Online Appendix for: Learning about Debt Crises

Radoslaw Paluszynski[*]

June 6, 2021

Contents

[*]Department of Economics, University of Houston, McElhinney Hall, 3623 Cullen Boulevard, Houston, TX 77204 (email: rpaluszynski@uh.edu).

A Supporting evidence

A.1 Dynamics of sovereign ratings over time

Figure 1 presents plots of the countries' real GDP, together with the "sovereign bond rating index".[1] 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.

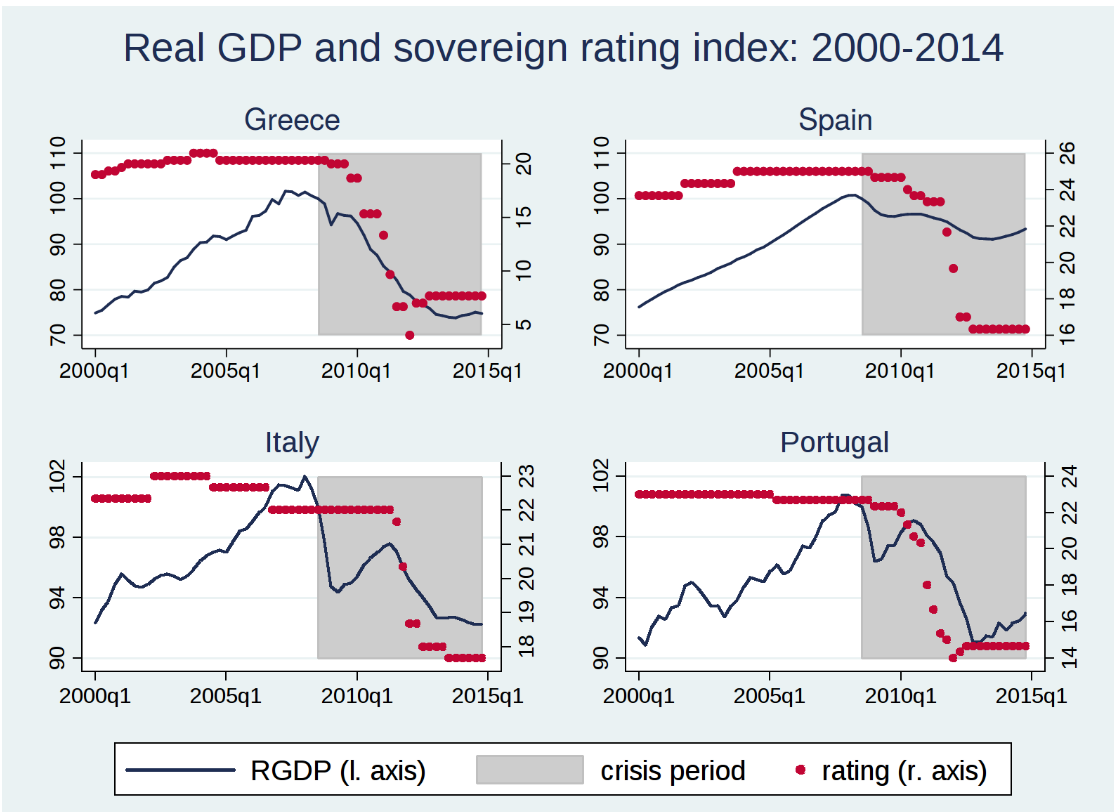

Note: The GDP series are in constant 2010 prices, and their values are normalized such that the third quarter of 2008 equals 100 (beginning of the financial crisis - shaded area). The bond rating index is constructed by converting the sovereign ratings of the three leading agencies - S&P, Moody's and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

Figure 1: Real GDP and sovereign bond rating index of the European economies: 2000-2014

[1]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 2 plots the evolution of GDP forecasts for Argentina, a representative emerging market economy, around the time of its crash in 2001.[2] Two differences stand out in comparison with Figure 2 in the paper 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

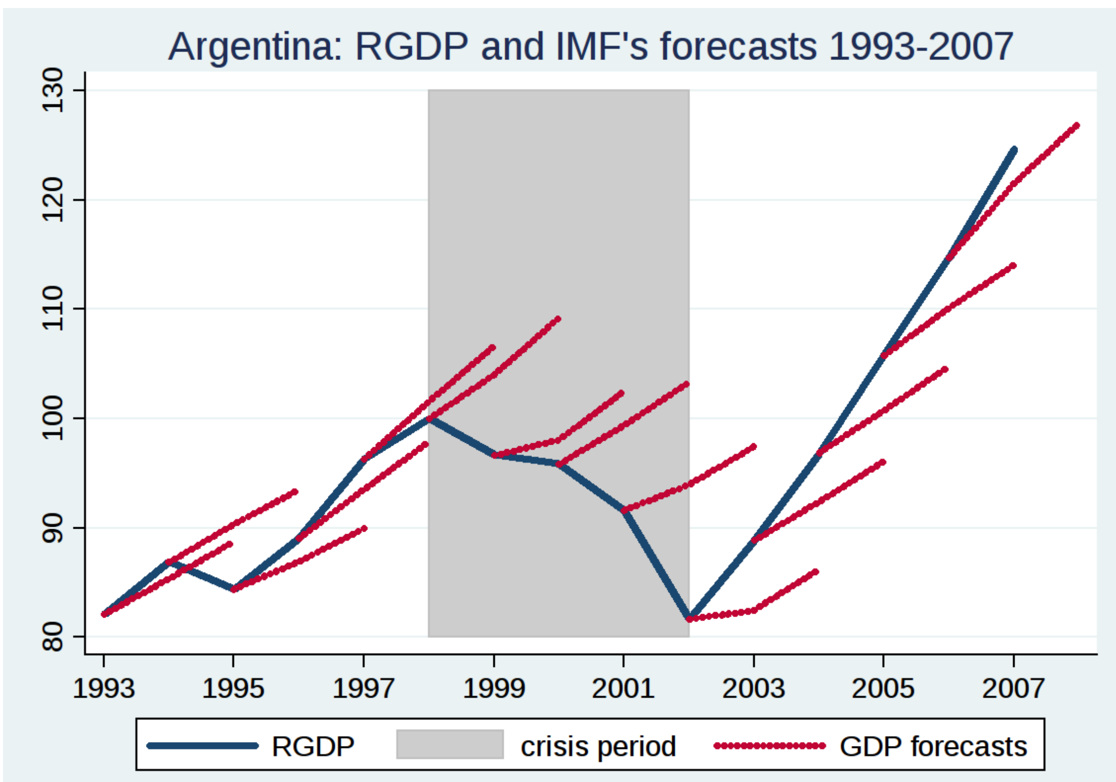

Note: The GDP series is annual and in constant prices; values are normalized so that it equals 100 in 1998. The red dotted lines represent one- and two-year ahead forecasts published in the fall of each year by IMF. The shaded area marks Argentina's debt crisis of 1998-2002.

Figure 2: Forecast and actual real GDP for Argentina

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

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 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 1 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 in the paper, 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 1: 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 3 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

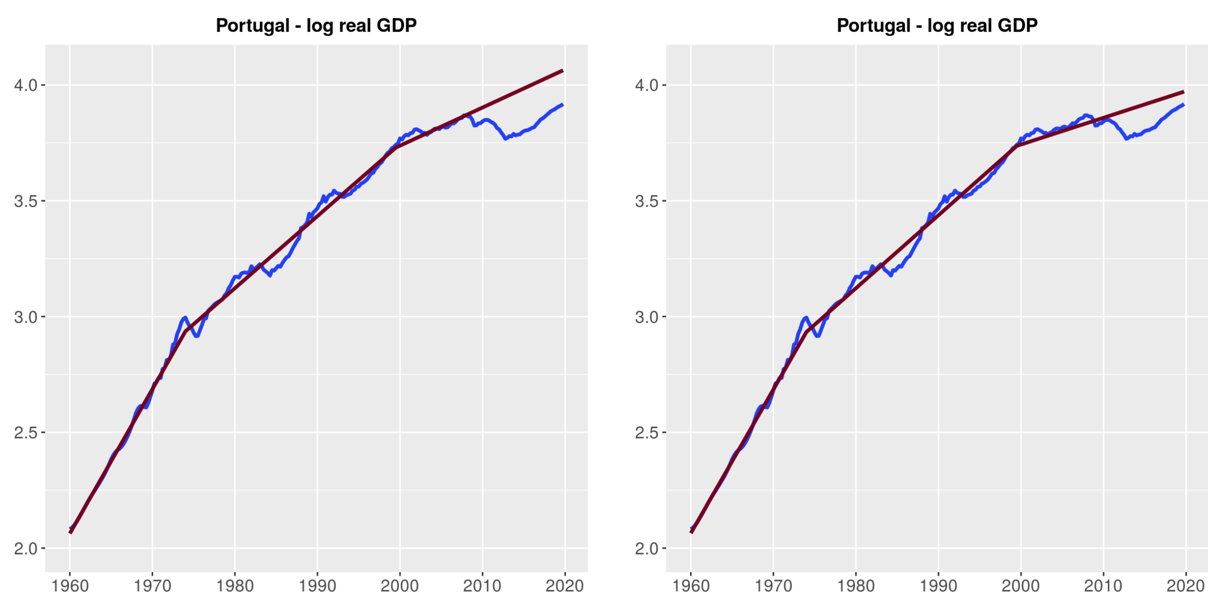

(a) Baseline case

(b) Simple AR1 case

Figure 3: 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 (formula 1 in the paper). 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 I.B in the paper).

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 $\mathcal{Y}_T = \{y_1, y_2, ..., y_T\}$, 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.[3] 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:

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

In formula (1), 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 (formula 2 in the paper). 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

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

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

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

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 (2). 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 III.B in the paper 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 2 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 4 overlays the paths of forecasts from the model and the data for all four countries,[4] in a detrended form, along with the paths of realized data 5 years later (the series ends

[4]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 in the paper, 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 2: 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 III.B in the paper. In each case, the regime-switching probabilities are the same as in Table 3 in the paper. The GDP data for each country is detrended in the same fashion as it is described for Portugal in Appendix B.1.

Figure 5 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 4). 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.

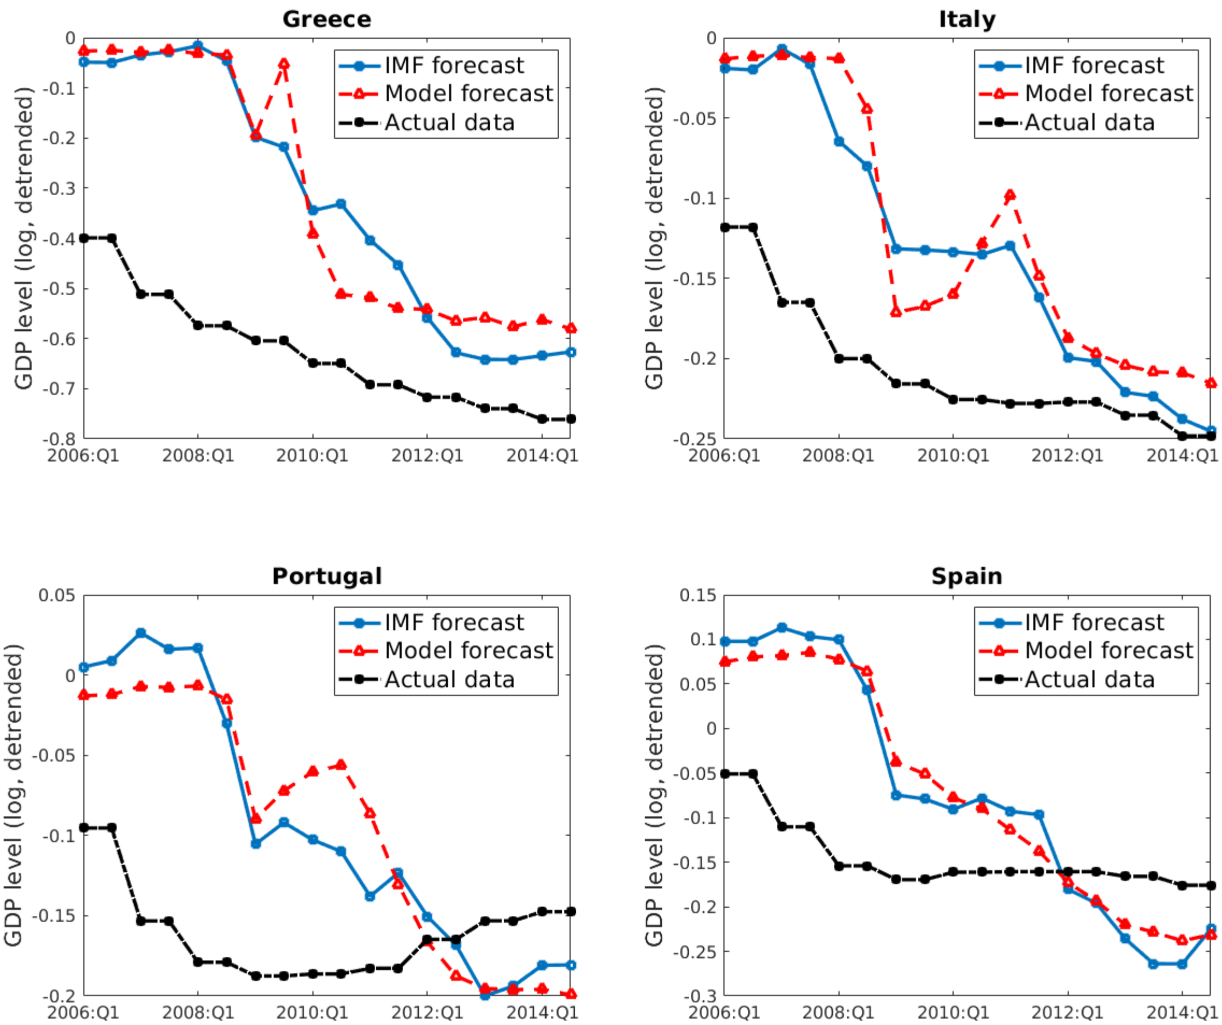

Note: Each point on the graph represents an annual detrended log-GDP level for five years ahead (for example, 2008:Q1 corresponds to the GDP level in 2013). The solid blue line represents the actual published forecasts (Q1 and Q3 refer to the spring and fall issues, respectively), while the dashed red 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 4: IMF- and model-generated projections for all countries

It is also worth reiterating, as Section III.A in the paper 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

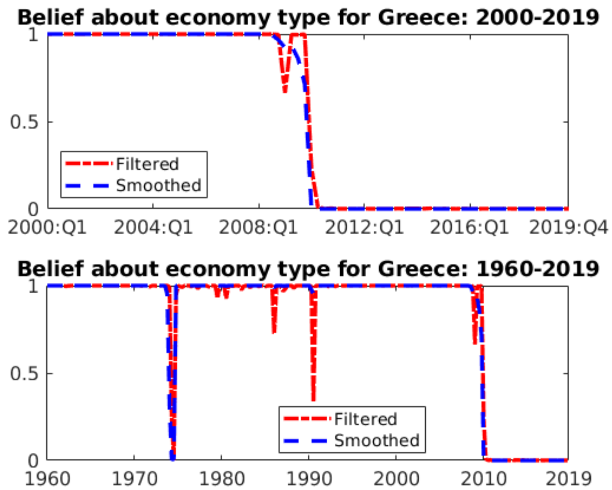

(a) Greece

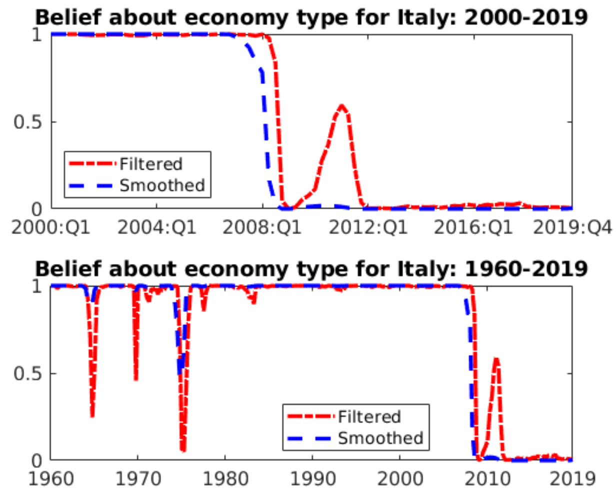

(b) Italy

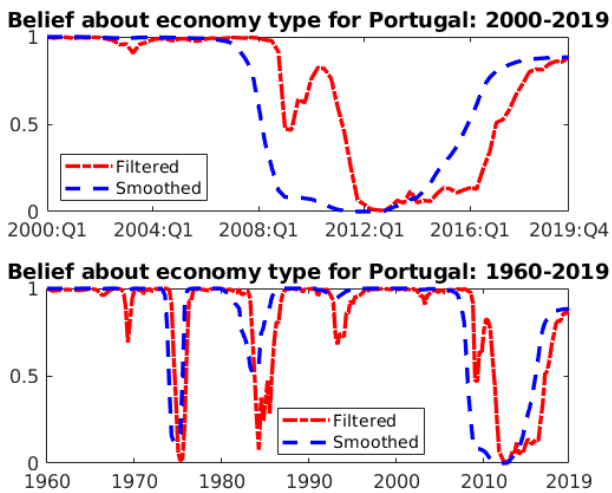

(c) Portugal

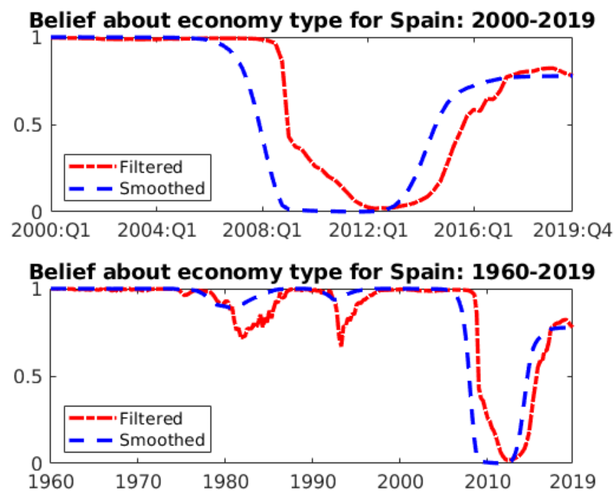

(d) Spain

Figure 5: Inferred paths of beliefs for each country

more promising 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 6 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 III.D in the paper. 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.

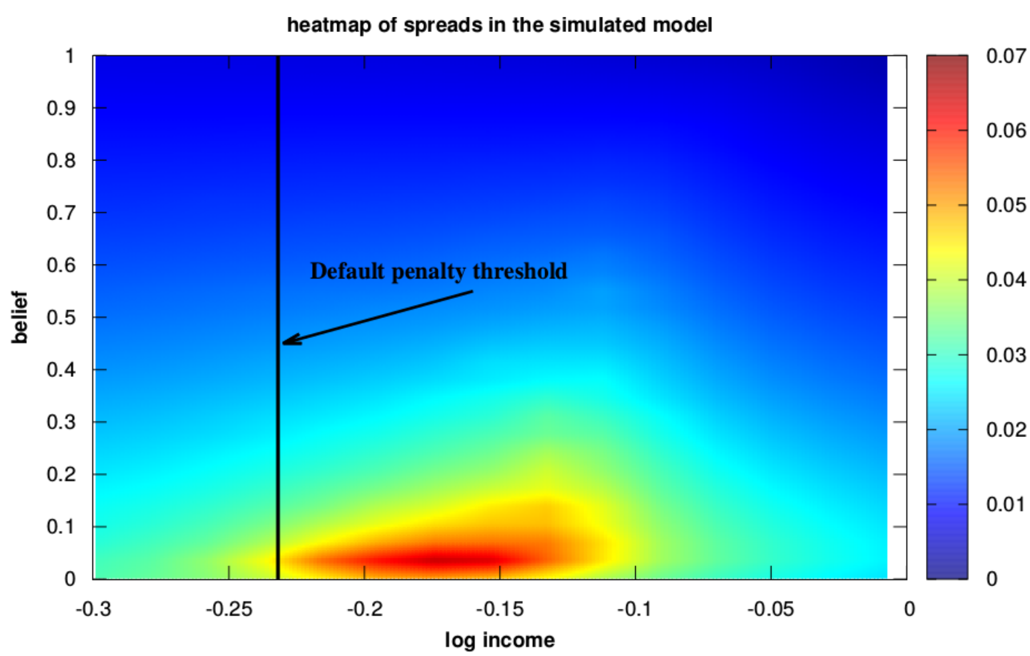

Figure 6: Heatmap of simulated spreads in the model

Figure 7 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.

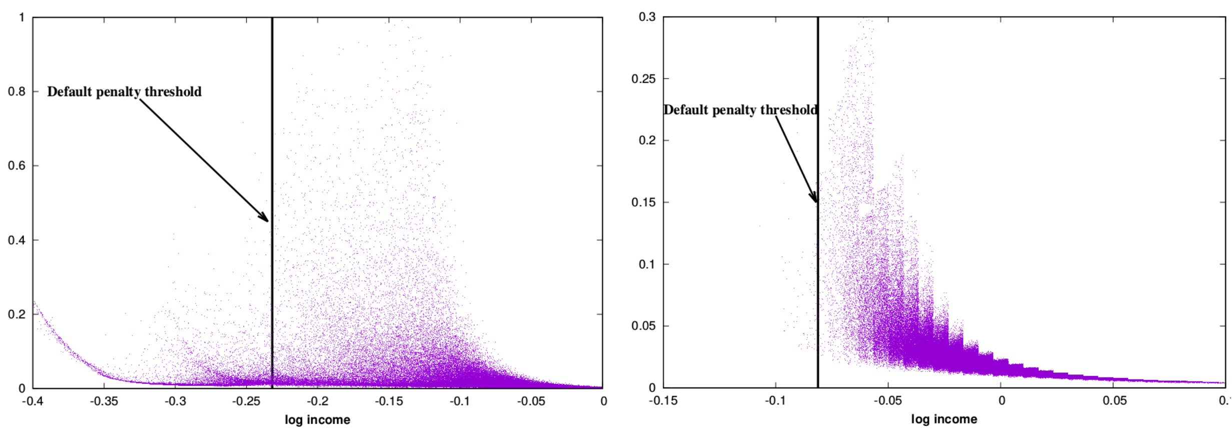

(a) Disasters and learning

(b) Simple AR1

Figure 7: 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 8 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.

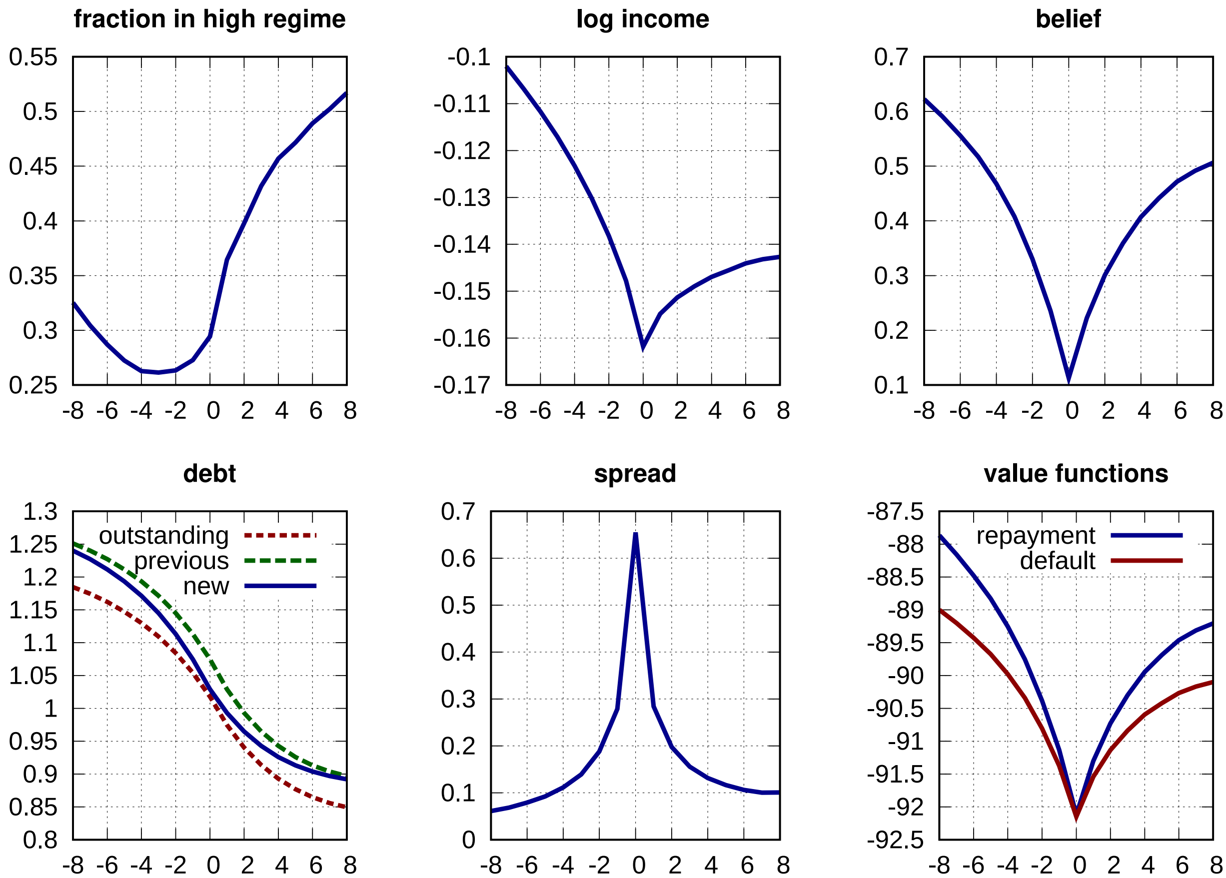

Figure 8: Analysis of episodes oh highest spreads

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 8, 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 $i \in \{y, p\}$:[5]

$$e_{s,i} = \frac{\partial s(b', y, p)}{\partial i} \times \frac{i}{s(b', y, p)} \tag{3}$$

Table 3 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 3), 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 3: 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 (3). Mean negative growth in state $i \in \{y, p\}$ is calculated as: $\mathbb{E}\left(\frac{i_{t+1}-i_t}{i_t} \middle| \frac{i_{t+1}-i_t}{i_t} < 0\right)$. Mean rise in spread due to state i is a product of the two.

[5] 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 8 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 (formula 11 in the paper) 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 4 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 4: 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 (3). 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 8) 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 9 is a transformation of Figure 7 in the paper 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.

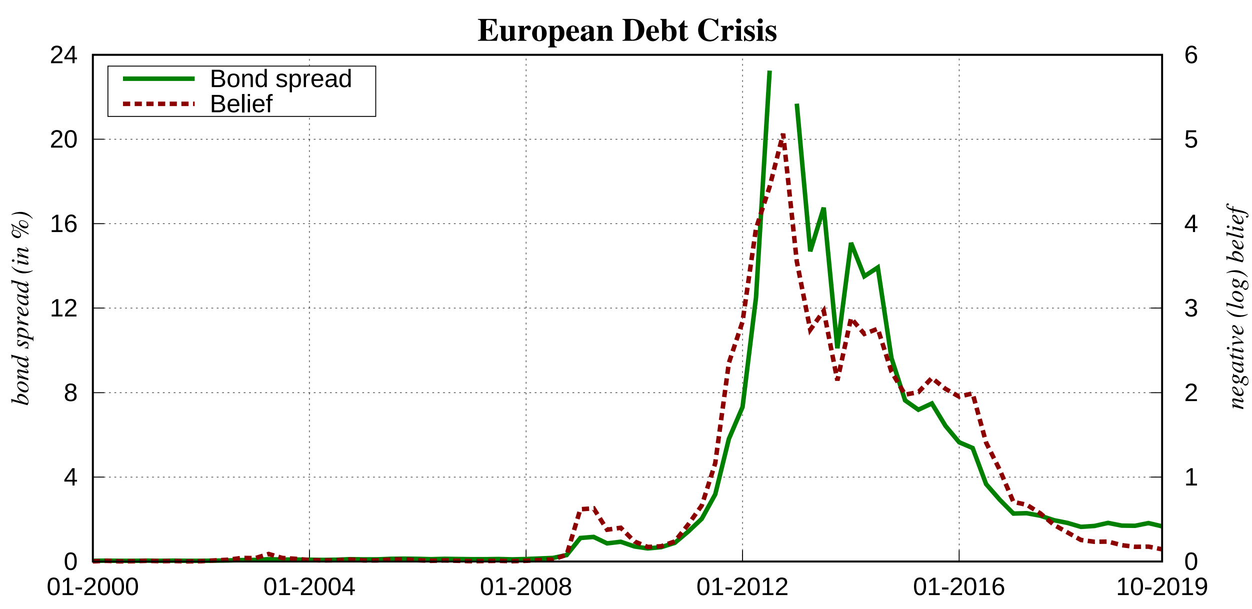

Figure 9: Belief and spread in the event analysis

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 III.E in the paper) using both variants of the model. The goal is to create a model counterpart of the data depicted in Figure 2.[6] Panel 10(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 10(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 10(b) reproduces the pattern of forecasts around Argentina's default in 2001 (Figure 2) more closely than the one around the European Debt Crisis.

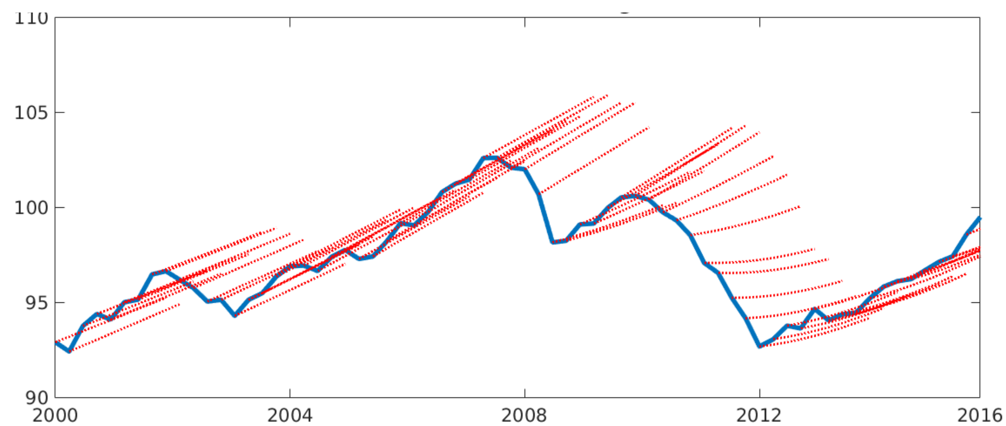

(a) Disasters and learning

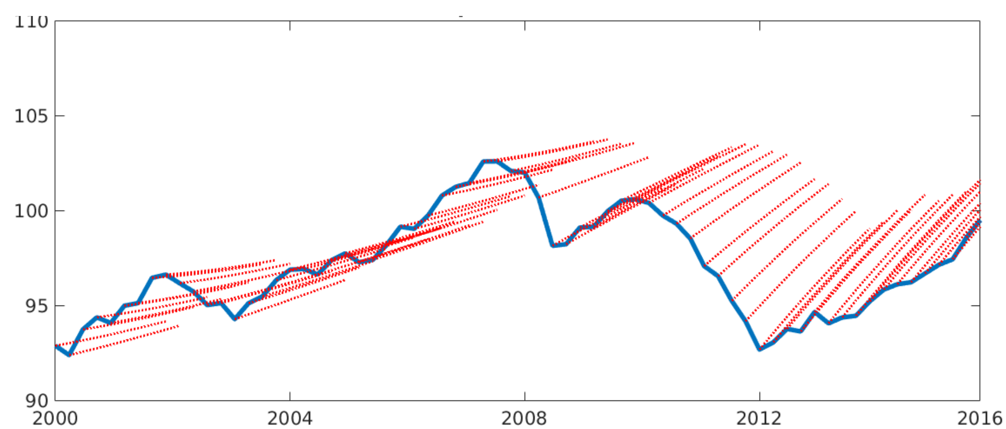

(b) Simple AR1

Figure 10: Evolution of GDP forecasts in event studies with the two versions of the model

[6]A notable difference between Figure 2 in the paper and Figure 10 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 III.E in the paper, 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%.

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 III.F in the paper. Note that the parameters of the income process are kept at the same level as in the main model. Table 5 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 5: 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 6 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%.

References

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

Table 6: 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.

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

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

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.

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

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

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