fraction of annualized GDP. The data come from the World Bank's Quarterly External Debt Statistics and start at different points in time for different countries. The shaded area starts from 2008:Q3 and represents the beginning of the crisis.

Figure 1: Government bond spreads and external debt of the peripheral European economies

[1]The bond spread is defined as the difference between the annual interest rates paid by the given country's

introduction of the euro, regardless of current economic performance. At the outset of the Great Recession, peripheral EU countries were hit by negative income shocks in the range between two to three standard deviations below their mean. Yet, despite the widespread expansion of external debt levels in 2009, government bond spreads temporarily rose above 1.5-2.5% in 2009:Q1, only to fall back below 1-2% by the second half of 2009.[2] The real stress did not start to build until mid-2010 when Greece experienced a sharp increase in borrowing costs, followed by other countries over the next two years. This slow evolution of the European crisis seems puzzling from the point of view of basic default models where interest rates are sensitive to income shocks and tend to comove positively with debt issuances.

To address this puzzle, I build on two observations about the European debt crisis. First, the experience of southern European economies in years 2008-2014 is much more than just a regular recession. The peak-to-trough decline in quarterly real GDP ranges from around 10% for Italy, Portugal and Spain to 30% for Greece. As such, the experiences of these countries qualify for what Barro (2006) refers to as a rare disaster, or what Kehoe and Prescott (2002) define as a great depression episode. Second, expectations about future income shocks evolved gradually among financial institutions in years 2008-2014, which I document using the real-time GDP forecast data. In particular, prior to 2008 the average forecast bias is at a similar level across all analyzed countries and forecasting agencies, below 1% of 2010 real GDP. Then, errors increase to 1.5-6% between 2008 and 2011, in all cases *overpredicting* future GDP growth. This indicates that negative output shocks at the time were perceived as a fairly typical recession, corresponding in size and persistence to Europe's post-war business cycle. Finally, in years 2012-2014 the forecasts become much more precise, with average bias falling back below 1% and often changing sign (i.e. underestimating future GDP). This increase in pessimism about the countries' economic outlook coincided with the dramatic spikes in interest rates on European governments' bonds.

Motivated by these two observations, I develop an otherwise standard model of sovereign debt that captures the two elements described above in a simple way. Starting with a long-term debt model as in Chatterjee and Eyigungor (2012) or Hatchondo and Martinez (2009), I first introduce a regime-switching income process where a "rare disaster" is represented by a large and negative shift in the long-run mean income of the economy. Then, I assume agents have incomplete information about the underlying switches between the regimes and

government bonds and a risk-free asset, in this case the German long-term government bonds.

[2] Appendix A.1 also shows that sovereign ratings assigned to these countries by the leading credit agencies closely mirrored the dynamics of bond spreads, with major downgrades only starting in 2010-2011.

learn about them over time in Bayesian fashion. As is typical for such a setting, bond yields carry a default premium which varies with the amount of outstanding debt and the expected fluctuations in future GDP. Crucially for my model, spikes in default risk coincide with rare disaster episodes rather than the recurring business cycle downturns as it is the case for emerging markets. In normal times, which can last for decades, there is little concern about a sovereign debt crisis in the foreseeable future and, consequently, bond prices carry a negligible default premium. When a rare disaster occurs, income is set on a downward trajectory; however, agents are not aware of this immediately. In other words, they cannot tell if the shocks they are observing are temporary or permanent in nature. In the presence of long-term debt, this information friction relieves the upward pressure on interest rates, because investors remain optimistic about the economy's long-run outlook. Over time, as income continues to decline, agents increasingly recognize the looming disaster and revise their forecasts. The result is a sudden, sharp spike in default risk that follows long periods of relative calmness in the bond markets.

I use Portugal as a quantitative case study and calibrate the parameters of the regime switching process using the aforementioned data on real-time GDP forecasts. The calibrated model exhibits interesting behavior in several ways. First, it features a highly volatile bond spread even though average spread is targeted at a low level. This result reflects the fact that the spread tends to be negligible for long periods of time while the good regime is in place, and then it shoots up and remains high when a disaster activates. Second, unlike in an off-the-shelf sovereign default model, my calibration requires a high value of the discount factor, reducing the typical high volatility of consumption or trade balance. This is due to the fact that defaults are generated by the occurrence of rare disasters, while the government behaves counter-cyclically during "normal times". Third, the government in the model sells bonds at steep discounts, resulting in high equilibrium interest rates, on average reaching 22% in the simulated crises. This occurs whenever the belief about an upcoming disaster increases suddenly, while the income level is still high and default is costly. The entire bond price schedule shifts downwards and the government is left with no other choice than to accept very low prices, until either it manages to deleverage or until income is low enough to justify a default.

In an event study of Portugal's debt crisis, I feed in the sequence of GDP data for years 1998-2019. Consistent with the data, the model predicts a negligible spread and slow debt accumulation prior to 2009. Then, the initial adverse income realizations cause an increase in the bond spread to 1.1% in the first quarter of 2009, which matches the data and contrasts with the increase to almost 5% for the standard, off-the-shelf AR1 model. This is because the agents are unsure if they are observing temporary shocks or a permanent regime switch. Over the next two years, the belief about the latter converges to certainty, and markets become convinced that the process has switched to a disaster. As a result, in 2012 we observe a delayed jump in the bond spread combined with a sharp reduction in government debt (and, in fact, an eventual sovereign default[3]). The model's predictions for the debt and spread also converge with the data in the post-bailout time period from 2016 until the end of 2019. More generally, the Portuguese bond spread closely tracks the (log) of the belief about the disaster realization for the entire period of 2000-2019. Finally, I show that learning is crucial to generate these predictions. In a counterfactual exercise where agents are fully aware of the upcoming disaster by the end of 2008, the bond spread increases in a similar fashion as in the standard AR1 model, while the government begins a drastic path of debt reduction.

1.1 Literature review

This paper is closely related to the quantitative sovereign debt literature, in particular it builds on the seminal work of Eaton and Gersovitz (1981) and, more recently, Aguiar and Gopinath (2006) and Arellano (2008). Chatterjee and Eyigungor (2012) and Hatchondo and Martinez (2009) introduce long-duration bonds to these models and show that they are important in accounting for the amounts of debt and average spreads observed in the data.

A recent branch of quantitative default literature investigates the ability of such models to match the volatility of sovereign spreads, with a focus on the European debt crisis. Aguiar et al. (2016) point out that calibrated long-term debt models often deliver a standard deviation of the spread an order of magnitude lower than what it is in the data. To address this issue, Aguiar et al. (2020), Ayres et al. (2019), and Lorenzoni and Werning (2019) revisit models with multiple equilibria to justify why bonds are often sold at large discounts, while Paluszynski and Stefanidis (2020) show that much of the missing spread volatility may be due to the frictions in adjusting government spending. Bocola and Dovis (2019) use the observed maturity choices to identify the rollover risk component of sovereign spreads. Bocola, Bornstein and Dovis (2019) emphasize the role of domestic debt in generating a high spread volatility relative to its mean. The model in this paper achieves a similar objective of generating volatile and high-peaking spreads through a mechanism of learning about rare disasters.

Chatterjee and Eyigungor (2019) present a model of sovereign default with political frictions that shares many similarities with the model in this paper, such as a regime-switching

[3]Recall that Portugal, together with other European countries, received official bailouts in excess of 40 percent of their GDP, to prevent them from defaulting, an element not present in my model.

income process. In particular, both models obtain volatile bond spreads with a relatively high value of the discount factor. What differentiates them is that their paper is interested in endogenous political turnover in emerging market economies (such as Mexico, Peru or Turkey), while I focus on the European debt crisis. The main advantage of my paper is that I discipline the key parameters of my model using the data of real-time forecasts, which allows me to capture the slow learning process. This is crucial for generating the correct predictions for the evolution of bond spreads at the outset of the Great Recession in Europe.

On a more general level, this paper is related to two strands of literature in macroeconomics and finance. The first strand introduces rare disasters to otherwise standard macroeconomic models, motivated by Rietz (1988) or Barro (2006). Several papers have recently used this concept in the context of sovereign debt.[4] The second strand incorporates learning about unobserved economic conditions in macroeconomics. Boz and Mendoza (2014) present a model where households learn about the probability of switching between credit cycles to produce a boom-bust cycle like the one observed around 2008 in the US. Boz, Daude and Durdu (2011) show that learning about the permanent vs. transitory nature of shocks can explain some of the observed differences in volatility between developed and emerging economies. My paper combines these two strands of literature and shows how the data on real time forecasts can be used to estimate the parameters of a rare disaster and to improve the model's predictions.

The remainder of this paper is structured as follows. Section 2 describes the motivating evidence regarding the European debt crisis. Section 3 introduces the main model. Section 4 calibrates the model and uses it to analyze the European debt crisis and contrast the results with those obtained using a benchmark version of the model. Section 5 concludes.

2 Empirical motivation

In this section, I document the two pieces of evidence that serve as the main motivation for the model in Section 3, namely the depth of the output declines and the evolution of forecast errors in years 2008-2016.

2.1 Depth of GDP drops

In order to highlight the magnitude of the decline in economic activity among the peripheral European countries, Table 1 lists the largest peak-to-trough drops since 2007. The numbers

[4]See, for example, Ayres et al. (2019), Aguiar and Amador (2020), or Rebelo, Wang and Yang (2019).

provided refer to the drops in real GDP both at face value, and in relation to a 2% trend. Notice that the former ranges from almost 10 percent for Spain, Italy and Portugal, up to over 30% percent in the case of Greece. The cutoff size for a contraction in face value that defines a rare disaster in Barro (2006) is 15 percent. He emphasizes, however, that using an alternative threshold of 10 percent delivers similar results in terms of solving the equity premium puzzle. On the other hand, Kehoe and Prescott (2002) define the "great depression" episode as a sustained negative deviation of at least 20 percent in the GDP level net of the 2 percent annual trend growth. Table 1 indicates that all the economies of interest are around this threshold or clearly above (Greece); the decline is also sustained in time.[5]

Table 1: Peak-to-trough GDP drops among the peripheral European economies

Country	**Largest decline (in %)**		**Quarters**
	face value	*detrended*	
Greece	30.6	39.9	23
Spain	9.1	19.0	21
Italy	9.4	21.4	19
Portugal	9.7	19.4	21

Note: This table shows the largest recorded declines in real GDP in the period 2007-2014, measured at face value and detrended with 2% trend as in Kehoe and Prescott (2002).

2.2 Market expectations during the recession

As a second piece of evidence, I investigate the paths of forecasts about GDP growth in real time. The distribution of future income shocks is a crucial element driving interest rate fluctuations in sovereign default models, and thus deserves particular attention.

Figure 2 presents the plot of real GDP over time for the four European countries, along with the GDP forecasts published every year by OECD.[6] As can be noticed for the period prior to 2008, while the European economies are still growing along a stable trend the observed forecast errors are small (with some overshooting for Italy and Portugal, whose economies experienced a slowdown in the early 2000s). When the financial crisis breaks out, the forecasts are still fairly optimistic, predicting a recovery in years 2008-2010. Over time however, as the GDP continues to plunge we also observe that the forecasts become flatter, indicating

[5] Using the countries' individual trends rather than the common 2 percent growth rate would make this conclusion similar, or even starker.

[6] For illustration, in this figure I use the two-year-ahead forecasts of real GDP growth from the Fall issues of OECD's Economic Outlook (the Spring version only provides one-year-ahead forecasts). In what follows, I also present the forecast errors from other institutions, public and private alike.

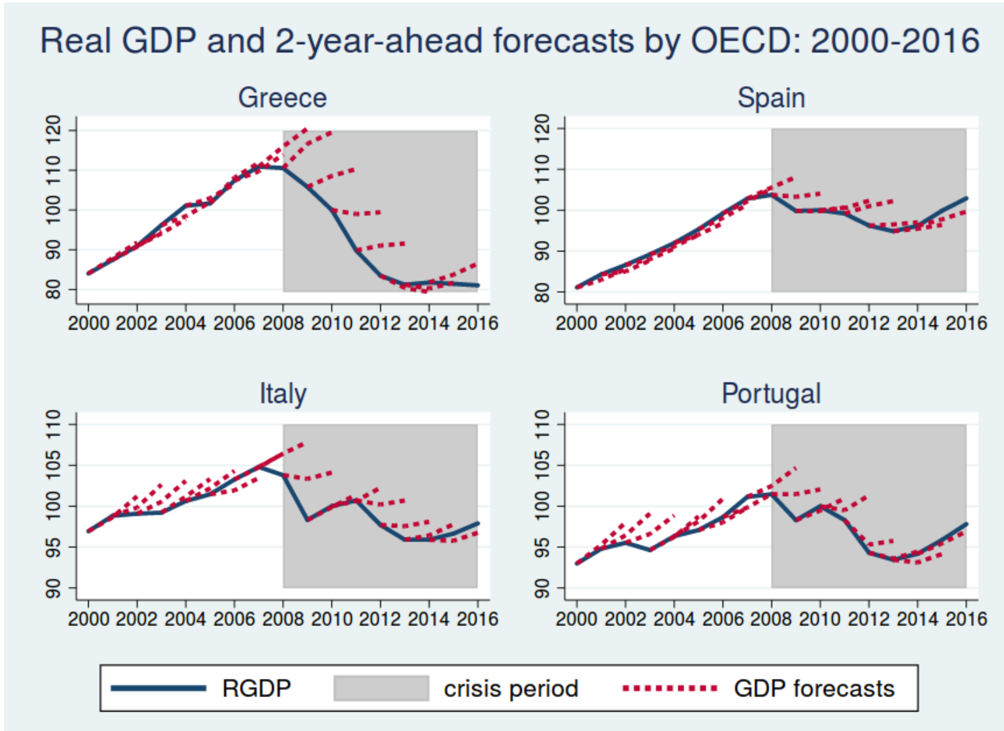

Note: The GDP series are annual and expressed in constant prices; their values are normalized such that the observation for 2010 equals 100. The red dotted lines represent one- and two-year ahead forecasts published by the OECD Economic Outlook (fall edition) and start in the year when each of them is made.

Figure 2: Forecast and actual real GDP for the peripheral European economies

that the markets have realized the recovery of output cannot be expected in the short and medium term. From 2012 on, the forecasts essentially line up again with the subsequently realized data for all of the depicted economies. This is also the time when the European bond markets undergo major turbulences that feature surging interest rates, sovereign bailouts and drastic reductions in debt levels, as documented in Figure 1.[7]

To analyze the forecast errors around the Great Recession more systematically, I acquire real-time predictions from three organizations: OECD, IMF and the European Commission. The historical forecasts of these institutions are available publicly and released every year in two vintages, Spring and Fall (roughly corresponding to May and November, respectively, so they are based on the knowledge of the data for first and third quarter of that year). In addition, I obtain the data from *Consensus Economics*, a survey of forecasters from private banks, government agencies, think-tanks and research centers. The queries are collected monthly and ask about a number of macroeconomic indicators, in particular the real GDP growth. A

[7]Appendix A.2 shows that an analogous learning process was absent around Argentina's episode in 2001, the most prominent case study of an emerging market sovereign default.

consensus forecast is defined as the average prediction across all participants of the survey. Broad literature documents that consensus forecasts do not suffer from many of the biases that may affect private and public institutions alike, and have traditionally performed better than any individual forecaster over the long run (see e.g. Batchelor (2001) and Cimadomo, Claeys and Poplawski-Ribeiro (2016)). In Tables 2 and 8 I test this conclusion for the southern European economies and investigate the size and direction of forecast errors.

Table 2: Average bias in real-time historical forecasts for different time frames

Average bias	**OECD**	**IMF**	**EC**	**CE**
(a) Pre-recession sample: 2000-2007				
Greece	−0.37	0.05	−0.35	0.00
Spain	0.10	0.03	0.07	0.16
Italy	−0.97	−1.18	−1.04	−0.97
Portugal	−0.97	−0.94	−0.75	−0.96
(b) Recession - first stage: 2008-2011				
Greece	−6.53	−6.85	−6.91	−6.73
Spain	−2.19	−2.42	−2.32	−2.33
Italy	−2.19	−1.98	−2.21	−2.05
Portugal	−1.53	−1.75	−1.31	−1.73
(c) Recession - second stage: 2012-2014				
Greece	−0.21	−1.09	−1.05	−0.16
Spain	0.90	0.83	0.65	0.88
Italy	−0.51	−0.70	−0.85	−0.64
Portugal	0.42	−0.07	−0.03	0.59

Note: The table presents average errors of one-year-ahead forecasts of real GDP level. Forecasts are acquired from four sources: OECD, IMF, European Commission, and Consensus Economics Inc. The bias is expressed as a percentage of the 2010 level of real GDP for each of the four countries. A positive value for the bias indicates that the forecasts underestimate the actual values, while a negative bias indicates that the forecasts overestimate them. All forecasts come in two vintages, Spring and Fall, which I use jointly. The number of forecasters participating in Consensus Economics surveys varies over time and across countries, with a minimum of four and a maximum of twenty in the entire sample.

I use one-year-ahead forecasts from each of the four institutions, combining both vintages in every year.[8] To ensure direct comparability, I consider the May and November issues of the Consensus Economics survey which roughly coincide in time with the Spring and Fall reports of the IMF, OECD and the European Commission. Table 2 reports the average bias in the GDP forecasts for the four countries and four sources of interest during three

[8]While OECD and the European Commission also publish up to two-year ahead, and IMF up to five-year-ahead forecasts, the Consensus Economics survey is limited to next-year predictions only.

separate time periods: a pre-recession sample of 2000-2007 and the two stages of the Great Recession, namely 2008-2011 and 2012-2014. The bias is expressed as a percentage of each country's 2010 real GDP level and takes negative values when forecasts *overshoot* the actual realizations. There are several interesting observations about this data. First, prior to 2008 the average bias is of similar size across the forecasting agencies, generally under 1% of 2010 real GDP. Second, in years 2008-2011 the errors increase sharply, ranging from $1.5-2.5\%$ for Spain, Italy and Portugal, up to 6.5% for Greece. Third, in years 2012-2014 the bias drops significantly for all analyzed countries, in most cases below the average level from before 2008 (some of the forecasts, in fact, underpredict the GDP level in that point). Finally, the Consensus Economics forecasts *do not* outperform the international organizations, especially during the first stage of the recession in 2008-2011. Interestingly, consensus forecasts do better prior to 2008, in line with the findings of earlier studies such as Batchelor (2001). This indicates that private-sector expectations were likely to exhibit excessive optimism especially at the beginning of the Great Recession.[9]

Large forecast errors at the height of the crisis became a subject of intense critique and led the OECD to publish a study to evaluate the source of mistakes. In a "Post Mortem", Pain *et al.* (2014) write:

> *GDP growth was overestimated on average across 2007-12, reflecting not only errors at the height of the financial crisis but also errors in the subsequent recovery. (...)*
>
> *The OECD was not alone in finding this period particularly challenging. The profile and magnitude of the errors in the GDP growth projections of other international organisations and consensus forecasts are strikingly similar.*

In their *ex post* reflection, the OECD points to the repeated expectation of a swift recovery as the main source of forecast errors, which suggests that a learning process was taking place. To not appear as the only culprit, the OECD also emphasizes that the overly optimistic forecasts have been common among other influential forecasters associated with international bodies and consensus measures. This claim is confirmed by Figure 2 and Tables 2 and 8.

[9]While average bias is a useful measure for evaluating the direction of forecast errors, it does not necessarily present their full magnitude. This is because errors of opposite signs may cancel each other out over time. Moreover, even if errors generally go in the same direction (just like during the Great Recession), the average bias aggregates them linearly, i.e. an error of 2% is equivalent to two errors of 1% each. As a result, it treats large errors during events like the Great Recession with a similar weight as several small errors combined in the period prior to 2007. To address this issue, Table 8 in Appendix A presents analogous calculations for the root mean square errors (RMSE). As a non-linear measure, RMSE punishes infrequent large errors more heavily than frequent small ones. Indeed, the analysis of RMSEs supports all four observations discussed above.

3 Model

In this section I present a model of sovereign debt that features an augmented specification of the income process and incomplete information about its realizations.

3.1 Economic environment

Consider a representative-agent small open economy with a benevolent sovereign government that borrows internationally from a large number of competitive lenders. Time is discrete and there is no production or labor. Instead, the economy faces a stochastic stream of endowment realizations. Markets are incomplete and the only asset available for trading is the multi-period non-contingent bond.

Endowment process Suppose the country's endowment follows an autoregressive regime-switching process. I assume that there are two possible regimes, High and Low, and each of them is characterized by its own long-run mean. For simplicity, the persistence and variance parameters are assumed to be constant across regimes. Specifically, the evolution of output, detrended with a deterministic long-run mean growth rate,[10] is given by

$$y_t = \mu_j(1-\rho) + \rho y_{t-1} + \eta\varepsilon_t \tag{1}$$

where $\varepsilon_t \sim \mathcal{N}(0,1)$ is an $i.i.d.$ random shock and $\rho, \eta, \{\mu_j\}_{j=L,H}$ are parameters of the two regimes. Regimes change according to a Markov process with the transition probability matrix given by

$$\Pi = \begin{bmatrix} \pi_{LL} & 1-\pi_{LL} \\ 1-\pi_{HH} & \pi_{HH} \end{bmatrix} \tag{2}$$

The specification of a bimodal stochastic process of endowment in formula (1) is non-standard in the sovereign debt literature.[11] It is motivated by the income pattern of European economies in the recent decade, illustrated in Table 1. Throughout the paper, I will consider the two regimes as highly asymmetric, with the low one having the interpretation of a rare disaster or a great depression.

[10] Appendix B.1 discusses details of the detrending method used in this paper.

[11] Simultaneously with the present paper, a bimodal income process was also used by Chatterjee and Eyigungor (2019) and Ayres et al. (2019). In contrast to these papers, the regimes here are meant to be highly asymmetric which can be interpreted as "normal times" and "rare disaster". In such a setting, learning about the underlying regime has a naturally powerful effect, as I show in Section 4, because agents tend to have a strong prior belief against a potential regime switch. I also show a new method of estimating this income process by incorporating real-time forecast data.

Preferences The representative household has preferences given by the expected utility of the form:

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{3}$$

where I assume the function $u(\cdot)$ is strictly increasing, concave and twice continuously differentiable. The discount factor is given by $\beta \in (0, 1)$.

Government In each period, the government chooses a consumption rule and the level of debt holdings to maximize the household's lifetime utility. The only asset available is the long-duration bond. In the spirit of Chatterjee and Eyigungor (2012), I assume that each unit of outstanding bonds matures probabilistically in every period, or pays a fixed coupon. The government may save at an international risk-free interest rate. If it decides to borrow, however, the government is not committed to repay the debt next period. Consequently, the bond is priced endogenously by risk-neutral lenders to account for the possibility of default as well as debt dilution in the future. As is commonly assumed in the sovereign debt literature, the government who refuses to honor its obligations faces an exogenous cost of default and is further excluded from borrowing in the financial markets, with a certain probability of being readmitted in every subsequent period.

Market clearing There is no storage technology which, under the aforementioned assumptions on the utility function, implies that the endowment is fully divided between current consumption and net borrowing. This market clearing condition is given by

$$c_t = y_t - b_t \left(\delta + (1 - \delta)\kappa\right) + q_t \left(b_{t+1} - (1 - \delta)b_t\right) \tag{4}$$

where q_t is the price of the debt stock b_{t+1} (to be repaid next period), δ is the rate at which bonds mature every period and κ is a fixed coupon.

Bond prices International lenders are perfectly competitive and have "deep pockets" in the sense that potentially even large losses do not affect their decisions. In equilibrium the lenders make expected zero profit and as a result, the bond pricing formula compensates them only for the default risk implied in the government's decisions.

3.2 Information structure

The two state variables mentioned so far, current bond holdings (b_t) and income (y_t), are standard in sovereign debt literature. In addition, this model features another exogenous stochastic variable, $z_t \in \{z_L, z_H\}$ representing the regime (Low or High) in which the economy is currently operating. While all agents know the latest income realization, they have incomplete information about the current regime. Instead of observing it directly, agents form a belief p_t defined as their perceived probability of being in the High regime, formally $p_t \equiv Prob(z_t = z_H)$. Intuitively, this variable can be thought of as market sentiment about the economy's expected future income path. As I show in Section 4, the belief about regime is quantitatively significant and appears to have fluctuated substantially in years 2008-2014.

3.3 Timeline

In every period, the timing of events is as follows:

1. The new regime $z \in \{z_L, z_H\}$ is drawn, with the probability distribution given by (2).
2. The new realization of endowment y is drawn, according to the newly updated regime z and conditional on its level from last period.
3. Agents observe y and mechanically form a new belief p about the regime, conditional on the previous and current endowment, as well as the last period's belief.
4. Default and redemption decisions take place:
   - The government that has recently defaulted on its debt draws a random number to determine whether it can be readmitted to the financial markets.
   - The government that has recently been current on its debt decides whether to repay or default this period.
5. Equilibrium allocations take place:
   - If the government defaults, it is excluded from financial markets this period and simply consumes its endowment, subject to a default penalty.
   - If the government repays, it chooses the new allocation of bonds b', while the lenders post the bond price $q(b', y, p)$.

3.4 Recursive formulation

In the following section I formalize the economic environment by stating the problems faced by market participants in recursive form. To begin, define the vector of aggregate state variables that are common knowledge as $\mathbf{s} = (b, y, p)$.

Government The government that is current on its debt obligations has the general value function given by

$$v^0(\mathbf{s}) = \max_{d\in\{0,1\}} \Big\{(1-d)v^r(\mathbf{s}) + dv^d(y,p)\Big\} \tag{5}$$

A sovereign who defaults ($d = 1$) is excluded from international credit markets and has probability θ of being readmitted every subsequent period with zero debt. The assumption that all debt is wiped out upon readmission is not necessary and can be relaxed at the expense of complicating the analysis. The associated default value is

$$\begin{aligned} v^d(y,p) = \; & u\Big(y - h(y)\Big) + \beta \sum_{z\in\{z_L,z_H\}} \sum_{z'\in\{z_L,z_H\}} Prob(z)\,\pi(z'|z)\times \\ & \int f_{z'}(y',y)\Big[\theta v^0(0,y',p') + (1-\theta)v^d(y',p')\Big]\; dy' \end{aligned} \tag{6}$$

subject to the law of motion for the belief

$$p'(y,p,y') = \frac{\Big[p\,\pi(z_H|z_H) + (1-p)\,\pi(z_H|z_L)\Big]\,f_{z_H}(y',y)}{\sum_{z'=z_L,z_H}\Big[p\,\pi(z'|z_H) + (1-p)\,\pi(z'|z_L)\Big]\,f_{z'}(y',y)} \tag{7}$$

In equation (6), $h(\cdot)$ is a reduced-form representation of the output cost of defaulting;[12] $f_{z'}(y'|y)$ denotes the probability density of transitioning from state y to state y' given that tomorrow's regime is z'. $Prob(z_H)$ is equal to p and $Prob(z_L)$ is $1-p$, $\pi(z'|z)$ is the probability of transitioning from regime z today to z' tomorrow. The next period belief p', described in equation (7), depends on the current and future income realization, as well as the current belief p. It is a simple application of Bayes' rule and takes into account a potential regime switch at the beginning of next period, according to the transition matrix given by (2).

The value of the government associated with repayment of debt is given by

$$v^r(\mathbf{s}) = \max_{c,b'} \left\{u(c) + \beta \sum_{z\in\{z_L,z_H\}} \sum_{z'\in\{z_L,z_H\}} Prob(z)\,\pi(z'|z)\int f_{z'}(y',y)v^0(\mathbf{s}')\; dy'\right\} \tag{8}$$

subject to the law of motion for the belief in formula (7) and

$$c = y - b\Big(\delta + (1-\delta)\kappa\Big) + q(b',y,p)\Big(b' - (1-\delta)b\Big) \tag{9}$$

where equation (9) is the budget constraint.

[12] Quantitative sovereign debt models typically assume an exogenous punishment in the case of default in order to facilitate calibration of the model to the data. For the specific functional form, see Section 4.3.

Having characterized the two value functions of the government, it is straightforward to derive the optimal default policy as a function of today's state variables

$$d(\mathbf{s}) = \begin{cases} 1, & \text{if } v^d(y,p) > v^r(\mathbf{s}) \\ 0, & \text{if } v^d(y,p) \leq v^r(\mathbf{s}) \end{cases} \tag{10}$$

International Lenders Every period the lenders only observe (b, y) and share a common market belief p. Although they do not see the current regime z, they know its distribution and independently update their belief about it, as described by the law of motion in formula (7). The denominator in those equations is always greater than zero and the resulting next period belief p' is strictly interior on the interval $(0, 1)$.

As is common in the quantitative models of sovereign debt, lenders are competitive and risk-neutral by assumption. The resulting equilibrium bond price is such that they make zero profit in expectation (according to their imperfect information). The bond price function is

$$\begin{aligned} q\left(b', y, p\right) = \ & \frac{1}{1+r^*}\Bigg(\sum_{z\in\{z_L,z_H\}}\ \sum_{z'\in\{z_L,z_H\}} Prob(z)\,\pi(z'|z)\times \\ & \int f_{z'}(y', y)\Big(1-d(\mathbf{s}')\Big)\Big[\delta + (1-\delta)\Big(\kappa + q\left(g(\mathbf{s}'), y', p'\right)\Big)\Big]\,dy'\Bigg) \end{aligned} \tag{11}$$

where $\mathbf{s}' = \Big(b', y', p'(y, p, y')\Big)$, $d(\cdot)$ and $g(\cdot)$ are the government's optimal decisions with respect to default and new debt, respectively, and r^* is the risk-free rate of interest.

Concluding this section, Definition 1 introduces the standard concept of a Markov Perfect Bayesian Equilibrium. In this equilibrium the posterior beliefs of agents must be specified at all states and for all strategies of other players (including those involving off-equilibrium actions). The agents' best responses must belong to the set of stationary Markov strategies.

Definition 1 *A Markov Perfect Bayesian Equilibrium for this economy consists of the government value functions $v^r(\mathbf{s})$, $v^d(y,p)$ and policy functions $c(\mathbf{s})$, $b'(\mathbf{s})$, $d(\mathbf{s})$ and the bond price schedule $q(b', y, p\Big)$ such that:*

1. *Policy function d solves the government's default-repayment problem (5).*
2. *Policy functions $\{c, b'\}$ solve the government's consumption-saving problem in (8).*
3. *Bond price function q is such that the lenders make zero expected profit (subject to their imperfect beliefs).*

4 Quantitative analysis

In this section I calibrate the model to Portuguese data and discuss its mechanics. I then present the simulated behavior of the model and use it to study the European debt crisis.

4.1 Data

I use data from the Portuguese economy as a case study for the theory developed in this paper. The model could also be calibrated to other European economies discussed in Section 2. The Portuguese episode is the most clear-cut case, however, because it does not coincide with other major economic events such as a banking crisis (like the one that occurred in Ireland) or Mario Draghi's "whatever it takes" speech and the introduction of the Outright Monetary Transactions (OMT) program in the summer of 2012, at the peak of the debt crises in Italy and Spain. Portugal is also a particularly relevant laboratory for this type of sovereign default model as it easily satisfies its core assumptions, i.e. it is arguably a small open economy and the vast majority of its debt securities were held externally (Andritzky, 2012). While in principle Greece is also a plausible candidate, the validity of its macroeconomic data is questionable. Nonetheless, Appendix B.3 extends the main estimation of the regime-switching income process to other countries mentioned in the introduction and discusses the usefulness of this theory in explaining these cases.

Quarterly data for real GDP are taken from OECD and cover the period 1960:Q1-2019:Q4. Consumption, trade balance and interest rates on long-term government bonds, also from OECD, span the time frame of 1998:Q1-2019:Q4. Government debt data is acquired from World Bank's Quarterly External Debt Statistics (I use debt securities only).

The identification strategy for the model's parameters is in line with the general approach in the literature. In what follows, I first estimate the income process using both historical GDP data and the real-time GDP forecast data introduced in Section 2.2. Then, I select the remaining structural parameters of the model partly from the literature and partly to match certain general characteristics of the Portuguese economy.

4.2 Estimation of the income process

I proceed to calibrate equation (1) in two steps. First, I fix the probabilities of switching into, and out of, the rare disaster regime based on recent historical experience. By all accounts, the recession in southern European countries has been the worst since the Great

Depression.[13] This gives us roughly 60 years of high regime duration.[14] On the other hand, the Great Depression lasted for about 10 years, which I use to pin down the expected low regime duration. The resulting probabilities of staying in the high and low regimes are therefore 0.996 and 0.975, respectively. The exact numbers behind these probabilities are not crucial for the results because, given their predetermined values, I use two other data sources below to discipline the remaining parameters of the income process. What matters, however, is to capture the right order of magnitude - intuitively, the low regime ought to be rare and severe enough so that it clearly stands out from a regular economic contraction.[15]

In the second step, I normalize the high-regime mean μ_H to zero and employ a variant of the Expectation-Maximization algorithm (Hamilton, 1990) to estimate the remaining coefficients. Importantly, to capture the slow-learning process documented in Section 2.2, I jointly use two data sources – the historical GDP series for years 1960:Q1-2019:Q4,[16] and the real-time GDP forecasts for years 1993-2014.[17] Appendix B.2 discusses the details of my estimation technique. Table 3 summarizes the calibrated parameters of the income process. It should also be emphasized that the estimated regime-switching model in Table 3 provides a better fit to Portuguese data than a single-mean AR1 process. To show this, I estimate the detrended AR1 process for years 1960-2011 to be 0.947 and 0.011, respectively.[18] Then,

[13]While Portugal was not heavily impacted by the Great Depression itself, it subsequently suffered during the 1934-1936 civil war in Spain, with a peak-to-trough decline in real GDP of 12.5%. Following the 1930s, the 2008-2014 episode is by far the most severe contraction in Portuguese economic history. It comes close to satisfying the defining criteria of a great depression established by Kehoe and Prescott (2002). Also Reis (2013) compares the Portuguese episode to the US Great Depression and Japan's lost decade.

[14]I ignore the second world war in my calculations as the model is not designed to account for such events.

[15]This type of approach to calibration is common for models with disaster risk. For example, pooling 60 episodes in 35 countries, Barro (2006) sets the probability of entering a disaster event at 1.7% annually, almost the same as the number I use. More recently, Coibion, Gorodnichenko and Wieland (2012) take a similar approach to approximate the frequency of interest rates in the US hitting the zero lower bound, which occurred for the first time since the second world war.

[16]Because of the growth trend changes for European economies over this relatively long sample, I follow Bai and Perron (1998) to identify statistically significant structural breaks in the growth rate of the Portuguese economy. Two breakpoints are detected, at 1974:Q2 and 1999:Q4, which is intuitive as they coincide with the democratic revolution in Portugal and accession to the Euro zone, respectively. The estimated quarterly trend growth rates for the three time windows are 1.6%, 0.8% and 0.4%. Including all of the available information since 1960 is important to capture the full scope of the variance of GDP while the high regime is in place, i.e. during "normal times". This is an issue especially for the European countries whose GDP data exhibits very little variance in years 1999-2008 alone. See Appendix B.1 for more details on the detrending.

[17]Specifically, I use 5-year-ahead projections published by the IMF, given that they provide the longest time series, and the longest forecasting horizon. I do not include the post-2014 forecasts in this estimation to avoid capturing the reversal of pessimism documented in Section 2.2. See also Paluszynski (2019) which uses long-term forecast data to calibrate the parameters of a stochastic process for the risk-free interest rate.

[18]Once again, I use a broken linear trend with two statistically significant breakpoints detected using the Bai and Perron (1998) test at 1974:Q2 and 1999:Q4. Note that using the GDP data all the way until 2019 produces similar estimates and mostly does not affect the results of the paper. Its interpretation is

Table 3: Parameters of the regime-switching endowment process

				Transition Prob.	
Regime	Mean μ	Persistence ρ	St. dev. η	Low	High
Low	-0.291	0.970	0.010	0.975	0.025
High	0.000	0.970	0.010	0.004	0.996

the likelihood ratio test statistic is 22.1 and has approximately a chi-squared distribution with three degrees of freedom (the number of additional parameters in the extended model), resulting in a p-value much smaller than 0.01. This suggest that we can reject the null model (single-regime AR1) at virtually all levels of significance.

Concluding this section, it is natural to ask about the results of a more standard estimation technique that does not use the real-time forecasts as an additional data source. Such an approach would yield the following estimates for the sample from 1960 to 2011: $\rho = 0.95$, $\eta = 0.01$, $\mu_H = 0.03$, $\mu_L = -0.32$, $\pi_{HH} = 0.96$, $\pi_{LL} = 0.71$. Notice that now the high regime mean is no longer normalized to zero, while the low regime has the unconditional mean similar to the one reported in Table 3. Crucially, switching is much more likely under this specification and the regimes do not last very long. We detect several instances of the low regime over time, in particular around 1969, 1974, 1983, 1992, 2008-2009 and 2011-2014. Because of such frequent switching, the learning process about an underlying switch is fast and does not align well with the evidence from real-time forecasts presented in Section 2.2.

4.3 Functional forms and calibration

To select the remaining structural elements of the model, I follow the general trends in the literature by fixing the value of some non-controversial parameters, and I use a moment-matching exercise to pin down the more problematic ones. The representative household's utility is a CRRA function of the form $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$, with risk aversion parameter set at the standard level of 2. The risk-free interest rate is set equal to 1% (quarterly value) and the probability of re-entry after default is fixed at 0.049, following Cruces and Trebesch (2013) who find that the average time to re-enter the credit market was 5.1 years in 1970-2010. Using OECD data I find that the average maturity of Portugal's debt in years 1996-2010 was 4.73 years, which translates into an average quarterly maturity rate of 0.053. The coupon payment is set to 1.25% following Salomao (2017), which implies an annual coupon of 5%.

problematic, however, because it implies that the slope of the trend has been essentially zero since 1999, and thus year 2008 appears to be the peak of an historic boom for the Portuguese economy.

The output cost of default is parametrized as $h(y) = \min\left\{y, \hat{y}\right\}$ following Arellano (2008). The parameter $\hat{y}$ is calibrated jointly with the discount factor β using the simulated method of moments. The economy's income path in years 1998-2019 is simulated 10,000 times, starting from the actual GDP and debt levels observed in 1998:Q1 and under the assumption that the regime switches from High to Low in 2008:Q3.[19] The idea behind the identification strategy is to match certain general characteristics of the Portuguese experience during that time period. To this end, I use information from two standard moments of Portugal's economy in years 1998-2019: average ratio of external debt securities to GDP of 38.6%;[20] and average 5-year bond spread of 1.75%.[21] The former is naturally an important piece of information to identify the relative impatience of the government and the punishment for default. The latter, in a model with risk-neutral lenders and zero recovery rate, simply reflects the government's average default probability. The target of 1.75% is thus reasonable given all the independent evidence on Portuguese default history.[22] Table 4 presents a summary of the parameter values that provide the closest match to the empirical moments. The model achieves an exact match in terms of both targeted moments. The calibration procedure results in a discount factor β of 0.988 and the default penalty parameter $\hat{y}$ of 0.793.

Finally, in order to make a meaningful comparison with a literature benchmark, I also calibrate a "standard" sovereign debt model with long-term debt, similar to Chatterjee and Eyigungor (2012) or Hatchondo and Martinez (2009), which uses a simple AR1 specification of the income process. As mentioned before, the estimated persistence and variance parameters are 0.947 and 0.011, respectively. Calibration of the structural parameters closely follows the strategy described above and is summarized in the AR1 column of Table 4. It is important to emphasize that this benchmark is not meant to claim that no model based on an AR1 process can produce relevant predictions for the European debt crisis. Indeed, several papers such as Bocola and Dovis (2019) or Salomao (2017) have managed to do so. The benchmark thus refers to an "off-the-shelf" model following a routine calibration approach.

[19] An alternative would be to simulate the economy over many years and calibrate to its ergodic (long-run) business cycle statistics outside of default, following Chatterjee and Eyigungor (2012). However, as it will become clear from Table 5 describing the simulation results, years 1998-2019 were a highly non-stationary period for the European economies, including slow accumulation of debt towards a steady state level.

[20] Because the model does not include post-default renegotiation, I follow Chatterjee and Eyigungor (2012) to calibrate only the true "unsecured" portion of the debt. While Portugal in the end did not default and it is difficult to know how much of its debt was in fact unsecured, the best guess is 0.535, the haircut rate in the case of Greek default of 2012.

[21] I use a 5-year spread, rather than 10-year as in the introduction, because Portugal's average debt maturity is 4.73 years. The data on 5-year spread is acquired from Bloomberg.

[22] Reinhart and Rogoff (2009) identify four sovereign defaults in Portugal's history since 1800, while Standard & Poor's (2014) identify three, implying an annual long-run probability of $1.5-2\%$

Table 4: Calibration of structural parameters of the model

Symbol	Meaning	Learning	AR1	Source
σ	Risk aversion	2	2	Literature
r^*	Risk-free rate	0.01	0.01	Literature
θ	Re-entry probability	0.049	0.049	Literature
δ	Probability of maturing	0.053	0.053	Data
κ	Coupon payment (in %)	1.250	1.250	Data
$\hat{y}$	Default cost par.	0.793	0.919	Calibration
β	Discount factor	0.988	0.980	Calibration
Calibration targets		**Learning**	**AR1**	**Data**
E (debt/GDP)		38.59	38.58	38.58
E (spread)		1.75	1.75	1.75

Note: Targeted moments are given in percentage points. Simulations are repeated 10,000 times for a period of 1998-2019.

An important conclusion from the moment-matching exercise is that it produces significantly different values of the discount factor in the model with disaster risk and learning, relative to the literature benchmark. In particular, in the "standard" model with a simple AR1 process, a value of β around 0.98 is needed to simultaneously generate high debt and defaults occurring with a desired frequency. On the other hand, in the model with learning about disasters the same targets are achieved with a discount rate of just under 0.99. This is because defaults here occur predominantly under the circumstances of a rare disaster, rather than due to myopic behavior of the borrower. As a result, the model features a government that mostly uses debt for consumption-smoothing purposes, but may occasionally default should the output collapse in a great-depression-like fashion. The following two sections illustrate the behavior of the government in this model in more detail.

4.4 Characterization of the equilibrium

In the following section I first characterize some of the key properties of the equilibrium, and then show how the model's simulated behavior compares with actual data. The model is solved numerically by value function iteration using a continuous choice of next period debt and cubic spline interpolation (Habermann and Kindermann, 2007) to evaluate off-grid points, similarly as described in Hatchondo, Martinez and Sapriza (2010). Expectations are approximated using Gaussian quadrature with 51 nodes and off-grid points for income and beliefs are linearly interpolated. I use 41 points for the grid of assets, income, and the belief.

4.4.1 Model mechanics

To understand how the model works, it is instructive to examine how the government's optimal decisions change with respect to state variables. Figure 3 shows the default and debt policies for different levels of prior belief. On the left-hand side panel, any combination of current debt and income above the line corresponding to some belief p indicates repayment, while a combination below the line indicates default. Not surprisingly, higher belief about being in the good regime induces the government to default in a smaller number of states. This relationship is strictly monotonic in the level of prior belief (but not necessarily linear). The right-hand side panel of Figure 3 shows that higher prior beliefs induce the government to borrow more. In this model, agents are impatient and would rather consume today than tomorrow. When making their debt decisions, however, they need to weigh their impatience against the expected income level in the future. A higher chance of being in economic depression next period implies that the government must restrict its consumption today and reduce foreign debt, in order to decrease the probability of defaulting tomorrow and to secure a high bond price today. Consequently, higher market belief has a strictly monotonic, increasing effect on the optimal debt level.

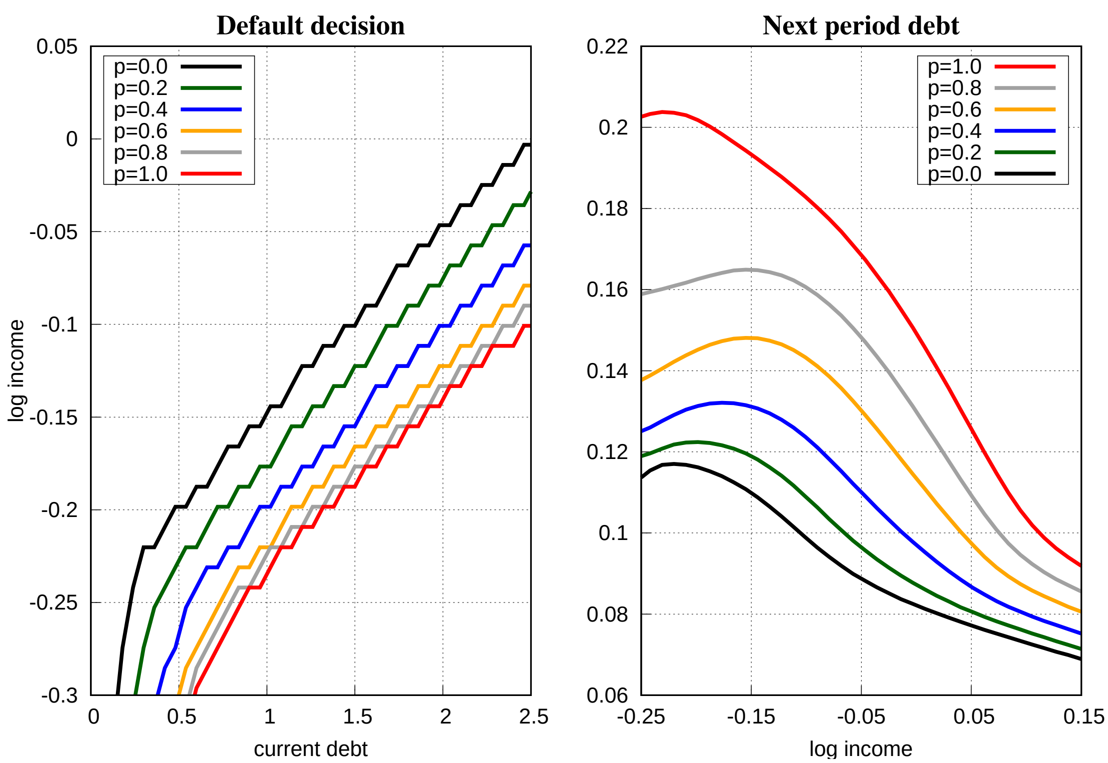

Figure 3: Default sets and bond price policy functions for different beliefs

Notice furthermore that policy functions are generally decreasing in income which implies a consumption smoothing behavior during normal times. It is only when income falls low enough, to a great-depression-like level, that the policy functions bend over and take an increasing shape, more common for this class of models, resulting in a procyclical fiscal policy.

Figure 4 plots government bond prices as functions of the next period debt choice, at several different levels of the belief.[23] The information about current regime is important in determining future default risk and leads to large differences in the offered bond prices. The highest (red) line represents the bond price schedule when markets are fully convinced the economy is in the high regime. As a result, the government is able to secure an almost maximum price for its bonds, regardless of its choice of next period debt (within reasonable bounds). By contrast, the lowest (black) line represents the schedule if the markets believe the economy is currently in a depression. Because default risk is much higher in such circumstances, the government is offered very low prices for its debt. Finally, the schedules in the mid-range are increasing monotonically as the belief of being in the high regime rises.

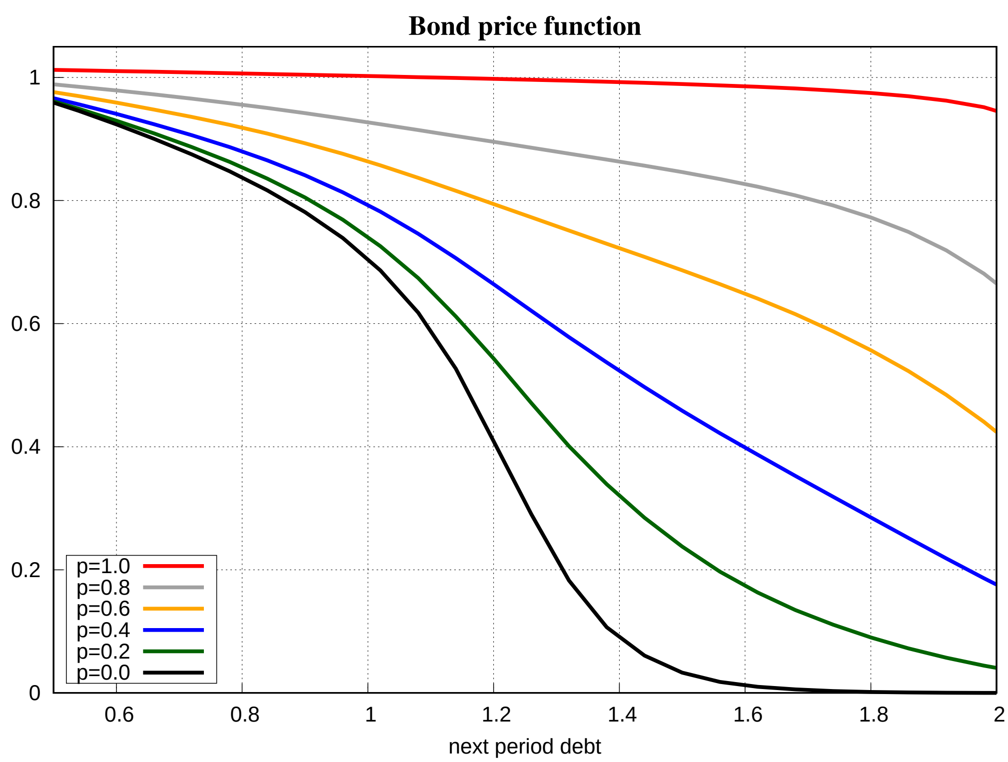

Figure 4: Bond price as function of next period debt for different beliefs

[23]This graph does not say anything about optimality of different debt choices, it merely depicts the possible price schedules.

4.4.2 Business cycle statistics

In the next step, I analyze the model's behavior in simulations. As discussed in Section 4.3, the model exhibits notably different behavior in the long run versus the short samples that aim at mimicking the period of 1998-2019 for the European economies.

Table 5 presents the simulated business cycle moments in the long- and short-run samples (i.e. ergodic and conditional distribution), along with the empirical ones. It can be noticed that the model simulated in the short run performs closer to actual data in terms of correlations between the main variables, and the moments of bond spread and debt. In particular, the average debt-to-GDP ratio for Portugal in the long run implied by the model is around 62%, much higher than in the data for 1998-2019. This is due to the fact that the government is gradually accumulating debt towards its steady-state level for most of the 2000s. Unlike the ergodic distribution, short-run simulations are able to capture this non-stationarity. Related to the policy functions discussed in Figure 3, notice that the government's ergodic behavior implies a counter-cyclical fiscal policy, with a positive correlation of income and trade balance (tb), and a low volatility of consumption relative to output. These features do not show up in the short-run statistics, however, due to the disproportionate presence of a rare disaster which makes the government behave more like a classic sovereign defaulter.

Table 5: Simulated behavior of the model

Statistic	**Data**	**Learning**		**AR1**	
		Ergodic	*Conditional*	*Ergodic*	*Conditional*
$E(s)$	1.75	1.12	1.75	1.93	1.75
$std(s)$	3.06	3.90	3.63	1.83	1.42
$std(c)/std(y)$	0.98	0.98	1.17	1.22	1.43
$std(tb)/std(y)$	0.55	0.26	0.39	0.35	0.78
$corr(y, c)$	0.98	0.97	0.95	0.97	0.86
$corr(y, tb)$	-0.96	0.21	-0.29	-0.52	-0.25
$corr(y, s)$	-0.42	-0.38	-0.71	-0.67	-0.69
$corr(s, tb)$	0.39	0.38	0.55	0.72	0.66
$E(debt/y)$	38.58	62.16	38.59	43.32	38.58

Note: Moments for the bond spread (annual), and debt-to-GDP ratio are given in percentage points. Ergodic (long-run) simulations extend to 10,000 quarters and are repeated 10,000 times, following closely Chatterjee and Eyigungor (2012). Conditional (short-run) simulations mimic the period of 1998-2019 (88 quarters) and are repeated 10,000 times starting from the actual levels of debt and GDP observed in 1998:Q1. Each short-run sample is constructed such that: i) the series start from the actual 1998:Q1 debt and income levels, and ii) the regime switches from good to bad in 2008:Q3. Consumption data is detrended using the common GDP trend.

4.4.3 Simulated behavior of spreads

Table 6 highlights the most notable difference between the two models. The first two moments of the simulated bond spreads are presented for each model, and contrasted with the data. The average spread is a targeted moment, so it is matched in both cases exactly. However, the standard deviation of the spread is not a calibration target, and for the AR1 model it falls short of the level observed in the data, producing a *coefficient of variation* of the spread smaller than 1 (reported in the last column). This result confirms the finding of Aguiar et al. (2016) who show that models of this type generally fail to deliver a realistic volatility of the bond spread. By contrast, the model with disaster risk and learning generates a standard deviation which actually exceeds what we observe in the data, resulting in a coefficient of variation above 1. Mirroring this result is the fact that the trade balance is less volatile in the disasters model than in the AR1 model, as evident in Table 5. By adjusting its trade balance more aggressively, the borrower in the standard model is able to target a desired level of bond spread with smaller variance.

Table 6: Bond spreads moments in the simulations and the data

	Bond spreads (in %)		
Country	μ	σ	c_v
Model-AR1	1.75	1.42	0.81
Model-learning	1.75	3.63	2.07
Data-Portugal	1.75	3.06	1.75

The intuition behind the result presented in Table 6 is the following. In the standard "off-the-shelf" AR1 model with long-term debt a sovereign default is possible most of the time, within the expected duration of an outstanding bond. Consequently, the spread never falls to zero, although it may on average be much smaller than the ones obtained for emerging market economies, as shown in Chatterjee and Eyigungor (2012) and other studies. On the other hand, in the model with disasters and learning, sovereign defaults occur almost exclusively when the economy has switched to the disaster regime.[24] Hence, most of the time while the economy is doing well and the market-wide belief is close to one, lenders do not fear that default is a possibility in any predictable future. As a result, bond spreads are very close to zero, and the average spread is low even if a debt crisis eventually does occur. Once that happens, the spread shoots up and can attain large values (which I explain in the subsequent paragraphs), resulting in an overall high standard deviation.

[24] In the long-run simulations summarized in Table 5, over 94% of all defaults occur in the low regime, and the average value of the belief at default is 0.06.

Table 7 generalizes this point by documenting the difference in bond spread moments of other peripheral European countries discussed in Section 2 and emerging market defaulters (which are the examples originally presented by Arellano (2008)). As can be noticed, the former tend to have a coefficient of variation of the bond spread above one, implying that spread volatility is high relative to its average. On the other hand, emerging market defaulters tend to have average spreads that exceed their standard deviations significantly, resulting in a coefficient of variation smaller than one. Consequently, as Table 6 shows, the standard AR1 model seems to be a better description of the debt crisis experienced by an emerging market economy, while the model with disaster risk and learning presented in this paper is a better description of the recent episode experienced by developed European nations.

Table 7: Bond spread statistics for European vs. emerging market economies

	Bond spreads (in %)		
European	μ	σ	c_v
Greece	3.90	5.94	1.52
Spain	0.98	1.33	1.36
Italy	1.05	1.21	1.15
Emerging	μ	σ	c_v
Argentina	10.25	5.58	0.54
Ecuador	16.91	10.72	0.63
Russia	19.41	17.60	0.91

Note: Bond spread moments for European countries are computed with OECD data covering 1999:Q1-2014:Q4, while the moments for emerging economies are taken from Arellano (2008).

Another interesting feature of the model is that the bond spreads can take much higher values on the equilibrium path than in the benchmark, i.e. the government sometimes sells bonds at deep discounts. To illustrate this point, Figure 5 compares the distributions of spreads realized in the simulations for the two models. In the model based on a simple AR1, spreads essentially do not carry any mass for values above 0.2, and the average maximum spread attained in the short-run simulations corresponding to the data sample is 9.2%. By contrast, the distribution of spreads in the model with disasters and learning features a long upper tail extending all the way to 100%, with the average maximum spread in the conditional distribution of 21%. The intuition behind this result is straightforward. In the former, the government targets a certain level of spread and only tolerates limited upward deviations. If the spread becomes too high, it must be the case that income is low enough and default becomes a more attractive option. In the latter, however, a sudden fall in the belief may cause a downward shift of the entire bond price schedule, as shown in Figure 4. Spreads may then shoot up, while income remains relatively high, making default unattractive because of the non-linear punishment function. As a result, the government sells bonds at steep discounts until one of the two outcomes occurs: either income falls enough to make default attractive, or debt is reduced enough and spreads return to a desirable level.

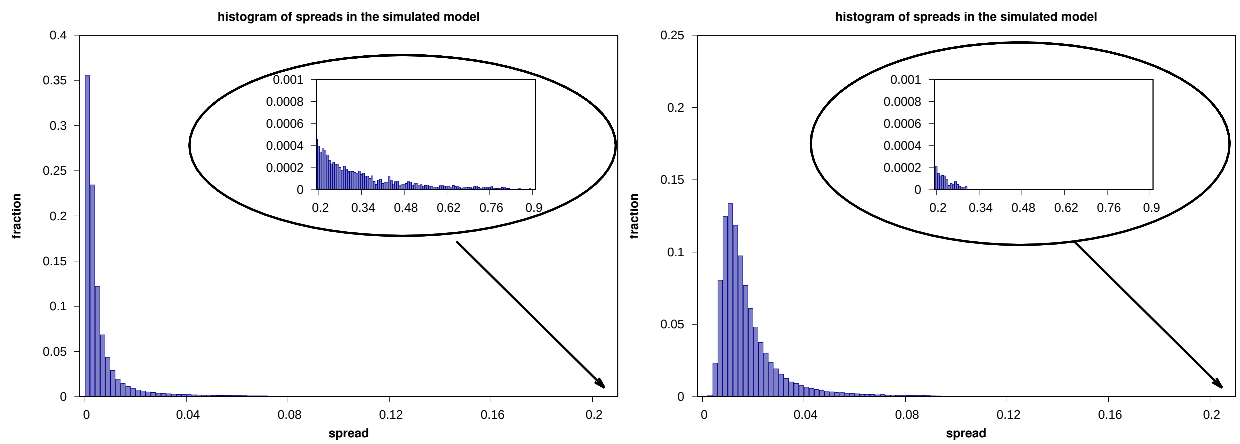

(a) Disasters and learning

(b) Simple AR1

Figure 5: Histogram of spreads in simulated models

The ability of the model to generate high values of the realized bond spread is akin to the result of Aguiar et al. (2020) who revisit the theory of rollover crises to rationalize the existence of "desperate deals". Crucially, here I show that a similar behavior can be obtained with fundamental factors. Appendix C provides a supplementary analysis of the drivers of equilibrium spreads. Specifically, it shows that movements in the belief are the main driver of spread volatility in the simulations. Moreover, this predominantly occurs through the belief's impact on the expected bond prices, rather than the next period default probability.

4.5 Event analysis of the European debt crisis

In this section, I use the calibrated model to conduct an event study of the debt crisis in Portugal. I start with the benchmark AR1 model, and then move on to the predictions of the main model. I conclude by showing a counterfactual where agents have full information.

4.5.1 Benchmark case - standard AR1

I start by feeding the actual detrended GDP observations for Portugal into the benchmark off-the-shelf AR1 version of the model. Figure 6 presents the predicted evolution of debt-to-GDP ratio and the bond spread in that model and in the data (this is the same time series as in Figure 1). The economy is started in the first quarter of 1998 with the actual debt and GDP levels from the data. I use the model-implied decision rules and bond prices to generate endogenous responses to the realized path of income shocks.

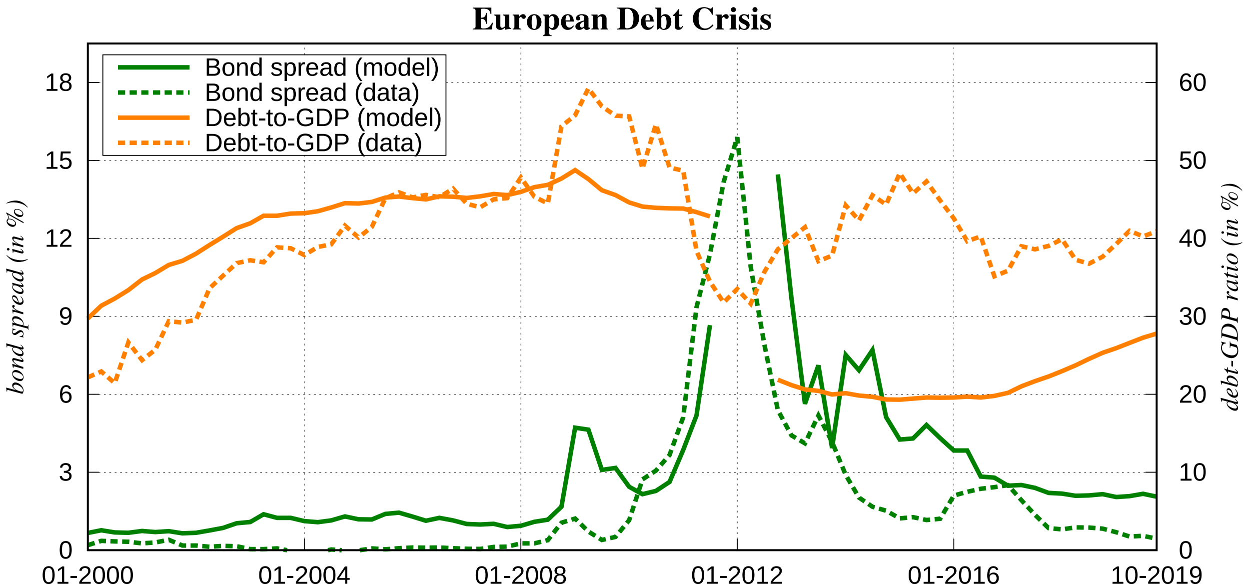

Note: To resolve the predicted default in this event study, I assume that the government re-enters the market with an exogenous debt write-off. This aims to mimic the aid implied in the emergency loans that Portugal received from the European Commission and the IMF. Quantitatively, the write-off amounts to about 20% *of the outstanding debt securities. Appendix D presents the details of this calculation.*

Figure 6: Event analysis in the model based on a standard AR1 process

Figure 6 highlights all of the problems with an off-the-shelf AR1 model that have been highlighted in the preceding sections. The government accumulates debt too fast, at least until 2003, due to the fact that the model requires a low value of the discount factor in order to hit the targets of average debt and average spread. Throughout this time, the bond spread is strictly positive at around 1% and actively responds to current income shocks, including the recession of 2002-2003. When the Great Recession starts, the spread jumps up to 4.7% in 2009:Q1 due to the extreme negative shock that hits Portuguese GDP. The government reduces its debt sharply which causes the spread to fall back. These predictions are at odds with the data - the bond spread was virtually zero up until 2009, and debt accumulation was slower. In 2009, the 5-year bond spread rose to 1.2%, while the government actually increased its external debt-GDP ratio. Finally, the model predicts a further rise in the spread and a sequence of defaults that start in the last quarter of 2011. Notably, the model falls short of replicating the observed peak spread level of 16% (the highest it can get before default is 8.7% in 2011:Q3), and mispredicts the key variables in the post-2016 recovery period.

4.5.2 Model with disaster risk and learning

Now I conduct an event analysis for my main model. The upper panel of Figure 7 presents the predicted evolution of debt-to-GDP ratio and bond spread for the actual path of Portuguese GDP realizations. The lower panel tracks the evolution of agents' belief about the economy type. The model offers an improvement on most fronts of the analysis relative to the benchmark in Figure 6. The pace of debt accumulation prior to 2008 is slower and consistent with the data due to the fact that we are able to use a much higher discount factor. At the same time, the predicted bond spread is essentially zero because investors have a near-certainty that the economy is operating in the high regime in which a default almost never happens. When the negative shock of 2009:Q1 hits, the belief drops on impact, but only partially, leading the spread to increase to 1.1% while debt is reduced in response, but then immediately picks up and stays on par with the data until the end of 2010. Finally, in 2011-2012 a new wave of low GDP shocks hits the economy which causes the belief to plunge. This contributes to a sharp increase in the bond spread which is predicted to reach 13% in 2012:Q2 and 23% in 2012:Q3, followed by a default. While in reality the Portuguese spread

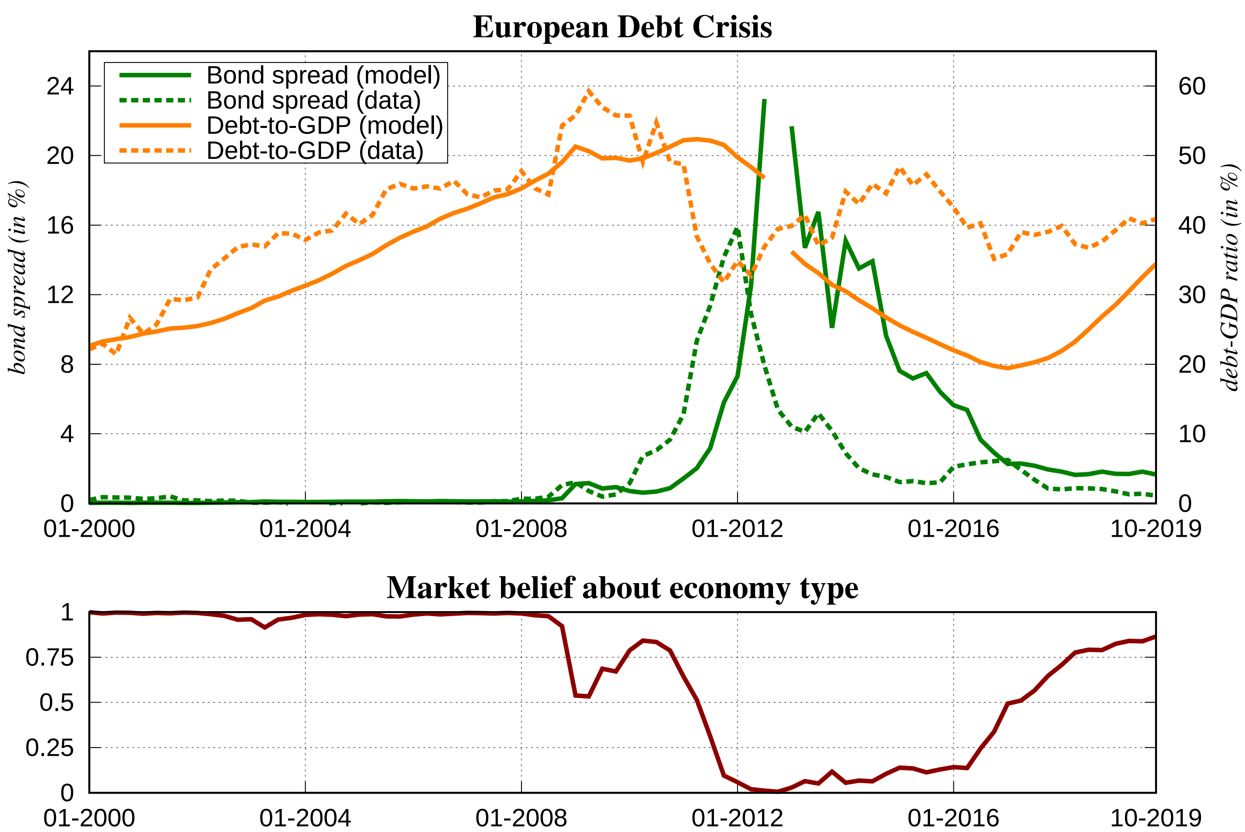

Note: To resolve the default, I apply the same assumption as described in Figure 6 and Appendix D.

Figure 7: Event analysis in the model with disaster risk and learning

did not go that high, its upward path was halted by the successful bailout provided by the IMF and the European Commission. By contrast, the Greek 10-year bond spread exceeded 27% in February 2012 which shows that such high values of the spread are feasible in equilibrium, and here I argue that they can be rationalized by a model based on fundamental factors.

In the aftermath of the default, the government is readmitted to the market and sharply reduces the debt further, due to the belief of being in the depression regime.[25] As the belief picks up again starting from 2016, the predicted bond spread and debt both converge towards their data counterparts for the same reasons as in the pre-2008 period. This shows that, outside the period of Portugal's participation in the bailout program, the theory can also fit the recovery that followed the crisis, especially after 2016.[26]

The contrast in the predictions captured by Figures 6 and 7 stems from the difference in the path of expectations about future income. Appendix C.6 constructs a counterpart to Figure 2 using both variants of the model. Unlike an off-the-shelf AR1 variant, the learning model generates the pattern of gradual forecast revisions that resembles the evidence in Section 2.2.

Naturally, the model still falls short of replicating the data in a few ways. First, the government in the model reduces its debt in response to the shocks in 2009, which is due to the elevated likelihood of an underlying regime switch. In reality, most European countries actively *increased* their debt during that time, as Figure 1 shows. Paluszynski and Stefanidis (2020) argue that such "borrowing into debt crises" behavior may be due to frictions in adjusting government expenditure. Second, the model predicts a sovereign default in 2012 which is not unreasonable, given that Portugal received a bailout from the European Commission and the IMF covering over 40% of its GDP. How to model a lender of last resort for sovereigns is a subject of active debate, and thus I leave this extension for future research.

4.6 Learning matters: a counterfactual with full information

Now I show that learning is crucial for the model to generate sensible predictions about the actual debt crisis in Europe. Figure 8 presents the event study with a variant of the model in which agents have complete information about the underlying regime switches (Appendix

[25]The lack of such debt reduction in the data may be due to Portugal's continual participation in the EU-IMF bailout program, a force that is absent from the model. Notice that as Portugal officially exits the program in mid 2014, it begins to reduce its debt securities. This lasts until mid 2016 when the government, in the model and in the data alike, starts accumulating debt again.

[26]A few crucial factors are missing in the model to explain years 2012-2016, as highlighted by many of the studies in the introduction, in particular the negotiations over EU bailouts and equilibrium multiplicity.

E describes the calibration details for this variant). As in the previous case, debt increases gradually prior to the Great Recession, while spreads remain at zero. In 2008:Q3, upon learning about the regime switch, spread increases to 3.3% while the government embarks on a drastic debt reduction path. Interestingly, even though the debt is much lower in the second stage of the European crisis relative to the predictions in Figure 6 and Figure 7, the spread still shoots up to 33% in 2012:Q4 without causing a default. This is because the agents are aware that the economy is on a downward trajectory, but there are no belief swings that would magnify the income shock and push the government into default.

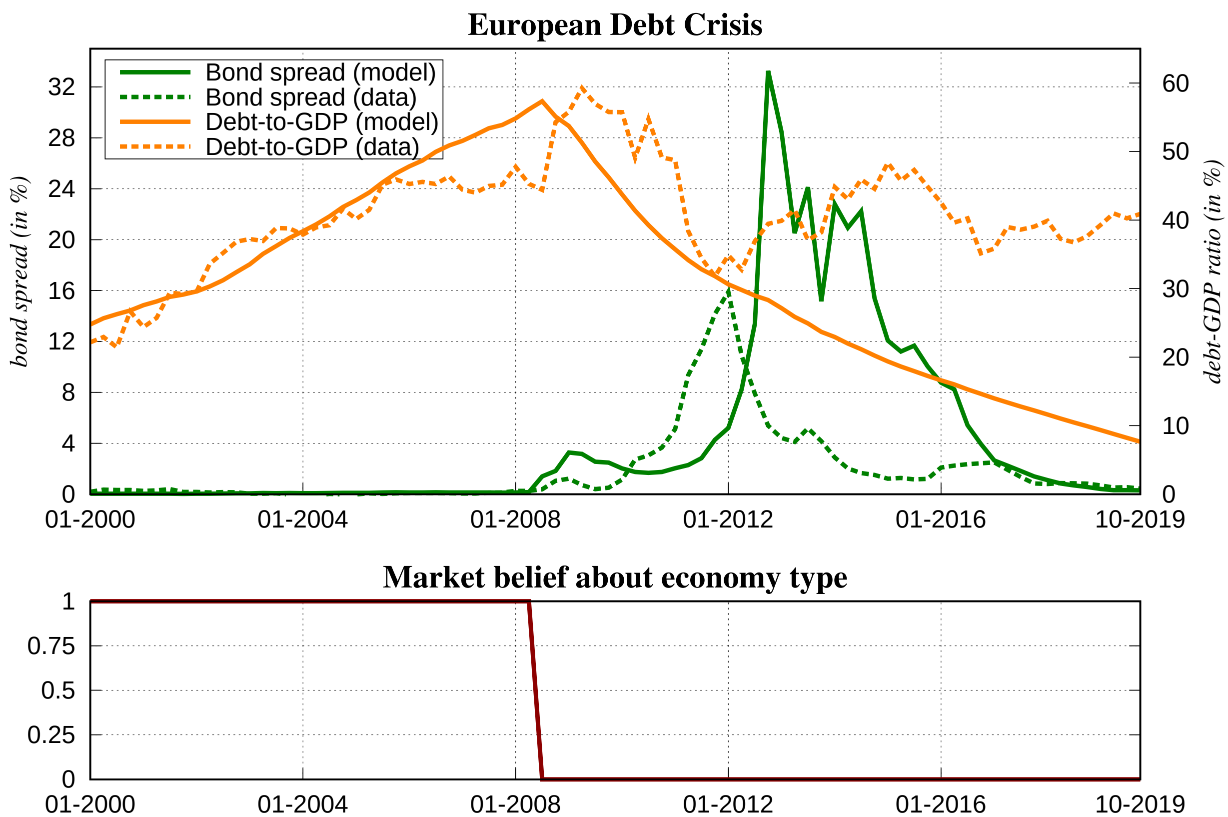

Figure 8: Event analysis in the model with full information about the disaster regime

5 Conclusion

In their seminal contribution, Lucas and Sargent (1979) make the following remark about general equilibrium macroeconomic models:

> *It has been only a matter of analytical convenience and not of necessity that equilibrium models have used the assumption of stochastically stationary shocks and the assumption that agents have already learned the probability distributions they face. Both of these assumptions can be abandoned, albeit at a cost in terms of the simplicity of the model.*

This paper shows that learning about the probability distributions of future income shocks was an important driver of the European debt crisis. It impacted not only the movements in asset prices, but also the real variables such as government debt. I show that an otherwise standard quantitative model of sovereign debt can be augmented to incorporate this learning process and match the evidence on the gradually evolving beliefs over time. As a result, we can obtain a delayed pattern of bond spread increases during the Great Recession in Europe.

References

Aguiar, Mark, Satyajit Chatterjee, Harold L. Cole, and Zachary Stangebye (2020): "Self-Fulfilling Debt Crises, Revisited", *unpublished manuscript*.

Aguiar, Mark, Satyajit Chatterjee, Harold L. Cole, and Zachary Stangebye (2016): "Quantitative Models of Sovereign Debt Crises", *Handbook of Macroeconomics*, Elsevier.

Aguiar, Mark and Gita Gopinath (2006): "Defaultable debt, interest rates and the current account", *Journal of International Economics* 69, 64-83.

Aguiar, Mark and Manuel Amador (2020): "Self-fulfilling Debt Dilution: Maturity and Multiplicity in Debt Models", *American Economic Review* 110, 2783-2818.

Andritzky, Jochen (2012): "Government Bonds and Their Investors: What Are the Facts and Do They Matter?", *IMF Working Paper*, 12/158.

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

Ayres, Joao, Gaston Navarro, Juan Pablo Nicolini and Pedro Teles (2019): "Self-Fulfilling Debt Crises with Long Stagnations", *unpublished manuscript*.

Bai, Jushan and Pierre Perron (1998): "Estimating and Testing Linear Models with Multiple Structural Changes," *Econometrica*, 66, 47-78.

Barro, Robert J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century", *Quarterly Journal of Economics*, 121, 823-866.

Batchelor, Roy (2001): "How useful are the forecasts of intergovernmental agencies? The IMF and OECD versus the consensus", *Applied Economics*, 33, 225-235.

Bocola, Luigi and Alessandro Dovis (2019): "Self-Fulfilling Debt Crises: A Quantitative Analysis", *American Economic Review*, 109, 4343-77.

Bocola, Luigi, Gideon Bornstein and Alessandro Dovis (2019): "Quantitative Sovereign Default Models and the European Debt Crisis", *Journal of International Economics*, 118, 20-30.

Boz, Emine, Christian Daude and C. Bora Durdu (2011): "Emerging markets business cycles: Learning about the trend", *Journal of Monetary Economics*, 58, 616-631.

Boz, Emine and Enrique Mendoza (2014): "Financial innovation, the discovery of risk, and the U.S. credit crisis", *Journal of Monetary Economics*, 62, 1-22.

Brunnermeier, Markus K., Harold James and Jean-Pierre Landau (2016): "The Euro and the Battle of Ideas", *Princeton University Press.*

Chatterjee, Satyajit and Burcu Eyigungor (2012): "Maturity, Indebtedness, and Default Risk", *American Economic Review*, 102, 2674-2699.

Chatterjee, Satyajit and Burcu Eyigungor (2019): "Endogenous political turnover and fluctuations in sovereign default risk", *Journal of International Economics*, 117, 37-50.

Cimadomo, Jacopo, Peter Claeys, and Marcos Poplawski-Ribeiro (2016): "How do experts forecast sovereign spreads?", *European Economic Review*, 87, 216-235.

Coibion, Olivier, Yuriy Gorodnichenko, and Johannes Wieland (2012): "Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?", *Review of Economic Studies*, 79, 1371-1406.

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

Eaton, Jonathan and Mark Gersovitz (1981): "Debt with potential repudiation", *Review of Economic Studies* 48, 289-309.

Habermann, Christian, and Fabian Kindermann (2007): "Multidimensional Spline Interpolation: Theory and Applications", *Computational Economics* 30, 153-169.

Hamilton, James D. (1990): "Analysis of time series subject to changes in regime", *Journal of Econometrics* 45, 39-70.

Hatchondo, Juan Carlos and Leonardo Martinez (2009): "Long-duration bonds and sovereign defaults", *Journal of International Economics*, 79, 1717-125.

Hatchondo, Juan Carlos, Leonardo Martinez and Horacio Sapriza (2010): "Quantitative properties of sovereign default models: Solution methods matter", *Review of Economic Dynamics*, 13, 919-933.

Kehoe, Timothy J. and Edward C. Prescott (2002): "Great Depressions of the Twentieth Century", edited volume of *Review of Economic Dynamics*, 5, 1-18.

Kim, Chang-Jin (1994): "Dynamic Linear Models with Markov-switching", *Journal of Econometrics* 60, 1-22.

Lorenzoni, Guido, and Iván Werning (2019): "Slow Moving Debt Crises", *American Economic Review*, 109, 3229-3263.

Lucas, Robert E. and Thomas J. Sargent (1979): "After Keynesian Macroeconomics", *Federal Reserve Bank of Minneapolis Quarterly Review*, 3, 1-16.

Pain, Nigel, Christine Lewis, Thai-Thanh Dang, Yosuke Jin, and Pete Richardson (2014): "OECD Forecasts During and After the Financial Crisis: A Post Mortem", *OECD Economics Department Working Papers*, 1107.

Paluszynski, Radoslaw (2019): "The Ultralong Sovereign Default Risk", *unpublished manuscript.*

Paluszynski, Radoslaw and Georgios Stefanidis (2020): "Borrowing into Debt Crises", *unpublished manuscript.*

Rebelo, Sergio, Neng Wang, and Jinqiang Yang (2019): "Rare Disasters, Financial Development, and Sovereign Debt", *unpublished manuscript.*

Reinhart, Carmen M. and Kenneth S. Rogoff (2009): "This Time is Different: Eight Centuries of Financial Folly", *Princeton University Press.*

Reis, Ricardo (2013): "The Portuguese Slump and Crash and the Euro Crisis", *Brookings Papers on Economic Activity*, 46, 143-193.

Rietz, Thomas (1988): "The Equity Risk Premium: A Solution", *Journal of Monetary Economics*, 22, 117-131.

Salomao, Juliana (2017): "Sovereign Debt Renegotiation and Credit Default Swaps", *Journal of Monetary Economics*, 90, 50-63.

Standard & Poor's (2014): "Sovereign Rating and Country T&C Assessment Histories", *McGraw Hill Financial*, February 5, 2014.

Appendix (for online publication)

A Supporting evidence

A.1 Dynamics of sovereign ratings over time

Figure 9 presents plots of the countries' real GDP, together with the "sovereign bond rating index".[27] Even though the economies entered a recession as early as in the second quarter of 2008, markets continued to perceive their bonds as relatively risk-free investments until about two years later. As a result, only around 2010-2011 do we observe a sequence of sovereign rating downgrades among peripheral European countries, indicating that market expectations about the sustainability of governments' debt had deteriorated significantly.

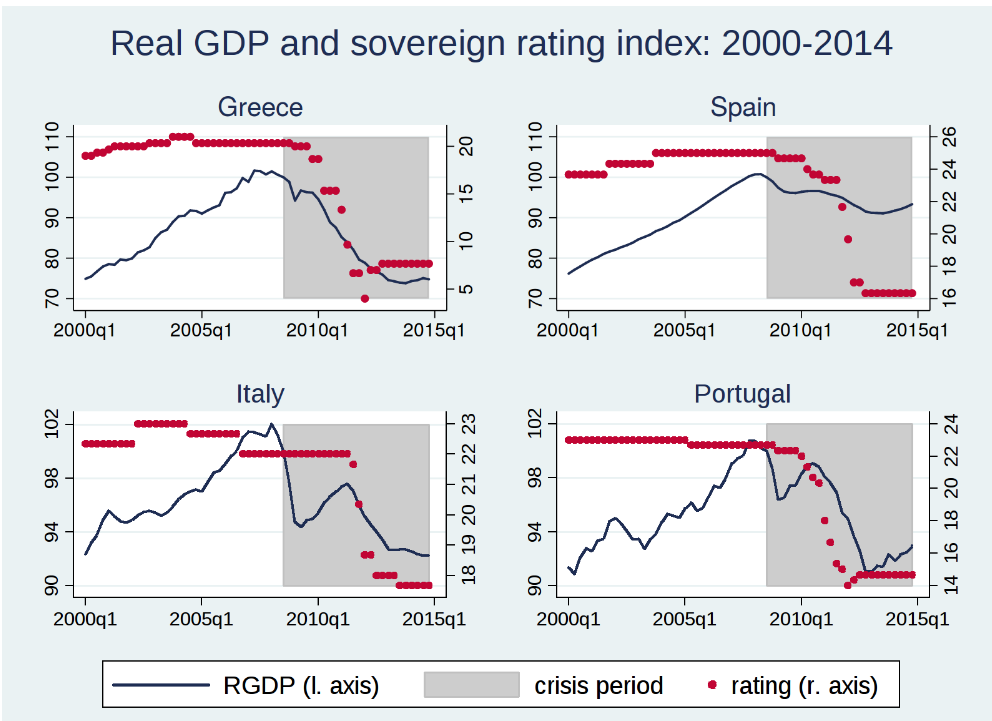

Note: The GDP series are in constant 2010 prices, and their values are normalized such that the third quarter of 2008 equals 100 (beginning of the financial crisis - shaded area). The bond rating index is constructed by converting the sovereign ratings of the three leading agencies - S&P, Moody's and Fitch - into a numerical scale from 0 to 25 and computing a simple average.

Figure 9: Real GDP and sovereign bond rating index of the European economies: 2000-2014

[27] The index is a simple weighted average of the three leading rating agencies (S&P, Moody's and Fitch), converted into a numerical scale from 0 to 25.

A.2 Learning about debt crises in emerging market economies

Aguiar and Gopinath (2006) show that emerging market bond spreads are better captured by a model with permanent income growth shocks, rather than transitory ones. It is natural to ask how the model developed in the present paper can be applied to the previous debt crises. Figure 10 plots the evolution of GDP forecasts for Argentina, a representative emerging market economy, around the time of its crash in 2001.[28] Two differences stand out in comparison with Figure 2 which plots analogous data for the European countries. First, even though Argentina's contraction is equally steep and much deeper than the ones of Portugal or Italy, a swift recovery follows. Unlike the European economies, Argentina returns to its peak output level of 1998 within six years. Second, forecasts for Argentina have an almost invariant slope over time, regardless of whether the economy is currently in a boom or a bust. As a result, large forecast errors arise most of the time, overestimating future GDP during a recession and underestimating it during a recovery. This suggests that forecasters have noisy information about Argentina's economy and form their projections based on the average long-run trend growth. Thus, while adding a regime-switching process with learning

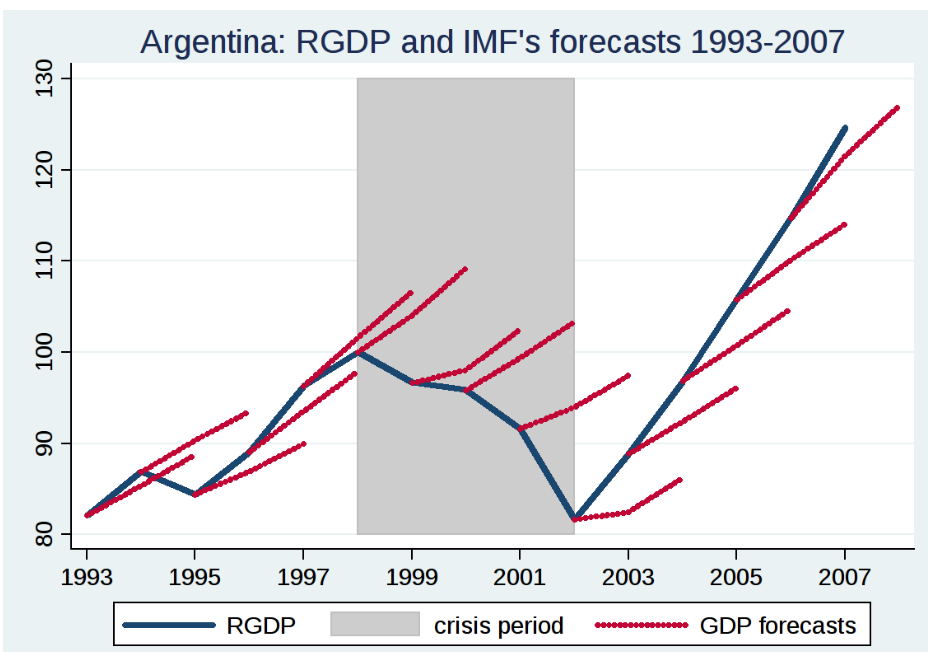

Note: The GDP series is annual and in constant prices; values are normalized so that it equals 100 in 1998. The red dotted lines represent one- and two-year ahead forecasts published in the fall of each year by IMF. The shaded area marks Argentina's debt crisis of 1998-2002.

Figure 10: Forecast and actual real GDP for Argentina

[28] I use IMF projections as it is the only source of forecast data out of the four I discuss in Section 2.2 that contains Argentina.

about its realizations to a model of emerging market debt crises is a possibility, there is relatively little information to be inferred from historical forecast data.

A.3 Further evidence on forecast errors

Table 8 documents the root mean square errors (RMSE) observed for each of the four countries of interest. As is evident, the RMSE in each case follows the same pattern as the average bias in Table 2, i.e. it increases significantly during the first stage of the recession (2008-2011), and then falls back (often below the pre-2007 level) during the second stage (2012-2014).

Table 8: Root mean square errors in real-time historical forecasts for different time frames

Root mean square error	**OECD**	**IMF**	**EC**	**CE**
(a) Pre-recession sample: 2000-2007				
Greece	2.18	2.25	2.18	2.13
Spain	0.77	0.94	1.02	0.97
Italy	1.43	1.56	1.47	1.43
Portugal	1.54	1.46	1.43	1.51
(b) Recession - first stage: 2008-2011				
Greece	6.69	7.01	7.12	6.90
Spain	2.98	3.33	3.30	3.15
Italy	3.39	3.48	3.57	3.45
Portugal	2.62	2.82	2.60	2.82
(c) Recession - second stage: 2012-2014				
Greece	1.60	1.88	1.89	1.45
Spain	1.36	1.63	1.23	1.35
Italy	0.79	0.85	1.13	0.94
Portugal	0.48	0.52	0.53	0.65

Note: The table presents root mean square errors of one-year-ahead forecasts of real GDP level. Forecasts are acquired from four sources: OECD, IMF, European Commission, and Consensus Economics Inc. The error is expressed as a percentage of the 2010 level of real GDP for each country. All forecasts come in two vintages, Spring and Fall, which I use jointly. The number of forecasters participating in Consensus Economics surveys varies over time and across countries, with a minimum of four and a maximum of twenty.

B Data Appendix

B.1 Detrending method

Figure 11 illustrates the detrending using a broken linear trend. For the baseline model, the trend is calculated until 2008:Q2 (the subsequent decline is assumed to be due to a regime shift, which is subsequently validated in Section B.3 with full-sample Bayesian inferences), while for calibration of the simple AR1 the trend includes data until 2011:Q4. In each case, two statistically significant breakpoints are detected using the Bai-Perron test (Bai and Perron, 1998), and continuity of the trend line is imposed. In both cases, the two breakpoints are detected at 1974:Q2 and 1999:Q4, coinciding with the democratic revolution in Portugal and adoption of the Euro, respectively. The estimated quarterly trend growth rates for the three time windows are 1.6%, 0.8% and 0.4% for the baseline case, and 1.6%, 0.8% and 0.3% for the simple AR1 case, respectively.

Detrending the data using a broken linear trend allows me to use a longer time series in the estimation, going back to 1960. This would not be possible with a single linear trend as the resulting residual would not be stationary. Including all the available information since 1960 is important to capture the full volatility of business cycles during the "normal times" regime, featuring regular expansions and recessions (the GDP data for European countries

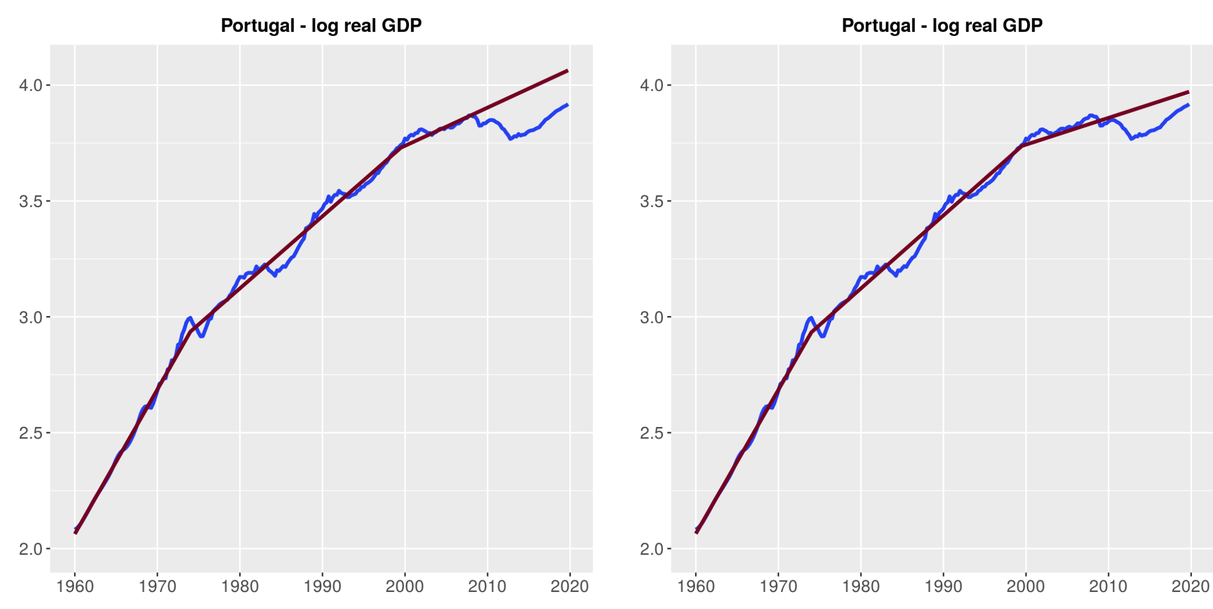

(a) Baseline case

(b) Simple AR1 case

Figure 11: Detrending the GDP using a broken linear trend

exhibits very little variance in years 1999-2008).

B.2 Estimation technique

In this section, I describe my approach to estimating the parameters of the regime-switching process (1). This is an extension of the Expectation-Maximization algorithm as in Hamilton (1990). The main novelty is that I use two separate data sources to inform the parameters: historical GDP series (as in the standard estimation), as well as the real-time GDP forecasts (to capture the evolution of market expectations documented in Section 2.2).

The procedure starts by fixing the regime-switching probabilities with recent historical experience. I also normalize the high-regime mean μ_H to zero, following the standard assumption. Denote the set of realized GDP data as $\mathcal{Y}_\mathcal{T} = \{y_1, y_2, ..., y_T\}$, and denote the set of observed forecasts as $\hat{\mathcal{Y}}_{T_f} = \{\hat{y}_1, \hat{y}_2, ..., \hat{y}_{T_f}\}$. Note that $T \neq T_f$ because the forecasts are not available for the same sample period as actual GDP data and only come in two vintages per year.[29] Notice also that while both y_t and $\hat{y}_t$ are measured as log deviations from trend, the former refers to quarterly GDP, while the latter refers to annual GDP. I assume that the forecasts observed in the data are made by the agents in my model, with noise (which is necessary to ensure that the historical GDP data, and the forecast data, are jointly consistent with Bayes' rule). Aggregating the model-generated forecasts to annual level we get:

$$\hat{y}_t = log\Big(\sum_{t_i=0}^{3} \frac{(1+g)^{t_i}}{\sum_{t_j=0}^{3}(1+g)^{t_j}} exp\Big(\mathbb{E}[y_{t+n+t_i}|\mathcal{Y}_t]\Big)\Big) + u_t \tag{12}$$

In formula (12), n denotes the number of quarters ahead until the first quarter of the year the forecast refers to, g represents the quarterly trend growth rate (the estimation of which is described in Appendix B.1), and $u_t \sim \mathcal{N}(0, \sigma_u^2)$ is an i.i.d. forecast error. Let $\mathbf{p}_t \equiv [1-p_t,\ p_t]$ be the agents' belief vector in period t, let $\mathbf{m} \equiv (1-\rho)[\mu_L,\ \mu_H]$, and let Π be the transition matrix as defined in formula (2). Making sure that all eigenvalues of $\rho\Pi^{-1}$ are smaller than 1, the agents' (detrended) forecast for i quarters ahead is given by

$$\mathbb{E}[y_{t+i}|y_t] = \rho^i y_t + \mathbf{p}_t'\Pi^i\Big[(I - \Pi^{-1}\rho)^{-1}\Big(I - (\Pi^{-1}\rho)^i\Big)\Big]\mathbf{m}$$

Denote $z_t \in \{L, H\}$ as a regime realization in period t. Taking as given a sequence of smoothed full-sample beliefs $Prob(z_t = i|\mathcal{Y}_T)$ (inferred by the econometrician), I then pose

[29] I associate the spring vintage with Q1, and the fall vintage with Q3.

an expected log-likelihood that takes into account the forecast errors

$$\mathbb{E}\Big(\ell_{Y,S,\hat{Y}}|\mathcal{Y}_T,\theta\Big) = \sum_{t=1}^{T}\sum_{i=L,H} Prob\Big(z_t = i|\mathcal{Y}_T\Big) \times \left\{-log\Big(\sqrt{2\pi\sigma^2}\Big) - \frac{(y_t - m(z_t) - \rho y_{t-1})^2}{2\sigma^2}\right\}$$
$$+\sum_{t=1}^{T_f}\left\{-log\Big(\sqrt{2\pi\sigma_u^2}\Big) - \frac{\Big[\hat{y}_t - log\Big(\sum_{t_i=0}^{3}\frac{(1+g)^{t_i}}{\sum_{t_j=0}^{3}(1+g)^{t_j}}exp\Big(\mathbb{E}[y_{t+n+t_i}|\mathcal{Y}_t]\Big)\Big)\Big]^2}{2\sigma_u^2}\right\} \tag{13}$$

As is standard, the algorithm then iterates on the following two steps until convergence:

1. *(maximization)* Taking as given the previous-iteration parameter vector $\boldsymbol{\theta}_0 \equiv \{\mu_{L,0}, \rho_0, \sigma_0^2, \sigma_{u,0}^2\}$ and the full-sample smoothed probabilities of the two regimes, $Prob(s_t = i|\mathcal{Y}_T, \theta_0)$ for $i \in \{L, H\}$, find the new parameter vector θ_1 that maximizes the expected log-likelihood function in (13). In particular, this involves solving numerically for ρ_1 and $\mu_{L,1}$.

2. *(expectation)* Given the new parameter vector $\boldsymbol{\theta}_1$, update the full-sample smoothed probabilities of the regimes as in Kim (1994).

Steps 1-2 are repeated until $|\boldsymbol{\theta}_0 - \boldsymbol{\theta}_1| < \varepsilon$ for some convergence criterion ε.

B.3 Estimation for other countries

In this section, I present the results of applying the estimation method described in Section 4.2 to all four southern European countries. The purpose is to check how applicable the model is to the remaining cases which jointly motivate the paper. For each country, I maintain the assumption that the expected duration for the high regime and low regime is 60 years and 10 years, respectively. Then, I estimate the remaining parameters using the sample of GDP data for 1960:Q1-2019:Q4, and the sample of GDP forecasts in years 1993-2014, using the algorithm described in Section B.2. Table 9 summarizes the obtained parameter values. The persistence and standard deviation parameters are in line with the common estimates for European economies. On the other hand, the estimated disaster regime means range from around 25% below trend for Portugal and Italy to roughly 60% below trend for Greece and Spain.

Figure 12 overlays the paths of forecasts from the model and the data for all four countries,[30] in a detrended form, along with the paths of realized data 5 years later (the series ends

[30] The discrepancies between the model-generated forecasts and the data forecasts are due to forecasting error.

on 2014 because, as of writing, complete GDP data is only available until 2019). The fit of model-generated forecasts is overall good, capturing the progression of pessimism in the projections over time. Similarly to Figure 2, all panels confirm that there was a clear learning process for each country. In particular, the gaps between the forecasts and the GDP data realized five years later tend to be large early on and shrink over time, possibly turning negative after 2012.

Table 9: Calibrated parameters of the regime-switching endowment process for all countries

Regime	Mean μ	Persistence ρ	St. dev. η	St. dev. σ_u
Greece	-0.779	0.924	0.021	0.056
Spain	-0.966	0.993	0.010	0.035
Italy	-0.286	0.955	0.008	0.045
Portugal	-0.291	0.970	0.010	0.049

Note: parameters are estimated independently for each country following the procedure described in Section 4.2. In each case, the regime-switching probabilities are the same as in Table 3. The GDP data for each country is detrended in the same fashion as it is described for Portugal in Appendix B.1.

Figure 13 presents the inferred paths of the belief about being in the disaster regime for each country. The filtered belief is the real-time Bayesian probability that the agents use in the model; while the smoothed probabilities refer to full-sample inferences which are calculated as in Kim (1994). The lower panels show the entire time period from 1960 to 2019, while the upper panels focus on the most recent episode of interest, since 2000. The figures confirm that fluctuations in the belief are generally rare and revert instantly for any time period before the Great Recession. The analysis also indicates that the model is likely to be applicable to Greece, in addition to Portugal. For both of these countries the belief about being in the high regime drops part ways on impact in 2009:Q1, and subsequently recovers before collapsing all the way to zero. By contrast, for Italy and Spain the belief drops most of the way in 2009:Q1 which possibly leads the agents in the model to *underpredict* future GDP level in that time period (Figure 12). This does not necessarily mean that the theory of gradual learning is not applicable to Italy or Spain, but rather that the model and the proposed calibration technique is not able to capture the slow learning process. One way to overcome this problem would be to augment the income process with a third regime. In this scenario, the first two regimes would have a standard expansion/recession interpretation, with frequent transitions between them, while the third regime would be a rare disaster. The “regular recession” regime would then help matching the forecasts in 2009:Q1 without an instant switch to the depression one.

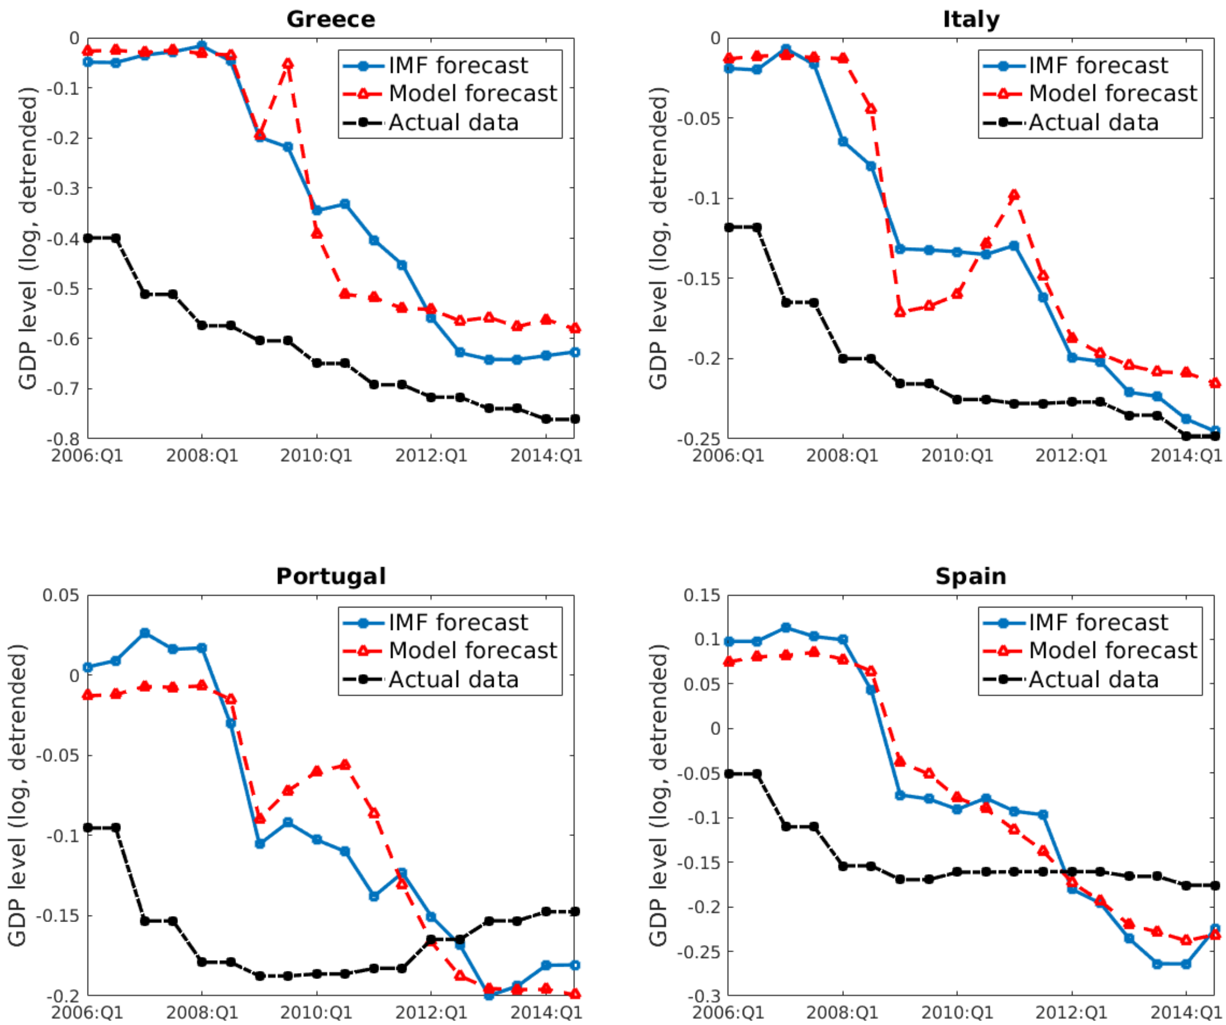

Note: Each point on the graph represents an annual detrended log-GDP level for five years ahead (for example, 2008:Q1 corresponds to the GDP level in 2013). The solid blue line represents the actual published forecasts (Q1 and Q3 refer to the spring and fall issues, respectively), while the dashed red line denotes the ones generated by the calibrated model. The dashed-dot black line shows the actual realized data that the corresponding forecasts refer to (only available until 2014:Q3, when the projection for year 2019 was made).

Figure 12: IMF- and model-generated projections for all countries

It is also worth reiterating, as Section 4.1 explains, that Spain and Italy are not necessarily the best countries to explain using the model in this paper. The reason is that both had large stocks of domestic, rather than external, debt and the height of their crisis coincided with the unprecedented actions of the European Central Banks in the summer of 2012. This makes explanations based on rollover crises, such as in Bocola and Dovis (2019) or Aguiar et al. (2020), or domestic default as in Bocola, Bornstein and Dovis (2019) a more promising

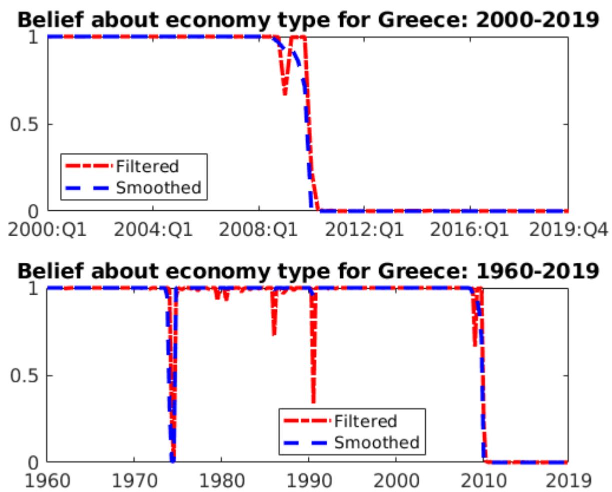

(a) Greece

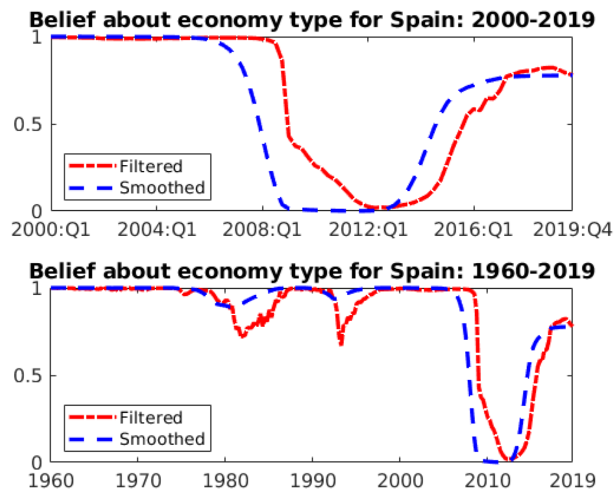

(b) Spain

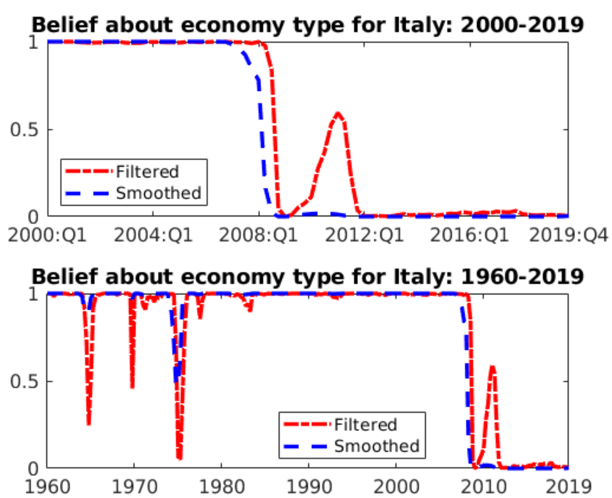

(c) Italy

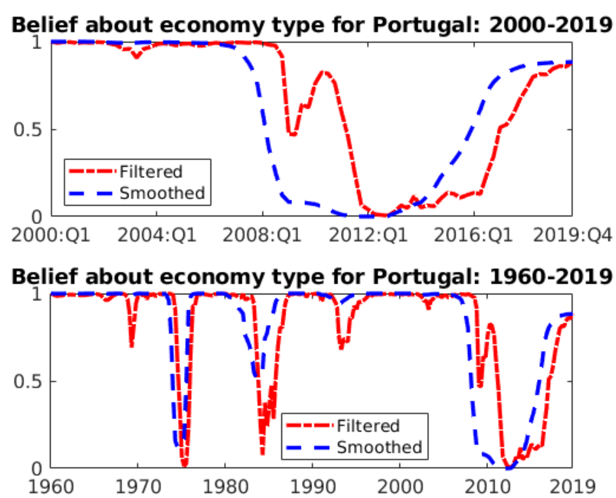

(d) Portugal

Figure 13: Inferred paths of beliefs for each country

avenue for rationalizing the events in these countries. On the other hand, countries such as Greece and Portugal are a better target for analysis using this model and the learning process naturally turns out to be more relevant for them.

C Supplementary analysis of bond spreads

In this section, I present additional results about bond spreads generated by the baseline model.

C.1 Determinants of simulated spreads

I start by showing that the highest spreads in the model are driven by the collapse of the belief. Figure 14 presents a heatmap of simulated spreads in the baseline model with respect to income and belief. The graph confirms the main points of the analysis in Section 4.4.3. As long as the belief is high enough (around 0.7 and above), the spread remains negligible. By contrast, the highest spreads in equilibrium are concentrated in the states where the belief is close to zero and (log) income is around -0.17. This number is noteworthy: it is considerably higher than $\hat{y} = -0.23$, the level of (log) income above which default imposes a direct cost to income (marked on the graph with a vertical line). As we move towards that threshold, the spreads remain elevated, but they no longer attain the highest values. This implies that the largest spreads that can be realized in this model occur when the belief is close to zero (such that agents have no doubts that the economy is in a depression regime), while income remains relatively high. A default is very costly in such a state, and the government sometimes chooses to tolerate unusually high spreads until it either manages to reduce its debt enough, or until income reaches a level at which the default cost is milder.

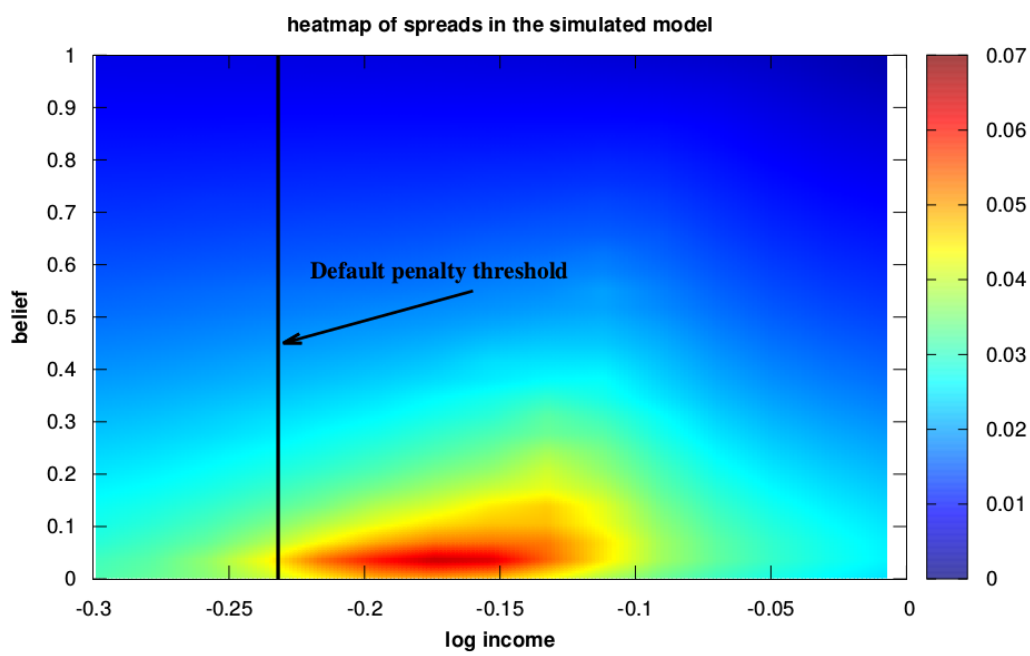

Figure 14: Heatmap of simulated spreads in the model

Figure 15 reinforces this point by comparing the scatter plots of simulated spreads with respect to (log) income in the baseline model to the standard AR1 model. In the latter, spreads tend to increase monotonically, as income declines, up to the default penalty threshold $\hat{y}$. In particular, the highest spreads on the equilibrium path tend to materialize for income realizations that are just above the threshold. In the former, on the other hand, there is no such apparent monotonicity. The highest spreads tend to occur *away* from the threshold $\hat{y} = -0.23$, starting already around (log) income of -0.1. Another way to appreciate this fact is to note that the unconditional correlation between spreads and (log) income in all of the simulations is negative, -0.35, which is unsurprising and in line with what the standard model would deliver. However, when we condition on the largest spreads, say greater than 0.5, that correlation becomes positive, 0.12, meaning that the largest spreads are expected for higher income levels.

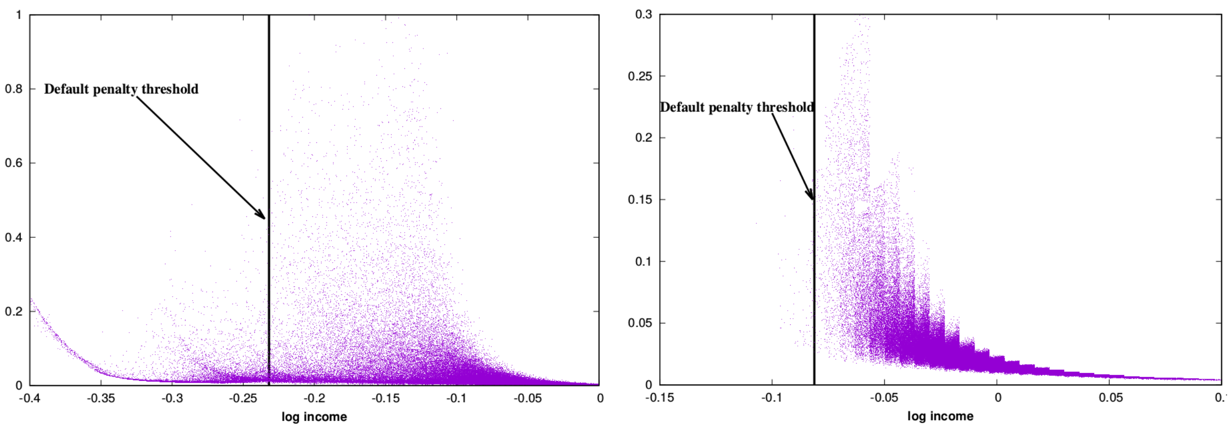

(a) Disasters and learning

(b) Simple AR1

Figure 15: Scatter plot of simulated spreads in the two models

C.2 Analysis of highest spreads

The ability of the baseline model to deliver high spread values on equilibrium path is unusual among the quantitative sovereign default models, and thus deserves more attention. Figure 16 plots the averaged simulated paths of key exogenous and endogenous variables around peak spreads (normalized to $t = 0$) that are greater than 0.5. On average, such a peak amounts to around 0.7 and coincides with a slump in the belief to 0.1. Consistent with the previous analysis, this tends to occur at rather high levels of (log) income around -0.16, significantly above the default penalty threshold. The government tends to deleverage sharply around the peak, on average reducing its debt throughout the episode. In particular, at peak spread, the government almost does not borrow any new debt above the current outstanding stock of long-term bonds.

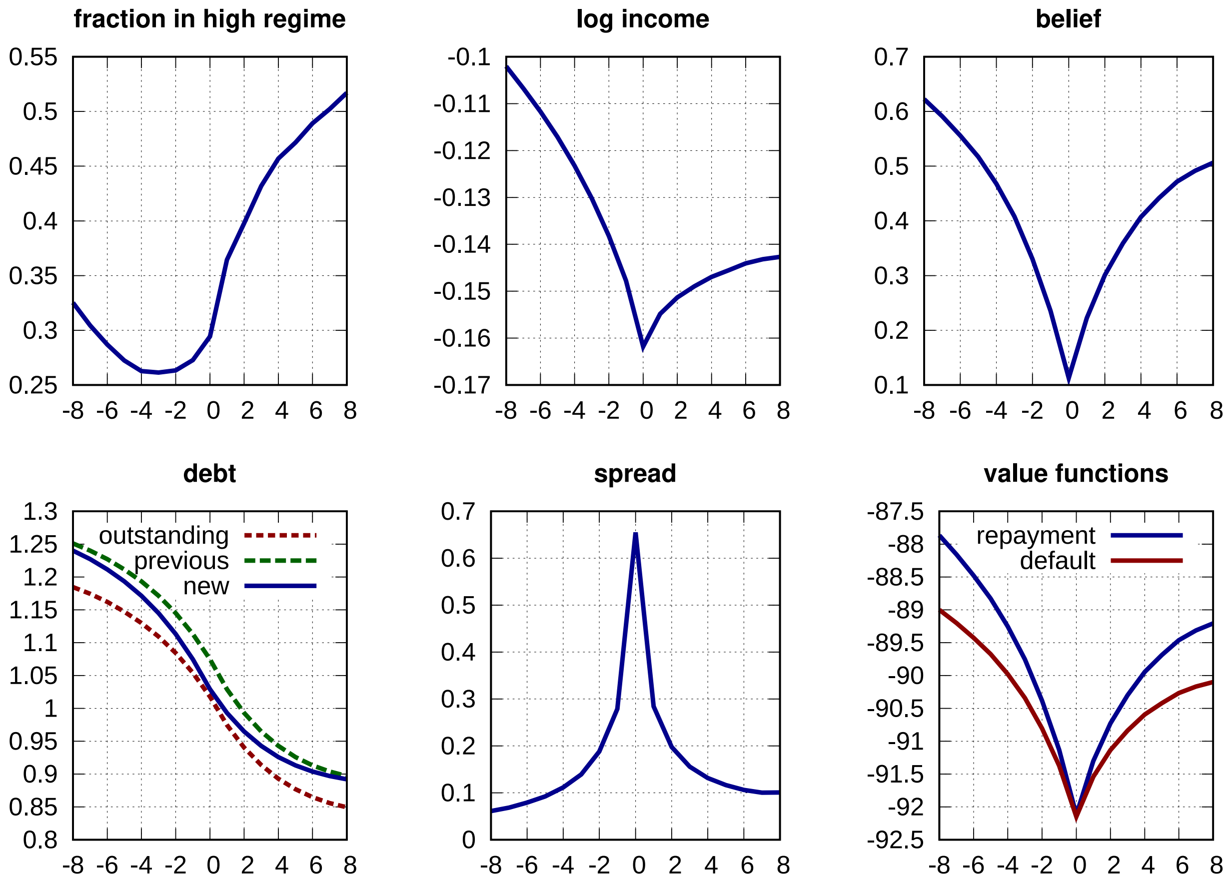

Figure 16: Analysis of episodes oh highest spreads

Note: the figure depicts averaged simulated paths of variables centered around peak spreads (at $t = 0$*) that are greater than 0.5. Outstanding, previous, and new debt refer to* $(1-\delta)b$*,* b*, and* b' *in the model, respectively.*

C.3 Assessing the elasticities of bond spreads

Analyzing Figure 16, we observe that the sharp spikes in bond spreads occur simultaneously with sharp reductions in both income and the belief. Separating the relative contributions of the two is not straightforward, as the latter derives from the former rather than being a shock of its own (which could be shut down). One way to disentangle the relative impact of income shocks and belief fluctuations is to compare the elasticities of the simulated spread with respect to these two states. To do so, I conduct the following exercise. At every simulated time period, I evaluate the bond price $q(b', y, p)$ at a perturbed state (income or belief) and calculate the resulting elasticity of the bond spread, for $i \in \{y, p\}$:[31]

$$e_{s,i} = \frac{\partial s(b', y, p)}{\partial i} \times \frac{i}{s(b', y, p)} \tag{14}$$

Table 10 reports the elasticities, averaged out across the simulated time periods, conditional on bond spreads being above several thresholds. In particular, column 2 reports elasticities for the unconditional sample, showing that a 1% reduction in income causes an 18% increase in the spread, while a 1% decline in the belief leads to an 7% increase. The difference between these elasticities is growing when we condition on higher spreads. To interpret these elasticities, we need to consider the typical relative variation in the two states. Of particular interest are the spread hikes, thus in the next two rows I report average negative growth rates exhibited by the two state variables. The size of income reductions is mostly determined by the standard deviation of the shock, around 1%. On the other hand, the belief on average decreases by anywhere between 6% (unconditionally) to almost 40% (for spreads higher than 6%). Taking the product of the elasticities and the mean growth rates (last two rows of Table 10), we find that the bond spread fluctuates much more in response to the typical changes in the belief than to the average changes in income for most of the simulated sample. Only in the case of the highest spreads, the impact of the two variables is quantitatively similar.

Table 10: Bond spread elasticities

	Spread greater than:			
	0.00	0.02	0.04	0.06
Elasticity w.r.t. y	-17.79	-29.10	-35.75	-40.24
Elasticity w.r.t. p	-7.14	-1.54	-0.89	-0.71
Mean negative growth in y	-0.01	-0.01	-0.01	-0.01
Mean negative growth in p	-0.06	-0.33	-0.37	-0.39
Mean rise in spread due to y	0.15	0.26	0.32	0.37
Mean rise in spread due to p	0.44	0.51	0.33	0.28

Note: The formula for elasticities is given in (14). Mean negative growth in state $i \in \{y, p\}$ is calculated as: $\mathbb{E}\Big(\frac{i_{t+1}-i_t}{i_t} | \frac{i_{t+1}-i_t}{i_t} < 0\Big)$. Mean rise in spread due to state i is a product of the two.

[31] A natural question regarding this calculation is whether we should hold the optimal choice b' constant or let it change optimally according to the perturbed state. The results I present take the former approach, but adjusting the policy function would not affect them significantly.

C.4 Unpacking the impact of the belief on bond spreads

The bottom-right panel of Figure 16 shows that towards the typical peak of the spread, default probability also spikes. With bonds having long durations, the impact of the belief on spreads evaluated in the previous section masks two separate forces: i) the effect on expectations about long-term movements of bond prices; ii) the effect on next-period default probability. Formally, consider the bond price equation (11) reformulated as follows:

$$q\left(b', y, p\right) = \frac{1}{1+r^*}\left(\sum_{z}\sum_{z'} Prob(z)\,\pi(z'|z)\int_{\tilde{y}(b',y,p,z')}^{\infty} f_{z'}(y',y)\Big[\delta + (1-\delta)\Big(\kappa + q\left(b'', y', p'\right)\Big)\Big]\,dy'\right)$$

In the formulation above, $\tilde{y}(b', y, p, z')$ is an income threshold that makes the government indifferent between repaying and defaulting. By perturbing p (say, reducing it), we impact the bond price q both by affecting the expected future price $q\left(b'', y', p'\right)$ (it goes down), and by changing the default threshold $\tilde{y}$ (it goes up).

To decompose the contribution of these two forces quantitatively, I modify the exercise from Section C.3 in the following way. I calculate the elasticity of the bond spread with respect to p while holding $\tilde{y}$ constant, i.e. using the same default threshold that was obtained for the unperturbed belief. In this way, we isolate the impact of the belief on the spread that is due to the change in long-term income expectations.

Table 11 shows the original and modified elasticities of the spread with respect to belief, averaged out across the simulations and conditional on different spread levels. The main finding is that the impact of a change in the belief on the default threshold $\tilde{y}$ is quantitatively

Table 11: Modified bond spread elasticities

	Spread greater than:			
	0.00	0.02	0.04	0.06
Elasticity w.r.t. p	-7.14	-1.54	-0.89	-0.71
Elasticity w.r.t. p (no change in $\tilde{y}$)	-7.13	-1.51	-0.83	-0.64
Mean negative growth in p	-0.06	-0.33	-0.37	-0.39
Mean rise in spread due to p	0.44	0.51	0.33	0.28
Mean rise in spread due to p (no change in $\tilde{y}$)	0.44	0.50	0.31	0.25

Note: The formula for elasticities is given in (14). Mean negative growth in p is calculated as: $\mathbb{E}\Big(\frac{p_{t+1}-p_t}{p_t}\,|\,\frac{p_{t+1}-p_t}{p_t} < 0\Big)$. Mean rise in spread due to p is a product of the two.

small for most of the simulated periods. This is not surprising, given that the default probability in this model is close to zero for long streaks associated with "normal times". In such cases, the default threshold may lie below the feasible range of income shocks to begin with, and hence its movements do not matter. It is only when spreads are high, the immediate default probability spikes (as can be seen in Figure 16) that the impact of p on $\tilde{y}$ matters more. But overall, it still tends to be dominated by the effect on future bond prices.

C.5 Association between belief and spread in the event study

Here, I turn my attention to the role of belief fluctuations in driving the bond spread during the European Debt Crisis. Figure 17 is a transformation of Figure 7 that overlays the predicted spread with the (negative) log of the belief. It is noteworthy that the two variables follow a very similar pattern and have a correlation of 0.95 (compared to the correlation of -0.69 of the spread with log income). Of course, this does not imply that income is unimportant for the determination of the spread; on the contrary, without the decline in income a debt crisis would not be possible in the first place. Instead, the point if that the precise timing of the movements in bond spread during the European Debt Crisis (the main object of interest in this paper) seems to be shaped by the relative changes in the belief.

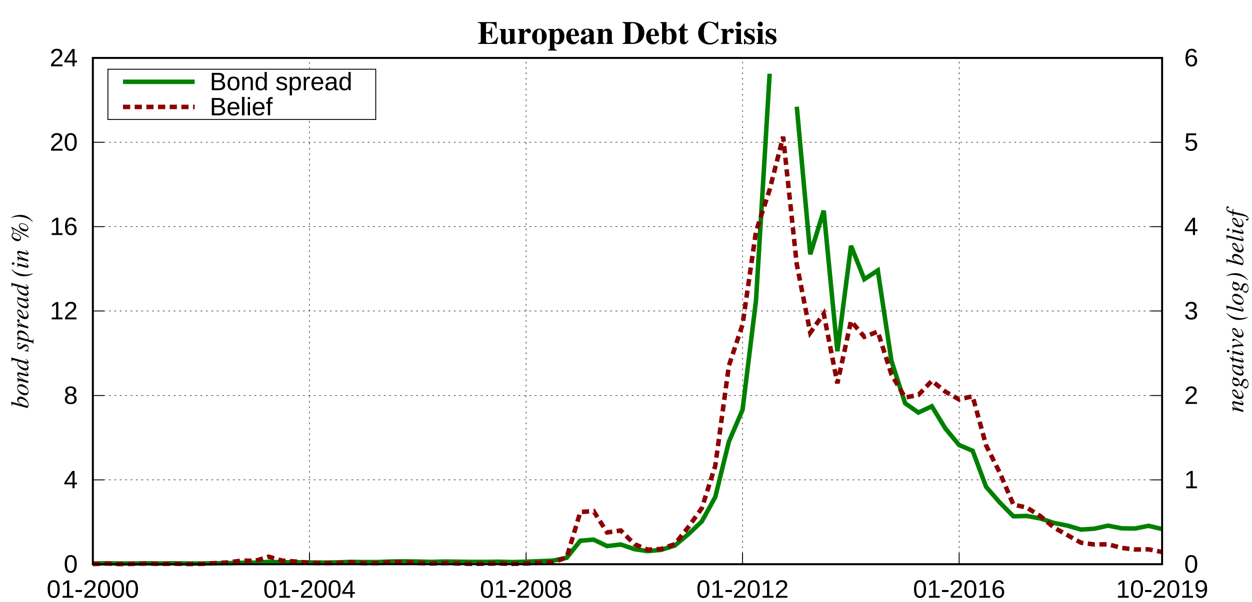

Figure 17: Belief and spread in the event analysis

C.6 Evolution of GDP forecast errors in model event studies

In this section, I analyze the evolution of GDP forecast errors from the event study of the European Debt Crisis (Section 4.5) using both variants of the model. The goal is to create a model counterpart of the data depicted in Figure 2.[32] Panel 18(a) reveals a similar pattern: forecasts roughly align with the data until 2008, and then start a gradual adjustment process (become flatter over time), resulting in sizable errors. By 2012, the forecasts are much more pessimistic and again align with the data. This contrasts with the pattern evident in panel 18(b), where no such downward revision in the forecasts ever takes place. This is because, with a simple AR1 process, output always mean-reverts in the same direction. In this way, panel 18(b) reproduces the pattern of forecasts around Argentina's default in 2001 (Figure 10) more closely than the one around the European Debt Crisis.

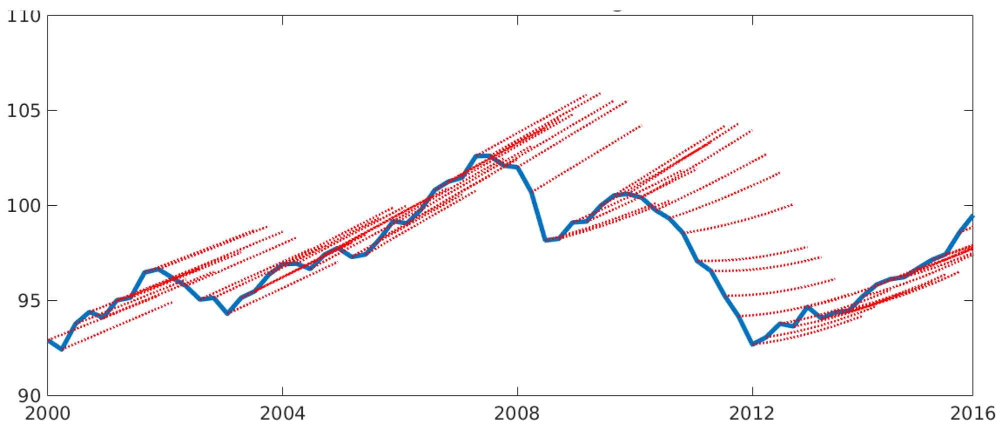

(a) Disasters and learning

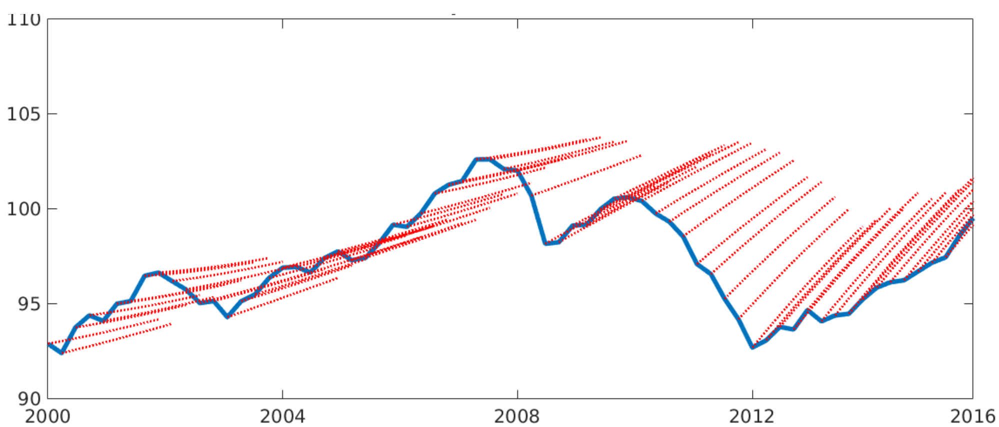

(b) Simple AR1

Figure 18: Evolution of GDP forecasts in event studies with the two versions of the model

[32]A notable difference between Figures 2 and 18 is that the former is more granular due to the fact that the data forecasts arrive twice a year and refer to annual GDP levels.

D Calculation of the debt write-off from bailouts

In the event studies presented in Section 4.5, I assume that upon defaulting the government re-enters the market with an exogenous debt write-off. This aims to mimic the bailout that Portugal received from the European Commission and the IMF at the end of 2011. In this Appendix, I show how such a write-off can be quantified.

I start with the facts. In May 2011, Portugal entered a joint emergency lending program by the IMF and the European Commission (EFSM+EFSF). The total extended credit amounted to 79 bn euro, out of which 76.8 was actually disbursed. At the time when the bailout was announced, the market value of Portuguese government's external debt securities was 68 bn euro.

The debt write-off stems from the fact that the bailout loans carried a low, essentially risk-free, interest rate at the time when the yields on Portugal's bonds were high. Hence, ignoring any differences in maturity structure of the two debt types, the total aid provided to the Portuguese government can be modeled as

$$\text{aid} = \ell \times [q^f - q]$$

where ℓ is the total face value of the emergency loans, q^f is the price of a risk-free bond, and q is the average price of a Portuguese bond. From the facts presented above, we know that ℓ can be expressed as

$$\frac{\ell}{q \times b} = \frac{76.8}{68} = 1.13 \Longrightarrow \ell = 1.13(q \times b)$$

where b represents the total outstanding debt securities of the government at the time of the bailout, and $q \times b$ is their market value. The bond prices can be expressed as

$$q^f = \frac{\delta + (1-\delta)\kappa}{r^f + \delta}$$
$$q = \frac{\delta + (1-\delta)\kappa}{r^f + E(s) + \delta}$$

where the average spread, $E(s)$, is 1.75% in the sample until 2019. Finally, we can calculate the debt write-off as

$$\text{write-off} = \frac{\text{aid}}{b} = 1.13\, q[q^f - q]$$

Using the parameters assumed in the model, the write-off amounts to 20.04%.

E Model with disasters and full information

In this section, I present the calibration and business cycle statistics for the model with rare disasters and full information about their realizations, which is used for the counterfactual exercise in Section 4.6. Note that the parameters of the income process are kept at the same level as in the main model. Table 12 summarizes the calibrated structural parameter values. It can be noticed immediately that the moments-matching exercise results in the values of β and $\hat{y}$ that are similar to the model with partial information.

Table 12: Calibration of structural parameters in the model with full information

Symbol	**Meaning**	**Full info**	**Source**
σ	Risk aversion	2	Literature
r^*	Risk-free rate	0.01	Literature
θ	Re-entry probability	0.049	Literature
δ	Probability of maturing	0.053	Data
κ	Coupon payment (in %)	1.250	Data
$\hat{y}$	Default cost par.	0.806	Calibration
β	Discount factor	0.987	Calibration
Calibration targets		**Full info**	**Data**
E (debt/GDP)		38.66	38.58
E (spread)		1.75	1.75

Note: targeted moments are given in percentage points. Simulations are repeated 10,000 times for a period of 1998-2019.

Table 13 presents the business cycle moments for this variant of the model. Note that the main differences relative to a benchmark AR1 model still hold in this case. It is also clear that learning about rare disasters has an important quantitative impact on real variables, contributing to a lower variance of consumption and trade balance relative to income. It is also worth noting that 99% of all defaults in this model occur in the disaster regime. The average maximum spread in the conditional distribution is 21%.

Table 13: Simulated behavior of the model with full information

Statistic	**Data**	**Full info** *Ergodic*	*Conditional*
$E(s)$	1.75	1.10	1.75
$std(s)$	3.06	5.52	3.61
$std(c)/std(y)$	0.98	0.95	1.19
$std(tb)/std(y)$	0.55	0.30	0.60
$corr(y,c)$	0.98	0.95	0.85
$corr(y,tb)$	-0.96	0.32	-0.03
$corr(y,s)$	-0.42	-0.29	-0.58
$corr(s,tb)$	0.39	0.19	0.45
$E(debt/y)$	38.58	71.89	38.66

Note: moments for the bond spread, debt-to-GDP ratio and the long-run default probability (annual) are given in percentage points. Ergodic (long-run) simulations extend to 10,000 quarters and are repeated 10,000 times, following closely Chatterjee and Eyigungor (2012). Conditional (short-run) simulations mimic the period of 1998-2019 (88 quarters) and are repeated 10,000 times starting from the actual levels of debt and GDP observed in 1998:Q1. Each short-run sample is constructed such that: i) the series start from the actual 1998:Q1 debt and income levels, and ii) the regime switches from good to bad in 2008:Q3. Consumption data is detrended using the common GDP trend.