A Continuous-Time Model of Sovereign Debt

Gideon Bornstein[*]

The Wharton School, University of Pennsylvania

June 2020

Abstract

I construct a continuous-time model of strategic default and provide a numerical algorithm that solves it. I compare the results and computation times to standard discrete-time models of sovereign debt. The method proposed here is faster than discrete-time computation methods while obtaining similar quantitative results. The few differences between the models can all be attributed to a *painful deleveraging* feature. When debt issuance happens at a higher frequency, the sovereign faces higher interest rate spreads along the deleveraging process. So rolling over its debt becomes more costly. This feature leads to a coefficient of variation for interest rate spreads that is higher and closer to the data relative to its discrete-time sovereign debt model counterpart, calibrated to quarterly frequency. I solve three variants of the model. The first includes short-term maturity bonds only and a constant risk-free interest rate. The second allows for stochastic fluctuations in the risk-free rate. Finally, I extend the model to allow for long-term bonds.

Keywords: Sovereign debt, Default, Business cycles, Continuous time, Numerical methods.

JEL codes: E44, F34, F41.

1 Introduction

One of the promising developments in macroeconomics during the past few years has been the application of continuous time methods to incomplete-markets models. Such methods have been applied in Brunnermeier and Sannikov (2014) to study financial frictions in a macroeconomic model, in Gabaix et al. (2016) to study the dynamics of inequality, and in Kaplan et al. (2016) to study monetary policy with heterogeneous agents. This paper extends these methods to problems of strategic default on unsecured debt. The contribution of this paper is twofold. First, I construct and solve a continuous time model which expands the quantitative capabilities of strategic default models as its solution is faster and can accommodate a large number of state variables. Second, I show that the solution of the model is very similar to its counterpart discrete time model, and that the few differences between the two can all be attributed to a *costly deleveraging* feature.

The benchmark model considered is a short-term debt model, a continuous-time version of Arellano (2008). A benevolent sovereign faces fluctuations in its domestic output and can save or borrow using short-term debt contracts. The stochastic process for the sovereign's output is not continuous over time but features jumps. At any point in time, the sovereign can choose to default on its debt obligations. If the sovereign defaults, it is excluded from financial markets for a stochastic period of time, during which it suffers an output loss. Risk-neutral international investors with access to a risk-free world interest rate buy the sovereign bond. In addition to the risk-free interest rate, the bond carries a premium corresponding to the risk of default. The sovereign takes into account the interest rate schedule when making its debt issuance decision. I propose a numerical method to solve the sovereign's problem, study the business cycle statistics of the economy, and compare them to the solution of the discrete time version of the model. Finally, I show how to extend the analysis to different environments, e.g., allowing the government to issue long-term maturity bonds.

I calibrate the benchmark model according to Arellano (2008) and study the differences between the business cycle statistics of the model in continuous time to the ones in discrete time. Both the discrete-time and continuous-time short-term debt models rely on the assumption that debt cannot be diluted. That is, sovereign bonds mature before more bonds can be issued. However, as opposed to discrete time short-term debt models, where the researcher can control the maturity length of sovereign bonds by choosing the length of a period in the model, short-term bonds in continuous time mature in high-frequency. So one drawback of modeling short-term sovereign debt in continuous time is that it abstracts from the maturity length of bonds. In the calibrations considered, I find that this drawback is rather minor as the quantitative results of the continuous-time and discrete-time models are quite similar.

While business cycle statistics of the discrete-time Arellano (2008) model and its continuous-time counterpart are overall similar, there are two notable differences regarding the behavior of the trade balance and of spreads. In continuous time, the trade balance is less counter-cyclical and

the average spread is lower. These differences are not driven by the different maturity length of sovereign bonds, but rather by one main feature of the continuous time model - painful deleveraging.

The main difference between discrete-time models of sovereign debt and the models presented in this paper is that in the latter, the deleveraging process is more costly. Consider a negative endowment shock which increases the sovereign's risk of default, as well as the spread, holding fixed the sovereign's debt level. In the discrete-time environment, the sovereign can immediately decrease its stock of debt which needs to be payed in the following period. This action decreases the probability of default, so that the sovereign does not face the high equilibrium spreads which it would if it kept its level of debt unchanged. In the continuous time version of the model, on the other hand, a negative jump in the stock of debt is not feasible. The negative endowment shock moves the sovereign into a region of higher spreads from which it cannot instantaneously escape. Deleveraging in continuous time is painful both because the sovereign faces higher spreads during the deleveraging process, and because quick deleveraging can only be done by cutting consumption considerably.

Painful deleveraging alters equilibrium outcomes along several dimensions. The continuous-time sovereign[1] is worse off at any point on the state space in which the discrete-time sovereign chooses to decrease its debt obligations. Denote by *the default frontier* the set of points on the state space for which the sovereign is indifferent between defaulting and repaying its debt. Whenever the discrete-time sovereign is indifferent between defaulting and deleveraging, the continuous time sovereign prefers to default. For low and medium endowment levels, the discrete-time sovereign chooses not to deleverage on the default frontier, and is constrained to hold its level of debt fixed.[2] For high endowment levels, the discrete-time sovereign chooses to deleverage when it is on the default frontier. Therefore, the default region of both sovereigns is similar for low and medium endowment levels but it is larger for the continuous-time sovereign for high endowment levels. In addition, policy functions of the sovereign in the two frameworks are remarkably similar for low and medium levels of endowment. For high levels of endowment and small levels of debt, both sovereigns choose to increase their debt. However, the discrete-time sovereign does so faster as the continuous-time sovereign has a lower incentive to reach a high level of debt associated with a high spread. So the continuous-time sovereign exhibits lower counter-cyclicality of the trade balance.

Since deleveraging is costly, the continuous-time sovereign tries to avoid regions of the state space associated with high spreads. This cautious behavior leads to a lower average spread in the continuous-time framework. However, negative endowment shocks do occasionally move the sovereign to regions of high spread from which it cannot immediately escape. So while the

---

[1]From here onwards I use 'continuous-time sovereign' to refer to the sovereign in the continuous time environment, and 'discrete-time sovereign' for the discrete time environment.

[2]On the default frontier, the sovereign cannot increase its debt obligations as this would lead to an immediate default. This is explained in details in section 2.

average spread is lower, the volatility of spreads is similar to the discrete-time environment. This implies that the coefficient of variation, the ratio between the volatility of the spread to its mean, is higher in the continuous-time environment. As Aguiar et al. (2016) explains, matching the high volatility of the spread in the data as well as the low probability of default is a difficult challenge for quantitative sovereign debt models. The continuous time environment helps overcoming the challenge for the same reason that deleveraging is costly. Avoiding high spreads following a negative endowment shock by instantaneously cutting the level of debt is infeasible.

It is important to note that the continuous-time version of the Arellano (2008) economy I study is not the limiting economy when taking the length of a period in the discrete-time model to zero. This is because I model the endowment process of the sovereign as a compound Poisson process rather than a Brownian motion. This modeling choice gives rise to defaults along the equilibrium path as in Arellano (2008).[3] The methodological contribution of this paper is to construct a continuous-time model that can be solved in a short amount of time and that its equilibrium is similar, both qualitatively and quantitatively, to workhorse discrete-time sovereign debt models. Nonetheless, the numerical methodology accommodates also a Brownian Motion for the endowment process of the sovereign.[4]

The numerical method proposed takes advantage of sparse matrix operations. It is between two to fifty times faster than discrete time solution methods, depending on the density of the grid. It's comparative advantage is in analyzing a dense asset grid as it does not add many non-zero elements to the sparse matrix used in computations. There is an additional computational advantage when solving long-term debt models. In discrete-time models, it is notoriously difficult to compute the interest rate schedule as it involves solving a fixed point problem. Chatterjee and Eyigungor (2012) discuss this difficulty and show how in order for the value functions to converge one needs to add an additional state variable to the problem, a noise to the endowment process of the sovereign. In the continuous time framework, on the other hand, finding the interest rate schedule is a simple task as it boils down to a simple sparse matrix inversion. Finally, I show how one can use the Kolmogorov Forward equations to obtain the ergodic distribution of the model. The Kolmogorov Forward equations have an additional role in the context of heterogeneous agents, e.g., in models of consumer bankruptcy. They can be used to track the evolution of assets traded by all agents.

In the final part of the paper, I solve a continuous-time version of the long-term debt model of Chatterjee and Eyigungor (2012). The quantitative results of the continuous-time and discrete-time models are very similar. With long-term maturity bonds, deleveraging is less costly than in the benchmark model because acquiring back debt is cheap when the spread is high. While low prices of long-term debt in bad states provide a welfare benefit also in lower debt-issuance

---

[3]In Section 2, I explain why the limiting economy of Arellano (2008), taking the length of a period to zero, has no defaults along the equilibrium path.

[4]The computing times for an economy with a Brownian motion endowment process are an order of magnitude smaller than when the endowment process follows a compound Poisson process

frequencies,[5] this benefit is larger in the continuous-time framework. Under Chatterjee and Eyigungor (2012) calibration, I find that the optimal debt maturity in the continuous-time framework is slightly longer than in the discrete-time framework. In discrete time, the optimal debt maturity is the short-term one, a one-quarter bond. In continuous time the optimal maturity length of the bond is, on average, 1.2 quarters.

This paper is related to several strands of the literature. First, it is related to the theoretical and quantitative literature on sovereign debt, building on the seminal work of Eaton and Gersovitz (1981). It includes Arellano (2008) and Aguiar and Gopinath (2006) for short-term debt, and Hatchondo and Martinez (2009), Arellano and Ramanarayanan (2012), Chatterjee and Eyigungor (2012), and Sánchez et al. (2018) for long-term debt. This paper is also related to models of consumer bankruptcy such as Chatterjee et al. (2007), Livshits et al. (2007), and Mitman (2016). While focusing on sovereign debt, the methods presented in this paper can also accommodate general equilibrium models of consumer bankruptcy.

I build upon the work of Achdou et al. (2020), which shows how to solve incomplete-markets heterogeneous-agents models in continuous time. They study an environment without default and impose an exogenous borrowing limit, which is independent of the current endowment of the agent. In contrast, I study an environment with a strategic default decision that implies an endogenous borrowing limit on the sovereign.

This paper is not the first one to study models of sovereign debt in continuous time. Nuño and Thomas (2019) studies the effects of monetary policy in a sovereign debt model using a continuous time framework, and Tourre (2017) studies the behavior of bond spreads with a non-separable utility function. Both papers use a Brownian motion for the exogenous process of the sovereign's endowment. While this assumption makes the analysis simpler, as a smooth pasting condition can be used on the default frontier, it makes the comparison to standard discrete time models difficult. For instance, these papers can only study long-term debt as there are no defaults in equilibrium if debt is of short maturity.[6] The model presented in this paper can accommodate short-term debt, and the types of defaults are similar to discrete-time models. Both Nuño and Thomas (2019) and Tourre (2017) do not compare the results of their models to ones they would have obtained in a discrete-time environment. Compared to Nuño and Thomas (2019) and Tourre (2017), this paper can be seen as a bridge for understanding the underlying mechanism and differences between continuous- and discrete-time models of sovereign debt.[7]

Aguiar et al. (2015) use a continuous-time environment to study fiscal and monetary policy in a monetary union with the potential of rollover crises on sovereign debt. As in the benchmark model I consider, they study short-term maturity bonds. Since no exogenous variable in their

---

[5] Arellano and Ramanarayanan (2012) refer to this welfare benefit of long-term debt as the 'hedging benefit'.

[6] This result is discussed in Section 2

[7] In terms of numerical strategy, Nuño and Thomas (2019) uses a finite difference method similar to the one in Achdou et al. (2020), and Tourre (2017) uses a Markov Chain Approximation Method. For a discussion about how the finite difference method compares to other numerical methods, I refer the reader to section 4.7 in Achdou et al. (2020).

model jumps over time, in the absence of rollover risk the equilibrium does not feature sovereign defaults. They adopt the timing convention of Cole and Kehoe (2000), which gives rise to rollover crises–the source of sovereign defaults in their model. In contrast to them, I abstract from rollover risk and study strategic sovereign defaults in the spirit of Eaton and Gersovitz (1981).

Lorenzoni and Werning (2019) and Aguiar and Amador (2019) use a continuous-time environment to study the existence of multiple equilibria in the long-term sovereign debt model. Lorenzoni and Werning (2019) study self-fulfilling equilibria in which high interest rate spreads, due to the fear of a future default, lead to a gradual but faster accumulation of debt, ultimately validating investors' fear. Aguiar and Amador (2019) analyze the necessary condition for such multiplicity to arise, and study whether "lender of last resort" policies can select a particular equilibrium. Similar to this paper, the income shock these papers consider includes a Poisson process.[8] Differently from this paper, in Lorenzoni and Werning (2019), upon arrival of the Poisson shock, all uncertainty is resolved. And in Aguiar and Amador (2019), the sovereign is assumed to be risk-neutral.

The paper proceeds as follows. Section 2 presents the benchmark model of short term debt. Section 3 contains the numerical method that solves the model. The calibration of the model, its equilibrium, and the concept of painful deleveraging are presented in Section 4. Section 5 considers the two extensions of the benchmark model, world interest rate fluctuations and long term debt. Section 6 concludes.

2   Benchmark Model

In this section, I present the benchmark model. It is a continuous-time version of Arellano (2008). A sovereign that faces fluctuations in its domestic output and a constant world interest rate optimally chooses the amount of debt it holds over time. The sovereign may choose to default on its debt obligations, in which case it is excluded from world financial markets for a stochastic period of time and faces output losses. The strategic default decision implies that the interest rate the sovereign pays on its debt depends on the probability of future default, and therefore on its current state.

Time is continuous and denoted by $t \in [0, \infty)$. I assume the sovereign is a small open economy so it takes the fixed world interest rate, r^f, as given. The sovereign receives utility from aggregate consumption which is denoted by c_t. Its preferences are given by

$$\mathbb{E}_0 \left\{ \int_0^\infty e^{-\rho t} u(c_t) dt \right\} ,$$

---

[8]Another paper which uses Poisson shocks in a continuous-time model of sovereign debt is Rebelo et al. (2018). They study how lack of financial development can lead to a "debt intolerance" phenomenon. Differently from this paper, and from workhorse discrete-time models of sovereign debt, they assume that the sovereign's endowment does not feature mean reversion. This assumption allows them to dispense with one state variable, and characterize the sovereign's behavior by its debt-to-income ratio.

where $\rho > 0$ is the discount factor, and the flow utility function, $u(c)$, is strictly increasing, strictly concave, differentiable, and satisfies Inada conditions. The sovereign receives a strictly positive flow endowment y_t at time t. I assume that y_t follows a compound Poisson process with arrival rate λ_y. That is, the flow endowment remains at its current level until a shock hits the economy. The probability of the economy experiencing a shock in a time interval dt is $\lambda_y dt$. Denote by y^- the flow endowment of the sovereign *just* before the shock hits the economy. When the shock hits the economy, the new endowment of the sovereign is drawn from the conditional distribution function $F(y, y^-)$.[9]

Figure 1 displays a sample path for the realization of endowment. It is critical that the endowment process features stochastic *jumps* in order to observe defaults on the equilibrium path. Consider, for example, a case in which endowment follows an Ornstein–Uhlenbeck process. This is a common stochastic process considered in continuous-time economic models, which moves continuously over time. In such case, if the sovereign is on the default frontier its probability of going into the default region in any interval of time, however small it may be, is equal to 1. This implies that the equilibrium interest rate on the default frontier is equal to infinity. Such interest rate schedule enforces the sovereign to never be on the default frontier. So there are no defaults on the equilibrium path.[10]

Figure 1: Example of an Endownment Process
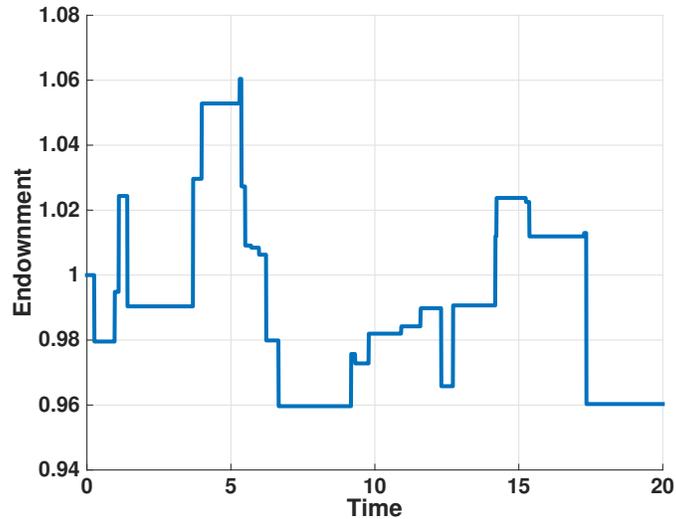
Notes: a realization of an endowment process with arrival rate $\lambda_y = 1$. Upon realization of the shock the new endowment is drawn from $\ln y = \rho \ln y^- + \epsilon$, with $\rho = 0.945$ and $\epsilon \sim \mathcal{N}(0, 0.025^2)$.

---

[9]I assume that this distribution function admits a joint density function denoted by $f(y, y^-)$ which satisfies $\int_0^\infty f(y, y^-) dy = 1$, for all y^-.

[10]This argument holds when the only asset available to the sovereign is short term debt. In case of long term debt, there may be defaults on equilibrium path even when the endowment process is continuous. Nuño and Thomas (2019), Tourre (2017), and Rebelo et al. (2018) study such a case.

The sovereign can save and borrow using short-term debt contracts. As in the discrete-time version of Arellano (2008), the key assumption is that debt cannot be diluted. That is, sovereign bonds mature before additional bonds can be issued. In continuous-time, this implies that bonds mature instantaneously.[11]

I denote the amount of assets held by the sovereign at time t by a_t. A negative value of a_t is interpreted as debt obligations. Conditional on repayment of debt obligations, the law of motion for a_t follows

$$da_t = (y_t - c_t)dt + r(\Omega_t)a_t dt , \qquad (1)$$

where $r(\Omega_t)$ is the interest rate the sovereign receives on its assets (or pays on its debt), which is a function of its current state denoted by Ω_t. As I explain below, the current state of the sovereign is fully characterized by its current level of output and assets, and whether it is excluded from international financial markets or not. The assumptions on $u(c)$ imply that it is never optimal for the sovereign to have a_t move in a discontinuous fashion. Notice that a discontinuous jump in a_t can occur only if $c_t \to \pm\infty$. Inada conditions imply setting $c_t \to \infty$ is sub-optimal, and $c_t \to -\infty$ is ruled out by non-negativity of consumption.

At any moment in time, the sovereign may choose to default on its debt obligations. In that case, its outstanding debt is reset to zero and it is excluded from international financial markets for a stochastic period of time. When excluded from financial markets, the sovereign suffers an output loss of $\phi(y)$ so that its flow endowment is given by $y_t - \phi(y_t)$. The sovereign regains access to international financial markets with Poisson intensity λ_D. So the average amount of time a sovereign is excluded from financial markets after a default is $\frac{1}{\lambda_D}$.

Given the interest rate schedule, I can write the Hamilton-Jacobi-Bellman (HJB) equations for the sovereign. Let us denote by $w(\cdot)$ the value function of a sovereign that is excluded from international financial markets, and by $v(\cdot)$ the value function of a sovereign that is not. When excluded from financial markets, the state of the sovereign is simply given by y_t. Its HJB equation is given by

$$\rho w(y) = u(y - \phi(y)) + \lambda_y \int_0^\infty \left(w(y') - w(y) \right) f(y', y)dy' + \lambda_D \left[v(0, y) - w(y) \right] . \qquad (2)$$

The HJB equation of a sovereign that has access to international markets is given by

$$\rho v(a, y) = \max_c \{ \rho w(y),$$
$$u(c) + v_a(a, y)(y - c + r(a, y)a) + \lambda_y \int_0^\infty \left(v(a, y') - v(a, y) \right) f(y', y)dy' \} . \qquad (3)$$

The first argument of the max represents the option of the sovereign to default. The second argument represents the option to repay debt obligations, choose the level of consumption and issue new bonds. The derivative of the value function with respect to the asset level is denoted by

---

[11] If debt maturity is equal to one quarter, as in Arellano (2008), but the government can issue bonds at any point in time, then debt pricing must take into account the possibility of debt dilution. That is, the model is effectively a long-term debt model.

8

$v_a(a, y)$. The term including $v_a(a, y)$ in equation (3) represents the effect of changing the bond holding position of the sovereign. Let the choice whether to default or not be denoted by $D(a, y)$, where it takes the value 1 for default, and 0 for no-default. Denote by $\mathcal{D} \subseteq \mathbb{A} \times \mathbb{Y}$ the region where the sovereign chooses to default. I assume that when the sovereign is indifferent between repaying its debt and defaulting, it repays its debt. Conditional on choosing to repay its debt, I can derive the necessary first order condition for optimal consumption decision,

$$u'(c(a, y)) = v_a(a, y) ,$$

where $c(a, y)$ denotes the optimal consumption level of a sovereign with asset position a, and current flow endowment y. In the appendix, I show that the value function of the sovereign in the region of no default is increasing in assets (decreasing in bond obligations), i.e., $v_a(a, y) > 0$. So the implicit level of consumption that satisfies the equation above is strictly positive. Since the value of a sovereign with no access to international markets does not depend on previous bond holdings, this immediately implies the following Proposition.

Proposition 1. *The default decision of the sovereign is of the threshold kind in bond obligations. That is, for every level of endowment y, there exists a level of assets $\underline{a}(y)$ such that*

$$D(a, y) = \begin{cases} 1, & if \quad a < \underline{a}(y) , \\ 0, & if \quad a \geq \underline{a}(y) . \end{cases}$$

The threshold level of bond holdings implies a borrowing constraint on the sovereign. Whenever the sovereign is at state $(\underline{a}(y), y)$, it cannot increase its level of outstanding bond obligations.[12] That is, $\dot{a}(\underline{a}(y), y) \geq 0$, for all y. This borrowing limit implies the following state boundary constraint.

$$v_a(\underline{a}(y), y) \geq u' \left[y + r(\underline{a}(y), y)\underline{a}(y) \right] , \qquad \forall y .$$

This condition ensures that for every level of endowment y,

$$c(\underline{a}(y), y) \leq y + r(\underline{a}(y), y)\underline{a}(y) ,$$

so that $\dot{a}(\underline{a}(y), y) \geq 0$. Thus, the model features an *endogenous borrowing limit*. The borrowing limit the sovereign faces depends on its current output and is equal to the level of debt such that the sovereign is indifferent between defaulting and repaying its debt obligations.

Let us now turn to the characterization of the interest rate schedule the sovereign faces. Assume risk neutral international banks have unlimited access to funds at the world interest rate, r^f. Perfect competition among themselves ensures that in expectation they make zero profits on sovereign bond holdings. The expected interest on sovereign bonds in the region of no default, $(a, y) \notin \mathcal{D}$, is given by

$$\mathbb{E}\left[dr(a, y) \right] = r(a, y)dt - \lambda_y \int_0^\infty D(a, y')f(y', y)dy'dt .$$

---

[12]The sovereign cannot increase its debt obligations as such decision would imply an infinite interest rate. This becomes clear when I derive the interest rate equation below.

The first term is the repayment in case of no default, while the second term indicates the loss of funds in case of a sovereign default. The zero profit condition in the region of no default then implies

$$r(a, y) = r^f + \lambda_y \int_0^\infty D(a, y') f(y', y) dy' . \tag{4}$$

This equation characterizes the interest rate in the region of no default. Notice that the interest rate does not depend on the savings decision of the sovereign but only on the current state of the sovereign. This is different than the interest rate schedule in discrete time. There, the interest rate depends on the amount of debt in the following period and not on the stock of debt in the current period. As debt obligations move continuously over time, there is no difference between "next period's" debt and the current level of debt in the model. In the default region,

$$r(a, y) = \infty, \qquad \forall (a, y) \in \mathcal{D} . \tag{5}$$

To see this, consider that the probability of receiving any return on debt in the default region is lower than $\lambda_y dt$ so that the return on debt is of magnitude dt^2.

Equilibrium definition. The solution concept is a Markov perfect equilibrium. A Markov perfect equilibrium is a set of value functions, $v(a, y)$ and $w(y)$, policy functions, $c(a, y)$ and $D(a, y)$, and interest rate schedule $r(a, y)$, such that:

1. The value functions $v(a, y)$ and $w(y)$ are the solution to the HJB equations (2)–(3), and the optimal control is $c(a, y)$.

2. Given the sovereign's default policy function $D(a, y)$, the interest rate schedule satisfies equations (4)–(5).

3. The default policy function is consistent with the value functions.

3 Numerical Solution Methodology

This section describes the numerical algorithm that solves the model described in the previous section. It can easily be extended to accommodate a more general structure of state variables and accommodate modifications to the model. I show two extensions of the model in section 5, where I consider long-term debt and world interest rate fluctuations.

The solution algorithm is based on a finite difference method using an upwinding scheme. The algorithm is an extension to the method used in Achdou et al. (2020). In both models, an agent with access to non-state-contingent assets who faces idiosyncratic risk optimally chooses its consumption level. However, in their model there is no default and an exogenous borrowing limit, which is independent of the current endowment of the agent, is imposed. The benchmark model, in contrast, features a strategic default decision that implies an endogenous borrowing limit. This section provides a comprehensive description of the numerical solution methodology, but I refer

the reader to Achdou et al. (2020) and its appendix for a detailed introduction on how to solve the HJB equation using a finite difference method.

3.1 Solving the HJB Equation with a Finite Difference Method

I solve the problem on a discrete state space equidistant grid with I different points for endowment levels, and J different points for asset levels. In its essence, this algorithm is a simple value function iteration procedure. I start with a guess for the value functions and default decision. In every iteration I obtain updated value functions using the current value functions and default decision. The updated value functions imply an updated default decision and an updated interest rate schedule. I stop whenever the updated and old value functions are sufficiently close.

Denote by $v_{i,j}$ the value function of a sovereign with access to financial markets, endowment y_i, and asset level a_j. Denote by w_i, the value function of a sovereign with no access to financial markets and endowment $y_i - \phi(y_i)$. Let the interest rate at state (a_j, y_i) be denoted by $r_{i,j}$. Rewriting the HJB equations (2) and (3) in their discrete form I have

$$\rho w_i = u(y_i - \phi(y_i)) + \lambda_y \sum_{i'=1}^{I} (w_{i'} - w_i) f(i', i) + \lambda_D (v_{i,j_0} - w_i) ,$$

$$\rho v_{i,j} = \max \left\{ w_i, \ \max_{c_{i,j}} u(c_{i,j}) + \partial_a v_{i,j} (y_i - c_{i,j} + r_{i,j} a_j) + \lambda_y \sum_{i'=1}^{I} (v_{i',j} - v_{i,j}) f(i', i) \right\},$$

where $a_{j_0} = 0$, and $f(i', i)$ denotes the probability of jumping from endowment y_i to $y_{i'}$ conditional on the arrival of an endowment shock. I use superscript n to denote the n^{th} guess of the value functions, consumption levels, interest rates, and default decision. The value function is updated in every iteration using a semi-implicit method. The next guess of the value functions satisfies the following equations

$$\frac{w_i^{n+1} - w_i^n}{\Delta} + \rho w_i^{n+1} = u (y_i - \phi(y_i)) + \lambda_y \sum_{i'=1}^{I} \left(w_{i'}^{n+1} - w_i^{n+1} \right) f(i', i) + \lambda_D \left(v_{i,j_0}^{n+1} - w_i^{n+1} \right) , \quad (6)$$

$$\frac{v_{i,j}^{n+1} - v_{i,j}^n}{\Delta} + \rho v_{i,j}^{n+1} = \max_{c_{i,j}} u \left(c_{i,j}^n \right) + \partial_a v_{i,j}^{n+1} \left(y_i - c_{i,j}^n + r_{i,j}^n a_j \right)$$

$$+ \lambda_y \sum_{i'=1}^{I} \left(v_{i',j}^{n+1} - v_{i,j}^{n+1} \right) f(i', i), \qquad \text{for } j \geq \underline{d}_i^n , \quad (7)$$

$$v_{i,j}^{n+1} = v_{i,\underline{d}_i^n}^{n+1}, \qquad \qquad \text{for } j < \underline{d}_i^n , \quad (8)$$

where Δ is a numerical updating parameter. The choice of the value of Δ is discussed at the end of this section. The endogenous borrowing limit for output level y_i is given by $a_{\underline{d}_i}$. The consumption $c_{i,j}^n$ at the no-default region ($j \geq \underline{d}_i$) is not chosen over a grid, but rather determined using the first order condition, as in Barillas and Fernández-Villaverde (2007). That is, consumption $c_{i,j}$ is computed using the derivative of the value function as follows.

$$c_{i,j} = u'^{-1} \left(\partial_a v_{i,j} \right)$$

For ease of notation, in the equation above and in the following equations I omit the guess iteration superscripts as all variables refer to the n^{th} iteration. The value of $\partial_a v_{i,j}$ is approximated using a finite difference method. When approximating this derivative, one can use either the forward or backward derivatives,

$$\partial_a^F v_{i,j} = \frac{v_{i,j+1} - v_{i,j}}{da}, \qquad \partial_a^B v_{i,j} = \frac{v_{i,j} - v_{i,j-1}}{da} ,$$

where $da = a_{j+1} - a_j$ is the density of the grid.[13] Notice that I cannot compute the backward derivative on the default frontier $(\underline{d}_i)$, as $v_{i,\underline{d}_i-1} = v_{i,\underline{d}_i}$ by construction. Instead of using the approximated derivative, I shall use the derivative that is implied by the state boundary constraint in case it is binding. That is,

$$\partial_a^B v_{i,\underline{d}_i} = u'^{-1} \left(y_i + r_{i,\bar{d}_i} a_{\underline{d}_i} \right) .$$

Additionally, I cannot obtain the forward derivative at the highest asset position on the grid . However, for high enough asset levels the forward derivative is not used.

Assume that the flow utility is of the CRRA form with coefficient of relative risk aversion equal to γ. This assumption makes the remaining exposition of the numerical method perspicuous, but the method extends to any increasing and concave utility function. Using the approximated derivatives, I can compute the consumption levels with the forward and backward derivatives as follows.

$$c_{i,j}^F = \partial_a \left(v_{i,j}^F \right)^{-\frac{1}{\gamma}} , \qquad c_{i,j}^B = \partial_a \left(v_{i,j}^B \right)^{-\frac{1}{\gamma}} .$$

The budget constraint implies the savings decision using the backward and forward derivatives.

$$s_{i,j}^X \equiv \dot{a}_{i,j}^X = y_i - c_{i,j}^X + r_{i,j} a_j , \qquad \text{for } X \in \{F, B\}.$$

The upwinding scheme implies using the forward derivative whenever savings computed using the forward derivative are positive $(s_{i,j}^F > 0)$, and the backward derivative whenever savings computed using the backward derivative are negative $(s_{i,j}^B < 0)$. Since the value function is concave in assets, savings computed using the forward derivative are lower than savings obtained by using the backward derivative. In the case savings are negative using the forward derivative, and positive using the backward derivative, I use the consumption level that implies $s_{i,j} = 0$.

Having the consumption levels in hand, I can use equations (6)–(8) to obtain the updated value functions, w^{n+1} and v^{n+1}. These equations are all linear, so updating the value functions narrows down to a simple matrix inversion. In the appendix, I explain how to efficiently do this using sparse matrices. If the difference between the updated and previous value functions is small enough, the value functions have been found. Otherwise, I need to update the default decision and the interest rate schedule for the next iteration. The default decision is obtained in the following way.

$$\underline{d}_i^{n+1} = \max \left\{ j : \quad v_{i,j}^{n+1} \leq w_i^{n+1} \right\} .$$

---

[13]There is no subscript on da as I use an equidistant gird.

The interest rate schedule in the region of no default, for the following iteration, can then be computed as follows.

$$r_{i,j}^{n+1} = r_f + \lambda_y \sum_{i'=1}^{I} \mathbb{1}_{j < \underline{d}_{i'}^{n+1}} f(i', i), \qquad \forall i, j \quad \text{s.t.} \quad j \geq \underline{d}_i^{n+1}.$$

The interest rates at the default region are equal to infinity, but are not needed for the purpose of updating the value function. Using the updated default decision and interest rate schedule I can use equations (6) - (8) to repeat the procedure and update the value function again. I do so until the value functions converge. The updating parameter in equations (6) - (7), Δ, plays an important role in whether the value functions converge or not. A higher level of Δ decreases the total number of iterations until convergence is reached at the risk of having the value functions oscillate without ever converging. I suggest starting with a high value of Δ (I used 2 in all computations below). Then, if the value functions do not converge, decrease the value of Δ until convergence is reached.

3.2 Finding Model Statistics using the Kolmogorov Forward Equations

The solution to the HJB equations yields the policy functions, $c(\cdot)$ and $s(\cdot)$. With those in hand, we may want to study business cycle statistics implied by the model. One way to do so is to simulate the model over a long period of time and look at statistics observed from the simulation. In this section, I present an alternative way to find the unconditional and various conditional statistics of the model. This is done by using the Kolmogorov Forward (KF) equations.

The benchmark model includes a single agent, the sovereign. However, it is important to note that the KF equations have an additional role in the context of a heterogeneous-agents model. Consider a model of consumer bankruptcy where households can choose to renege on their debt obligations. The problem of the household in that context is almost identical to the problem of the sovereign. The risk-free interest rate in the framework of consumer bankruptcy, as opposed to a small open economy, may be endogenous and set such that markets clear. In such environment one needs a way to track the evolution of assets traded by all agents. This can be achieved using the Kolmogorov Forward equations presented in this section.

Generally, the KF equations characterize how the probability that a stochastic process is in a certain state changes over time. They are often used in heterogeneous agents models in continuous time to describe the evolution of the distribution of agents across different states over time. The KF equations in the benchmark model are given by

$$\partial_t \Gamma^{ND}(a, y, t) = -\partial_a \left[s(a, y) \Gamma^{ND}(a, y, t) \right] - \lambda_y \Gamma^{ND}(a, y, t) + \lambda_y \int_0^\infty f(y, \tilde{y}) \Gamma^{ND}(a, \tilde{y}, t) \, d\tilde{y}$$
$$+ \lambda_D \Gamma^D(y, t) \mathbb{1}(a = 0) , \tag{9}$$
$$\partial_t \Gamma^D(y, t) = -\lambda_y \Gamma^D(y, t) + \lambda_y \int_0^\infty D(a, y) f(y, \tilde{y}) \Gamma^{ND}(a, \tilde{y}, t) \, d\tilde{y} - \lambda_D \Gamma^D(y, t) , \tag{10}$$

where $\Gamma^{ND}(a, y, t)$ is the probability the economy has access to international markets and its state is (a, y) at time t. Equation (9) is defined for (a, y) such that the sovereign chooses to repay its debt, i.e., $D(a, y) = 0$. For states in which the sovereign chooses to default $\Gamma^{ND}(a, y, t) = 0$ for all t. $\Gamma^{D}(y, t)$ denotes the probability the economy has no access to international markets and its endowment is equal to y at time t.

The first term in equation (9) describes movement into state (a, y) due to the savings decision of the sovereign. If the sovereign always decides to hold assets constant, i.e., $s(a, y) = 0$, for all (a, y), then this term is equal to 0. The second term in equation (9) and the first term in equation (10) describe the probability of moving away from the current state due to an endowment shock. The following term describes movement into the state (a, y) from other endowment levels, $\tilde{y}$, due to an endowment shock. $f(y, \tilde{y})$ is the joint density function of switching from endowment $\tilde{y}$ to endowment y upon arrival of the endowment shock. Notice I integrate over $\tilde{y}$ and not y as this term describes movement *into* endowment y and not *from* endowment y. The last term in both equations describes the probability of regaining access to international financial markets.

Consider the unconditional distribution of the state of the economy of the sovereign. Define them by $\bar{\Gamma}^{ND}(a, y)$ and $\bar{\Gamma}^{D}(y)$. This distribution must satisfy $\partial_t \bar{\Gamma}^{ND}(a, y) = 0$, and $\partial_t \bar{\Gamma}^{D}(y) = 0$. So the distribution which satisfies the KF equations with the following two conditions is the unconditional distribution of the state of the economy.

After solving the value functions, finding the unconditional distribution boils down to inverting a sparse matrix as I show in the appendix. This implies that this method for obtaining the unconditional distribution is faster than simulating the model many times. The unconditional distribution provides us, among other statistics, the unconditional probability of the sovereign having no access to financial markets and the unconditional variance of the trade balance. The Kolmogorov Forward equations can also be used to obtain some conditional moments, such as the correlation of the trade balance and endowment conditional on the sovereign having access to financial markets.

4 Results and Discussion

In this section, I calibrate the model according to Arellano (2008). The results of the continuous-time model are overall quantitatively similar to the results of the discrete-time model. The few differences can be attributed to a feature that arises when debt issuance happens at higher frequency. This feature is *painful deleveraging*. In continuous time, assets move continuously across time so a bad endowment shock can move the economy into a region with high spreads for a short amount of time. In discrete time, the level of assets can jump between two consecutive periods. So a high spread due to a negative endowment shock can be avoided by an instant decrease in the level of debt. Painful deleveraging induces the sovereign to refrain from high-spread regions. Accordingly, the average spread in the benchmark model is lower than in the discrete-time Arel-

lano model. Nevertheless, the volatility of spreads in the benchmark model is similar to the one in discrete time. Negative shocks can move the sovereign into regions of high spreads from which it cannot instantaneously escape. So the coefficient of variation of spreads in the benchmark economy, the ratio between spread volatility to the average level of spreads, is about two times higher than in the discrete-time economy.

The computing time for the benchmark grid considered is between two to fifty times faster than standard discrete time computation methods. The continuous time methodology is relatively more efficient in studying a denser grid of asset levels, and relatively less efficient in considering denser endowment grids. In section 4.3, I discuss the source of the comparative advantage of the continuous time method. Finally, I study the difference between statistics computed using Monte Carlo methods as opposed to ones computed using the Kolmogorov Forward equations.

4.1 Calibration

Following Arellano (2008), the risk free interest rate, r^f, is calibrated to 1.7%. The coefficient of relative risk aversion, γ, is equal to 2. The discount factor, ρ, is calibrated to 4.81% so that it corresponds to the quarterly discount factor in discrete time, $\beta = 0.953$. The probability of switching from the bad credit state to the good one, λ_D, is equal to 0.282. The cost of default is modeled as a maximum level of output that can be reached in the bad credit state. That is,

$$\phi(y) = \begin{cases} 0 & \text{if} \quad y < \bar{y}, \\ y - \bar{y} & \text{if} \quad y \geq \bar{y}, \end{cases}$$

where $\bar{y}$ is calibrated to $0.969\mathbb{E}(y)$. Finally, I need to calibrate the stochastic process for output. Assume that $\lambda_y = 1$ so that on average a shock occurs once a quarter. When a shock occurs, a new productivity is drawn from a log-normal distribution with mean $\rho_y \ln y^-$. The persistence parameter, ρ_y, is calibrated to 0.945 and the standard deviation of the shock is set to 0.025.

The benchmark grid consists of 601 asset grid points uniformly distributed between -1 and 0.5, and 51 output grid points uniformly distributed between 3 standard deviations below and above the unconditional mean. The model is also solved on a coarse grid which consists of 301 and 25 asset and output grid points, respectively, and a dense grid with 1,801 and 151 asset and output grid points. For the dense grid, the output grid points are uniformly distributed between 4 standard deviations above and below the unconditional mean.

4.2 Results and Painful Deleveraging

The left panel of Table 1 displays business cycle statistics of the sovereign for the different grids. To be consistent with Arellano (2008), I simulate the model for 500,000 quarters and compute statistics across episodes of 74 quarters prior to default. The right panel of this table contains the corresponding business cycle statistics in the discrete time model. Even though the two models

are not isomorphic, the qualitative and quantitative results are quite similar across the two. There are four differences between the two models: (i) the volatility of the trade balance and the correlation between the trade balance and output are about 20%-30% lower in continuous time, (ii) the average spread is about 25% lower in continuous time, (iii) the debt to output ratio is about 15% lower in continuous time, and (iv) the coefficient of variation for the spread is about 40% higher in continuous time. In this section, I explain the main feature driving these differences, painful deleveraging.

Table 1: Business Cycle Statistics

	Continuous time					Discrete time				
	Benchm.	Coarse	A-fine	Y-fine	Fine	Benchm.	Coarse	A-fine	Y-fine	Fine
Asset grid points	601	301	1801	601	1801	601	301	1801	601	1801
Output grid points	51	25	51	151	151	51	25	51	151	151
$corr(c,y)$	0.99	0.99	0.99	0.99	0.99	0.98	0.98	0.98	0.98	0.98
$corr(TB/Y,y)$	-0.18	-0.23	-0.16	-0.18	-0.15	-0.23	-0.23	-0.24	-0.23	-0.23
$corr(r_s,y)$	-0.31	-0.32	-0.31	-0.32	-0.29	-0.26	-0.16	-0.40	-0.30	-0.38
$corr(r_s,TB/Y)$	0.71	0.71	0.70	0.70	0.71	0.53	0.35	0.60	0.54	0.65
$\sigma(y)$	5.27	5.20	5.37	5.24	5.42	5.72	5.79	5.70	5.67	5.67
$\sigma(c)$	5.46	5.45	5.54	5.43	5.59	6.11	6.23	6.10	6.05	6.04
$\sigma(TB/Y)$	0.72	0.80	0.71	0.73	0.70	1.15	1.25	1.15	1.09	1.09
$\sigma(r_s)$	4.76	5.88	5.06	4.75	4.72	4.44	6.01	4.52	4.37	3.60
$\mu(r_s)$	2.60	2.63	3.01	2.59	2.82	3.57	3.86	4.09	3.51	3.54
$\sigma(r_s)/\mu(r_s)$	1.88	2.30	1.75	1.88	1.73	1.32	1.63	1.17	1.31	1.08
Average debt-to-output (%)	3.65	4.67	3.33	3.64	3.18	4.23	4.74	4.42	4.02	4.04
Defaults per 500,000 quarters	2,254	2,024	2,268	2,254	2,222	2,229	2,262	2,239	2,183	2,157
Time to converge	23″	4″	1′:33″	2′:52″	10′:04″	2′:01″	12″	22′:03″	5′:14″	9h 18′:51″

The main source of difference between the two models is the result of the law of motion for assets. In the discrete-time model, assets move discretely across time. In continuous time, the optimal evolution of assets across time is continuous. Consider a negative endowment shock that moves the sovereign into a region with higher risk of default. In discrete time, the sovereign can choose to decrease its debt obligations so that the spread it pays on its debt does not increase substantially with respect to the previous period. In continuous time, the negative endowment shock moves the sovereign into a region of higher spreads from which the sovereign cannot instantaneously escape. Deleveraging in that case is costly both because the sovereign has to face higher spreads during the deleveraging process, and because quick deleveraging can only be done by cutting consumption considerably.

I illustrate this difference in Figure 2. Let's consider two economies, a continuous-time one (CTE) and a discrete-time one (DTE) and assume both start at the same point, point A. For simplicity I shall consider the case in which the interest rate schedule in both economies is equivalent. Suppose a negative endowment shock occurs which shifts the economy from point A to point B. As the probability of default is decreasing in the level of endowment, the interest rate which cor-

Figure 2: painful deleveraging - An Illustration

I. Prior to Shock

II. After Shock
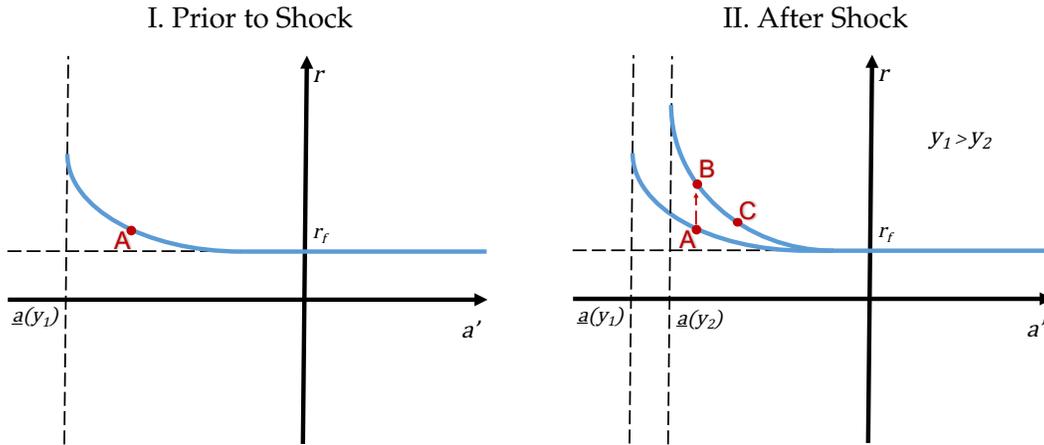
Notes: this figure illustrates why deleveraging is more costly in continuous time. The line in the left panel corresponds to the interest rate faced by a sovereign with endowment y_1, for different levels of assets. The additional line in the second panel corresponds to the interest rate faced by a sovereign with endowment $y_2 < y_1$. In discrete time, the sovereign can jump from point B to C, while in continuous time the sovereign moves continuously from B to C.

responds to point B is higher than in point A. Suppose both economies wish to deleverage so that in one period they arrive to point C. Since point C features a lower level of debt, it also corresponds to a lower interest rate. Let's denote these interest rates by $r_C < r_B$. In the DTE, the interest rate immediately jumps to r_C as it depends on the probability of default in the following period. In the CTE, assets move continuously from point B to point C. So the average interest rate paid by the sovereign in the CTE is greater than r_C.

Note that the sovereign can choose to move assets from point A to point B in a short amount of time by setting consumption close to 0, avoiding some of the higher interest rate burden. However, concavity of the utility function incentivizes it to smooth consumption over time. To conclude, a negative endowment shock of the same magnitude is more welfare deteriorating in the CTE when the sovereign is in debt.

Painful deleveraging induces the sovereign to default in the continuous-time framework at regions where it would choose to repay its debt under discrete time. This difference occurs only when the endowment of the sovereign is high. The endowment level of the sovereign needs to be high enough so that in discrete time it would choose to decrease its debt level and avoid the high spread. We can see this in Figure 3 which plots the default region in the two economies. The default behavior of the sovereign in both economies is very similar for lower levels of endowment, but for higher levels of endowment the default region is larger in the continuous-time environment.

Intuitively, one may suspect that the lower threshold for default at high endowment levels in

Figure 3: Default Region
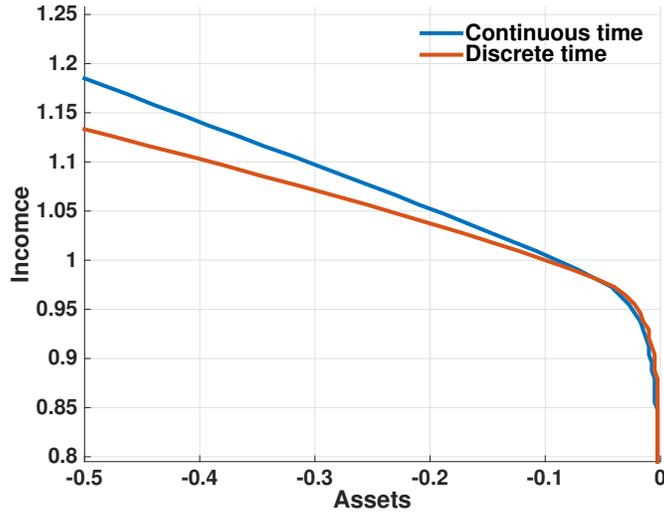
Notes: this figure depicts the default frontier of both economies. The sovereign chooses to default whenever its state is below the default frontier.

the continuous time model leads to overall more defaults in continuous time. This assertion is wrong. Since reaching areas of higher spread is costly for the sovereign, it tries to avoid these regions. So while, in discrete time, a high endowment is often associated with a large negative trade balance, it is less so the case in continuous time. This can be seen in the right panel of Figure 4, which plots the savings decision of the sovereign in the continuous- and discrete-time models, for three levels of endowments.[14] The left and center panels include the savings decisions of a sovereign with 20^{th}-quantile and median income, respectively. The similarity of the policy functions at those levels of endowment between the discrete- and continuous-time economies is striking. There is almost no difference between the two. The right panel displays the savings decision for a sovereign with the 80^{th} quantile endowment. Here we see that when endowment is high, the painful deleveraging motive kicks in. In discrete time, the sovereign chooses to decrease its assets and face a higher spread. In continuous time, the fear of painful deleveraging induces the sovereign to borrow less. So the trade balance is less counter-cyclical.

The cautious behavior of the continuous-time sovereign is apparent in its lower average debt-to-output ratio. This behavior is also the reason for the lower average spread in the continuous-time environment. Despite the attempt of the sovereign to avoid high spreads, negative endowment shocks do occasionally cause spreads to rise substantially. Such shocks are the cause for the higher coefficient of variation of spreads, the ratio between the volatility of the spread to its mean.

Aguiar et al. (2016) discuss the difficulty of discrete time quantitative models of sovereign debt

---

[14]The quarterly savings of the sovereign in continuous time is defined to be the change in assets that would occur in one quarter if no endowment shock hits the economy.

Figure 4: Policy Functions - Savings
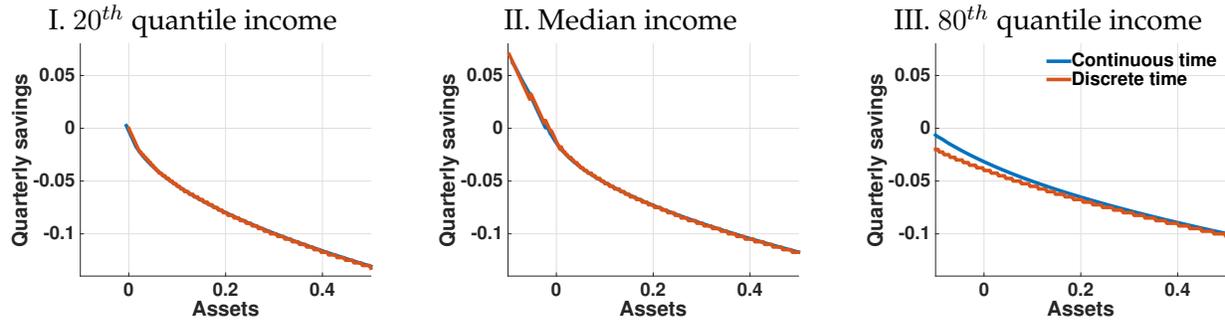
Notes: this figure includes the savings policy function of the sovereign in the continuous and discrete time economies. Three levels of endowment are considered: the 20th quantile (left panel), median income (center panel), 80th quantile (right panel).

to simultaneously match the low probability of default in the data as well as the high volatility of spreads. The continuous-time environment helps overcome this challenge for the same reason deleveraging is costly. High spreads which follow negative endowment shocks cannot be avoided by immediately cutting the stock of debt.[15]

4.3 Speed of Convergence and Comparative Advantage of the Continuous-Time Model

The final row of Table 1 presents the time which takes the algorithm to find the value functions for the different grids considered.[16] Convergence time for the benchmark grid is 4 times faster using the continuous-time methodology. The coarse grid is solved three times faster and the fine grid is solved more than fifty times faster using the continuous-time method. For fine grids, the performance of Matlab in solving the discrete-time model is much worse than high-performance programming languages such as Julia (it takes more than 9 hours on Matlab as opposed to 30 minutes on Julia). It is important to note that the continuous-time method solved on Matlab is also about 2 to 4 times faster when comparing it to the discrete-time model solved on Julia.[17] These differences are substantial for estimation purposes which require solving the model thousands of times.

The larger grid size provides insight into the comparative advantage of the continuous-time methodology. The time consuming operation in the numerical methodology is the inversion of

---

[15]Other papers that proposed alternative mechanisms to capture the higher coefficient of variation of spreads observed in the data include Pouzo and Presno (2016), Martinez and Hatchondo (2017), Passadore and Xu (2018), Bocola and Dovis (2019).

[16]The models were solved on a MacBook Pro (Retina, 15-inch, Mid 2014) with a 2.8 GHz Intel Core i7 processor and 16 GB 1600 MHz DDR3 memory. Both models were solved on Matlab (2014b).

[17]The convergence times of the discrete-time method for the five grid specifications solved on Julia, in the order in which they appear in Table 1, are 1':19", 9", 10':17", 3':42", and 30':35".

large sparse matrices. More than the size of the matrix itself, the foremost factor in the time it takes Matlab to inverse the sparse matrix is the number of non-zero elements (nnz). Consider a grid with J asset points and I endowment points. The number of non-zero elements in the large sparse matrix which Matlab inverts is approximately equal to

$$nnz = I^2 J + 2IJ . \tag{11}$$

Consider a given row in the sparse matrix which corresponds to a state (a_j, y_i). This row has non-zero elements in (a_{j-1}, y_i), (a_j, y_i), and (a_{j+1}, y_i), which are used in the computation of the numerical backward and forward derivatives with respect to assets. In addition, it has an additional I non-zero elements which correspond to states the sovereign can jump to due to an endowment shock. So in every row corresponding to states in which the sovereign has access to international financial markets the matrix has $I + 2$ elements. There are IJ rows in the matrix, so the maximal number of non-zero elements is $I^2 J + 2I$. Notice this is an upper bar on the number of non-zero elements since, as I explain in the appendix, states of default are not considered in the large sparse matrix.

This implies that in continuous time, adding more points to the endowment grid is much more time consuming than adding points to the asset grid. As seen in Table 1, the time it takes to converge for the grid with many asset point (A-fine) is about 15 times faster in continuous time than discrete time. For the endowment fine grid (Y-fine), on the other hand, the computing time is only 2 times faster.

Some of the non-zero elements in the sparse matrix correspond to very low probability events. Consider for example an endowment shock that moves endowment from its highest level to its lowest level. The probability for such an event is close to zero, but not equal to it. In order to increase the speed of the algorithm, we can manually set such low probability events to zero and rescale the transition matrix so its rows sum to 1. For the continuous time model presented in Table 1, I set endowment shocks with probability of less than 0.01% to 0.[18]

Table 2 presents business cycle statistics and computation times for four different thresholds of low probability events: 1%, 0.01%, 10^{-6}% and 0%.[19] All computations are done on the benchmark grid. We can see that business cycle statistics computed using a threshold less than or equal to 0.01% are all very similar. Computations time, on the other hand, are substantially different. Setting the threshold to 10^{-6}% reduces computation time by 20%. Setting it to 0.01%, the threshold level used in Table 1, reduces it by an additional 35%. When setting the threshold level to 1% the business cycle statistics are slightly different. In particular, the average spread is lower. Large negative endowment shocks with probability less than 1%, which move the sovereign into default, are quantitatively important for the level of the spread. So eliminating such events decreases the average equilibrium spread.

---

[18]This is the probability conditional on the arrival of an endowment shock. i.e., $f(i', i) < 0.01\%$ is set manually to 0.
[19]0% implies the transition probability matrix is unchanged.

Table 2: Speeding Computation by Eliminating Low Probability Events

Threshold probability	0%	0.000001%	0.01%	1%
$corr(c, y)$	0.99	0.99	0.99	0.99
$corr(TB/Y, y)$	-0.18	-0.18	-0.18	-0.18
$corr(r_s, y)$	-0.31	-0.31	-0.31	-0.19
$corr(r_s, TB/Y)$	0.71	0.71	0.71	0.70
$\sigma(y)$	5.28	5.28	5.27	4.98
$\sigma(c)$	5.47	5.47	5.46	5.19
$\sigma(TB/Y)$	0.72	0.72	0.72	0.80
$\sigma(r_s)$	4.75	4.75	4.76	5.48
$\mu(r_s)$	2.62	2.62	2.60	2.02
$\sigma(r_s)/\mu(r_s)$	1.86	1.86	1.88	2.78
Average debt-to-output (%)	3.65	3.65	3.65	3.90
Defaults per 500,000 quarters	2,258	2,258	2,254	2,291
Time to converge	45″	35″	23″	18″

The number of non-zero elements in the sparse matrix can also be used to assess how costly it is to add additional state variables. Suppose, for example, that the endowment of the sovereign depends on two exogenous state variables. Let I_1 denote the number of grid points for the first state variable, and assume that it follows a compound Poisson process so that the transition between any two income points on the grid is possible. Let I_2 denote the number of grid points for the second state variable.

If the second state variable also follows a compound Poisson process, the number of non-zero elements is equal to $I_1^2 I_2 J + (I_2 - 1)I_1 I_2 J + 2I_1 I_2 J$. The first two terms correspond to shocks that move the first and second state variables, respectively. The third term corresponds to the sovereign's savings decision. If, instead, the second exogenous state variable follows a Brownian motion, a process that moves continuously over time, then the matrix becomes much more sparse. The second term goes down from $(I_2 - 1)I_1 I_2 J$ to $2I_1 I_2 J$. This is because we only need to consider movement across adjacent grid points. A researcher can assess how computationally consuming it is to add additional state variable, by calculating the relative increase in the number of non-zero elements in the spare matrix.

4.4 Kolmogorov Forward Statistics vs. Simulation Statistics

The business cycle statistics in Table 1 are based on simulating the model for a long time period. In section 3.2, I presented a method to find the ergodic distribution of the economy using the Kolmogorov Forwards equations. This section compares business cycle statistics computed using the ergodic distribution of the model, to statistics computed using simulating the model and then averaging observations across quarters.

Table 3 presents business cycle statistics of the sovereign, conditional on having access to international financial markets. The first column presents the ergodic distribution of the model com-

puted using the Kolmogorov Forward equations. The last two columns present statistics computed using a simulation at a frequency of a business day and then averaged across quarters. Column two includes statistics computed using all periods in which the sovereign had access to international financial markets. Column three corresponds to the statistics reported in Table 1. It reports the average business cycle statistics computed across periods of length 74 quarters prior to the default of the sovereign. So the difference between columns one and two are the result of averaging observation across quarters.[20] The difference between columns two and three is that column three limits the sample to periods before default as well as computing statistics across small samples.

The first substantial difference between the three computation methods is the volatility of output. Since output is quite persistent, computing it across 74 quarters leads to a small sample bias. As output and consumption are highly correlated this implies the volatility of both output and consumption is lower using the third method (column three). The second notable difference is the much higher volatility of the spread when using the Kolmogorov Forward method. This difference implies there are many high frequency changes in the spread that are smoothed out when averaging across quarters. In addition, the average spread is higher using the Kolmogorov Forward equation. This is a mechanical result. If the sovereign had no access to financial markets during any time in the quarter, that quarter is flagged as a no-access quarter. This leads to the exclusion of observations that occur at the same quarter of a default, prior to the default, in columns two and three. In addition to explaining the higher spreads, this exclusion also reduces the volatility of the trade balance in the quarterly simulation.

5 Two Extensions

In this section, I present two extensions of the benchmark model. First, I relax the assumption that debt maturity is instantaneous. Instead, I assume the government issues long term debt contracts that mature probabilistically and calibrate the model according to Chatterjee and Eyigungor (2012). Business cycle statistics of the model are similar to its discrete-time counterpart. The long term maturity of the bond makes deleveraging less costly compared to the benchmark model, as repurchasing debt in regions of high spreads is cheap. I find that the optimal debt maturity in the continuous-time model is 1.2 quarters on average. In comparison, Chatterjee and Eyigungor (2012) finds that in the discrete-time model the optimal debt maturity is the short-term one, a one-quarter bond.

---

[20] A quarter in which the sovereign has access to international financial markets is considered such if the sovereign has access to these markets throughout the quarter. So another difference between columns one and two is that column two "leaves out" observations before default if they are at the same quarter as a default. It also does not include observations after regaining access to financial markets, at the same quarter in which the sovereign regained access to these markets.

Table 3: Different methods for computing statistics

| | Kolmogorov Forward | Simulation + averaging across quarters | |
	Ergodic distribution	Ergodic distribution	74 Q before default
$corr(c,y)$	0.99	1.00	0.99
$corr(TB/Y,y)$	-0.11	-0.08	-0.18
$corr(r_s,y)$	-0.12	-0.22	-0.31
$corr(r_s,TB/Y)$	0.61	0.66	0.71
$\sigma(y)$	7.56	7.33	5.27
$\sigma(c)$	7.71	7.42	5.46
$\sigma(TB/Y)$	0.90	0.68	0.72
$\sigma(r_s)$	13.32	4.82	4.76
$\mu(r_s)$	3.14	2.54	2.60
$\sigma(r_s)/\mu(r_s)$	4.24	1.90	1.88
Average debt-to-output (%)	3.03	3.14	3.65

The second extension of the model relaxes the assumption that the world interest rate is constant and assumes it follows a stochastic process instead. I find that under the calibrated process for the risk-free rate, business cycle statistics are very similar to the benchmark model. This is not because the sovereign does not respond to the risk-free rate. Rather, it is because the calibrated volatility of the interest rate is small relative to the volatility of endowment. In the model, interest rate fluctuations are not responsible for any of the sovereign's defaults in equilibrium.

The main purpose of this section is to illustrate that the continuous-time sovereign debt model can accommodate a large variety of modifications to the benchmark model. In addition, there is a computational advantage in solving the model with long-term debt. While in discrete time, finding the pricing schedule is a daunting task which requires solving a fixed point problem, in continuous time, computing the bond price schedule amounts to solving a sparse system of linear equations.

5.1 Long Term Debt

This section presents a version of the benchmark model which includes long-term duration bonds instead of short term ones. The model is a continuous-time version of Chatterjee and Eyigungor (2012).

Assume the sovereign can save and borrow using long-term debt contracts that mature probabilistically. A bond matures with Poisson intensity λ_b and pays a flow payment z until it matures. Thus, a bond is characterized by the pair (z, λ_b) and its price does not depend on the origination date. Denote by a_t the number of long term bonds held by the sovereign at time t, so that $-a_t$ is the number of outstanding bonds of the sovereign. I assume that bonds can be infinitesimally small so that in a time interval dt, a fraction $\lambda_b dt$ of the outstanding bonds matures with certainty. I consider a Markov perfect equilibrium so that the equilibrium price of a bond in period t, $q_t(\cdot)$,

23

depends on the current state of the sovereign, (y_t, a_t). The sovereign flow budget constraint is given by

$$c_t dt = y_t dt + z a_t dt + \lambda_b a_t dt - q(a_t, y_t)(\lambda_b a_t dt + da_t) .$$

Rearranging, I get

$$s_t \equiv \dot{a}_t = \frac{y_t + (z + \lambda_b)a_t - c_t}{q(y_t, a_t)} - \lambda_b a_t .$$

Consider the expected return on a long term sovereign bond in the region of no-default,

$$\mathbb{E}\left[dr(a, y)\right] = \frac{z + \lambda_b(1 - q(a, y)) + \lambda_y \int_0^\infty \left[q(a, y') - q(a, y)\right] dF(y'|y) + s(a, y)q_a(a, y)}{q(a, y)} dt, \quad (12)$$

where $q_a(b, y)$ is the derivative of $q(a, y)$ with respect to bond holdings a. The first two terms in the denominator are, respectively, the flow income of holding the bond, and the gain due to a share of the bonds maturing. The third term corresponds to changes in the bond price due to endowment shocks. Note that the risk of default is taken into account in this term. In the region of default the price of the bond is equal to zero. So endowment shocks that move the sovereign into default appear in equation (12) as a loss of the current price of the bond.

Finally, the last term of equation (12) corresponds to change in the bond price due to the savings behavior of the sovereign. Intuitively, the price of the bond is weakly increasing in the level of assets of the sovereign, a. The sovereign has more incentive to default if its debt is higher ($a \downarrow$), so the price of the bond is decreasing in the sovereign's debt ($-a$). Zero profit condition on sovereign bond holdings implies that in the region of no-default,

$$q(a, y)r_f = z + \lambda_b(1 - q(a, y)) + \lambda_y \int_0^\infty \left[q(a, y') - q(a, y)\right] dF(y'|y) + s(a, y)q_a(a, y) . \quad (13)$$

The main difference between long-term and short-term debt models is the possibility of *debt dilution*. Additional issued debt reduces the value of current outstanding debt obligations. So, as opposed to short-term debt pricing, the price of long-term maturity bonds depends on the savings behavior of the sovereign. Even if we know the savings behavior of the sovereign, finding the bond pricing schedule amounts to solving a fixed point problem. In principal, this dependence makes the computational problem harder. In continuous time, however, we can use a finite difference method so that finding the bond price schedule given the savings behavior amounts to a single sparse matrix inversion. This is explained in details in the appendix.

In comparison to the benchmark model, deleveraging is less costly with long term debt. Consider the law of motion for assets in both models,

$$\overbrace{\dot{a}_t = a_t r(a_t, y_t) + (y_t - c_t)}^{\text{Benchm. model}}, \qquad \overbrace{\dot{a}_t = a_t \left(\frac{z + \lambda_b}{q(a_t, y_t)} - \lambda_b\right) + \frac{1}{q(a_t, y_t)}(y_t - c_t)}^{\text{Long-term debt model}} .$$

The first term in both models corresponds to debt service payments, while in long-term debt this term also includes the maturing of debt. When the sovereign has outstanding bond obligations ($a < 0$), the level of consumption needed to sustain a certain level of savings is decreasing in

24

the level of the spread. In other words, the higher the spread, the more the sovereign has to pay on its debt. Since debt moves continuously over time, the sovereign cannot instantaneously escape high spread regions by deleveraging. The difference lies in the second term. This is the change in debt position due to the trade balance. A trade surplus ($y - c > 0$) leads to an increase in a, i.e., a decrease in the sovereign's debt level. We can see that with long term debt, this term is divided by the price of debt. In regions of high spreads, the price of sovereign debt is low. This implies the sovereign's cost of buying back its debt is cheaper. i.e., for the same level of trade surplus the sovereign is able to deleverage faster when debt does not mature instantaneously. Therefore, the painful deleveraging motive is weaker with long-term debt. This welfare benefit of long term debt is also present in the discrete time environment and has been studied in Arellano and Ramanarayanan (2012), which refers to this benefit as the hedging benefit of long-term debt. In comparison to discrete time models of sovereign debt, the hedging benefit is higher in continuous time since it ameliorates the deleveraging process following a negative endowment shock.

I follow the calibration of Chatterjee and Eyigungor (2012). The Poisson intensity of bond maturity is set to 0.05, so that on average a bond matures once every 20 quarters. The flow payment, z, is set to 0.03. The interest rate is set to 1% and the coefficient of relative risk aversion to 2. Endowment follows a similar process to the benchmark model with $\rho_y = 0.948$ and $\sigma_y = 0.027$. Finally the discount factor parameter, ρ, is set to 0.047, so it corresponds to a discrete time discount factor of 0.954.

Table 4 presents business cycle statistics of the continuous-time and discrete-time models.[21] Once again, the qualitative and quantitative results of the two models are remarkably similar. In contrast to the benchmark model, the average spread and its volatility are higher in continuous time. Since the painful deleveraging motive is weaker, the sovereign is less cautious in staying away from regions of high spreads. While with short-term debt the average debt-to-output ratio is lower in the continuous-time environment, with long-term debt the average debt-to-output ratio in the continuous-time framework is similar to the discrete-time one. Nevertheless, when the sovereign suffers from a negative endowment shock that moves it to a region with high spreads, it cannot instantaneously escape it. So the average spread, as well as the volatility of spreads, in the continuous-time environment are higher.

Table 5 compares the business cycle statistics and convergence times of the continuous time model for a variety of grids. The differences between having 25, 51, or 201 endowment grid points are very small. The only notable difference between those grids is the volatility of the spread. It declines as the number of endowment grid points increases.[22] In the benchmark model, adding asset grid points is relatively faster than adding endowment grid points. With long term debt,

---

[21]Following Chatterjee and Eyigungor (2012), I compute business cycle statistics by simulating the model for 5,000 quarters, excluding quarters in which the sovereign has no access to financial markets together with the first 20 quarters after regaining access to markets. I repeat this simulation 300 times and report the average statistics.

[22]Hatchondo et al. (2010) show a similar relationship between the number of grid points and the volatility of the spread for discrete time short term debt models solved on an equidistant grid.

Table 4: Business Cycle Statistics with Long-Term Debt

	Continuous time model	Chatterjee and Eyigungor (2012)
$corr(c,y)$	0.99	0.99
$corr(TB/Y,y)$	-0.47	-0.44
$corr(r_s,y)$	-0.75	-0.65
$\sigma(c)/\sigma(y)$	1.09	1.11
$\sigma(TB/Y)/\sigma(y)$	0.17	0.17
$\sigma(r_s)$	5.11	4.43
$\mu(r_s)$	10.22	8.15
Average debt-to-output (%)	70	70
Average debt service payments	5.60	5.55

Notes: continuous time model solved on a grid of 401 asset points equally distributed between -2 and 0, and 25 endowment grid points.

this is no longer the case. Endowment grid points still make the matrix less sparse than asset grid points. However, in the long-term debt model, a dense asset grid can lead to oscillations in the default frontier. The cause lies in the optimality condition for consumption,

$$\overbrace{u'(c(a,y)) = v_a(a,y)}^{\text{Benchm. model}}, \qquad \overbrace{u'(c(a,y)) = v_a(a,y)/q(a,y)}^{\text{Long-term debt model}} .$$

The state boundary constraint therefore depends not only on the level of assets and endowment, but also on the endogenous price of the bond at that point. This dependence makes the algorithm less stable as it is prone to oscillations in the default frontier. The last two columns of table 5 compare business cycle statistics between a 401 and 1201 asset grid points, fixing the endowment grid to 7 points. We can see that there is very little difference between the two columns. This suggests we do not lose accuracy by restricting the benchmark grid to 401 points.

Table 5: Business Cycle Statistics and Convergence Times for Different Grids

Asset grid points	401	401	401	401	1201
Endowment grid points	25	51	201	7	7
$corr(c,y)$	0.99	0.99	0.99	0.99	0.99
$corr(TB/Y,y)$	-0.47	-0.46	-0.46	-0.29	-0.29
$corr(r_s,y)$	-0.75	-0.76	-0.75	-0.61	-0.61
$\sigma(c)/\sigma(y)$	1.09	1.09	1.09	1.06	1.06
$\sigma(TB/Y)/\sigma(y)$	0.17	0.17	0.17	0.15	0.15
$\sigma(r_s)$	5.11	4.75	4.60	5.05	5.04
$\mu(r_s)$	10.22	10.21	10.11	8.58	8.61
Average debt-to-output (%)	70	69	69	84	84
Average debt service payments	5.60	5.54	5.54	6.73	6.70
Time to converge	6':14"	12':13"	61':50"	4"	35"

There are pros and cons for having a longer maturity length of sovereign bonds. As Arellano

and Ramanarayanan (2012) explain, longer maturity length is welfare-enhancing as its endogenous price acts as a hedge against bad shocks. On the other hand, since the government cannot commit to future debt issuance decisions, the endogenous price of long-term bonds makes it a worse instrument to raise funds. Arellano and Ramanarayanan (2012) refer to this welfare benefit of short-term duration bonds relative to long-term ones as the incentive benefit, as short-term provides the sovereign a higher incentive to repay its debt. Chatterjee and Eyigungor (2012) compare the welfare of a sovereign with no debt obligations, and find that the sovereign is better off with short-term debt than with long-term debt. In continuous time, as explained above, there is an additional welfare gain in having long-term maturity bonds. Deleveraging is less costly with long-term debt. Figure 5 displays the level of consumption that makes $c^{1-\gamma}/\rho$ equal to $\sum_y v(0, y)\pi(y)$, where $\pi(y)$ is the ergodic distribution of endowment. Compared to Chatterjee and Eyigungor (2012), the maximum level of welfare is not obtained with short term maturity bonds, which corresponds to $\lambda_b \to \infty$.[23] The welfare-maximizing value of λ_b is 0.85. This value leads to an additional 0.12% certainty consumption equivalence in comparison to short term debt, and to 0.85% increase compared to the calibrated level of λ_b, 0.05. The optimal value of λ_b corresponds to an average maturity length of 1.2 quarters. It is similar to the average maturity length which Chatterjee and Eyigungor (2012) finds to be optimal, a one-quarter bond. This finding shows that the incentive benefit is not always greater than the hedging benefit. In other words, *shorter is not always better*.

**Discussion of application to long-term debt models with two endogenous state variables.** The Chatterjee and Eyigungor (2012) model features a single endogenous state variable, the number of long term bonds held by the sovereign. The maturity length of the sovereign bond is taken as exogenous. There are other papers that study the optimal maturity length of sovereign debt in discrete time, but allow the sovereign to choose the maturity length in every period. In Arellano and Ramanarayanan (2012), the sovereign can issue debt using a combination of two bonds, a short-term and a long-term maturity one. The composition of the two imply that the average maturity length is endogenous. In Sánchez et al. (2018), the sovereign chooses in every period the amount of its outstanding debt obligations as well as the amount of time in which these obligations are to be repaid.

Solving a model with two endogenous state variables for the sovereign's debt position poses a challenge to the finite difference method. The approach relies on the continuity of the endogenous state variables over time. That is, endogenous state variables should not feature jumps. In models with a single debt variable, as the ones studied in this paper, this is not a problem. Any finite choice of the consumption flow implies that assets move continuously over time. However, with two choices regarding the sovereign's debt position, this is no longer the case. For example, in Arellano and Ramanarayanan (2012), the sovereign may decide to change its debt portfolio discontinuously

---

[23]The long term debt model was also solved for $\lambda_b = 10^6$, in which case the certainty consumption equivalence is approximately equal to the one of short term debt.

Figure 5: Welfare and Bond Maturity Length
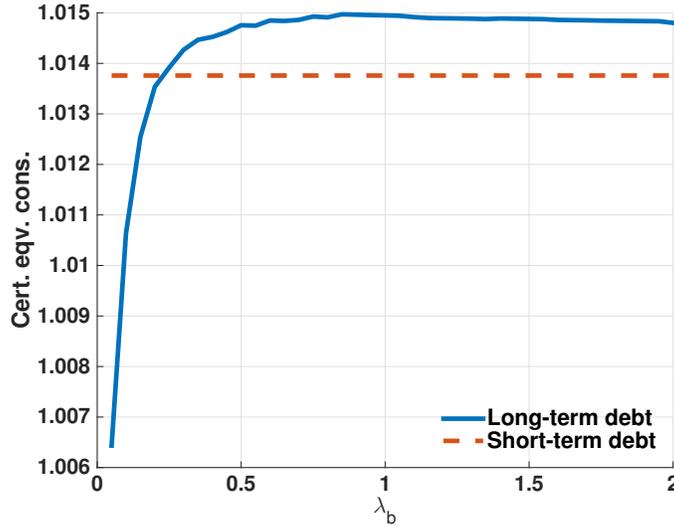
Notes: this figure displays the certainty consumption equivalence of a sovereign with no debt for different maturity lengths. The dashed line is the certainty consumption equivalence when debt is of short maturity. i.e., when $\lambda_b \to \infty$.

in response to an endowment shock. In order to apply the methodology presented in this paper to models with two endogenous state variables, their continuous-time versions need to be modified so that their endogenous state variables move continuously over time.

A simple modification to two endogenous state variables model that ensures their continuity over time is the introduction of adjustment costs. In Arellano and Ramanarayanan (2012), one can assume that changes in the amount of long-term debt obligations are subject to adjustment costs. In Sánchez et al. (2018), one can assume that changes to the amount of time left to repay debt obligations is subject to adjustment costs. These assumptions ensure that all endogenous state variables move continuously over time, and that the numerical methodology presented in this paper can be used to solve them.[24]

5.2 World Interest Rate Fluctuations

Consider again the case in which the government only has access to short term maturity bonds as in the benchmark model. Assume that instead of a constant world interest rate, r_f, the world interest rate is stochastic. In particular, I assume the world interest rate follows a compound Poisson process. The arrival rate of the shock is λ_r and upon the realization of the shock, a new

---

[24]Note that the limitation of the finite difference method in dealing with discontinuous endogenous state variables is not due to the endogenous default decision studied in this paper, but is more general. Kaplan et al. (2016), for example, solve a continuous-time heterogeneous-agent model with two types of assets, a liquid and an illiquid one, using a finite difference method. The form of adjustment costs they assume ensures that both assets move continuously over time.

world interest rate is drawn from the distribution $F_r(r^f, r_-^f)$. The HJB equations corresponding to this case are relegated to the appendix.

I assume that upon realization of an interest rate shock, the new interest rate is set according to the following law of motion.

$$r^f = (1 - \rho_r)\bar{r}^f + \rho_r r_-^f + \epsilon_r \,,$$

where ϵ_r is iid normally distributed with mean 0 and variance σ_r^2, and $\bar{r}^f$ is the unconditional mean of the process. I maintain the parameter values calibrated in the previous section which correspond to Arellano (2008). Arellano calibrates the risk free interest rate to the average quarterly yield on 5-year US treasury bonds between 1983Q3-2001Q4, setting it to 1.7%. I set the arrival rate of an interest rate shock, λ_r, the persistence parameter, ρ_r, and the variance of the interest rate shock, σ_r^2, so it matches the quarterly variance, quarterly autocorrelation, and kurtosis of the quarterly change in 5-year US treasury bond yield during this period. This amounts to setting $\lambda_r = 0.96$ so that an interest rate shock arrives on average approximately once a quarter. The persistence parameter is calibrated to 0.961, and the volatility of the iid shock is 13 basis points, i.e., $\sigma_r = 0.0013$.[25]

Table 6 displays the business cycle statistics of the model with world interest rate fluctuations in comparison to the benchmark model. Under our calibration, the business cycle statistics of the two models are very similar. Even though, theoretically, defaults can now occur either due to an endowment shock or to a world interest rate shock, approximately zero of the defaults under this calibration are due to world interest rate shocks. Interest rate shocks are not responsible for any of the sovereign's default as they are small in size and quite frequent. The sovereign does respond to interest rate fluctuations, as I show below. However, the volatility of the interest rate is so low with comparison to the volatility of output that it does not substantially change the business cycle statistics of the model.

Figure 6 shows the savings decision for three different levels of income and three different levels of world interest rate. A higher interest rate incentivizes the sovereign to increase its assets and save more. However, the middle panel shows that a sovereign with median income and debt (negative assets) is not responsive to the risk-free rate. This is due to the high spreads that characterize that region. The sovereign responds to the interest rate it faces, which is the sum of the risk-free rate and the spread. If the spread is substantially higher than the risk-free rate, changes to the risk free rate do not change the behavior of the sovereign much.

---

[25]Table 7 in the appendix shows this calibration matches well also higher frequency moments which were not targeted. In particular, it does a good job at matching the volatility, kurtosis, autocorrelation, and the volatility of changes at a daily, weekly, and monthly frequency. It overestimates the kurtosis of changes at lower frequencies. The latter is likely due to high frequency shocks that are small in size, which the calibrated process does not capture.

Table 6: Business Cycle Statistics with World Interest Rate Fluctuations

	Benchm. model	WIR fluctuations
$corr(c, y)$	0.99	0.99
$corr(TB/Y, y)$	-0.23	-0.23
$corr(r_s, y)$	-0.32	-0.33
$corr(r_s, TB/Y)$	0.71	0.70
$\sigma(y)$	5.20	5.26
$\sigma(c)$	5.45	5.51
$\sigma(TB/Y)$	0.80	0.80
$\sigma(r_s)$	5.88	5.74
$\mu(r_s)$	2.63	2.59
$\sigma(r_s)/\mu(r_s)$	2.30	2.29
Average debt-to-output (%)	4.67	4.59
Defaults per 500,000 quarters	2,024	2,004
Time to converge	4″	2′:34″

Notes: all computations are performed on a grid with 301 asset points, 25 endowment points, and 11 world interest rate points.

Figure 6: Savings with World Interest Rate Fluctuations
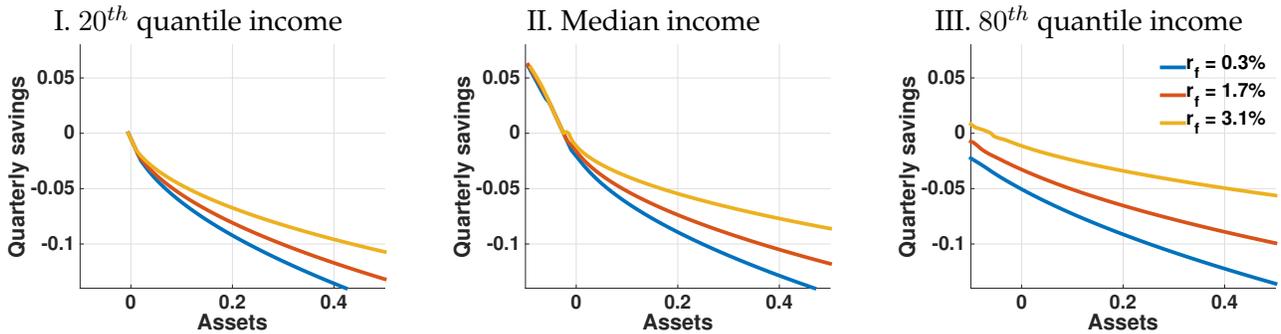
Notes: this figure presents the savings policy function of the sovereign for different world interest rates and different endowment levels.

6 Conclusion

In this paper, I develop a continuous time model of sovereign debt and provide a numerical method for solving it. I show that the quantitative results are similar between the model and its discrete-time counterpart. The few differences between the two, most notably the lower average spread in the continuous-time model, are a result of the different spreads faced by the sovereign along the deleveraging process. In continuous time, deleveraging is more costly.

Two extensions of the benchmark model are considered. The first extension includes long term debt. I show that, as opposed to a discrete-time environment, computing the pricing schedule of

the long-term debt in continuous time is quite simple. In continuous time, longer maturity of debt makes deleveraging less costly in comparison with the short-term debt model. I calibrate the model according to Chatterjee and Eyigungor (2012) and find that the business cycle statistics in the discrete- and continuous-time models are very similar. In addition, I find that the optimal average bond maturity is 1.2 quarters. This result is similar to the optimal average debt maturity in Chatterjee and Eyigungor (2012), a one-quarter bond. The second extension includes stochastic world interest rate fluctuations. I find that under the calibrated process for the world interest rate, average business cycle statistics of the model remain unchanged and sovereign defaults are accounted for almost entirely by endowment shocks and not by world interest rate shocks.

The methods presented in this paper lay the groundwork for future quantitative research on sovereign debt and consumer bankruptcy, which were so far infeasible to compute and estimate in a timely manner.

References

ACHDOU, Y., HAN, J., LASRY, J.-M., LIONS, P.-L. and MOLL, B. (2020). Income and wealth distribution in macroeconomics: A continuous-time approach.

AGUIAR, M., AMADOR, M., FARHI, E. and GOPINATH, G. (2015). Coordination and crisis in monetary unions. *Quarterly Journal of Economics*, **130**.

AGUIAR, M., CHATTERJEE, S., COLE, H. and STANGEBYE, Z. (2016). Quantitative models of sovereign debt crises. *Handbook of Macroeconomics*, **2** 1697–1755.

AGUIAR, M. and GOPINATH, G. (2006). Defaultable debt, interest rates and the current account. *Journal of international Economics*, **69** 64–83.

AGUIAR, M. A. and AMADOR, M. (2019). Self-fulfilling debt dilution: Maturity and multiplicity in debt models. Tech. rep., National Bureau of Economic Research.

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

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

BARILLAS, F. and FERNÁNDEZ-VILLAVERDE, J. (2007). A generalization of the endogenous grid method. *Journal of Economic Dynamics and Control*, **31** 2698–2712.

BOCOLA, L. and DOVIS, A. (2019). Self-fulfilling debt crises: A quantitative analysis. *American Economic Review*, **109** 4343–77.

BRUNNERMEIER, M. K. and SANNIKOV, Y. (2014). A macroeconomic model with a financial sector. *The American Economic Review*, **104** 379–421.

CHATTERJEE, S., CORBAE, D., NAKAJIMA, M. and RÍOS-RULL, J.-V. (2007). A quantitative theory of unsecured consumer credit with risk of default. *Econometrica*, **75** 1525–1589.

CHATTERJEE, S. and EYIGUNGOR, B. (2012). Maturity, indebtedness, and default risk. *The American Economic Review*, **102** 2674.

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

EATON, J. and GERSOVITZ, M. (1981). Debt with potential repudiation: Theoretical and empirical analysis. *The Review of Economic Studies*, **48** 289–309.

GABAIX, X., LASRY, J.-M., LIONS, P.-L. and MOLL, B. (2016). The dynamics of inequality. *Econometrica*, **84** 2071–2111.

HATCHONDO, J. C. and MARTINEZ, L. (2009). Long-duration bonds and sovereign defaults. *Journal of international Economics*, **79** 117–125.

HATCHONDO, J. C., MARTINEZ, L. and SAPRIZA, H. (2010). Quantitative properties of sovereign default models: solution methods matter. *Review of Economic Dynamics*, **13** 919–933.

KAPLAN, G., MOLL, B. and VIOLANTE, G. L. (2016). Monetary policy according to hank. Tech. rep., National Bureau of Economic Research.

LIVSHITS, I., MACGEE, J. and TERTILT, M. (2007). Consumer bankruptcy: A fresh start. *The American Economic Review*, **97** 402–418.

LORENZONI, G. and WERNING, I. (2019). Slow moving debt crises. *American Economic Review*, **109** 3229–63.

MARTINEZ, L. and HATCHONDO, J. (2017). Credit risk without commitment. In *2017 Meeting Papers*. 1326, Society for Economic Dynamics.

MITMAN, K. (2016). Macroeconomic effects of bankruptcy and foreclosure policies. *The American Economic Review*, **106** 2219–2255.

NUÑO, G. and THOMAS, C. (2019). Monetary policy and sovereign debt vulnerability.

PASSADORE, J. and XU, Y. (2018). Illiquidity in sovereign debt markets. *Available at SSRN 3195968*.

POUZO, D. and PRESNO, I. (2016). Sovereign default risk and uncertainty premia. *American Economic Journal: Macroeconomics*, **8** 230–66.

REBELO, S., WANG, N. and YANG, J. (2018). Rare disasters, financial development, and sovereign debt. Tech. rep., National Bureau of Economic Research.

SÁNCHEZ, J. M., SAPRIZA, H. and YURDAGUL, E. (2018). Sovereign default and maturity choice. *Journal of Monetary Economics*, **95** 72–85.

TOURRE, F. (2017). A macro-finance approach to sovereign debt spreads and returns.

A   Mathematical Appendix

I'll first prove that the value function is strictly increasing in assets in the region of no-default. This result would immediately validate Proposition 1.

Lemma 1. *The value function in the region of no default is strictly increasing in assets. That is,*

$$v(a_1, y) < v(a_2, y) , \qquad \forall (a_1, a_2, y) : \quad a_1 < a_2 , \quad D(a_1, y) = D(a_2, y) = 0 .$$

Proof. Let c_t^1 denote the optimal consumption path for the sovereign at state (a_1, y) at time $t = 0$. A feasible consumption level for the sovereign with assets (a_2, y) at time $t = 0$ is $c_t^2 = M c_t^1$ for an arbitrarily large M, and specifically $M > 1$. Since assets move continuously over time, after some amount of time, Δ, the assets of the sovereign which started with (a_2, y) would be equal to the assets of the sovereign that started with (a_1, y). The amount of time it takes the two asset levels to converge can be arbitrarily small by choosing a larger M. Assume that at time $t + \Delta$ the sovereign which started with asset level a_2 mimics the behavior of the sovereign which started with a_1. Denote the value function of such behavior as $v^M(a_2, y)$.

At time $t + \Delta$, the expected value of the two sovereigns are equal as their state variables are the same. Between time t and $t + \Delta$, the utility of the sovereign which started with (a_2, y) and follows the strategy above is larger since utility is strictly increasing in consumption. I get that

$$v(a_1, y) \leq v^M(a_2, y) .$$

Finally, note that $v^M(a_2, y) \leq v(a_2, y)$ as following the M-strategy is a feasible path. Therefore, I conclude that

$$v(a_1, y) < v(a_2, y) .$$

$\square$

Proposition 1. *The default decision of the sovereign is of the threshold kind in bond obligations. For every level of endowment y, there exists a level of assets $\underline{a}(y)$ such that*

$$D(a, y) = \begin{cases} 1, & \text{if} \quad a < \underline{a}(y) , \\ 0, & \text{if} \quad a \geq \underline{a}(y) . \end{cases}$$

Proof. Suppose a sovereign at state (a_1, y) chooses to repay its debt. It must be the case then that

$$v(a_1, y) \geq w(y) .$$

Now consider the problem of a sovereign with the same endowment level but higher level of assets, $a_2 > a_1$. Lemma 1 implies that

$$v(a_2, y) > v(a_1, y) ,$$

so that

$$v(a_2, y) > w(y) \, .$$

That is, if a sovereign chooses not to default, all sovereigns with the same endowment level and more assets would also choose not to default. □

B Numerical Method using Sparse Matrices

This section shows how to implement the numerical method using sparse matrices. For the most part, it follows the numerical appendix of Achdou et al. (2020). There are two differences between their method and the one presented here. First, the interest rate schedule is endogenous. Second, default adds another step to the procedure, which I discuss in details below.

Each iteration involves three steps:

1. Computing the interest rate schedule given the default decision.

2. Constructing the sparse matrix used to compute the HJB equations, updating the guess for v, and then for w.

3. Updating the default decision, and check for convergence of value functions.

Updating the interest rate schedule is described in the main text. Here I shall elaborate on how to do the second step efficiently using sparse matrices.

B.1 HJB Equations in Matrix Form

Suppose we are currently at iteration n of the value functions, and the value functions have not yet converged. We want to update the value functions v^{n+1} and w^{n+1}. We do so using the implicit method approach. The value function of the sovereign in the no-default region is implicitly defined by the following equation.

$$\frac{v_{i,j}^{n+1} - v_{i,j}^n}{\Delta} + \rho v_{i,j}^{n+1} = u\left(c_{i,j}^n\right) + \partial_a v_{i,j}^{n+1}\left(y_i - c_{i,j}^n + r_{i,j}^n a_j\right) + \lambda_y \sum_{i'=1}^{I}\left(v_{i',j}^{n+1} - v_{i,j}^{n+1}\right)f(i',i), \text{ for } j \geq \underline{d}_i^n,$$

where $\underline{d}_i^n$ is the default frontier from the previous iteration of the value function. Note the $n+1$ superscripts in the right-hand side of the equation.

Recall that we use a finite difference method approach with an upwinding scheme to obtain the policy functions. The forward and backward difference approximations $\partial_a v_{i,j,F}$ and $\partial_a v_{i,j,B}$ are defined as:

$$\partial_a v_{i,j,F} = \frac{v_{i,j+1} - v_{i,j}}{da}, \qquad \partial_a v_{i,j,B} = \frac{v_{i,j} - v_{i,j-1}}{da},$$

where da is the difference between two adjacent asset grid points, and iteration superscripts are omitted to simplify notation. To obtain the policy functions, we first compute savings according to both the forward and backward difference approximations:

$$s_{i,j,F} = y_i + r_{i,j}a_j - \left(u'\right)^{-1}\left(\partial_a v_{i,j,F}\right), \qquad s_{i,j,B} = y_i + r_{i,j}a_j - \left(u'\right)^{-1}\left(\partial_a v_{i,j,B}\right),$$

We want to use the forward derivative whenever it implies a positive savings behavior and a negative one whenever that implies a negative savings behavior. The concavity of the value function implies that, in general, only one of the two conditions would be true. We use the following

approximation,

$$\partial_a v_{i,j} = \partial_a v_{i,j,F} \mathbb{1}_{s_{i,j,F}>0} + \partial_a v_{i,j,B} \mathbb{1}_{s_{i,j,B}<0} + u'\left(y_i + r_{i,j}a_j\right) \mathbb{1}_{s_{i,j,F}\leq 0 \leq s_{i,j,B}}, \tag{14}$$

where the third term corresponds to the savings policy equal to zero when the forward derivative is negative and the backward positive.

Using the implicit method together with the upwinding scheme, the value function is obtained from the following equation,

$$\frac{v_{i,j}^{n+1} - v_{i,j}^n}{\Delta} + \rho v_{i,j}^{n+1} = u\left(c_{i,j}^n\right) + \partial_a v_{i,j,F}^{n+1}\left[s_{i,j,F}^n\right]^+ + \partial_a v_{i,j,B}^{n+1}\left[s_{i,j,B}^n\right]^- + \lambda_y \sum_{i'=1}^{I}\left(v_{i',j}^{n+1} - v_{i,j}^{n+1}\right) f(i',i),$$
$$\tag{15}$$

where the superscripts $+$ and $-$ indicate the positive and negative parts of the value in parenthesis, respectively. The consumption value from iteration n is given by $c_{i,j}^n = (u')^{-1}\left(\partial_a v_{i,j}^n\right)$.

Equation (15) consists of $J \times I$ linear equations, and it can be written in matrix notation using the following steps. First, use the definition of the forward and backward derivatives to rewrite the equation as follows,

$$\frac{v_{i,j}^{n+1} - v_{i,j}^n}{\Delta} + \rho v_{i,j}^{n+1} = u\left(c_{i,j}^n\right) + \frac{v_{i,j+1}^{n+1} - v_{i,j}^{n+1}}{da}\left[s_{i,j,F}^n\right]^+ + \frac{v_{i,j}^{n+1} - v_{i,j-1}^{n+1}}{da}\left[s_{i,j,B}^n\right]^-$$
$$+ \lambda_y \sum_{i'=1}^{I}\left(v_{i',j}^{n+1} - v_{i,j}^{n+1}\right) f(i',i),$$

Collecting terms with the same subscripts on the right-hand side, we obtain

$$\frac{v_{i,j}^{n+1} - v_{i,j}^n}{\Delta} + \rho v_{i,j}^{n+1} = u\left(c_{i,j}^n\right) + v_{i,j-1}^{n+1} L_{i,j} + v_{i,j}^{n+1} M_{i,j}^n + v_{i,j+1}^{n+1} R_{i,j}^n + \lambda_y \sum_{i'=1}^{I} v_{i',j}^{n+1} f(i',i), \tag{16}$$

where

$$L_{i,j}^n = -\frac{\left(s_{i,j,B}^n\right)^-}{da},$$

$$M_{i,j}^n = \left[\frac{\left(s_{i,j,B}^n\right)^- - \left(s_{i,j,F}^n\right)^+}{da} - \lambda_y\right],$$

$$R_{i,j}^n = \frac{\left(s_{i,j,F}^n\right)^+}{da}.$$

We are almost ready to write the equation in matrix form. Let $\mathbf{v}^n$ be a $JI \times 1$ vector of the value function of iteration n. The first element corresponds to the state (a_1, y_1), the second element to (a_2, y_1), the $J^{th} + 1$ element to (a_1, y_2), etc. Recall that in this step, we only want to find the value functions where $j \geq \underline{d}_i^n$. That is, we want to find it in the no-default region. Let $\mathbf{d}^n$ be a vector of

length D^n which contains the locations of $\mathbf{v}^n$ where the sovereign chooses to default. Denote by $\bar{\mathbf{v}}^{n+1}$ the value function in the no-default region, which is a vector of size $JI - D^n$. Let $\bar{\mathbf{u}}^n$ denote the vector of utilities of the sovereign at different no-default states, based on the consumption levels computed using the backward/forward derivatives of the value function. Equation (16) can be written in matrix form as follows,

$$\frac{\bar{\mathbf{v}}^{n+1} - \bar{\mathbf{v}}^n}{\Delta} + \rho\bar{\mathbf{v}}^n = \bar{\mathbf{u}}^n + \bar{\mathbf{A}}^n\bar{\mathbf{v}}^{n+1} + \lambda_y\bar{\Lambda}^n\bar{\mathbf{v}}^{n+1} \,, \tag{17}$$

where both $\bar{\mathbf{A}}^n$ and $\bar{\Lambda}^n$ are $(JI - D^n) \times (JI - D^n)$ sparse matrices.

Construction of the $\mathbf{A}^n$ matrix. The $\mathbf{A}^n$ matrix is a block diagonal matrix that contains $L_{i,j}^n$, $M_{i,j}^n$, and $R_{i,j}^n$. There are I blocks corresponding to the I endowment levels. Each block is tridiagonal and has $J - D_i^n$ rows and columns, where D_i^n is the number of asset levels in which the sovereign with endowment i chooses to default. Following the notation in section 3.1, $D_i^n = \underline{d}_i^n - 1$.

Let $\mathbf{A}_i^n$ denote the i-th block in the matrix $\mathbf{A}^n$. Its structure is as follows. The center diagonal in block i is given by $M_{i,j}^n$, where the first element is $M_{i,\bar{d}_i^n}$. Similarly, the left diagonal is given by $L_{i,j}^n$, and the right diagonal by $R_{i,j}^n$. That is,

$$\mathbf{A}_i^n = \begin{bmatrix} M_{i,\bar{d}_i^n} & R_{i,\bar{d}_i^n} & 0 & 0 & \cdots & 0 \\ L_{i,\bar{d}_i^n+1} & M_{i,\bar{d}_i^n+1}^n & R_{i,\bar{d}_i^n+1}^n & 0 & \cdots & 0 \\ 0 & L_{i,\bar{d}_i^n+2}^n & M_{i,\bar{d}_i^n+2}^n & R_{i,\bar{d}_i^n+2}^n & \mathbf{0} & 0 \\ 0 & 0 & \ddots & \ddots & \ddots & \mathbf{0} \\ 0 & 0 & \mathbf{0} & L_{i,J}^n & M_{i,J}^n & R_{i,J}^n \end{bmatrix},$$

and the $\mathbf{A}^n$ matrix is given by

$$\mathbf{A}^n = \begin{bmatrix} \mathbf{A}_1^n & 0 & \cdots & 0 \\ 0 & \mathbf{A}_2^n & \mathbf{0} & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \mathbf{0} & \mathbf{A}_I^n \end{bmatrix}$$

Note that $\mathbf{A}^n$ is very sparse, and only has $[3(J - D^n) - 2] I$ non zero elements.

Construction of the $\bar{\Lambda}_n$ matrix. The matrix $\bar{\Lambda}_n$ corresponds to transition due to endowment shocks. It may be useful to first consider the transition matrix in the case of no default. Denote by Π_y the $I \times I$ transition matrix conditional on an endowment shock, which columns sum to 1. Define

$$\Lambda = (\Pi_y - \mathbb{I}_I) \otimes \mathbb{I}_J \,, \tag{18}$$

where $\otimes$ is the Kronecker product, and $\mathbb{I}$ is the identity matrix. Λ is a $JI \times JI$ matrix with I^2J non-zero elements. $\bar{\Lambda}^n$ is a $(JI - D^n) \times (JI - D^n)$ matrix, and is quite similar to Λ. The difference between the two is that in $\bar{\Lambda}^n$, an endowment shock that would move the sovereign into default

moves it instead to default frontier. Recall that the value at the default frontier is equal to the value of default, so that for calculating the value function it is equivalent whether we consider the value w_i when the sovereign defaults or $v_{d_i,i}$. $\bar{\Lambda}^n$ is constructed of I^2 blocks. The number of rows and columns of block $b_{i1,i2}$ are $J - D_{i1}^n$ and $J - D_{i2}^n$, respectively. The last $\min\{J - D_{i1}^n, J - D_{i2}^n\}$ rows and columns of block $b_{i1,i2}$ are a diagonal matrix with $\pi_{i1,i2}^y$ in all the elements on the diagonal. It may be simpler to understand by looking at Figure 7. Note that endowment shocks that move the sovereign into default appear only in the right panel, these are the non-zero elements that are at the first column but not on the diagonal.

Figure 7: Block $b_{i1,i2}$ of the $\bar{\Lambda}^n$ matrix

I. # rows < # columns

$$\begin{bmatrix} \mathbf{0} & \pi_{i1,i2}^y & 0 & \cdots & 0 \\ \mathbf{0} & 0 & \pi_{i1,i2}^y & \mathbf{0} & 0 \\ \mathbf{0} & 0 & 0 & \ddots & 0 \\ \mathbf{0} & 0 & 0 & \mathbf{0} & \pi_{i1,i2}^y \end{bmatrix}$$

II. # rows = # columns

$$\begin{bmatrix} \pi_{i1,i2}^y & 0 & \cdots & 0 \\ 0 & \pi_{i1,i2}^y & \mathbf{0} & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \mathbf{0} & \pi_{i1,i2}^y \end{bmatrix}$$

III. # rows > # columns

$$\begin{bmatrix} \pi_{i1,i2}^y & 0 & \cdots & 0 \\ \vdots & 0 & \cdots & 0 \\ \pi_{i1,i2}^y & 0 & \cdots & 0 \\ 0 & \pi_{i1,i2}^y & \mathbf{0} & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \mathbf{0} & \pi_{i1,i2}^y \end{bmatrix}$$

Finding the value function using sparse matrix inversion. We can rearrange equation (17) to have

$$\bar{\mathbf{v}}^{n+1} = \left[\left(\frac{1}{\Delta} + \rho\right)\mathbb{I} - \bar{\mathbf{A}}^n - \lambda_y\bar{\Lambda}^n\right]^{-1}\left(\frac{1}{\Delta}\bar{\mathbf{v}}^n + \bar{\mathbf{u}}^n\right).$$

So to update the value function we simply need to invert a highly sparsed matrix.

Let us now find the value function when the sovereign has no access to financial markets. Let $\mathbf{v}_0^{n+1}$ be a vector of length I containing the value function when the sovereign has access to financial markets and assets are equal to 0. We can rewrite equation (6) as follows

$$\frac{\mathbf{w}^{n+1} - \mathbf{w}^n}{\Delta} + \rho\mathbf{w}^{n+1} = \mathbf{u}_w + \lambda_y\left(\Pi_y - \mathbb{I}\right)\mathbf{w}^{n+1} + \lambda_D\left(\mathbf{v}_0^{n+1} - \mathbf{w}^{n+1}\right),$$

where u_w is simply given by the flow utility of the endowment of the sovereign, when it has no access to financial markets. Rearranging we get

$$\mathbf{w}^{n+1} = \left[\left(\frac{1}{\Delta} + \rho + \lambda_D\right)\mathbb{I} - \lambda_y\left(\Pi_y - \mathbb{I}\right)\right]^{-1}\left(\frac{1}{\Delta}\mathbf{w}^n + \mathbf{u}_w + \lambda_D\mathbf{v}_0^{n+1}\right).$$

So we immediately obtain $\mathbf{w}^{n+1}$ once we have $\mathbf{v}^{n+1}$. The final step is to update the default decision as described in section 3.1, check whether the value functions converged, and if not repeat the procedure.

B.2 Solving the Kolmogorov Forward equations

Consider the stationary Kolmogorov Forward equations:

$$0 = - \partial_a \left[s(a,y) \bar{\Gamma}^{ND}(a,y) \right] - \lambda_y \bar{\Gamma}^{ND}(a,y) + \lambda_y \int_0^\infty f(y,\tilde{y}) \bar{\Gamma}^{ND}(a,\tilde{y}) \, d\tilde{y} + \lambda_D \bar{\Gamma}^D(y) \mathbb{1}(a=0) \ ,$$

$$0 = - \lambda_y \bar{\Gamma}^D(y) + \lambda_y \int_0^\infty D(a,y) f(y,\tilde{y}) \bar{\Gamma}^{ND}(a,\tilde{y}) \, d\tilde{y} - \lambda_D \bar{\Gamma}^D(y) \ .$$

The corresponding discretized version of the Kolmogorov Forwards equations are given by,

$$0 = - \partial_a \left[s_{i,j} \bar{\Gamma}^{ND}_{i,j} \right] - \lambda_y \bar{\Gamma}^{ND}_{i,j} + \lambda_y \sum_{i'} \pi(y_i, y_{i'}) \bar{\Gamma}^{ND}_{i',j} + \lambda_D \bar{\Gamma}^D_i \mathbb{1}(a_j = 0) \ , \tag{19}$$

$$0 = - \lambda_y \bar{\Gamma}^D_i + \lambda_y \sum_{i'} D_{i,j} \pi(y_i, y_{i'}) \bar{\Gamma}^{ND}_{i',j} - \lambda_D \bar{\Gamma}^D_i \ . \tag{20}$$

As in Achdou et al. (2020), I use an upwinding scheme to obtain $\partial_a \left[s_{i,j} \bar{\Gamma}^{ND}_{i,j} \right]$. In particular,

$$\partial_a \left[s_{i,j} \bar{\Gamma}^{ND}_{i,j} \right] = \frac{\bar{\Gamma}^{ND}_{i,j} (s_{i,j,F})^+ - \bar{\Gamma}^{ND}_{i,j-1} (s_{i,j-1,F})^+}{da} + \frac{\bar{\Gamma}^{ND}_{i,j+1} (s_{i,j+1,B})^- - \bar{\Gamma}^{ND}_{i,j} (s_{i,j,B})^-}{da} \ .$$

Using the equation above, I can rewrite equations (19) and (20) in matrix notation as follows,

$$\mathbf{T} \bar{\Gamma} = \mathbf{0} \ , \tag{21}$$

where T is a $(JI - D + I) \times (JI - D + I)$ transition matrix, and $\bar{\Gamma}$ is a vector of length $(JI - D + I)$. D is the number of states in which the sovereign chooses to default. The first $(JI - D)$ elements of $\bar{\Gamma}$ correspond to states in which the government has access to international financial markets and chooses to repay its debt. The last I elements of $\bar{\Gamma}$ correspond to the I endowment states when the sovereign is in the state of default. Each column of transition matrix $\mathbf{T}$ sums to 1. This transition matrix is constructed as follows,

$$\mathbf{T} = \begin{bmatrix} T_{NW} & T_{NE} \\ T_{SW} & T_{SE} \end{bmatrix} \ .$$

The matrix T_{NW} is a $(JI - D) \times (JI - D)$ matrix which corresponds to movements between states of no-default to other states of no-default. It is combined of two matrices,

$$T_{NW} = \mathbf{A}^T + \lambda_y \tilde{\mathbf{\Lambda}}^T \ ,$$

where $\mathbf{A}$ is the one from equation (17), and represents moving across states due to savings behavior. $\tilde{\mathbf{\Lambda}}$ represents the transition between no-default states due to endowment shocks. To construct it we use the full transition matrix Λ from equation (18), and truncate it in all rows and columns which corresponds to a state of default (so that $\tilde{\mathbf{\Lambda}}$ is a $(JI - D) \times (JI - D)$ matrix).

The matrix T_{NE} is of size $(JI - D) \times I$ and corresponds to movements from the default states to the no-default states. Recall that the Poisson intensity of regaining access to international financial markets is Λ_D. T_{NE} is constructed as follows. First, construct

$$\tilde{T}_{NE} = \lambda_D \left(\mathbb{I}_{I \times I} \otimes e_{j_0} \right) ,$$

where e_{j_0} is a $J \times 1$ vector which has a single non-zero element in the j_0'th row where it is equal to 1. j_0 is the index which corresponds to $a_{j_0} = 0$. $\tilde{T}_{NE}$ is therefore a $JI \times I$ matrix. Then, T_{NE} is equal to the truncation of $\tilde{T}_{NE}$, where the rows corresponding to states of default are deleted.

T_{SE} corresponds to movements between default states due to endowment shocks and is given by

$$T_{SE} = \lambda_y \left(\Pi_y - \mathbf{I} \right) .$$

Finally, the matrix T_{SW} corresponds to movements between states of no-default into states of default. Such movements can arise only due to endowment shocks that move the sovereign into the default region. The size of T_{SW} is $I \times (JI - D)$. It is easy to compute it using $\tilde{\boldsymbol{\Lambda}}$. Let $\left(\sum \tilde{\boldsymbol{\Lambda}} \right)$ denote the summation of the $\tilde{\boldsymbol{\Lambda}}$ across columns, so that $\left(\sum \tilde{\boldsymbol{\Lambda}} \right)$ is of size $(JI - D) \times 1$. While the columns of Λ sum to 0, the columns of $\tilde{\boldsymbol{\Lambda}}$ sum to a negative number as some columns with only positive elements have been truncated. So $- \left(\sum \tilde{\boldsymbol{\Lambda}} \right)$ represents the probability of each no-default state to default conditional an endowment shock. We still need to rearrange the elements of $- \left(\sum \tilde{\boldsymbol{\Lambda}} \right)$ into a $I \times (JI - D)$ matrix. Define $\tilde{\boldsymbol{\Theta}} = (I_I \otimes \mathbf{1}_{J \times 1})$, and let $\boldsymbol{\Theta}$ be the truncation of $\tilde{\boldsymbol{\Theta}}$ where the rows corresponding to default states are truncated. Then we have that,

$$T_{SE} = -\lambda_y \left\{ \left[\left(\sum \tilde{\boldsymbol{\Lambda}} \right) \times \mathbf{1}_{1 \times I} \right] \odot \boldsymbol{\Theta} \right\}^T ,$$

where $\mathbf{1}_{M \times N}$ is an $M \times N$ matrix of ones, and $\odot$ is elementwise multiplication.

Note that the rows of matrix $\mathbf{T}$ sum to 0, so it is not invertible. In addition, notice that $\bar{\boldsymbol{\Gamma}}$ is pinned down only to a normalization as multiplying it by a constant would satisfy equation (21). I normalize the sum of $\bar{\boldsymbol{\Gamma}}$ to 1. While there are many ways to compute $\bar{\boldsymbol{\Gamma}}$, I follow the one proposed in the numerical appendix of Achdou et al. (2020). I set the element of $\bar{\boldsymbol{\Gamma}}$ which corresponds to the median income and zero assets to 1, call it $\bar{n}$. This means solving equation (21) but instead of a vector of zeros on the right hand side, I set the $\bar{n}$'th element to 1. In addition, I set the $\bar{n}$'th row of $\mathbf{T}$ to zeros except for the $\bar{n}$'th element which is 1. The transformed $\mathbf{T}$ matrix can now be inverted. So we can solve for $\bar{\boldsymbol{\Gamma}}$ by inverting the sparse transformed $\mathbf{T}$ matrix. Finally, we divide each element of the resulting $\bar{\boldsymbol{\Gamma}}$ by the sum of $\bar{\boldsymbol{\Gamma}}$ so that it sums to 1.

B.3 Algorithm for the Long Term Debt Model

The algorithm for solving the long-term debt model presented in section 5.1 is very similar to the one of short-term debt. The main difference is in the computation of the bond pricing schedule. While in the short-term debt model the pricing schedule depends only on the default policy function, with long-term debt it also depends on the debt issuance policy of the sovereign.

As in the short-debt model, the algorithm consists of three main steps: (i) computing the bond pricing schedule, (ii) solving the HJB equation of the sovereign, (iii) updating the default decision and check for convergence. However, as Chatterjee and Eyigungor (2012) explains, the long-term debt model is prone to oscillations which can prevent the convergence of the model. To remedy this risk, the main part of the algorithm consists of two nested loops instead of one loop. In the inner loop the default decision of the sovereign is kept fixed so that only the debt issuance policy function of the sovereign and the bond pricing schedule are being updated. Once the latter two converge, the default decision of the sovereign is being updated and the inner loop starts again. The algorithm is solved when the default policy function is equivalent to one of the previous default policy functions. This convergence criteria is discussed in details below. To start the outer loop from a good guess for the default policy function, the algorithm first solves 200 iterations while updating the default policy function after every inner loop iteration. In summary, the algorithm for the long-term debt model is as follows:

1. To start from a *good* guess for the default policy function, I repeat the following loop for 200 iterations:[26]

   (a) Compute the bond pricing schedule.

   (b) Solve the HJB equations.

   (c) Update default and debt issuance policy function, and repeat from step (a).

2. Main part of the algorithm. The outer loop:

   (a) Update default policy function. The inner loop:

       i. Update bond pricing schedule.

       ii. Solve HJB equations.

       iii. Update debt issuance policy function and check for convergence.

   (b) Check whether default policy function converged, if not repeat.

Computing the bond pricing schedule with long-term debt. Recall that this pricing schedule depends not only on the default decision but also on the savings behavior of the sovereign. The matrix form of equation (13) is

$$\bar{\mathbf{q}}^{n+1} r_f = z + \lambda_b - \lambda_b \bar{\mathbf{q}}^{n+1} + \lambda_y \tilde{\Lambda}^{n+1} \bar{\mathbf{q}}^{n+1} + \bar{\mathbf{A}}^n \bar{\mathbf{q}}^{n+1} ,$$

where $\bar{\mathbf{q}}^{n+1}$ is a vector of length $JI - D^{n+1}$. The matrix $\bar{\mathbf{A}}^n$ is the same as in section B.1 so we do not need to compute it. Finally, the matrix $\tilde{\Lambda}^n$ is the Λ matrix where the rows and columns corresponding to default states, D^{n+1}, are truncated so that $\tilde{\Lambda}^n$ is $(JI - D^n) \times (JI - D^n)$. Due to

---

[26]The first iteration starts from a default policy function of always repaying debt (never defaulting), and debt issuance policy function of always not changing the debt position.

the truncation the columns of $\tilde{\Lambda}^n$ do not sum to 1. This represents the possibility of an endowment shock moving the sovereign to default where the value of the bond is equal to zero. We can rearrange the equation above to obtain

$$\bar{\mathbf{q}}^{n+1} = \left[(r_f + \lambda_b)\, \mathbb{I} - \lambda_y \tilde{\Lambda}^{n+1} - \bar{\mathbf{A}}^n \right]^1 (z + \lambda_b) \ .$$

So finding the bond price schedule amounts to a simple sparse matrix inversion. Note that in theory, this method for finding the bond price schedule can also be applied in discrete time. However, in the continuous time formulation the matrix inverted in the procedure is sparse and not very different from a tri-diagonal matrix. This is not the case in discrete time. So, while theoretically possible, inverting the matrix above in the discrete time environment is very computationally costly. To prevent the bond pricing schedule from oscillating, instead of using $\bar{\mathbf{q}}^{n+1}$ in the following iteration I use a convex combination of $\bar{\mathbf{q}}^n$ and $\bar{\mathbf{q}}^{n+1}$. In particular, the bond pricing schedule is updated according to

$$\tilde{\mathbf{q}}^{n+1} = \Delta_q \bar{\mathbf{q}}^{n+1} + (1 - \Delta_q)\tilde{\mathbf{q}}^n \ ,$$

where Δ_q is an updating scalar.

Discussion of convergence in discrete vs. continuous time. The discreteness of the asset grid is the source of convergence problems both in the discrete-time long-term debt model studied in Chatterjee and Eyigungor (2012) and in the continuous-time long-term debt model studied in this paper. In discrete time, as discussed in Chatterjee and Eyigungor (2012), the discreteness of the asset grid leads to a convergence problem in two ways:

1. *Discrete borrowing decision*. The choice over the borrowing decision is discrete. Since the debt pricing schedule depends on the savings behavior, discrete changes in the savings behavior can lead to changes in debt pricing, which in turn would alter the optimal savings behavior of the sovereign. So, even without a change to the default frontier, the algorithm can oscillate without convergence. See Figure 2 and its discussion in Chatterjee and Eyigungor (2012).

2. *Discrete default frontier*. The maximal amount of debt the sovereign can sustain without defaulting, for each endowment level, is discrete. Since the debt pricing schedule depends on the future probability of default, which depends on the default frontier, discrete changes to the default frontier can lead to substantial changes in debt pricing. Such changes can alter the sovereign's default behavior, leading to oscillations in the default frontier.

The continuous-time approach presented in this paper overcomes the first problem, but not the second. The first problem is resolved because the savings decision is no longer discrete. The sovereign does not choose the discrete level of debt in the following period, but rather the intensity of the savings/borrowing ($\dot{a}$). The savings policy is recovered using the upwinding scheme from the value function, and is not restricted to take on discrete values.

The continuous-time approach, however, does not resolve the second convergence problem. The default frontier is restricted to be on the discrete asset grid, so oscillations in the default frontier remain to be a challenge for convergence. The outer loop of the algorithm may not converge. Below I discuss the default convergence criterion I employ in order to overcome the convergence problem.

Default convergence criterion. When the endowment and asset grids are dense there is a higher probability the outer loop of the default policy function may still oscillate despite the nested loop structure of the problem. Consider a case in which the asset grid is fine so that the default frontier varies between (only) two adjacent asset grid points for every endowment level. Notice that since the grid is fine, whether the default frontier is in one asset point or the adjacent one makes almost no difference at all in the quantitative solution of the model. However, this can be a major problem for convergence. If the endowment grid consists of 25 points, there are more than 33 million possible default frontiers. So the outer loop may not converge in a reasonable amount of time.

To avoid such oscillations, the first convergence criteria of the outer loop is to stop whenever the default frontier is equivalent to the default frontier in any previous iteration. This avoids "hard oscillations" - oscillations that may create infinite loops. The difficult challenge is to avoid "soft oscillations" - oscillations that do not cause infinite loops but rather that do not converge for a very long period of time, such as the 33 million default frontier possibilities discussed above. To avoid such oscillations, which may be costly when performing estimations, the researcher can put an upper limit on the number of outer loop iterations. For all computations in this paper, such upper limit was not imposed. As a rule of thumb, numerical simulations suggest that after 50 iterations the default frontier has only extremely small changes so that an upper limit of 50 iterations may suffice.

An alternative approach to avoid "soft oscillations" is to set the convergence criterion such that the default frontier is not changing more than some level of tolerance from one iteration to the other. Such criterion would significantly reduce the maximum possible number of default frontier options. It is worth noting again that this problem occurs only for dense asset and endowment grids, and that considering very fine grids does not change the quantitative results of the model substantially.

Figure 8 presents an example of such soft oscillation. The endowment grid consists of 25 points and the asset grid consists of 801 points. The convergence criteria is met only after almost 300 outer loop iterations. The left panel presents the default thresholds for different endowment levels (colors) across iterations. As can be seen in the figure, the default frontier for each endowment level does not vary much across iterations. The solution to the default frontier corresponds to the points with the highest outer loop iteration (the extreme right of the x-axis). This indicates that the initial guess, after the first 200 iterations, is quite close to the actual solution of the problem.

The right panel presents the difference in each iteration, for every endowment level, between the asset default threshold in each iteration to the solution of the default frontier. The difference

Figure 8: Example of Soft Oscillation
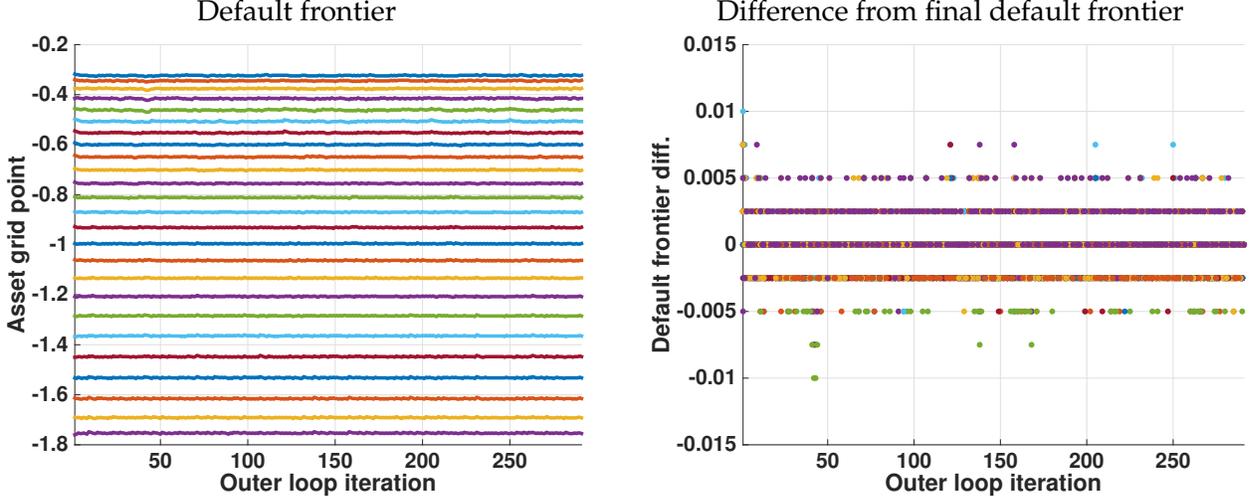
This figure presents an example in which the algorithm takes a long period of time to converge due to a *soft oscillation* problem. The left panel presents the asset level point in which the sovereign is indifference between default and repaying its debt for each outer loop iteration, for different endowment levels. In the right panel each dot represents the difference between the asset default threshold at iteration i to the fixed point default threshold. Each color represents a different endowment level.

is quite small and almost all points fall between 2 asset grid points below or above the solution to the default frontier (the difference between every two asset grid points is 0.0025).

This example suggests that stopping the algorithm prior to the convergence criterion would result in very minor quantitative differences.

C   World Interest Rate Fluctuations

C.1   HJB Equations

$$\rho w(y,r) = u(y - \phi(y)) + \lambda_y \int_0^\infty \big(w(y',r) - w(y,r)\big) \, dF_y(y',y)$$

$$+ \lambda_r \int_0^\infty \big(w(y,r') - w(y,r)\big) \, dF_r(r',r) + \lambda_D \left[v(0,y,r) - w(y,r)\right] \, ,$$

$$\rho v(a,y,r) = \max_c \{ \rho w(y,r), u(c) + v_a(a,y,r) \left(y - c + r(a,y,r)a\right)$$

$$+ \lambda_y \int_0^\infty \big(v(a,y',r) - v(a,y,r)\big) \, dF(y',y) + \lambda_r \int_0^\infty \big(v(a,y,r') - v(a,y,r)\big) \, dF_r(r',r) \Big\} \, .$$

C.2 Fit of Exogenous Interest Rate Process

Table 7: World Interest Rate Fluctuations - Targeted and Non-targeted Moments

Moment	Frequency	Data	Model	Targeted
Volatility	Daily	0.472	0.472	✗
	Weekly	0.473	0.472	✗
	Monthly	0.473	0.472	✗
	Quarterly	0.472	0.472	✓
Kurtosis	Daily	3.42	3.00	✗
	Weekly	3.40	3.00	✗
	Monthly	3.42	3.00	✗
	Quarterly	3.46	3.00	✗
Auto-correlation	Daily	0.999	0.999	✗
	Weekly	0.998	0.997	✗
	Monthly	0.989	0.988	✗
	Quarterly	0.963	0.963	✓
Volatility of changes	Daily	0.016	0.017	✗
	Weekly	0.031	0.037	✗
	Monthly	0.070	0.074	✗
	Quarterly	0.127	0.129	✗
Kurtosis of changes	Daily	8.02	187.42	✗
	Weekly	4.75	37.48	✗
	Monthly	2.96	9.37	✗
	Quarterly	3.12	3.12	✓

This table includes data moments and the moments of the exogenous compound Poisson process for the world interest rate. The data used is 5-year treasury constant maturity rate of the US between 1983:Q3 and 2001:Q4. The last column indicates the three moments that were targeted to calibrate the three exogenous process parameters.