THE UNIVERSITY OF CHICAGO

# REPUTATION AND SOVEREIGN DEFAULT

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF ECONOMICS

BY
PHILIP BARRETT

CHICAGO, ILLINOIS
JUNE 2016

# TABLE OF CONTENTS

# SUMMARY OF THE THESIS

This thesis has two central contributions. One is economic; it explains the behavior of sovereign interest rates after a default. The other contribution is principally technical; it defines and provides an accurate closed-form approximation to a common filtering problem. The main steps in the argument are as follows:

- Standard economic theory predicts that after conditioning on macroeconomic conditions, the cost of sovereign borrowing is independent of a country's past default history.
- Empirical evidence contradicts this prediction. Countries which have recently defaulted pay more to borrow than would be expected given their observable macroeconomic conditions. This extra cost of borrowing is the *conditional default premium*.
- I propose a model of sovereign default where the government's immediate cost of default is governed by a hidden type.
- The ability of the model to generate high post-default spreads depends crucially on the continuous nature of the government' type.
- The evolution of investors beliefs of the government's type given the past default history is a *threshold filtering problem*. This is a common problem in many models with persistent hidden types and binary actions.
- I propose an approximate solution to the threshold filtering problem, which I term the *threshold filter*. This provides a closed-form solution to the approximate evolution of the first two moments of the distribution of the hidden state. This is a substantially more accurate solution to the problem than standard filtering methods, such as the unscented Kalman filter or exact Gaussian filter.
- I use the threshold filter to model investor's beliefs and so solve for equilibrium in the sovereign default model.

- A calibrated version of the model generates quantitatively accurate high post-default spreads.
- I show that the threshold filtering problem shows up in a wide variety of other economic problems, including market entry and long-term unemployment. The threshold filter can be applied to these scenarios as well.

# ACKNOWLEDGEMENTS

The support and advice of my thesis committee - Lars Hansen, Fernando Alvarez (co-chairs), and Mark Wright - has been invaluable. As has the input of my colleagues and friends both at Chicago and elsewhere, particularly: Jon Adams, Chanont Banternghansa, Martin Beraja, Seth Blumberg, Bon Guen Choi, Thorsten Drautzburg, Richard Harrison, Alex MacKay, Balasz Szentes, Willem Van Vliet, and participants of the Economic Dynamics, Applied Macro Theory & Capital Theory seminars. Finally, this thesis would not have been possible without the love and support of my wife, Scarlett.

# LIST OF FIGURES

# LIST OF TABLES

# CHAPTER 1

# INTRODUCTION

## 1.1 Motivation

This thesis is motivated by an empirical fact: countries pay more to borrow from foreign creditors in periods shortly after a recent default. Crucially, this holds even after accounting for the usual co-movements of debt prices with a wide range of domestic and international macroeconomic variables and fixed effects.

This fact, which I term the which I term the conditional default premium[1], represents a serious challenge to our understanding of sovereign default. The answer to many fundamental questions about sovereign default (including: why do sovereigns default so rarely? and why can they borrow so much?) depend critically on the consequence of default. The more punitive these consequences, the greater the deterrent to default.

In the canonical model[2], default is deterred by a consequent temporary loss of output and exclusion from international capital markets. The existence of the conditional default premium suggests that there is a third factor which may also deter default: after a default, borrowing is simply more expensive. This phenomenon is not something that standard models can generate, as their state space[3] is a subset of the conditioning variables.

Empirical estimates of the conditional default premium suggest that this effect is large. Spreads are usually around two to four percentage points higher in the first year after borrowing resumes, and the effect only dissipates after three to five years. Given that countries typically exit default with high debt burdens[4], this can represent a sizable payment to overseas creditors. For example, Venezuela had public debts of around 45% of GDP following

1. "Conditional" because it exists even after conditioning on other predictors of the price of borrowing.

2. Meaning: Arellano (2008) and successors.

3. Usually two variables: one summarizing past obligations and one future income.

4. Indeed, Benjamin and Wright (2013) show that governments very often exit default with debt burdens that are higher than when they entered default.

the end of their 1998 default. So an extra three percentage points of interest payments would have cost nearly 1.5% of GDP; roughly equivalent to Venezuela's entire public health expenditure in 1999.

## 1.2 Aims of the Thesis

In this paper I seek to answer two questions, motivated by the preceding discussion.

**Question 1:** What theory can explain the conditional default premium?

**Question 2:** What does the answer to question 1 tell us about sovereigns' motivations for default?

The answers to these questions depend critically on whether high post-default spreads are a direct consequence of default, or if they are merely correlated

If the conditional default premium is a direct consequence of default then higher spreads enter the calculus of a forward-looking sovereign contemplating default, acting as a deterrent to default. Alternatively, spreads and incentives for default might both be driven by an independent third factor. This could also generate a conditional default premium in equilibrium, as default and subsequent spreads would be correlated. But because they are driven by the third factor, future spreads would be high whether or not the sovereign defaulted.

The distinction between these two stories is important for understanding sovereigns' default decisions. In the first one, the conditional default premium influences debtors' behavior by acting as a deterrent to default. In the second story, default precedes high spreads, but neither causes them nor affects the sovereign's default decision. Which of these is true has a fundamental impact on our understanding of why and when sovereign default happens.

## 1.3 Methodology and contributions

To answer these question I present a simple model of sovereign default where the government has private incentives to default. I consider two candidate information structures: one where post-default spreads are a direct function of recent history; and one where default and subsequent spreads are simply correlated.

When information about the government's private incentives is asymmetric (known only be the government, and not investors), the default reveals information about the government's hidden type. And so if this model were to generate a conditional default premium, it would act as a deterrent to default.

On the other hand, when information is symmetric (known to both investors and the government), then future interest rates are independent of the current repayment or default decision. And so a conditional default premium in this model would merely be a correlated phenomenon, and not a deterrent to default.

Although both models could, in principle, generate a conditional default premium, only the one with asymmetric information does. Figure 1.1 plots empirical estimates of the conditional default premium (in blue) and overlays this with the results of the same calculation in the symmetric (broken line) and asymmetric (solid line) information models[5]. Clearly, the symmetric information model predicts a negative conditional default premium for all but the first period after default. This is at odds with the data. But when information is asymmetric, the model can match the data, both qualitatively and quantitatively[6].

As a result, I reject the model with symmetric information and conclude that:

**Answer 1:** Conditional default premia are consistent with a dynamic game of asymmetric information where default reveals that the government's future likelihood of default is high.

5. See Figure 3.10 for a version of this chart with standard errors on the empirical estimates

6. In contrast, standard models would produce conditional default premia of zero at all horizons.

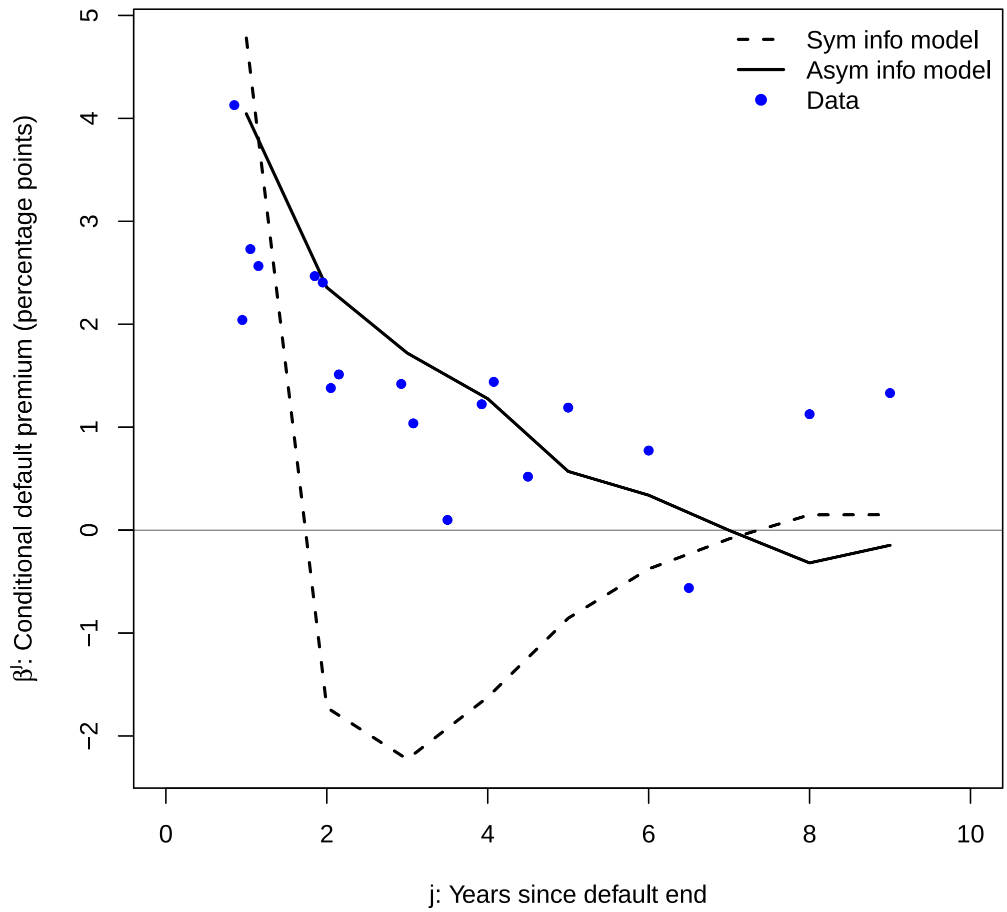


Figure 1.1: Data and model measures of the conditional default premium

**Answer 2:** The sovereign's current default decision determines the future price of borrowing, as default today reveals information about the likelihood of default in future. The conditional default premium is therefore a deterrent to default.

A key aspect of the asymmetric information model is that the government's type is continuous. Models with a discrete number of types present the modeler with a very tough trade-off between tractability and ability to fit the data. Because the dimension of the state space grows with the number of types, models with more than three types are effectively impossible to solve numerically[7]. Models with two or three states, however, cannot generate a persistent conditional default premium. This is because a persistent conditional default premium requires investors to be uncertain about the borrower quality for an extended period of time. But with few types, actions very quickly reveal type almost completely.

Using a continuum of types eliminates this problem. With added uncertainty, revelation of types (and so movements in default premia) is slow. But this also introduces a difficulty of its own: how to parsimoniously encapsulate the distribution of investor beliefs over a continuum of types. Developing a method to address this issue is the main technical contribution of this thesis. I show that the filtering problem that investors face is one that standard nonlinear filtering techniques, such as the unscented Kalman filter, fail to deal with well. I propose an alternative approach, which I call the threshold filter, and show that it approximates well the full distribution of investors' beliefs. Although the threshold filter is motivated by the sovereign default problem, it is also applicable to a wide range of other economic situations where binary signals reveal information about a hidden type, such as market entry or employment.

## 1.4 Discussion: Economic mechanisms

The difference between the symmetric and asymmetric information models is due to the offsetting effects of two mechanisms. First, a composition effect pulls spreads down after return to capital markets, whether information is symmetric or not. Observations $n$ periods after the most recent default are, by definition, conditioned on observing repayment for the preceding $n-1$ periods. So the composition of types of governments that repay after default is relatively good, even though defaulters themselves are usually bad. The worst types "default out" of the unit of analysis very soon after default. This mechanism produces

7. For example, with four types the description of beliefs is three-dimensional (as the vector of probabilities is constrained to be unit-sum). Plus the two observed states (income and debt) and the government's realized type makes a six-dimensional problem. Furthermore, because off-equilibrium outcomes influence the equilibrium, the model must be solved accurately on the entire state space, not just the equilibrium.

low post-default spreads in the model with symmetric information.

The composition effect acts in the asymmetric information model too, although it is offset by a credibility effect. Here, credibility refers to investors' beliefs about the government's type[8]. Following a default, investors already believe that a government is of a poor type. The government therefore has little incentive to repay in order to preserve its credibility. Accordingly, the model with asymmetric information produces high default rates and spreads after a default, above and beyond the variation due to changes in macroeconomic conditions. That is, it generates a conditional default premium.

8. Credibility in this sense is distinct from reputation. While credibility pertains to beliefs about type, reputation relates to beliefs about actions. As I consider only Markovian equilibria, reputational concerns are beyond the scope of this paper. For more on the distinction between reputation and credibility see Drazen (2004)

## 1.5 Related Literature

This thesis is related to four distinct but related areas of economic research, and a technical literature on nonlinear filtering.

### *1.5.1 The incentives for repayment and default*

At its most basic, this paper is an attempt to understand the factors that induce or deter default, and so it touches on a broad literature on this fundamental subject. In their classic paper, Bulow and Rogoff (1989) argue that the ability of sovereigns to save means that exclusion from credit markets cannot be a sufficient punishment to deter default on any outstanding obligations. However, Wright (2002) shows that collusion by a finite number of lenders can undermine this argument. Amador (2003) also describes an environment where Bulow and Rogoff's argument fails. In this paper, frequent political change eliminates the necessary incentives for saving.

More recent work in the vein of Arellano (2008) (including Chatterjee and Eyigungor (2012) and Arellano and Ramanarayanan (2012)) adopts a hybrid approach, where default is costly both because it exposes the sovereign to income fluctuations, and because it entails a fall in output. My model is in the spirit of these in that default involves exclusion from capital markets, income risk due to autarky, and a temporary "level" cost (such as an output loss). But I augment this strand of the literature with an informational cost of default due to persistently higher post-default spreads.

### *1.5.2 Empirical work*

The empirical assessment of the model owes much to the work of Catão and Mano (2015), Borensztein and Panizza (2009) and Cruces and Trebesch (2013). These papers provide evidence that spreads after a default are persistently elevated, work that I develop to generate empirical targets for my model. The magnitude of the effect varies, but in the first year after returning from default are an increase of between two and four percentage points. Point estimates are almost always positive for around five years, although significance is harder to establish at horizons longer than three years.

### *1.5.3 Private government type*

From a modeling perspective, the most similar group of papers to mine feature information about government type that is not public and so default and repayment are informative about type. The seminal work in this field is Cole, Dow, and English (1995), where two government types alternate stochastically. Here, the role of signaling is quite limited, as the government's type is fully revealed by their action: good types repay and bad ones default.

Subsequent attempts to allow for more sophisticated signaling have faced two related technical challenges: how to maintain a parsimonious representation of investor beliefs over type, and how to model the type information conveyed by government's decision about how much debt to issue. Alfaro and Kanczuk (2005) and D'Erasmo (2008) are two recent models of sovereign default with asymmetric information and so have to confront these two issues.

Both models feature only two types of government, and so actions either completely reveal type or are totally uninformative. As a result, neither model can generate a conditional default premium; borrowing is either completely risk free (both types repay) or repayment fully reveals the better type (when only one type repays). In neither case are post-default spreads persistently elevated. In addition, D'Erasmo allows debt issuance to be a government choice, chosen to smooth consumption. As a result the good type can almost always differentiate themselves by issuing less debt, and so the signaling role of default (or repayment) is minimal.

I build on these approaches by allowing for a continuum of types. As opposed to a finite-type model, actions are never completely informative (or wholly uninformative). And because surpluses are exogenous, debt issuance reveals no information about type beyond that contained in the default-repayment decision. This allows for slow-moving beliefs about type that generate behavior of spreads around default which match empirical patterns.

Similar in spirit to my work are Sandleris (2008) and Onder (2013). These papers both develop environments where the government has private information about the state of the economy, the revelation of which influences the default decision. However, the neither model is designed to investigate the source or implications of default premia. Sandleris's model has only three periods, and so cannot speak to persistently elevated post-default spreads. And Onder focuses on cross-country patterns of spreads, not their dynamics evolution after default. Phelan (2006) also touches upon themes of information asymmetry and and time consistency in government policy. Here, the government may confiscate output ex post, and so the government's reputation for confiscation (or not) influences the private sector's production. The main difference between this and a model of borrowing is that there is no price which reacts to private sector beliefs, only a quantity response.

One final alternative in this category is the narrative approach, best exemplified by Tomz (2007). Tomz suggests a similar mechanism to that which I propose, presenting supporting data from a number of historical episodes to show that spreads are high following first-time entry or post-default return to international capital markets. He combines this with media reports and statements by public officials to argue that spreads fall after (re-)entry as sovereigns establish a reputation for repayment.

### *1.5.4 Public government type*

The last class of papers related to the economics of this thesis are those where the government's type is public. Cuadra and Sapriza (2008) show that a model of this type can explain why spreads are higher when there is more political instability. Similarly, Hatchondo, Martinez, and Sapriza (2009) show that political shocks in a complete-information model can reduce the correlation between output and default produced by many standard models of sovereign default (which Tomz and Wright (2007) document is weak).

### *1.5.5 Nonlinear filtering*

In solving the government's signaling problem, this thesis also makes a novel contribution to the literature on nonlinear filtering. The specific problem that I study is one where the law of motion of the system is linear by the signal is discrete in the state, and so highly nonlinear.

Filtering problems arise in trying to recover the distribution of a hidden state from a sequence of partially informative signals. This is exactly the problem that investors face when trying to uncover the government's state from their history of past repayments and defaults. In their seminal papers, Kalman (1960) and Kalman and Bucy (1961) provide *inter alia* the exact solution to this problem when the law of motion of the state and the noise process are both linear, and the noise is Gaussian.

A large body of subsequent work has developed Kalman's ideas to nonlinear problems. These fall into two broad classes. In the first class are those that linearize the system at each step and then apply the basic Kalman filter machinery to the linearized problem. This includes the extended Kalman filter, developed in Jazwinski (1966) and Jazwinski (1970), as well as the statistically linearized filter (see Gelb (1974)). The linearization process in these filters make them unsuitable for the filtering problem in the default model I study. Applying them blindly leads to predicted variances of either zero or infinity.

The second class of methods are those that attempt to address the nonlinearity in the signal by computing the variance of the signal (and hence the nonlinear Kalman gain) without resort to linear approximation. Examples include the unscented Kalman filter (Julier and Uhlmann (1997), Wan and Van Der Merwe (2000)) and the various flavors of the Gaussian filter (Ito and Xiong (2000), Särkkä (2013)).

These latter methods are applicable to the threshold problem that I study, but generate poorer estimates of the distribution of the hidden state because they a) approximate the distribution of the discrete signal with a normal, and b) use only approximate integration to compute the signal variance.

Last, a note on a literature not related to this paper. Discrete-element methods can, in general, be used to generate superior approximations to distributions of hidden state variables. Examples include finite-state Markov chains and particle filters. However, because the focus in this thesis is on being able to use the output of a filter as a state variable in a strategic economic model, there is a high premium on parsimony. Increasing the dimension of the description of investor beliefs very quickly makes solving the model computationally intractable. Discrete-element methods typically do not have simple representations. For example, the value of an $n$-state Markov chain is given by an element of an $n-1$ dimension simplex. In contrast, the threshold filter tracks only the first two moments of an approximating distribution. As a result, discrete-element methods are ruled out of consideration as a solution to the filtering problem purely on the grounds of excess complexity.

## 1.6 Recapitulation

In summary, the main economic contribution of this thesis is to show that variations in a government's private distaste of default can explain why default predicts future spreads above and beyond variation due to macroeconomic observables. In doing so, I distinguish between two competing mechanisms: that default is genuinely informative about future borrowing costs, or that default and spreads are simply driven by a third factor. I conclude that the former mechanism best explains the data. Default itself is a determinant of future spreads. So high spreads are a direct and long-lasting consequence of default and impact the government's repayment-default decision.

This paper also makes a further, methodological contribution: it implements a new approximate solution to the of a threshold filtering problem. This method allows for a simple compression of the state space of beliefs about government type to just two variables. This approximation is potentially widely applicable to other models with asymmetric information and a continuum of persistent types.

# CHAPTER 2

# A MODEL OF SOVEREIGN DEFAULT AND REPUTATION

## 2.1 Empirical evidence

In models of sovereign default with fixed government preferences, the state variables alone (typically outstanding government debt, and some variable correlated with permanent income, such as total output) predict default risk and therefore spreads. The incidence of default is therefore not informative about default risk beyond its correlation with these states.

This section puts these predictions to the test. I run panel regressions which show that after accounting for the impact of outstanding debt, government surpluses and other observable economic factors, spreads are around 200-250 basis points higher for at least two years after the end of a default. This difference is economically significant. These results are very similar to those found in Catão and Mano (2015), Cruces and Trebesch (2013) and Borensztein and Panizza (2009). The purpose of these results is to establish that the particular specification that I use when running the model regressions also generates a conditional default premium consistent with the literature

### *2.1.1 Data*

My sample covers 27 emerging market economies between 1980 and 2013. During this time, 19 of the countries default on a total of 27 separate occasions. My variables of interest are: spreads from JP Morgan's Emerging Market Bond Index (EMBI), and data on dates and sizes of defaults from Benjamin and Wright (2013). I supplement this with cross-country regression controls from the World Bank WDI database, Moody's and the EUI International Country Risk Guide. I describe the data further in Appendix B.

### 2.1.2 *Panel regressions*

I start by using the data I have collected to test empirically whether the interest rate that countries face rises after a default. Each period is one month. I construct a vector $v_{it}$ indicating the time elapsed since the conclusion of the last default. Successive elements of $v_{it}$ are 1 if country $i$ finished its last default, respectively: after $t-12$ (one year prior); between $t-13$ and $t-24$ (two years prior); between $t-25$ and $t-48$ (three or four years prior); and between $t-49$ and $t-84$ (five to seven years priors).

I estimate three specifications for regressing spreads on $v_{it}$.

**Specification 1:** Regresses spreads $y_{it}$ on $v_{it}$ and year & region fixed effects $z_{it}$.

$$y_{it} = \beta' v_{it} + \delta z_{it} + \epsilon_{it}$$

**Specification 2:** Regresses spreads $y_{it}$ on $v_{it}$, lagged government primary balance $s_{i,t-1}$, a second-order polynomial in external debt-GDP ratio $b_{i,t-1}$, and year & region fixed effects. This accounts for nonlinear effects of the conventional state variables: surpluses and debt.

$$y_{it} = \beta' v_{it} + \alpha_1 b_{i,t-1} + \alpha_2 b_{i,t-1}^2 + \gamma s_{i,t-1} + \delta z_{it} + \epsilon_{it}$$

I include the square of debt with an eye to capturing debt-driven increases in spreads in the analogous model regressions. In my model (as in many others), default probabilities are nonlinear in the outstanding debt stocks. A small increase in the debt will usually trigger much larger responses in spreads when the debt burden is high.

**Specification 3:** The same as specification 2, but adds a battery of lagged economic controls $x_{it}$.

$$y_{it} = \beta' v_{it} + \alpha_1 b_{i,t-1} + \alpha_2 b_{i,t-1}^2 + \gamma s_{i,t-1} + \eta x_{i,t-1} + \delta z_{it} + \epsilon_{it}$$

The battery of economic controls are: GDP growth, reserves/imports ratio, inflation, Moody's Baa bond yield index, and the current account. Estimations exclude country-year pairs where the country is in default, as then spreads measure the expected investor losses on old debt rather than the cost of issuing new debt.

Table 2.1 shows the results of these regressions. The first 5 numbers are the vector of coefficients $\beta$, the parameter of interest. This shows that spreads *are* higher after defaults and that this cannot be accounted for by changes in the total amount of debt or the government's primary balance, even allowing for nonlinear effects of outstanding debt and a wide range of controls. The increase in spreads can be economically meaningful. An increase in interest rates of 2pp on a debt of just 40% of GDP would correspond to an increase in total interest payments of 0.8% of GDP - the same as a small recession. This result is robust to excluding the single largest default (Argentina) and including third-order polynomials of debt and surpluses to allow for further nonlinearity in spreads (not shown).

Figure 2.1 compares specification 3 to the most similar specifications in Catão and Mano (2015), Cruces and Trebesch (2013) and Borensztein and Panizza (2009). The interpretation of the coefficients is that they measure the average spread relative to what would have otherwise been predicted using movements in the controls alone. Positive numbers mean that spreads were higher than the controls alone would have suggested, negative numbers that spreads were lower than expected.

Even though the samples and conditioning variables differ across the four estimates[1] the coefficients are much the same. Spreads following default are typically elevated by around 200-250 basis points for at least the two years, and declining thereafter.

The main difference between my work and past studies of post-default spreads is that I choose the specification of the data regressions to be directly comparable to those that I will run later on the model simulations. This is pertinent in two regards. First, unlike in Cruces and Trebesch (2013) I use the "consolidated" default definitions of Benjamin and Wright

1. Both studies also include the government primary balance and outstanding debt as controls though. Where coefficients measure an average over several years, I plot the point at the average number of years before or after a default. For example, the coefficient for 3-4 years post-default is plotted at 3.5 on the x axis.

(2013), which aggregate very close defaults into one default event. Although this reduces the number of default events in the sample, it is appropriate here because my model will typically produce sequences of consecutive defaults that I interpret as a single default event. Second, in my model (as in many others) the impact of debt on spreads is nonlinear. An increase in debt levels has a large effect on default risk (and so spreads) when the debt is high and so default is likely, but only a small effect when debt is small. I need to account for this in my data work too, and so experiment with polynomial terms in the debt.

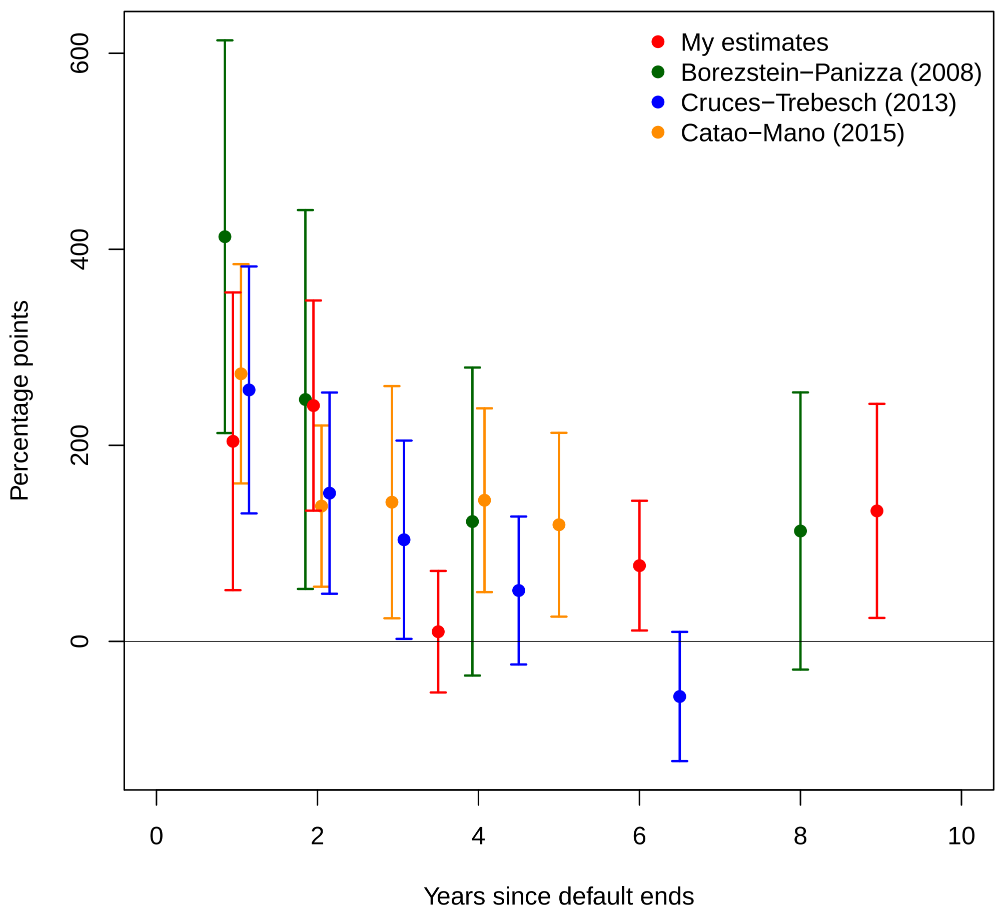


Figure 2.1: Excess spread $\beta$ in specification (3) and other papers' estimates. Error bars show 90% confidence intervals.

Table 2.1: Spreads after return from default

| | *Dependent variable:* | | |
|---|---|---|---|
| | Spread | | |
| | (1) | (2) | (3) |
| 1 yr post-default | 248.10 | 136.92 | 204.12** |
| | (180.11) | (108.39) | (92.33) |
| 2 yrs post-default | 89.71 | 148.64 | 240.57*** |
| | (111.04) | (117.32) | (65.18) |
| 3-4 yrs post-default | −0.02 | −25.33 | 9.87 |
| | (41.06) | (49.42) | (37.69) |
| 5-7 yrs post-default | 34.18 | 54.55 | 77.27* |
| | (59.94) | (43.47) | (40.23) |
| 8-10 yrs post-default | 43.49 | 100.93* | 133.11** |
| | (90.44) | (53.78) | (66.39) |
| Public Debt (% GDP) | | 4.88 | 6.25 |
| | | (7.29) | (10.65) |
| Public Debt squared | | −0.02 | −0.10 |
| | | (0.09) | (0.20) |
| Government primary balance (% GDP) | | −24.19*** | −25.93*** |
| | | (8.41) | (7.52) |
| GDP growth (% p.a.) | | | −16.92*** |
| | | | (4.33) |
| Reserves / imports (%) | | | −1.69 |
| | | | (1.18) |
| Inflation (CPI, %) | | | 0.18 |
| | | | (0.19) |
| Baa corporate bond yield | | | −0.03 |
| | | | (0.13) |
| Current account (% GDP) | | | −1.71 |
| | | | (8.10) |
| Observations | 2,734 | 2,154 | 2,061 |
| $R^2$ | 0.45 | 0.56 | 0.61 |
| Adjusted $R^2$ | 0.45 | 0.55 | 0.60 |

*Note:* *p<0.1; **p<0.05; ***p<0.01
All specifications include region & year fixed effects and a constant.
Standard errors clustered at regional level in parentheses.

### *2.1.3 Discussion*

So what might explain the empirical pattern identified above? The obvious answer is that as the start of the default event recedes into history, the forces which brought about the default also weaken. These forces can be broadly divided into two categories: changing economic circumstances, and changing government priorities. The evidence presented here suggests that the latter is important in explaining post-default spreads, as I already control for the variation in economic circumstance via the battery of controls and fixed effects. If spreads were higher following longer defaults because, for example, output simply had more time to recover, then this would show up the coefficient on GDP.

The alternative explanation, consistent with the data, is that changing government priorities make a significant contribution to the to default decision. However, this empirical work alone cannot establish whether conditional default premia occur as a result of default, or if they are merely correlated. This requires a model, which the subsequent sections outline.

## 2.2 A model of sovereign default

### *2.2.1 Symmetric information*

I start by presenting my benchmark model, where investors and the government have symmetric information about repayment probabilities. If this model can generate a conditional default premium, then this will be purely a correlation, not a consequence of default.

#### Timing

Time is discrete. Periods are labeled $t = 1, 2, , \ldots$[2]. There are two types of actor: a government, and foreign investors. The government starts out with initial primary surplus

2. Time in this model is infinite. In dynamic games such as this, infinite time can give rise to difficult issues of multiplicity and equilibrium selection. To avoid such problems this model should be thought of as the limit of a finite time model $t = 1, \ldots, T$ as $T \to \infty$.

$s_1$, type $\theta_1$ and outstanding debt $b_1$. For all subsequent periods, the timing is as follows: outstanding debt $b_t$ is inherited from the preceding period; next the government's type $\theta_t$ and primary surplus $s_t$ are realized; after which the the government announces its period $t$ default/repayment decision $x_t \in \{0, 1\}$ where $x_t = 1$ signals repayment; the government then issues new debt $b_t$ and follows through on its default/repayment decision. The period then ends. The within-period timing in shown in Figure 2.2.

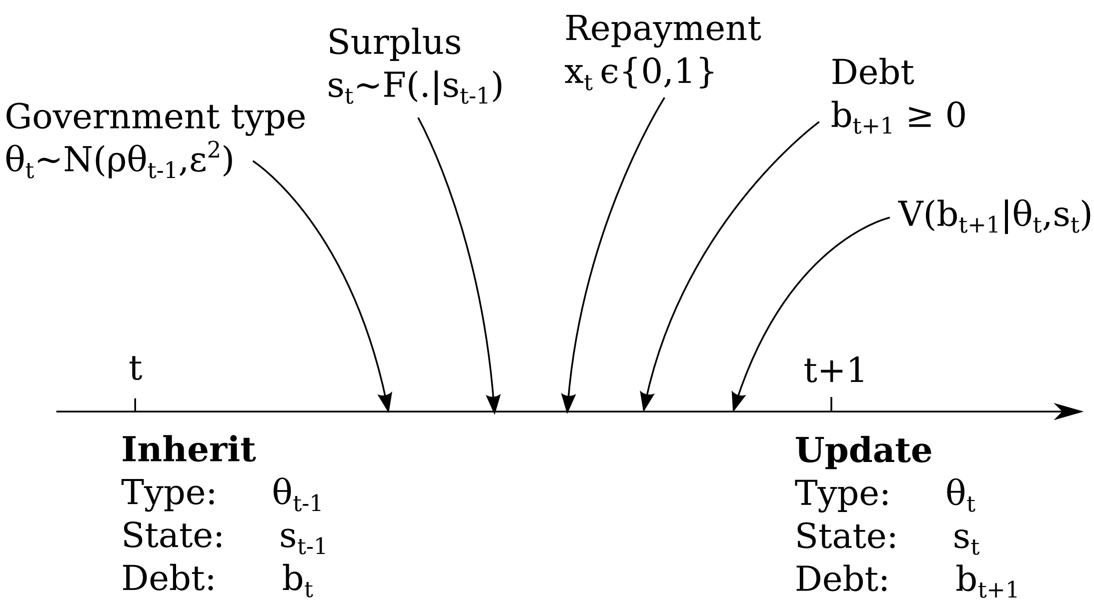


Figure 2.2: Timeline for the symmetric information model

The key aspect of the timing convention in this paper is that the government can commit to its default/repayment decision within the period. This is in common with many papers in the literature and rules out sunspot equilibria of Cole and Kehoe (2000) where investors' expectations of default become self-fulfilling.

## Exogenous processes

Two persistent exogenous processes drive the model. The first is the process for primary government surpluses $\{s_t\}$. This is a discrete first-order Markov process with distribution $s_{t+1} \sim F(\cdot|s_t)$. The surplus is expressed as a fraction of total output in the economy, which grows at a secular rate $g$, so realized surpluses are $s_1, s_2(1+g), s_3(1+g)^2, \ldots$. Including growth in surpluses is important as by increasing the net present value of future incomes it helps raise the debt thresholds at which default occurs. Without this, default occurs at very low levels of debt. Crucially, $s_t$ can be negative; the government may run a primary deficit. The second process is that which governs the government's type, which is a (continuous) AR(1) process, with persistence $\rho$ and innovation variance $\epsilon^2$, so $\theta_{t+1} \sim N(\rho\theta_t, \epsilon^2)$. Where convenient, I will occasionally denote this process by $G(\theta_{t+1}|\theta_t)$.

Here, I assume that $\theta_t$ is uncorrelated with $s_t$. In practice, governments' default preferences and private beliefs are likely correlated. However, only the portion of default preferences that is uncorrelated with the primary balance is relevant for generating a conditional default premium. I therefore restrict the model to only allow for variation in type independent of $s_t$.

In this, the symmetric information version of the model, both $s_t$ and $\theta_t$ are perfectly known to both the government and investors. Because both these variables are persistent, the payoff-relevant state variables at the end of period $t$ are $(s_t, \theta_t, b_{t+1})$.

The introduction of exogenous surpluses in a model of strategic default is uncommon. The benefit of this modeling choice is that it makes the asymmetric information model (in the next section) tractable. A more typical approach is to allow endogenous surpluses, determined by the government's tax and spending choices. To admit this would make the signaling problem in the asymmetric information version of the model very complicated. Governments of different types would potentially be able to signal their type by their choice of surplus (and so new debts). Working with two signals - a discrete repayment/default decision and a continuous surplus choice - makes the problem much more difficult. Assuming exogenous surpluses avoids this problem.

## The government budget constraint

The government has one choice in this model: default or repayment. The evolution of debt differs according to this choice.

*Under repayment*

The government repays its debt through a combination of its primary surplus and revenue raised via new debt issuance. When $s_t = s$, the government of type $\theta$ faces a demand curve for new debt given by $q(b'|s,\theta)$. In equilibrium, investors' payoffs will put restrictions on this function. But for now, we treat it as an unknown function that is decreasing in $b'$, increasing in $\theta$ and that the government takes as given when issuing debt.

When the government repays its debts, it issues $b_{t+1}$ satisfying:

$$b_t - s_t = (1+g)b_{t+1}q(b_{t+1}|\theta_t, s_t) \tag{2.1}$$

Recall that $g$ is the secular growth rate of the economy. This means that the units of $b_t$ and $s_t$ are as a fraction of output (which grows at a constant rate $g$). The price $q$ is then that of a bond which in period $t+1$ pays out an amount equal to one unit of period $t$ output (so, of course, the government typically sells fractions of this bond, and $b_t < 1$).

The government has no savings technology, so $b_t \geq 0$ always. If $s_t > b_t$, then $b_{t+1} = 0$ and the excess surplus $s_t - b_t$ is rebated back to households (more on this in the following section, "Government payoffs").

*In default*

If the government fails to meet its obligations $b_t$, it partially reschedules the debt, paying investors nothing in period $t$ but issuing them with claims $\psi b_t$ due in period $t+1$ where $\psi \in [0,1)$ is a model parameter. Therefore, in default:

$$b_{t+1} = \frac{\psi b_t}{1+g} \tag{2.2}$$

In default, outstanding deficits are unfunded (and so some spending cannot be met). Likewise, any surpluses are rebated to households. Partial default of this kind is desirable for three reasons. First, loss given default in the data is typically less than 100%. Using a large sample of debts, Benjamin and Wright (2013) estimate the average net present value haircut to be around 50%. Second, that partial losses will allow empirically plausible default thresholds. When $\psi = 0$, the loss in default is total, which in equilibrium means that debt prices fall very sharply as debt increases, and so in equilibrium causes the government to default at very low debt thresholds. Third, writedowns of this form typically generate a sequence of multi-period defaults that look very similar to real-world default events. Countries are excluded from capital markets for a number of periods and negotiate with their creditors for relief until they eventually settle, typically still owing their creditors a fraction of the original debt. As also shown by Benjamin and Wright, longer defaults are typically associated with larger writedowns. The assumptions on default allow the model to capture this in reduced form.

## The Laffer Curve and the Inverse Revenue Function

Revenue raised by a given debt issuance $b'$ in state $(s, \theta)$ is:

$$Rev(b'|s,\theta) = (1+g)b'q(b'|s,\theta) \tag{2.3}$$

If the net present value of future surpluses is finite (as they will be if $1+g$ is less than the gross risk-free rate), then this is bounded. Figure 2.3 shows a stylized revenue Laffer curve. As debt increases, the price falls, slowing the rate at which further debt issuance increases revenue raised. Eventually, revenue declines as more debt is issued. I also assume that:

$$\arg\max_{b'>0} Rev\left(b'|s,\theta\right) < \infty$$

I verify in section 2.2.1 that this holds in equilibrium. So the revenue function looks like a Laffer curve, with peak revenue raised with finite debt issuance. Figure 2.3 shows a stylized representation this curve. The maximum revenue is:

$$Rev_{max}(s,\theta) = \max_{b'>0} Rev(b'|s,\theta) \tag{2.4}$$

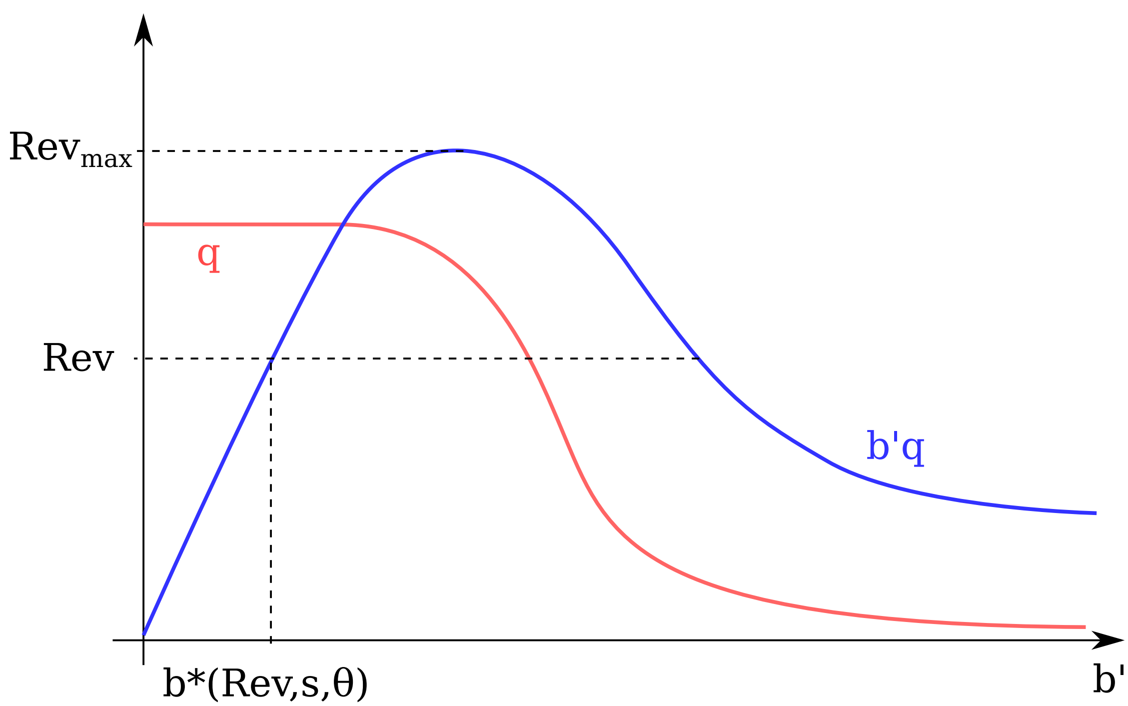


Figure 2.3: Stylized Laffer curve

I assume that the government can commit to a given debt issuance (thereby sidestepping the issues of dynamic multiplicity raised in Lorenzoni and Werning (2014)). Given a target revenue $y$, the government issues $b^*(y,s,\theta)$ where:

$$b^*(y,s,\theta) = \min_{b'\geq 0} b' \qquad \text{s.t.} \qquad Rev(b'|s,\theta) = y \tag{2.5}$$

### Government Payoffs

The government has two possible actions: repay or default. However, repayment might not be feasible, as a low debt demand curve may mean that the government is simply unable to raise the revenue required. So the set of feasible actions for the government is:

$$A(b, s, \theta) = \begin{cases} \{Repay,\ Default\} & \text{if} \quad b \leq s + Rev_{max}(s, \theta) \\ \{Default\} & \text{if} \quad b > s + Rev_{max}(s, \theta) \end{cases}$$

Appendix A.3 derives primitive assumptions under which the government's payoff is the household's indirect utility. This payoff is:

$$\begin{aligned} &\text{Default:} && u_d(b, s, \theta) = \min(s, 0) - (k + \theta) \\ &\text{Repayment:} && u_r(b, s, \theta) = 0 \end{aligned}$$

And the government discounts future payoffs at rate $\beta$.

In this set-up, the costs of default come through two channels: the private cost $(k + \theta)$; and the exposure to surplus variability in $\min(s, 0)$. Set against this is the benefit of default: by lowering $b'$ default reduces interest payments and defers the likely date of the next default.

Although the exact payoffs here are different from in a "standard" sovereign default model (such as Arellano (2008) ), the main mechanisms are the same. There, the costs of default come from two sources: an exogenous loss of output, and consumption variability due to income fluctuations in autarky. The payoffs here can be thought of mimicking these two channels. There is a cost of default independent of fundamentals, $k + \theta$. And there is a loss due to unsmoothed variation in permanent income, $\min(s, 0)$.

The government value function

At the end of a period with state $(b', s, \theta)$, the government's expected future payoff is:

$$V(b'|s,\theta) = \int_{s'}\int_{\theta'} \max \begin{bmatrix} \beta V\left(b^*\left(\max\left(0, b'-s'\right), s', \theta'\right)|s',\theta'\right) 1_{A(b',s',\theta')}, \\ u_d(b',s',\theta') + \beta V\left(\psi b'/(1+g)|s',\theta'\right) \end{bmatrix} dG(\theta'|\theta)dF(s'|s) \tag{2.6}$$

$$V(0|s,\theta) = \int_{s'}\int_{\theta'} \beta V\left(b^*\left(\max\left(0, -s'\right), s', \theta'\right)|s',\theta'\right) dG(\theta'|\theta)dF(s'|s) \tag{2.7}$$

Where $1_{A(b',s',\theta')}$ is 1 if repayment is feasible and $-\infty$ if not, and the second line is a special case - default is impossible when $b' = 0$.

**Proposition 1** (Monotonicity of the value function)**.** *The government's value function under symmetric information is decreasing in $b'$ for each $(s, \theta)$.*

*Proof.* Set $b_1 = b$, $s_1 = s$, $\theta_1 = \theta$. For any future realization of the exogenous processes $\{\theta_t, s_t\}$, let $x_t = \hat{x}_t(\theta_t, s_t)$ be the optimal default rule in each exogenous state of the world, and $b_t = \hat{b}_t(\theta, s_t)$ be the associated debt level generated by the budget constraints (2.1) and (2.2). Define $\tau = \arg\min_{t>0} t$ s.t. $\hat{b}_t(\theta_t, s_t) = 0$ as the first time that the debt level is zero.

Now imagine that $b_1 = \tilde{b} < b'$. If the government follows the same default rule $\hat{x}_t(\theta_t, s_t)$, then for any realization of $\{\theta_t, s_t\}_{t=0}^{\infty}$, $b_t \leq \hat{b}_t(\theta_t, s_t)\ \forall\ t < \tau$ and $b_t = \hat{b}_t(\theta_t, s_t)\ \forall\ t \geq \tau$. As a result the period payoff is the same whenever $x_t = 0$ and weakly larger when $x_t = 1$. So for any realization of $(s_t, \theta_t)$ the payoff must be weakly larger. The same must also hold true for the integral. Then we have that:

$$V(\tilde{b}, s, \theta) \geq V(b', s, \theta)$$

□

### Default decision

There are three regions of $(s, \theta, b)$ space:

1. If $b > s + Rev_{max}(s, \theta)$ the government always defaults.

2. If $s < 0$ or $s < b < s + Rev_{max}(s, \theta)$ the government defaults if:

$$\min(s, 0) - (k + \theta) + \beta V(\psi b/(1+g)|s, \theta) > \beta V(b^*(b - s, s, \theta)|s, \theta)$$

3. And if $s > b > 0$, the government defaults if:

$$b - (k + \theta) + \beta V(\psi b/(1+g)|s, \theta) > \beta V(0|s, \theta)$$

I look for a threshold solution in $\theta$. That is, $\hat{\theta}(s, b)$ such that the government of type $\theta$ repays in when the surplus is $s$ and the debt level is $b$ if $\theta > \hat{\theta}(s, b)$. This is the solution when $\beta$ is sufficiently small. Because $q(b', s, \theta)$ is increasing in $\theta$, region 1 is shrinking in $\theta$. And for small enough $\beta$, changes in the continuation in regions 2 & 3 are dominated by the changes in the period payoff. Saying more than this is difficult, as the effect of changes in $\theta$ on the value function are ambiguous. Lower $\theta$ reduces the cost of default in current and future states of the world, increasing the period payoff in default. But lower $\theta$ also reduces the price at which government debt is sold, increasing debt issuance required to meet a given level of revenue and depressing the continuation value under repayment. In any numerical solution, I assume that such a threshold equilibrium holds and verify ex post that this is the optimal strategy. A stylized default boundary is shown in Figure 2.4

### Prices

External investors are atomistic and risk-neutral. They have an outside gross rate of return on funds of $R$. The government sells $(1 + g)b'$ units of debt. The next-period payoff in state

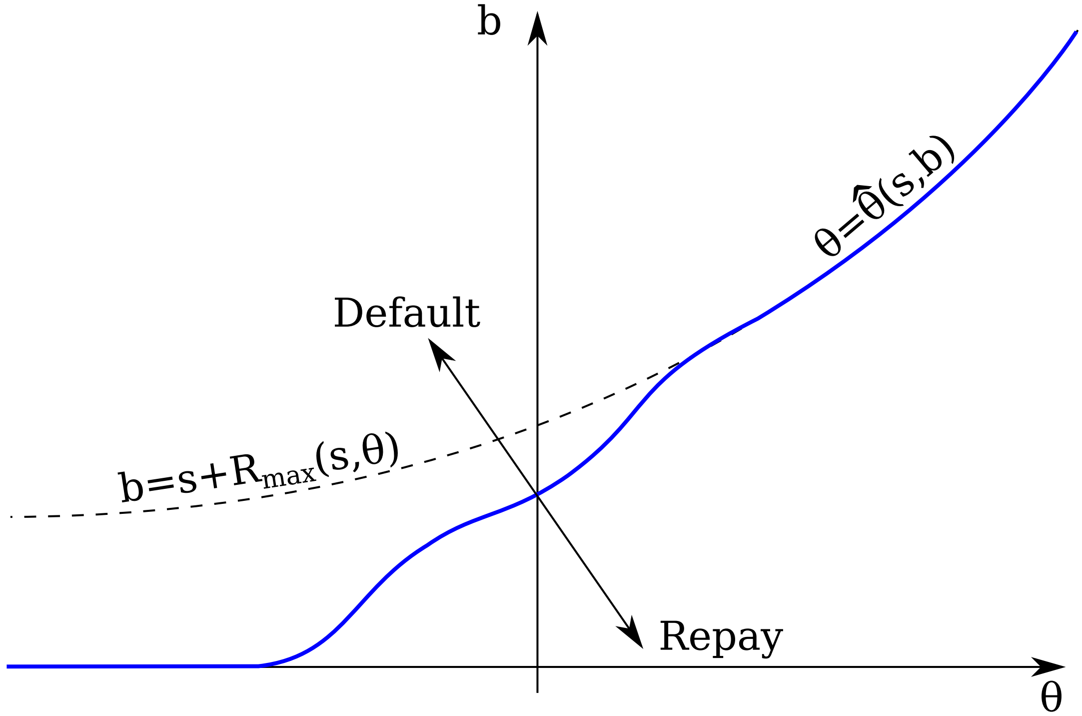


Figure 2.4: Stylized default boundary under symmetric information (for fixed $s$)

$(s', \theta', b')$ is then:

$$\begin{cases} 1 & \text{if } \; x' = 1 \\ \dfrac{\psi}{1+g} q\left(\dfrac{\psi b'}{1+g} \,\middle|\, s', \theta'\right) & \text{if } \; x' = 0 \end{cases}$$

Investors decide whether to lend to the government after the government has announced their period default decision $x_t$ and their new debt issuance $b'$. Investors have no cost other than the opportunity cost of funds $R$. So if issuance is $b' \in (0, \infty)$, then the supply of loans equals this exactly only if a zero-profit condition holds:

$$q(b'|s,\theta) = \frac{1}{R} \int_{s'} \left\{ \Phi\left(\frac{\rho\theta - \hat{\theta}(b', s')}{\epsilon}\right) + \int_{\theta' < \hat{\theta}(b', s')} \left(\frac{\psi}{1+g}\right) q\left(\frac{\psi b'}{1+g} \,\middle|\, s', \theta'\right) dG(\theta'|\theta) \right\} dF(s'|s) \tag{2.8}$$

Where $\hat{\theta}(b', s')$ is the cutoff type and $\Phi(\cdot)$ is the CDF of the standard normal. The formula for prices highlights the two sources of risk that investors bear: from the realization of the surplus, and of the government type. However, governments face exactly the same uncertainties. There is no information advantage of either party over the other.

I now show that the revenue-maximizing debt issuance is always finite. This is important as it means that the state space is bounded.

**Proposition 2** (Finite revenue-maximizing debt)**.** $\arg\max_{b' \geq 0} Rev_{max}(b'|s,\theta) < \infty \ \forall \ (s,\theta)$

*Proof.* Fix $(s,\theta)$. Because $s'$ and $Rev_{max}(s',\theta')$ are finite, then there exists $b'$ such that $b^* > s' + Rev_{max}(s',\theta') \ \forall \ (s',\theta')$. For such $b^*$:

$$q(b^*|s,\theta) = \frac{1}{R} \int_{s'} \int_{\theta' < \hat{\theta}} \left( \frac{\psi}{1+g} \right) q \left( \frac{\psi b^*}{1+g} \,\middle|\, s', \theta' \right) dG(\theta'|\theta) dF(s'|s)$$

$$\begin{aligned} \Rightarrow \qquad Rev(b^*|s,\theta) &= (1+g) b^* q(b^*|s,\theta) \\ &= \left( \frac{1+g}{R} \right) \int_{s'} \int_{\theta' < \hat{\theta}} \left( \frac{\psi b^*}{1+g} \right) q \left( \frac{\psi b^*}{1+g} \,\middle|\, s', \theta' \right) dG(\theta'|\theta) dF(s'|s) \\ &= \left( \frac{1+g}{R} \right) \int_{s'} \int_{\theta' < \hat{\theta}} Rev \left( \frac{\psi b^*}{1+g} \,\middle|\, s', \theta' \right) dG(\theta'|\theta) dF(s'|s) \end{aligned}$$

If $G^n(\cdot|\theta)$ and $F^n(\cdot|s)$ are the distributions of $\theta_{t+n}$ and $s_{t+n}$ conditional on $\theta_t = \theta$ and $s_t = s$, then:

$$Rev(\psi^{-n} b^*|s,\theta) = \left( \frac{1+g}{R} \right)^n \int_{s'} \int_{\theta' < \hat{\theta}} Rev \left( \frac{b^*}{1+g} \,\middle|\, s', \theta' \right) dG^n(\theta'|\theta) dF^n(s'|s)$$

As $n \to \infty$, then since $G$ is ergodic revenue tends to zero at rate $(1+g)/R$. And because revenue is strictly positive for some values of $b'$, then it must be maximized for finite $b'$ $\square$

An obvious corollary of this proposition is that there exists $\bar{b}$, an upper bound on debt such that $b_t \leq \bar{b} \ \forall \ t$.

## Default events

Because default is only partial and primary balance is serially correlated, the model will often generate sequences of back-to-back default. In these cases, debt repayment is postponed multiple periods, written down by a fraction $\psi$ at each postponement. For example, if $x_t =$ for $t \in \{\tau, \tau+1, \tau+2\}$ and $x_{\tau+3} = 1$, then investors who lent in period $\tau - 1$ have their debts written down and postponed three times. In the end, creditors owed 1 at $t = \tau$ receive $\psi^3$ at $t = \tau + 3$.

In these cases, I consider periods $\tau, \tau+1, \tau+2$ a single *default event*, lasting three periods and with total loss-given default of $(\psi/R)^3$ (accounting not only for the write down of the debt's face value, but also the postponement of the eventual repayment). As a result, the loss-given-default and period of exclusion are endogenous, with write-downs and exclusion continuing until the government decides to final repay.

## Equilibrium

A threshold Markovian equilibrium is given by functions $V(b'|s,\theta)$, $\hat{\theta}(b,s)$, $x(b,s,\theta)$, $b'(b,s,\theta)$, $Rev(b'|s,\theta)$, $Rev_{max}(b'|s,\theta,$ $b^*(y,s,\theta)$ and $q(b'|s,\theta)$ such that:

1. The cutoff $\hat{\theta}(b,s) = \max_\theta \theta$ s.t. $x(b,s,\theta) = 0$ for all $(b,s)$

2. Prices $q(b'|s,\theta)$ satisfy the zero-profit condition equation (2.8) for all $(b',s,\theta)$ given the cutoff function $\hat{\theta}(b,s)$.

3. Given $q(b'|s,\theta)$, the revenue, maximum revenue and inverse revenue functions $Rev(b'|s,\theta)$, $b^*(y,s,\theta)$ satisfy equations (2.3), (2.4) and (2.5) for all $(b',s,\theta,y)$.

4. For all $(b,s,\theta,y)$, $b'(b,s,\theta)$ satisfies equation (2.2) when $x(b,s,\theta) = 0$ and equation (2.1) when $x(b,s,\theta) = 1$.

5. For all $(b,s,\theta)$, the repayment indicator $x(b,s,\theta)$ is 0 whenever the maximand in equation (2.6) selects the first term, and 1 otherwise given $b^*(y,s,\theta)$, $V(b'|s,\theta)$.

6. Given $b^*(y, s, \theta)$, the value function satisfies equations (2.7) for all $(s, \theta)$ whenever $b' = 0$ and (2.6) otherwise.

I now extend the model to the case where the type is not directly observed by investors. Instead it is inferred from the government's history of repayment-default decisions.

### *2.2.2 Asymmetric information model*

I now consider what happens when $\theta$ is not publicly known. The government cannot signal their type except through their repayment-default decision. This is rather strict. In reality, governments can probably signal their default preferences through a variety of means. However, any signals that are correlated with fundamentals (for example, via the choice of the surplus or debt) will not help explain the conditional default premium because this is orthogonal to fundamentals by construction. Any model which hopes to generate a conditional default premium must leave some room for default to communicate information about type. A model where only default reveals type is the simplest one.

With asymmetric information, the distribution of beliefs over types becomes a state variable. With a continuum of types this is a potentially infinite-dimensional object. I simplify this problem by exploiting the threshold nature of the equilibrium. This allows me to write down a good second-order approximation to the distribution of beliefs, compressing the state space to only two dimensions. I start by introducing the general problem, and then simplify the model using this approximation.

#### Set up

The timeline is similar to the case with symmetric information, except now investor beliefs at the start of each period are given by the endogenous object $p_t(\theta_{t+1})$. Relevant beliefs are those over the next period type as investors lend at a price informed by the probability of default in the following period. I use the notation $p' = \mathcal{P}(s, b, x, p)$ to denote the dependence of updated beliefs $p'$ on the observed states $(s, b)$, the action $x$, and the prior $p$. In equilibrium, Bayes' law will restrict $\mathcal{P}(\cdot)$.

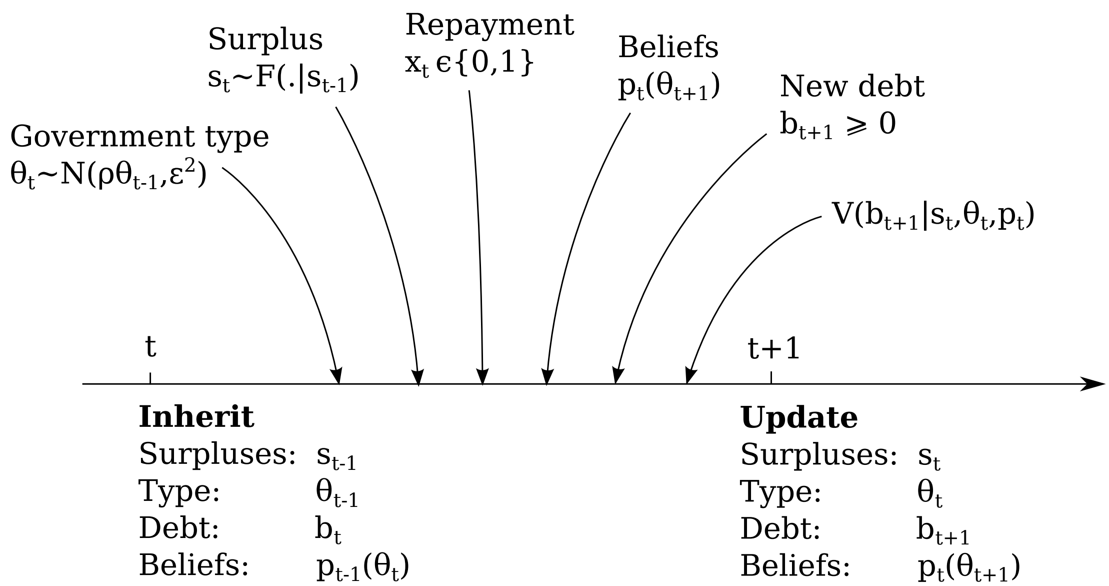


Figure 2.5: Timeline for the symmetric information model

## Revenue

When information is asymmetric, prices are given by[3] $q(b'|s, p')$, which is:

$$q(b'|s,p') = \frac{1}{R}\int_{s'} \left\{ \begin{array}{c} \left(\frac{\psi}{1+g}\right) q\left(\frac{\psi b}{1+g}\middle| s', \mathcal{P}(s', b', 0, p')\right) P'\left(\hat{\theta}(b', s', p')\right) \\ +1 - P'\left(\hat{\theta}(b', s', p')\right) \end{array} \right\} dF(s'|s) \quad (2.9)$$

Where $P'(\cdot)$ is the CDF of the belief distribution $p'$.

Asymmetric information influences the debt price in three ways. There is a direct effect of $p'$ on the probability of repayment. For example, for a given cutoff $\hat{\theta}$, the probability of repayment increases. There is also an indirect effect of $p'$ on repayment probabilities through

3. To save on notation, I do not differentiate between the symbols representing analogous objects in the symmetric and asymmetric cases. For example, prices in both cases are denoted $q$. Which object is in use should be obvious from the context and the function arguments.

its impact on $\hat{\theta}$. Typically, if higher $\theta$ is more likely then $\hat{\theta}$ will decline, as high reputation is valuable, and hence worth preserving. Finally, there is an indirect effect via the resale price. When default is particularly indicative of a low type (for example, b is small), then $\mathcal{P}(s', b', p', 0)$ will put more weight on low $\theta$, reducing the value of the written-down debt in default.

Note that prices are actually easier to compute in this setting than in the symmetric information case. Because continuation prices depend only on repayment and default directly, and not on $\theta'$, one integral is removed.

The revenue and inverse revenue functions are now:

$$Rev(b'|s, p') = b'q(b'|s, p') \tag{2.10}$$

$$Rev_{max}(s, p') = \max_{b'>0} Rev(b'|s, p') \tag{2.11}$$

$$b^*(y, s, p') = \min_{b'>0} b' \quad \text{s.t.} \quad Rev(b'|s, p') = y \tag{2.12}$$

## Government payoffs & value function

The set of feasible actions is conceptually the same as in the symmetric information case, but now the Laffer Peak depends explicitly on beliefs.

$$A(b, s, p) = \begin{cases} \{Repay,\ Default\} & \text{if} \quad b \leq s + R_{max}(s, \mathcal{P}(s, b, 1, p)) \\ \{Default\} & \text{if} \quad b > s + R_{max}(s, \mathcal{P}(s, b, 1, p)) \end{cases}$$

The period payoffs are otherwise unchanged:

$$u_d(b, s, \theta) = \min(s, 0) - (k + \theta)$$
$$u_r(\theta, b, s, p) = 0$$

Then the government's value function is now:

$$V(b'|s,\theta,p') = \int_{s'}\int_{\theta'} \max\left[\begin{array}{c}\beta V\left(b^*\left(\max(0,b'-s'),s',p'\right)|\theta',s',p'_r\right)1_{A(b',s',p')},\\ u_d(\theta',b',s')+\beta V\left(\psi b/(1+g)|s',\theta',p'_d\right)\end{array}\right]\dots$$

$$\dots dG(\theta'|\theta)dF(s'|s) \tag{2.13}$$

$$V(0|s,\theta,p') = \int_{s'}\int_{\theta'} \beta V\left(b^*\left(\max(0,b'-s'),s',p'_0\right)|\theta',s',p'_0\right)dG(\theta'|\theta)dF(s'|s) \tag{2.14}$$

Where:

$$p'_r = \mathcal{P}(s,b,1,p) \qquad p'_r = \mathcal{P}(s,b,0,p) \qquad p'_0 = \mathcal{P}(s,0,1,p)$$

And so the government defaults if $b > s + R_{max}(s,\mathcal{P}(s,b,1,p))$, or if:

$$u_r(\theta',b',s',p')+\beta V\left(b^*\left(\max(0,b'-s'),s',p'_r\right)|\theta',s',p'_r\right) <$$
$$u_d(\theta',b',s')+\beta V\left(\psi b'/(1+g)|s',\theta',p'_d\right)$$

Here default influences the continuation value of default beyond simply the debt issuance choice. Now, the continuation values under repayment and default also differ due to the effect that this has on beliefs about type. Under default, future beliefs about type are worse, pushing debt prices down, increasing future debts and so lowering payoffs. The response of future prices to default acts as an extra, endogenous cost of default. As a result, default boundaries are typically much higher under asymmetric information. The threat of loss of credibility and the added cost of borrowing that this entails allows the government to sustain higher debts than it otherwise would. In the symmetric information model, there was no such channel; future price functions depended only on $s$ and $\theta$, and not on the default-repayment decision. This is discussed further later.

Equilibrium

A threshold Markovian equilibrium of the asymmetric information model is given by functions $V(b'|s,\theta,p')$, $\hat{\theta}(b,s,p')$, $x(b,s,p')$, $b'(b,s,p')$, $Rev(b'|s,p')$, $Rev_{max}(b'|s,p')$, $b^*(y,s,p')$, $q(b'|s,p')$ and $\mathcal{P}(s,b,x,p)$ such that:

1. The cutoff $\hat{\theta}(b,s,p') = \max_\theta \theta$ s.t. $x(b,s,p') = 0$ for all $(b,s,p')$

2. Prices $q(b'|s,p')$ satisfy the zero-profit condition equation (2.9) for all $(b',s,p')$ given the cutoff function $\hat{\theta}(b,s,p')$.

3. Given $q(b'|s,p')$, the revenue, maximum revenue and inverse revenue functions $Rev(b'|s,p')$, $Rev_{max}(b'|s,p')$, and $b^*(y,s,p')$ satisfy equations (2.10), (2.11) and (2.12) for all $(b',s,p',y)$.

4. For all $(b,s,p',y)$, new debt issuance satisfies the government budget constraint.

$$b'(b,s,\theta) = \begin{cases} b^*(b-s,s,p') & \text{if } \; x(b,s,p') = 1 \\ \frac{\psi b}{1+g} & \text{if } \; x(b,s,p') = 0 \end{cases}$$

5. For all for all $(b,s,p')$, the repayment indicator $x(b,s,p')$ is 0 whenever the maximand in equation (2.13) selects the first term, and 1 otherwise given $b^*(y,s,p')$, $V(b'|s,p')$.

6. Given $b^*(y,s,p')$, the value function satisfies equations (2.14) for all $(s,p')$ whenever $b' = 0$ and (2.13) otherwise.

7. Given $\hat{\theta}(b,s,p')$, for all $(b,s,p',x)$ the belief updating functions satisfy Bayes' law:

$$\mathcal{P}(s,b,x,p) = \begin{cases} \int_{-\infty}^{\infty} p(\theta)\phi\left(\frac{\theta'-\rho\theta}{\epsilon}\right) d\theta & \text{if } \; b = 0 \\ \left(1 - P(\hat{\theta}(s,b,p))\right)^{-1} \int_{\hat{\theta}(s,b,p)}^{\infty} p(\theta)\phi\left(\frac{\theta'-\rho\theta}{\epsilon}\right) d\theta & \text{if } \; b > 0 \;\&\; x = 1 \\ P(\hat{\theta}(s,b,p))^{-1} \int_{-\infty}^{\hat{\theta}(s,b,p)} p(\theta)\phi\left(\frac{\theta'-\rho\theta}{\epsilon}\right) d\theta & \text{if } \; b > 0 \;\&\; x = 0 \end{cases}$$

We can guarantee the existence of threshold Markovian equilibria (and that this is the only type of Markovian equilibrium) if either $\beta$ or $\rho$ are not too large.

**Proposition 3** (Existence of threshold equilibria)**.** *If* $\beta < 1/(2\rho)$*, then:*

*a) There exists a Markovian equilibrium*

*b) Any Markovian equilibrium can be characterized by a threshold rule* $\hat{\theta}(b, s, p)$

The proof is in Appendix A.1. The intuition for this result, especially part b), is that the existence of a threshold rule depends on the comparison of the (linear) effect of $\theta$ on the current period payoff against the (nonlinear) change in the difference in continuation values under default and repayment. In the proof I show that the change in the difference in continuation values is less than one-for-one (and hence a cutoff rule exists) if $\beta$ is small enough. The reason is simply that lower $\beta$ reduces the importance of the continuation values, and so allows the period payoff to dominate. The role of $\rho$ is to provide a lower bound on what "small enough" means. As $\rho$ shrinks, differences in the current value of $\theta$ cause ever smaller differences in expected future payoffs. And so smaller $\rho$ means smaller changes in the continuation value functions with respect to $\theta$. As a result, if $\rho$ is small, a larger value of $\beta$ is required to dampen the effect of future payoffs and so permit a threshold rule.

## Approximating beliefs

In Section 3.1, I consider an abstract signaling problem where a hidden state $\theta_t$ follows an AR(1) process which is partly revealed by a sequence of signals $\{\hat{\theta}_t, x_t\}$, where $x_t = 1$ if and only if $\theta_t \geq \hat{\theta}_t$. I show that if prior beliefs about the hidden state are distributed according to $N(\mu_t, \sigma_t^2)$ at the start of period $t$, then if beliefs about $\theta_t$ are updated according to Bayes' rule, at the start of period $t+1$ they can be well-approximated[4] by $N(\mu_{t+1}, \sigma^2_{t+1})$, where:

$$\mu_{t+1} = \begin{cases} \rho\left(\mu_t + \sigma_t h\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t}\right)\right) & \text{if } x_t = 1 \\ \rho\left(\mu_t - \sigma_t r\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t}\right)\right) & \text{if } x_t = 0 \end{cases}$$

$$\sigma^2_{t+1} = \begin{cases} \epsilon^2 + \rho^2\sigma_t^2\left[1 + h\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t}\right)\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t} - h\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t}\right)\right)\right] & \text{if } x_t = 1 \\ \epsilon^2 + \rho^2\sigma_t^2\left[1 - r\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t}\right)\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t} + r\left(\frac{\hat{\theta}_t - \mu_t}{\sigma_t}\right)\right)\right] & \text{if } x_t = 0 \end{cases}$$

Where $h(\cdot)$, $r(\cdot)$ are the hazard and reverse hazard rates for the normal distribution:

$$h(z) = \frac{\phi(z)}{1 - \Phi(z)} \qquad\qquad r(z) = \frac{\phi(z)}{\Phi(z)}$$

There is some intuition for these expressions. Imagine that $\hat{\theta}_t$ is very large[5]. Then $x_t = 1$ is the most likely signal. In this case, $(\mu_{t+1}, \sigma^2_{t+1}) \simeq (\rho\mu_t, \rho^2\sigma_t^2 + \epsilon^2)$, then unconditional update. In contrast, if $x_t = 0$, then the signal is highly informative. The posterior belief puts high weight on $\theta_t \simeq \hat{\theta}_t$ and so $(\mu_{t+1}, \sigma^2_{t+1}) \simeq (\rho\hat{\theta}_t, \epsilon^2)$. The exact formulae for $(\mu_{t+1}, \sigma^2_{t+1})$ smooth these two extremes.

This problem is, of course, the exact problem that investors face when inferring the government's type from their history of default/repayment decisions when there is a threshold equilibrium. The threshold $\hat{\theta}$ varies with the state but the updating formulae are unchanged. In solving the model, I assume that beliefs follow this updating scheme. This makes the state space of the model manageable. Instead of having to keep track of the entire distribution $p(\theta)$, the space of beliefs can be fully captured by two variables: the mean and variance of beliefs $\mu$ and $\sigma^2$.

4. The max average error on the CDF is typically less than 0.2pp.

5. The logic for the case with small $\hat{\theta}_t$ is identical but of opposite sign.

# CHAPTER 3

# DEFINITION AND APPLICATION OF THE THRESHOLD FILTER TO THE SOVEREIGN DEFAULT MODEL

## 3.1 The threshold filter

### *3.1.1 Introduction*

#### Motivation

This chapter defines the threshold filtering problem and proposes a solution, the threshold filter. In general, a filtering problem is any situation where an actor (either an econometrician or an agent inside an economic model) attempts to infer the distribution of a hidden type from a sequence of partially informative signals. In the context of the sovereign default model I have presented in the preceding chapters, the threshold filtering problem describes the situation of investors trying to infer the distribution of government types from their repayment history. However, the threshold filtering problem also has much more general interest because it can shows up in a wide variety of other economic problems.

For example, an employer evaluating a prospective employee's hidden quality may see their work history. This is a sequence of binary signals; in each period the individual was either employed or unemployed. Or a bank assessing a potential borrower's credit quality may see their past repayment history.

#### Approximation and filtering

The general solution to any filtering problem is conceptually simple: use the laws for the state and signal process to apply Bayes's rule to the set of signals to recover a distribution of the hidden type at each point in time. In practice, though, this is challenging. Unless the law of motion for the state is conjugate to the distribution of the signals, the distribution at time $t$ is a function of the time-zero prior and *all* subsequent signals. The dimension of the description of the period-$t$ distribution therefore grows linearly with $t$. The propensity of the state description to grow without bound is problematic even in the original application of filtering - engineering problems. But it is doubly difficult in economic settings where the model is driven by agents' beliefs about the hidden state. Linear growth in the dimension of the state space precludes any hope of expressing the model recursively.

As a result, nonlinear filtering problems are really approximation problems: how best can one represent the posterior distribution with as few parameters as possible. The filtering technique that I develop approximates the true distribution with a two-parameter distribution (the normal). One of the aims of this section is therefore to provide metrics against which the accuracy of this approximation be assessed. I use an $L^0$ norm to compare the accuracy of the threshold filter to the true distribution derived from Bayes's law. The $L^0$ normal is appealing for a number of reasons: it is a whole-distribution measure of the error; it is easily interpreted[1]; and it is scale-invariant.

Of course, computing the true distribution in order to make this comparison is itself very difficult. This is one of the reasons for developing the threshold filter in the first place. I therefore use two benchmark distributions in place of the true distribution, which capture different aspects of threshold filter error: a very large discretized Markov approximation, and the exact one-step distribution.

My findings are that for a moderately persistent state the threshold filter is an excellent approximation to the true distribution. In a very long simulation the threshold filter has average absolute (the $L^0$ norm) of 0.2pp, and maximal one of 0.5pp. In contrast, the best alternative filter in the literature (the exact Gaussian filter) has much larger errors, particularly in the worst case, with average and maximum errors of 0.4pp and 3pp respectively. For high $\rho$ the approximation does deteriorate, but still produces low average errors for

1. The $L^0$ norm is the average absolute error on measuring events using the approximate distribution when they are drawn from the true distribution

persistence less than around 0.9.

### *3.1.2 The threshold filter*

## The threshold filtering problem

The general nonlinear filtering problem is given by:

$$x_k = f(x_{k-1}) + u_k \qquad\qquad u_k \sim N(0, Q_k)$$
$$y_k = g(x_{k-1}) + v_k \qquad\qquad v_k \sim N(0, R_k)$$

For $k = 0, 1, \ldots$, where $x_k$ is the hidden state, $y_k$ are the signals, and $u_k$, $v_k$ are independent of each other and $\{x_l\}_{l=0}^{k}$. The threshold filtering problem is a nonlinear filtering problem where $x$ follows a known AR(1) process and $y_k = 1$ if $x_k > \hat{x}_k$ and is zero otherwise.

That is:

$$f(x) = \rho x \qquad\qquad Q_k = \epsilon^2$$
$$g(x) = 1_{\{x_k > \hat{x}_k\}} \qquad\qquad R_k = 0$$

And the threshold $\hat{x}_k$ is both known and independent of the full history of states $\{x_l\}_{l=0}^{k}$. Here $f(\cdot)$ is linear, but the signal process is highly nonlinear.

Figure 3.1 illustrates the evolution of the state and the approximation within a period. Panel a) displays the prior $\pi(\cdot)$ over the hidden state $x$. Panel b) contains a realization of the threshold signal. In this case, the threshold is a lower bound on the hidden state, so the econometrician learns that $x > \hat{x}$ (i.e. in this period, $y = 1$). Panel c) then describes how the exact posterior is formed (that is, the distribution of the type at the start of the next period, denoted $x'$). The posterior is built up by integrating the distribution of $x'$ conditional on $x$ over the prior measure restricted to $x > \hat{x}$. The resulting posterior is typically skewed, as there is no likelihood contribution from the conditional distributions with $x < \hat{x}$. The approximating distribution, however, is symmetric. A central aim of this section is to argue that the inaccuracies induced by approximating the true skewed distribution with a symmetric one are very small.

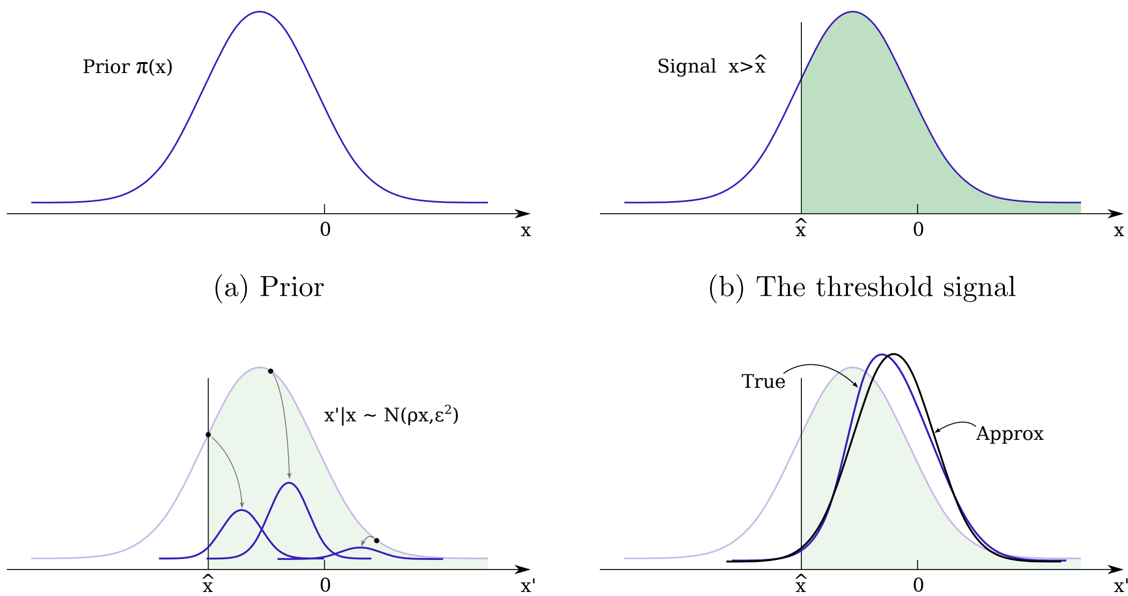


Figure 3.1: The evolution of the state

## The threshold filter

The threshold filter generates a Gaussian approximation to the state of the next-period distribution of $x$.

**Algorithm** (The threshold filter)**.** *If the distribution of the hidden state in period $k$ conditional on the signals $\{y_l\}_{l=0}^{k-1}$ is given by $x_{k|k-1} \sim N(\mu_t, \sigma_t^2)$, then the threshold filter predicts* *an updated distribution in period $k+1$ of $x_{k+1|k} \sim N(\mu_{t+1}, \sigma^2_{t+1})$:*

$$\mu_{k+1} = \begin{cases} \rho\left(\mu + \sigma h(\psi_k)\right) & \text{If } \; y_k = 1 \\ \rho\left(\mu - \sigma r(\psi_k)\right) & \text{If } \; y_k = 0 \end{cases}$$

$$:= \mu'(\mu_k, \sigma_k^2, \psi_k, y_k)$$

$$\sigma^2_{k+1} = \begin{cases} \epsilon^2 + \rho^2\sigma_k^2\left(1 - h(\psi_k)^2 + \psi h(\psi_k)\right) & \text{If } \; y_k = 1 \\ \epsilon^2 + \rho^2\sigma_k^2\left(1 - r(\psi_k)^2 - \psi r(\psi_k)\right) & \text{If } \; y_k = 1 \end{cases}$$

$$:= \sigma'^2(\mu_k, \sigma_k^2, \psi_k, y_k)$$

*Where $\psi_k = (\hat{x}_k - \mu_k)/\sigma_k$ is the normalized threshold and $h(\cdot)$ and $r(\cdot)$ are the hazard and reverse hazard rates:*

$$h(\psi) = \frac{\phi(\psi)}{1 - \Phi(\psi)} \qquad\qquad r(\psi) = \frac{\phi(\psi)}{\Phi(\psi)}$$

The threshold filter is exact up to the second moment if the true prior is normal.

**Proposition 4.** *If $x_{k|k+1} \sim N(\mu_k, \sigma_k^2)$, then:*

$$\mathbb{E}(x_{k+1}|y_k) = \mu'(\mu_k, \sigma_k^2, \psi_k, y_k) \qquad\qquad Var(x_{k+1}|y_k) = \sigma'^2(\mu_k, \sigma_k^2, \psi_k, y_k)$$

*Where expectations are taken using Bayes' rule.*

*Proof.* See appendix A □

An interesting consequence of Proposition 4 is that it means that the integration of quadratic functions is exact under the threshold filter.

**Corollary.** *Let $x_k \sim N(\mu, \sigma^2)$ and $V(x) = a + bx + cx^2$. Then if $F(\cdot)$ and $\tilde{F}(\cdot)$ are the exact* *and threshold filter approximation of the one-step ahead distribution of* $x_{k+1}$ *conditional on* $y_k$*:*

$$\int_{-\infty}^{\infty} V(x)dF(x) = \int_{-\infty}^{\infty} V(x)d\tilde{F}(x)$$

*Proof.* Follows from the additive separability of integration and Proposition 4 □

This property is particularly relevant in economic models where agents' decisions are influenced by expectations of the next period's outcomes. If the continuation value function is near-quadratic and then the errors from integrating over beliefs about next period's state are minimal.

## Behavior of the threshold filter

The behavior of $\mu'(\cdot)$ and $\sigma'^2(\cdot)$ is summarized in Proposition 5.

**Proposition 5** (Limiting behavior of the threshold filter)**.** *For all* $(\mu, \sigma^2, \psi)$*:*

$$\mu'(\mu, \sigma^2, \psi, 1) = -\mu'(\mu, \sigma^2, -\psi, 0) \qquad \sigma'^2(\mu, \sigma^2, \psi, 1) = \sigma'^2(\mu, \sigma^2, -\psi, 0)$$

*And the limits are given by:*

$$\begin{aligned} \psi \to \infty: \quad & \mu'(\mu, \sigma^2, \psi, 1) \sim \rho\hat{x} + O(\psi^{-1}) \quad & \sigma'^2(\mu, \sigma^2, \psi, 1) \sim \epsilon^2 + \left(\frac{\rho\sigma}{\psi}\right)^2 + O(\psi^{-4}) \\ \psi \to -\infty: \quad & \mu'(\mu, \sigma^2, \psi, 1) \sim \rho\mu & \sigma'^2(\mu, \sigma^2, \psi, 1) \sim \epsilon^2 + \rho^2\sigma^2 \end{aligned}$$

*Proof.* The first part is immediate. But for the limits we need to use the asymptotic expansions for the hazard and reverse hazard rates:

$$\begin{aligned} \psi \to \infty: \qquad & h(\psi) = \psi + \psi^{-1} - 2\psi^{-3} + 10\psi^{-5} + O(\psi^{-7}) \\ \psi \to -\infty: \qquad & r(\psi) = -\psi - \psi^{-1} + 2\psi^{-3} - 10\psi^{-5} + O(x^{-7}) \end{aligned}$$

Substituting these into the definitions for $\mu'(\cdot)$ and $\sigma'^2(\cdot)$ gives the proof. $\square$

Figure 3.2 illustrates the bounds on the posterior parameters. As $\psi$ gets very small, the cutoff is low relative to the prior. As a result, an upper threshold signal is not informative, and so the updated distribution is almost exactly the uninformed one. The updated mean $\mu'_h \simeq \rho\mu$ and the variance $\sigma'^2_h \simeq \epsilon^2 + \rho^2\sigma^2$ are therefore very close to their uninformed values. This is illustrated by the blue lines in Figure 3.2 illustrates this when $(\mu, \sigma^2) = (0, 1)$, $\rho = 0.95$ and $\epsilon^2 = 1 - \rho^2$ (so the prior is the long-run variance of the AR(1) process).

Given such a low $\psi$, though, a lower threshold signal ($y = 0$) is very informative, and causes a large shift in the distribution of beliefs. The mean will fall almost to the threshold (multiplied by the persistence $\rho$), and the variance will shrink almost to the innovation variance $\epsilon^2$. The red lines in Figure 3.2 describe this case for $\psi$ small. In addition, the posterior variance falls, as the low signal is a rare event when $\psi$ is large and negative.

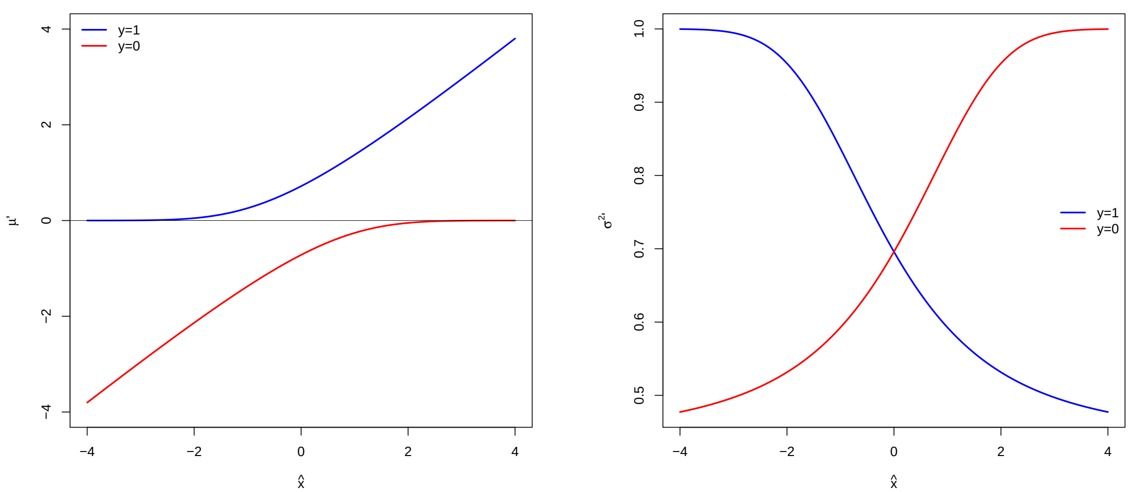


(a) Updated mean $\mu'$ (b) Updated variance $\sigma'^2$

Figure 3.2: Threshold filter when $\rho = .9$, $\sigma^2 = 1$, and $\epsilon = 1 - \rho^2$

The role of the asymptotic expansion of the hazard and reverse hazard rates in these results should be emphasized. At extreme values, the hazard and reverse hazard rates asymptote to the ratio of zero over zero. Simply applying L'Hôpital's rule to the hazard rate will give the correct limit for the mean (although convergence to this limit is very slow). But substituting this into the expression for the variance will incorrectly imply that $\lim_{\psi \to \infty} \sigma_h'^2 = \epsilon^2 + \rho^2\sigma^2$, instead of the true limit $\epsilon^2$.

Furthermore, the asymptotic expansion also matters for the numerical calculation of the hazard rate for extreme $\psi$. Computing $h(\cdot)$ as the ratio of the density and the upper cdf gives wildly inaccurate results for large $\psi$ because the denominator approaches zero.

If the threshold $\hat{x}_k$ is constant then we can also characterize the dynamic behavior of the threshold filter. The mean and variance converge to fixed levels at constant rates following sequences of constant signals. Figure 3.4 illustrates this with $\hat{x} = 0$. It shows how for sequences of $y = 1$, the mean of the threshold filter converges to a long-run level $\bar{\mu}_1 > 0$. And for extended sequences $y = 0$ it converges to $\bar{\mu}_0 < 0$. Proposition 6 expresses this formally. The proof is trivial.

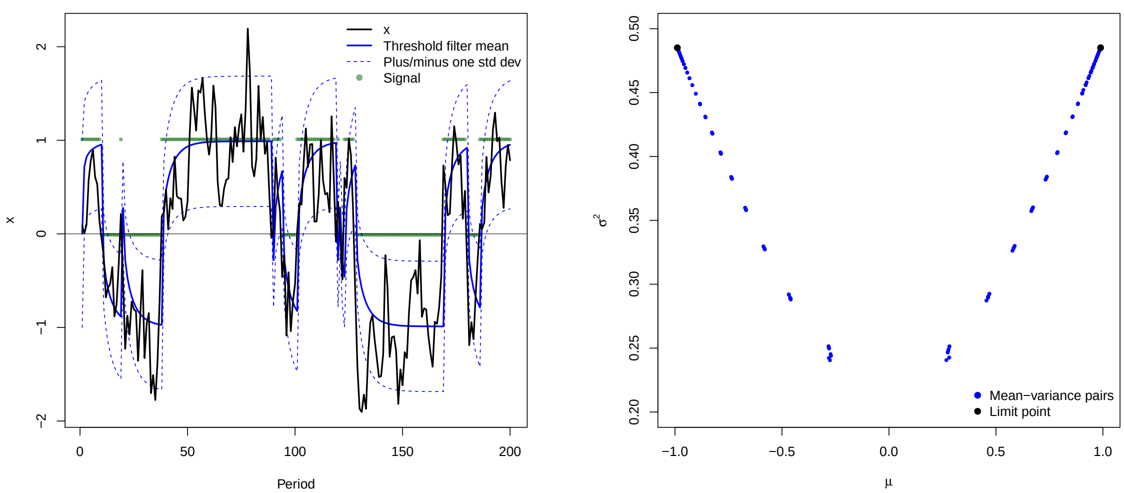


(a) Dynamic behavior

(b) Cross section

Figure 3.3: Threshold filter when $\hat{x} = 0$ always

**Proposition 6.** *When $\hat{x}$ is constant then during extended sequences of $y_k = 1$, $(\mu_k, \sigma_k) \to$*

| **Filter** | **Signal variance** | **Use on threshold problem** | **Notes** |
|---|---|---|---|
| Extended Kalman filter | Taylor series linearization | ✗ | Linearization not invertible |
| Statistically linearized filter | Min square error linearization | ✗ | Linearization implies $\sigma'^2 = 0$ |
| Unscented Kalman filter | Monomial approximation | ✓ | Assumes signal Gaussian |
| Exact Gaussian filter | Analytic | ✓ | Assumes signal Gaussian |

Table 3.1: Summary of alternative filtering methods

$(\bar{\mu}_1, \bar{\sigma}_1)$, *where:*

$$\bar{\mu}_1 = \frac{\rho\bar{\sigma}_1 h(\bar{\psi})}{1-\rho}$$
$$\bar{\sigma}_1^2 = \frac{\epsilon^2}{1-\rho^2\left(1-h(\bar{\psi}_1)^2+\bar{\psi}_1 h(\bar{\psi}_1)\right)}$$
$$\bar{\psi}_1 = (\hat{x}-\bar{\mu}_1)/\bar{\sigma}_1$$

*A similar result holds for sequences* $y_k = 0$

### *3.1.3 Alternative filters*

Several nonlinear filtering techniques already exist. Table 3.1 summarizes the main alternatives.

The alternative methods are all Kalman-type methods in that they use the Kalman updating formulae to translate the mean and variance of the signal and lagged state into a mean and variance of the next-period state. They differ in how they compute the variance of the signal. The Extended Kalman filter (Jazwinski (1966)) and the statistically linearized filter (Gelb (1974)) both use linear approximation of the signal combined with the delta method to compute this. This fails in both cases. In the case of the extended Kalman filter this is because local linearization via a Taylor approximation is non-invertible (as $g'(x) = 0$ almost everywhere). And the statistically linearized filter fails because it falsely interprets

$R = 0$ as implying that $\sigma'^2 = 0$. Accordingly, neither the extended Kalman filter nor the statistically linearized filter is appropriate for use in the threshold filtering problem.

More promising as the unscented Kalman filter and the exact Gaussian filter. These use respectively a monomial integration rule and analytical formulae to generate the variance of the signal. but both methods introduce error in assuming that the signal $g(x)$ is distributed normally. The unscented Kalman filter introduces further errors because it uses only approximate integration to evaluate $Var[g(x)]$, which are prone to inaccuracy when $g(\cdot)$ is very non-normal (as it is here).

The exact Gaussian filter is a Kalman-style filter that computes the variance of $g(\cdot)$ analytically. It is given by:

**Algorithm** (The exact Gaussian filter)**.** *If the distribution of the hidden state in period $k$ conditional on the signals $\{y_l\}_{l=0}^{k-1}$ is given by $x_{k|k-1} \sim N(m_k, s_k^2)$, then the exact Gaussian filter predicts an updated distribution in period $k+1$ of $x_{k+1|k} \sim N(m_{k+1}, s_{k+1}^2)$ where:*

$$
\begin{aligned}
m_{k+1} &= \begin{cases} \rho\left(m_k + s_k h(\psi_k)\right) & \text{If } \; y_k = 1 \\ \rho\left(m_k - s_k r(\psi_k)\right) & \text{If } \; y_k = 0 \end{cases} \\
&:= m'(m_k, s_k^2, \psi_k, y_k) \\
s_{k+1}^2 &= \epsilon^2 + \rho^2 s_k^2 \left(1 - r(\psi_k) h(\psi_k)\right) \\
&:= s'^2(m_k, s_k^2, \psi_k, y_k)
\end{aligned}
$$

*Where, as before, $\psi_k = (\hat{x}_k - m_k)/s_k$ is the normalized threshold and $h(\cdot)$ and $r(\cdot)$ are the hazard and reverse hazard rates*

The derivation of the exact Gaussian filter can be found in Appendix A. Note that though the exact Gaussian filter is derived using the assumption that $g(\cdot)$ is normal, becuase it is still essentially a moment-matching method it produces the same updating formula for the mean $m$ as the threshold filter. However, the variance updating formula is different. So not only does the exact Gaussian filter fail to satisfy Proposition 4, but the evolution of the mean for $k > 1$ will also differ from that implied by the threshold filter, as $m$ is a function of $s$. Furthermore, the evolution of the variance does not directly depend on the realization of the signal, a property inherited from the linear Kalman filter. This will potentially introduce sizable errors in the second moment, as the variance of $x'$ is a function of the realization of $y$.

We can actually describe the evolution of the exact Gaussian filter further, analogously to Proposition 6.

**Proposition 7.** *When $\hat{x}$ is constant then during extended sequences of $y_k = 1$, $(m_k, s_k) \to (\bar{m}_1, \bar{s}_1)$, where:*

$$\begin{aligned}
\bar{m}_1 &= \frac{\rho \bar{s}_1 h(\bar{\psi})}{1-\rho} \\
\bar{s}_1^2 &= \frac{\epsilon^2}{1-\rho^2\left(1-r(\bar{\psi}_1)h(\bar{\psi}_1)\right)} \\
\bar{\psi}_1 &= (\hat{x}-\bar{m})/\bar{s}
\end{aligned}$$

*A similar result holds for sequences $y_k = 0$*

### *3.1.4 Evaluating the filter*

## Measuring errors on distributions

The aim of this section is to provide a measure of the accuracy of the threshold and exact Gaussian filters compared to the true distribution of the hidden state. The measure of similarity that I use to compare distributions is the $L^0$ measure with respect to the true

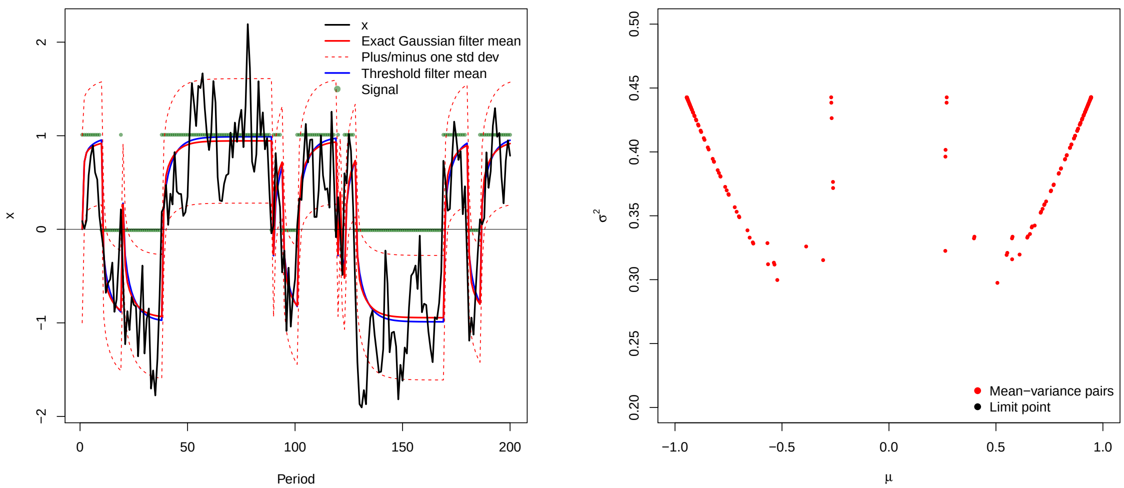


(a) Dynamic behavior (b) Cross section

Figure 3.4: Threshold filter when $\hat{x} = 0$ always

distribution. If $P(x_{t+1}|\{y_s\}_{s=0}^t)$ denotes the true distribution of the state conditional on the signals (i.e. computed using Bayes' law), and $Q(x_{t+1})$ is an approximating distribution, then this is:

$$d(P, Q) = \int_{x'} \left| P(x') - Q(x') \right| dP(x')$$

This measure has a number of appealing characteristics. First, it is easily interpretable. It is the average absolute error arising when using the approximate CDF to measure the probability of events generated by the true distribution. Second, it scale-independent. That is if $\epsilon = \epsilon_1$ and $\epsilon = \epsilon_2$ generate true distributions $P_1(x_{t+1}|\{y_s\}_{s=0}^t)$ and $P_2(x_{t+1}|\{y_s\}_{s=0}^t)$, and if $Q_1(x_{t+1}/\epsilon_1|\{y_s\}_{s=0}^t) = Q_2(x_{t+1}/\epsilon_2|\{y_s\}_{s=0}^t)$ then $d(P_1, Q_1) = d(P_2, Q_2)$[2]. This means that in any numerical assessment of these errors $\epsilon$ can simply be set to unity. Finally, this particular error measure is likely to be particularly important in economic problem with a continuum of types. In problems where a threshold filtering problem arises, agents typically need to form expectations over the *entire* distribution of future outturns. So a measure of

2. This follows from the fact that $P_1(x_{t+1}/\epsilon_1|\{y_s\}_{s=0}^t) = P_2(x_{t+1}/\epsilon_2|\{y_s\}_{s=0}^t)$ and the chain rule.

the error that is affected by difference in the CDF at all points is particularly appealing.

### Benchmark distribution

In order to measure the errors on the filters, we must have a benchmark distribution to compare them too. Ideally, this would be the exact distribution computed from applying Bayes' rule. But this is not easy; the whole reason for using an approximate filter is that the exact distribution is not an easy object to work with. So I use two benchmark distributions, against which I compare the distributions implied by the threshold and Gaussian filters: a very high-dimension discrete Markov approximation, and the exact one-step distribution.

The discretized Markov approximation to the true distribution at period $k$ is given by a vector $\mathbf{m_k}$ of probabilities on a grid of points $\mathbf{b}$. If $y_k = 1$, then $\mathbf{m'_{k+1}} = \mathbf{\tilde{m}'_k T}$ where $\mathbf{\tilde{m}_k}$ is a unit-sum vector which is zero for elements of $\mathbf{b}$ less than $\hat{x}$ and proportional to $\mathbf{m_k}$ otherwise, and $\mathbf{T}$ is the transition matrix for a discretized AR(1) process.

By approximating the true distribution with a very high-dimensional method, the discrete Markov approximation measures the propagation of the approximation errors over many periods, so is the most general method. This measure is therefore the "gold standard" for assessing the accuracy of the threshold filter. However, even with a large number of states the discretization will be subject to some error. And so as an extra check I also compute a secondary measure of the error.

The exact one-step distribution describes the true distribution of the hidden state if it *were* distributed normally in the preceding period. Given a normal prior with mean and variance $(\mu, \sigma^2)$, we can apply Bayes' rule to compute the exact one-step posterior. For

$y = 1$, this is:

$$\begin{aligned} p_1(x'|\mu,\sigma^2,\hat{x}) &= \int_{\hat{x}}^{\infty} \frac{1}{\epsilon\sqrt{2\pi}} e^{-\frac{1}{2}\frac{(x'-\rho x)^2}{\epsilon}} \frac{1}{\sigma\sqrt{2\pi}} \frac{1}{\Phi\left(\frac{\mu-\hat{x}}{\sigma}\right)} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2} dx \\ &= \frac{1}{\epsilon\sqrt{2\pi}} e^{-\frac{1}{2}\frac{(\rho\mu-x')^2}{(\rho^2\sigma^2+\epsilon^2)}} \int_{\hat{x}}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \frac{1}{\Phi\left(\frac{\mu-\hat{x}}{\sigma}\right)} e^{-\frac{1}{2}\frac{\rho^2\sigma^2+\epsilon^2}{\sigma^2\epsilon^2}\left(x-\frac{\sigma^2\mu+\sigma^2\rho x'}{(\rho^2\sigma^2+\epsilon^2)}\right)^2} dx \\ &= \frac{\Phi\left(\frac{\rho\sigma^2 x' - (\gamma^2\hat{x}-\epsilon^2\mu)}{\sigma\epsilon\gamma}\right)}{\Phi\left(\frac{\mu-\hat{x}}{\sigma}\right)} \phi\left(\frac{x'-\rho\mu}{\gamma}\right) \end{aligned}$$

Where $\gamma^2 = \rho^2\sigma^2 + \epsilon^2$ is the uninformed variance of the posterior (i.e. that which would arise if there were no signal about the type). Likewise, the distribution following a lower threshold signal is:

$$p_0(x'|\mu,\sigma^2,\hat{x}) = \frac{\Phi\left(\frac{(\gamma^2\hat{x}-\epsilon^2\mu)-\rho\sigma^2 x'}{\sigma\epsilon\gamma}\right)}{\Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)} \phi\left(\frac{x'-\rho\mu}{\gamma}\right)$$

The comparison to the exact one-step distribution measures exactly the errors introduced at the first step of the filter. It makes no attempt to measure whether small errors propagate when the approximation is used in multiple periods but suffers from no inaccuracy in its own right.

## Results

I compute summary statistics for the errors in the filters by the following algorithm:

**Algorithm** (Measuring filter errors)**.**

- *Fix* $\rho$*, and set* $\epsilon = 1 - \rho^2$*.*
- *Generate a simulation of 20000 points for the underlying state* $x_t$*, drawing thresholds*

$\hat{x}_t$ *from* $N(0,1)$ *(which is the ergodic distribution for* $x_t$*), and produce signals* $y_t$.

- *Calculate the mean and variance for the threshold and Gaussian filters for each* $t$.
- *At each* $t$ *compute the one-step distribution assuming that the threshold and Gaussian filters produce the true priors.*
- *Calculate a 600 state discrete Markov approximation to the true state using the simulated signals.*
- *For each* $t$ *compute the* $L^0$ *distance between the benchmark distributions and those implied by the filters.*
- *Report the mean and maximum error for each simulation.*

Figures 3.6 and 3.5 show the errors on the two filters compared to the two benchmarks. The first point to note is that the overall magnitude of the errors is smaller for the discrete-state Markov approximation. This suggests that approximation errors do not propagate, but in fact may dissipate over time. As $\rho$ gets larger, however, the errors can be quite large, particularly for the one-step-ahead case. In the worst case almost as much 50pp from $\rho = 0.95$, although mean errors are around a fifth as large.

The threshold filter has notably superior performance when compared to the discrete-state Markov approximation. When $\rho = 0.6$ the average error is less than 0.2pp, and even in the worst case is never more than 0.5pp. In contrast, the maximum error on the exact Gausian filter is nearly 3pp for the same value of $\rho$. For all $\rho$, the mean error on the threshold filter is approximately half that of the Gaussian distribution. However, the worst-case error often many time smaller.

The threshold and Gaussian filters have very similar one-step performance. This is to be expected, as the formula for the means are the same across the two filters, so one-step-ahead filter predictions will have the same mean (although the variance will differ). But over longer simulations, the inaccuracies in the Gaussian filter do not dissipate as fast.

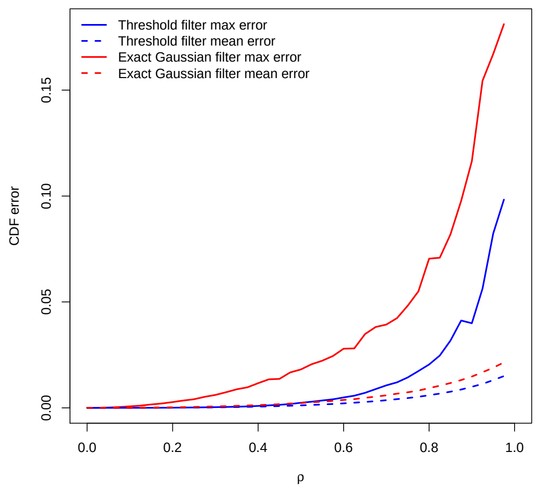


Figure 3.5: Filter error vs. discrete-state Markov distribution

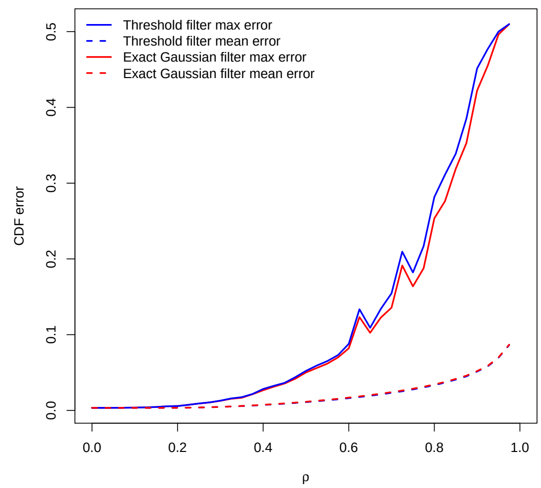


Figure 3.6: Filter error vs. one-step exact distribution

### *3.1.5 Conclusions on the threshold filter*

In this section I presented the threshold filtering problem. This arises in any economic model where there is a continuum of hidden types and a threshold defining the cutoff point between two binary actions (which generates a signal). Solving such a model requires a description of the distribution of the hidden type conditional on the signals. Parsimony is key in this description, because each extra parameter of the distribution becomes an extra dimension of the problem's state space.

To solve this problem, I define the threshold filter, and showed that it is accurate up to the second moment if the prior is correct. I also derived the dynamic behavior when the threshold is fixed. I discussed other alternative filters and compare the threshold filter to the best of these - the exact Gaussian filter. I then showed that in a long simulation with random thresholds, the threshold filter has very small errors relative to a high-dimension discrete approximation of the underlying state when persistence of the underlying state is not too high; the average error on the cdf is never more than around half of one percentage point. The Gaussian filter also behaves impressively well, but for moderate state persistence, the worst case errors in a long simulation can be up to six times as large as those produced by the threshold filter.

## 3.2 Using the threshold filter to solve the model

I now calibrate and solve the model using the threshold filter to track the evolution of investors' beliefs about default probabilities. I separately calibrate the symmetric and asymmetric models to hit the same targets.

### *3.2.1 Calibration*

#### Surplus process

I calibrate the surplus process in the model from the WEO data on annual surpluses from a sample of eight Latin American countries for which long-run data is available. There I find that surpluses are typically small and positive on average (1 or 2 % of GDP), fluctuate by around 2% of GDP, and are moderately persistent. See appendix C for further details of the calibration of the surplus process. Informed by this, I assume that surpluses follow a 3-state Markov process with mean 0.015, standard deviation 0.02 and persistence 0.56. This gives:

$$s = \begin{bmatrix} -0.02 \\ 0.02 \\ 0.06 \end{bmatrix} \qquad M = \begin{bmatrix} 0.55 & 0.40 & 0.05 \\ 0.15 & 0.70 & 0.15 \\ 0.05 & 0.40 & 0.55 \end{bmatrix}$$

#### Other parameters

I choose the other parameters to hit a variety of unrelated empirical targets. I pick $g = 0.04$, as this is the average rate of growth in the middle-income countries from 1980 to 2010. I set $R = 1.05$ to give an annual risk-free rate of return of 5%, and $\beta = 0.95$, consistent with the risk-free rate.

The remaining parameters, $(\psi, k, \rho, \epsilon)$ are chosen to match other, unrelated moments of the data. In particular, I target a default rate of 2%, which is in line with historical source on long-run default rates (see Tomz and Wright (2013) for a survey). For loss-given default, I use the results of Benjamin and Wright (2013), who show that the average market loss-given-default is 40%. Here, loss-given-default is calculated as the realized writedown of a bondholder who holds a bond through a default event, discounted at the risk-free rate. For example, imagine that a government defaults in period $t + 1$ on debts issued in $t$, and then defaults again on the newly issued debt before repaying in period $t + 3$. Then the realized loss-given-default is $\psi^2 q_{t+3}$. Instead of receiving a payoff of 1 in period $t + 1$, the investor has their claim written down and twice and exchanged for new debt.

The target for the average debt service to GDP ratio at default (8.4%) are taken from a subset of the events used in the panel regressions in section 2.1 for which there are consistent measures of debt service available. See Table B.2 in Appendix B for further details. As is common in the literature, I aim to match debt service rather than the total debt burden, as the model features only short-term debt. Total debt levels are very hard to match without longer-maturity bonds. And I choose to target debt service at default (and not, say, average debt service) because when default is rare (as it is here) the main driver of debt dynamics will be the process for surpluses, which is already set. So choices of $(\psi, k, \rho, \epsilon)$ will have little effect on the path of debts in general. But they will have much more of an impact on the default boundaries and so debt at default. Moreover, given that this paper is about spreads after default it seems important to get debt levels correct at the start of default, rather averaged over predominantly non-default times.

The parameters I choose and their success in hitting the targets are shown in Tables 3.2 and 3.3. The choices of $(\psi, k, \rho, \epsilon)$ do well in matching the target moments in the symmetric information version of the model. The haircut-exclusion correlation ratio is a little higher than in the data, but the calibrate broadly matches the data, even with a much higher choice of $\beta$ than most other models use.

2. Not a target: the variance of $\theta$ is held fixed across the two calibrations instead

| Parameter | Symmetric info | Asymmetric info |
|---|---|---|
| $g$ | 0.04 | 0.04 |
| $R$ | 1.05 | 1.05 |
| $\beta$ | 0.95 | 0.95 |
| $\epsilon$ | 0.10 | 0.11 |
| $\rho$ | 0.85 | 0.80 |
| $\psi$ | 0.75 | 0.70 |
| $k$ | 0.45 | 0.35 |

Table 3.2: Calibration parameters

| Target moment | Target Value | Sym info | Asym info |
|---|---|---|---|
| Ave default rate | 2% | 1.6% | 1.0% |
| Ave market loss-given-default | 40% | 40% | 42% |
| Ave debt service at default | 8.4% | 10.2% | 11.0% |
| Haircut-exclusion correlation | 0.62 | 0.87 | 0.94 |

Table 3.3: Calibration targets for $\psi, k, \rho, \epsilon$ (see text for targets for $R, \beta, g$)

To calibrate the asymmetric model, I fix $\epsilon^2/(1-\rho^2)$ at the value in the symmetric calibration, so that the variance of $\theta$ is fixed. Changes in $\theta$ therefore have the same units across the two versions of the model. To impose this restriction, I drop the average loss-given-default from the list of targets. This calibration is shown in the third column of Table 3.2 and the fourth column of Table 3.3.

### *3.2.2 Default boundaries*

Figure 3.7 shows the default thresholds $\hat{b}(\theta, s)$ for the symmetric information model. The three solid lines show the default thresholds for each surplus level. For a given surplus, a higher private cost $\theta$ increases the debt burden that can be sustained. The default boundaries reflect the calibration target for debt levels at default, ranging from about 7% of output to nearly 25%.

The dashed lines show the maximum debt that could be repaid in a given state, $s + Rev_{max}(s, \theta)$. When $\theta$ is large, the dashed and solid lines overlap. Governments with high private cost always repay if they can, only defaulting when the net present value of future surpluses is greater than their total debt. For intermediate $\theta$, the dashed and solid lines overlap only when $s$ is large, implying that there is strategic default at low $s$. When $s$ is low, repayment is less valuable as future default is likely, causing very low debt prices. And so the government must issue much more debt to pay outstanding creditors, reducing the value of repayment.

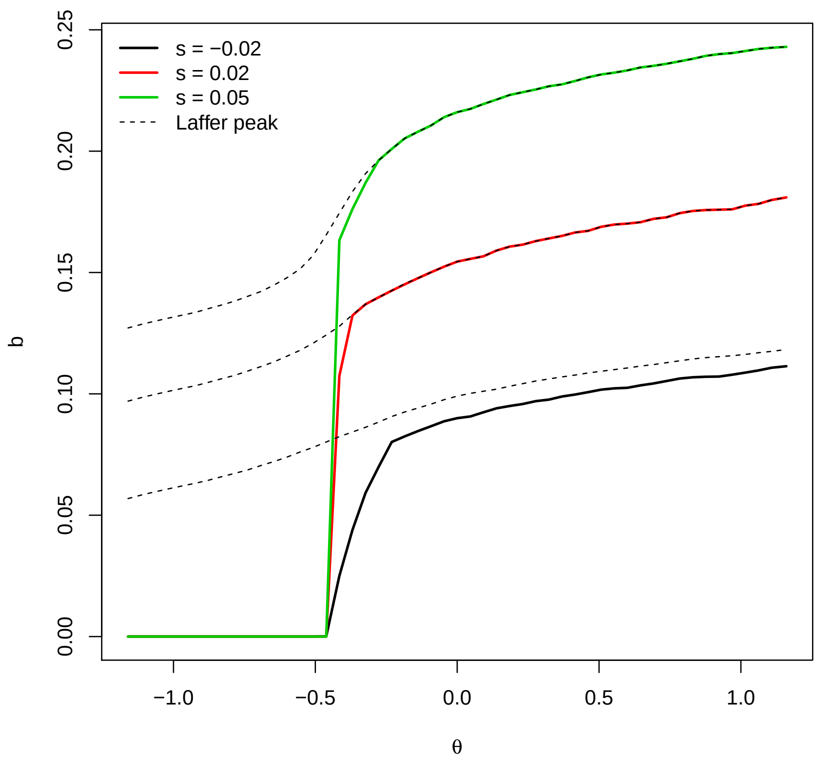


Figure 3.7: Default thresholds for symmetric information model.

When information is asymmetric, the default boundaries increase. Default signals likely future default, and so pushes up subsequent interest rates. This acts as a deterrent to default. Figure 3.8 shows the default boundaries for the asymmetric information model. Each line represents a value of $s$ and $\mu$ (the values of $\mu$ range from -0.8 to 0.8, covering around six times their simulation standard deviation). Higher $\mu$ produces higher default boundaries, as debt prices are higher because investors expect lower default probabilities; $\sigma^2$ is fixed[3] at

3. This is also the most common value in the simulation, and changing $\sigma^2$ makes little difference to the

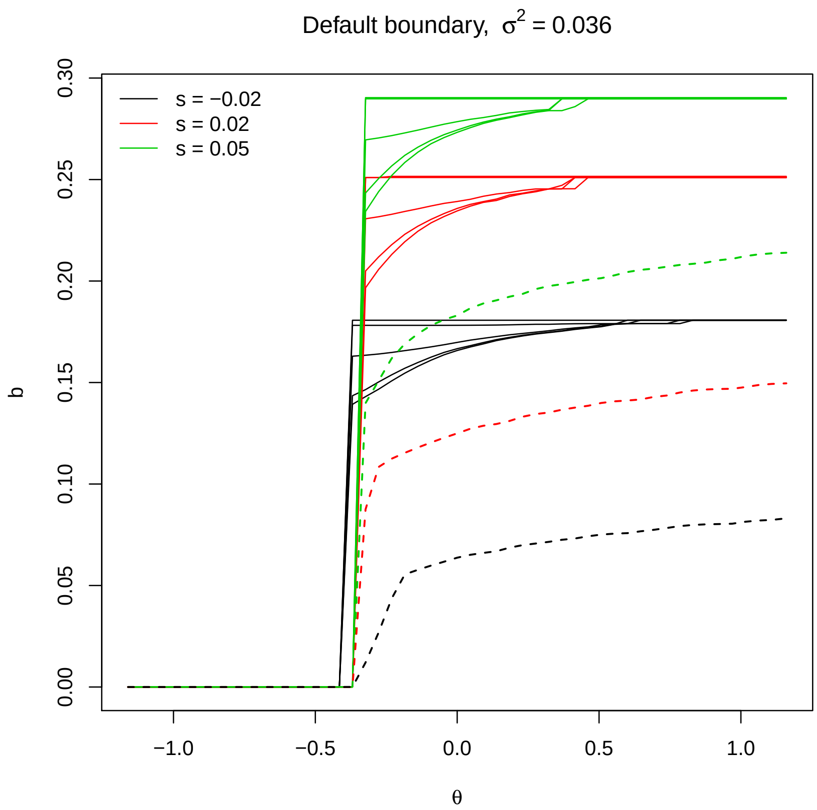


Figure 3.8: Default thresholds for the asymmetric information model (dashed lines symmetric boundaries from same calibration).

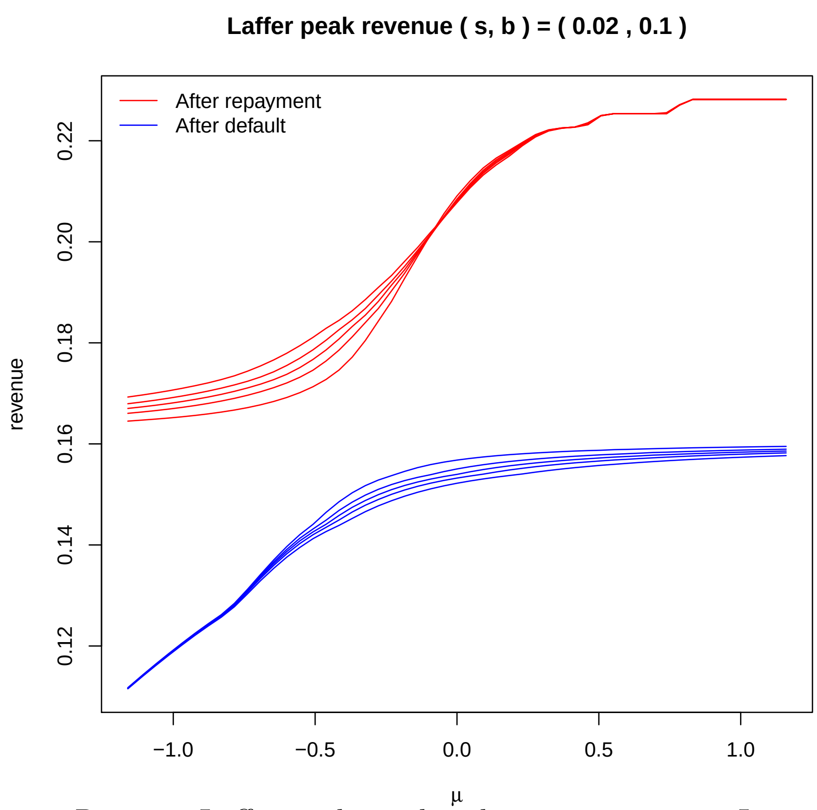


Figure 3.9: Revenue Laffer peaks under alternative actions. Lines vary by $\sigma$.

$\sigma^2_{max} = \epsilon^2/(1-\rho^2)$.

The broken lines in Figure 3.8 show the default boundaries for the symmetric information model under the same calibration. These are much lower than the same-calibration asymmetric information boundaries. For example, the default boundary for the lowest level of $s$ (in black) is nearly three times as high under asymmetric information as when the government's type is known. The reason is that in the asymmetric information model, default now conveys (negative) information about the government's type, pushing up future interest rates.

This effect is very strong. The flattest parts of the default boundaries are almost exactly what would arise if the government where able to commit to repayment whenever they were solvent. In other words, the threat of lost reputation can come close to replicating the commitment solution over much of the state space.

Figure 3.9 illustrates this mechanism. The vertical axis shows for a fixed level of surpluses and debt the maximum revenue that a government can raise immediately following default and repayment. The horizontal axis displays investors' prior mean belief about type, $\mu$. Each red line indicates the maximum revenue which can be repaid for a fixed value of the prior variance $\sigma$. Note that the red lines all intersect at the same value of $\mu$, around $\mu = -0.1$. Smaller $\sigma$ corresponds to higher revenue when $\mu$ is large, and lower revenue when $\mu$ is high. This is intuitive; certainty (low $\sigma$) improves the terms on which government can borrow if their reputation is generally good ($\mu$ is high) and worsens them if their reputation is generally bad ($\mu$ is low).

Note also that for high $\mu$, the difference between the Laffer peak under default and repayment is large, but as $\mu$ falls this difference shrinks. This is because if $\mu >> \hat{\theta} \simeq -0.4$, repayment is expected, and so default is highly informative about type. Default causes a large, discrete fall in the mean belief: $\mu'_d \simeq \hat{\theta}$. Repayment, on the other hand, is almost entirely uninformative. The expected action occurred, so beliefs are mostly dominated by

default boundaries.

the underlying process for $\theta$; $\mu_r' \simeq \rho\mu$. As this is large if $\mu$ is large, then borrowing occurs at almost the risk-free rate.

When $\mu$ is lower, neither repayment nor default are terribly informative. This is particularly true when $\mu \simeq \hat{\theta}$, because then the default-repayment choice simply shifts mean beliefs a little either side of $\hat{\theta}$. This is reflecting in the narrowing of distance between the Laffer peaks near $\mu \simeq \hat{\theta} \simeq -0.4$. Because beliefs following repayment and default are not very different, loss of reputation is less of a deterrent to default when $\mu \simeq \hat{\theta}$.

The changing difference between these curves also helps to explain why the asymmetric information model can generate a conditional default premium. Defaults tend to occur when $\mu \simeq \hat{\theta}$, simply because $\theta$ is usually a little smaller than $\hat{\theta}$ at default. Here, the difference between the Laffer peaks under repayment and default is relatively small. And so the loss from default relative to repayment is lower than at other times. This pushes up on default rates and spreads following a default.

### *3.2.3 Comparing the model to post-default spreads*

I now show that the model can replicate the stylized fact in the data - that spreads are high following a default - but only if information is asymmetric. The intuition is that when information is symmetric, a compositional effect is at work. Those governments that repay for multiple periods following a default are typically better-than-average types. And so spreads after a default are low, which is at odds with the data. When information is asymmetric, a credibility force is also at work, pushing up on spreads. Default signals that the government is a poor type. This loss of credibility acts as a deterrent to default. As result, defaulting a second time would cause little deterioration in subsequent beliefs about type, mitigating the deterrent effect and increasing post-default spreads.

### Model regression

To compare the model to these empirical targets, I run the analogous regression on a simulation of 100,000 periods. The regressions are:

$$y_t = \alpha + \gamma s_t + \alpha_1 b_t + \alpha_2 b_t^2 + \beta' v_t \tag{3.1}$$

Where:

- $y_t = 1/\sqrt{q_{t,2}} - R$ is the annual spread on a synthetic 2-period bond;
- $\alpha$ is a constant corresponding to the average spread in the periods more than $n = 15$ periods before or after a default;
- $s_t, b_t$ are the primary surplus and start-of-period debt;
- $v_t$ is a length $n$ vector of indicators where element $i$ is one if period $t$ is i periods after the end of a default event (recall that default event is a sequence of one or more successive periods of default);

The coefficients of interest are the elements of $\beta$. Figure 3.10 overlays the empirical estimates with the results from the two versions of the model. The solid black line shows the results from the asymmetric information model, and the broken one those from the symmetric version.

The asymmetric information model does a much better job of matching the stylized fact. Under symmetric information, spreads are counterfactually low following the conclusion of a default. In contrast, the asymmetric information model features a persistent increase in the cost of borrowing. The reasons for these patterns are subtle, and require a little explanation.

There are two effects at work: a composition effect and a reputation effect. The composition effect is at work in both models. As the x-axis variable is periods since *last* default, then observations $n$ periods after default are conditioned on repayment in the preceding $n - 1$

periods. This typically selects for particularly high-type governments. For example, the sample of observations three periods after the last default contains those types which have repaid for three consecutive periods following a default. Given that debt is typically high and surpluses often negative, only the best government types repay and therefore continue in the sample. This causes post-default spreads in the symmetric information model to be *lower* than would be predicted by surpluses and debt outstanding alone.

The asymmetric information model is subject to this effect too, but it is mitigated by a credibility effect. Following default, the belief about type falls sharply. For a given default rule, this would reduce the price of debt. In equilibrium, though, this results in higher default thresholds. As discussed in section 3.2, better beliefs makes borrowing cheaper and repayment more attractive. As a result, when beliefs about type decline following a default, the spreads on debt rise to reflect the increased default risk.

### *3.2.4 Comparison to models with discrete types*

A natural alternative to my model is one with a discrete number of types. This section explains why producing persistently elevated spreads after a default requires as state space at least as big as that in my model. In fact, any discrete-type a model that is more parsimonious than mine is guaranteed to generate counterfactual spread behavior.

Consider first the simplest possible cases where there are just two government types[4]. Following a default there are two relevant possibilities. First, that both types repay. In which case there is no default risk; the government can borrow at the risk-free rate. Given that debt is typically high and surpluses low after a default, then this would lead to abnormally *low* spreads after default, in contradiction to the data. The second case is that one type repays and the other defaults. If this happens, spreads will be high in the first period after default, but becuase actions are fully revealing of type in this situation, repayment signals with certainty that the type is the very best it can be. In effect, the strength of the compositional

4. This is the set-up in both Alfaro and Kanczuk (2005) and D'Erasmo (2008)

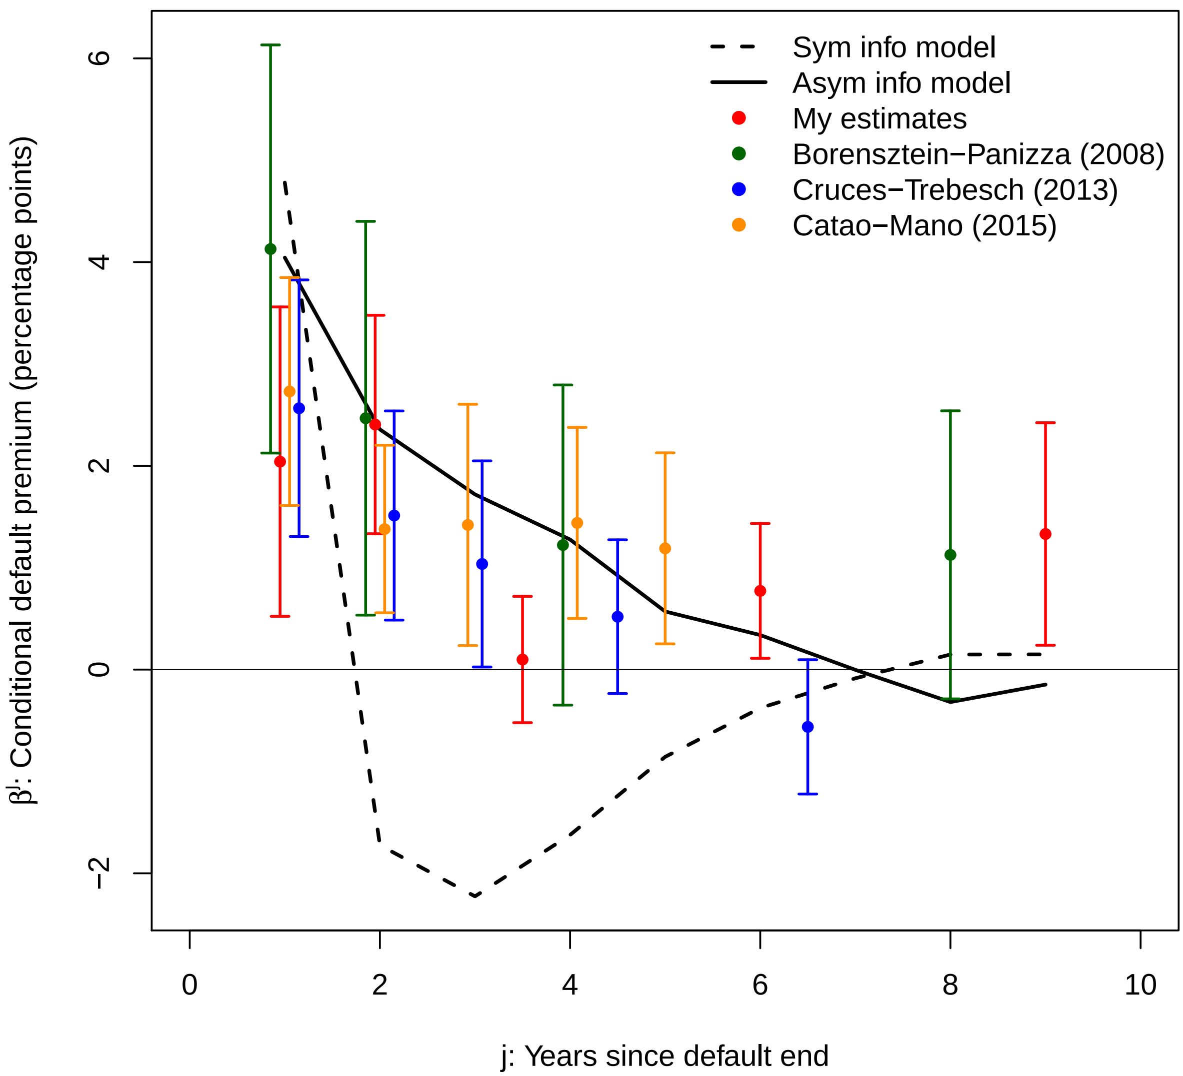


Figure 3.10: Spreads from the panel regressions and symmetric information model.

effect would leave no room for the reputational effect to operate. Accordingly, spreads more than one period after default would be abnormally low if type is persistent[5].

Even expanding the model to three types is unlikely to help much, as it would generate persistently higher spreads only when two types repaid and one defaulted. Furthermore, a probability distribution over three types is a two-dimensional object. And so the model would be no simpler than the second-order approximation that I use. Introducing more discrete types may eventually fix this, but only at the cost of a greatly expanded state space.

5. There is a third case, of course, which is that both types default. This is not relevant as then the model would not produce repayment sequences of more than one period.

# REFERENCES


Alfaro, L. and F. Kanczuk (2005). Sovereign debt as a contingent claim: a quantitative approach. *Journal of International Economics 65*(2), 297–314.

Amador, M. (2003). A political economy model of sovereign debt repayment. *Manuscript, Stanford University*.

Arellano, C. (2008). Default risk and income fluctuations in emerging economies. *The American Economic Review*, 690–712.

Arellano, C. and A. Ramanarayanan (2012). Default and the maturity structure in sovereign bonds. *Journal of Political Economy 120*(2), 187–232.

Benjamin, D. and M. L. Wright (2013). Recovery before redemption: A theory of delays in sovereign debt renegotiations. *Unpublished*.

Borensztein, E. and U. Panizza (2009). The costs of sovereign default. *IMF Staff Papers 56*(4), 683–741.

Bulow, J. and K. Rogoff (1989). Sovereign debt: Is to forgive to forget? *The American Economic Review*, 43–50.

Catão, M. L. and R. Mano (2015). Default premium. *IMF Working Paper 15*(167).

Chatterjee, S. and B. Eyigungor (2012). Maturity, indebtedness, and default risk. *The American Economic Review 102*(6), 2674–2699.

Cole, H. L., J. Dow, and W. B. English (1995). Default, settlement, and signalling: lending resumption in a reputational model of sovereign debt. *International Economic Review*, 365–385.

Cole, H. L. and T. J. Kehoe (2000). Self-fulfilling debt crises. *The Review of Economic Studies 67*(1), 91–116.

Cruces, J. J. and C. Trebesch (2013). Sovereign defaults: The price of haircuts. *American Economic Journal: Macroeconomics 5*(3), 85–117.

Cuadra, G. and H. Sapriza (2008). Sovereign default, interest rates and political uncertainty in emerging markets. *Journal of International Economics 76*(1), 78–88.

D'Erasmo, P. (2008). Government reputation and debt repayment in emerging economies. *Manuscript, University of Texas at Austin 12*, 17–28.

Dias, D. A., C. J. Richmond, and M. L. Wright (2011). The stock of external sovereign debt: Can we take the data at face value? *Working paper*.

Drazen, A. (2004). *Political economy in macro economics.* Orient Blackswan.

Gelb, A. (1974). *Applied optimal estimation.* MIT press.

Hatchondo, J. C., L. Martinez, and H. Sapriza (2009). Heterogeneous borrowers in quantitative models of sovereign default*. *International Economic Review 50*(4), 1129–1151.

Ito, K. and K. Xiong (2000). Gaussian filters for nonlinear filtering problems. *Automatic Control, IEEE Transactions on 45*(5), 910–927.

Jazwinski, A. (1966). Filtering for nonlinear dynamical systems. *Automatic Control, IEEE Transactions on 11*(4), 765–766.

Jazwinski, A. H. (1970). *Stochastic processes and filtering theory*. Courier Corporation.

Julier, S. J. and J. K. Uhlmann (1997). New extension of the kalman filter to nonlinear systems. In *AeroSense'97*, pp. 182–193. International Society for Optics and Photonics.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of basic Engineering 82*(1), 35–45.

Kalman, R. E. and R. S. Bucy (1961). New results in linear filtering and prediction theory. *Journal of basic engineering 83*(1), 95–108.

Lorenzoni, G. and I. Werning (2014). Slow moving debt crises.

Onder, Y. K. (2013). *Essays on sovereign debt*. Ph. D. thesis, Georgetown University.

Orjebin, E. (2014). A recursive formula for the moments of a truncated univariate normal distribution. *Manuscript, The University of Queensland*.

Phelan, C. (2006). Public trust and government betrayal. *Journal of Economic Theory 130*(1), 27–43.

Sandleris, G. (2008). Sovereign defaults: Information, investment and credit. *Journal of International Economics 76*(2), 267–275.

Särkkä, S. (2013). *Bayesian filtering and smoothing*, Volume 3. Cambridge University Press.

Tomz, M. (2007). *Reputation and international cooperation: sovereign debt across three centuries*. Princeton University Press.

Tomz, M. and M. L. Wright (2007). Do countries default in bad times? *Journal of the European Economic Association 5*(2-3), 352–360.

Tomz, M. and M. L. Wright (2013). Empirical research on sovereign debt and default. *Annual Review of Economics 5*(1), 247–272.

Wan, E. A. and R. Van Der Merwe (2000). The unscented kalman filter for nonlinear estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000*, pp. 153–158. Ieee.

Wright, M. L. (2002). Reputations and sovereign debt. *Manuscript, Stanford University*.

# APPENDIX A

# PROOFS

## A.1 The sovereign default model

Proof of proposition 3:

*Proof.* Existence of Markovian equilibria is straightforward. We can apply backward induction to a version of the game truncated after $T$ periods and show that a Markovian equilibrium exists there. Letting $T \to \infty$ generates payoffs arbitrarily close to those in infinitely repeated game. The main challenge is in establishing the conditions under which we can guarantee that such an equilibrium can be characterized by a threshold rule.

First, we establish some bounds on $V_1(b'|\theta, s, p')$, the derivative of $V(b'|\theta, s', p')$ with respect to $\theta$. Fix the predetermined type $\theta_0 = \theta$ and prior surplus $s_0 = s$, as well as initial debt $b_1 = b'$ and beliefs $p_1 = p'$. Let $x_t(\theta, \{\epsilon_\tau, s_\tau\}_{\tau=1}^t)$ be the period $t$ default rule for a government starting in period 0 with initial type $\theta$ and facing subsequent exogenous shocks $\{\epsilon_\tau, s_\tau\}_{\tau=1}^\infty$.

Now consider a government with initial type $\tilde{\theta} \neq \theta$. Then following the same default rule is feasible for them, becuase at each point in time it will produce the same beliefs, prices and debts as for the government with initial type $\theta$. Therefore denote by $V_\theta(b'|\tilde{\theta}, s, p')$ the value to type $\tilde{\theta}$ of implementing the default rule $x_t(\theta, \{\epsilon_\tau, s_\tau\}_{\tau=1}^t)$ for all future $t$. Then because this is feasible, it must be weakly worse than the value from following the optimal policy. That is:

$$V(b'|\tilde{\theta}, s, p') \geq V_\theta(b'|\tilde{\theta}, s, p') \tag{A.1}$$

Becuase the values $V_\theta(b'|\tilde{\theta}, s, p')$ and $V(b'|\theta, s', p')$ differ only in their valuation of expected future default, we have that:

$$V_\theta(b'|\tilde{\theta}, s, p') = V(b'|\theta, s, p') - \sum_{t=0}^{\infty} \beta^t z_t(\theta, b', s, p') \rho^{t+1} (\tilde{\theta} - \theta) \tag{A.2}$$

Where $z_t(b', s, p')$ is the probability of default in period $t$ under the default rule $x_t(\theta, \{\epsilon_\tau, s_\tau\}_{\tau=1}^t)$ (so $z_t$ is formed by integrating $x_t$ over realizations of $\{\epsilon_\tau, s_\tau\}_{\tau=1}^t$). We can use the fact that $z_t \in [0, 1]$ to bound the derivative of the value function with respect to $\theta$. First, when $\tilde{\theta} = \theta + \epsilon$ for $\epsilon > 0$. In this case the lower bound on $V_\theta(b'|\tilde{\theta}, s, p')$ is attained when $z_t = 1 \; \forall \; t$. Combining (A.1) with (A.2) gives:

$$V(b'|\theta + \epsilon, s, p') \geq V(b'|\theta, s, p') - \frac{\epsilon\rho}{1 - \beta\rho}$$

Dividing by $\epsilon$ and taking the limit as $\epsilon \to 0$ gives a lower bound on the derivative:

$$V_1(b'|\theta, s, p') \geq \frac{-\rho}{1 - \beta\rho}$$

If $\tilde{\theta} = \theta - \epsilon$ then the lower bound occurs when $z_t = 0$. Rearranging and taking limits again gives:

$$V_1(b'|\theta, s, p') \leq 0$$

Which should have been obvious anyway - any type can always follow the same decision rule as that with higher costs and achieve a weakly superior outcome. We can now apply these bounds to the default decision equation. As shown in the main body of the paper, default occurs when

$$u_r(\theta', b', s', p') + \beta V\left(b^*\left(\max(0, b' - s'), s', p'_r\right)|\theta', s', p'_r\right) < \\ u_d(\theta', b', s') + \beta V\left(\psi b'/(1+g)|s', \theta', p'_d\right)$$

Substituting in for the period payoffs and rearranging we have that:

$$\beta\left[V\left(b^*\left(\max(0, b' - s'), s', p'_r\right)|\theta', s', p'_r\right) - V\left(\psi b'/(1+g)|s', \theta', p'_d\right)\right] < \min(0, s) - (k + \theta)$$

The derivative of the right hand side of this equation is -1. The bounds on $V_1(b'|\theta, s, p')$

imply that the derivative of the left hand side is bounded above by $-\beta\rho/(1-\beta\rho)$. We therefore have a threshold rule if:

$$\begin{aligned}\frac{\beta\rho}{1-\beta\rho} &< 1 \\ \Rightarrow \beta &< \frac{1}{2\rho}\end{aligned}$$

□

## A.2 The threshold filter

### *A.2.1 Proof of Proposition 4*

I prove the proposition only for the case where $y_k = 1$. When $y_k = 0$ the proof is identical. Let $M_n$ be the $n^{th}$ non-central moment of $x_{k+1}$. then:

$$\begin{aligned}M_n &= \mathbb{E}\left((x')^n \;\middle|\; \hat{x}, \mu, \sigma^2\right) \\ &= \mathbb{E}\left((\rho x + u)^n \;\middle|\; \hat{x}, \mu, \sigma^2\right) \\ &= \mathbb{E}\left(\sum_{i=0}^{n} \binom{n}{i} \rho^i x^i u^{n-i} \;\middle|\; \hat{x}, \mu, \sigma^2\right) \\ &= \sum_{i=0}^{n} \binom{n}{i} \rho^i \mathbb{E}\left(x^i \;\middle|\; \hat{x}, \mu, \sigma^2\right) \mathbb{E}u^{n-i}\end{aligned}$$

Where $u \sim N(0, \epsilon^2)$. Then the uncentered moments of $u$ are:

$$\mathbb{E}u^k = \begin{cases} 0 & k \text{ odd} \\ \epsilon^k (k-1)!! & k \text{ even} \end{cases}$$

Where $n!! = n(n-2)(n-4)\ldots$ is the semifactorial of $n$. Denote by $m_k$ the $k^{th}$ noncentral moment of $x'|\hat{x}, \mu, \sigma^2$. Then:

$$M_n = \sum_{i=0}^{\lfloor n/2 \rfloor} \binom{n}{2i} \rho^{n-2i} m_{n-2i} \epsilon^{2i} (2i-1)!! \tag{A.3}$$

Because $x|\hat{x}, \mu, \sigma^2$ is just a normal with mean $\mu$ and variance $\sigma^2$ left-truncated at $\hat{x}$, then the moments $m_k$ can be generated recursively via[1]:

$$\begin{aligned} m_{-1} &= 0 \\ m_0 &= 1 \\ m_k &= (k-1)\sigma^2 m_{k-2} + \mu m_{k-1} + \sigma \hat{x}^{k-1} h((\hat{x}-\mu)/\sigma) \end{aligned} \tag{A.4}$$

Where $h((\hat{x}-\mu)/\sigma)$ is the hazard rate of the normal distribution:

$$h((\hat{x}-\mu)/\sigma) = \frac{\phi\left(\frac{\hat{x}-\mu}{\sigma}\right)}{1 - \Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)}$$

In particular, the first and second moments of $x|\hat{x}, \mu, \sigma^2$ are:

$$\begin{aligned} m_1 &= \mu + \sigma h((\hat{x}-\mu)/\sigma) \\ m_2 &= \mu^2 + \sigma^2 + \sigma(\mu + \hat{x}) h((\hat{x}-\mu)/\sigma) \end{aligned}$$

Recall that the first two non-central moments of $x'|\hat{x}, \mu, \sigma^2$ are:

$$\begin{aligned} M_1 &= \rho m_1 \\ M_2 &= \epsilon^2 + \rho^2 m_2 \end{aligned}$$

Substituting in for $m_1$, $m_2$ gives the results.

1. See Orjebin (2014) for the derivation of these formulae.

### *A.2.2 Derivation of Algorithm 3.1.3*

The nonlinear Kalman gain in period $k$ is given by:

$$K_k = \text{Cov}[x, g(x)] \left(\text{Var}[g(x)] + R\right)^{-1}$$

Conditional on $(\mu, \sigma, \hat{x})$:

$$\begin{aligned}
\text{Var}[g(x)] &= \Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)\left(1-\Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)\right) \\
\text{Cov}[x, g(x)] &= \sigma\phi\left(\frac{\hat{x}-\mu}{\sigma}\right)
\end{aligned}$$

And becuase $R = 0$ in this example, the gain simplifies to:

$$K_k = \frac{\sigma h\left(\frac{\hat{x}-\mu}{\sigma}\right)}{\Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)}$$

The updated mean and variance are then given by:

$$\begin{aligned}
m_{k+1} &= \mu + K_k(y - \mathbb{E}y) \\
&= \begin{cases} \rho\left(\mu + \sigma h(\psi_k)\right) & \text{If } \; y_k = 1 \\ \rho\left(\mu - \sigma r(\psi_k)\right) & \text{If } \; y_k = 0 \end{cases} \\
s_{k+1}^2 &= \sigma^2 - K_k \text{Var}[g(x)] K_k \\
&= \sigma^2 - \frac{\sigma^2\left(\phi\left(\frac{\hat{x}-\mu}{\sigma}\right)\right)^2}{\Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)\left(1-\Phi\left(\frac{\hat{x}-\mu}{\sigma}\right)\right)} \\
&= \sigma^2\left(1 - r(\psi_k)h(\psi_k)\right)
\end{aligned}$$

## A.3 Primitive assumptions that generate government preferences

This appendix outlines a set of assumptions that would generate the government's period payoff function defined in Section 2.2.1.

The government's tax income is a constant fraction of output $T$. Output grows at constant rate $g$. The government has two forms of outstanding obligations: previously incurred incurred debts $b_t$, and required expenditure $\bar{G}_t$ (both measured as fractions of output in period $t$). Expenditure is required in the sense that:

1. Required expenditures are senior to past debt. So $G_t \leq \bar{G}_t$ only if the government defaults on $b_t$.

2. Households' value expenditure up to $\bar{G}_t$ but not thereafter. Households have preferences over realized expenditure $G_t$ given by:

$$u^H(G_t, \bar{G}_t) = \min(G_t, \bar{G}_t, 0)$$

I make one further modeling assumption:

4. The government cannot borrow to fund $G_t > \bar{G}_t$

Assumption 3 can be interpreted as a modified balanced budget rule. To government can issue some debt, but this is limited to only required expenditures. This rules out situations where the government can signal its type by issuing enough debt to make $G_t \geq \bar{G}_t$ even though this provides no extra period payoff.

Assumptions 2 & 3 imply that the realized government expenditure is:

$$G_t = \begin{cases} \max(\bar{G}_t, T - b_t) & \text{under repayment} \\ T & \text{under default} \end{cases}$$

The maximand in the first expression arises becuase the government may pay down all its debts. In that case, $b_{t+1} = 0$ and $G_t$ is the fraction of $T$ which is left over after paying for $b_t$. Usually, $b_{t+1} > 0$ and so $G_t = \bar{G}_t$. Of course, in both cases the household payoff is zero, as $G_t \geq \bar{G}_t$. The solution under default follows from exclusion from international capital markets alone, and not any of the extra assumptions in this section.

Letting $s_t = T - \bar{G}_t$, then the household's indirect utility is:

$$\text{Repayment:} \qquad u^H(G_t, \bar{G}_t) = u^H(\max(\bar{G}_t, T - b_t), \bar{G}_t) = 0$$

$$\text{Default:} \qquad u^H(G_t, \bar{G}_t) = u^H(T, \bar{G}_t) = \min(s_t, 0)$$

Which is the government's payoff in Section 2.2.1 (excluding the private taste for default).

# APPENDIX B

# DATA ON SPREADS & DEFAULTS

## B.1 Spreads

The EMBI stripped spreads are published by JP Morgan and are a measure of the premium that emerging market governments pay on their borrowing above that paid by the US government. For an individual country the EMBI spread is an average of the spread on each of their eligible liabilities. The set of eligible liabilities is comprised of all US$ denominated bonds[1] and loans with remaining maturity of at least two and a half years for which prices are readily available to JP Morgan dealers each day.

For each liability, the spread is given by the size of the uniform increase in the risk-free US$ yield curve which would be required to equate the value of the discounted cashflow of the debt to its market price. For collateralized bonds (such as Brady bonds), the value of the collateral is subtracted from the price of the debt to give a price for only the unsecured part. Individual liabilities are then averaged to a country-level spread index using weights proportional to the market value of the outstanding obligation.

Daily EMBI spreads are available for 63 countries with start dates varying from 1993 through to 2013. As data on defaults cover only those settled in 2006 or earlier, I restrict the sample to those with at least 3 years of data before this cutoff. The 27 remaining countries and their start dates are shown in Table B.1. Between them, these countries produced some 90% of the IMF's definition of emerging economy output in 2006. The daily time series for spreads are shown in Figure B.1

1. In theory, this includes both external and domestic US$ borrowing In practice, only Argentina markets any substantial amount of foreign currency debt to home investors.

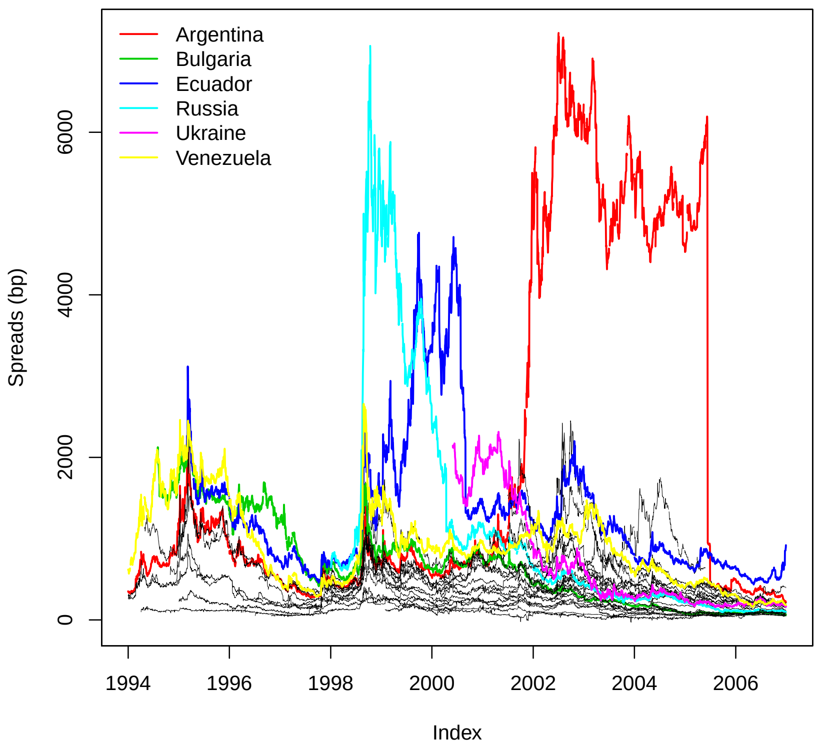


Figure B.1: Daily time series data on spreads. Select countries in color.

| Argentina | Brazil | Bulgaria | Chile | China |
|---|---|---|---|---|
| 1993 | 1994 | 1994 | 1999 | 1994 |
| Colombia | Dominican Republic | Ecuador | Egypt | El Salvador |
| 1997 | 2001 | 1995 | 2001 | 2002 |
| Hungary | Indonesia | Lebanon | Malaysia | Mexico |
| 1999 | 2004 | 1998 | 1996 | 1993 |
| Pakistan | Panama | Peru | Philippines | Poland |
| 2001 | 1996 | 1997 | 1993 | 1994 |
| Russia | South Africa | Tunisia | Turkey | Ukraine |
| 1997 | 1994 | 2002 | 1996 | 2000 |
| Uruguay | Venezuela | | | |
| 2001 | 1993 | | | |

Table B.1: Start dates for spreads

## B.2 Defaults

I use Benjamin and Wright's data on sovereign default and settlement. They provide data on 90 sovereign default episodes in 73 countries where settlement occurred between 1980 and 2006. Start and end dates of the episodes are based on S&P's default criteria, augmented with other academic work and information from original sources. These start and end dates are then used to compute length of default. Benjamin and Wright also compute the haircuts on outstanding debt at settlement. These measure investors' losses on their outstanding debts when they accept the defaulting sovereigns' settlement. This is computed using the zero coupon equivalent (ZCE) measure of debt. Proposed by Dias, Richmond, and Wright (2011), this is a measure of outstanding sovereign liabilities which includes coupon payments and is unaffected by default probabilities.

I use the subset of Benjamin and Wright's data for which I have sovereign spreads data for at least three years. This leaves 27 default episodes in 19 countries. These defaults are shown in Table B.2.

| Country | Length (years) | Haircut (%) | Debt Service at Default |
|---|---|---|---|
| Argentina | 11.2 | 30 | |
| Argentina | 3.6 | 63 | |
| Brazil | 11.2 | 21 | 6.6 |
| Bulgaria | 4.3 | 46 | 6.6 |
| Chile | 7.4 | 46 | |
| Colombia | 5.3 | 2 | 5.7 |
| Dominican Republic | 10.9 | 47 | 4.4 |
| Ecuador | 12.3 | 23 | 10.8 |
| Ecuador | 1.7 | 34 | 11.4 |
| Ecuador | 1.1 | 0 | 10.3 |
| El Salvador | 15 | 64 | 2.9 |
| Mexico | 7.9 | 34 | 9.0 |
| Pakistan | 1.6 | 29 | 3.7 |
| Panama | 12.7 | 34 | 10.0 |
| Peru | 0.9 | 0 | 12.3 |
| Peru | 14.4 | 29 | 7.8 |
| Philippines | 9.6 | 35 | 9.1 |
| Poland | 12.9 | 42 | |
| Russia | 6 | 32 | |
| Russia | 2.3 | 32 | |
| South Africa | 0.7 | 0 | |
| Ukraine | 1.4 | 1 | 4.8 |
| Uruguay | 1.1 | 16 | |
| Uruguay | 0 | 0 | |
| Venezuela | 1 | 14 | |
| Venezuela | 2 | 2 | |
| Venezuela | 0 | 0 | |

Table B.2: Default episodes

To check that my restriction on the sample of defaults does not look fundamentally alter its character, I compare the mean haircut and duration of default in the full and restricted samples, as well as their correlation. Table B.3 displays these statistics. Although delays and haircuts are smaller in the restricted, the difference is not statistically significant at the 5% level. The haircut and correlation between the two series are almost identical in the two samples.

| | Defaults | Countries | Delay (years) | Haircut | Correlation |
|---|---|---|---|---|---|
| Benjamin-Wright (2013) | 90 | 73 | 7.38 | 37.67 | 0.66 |
| This paper | 27 | 19 | 5.87 | 25.04 | 0.59 |
| Difference p-value | | | 0.15 | 0.89 | 0.58 |

Table B.3: Mean delays and haircuts in restructuring

# APPENDIX C

# ESTIMATING THE DYNAMIC PROPERTIES OF PRIMARY SURPLUSES

I take a sample of 8 Latin American countries starting mostly in 1990 or before and ending in 2014. Surpluses for more middle income countries are available, but almost all with much later start dates. Including these results in estimates for surplus processes that are driven almost entirely by the crisis from 2008 onwards.

Figure C.1 shows the partial autocorrelation function for the demeaned annual surplus/GDP ratio for these countries. These show that the partial autocorrelation is statistically significant for at most one lag, suggesting that a first-order Markov process well-describes the surplus process in most middle-income countries. Table C.1 presents (amongst other things) the optimal lag length as determined by the Akiake Information Criterion. This suggests that a short-lag Markov process is appropriate.

Give that a first-order Markov process seems to be a good fit to the data, I estimate for each country:

$$s_t = (1-\rho)\mu_s + \rho s_{t-1} + \epsilon_t$$

The estimates of $\mu_s$, the sample mean, the autocorrelation $\rho$, and the variance of $s_t$ are included in Table C.1. The units are in percent of GDP. In general, surpluses are on average positive, moderately persistent, and fluctuate annually by a few percentage points of GDP.

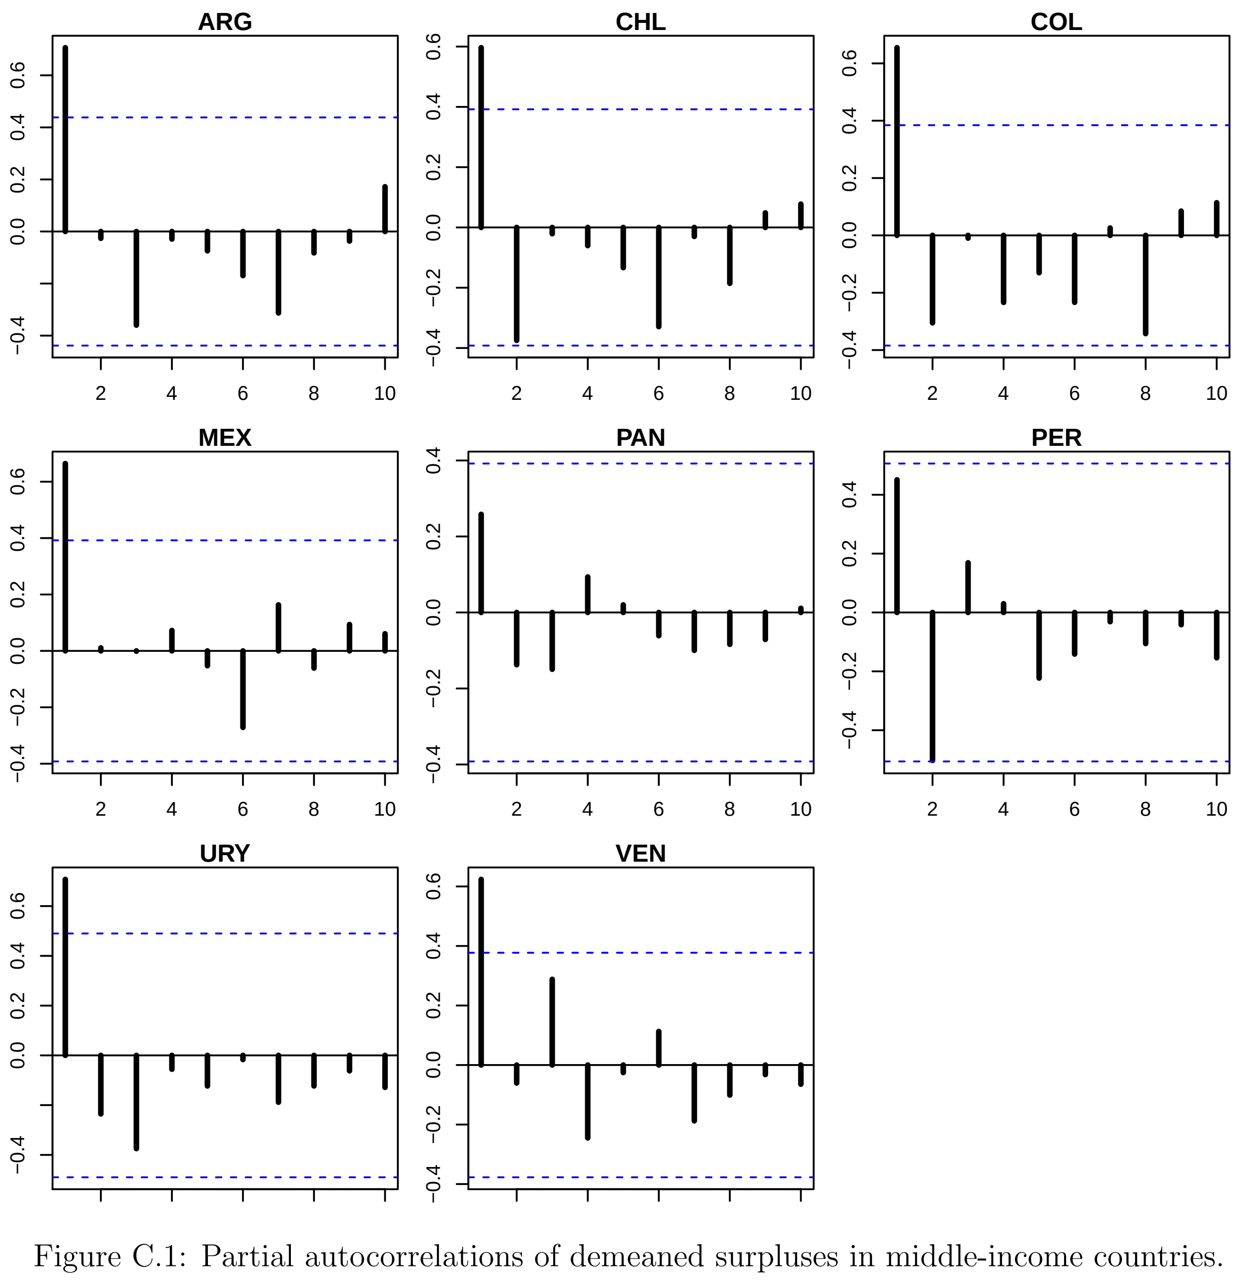


Figure C.1: Partial autocorrelations of demeaned surpluses in middle-income countries.

Table C.1: Summary statistics of surplus processes in middle income countries

| | Country | Sample years | N | AIC lag length | $\mu$ | $\sigma^2/(1-\rho^2)$ | $\rho$ |
|---|---|---|---|---|---|---|---|
| ARG | Argentina | 1995-2014 | 20 | 1 | 0.81 | 1.70 | 0.71 |
| CHL | Chile | 1990-2014 | 25 | 2 | 1.97 | 2.85 | 0.60 |
| COL | Colombia | 1989-2014 | 26 | 2 | 0.89 | 1.40 | 0.65 |
| MEX | Mexico | 1990-2014 | 25 | 1 | 1.00 | 1.87 | 0.66 |
| PAN | Panama | 1990-2014 | 25 | 0 | 2.21 | 4.10 | 0.26 |
| PER | Peru | 2000-2014 | 15 | 2 | 1.78 | 1.68 | 0.45 |
| URY | Uruguay | 1999-2014 | 16 | 1 | 1.29 | 1.90 | 0.71 |
| VEN | Venezuela | 1988-2014 | 27 | 1 | $-0.38$ | 6.93 | 0.62 |
| | Average | | 22.4 | 1.15 | 1.23 | 2.98 | 0.58 |