

Interest-growth differentials and debt limits in advanced economies*

Philip Barrett[†]

First draft: April 2018
This draft: October 2019

Abstract

Do persistently low nominal interest rates mean governments can safely borrow more? I argue standard models of debt sustainability cannot address this issue. The key model parameter in a wide class of models is the long-run difference between the nominal risk-free interest rate and the nominal growth rate. If negative, maximum sustainable debts are unbounded. Data from five advanced economies suggest that this differential is indeed likely negative; different approaches all produce negative point estimates and confidence intervals implying low probabilities of positive values. And even if the long-run differential were positive, the quantitative impact of short-term fluctuations is tiny.

JEL-Codes: E62, H62, H63

Keywords: Fiscal policy, Government Debt, Interest Rates

*This paper reflects my own personal views and does not represent those of the IMF, its board, or management. I am grateful for helpful comments from Adrian Peralta Alva, Tamon Asonuma, Olivier Blanchard, Romain Duval, Raphael Espinoza, Luc Eyraud, Vitor Gaspar, Klaus Hellwig, Jun Kim, Divya Kirti, Leonardo Martinez, Catherine Pattillo, Ricardo Reis, Abdel Senhadji, Mark Watson, Philippe Wingender, Jermoe Zettelmeyer, and participants of the MFM Summer Session.

[†]International Monetary Fund. Website: www.pobarrett.com Email: pbarrett@imf.org

1 Introduction

Following the global financial crisis, central banks in advanced economies loosened monetary policy aggressively, pushing rates to historic lows (Figure 1). There, interest rates have mostly stayed. Even in the most notable exception – the United States – policy appears to be becoming increasingly accommodative.

The lack of room for further monetary stimulus – and, perhaps, its diminishing effectiveness – has led policymakers and academics alike to ask whether there is scope for fiscal policy to take a more active role in stabilizing business-cycle fluctuations.¹ Low interest rates, the argument goes, make the present a particularly appealing time for fiscal expansion simply because finance is so cheap, despite very high public debt levels (Figure 2).

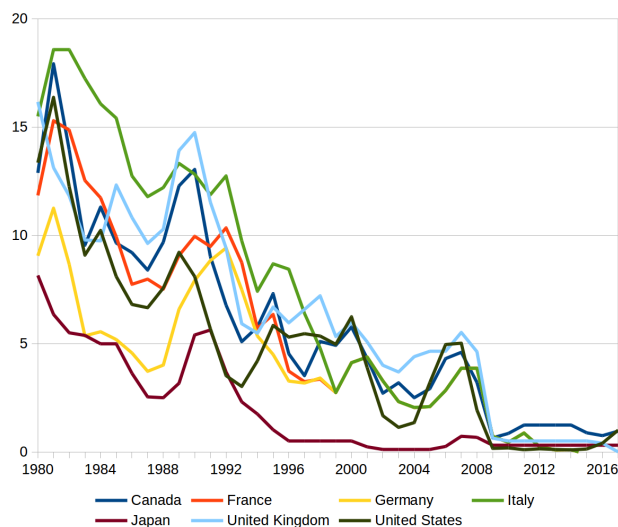


Figure 1: Central Bank interest rates in G7 economies 1980-2017 (annual data)

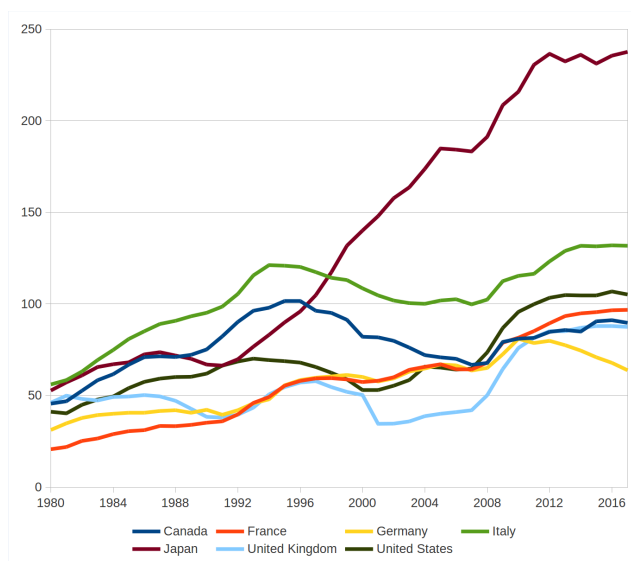


Figure 2: General government gross debt to GDP ratio, percent, in G7 economies 1980-2017 (annual data)

The contribution of this paper is to assess this argument. In doing so, I make three central points. First, a point of *theory*, showing that the conceptual framework on which this argument rests is valid only if a parameter restriction holds. If the framework fails, so too does the argument. Second, a point of *measurement*. I test this parameter restriction probabilistically and show that it most likely fails, casting doubt upon both the framework and the argument. However, there a degree of statistical uncertainty over this result. Given this, the third point is an exploration of

1. For a recent policy example, see Sajid Javid’s statement to the House of Commons (HC Deb, 4 September 2019) which cites low interest rates as a primary motive for increased investment expenditures. For an academic one, see (Blanchard 2019).

the *quantitative implications* of this result. I show that even if this restriction did not fail, this uncertainty has a sufficiently large impact on sustainable debt levels as to totally dwarf the effect of short-term fluctuations in interest rates.

The rest of the introduction introduces these contributions in a little more detail..

Theory: In section 2, I show that in a wide class of models – including in almost all contemporary models of sovereign default and debt sustainability – the parameter governing the knife-edge between finite and infinite maximum sustainable debt levels (or, more succinctly, debt limits) is the sign of the long-run average difference between the *risk-free* nominal interest rate and the nominal growth rates. If this is positive, sustainable debts are necessarily bounded. If this differential is negative, there is no limit on the size of a sustainable debt.² If true, this would cast doubt the idea that currently low interest rates increase public borrowing capacity; a temporary reduction in borrowing costs has no impact on a government that can already borrow as much as it wishes.

This is an important result because it corrects an implicit assumption of the empirical literature: that the relevant interest rate for assessing sustainability of large debts is something other than the risk free rate. For example, Escolano (2010) and Mauro et al. (2015) use the effective rate (i.e. the sum of budgetary interest payments divided by the face value of the debt stock) when assessing debt sustainability. Likewise, Blanchard (2019) uses a maturity-weighted average of US government yields.

Of course, the usefulness of theory for guiding empirics is bounded by the extent to which the theory is reasonable, or at least well-accepted. This is why I describe carefully the set of models in which the result holds and show that it nests almost all conventional models of debt sustainability. This includes both those where default is strategic default and those where affordability alone drives default.³ In models where the result does not hold, such as Blanchard (2019) or Aguiar et al. (2017), the theory also helps explain why.

Measurement: In Section 3, I test the findings of Section 2 empirically, and examine historical risk-free interest-growth differentials in Advanced Economies in two samples, one annual since 1880 using data from Jordà, Schularick, and Taylor (2017), and one quarterly since 1960.

In common with other studies, I find that point estimates of the long run average interest-growth differential in advanced economies are frequently negative.⁴ For example, Figure 3 shows

2. This is not to say that high debt levels are necessarily desirable, only that they are possible. Not all things that can be done should be done.

3. Those in the former class include Arellano (2008), Aguiar and Gopinath (2006), and Chatterjee and Eyigungor (2012); those in the latter include Bohn (1998) and Ghosh et al. (2013). In proving the central result, I also describe conditions for existence and uniqueness of equilibrium, thereby extending the results of Auclert and Rognlie (2016) to a setting where default is non-strategic.

4. This observation is not new. Ball, Elmendorf, and Mankiw (1998), Escolano (2010), and Mauro et al. (2015)

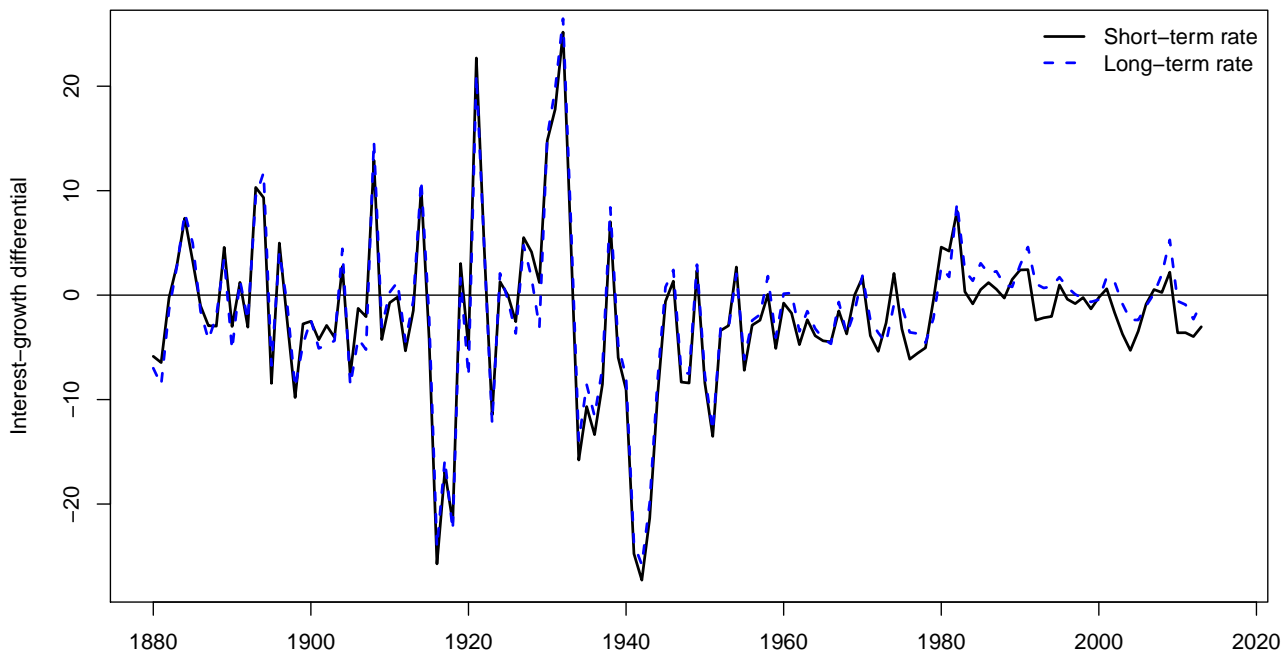


Figure 3: US interest-growth differentials 1880-2016

two measures of the interest-growth differential for the USA using both short- and long-term interest rates from the annual sample.⁵ Both measures are negative in around two-thirds of years, and negative on average over both the whole sample, the post-war period, and since 1990. Nor is this a US-specific phenomenon; a similar pattern holds in other advanced economies.

Yet given the potentially very large policy implications of this result, it seems wise to test statistically whether we can eliminate the possibility that the long-term interest-growth differential may, in fact, be positive. This is not something that previous authors have attempted.

I therefore employ two statistical tests to determine the sign of the long-run interest-growth differential. The first is a likelihood ratio approach based on combinations of VAR coefficients. For both datasets, the point estimates from the VAR are negative,⁶ but the associated likelihood ratio test cannot reject the possibility that the long-run interest-growth differential is positive. This upper bound is typically in the order of zero to one percent per year.

all note this point, albeit using different interest rates.

5. The short rate is a policy rate, equal to the Federal Funds Rate in the post-war period, and the long-term rate is that on long-term government debt (typically a ten-year bond).

6. This is not inconsistent with the findings of Piketty (2014), who shows that the difference between the rate return of capital and economic growth is typically positive at standard confidence levels. That the same difference for risk-free rates is negative is simply a product of the large spread between risky and risk-free rates of returns. The surprisingly large extent of this difference is just another manifestation of the equity premium puzzle.

The second approach is a spectral one, developed by Müller and Watson (2016). Although this approach produces similar, negative, point estimates for the long-run interest-growth differential, it is a less restrictive statistical method, and therefore generates larger confidence intervals and is more sensitive to outliers. But for countries with long, uninterrupted data, it provides upper bounds on long-run interest-growth differentials that are around one to two percent per year. Although they disagree on the details, both methods suggest that small and positive long-run interest-growth differentials are unlikely, but not impossible, and certainly not very large.

The difference therefore between this contribution and other past empirical work is twofold. First, the results of the theory section provide guidance about *which* interest rate to use in the calculation of the differential. Second, I go beyond just point estimates of this key parameter and put confidence intervals around estimates of the long-run differential .

Quantitative implication: Although the results of the empirical analysis in Section 3 suggest that most models of debt sustainability do not place bounds on sustainable debt levels, there is some uncertainty about this conclusion. In Section 4, I therefore investigate whether the class of standard models might still have something useful to say about the motivating policy question: how shocks to interest-growth differentials affect maximum sustainable debt levels.

To do this, I use a version of one such model to generate conservative estimated historical debt limits for the USA. The limits are conservative in that long-run interest-growth differentials in the model are positive, guaranteeing finite debt limits. But they are plausible in that this restriction is one that is consistent with the data; that is to say, small. This leads to debt limits which move little in response to variation in the interest-growth differential, increasing around 5 percent of GDP following the 2008-9 financial crises.

The overall level of these limits, though, are highly sensitive to the assumptions about the long run interest-growth differential. The statistical analysis of this differential in the empirical part of the paper allows for a probabilistic quantification of the impact of this uncertainty on debt limits. For example, solving the model with a long-run differential at the 4% critical value for the VAR-based estimates (instead of the 5% value) leads to maximum sustainable debt-GDP ratios which are about 35 percentage points lower. More intuitively, the increase in fiscal space due to the large declines in interest rates since 2008 would be entirely offset by a microscopically more conservative approach to uncertainty from just one key parameter: the long-run interest-growth differential.

Section 5 concludes, discussing the implications of this paper for future work on debt sustainability and suggesting that the mechanisms that current theory might lack.

1.1 Related Literature

This paper is most closely related to three main literatures: one on models of debt repayment capacity; one on sovereign default in a stochastic environment; and one on interest-growth differentials.

The model used in this paper to compute debt limits is one of capacity to repay. One commonly used approach in this literature is to compute the level to which debt converges, usually given a linear fiscal rule. Papers employing this approach include Bohn (1998), Mauro et al. (2015), D’Erasmus, Mendoza, and Zhang (2016), and Collard, Habib, and Rochet (2015). These approaches do not yield a *maximum* debt level, just an average to which debt converges in the long run. This is either finite or it is not, and the long-run dynamics are independent of the starting debt level. Ghosh, Kim, Mendoza, Ostry, and Qureshi (2013) go further, employing a nonlinear fiscal reaction function to produce debt dynamics which are only stable if debt starts below some threshold. I use a similar model as this paper, but extend it to allow for stochastic variation in interest and growth rates, showing how this can represent the key mechanism in a large class of models.⁷ Ostry and Kim (Forthcoming) and Kim and Asonuma (Forthcoming) also develop versions of this model with uncertainty over growth, but without persistent shocks or uncertainty over the risk-free interest rate. A further important difference relative to Ghosh et al. (and Kim and Asonuma) is that I consider only bounded fiscal reaction functions.

The second related literature is that on models of strategic sovereign default. These papers combine the questions of what debt levels are affordable with what are desirable. In Arellano (2008) and Aguiar and Gopinath (2006), real growth shocks drive default. More recently, Tourre (2016) extends this literature by allowing for stochastic shocks to investors’ pricing kernel (and hence the cost of borrowing). There, as in the model in this paper, the correlation between growth and risk-free interest rates is the key determinant of sustainable debt levels. And Hatchondo, Martinez, and Roch (2017) echo the basic message of this paper in a model of strategic default, producing default thresholds that are highly sensitive to parameters. Instead, they argue, elevated spreads are a more robust measure of limited fiscal space.

Last, this paper is related to previous works which have on measured interest-growth differentials in advanced economies. Ball, Elmendorf, and Mankiw (1998) find that interest-growth differentials have, on average, been negative in the USA since the late 19th century. This finding is robust to different subsamples and measures of interest rates. Expanding to advanced economies, Escolano (2010) finds interest-growth differentials are typically small and positive.⁸

7. A prototype version of this stochastic model underpins the analysis in Box 1.4 of IMF (2017).

8. Although is largely a function of the sample period; interest-growth differentials were abnormally high by historical standards during 1991-2008. Furthermore, these authors find considerable cross-country heterogeneity and many countries are estimated to have a negative differential.

This point is also discussed in Blanchard (2019), whose discussion also extends to an analysis of capital accumulation and welfare. And Mehrotra (2017) also identifies negative average differentials, using data from advanced economies in the post-war period. The fundamental difference between the empirical analysis of those papers and that pursued here is that I go beyond point estimates and formulate statistical tests of the long-run interest-growth differential. This puts confidence intervals around the key model parameter in a wide class of models.

Other papers, such as Turner and Spinelli (2011) and Woo, Shabunina, and Escolano (2017), have sought to understand *why* interest-growth differentials have declined since the 1980s.⁹ In contrast, I take the behavior of interest-growth differentials as a given, purely statistical phenomenon and have little to contribute to economic explanations of past movements of this variable. Likewise, this paper takes no stand on whether higher debt limits are desirable, only if they are possible. For discussion on this point, particularly in the context of negative average interest-growth differentials, see Blanchard and Weil (2001).

2 Theories of debt sustainability

In this section, I present a general framework for thinking about debt sustainability which nests a large fraction of models in the literature. I show conditions for existence are uniqueness of equilibrium, and demonstrate an important knife-edge result: that debt is sustainable at an arbitrarily high level if risk-free interest rates exceed nominal growth rates on average over the long run.

As discussed in the preceding section, there are many models of debt sustainability. Some consider only how much debt a government can afford. Others integrate both the debt default decision with some notion of optimal fiscal policy. Yet a common feature unites these otherwise disparate models: that the government faces a budget constraint where debt is issued at a price which reflects the probability that it will be repaid in future. In its simplest form, this can be expressed in the form:

$$s_t + q_t(b_{t+1})b_{t+1} = (1 - d_t)b_t \tag{1}$$

where s_t is government's surplus, b_t and b_{t+1} are the government's maturing and newly issued debt respectively, d_t is the default indicator, and $q_t(\cdot)$ is the price of newly issued debt.¹⁰ Crucially,

9. And a yet further group of papers attempts to explain why real rates have fallen during the same period, usually without relating to simultaneous movements in growth rates. These include Caballero, Farhi, and Gourinchas (2008), Rachel and Smith (2015), Sajedi and Thwaites (2016), Carvalho, Ferrero, and Nechio (2016), Lisack, Sajedi, and Thwaites (2017), Borio et al. (2017), and Lunsford (2017)

10. In this motivating example, I suppress other quantitatively important features such as debt maturity, partial repayment of defaulted debt, the evolution of state variables and so forth. I am specific about these in the following

$q_t(\cdot)$ is a decreasing function, reflecting the fact that equilibrium debt prices feature a discount for default risk and that the probability of default increases with the size of the debt.

This equation produces an inter-temporal trade-off for a government contemplating default. Policies which increase the probability of default in period $t + 1$ increases, raise the cost of borrowing in period t . This mechanism plays a fundamental role in all models with some version of equation (1), and is present no matter how the surplus function is determined. For example, in models of strategic default a government will occasionally endure a painful default in the present only because meeting the budget constraint in future would be more onerous.

I therefore work in a framework where a generalized form of equation (1) holds, and am agnostic about the other elements in the model. The key results are therefore applicable to a wide class of models.

2.1 Environment

I consider a class of models of fiscal policy that can be described by particular states, exogenous variables, policies, and stochastic processes.

States: The states are x and b . The variable b represents the government's debt-to-GDP ratio. The variable x is a vector of all other stochastic states needed to summarize the state of the economy, such as productivity, capital, and the like. For notational simplicity, I assume that x is discrete with support \mathcal{X} , although this is not essential.

Exogenous variables: The government's state-dependent tax and spending policies are taken as given, resulting in surplus characterized by the rule $s(x, b)$. The state variable x also determines the risk-free nominal interest rate $R(x)$ and the nominal growth rate of the economy $G(x)$.

Policies: The government chooses first whether to default then its debt issuance. Default is denoted by δ^{11} and new debt by b' .

Stochastic processes: The continuation state then evolves according to:

$$x' \sim \pi(x'|x, b', \delta) \tag{2}$$

Where the function $\pi(\cdot|x, b', \delta)$ captures all the behavior of all the other model variables. I assume 1) $\pi(x'|x, b', \delta) > 0 \forall (x, b', \delta)$;¹² and 2) $\pi(x'|x, b', \delta)$ is continuous in b' .

section.

11. Repayment is $\delta = 0$, default $\delta = 1$

12. This guarantees that if continuation debt b' is bounded, then (x, b) is an irreducible Markov chain for any default rule $\delta = d(x, b)$.

Described this way, the evolution of the non-debt state can capture by a very wide range of economic models. The dependence of x' on δ means that this framework can represent optimal default models where output declines or governments are excluded from credit markets (modelled as states where $R(x) = \infty$) for some stochastic interval following a default. And the dependence of x' on b' admits a meaningful economic role for the government's debt choice, such as when debt issuance crowds out private capital accumulation. In Appendix B I show explicitly how three standard models of debt sustainability can be mapped into this framework: Arellano (2008), Ghosh et al. (2013), and a simple optimal tax model.

While this framework does rule out models where current debt b directly affects risk-free interest or growth rates, the future evolution of the state x is determined by next-period debt b' . And so there is still scope for debt to impact risk-free nominal interest or growth rates, just with a lag.

Given this environment, the model is closed by rules for policies $\delta = d(x, b)$ and $H(x, b, \delta)$, and prices $q(x, b', \delta)$. In general, one could imagine that the policies might be chosen to optimize some social welfare criterion. That is not ruled out here, but it is not imposed either. Instead, I simply require that these functions collectively satisfy equilibrium, which is defined by three conditions. Intuitively, these are: that investors are rational in that the price of debt reflects the model-implied default probabilities; that the government always satisfies a period budget constraint with finite debt; and that debt is in a sense to be defined later “stable”.

I also assume that the long-run behavior of $R(x)$ and $G(x)$ is unaffected by the government's default choices:

Assumption 1 *For all default rules $\delta = d(x, b)$ and debt issuance rules $H(x, b, \delta)$, the long-run expectation of the interest-growth different is constant, that is:*

$$\mathbb{E} \left[\log \left(\frac{R(x)}{G(x)} \right) \right] = \rho$$

This assumption is a crucial to the results that follow, so a brief discussion of its economic underpinnings is in order.

Standard models of debt sustainability typically guarantee Assumption 1 via fairly restrictive assumptions on the process for $R(x)$ and $G(x)$. For example, in Ghosh et al. (2013) $R(x)$ and $G(x)$ are both constant. And in Aguiar and Gopinath (2006) and Arellano (2008), $R(x)$ is constant but $G(x)$ is stationary and independent of fiscal policy.

Yet Assumption 1 holds in much more general models too. For instance, if debt is priced by

domestic investors with CRRA preferences and real growth and inflation are lognormal, then:

$$\begin{aligned} \log R_t - \log G_{t+1} = & -\log \beta + \underbrace{(\sigma \mathbb{E}_t g_{t+1}^c - g_{t+1})}_{\text{Growth surprise}} + \underbrace{(\mathbb{E}_t \pi_{t+1} - \pi_{t+1})}_{\text{Inflation surprise}} \\ & + \underbrace{\frac{\sigma^2}{2} \text{Var}_t g_{t+1}^c + \frac{1}{2} \text{Var}_t \pi_{t+1} + \sigma \text{Cov}_t(g_{t+1}^c, \pi_{t+1})}_{\text{Risk premium}} \end{aligned} \quad (3)$$

Where β is the subjective discount rate, σ the inverse intertemporal elasticity of substitution, g_{t+1}^c and g_{t+1} the log growth rates of per capita consumption and output respectively, and π_t the log inflation rate.

Taking expectations over the long-run stationary distribution, inflation drops out:

$$\mathbb{E} \left[\log \left(\frac{R(x)}{G(x)} \right) \right] = -\log \beta + \mathbb{E}(\sigma g^c - g) + \mathbb{E}(\text{Risk premium})$$

In any model with an Euler equation of this type, therefore, Assumption 1 is violated in only three cases. First, when fiscal policy can alter the long-run wedge between growth rates of per capita consumption g^c and output g . If the marginal investor is the representative household, though, this is impossible. Then, this wedge must be constant (and equal to the long-run rate of population growth) otherwise the long run consumption-output ratio will be zero or one. But if the marginal investor is not representative, this condition can be violated. Indeed, this is the key mechanism in Blanchard (2019), which shows that in an overlapping generations model fiscal policy can induce high lifetime consumption growth for the marginal investor without a corresponding effect on aggregate economic growth. This results in an average risk-free interest-growth differentials in the long run which depends on fiscal policy.

The second case where Assumption 1 might be violated is through the long-run risk premium. A risk-free nominal asset can feature such a premium because inflation generates volatile real returns. Fiscal policy which reduces the volatility of inflation (or its covariance with real growth) can therefore affect the interest-growth differential in the long run. While theoretically possible, this seems empirically implausible, at least for advanced economies. Inflation and consumption growth volatility are typically too small for there to be much scope for fiscal policy to act through this channel (see Table 8 in Appendix C.1 for data on this point).

The third way that Assumption 1 might be violated is that an Euler equation of this form simply does not hold at all. For example, financial repression might cause intertemporal distortions which prevent investors from optimizing their consumption-investment decisions, and thus violating their Euler equation. Or, foreigners may be the important marginal investors, and their

consumption growth may diverge drastically from domestic growth. Nevertheless, it seems unlikely that such factors would be affected by fiscal policy, and so can reasonably be considered as part of the exogenous behavior of $R(x)$ and $G(x)$.

Equation (3) also has something to say about the short-run dynamics of interest-growth differentials; specifically that the auto-correlation is likely to be rather small. This is because nominal interest rates are, to first order, compensation for deferred consumption and expected inflation. Nominal growth is the sum of real output growth – which is likely to be highly correlated with expected consumption growth – and realized inflation. The persistence of inflation and growth surprises is therefore likely to be very short. Indeed in the common special case of log preferences (i.e. $\sigma = 1$) and a constant risk premium, the risk-free interest-growth differential has an auto-correlation of zero. This is important for short-run fluctuations in sustainable debt levels in any theory where solvency plays a key role, as the impact of short-run fluctuations in interest-growth differentials are amplified by the persistence of the shock.

Of course, equation (3) may not hold for many other reasons: preferences may not be CRRA or shocks lognormal. The purpose of the foregoing is not to argue that this is the one way to think about long-run interest-growth differentials. Rather, it is to highlight that Assumption 1 is a property not of most relevant models in the literature, but also of the most obvious extensions.

2.2 The surplus function

I assume that the surplus function satisfies the following regularity conditions:

Assumption 2 (Regularity conditions for $s(x, b)$)

1. $s(x, b)$ is continuous and twice differentiable in b
2. For all x , $\exists b_s(x)$ such that $s(x, \hat{b}(x)) > \hat{b}_s(x)/G(x)$
3. $s(x, b)$ is bounded for all b , and $\exists \bar{s}(x)$ s.t. $\lim_{b \rightarrow \infty} s(x, b) = \bar{s}(x)$
4. $s(x, 0) = -\epsilon \leq 0$ and $s''(x, 0) < 0$
5. $s'(x, b)$ and $s''(x, b)$ change sign at most once

Intuitively, this means that the government's surpluses are bounded and either strictly concave or single-peaked and that in each state gross financing needs are negative for at least some level of the debt. Figure 4 illustrates the two types of permissible surplus rules. These are pertinent

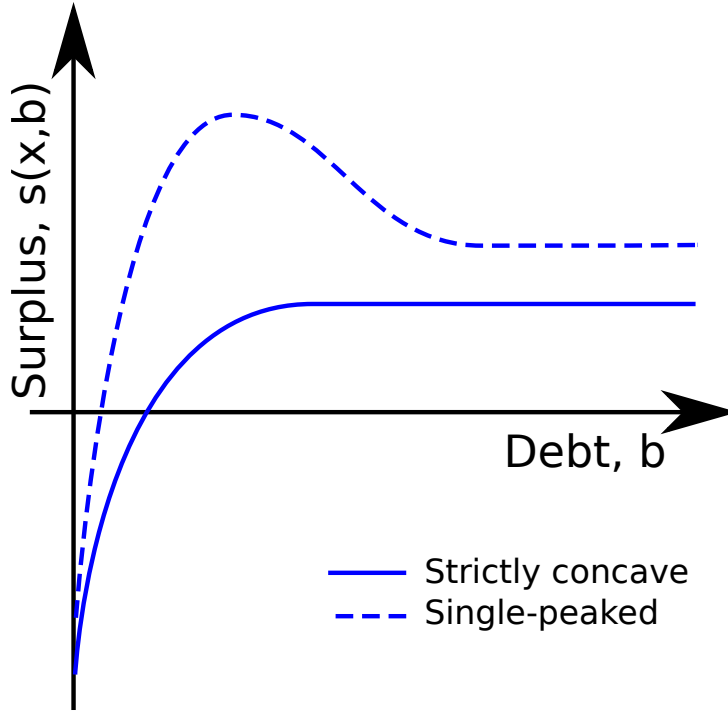


Figure 4: Stylized surplus function

because the surplus function for the canonical strategic default model is strictly concave (see Appendix B) and single-peaked in the “fiscal fatigue” model of Ghosh et al. (2013).¹³

Taking the surplus functions as given is not as restrictive as it might at first seem. Even in models where the surplus, default, and debt issuance choices are made simultaneously – perhaps in order to maximize some welfare criterion – the debt and default rules must still satisfy investor rationality, period feasibility, and long-run stability *conditional* on the surplus rule.

2.3 Investor rationality

The government issues long-maturity debt, in contrast to the motivating example in equation (5), this setting allows for long-maturity debt. Here, this is modeled with a declining coupon bond, as in Hatchondo and Martinez (2009). The payoff of one unit of outstanding debt in each period is a coupon of size $\lambda < 1$ and $1 - \lambda$ units of new debt. The resulting bond has average maturity $1/\lambda$.

If the government defaults, investors receive a partial payment equivalent to the fraction $\psi < 1$

13. Although in Ghosh et al. (2013) the surplus function is not bounded below, meaning that they have finite debt limits under negative interest-growth differentials but only because the government commits to an infinitely large deficit.

of the value of the bond. When debt is short-term (i.e. $\lambda = 0$) this is equivalent to a repayment in default proportional to the principal but not the interest. Investors are risk-neutral but risk-aversion can be mapped into this setting simply by replacing the true probability of default with investors' risk-neutral probabilities (discussed further below).

Definition 1 (Investor rationality) *Given functions $d(x, b)$, $H(x, b, \delta)$, investors are rational if the pricing function $q : \mathcal{X} \times \mathbb{R} \times \{0, 1\} \rightarrow \mathbb{R}$ satisfies:*

$$q(x, b', \delta) = \frac{1}{R(x)} \sum_{x'} \pi(x'|x, b', \delta) (1 - d(x', b')) (\lambda + (1 - \lambda)q(x', H(x', b', d(x', b')), d(x', b'))) \\ + \psi \left(\sum_{x'} d(x', b') \pi(x'|x, b', \delta) \right) q(x, b', \delta) \quad \forall (x, b', \delta)$$

The first term on the right-hand side is the payoff under repayment, which itself is the sum of the coupon and a new bond. The second term is the payoff under default. The investor rationality condition can be written more compactly as:

$$q(x, b', \delta) = \frac{\lambda(1 - p(x, b', \delta)) + (1 - \lambda) \sum_{x'} \pi(x'|x, b', \delta) (1 - \lambda)q(x', H(x', b', d(x', b')), d(x', b'))}{R(x) (1 - \psi p(x, b', \delta))} \quad (4)$$

Where: $p(x, b', \delta) = \sum_{x'} \pi(x'|x, b', \delta) (1 - d(x', b'))$ is the next-period default probability

Investor rationality in this form is a standard assumption in almost all models of debt sustainability. This even extends to models of self-fulfilling default, such as Cole and Kehoe (2000). The only modification in this case is that the state x also includes a randomization variable determining whether investors co-ordinate on “running” on the debt.¹⁴

With long-run debt, the pricing function is the solution to a fixed-point problem; the price of the debt in the next period enters the price of the debt today. In Appendix A I prove that this fixed point problem has a unique solution given $D(x, b)$ and $H(x, b, \delta)$.

Proposition 1 *If $\min_x R(x) > 1 - \lambda$ then for any $d(x, b)$ and $H(x, b, \delta)$:*

1. *There exists a pricing function $q_{|d,H}$ satisfying investor rationality.*
2. *If $H(x, b, \delta)$ is weakly increasing in b' then $q_{|d,H}$ is weakly decreasing in b'*

14. A notable exception is Bi and Leeper (2013), in which investors price debt not according to the probability of default in the next period but at some point over the entire future.

Note that this is a very weak restriction as λ is bounded below by 0 and $R(x)$ is rarely much less than 1.

2.4 Period feasibility

Period feasibility is the generalized equivalent of equation (1). It means that the government's default rule satisfies its period budget constraint given prices consistent with default.

Definition 2 (Period Feasibility) *A default rule $d(x, b)$ is period feasible if there exists finite $H \in \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$ such that $\forall (x, b) \in \mathcal{X} \times \mathbb{R}$:*

$$(H(x, b, d(x, b)) - (1 - \lambda)(1 - d(x, b))b/G(x)) q_{|d, H}(x, H(x, b, d(x, b)), d(x, b)) = \max \left(\frac{\lambda b}{G(x)}(1 - d(x, b)) - s(x, b), 0 \right) \quad (5)$$

The right hand side of equation (5) is the government's gross financing needs, i.e the primary deficit plus debt coming due; the left hand side is revenue from debt issuance. Although prices and debt issuance are defined for all δ , period feasibility requires only that gross financing needs are covered by debt issuance only under the default policy defined by the government's policy rule $\delta = d(x, b)$. That is, period feasibility need only hold in equilibrium.

I focus on a particular set of default rules, where default occurs if and only if debt is above a certain (state-dependent) threshold. The reason for this is that, as I show in Appendix A, if H is weakly increasing and solves equation (5), then $H(x, b, 1)$ is finite only for b less than some state-dependent threshold $\bar{b}(x)$. That is, the continuation debt rule itself implies a threshold rule. This is also why models of strategic default also produce on threshold default rules.

Definition 3 (Threshold default rules) *A default rule d is a threshold default rule if there exists $\bar{b}(x)$ such that*

$$d(x, b) = \begin{cases} 1 & \text{if } b > \bar{b}(x) \\ 0 & \text{if } b \leq \bar{b}(x) \end{cases}$$

If $\bar{b}(x) < \infty$ for all x then we say that d is a finite threshold rule.

In which case we use the notation $d = d_0(\bar{b})$, and with slight abuse of notation use $q_{|\bar{b}, H}$, as shorthand for $q_{|d_0(\bar{b}), H}$. Such rules are period feasible under weak restrictions:

Theorem 1 (Existence of period-feasible default threshold rules) *If $\min_x R(x) > 1 - \lambda$ and $s(x, 0) = 0$ then there exists period-feasible threshold default rule $d(\bar{b})$ with associated debt issuance rule H weakly increasing in b .*

The proof is in Appendix A.

Corollary 1 *If a threshold default rule with thresholds \bar{b} is period feasible given a surplus rule $s(x, b)$ such that $s(x, 0) = 0$, then it is also feasible under the surplus rule $\tilde{s}(x, b)$ so long as for all x :*

1. $\tilde{s}(x, \bar{b}(x)) = s(x, \bar{b}(x))$
2. $\tilde{s}(x, 0) > \lambda b_s(x)/G(x) - s(x, b_s(x))$

Note that the period-feasible default rule may not be unique. It is straightforward to show that if $d_0(\bar{b})$ is period-feasible, then so too is $d_0(\tilde{b})$ for some \tilde{b} such that $\tilde{b}(x) \leq \bar{b}(x) \forall x$. That is, a government can always commit to defaulting more often. A more challenging question, discussed further next, is whether a government can promise to default less.

A useful object in understanding this is the price of debt when the government never defaults, given by:

Definition 4 (Maturity-adjusted risk-free debt price) *The maturity-adjusted risk-free debt price is a vector of prices $q^{rf}(x)$ satisfying:*

$$q^{rf}(x) = \frac{1}{R(x)} \left[\lambda + (1 - \lambda) \sum_{x' \in \mathcal{X}} \pi(x', |x, 0, \delta) q^{rf}(x') \right]$$

This can be solved via a matrix equation.

$$\mathbf{q}^{rf} = \lambda (I - (1 - \lambda) \text{diag}(\mathbf{R}^{-1}) \Pi)^{-1} \mathbf{R}^{-1}$$

Where \mathbf{R}^{-1} is the vector of inverse interest rates and Π the matrix of transition probabilities. Equipped with this, we can now ask if there is an upper limit on period feasible default rules. The next theorem describes the condition under which such a limit exists.

Theorem 2 *If $\psi = 0$ and $\pi(y|x, b', 0)$ is independent of b' and if there exists no subset Y of states such that:*

$$\sum_{y \in Y} \pi(y|x, 0) \geq \frac{\lambda + (1 - \lambda) q^{rf}(x)}{G(x) q^{rf}(x)} \quad \forall x \in Y \quad (6)$$

Then there exists some finite \bar{b}_+ such that any period-feasible threshold default rule $d(\bar{b})$ satisfies $\bar{b}(x) < \bar{b}_+(x) \forall x$.

The proof is in Appendix A. The restrictions on ψ and $\pi(y|x, b', 0)$ are not essential. Removing them replaces condition (6) with a much more complicated expression but does not alter the underlying point.

The intuition behind the condition in equation (6) is relatively simple. If there is a subset of states where the gain from a low interest-growth differential exceeds the probability of transiting to other states, then the government can commit to repay arbitrarily large debts in those states of the world, even if default is guaranteed in the other states. This is clearest when $\lambda = 0$. In this case, the requirement is that there exists no subset Y of state such that :

$$\sum_{y \in Y} \pi(y|x, 0) \geq \frac{R(x)}{G(x)} \quad \forall x \in Y$$

Then imagine that there is some state y where the probability of remaining in that state conditional on repayment is is very high and the interest-growth differential is very low, for example $\pi(y'|y, 0) = 0.96$, $\log(R(y)/G(y)) = -0.05$. Then condition (6) is satisfied. Even if default is guaranteed in all other states, the interest-growth differential net of default in state y will still be -1% . And so in this state of the world, the natural debt dynamics will always be to drift down, allowing the government to promise guaranteed repayment at all debt levels in state y , i.e. $d(y, b) = 0 \forall b$. With sufficiently low probability of transition *into* state y , the long-run average interest-growth differential can still be positive. In other words, there is no overall limit on the total amount of debt that a government can borrow, despite $\rho > 0$.

More broadly, the results in this section can be viewed as an extension of Auclert and Rognlie (2016), who derive conditions for existence and uniqueness of threshold equilibria in the canonical strategic default model. In that setting, the concavity of government preferences in the strategic default rule guarantees uniqueness. By relaxing this assumption I can only put an upper bound on threshold default rules.

2.5 Long-run stability

Period feasibility is a fairly simple requirement: that debt prices are consistent with the government's default and issuance rules, and that issuance is consistent with a period budget constraint given prices. Yet this does not fully capture standard notions of what it means for government debt to be sustainable. This is where the idea of long-run stability comes in. It is the stochastic equivalent of the idea that debt levels should not be explosive.

If d is period-feasible, then the mappings $\pi(\cdot)$ and $H(\cdot)$ together define a Markov process for (x, b') . Given (x_1, b_0) , recursive application of their one-period evolution generates a sequence of debt levels $\{b_{t+1}\}_{t=0}^{\infty}$. The sequence so-generated is long-run stable if it satisfies the following

condition:

Definition 5 (Long-run stability) *A stochastic sequence of debt levels $\{b_{t+1}\}_{t=0}^{\infty}$ is long-run stable if for any b_0 , the distribution of b_{t+1} converges to some limiting invariant distribution as $t \rightarrow \infty$.*

In what follows, I show that any threshold finite default rule with finite thresholds is long-run stable. However, the default rule where the government always repays is only long-run stable if the long-run average interest-growth differential is negative.

Intuitively, long-run stability is the stochastic analogue of debt being “non-explosive”. If debt diverges to infinity with probability greater than zero, then the debt has no limiting distribution and long-run stability fails. Note that we do not require that the limiting distribution need not be unique, just that it must exist for any initial b_0 .

One obvious alternative to long-run stability is a no Ponzi game condition, which prohibits a strictly positive net present value of the debt in the long run.¹⁵ Yet this admits explosive debt, and is thus too weak to meet any reasonable standard of debt sustainability. In particular, if the long risk-free interest-growth differential ρ is positive, then a no Ponzi game criterion is satisfied by any fiscal policy which produces a debt ratio that grows explosively so long as it grows at a rate strictly less than ρ . Indeed, computational solutions to strategic default models typically acknowledge this point – albeit often implicitly – by imposing a bounded state space for government debt. This is a much tighter restriction than a no Ponzi game condition, as it categorically rules out slowly-exploding debt paths. Indeed, this is one way to guarantee a invariant distribution, in this case one bounded by the range of the state space.¹⁶

If default thresholds are finite, then long-run stability always holds, as debt is bounded by the maximum default threshold. Thus, it is evident that:

Theorem 3 *All finite threshold default rules are long-run stable.*

This means that if we want to assess whether a *finite* threshold rule is both long-run stable and period feasible, we need only show that it is period feasible.

If a threshold rule is not finite then it is the universal repayment rule, which dictates repayment in all states:

15. This is usually motivated as the counterpart to an investor’s transversality condition. The logic is if investors are optimizing then cannot be willing to leave some part of their lifetime wealth unused, as would be implied if the net present value of long-run debt were positive.

16. And if $\rho < 0$ then the no Ponzi game condition is far too restrictive to be useful, as it requires *all* debt parts to converge to zero with probability one.

Definition 6 (The universal repayment rule) *The universal repayment rule is $d(x, b) = 0 \forall x$*

The universal repayment rule is interesting because it implies that $p(x, b') = 0 \forall (x, b')$ in any model where $\lim_{b \rightarrow \infty} s(x, b) > -\infty$. As a result, the cost of borrowing is minimized over all default rules; any level of debt can be issued risk-free. Furthermore, this rule has a realized default frequency of zero, which is rather appealing if default is costly.

The reason that period feasibility of the universal default rule does not violate Theorem 2, even if condition (6) holds, is that Theorem 2 only covers finite threshold default rules. The proof relies on a finite upper bound $\bar{b}_0(x)$ on the repayment region, and so fails for the universal repayment rule.

Theorem 4 *The universal-repayment rule is long-run stable if and only if the long-run average of $\log(R(x)/G(x))$, ρ is negative*

The proof is in Appendix A. Crucially, it is the *risk-free* long-run interest-growth differential which determines whether a policy of never defaulting produces stable or unstable debt paths. If investors are not risk-neutral, then the probabilities which determine the price q^{rf} of the risk-free long-duration bond are no longer the probabilities which determine the long-run evolution of the state, π , but rather investors' risk-neutral probabilities $\tilde{\pi}$. In Appendix A I show to first order that this changes the condition in Theorem 4 to:¹⁷

$$\begin{aligned} \rho &< (1 - \lambda) \mathbf{v}' (\Pi - \tilde{\Pi}) \mathbf{q}^{rf} \\ &= (1 - \lambda) (\mathbf{v} - \tilde{\mathbf{v}})' (I - \tilde{\Pi}) \mathbf{q}^{rf} \end{aligned} \tag{7}$$

Where: \mathbf{v} is the stationary distribution of the state; Π is the state transition matrix; $\tilde{\Pi}$ is the matrix of investors' risk-neutral state probabilities; and $\tilde{\mathbf{v}}$ the corresponding stationary distribution. Thus departures from risk-neutrality can have an impact on the conditions for long-run stability, but only to the extent that 1) debt is long-run, and 2) the long-run distribution of states \mathbf{v} and $\tilde{\mathbf{v}}$ differ.

So if some factor outside the model compels the government to issue long debt (e.g. demand from institutional investors such as pension funds), and investors are sufficiently risk-averse then the relevant threshold for ρ is not quite zero but something possibly a little larger. But this raises an obvious question: what other reason could possibly be so compelling as to mean that

17. In contrast, Theorem 3 still holds, just with the state probabilities π replaced by the risk-neutral probabilities $\tilde{\pi}$.

governments to choose to issue long-term debt when short-term debt has no fiscal cost? If $\rho < 0$, then the government could make a transfer to institutional investors to compensate them, financed by short-term debt and no future tax increases. For this reason, I focus on the simple case where the threshold for ρ is zero.

To summarize, Assumptions 1 and 2 imply that:

1. So long as $s(x, 0)$ is not too negative, then there exists a period-feasible threshold default rule d (Theorem 1);
2. We can compute such a default rule using an iterative operator T (shown in Appendix A);
3. If equation (6) as holds, then there exists an upper bound on period feasible default thresholds for s (Theorem 2);
4. All threshold rules with finite thresholds are long-run stable (Theorem 3);
5. The universal-repayment rule is long-run stable if and only if $\mathbb{E} \log R(x) - \log G(x) < 0$ (Theorem 4).

In the next section I address the key empirical issues for this model. What is the process defining $R(x)$ and $G(x)$? And does it imply that universal repayment might be possible?

3 Measuring interest-growth differentials in the long run

This section has two purposes. First, given the critical role that the sign of the long-run interest-growth differential plays in determining the existence of a debt limit, to test whether the long-run interest-growth differential is positive or negative. Second, to estimate a dynamic stochastic process for $R(\cdot)$ and $G(\cdot)$ that can be fed into the model to produce predicted debt limits.

Obviously, these two steps are related; estimation of a stationary process for interest and growth rates will generate a long-run mean for their difference. So I first estimate VARs for nominal interest and growth rates in seven advanced economies, using two different datasets: annual data from 1880, and quarterly data from 1956. I then use the estimated VAR coefficients and their variance-covariance matrix to form confidence intervals around the long-run interest-growth differential.

Of course, this approach fits the data to a particular statistical process, namely a VAR, which may be quite restrictive. To relax this assumption, I also use the technique of Müller and Watson (2016), who develop a spectral method for estimating confidence sets for long-run averages which

are robust to data generating processes with a wide variety of long-run properties. An alternative approach is that of Mehrotra (2017), who uses a probit regression to infer the probability that future interest rates will be positive at a given finite horizon, and documents the determinants of such events.

Accurately estimating confidence intervals is critical for the policy implications of this paper. Point estimates of the long-run average are negative for all countries, periods, and estimation methods. But the confidence intervals around these estimates typically include small positive numbers. In most countries, VAR-based estimates cannot reject at the 5% significance level the possibility that the long-run average interest-growth differential is greater than zero, based on both annual data from 1880 and quarterly data from 1956. Using annual data, the spectral method generates confidence intervals which are centered on negative values.¹⁸ But because this is a much more general approach, these confidence sets are wider.

3.1 Data on interest and growth rates

I use two separate datasets, each covering a set of major Advanced Economies,¹⁹ but differing in their periods and frequencies. The first dataset is that of Jordà, Schularick, and Taylor (2017). This is annual, and covers many currently advanced economies. From this, I select four: the USA, the UK, France, and Germany. These are the four major world economies during most of this period, and are the only members of the G7 for which there is sufficient uninterrupted data.²⁰ For all countries the data covers 1880-2015. The benchmark interest-growth differential is the difference between the short-term nominal rate – equal to the central bank policy rate – and the annual nominal growth rate, as implied by the results of Section 2.²¹ Years where the interest-growth differential is more than six standard deviations from the mean are excluded as outliers. Table 1 specifies the sample years remaining and Figure 5 plots the time series for all G7 countries.

The second dataset is quarterly, using average overnight central bank interest rates and national accounts data for growth. Quarterly data is typically only available relatively recently,

18. The quarterly sample is too short to use the spectral method.

19. Emerging and low-income countries are excluded, as interest-growth differentials there seem to exhibit secular trends, slowly increasing over time. This makes it hard to fit a stationary statistical model, or to think that the data has anything meaningful to say about long-run average levels.

20. Italy, Japan and Canada all have too many missing data points to produce reliable estimates.

21. In the long run, of course, the difference between interest and growth rates should be the same whether nominal or real rates are used; the nominal and real interest-growth differential are separated only by errors on expected inflation, which should be zero in a long enough sample. However, it is not entirely obvious how one would construct a long-run dataset of real rates, except by subtracting realized inflation from the nominal rate, which is exactly the same as computing the long-run *nominal* interest-growth differential.

Country	Observations	Periods
<i>Annual data</i>		
USA	134	1880-2015
UK	134	1880-2015
France	114	1880-1913, 1922-1938, 1951-2015
Germany	128	1880-1921, 1925-1944, 1950-2015
<i>Quarterly data</i>		
USA	228	1960Q1-2016Q4
UK	228	1960Q1-2016Q4
France	228	1960Q1-2016Q4
Canada	223	1961Q4-2016Q4

Table 1: Sample periods for interest and growth rates (after removing outliers)

and for some countries only very recently. This dataset therefore covers only a subset of the G7: Italy and Japan are excluded as quarterly nominal growth data is only available from the 1990s. Likewise Germany, for which quarterly data starts only in 1976. The other remaining G7 countries (the USA, UK, France, and Canada) do have sufficient quarterly data, starting in 1956 (1961 for Canada).

One important feature of this approach is that, guided by the theory in Section 2, I use risk-free rates to construct the interest-growth differential. This is in contrast to past empirical studies, which use variously the effective rate (Escolano (2010), Kim and Asonuma (Forthcoming), Mauro et al. (2015)), a long-run rate less average inflation (Mehrotra (2017), in a real model), or a maturity-weighted average of government bond yields (Blanchard (2019)).²²

The problem with these approaches is that the resulting interest-growth differential are potentially endogenous to the debt level, as they include default and term premia. Using this to assess whether debt is sustainable will therefore deliver results which depend on the current level of the debt. This cannot be right, as debt limits are an upper bound on debt, not a function of it. To explore the sensitivity of the empirical results to this point in Appendix C.3, I re-run the analysis using interest rates on long-term government bonds and effective interest rates. These affect the results materially. Using long-term rates, point estimates of interest-growth differentials are higher, presumably due to default and term premia, and are harder to distinguish from zero. With effective rates, however, interest-growth differentials as a result of unexpectedly and persistently high inflation in the 1970s and 1980s. As perpetual surprise inflation is not a plausible strategy for managing debt ratios in the long run, the effective rate also gives a distorted picture

²². Blanchard also adjusts the rate of return to reflect special tax treatment of debt. However, this is not a saving to the government; the corresponding tax expenditure merely increases the primary deficit by an identical amount. And so tax-adjustment is misplaced when measuring the sign of the interest-growth differential for debt sustainability purposes.

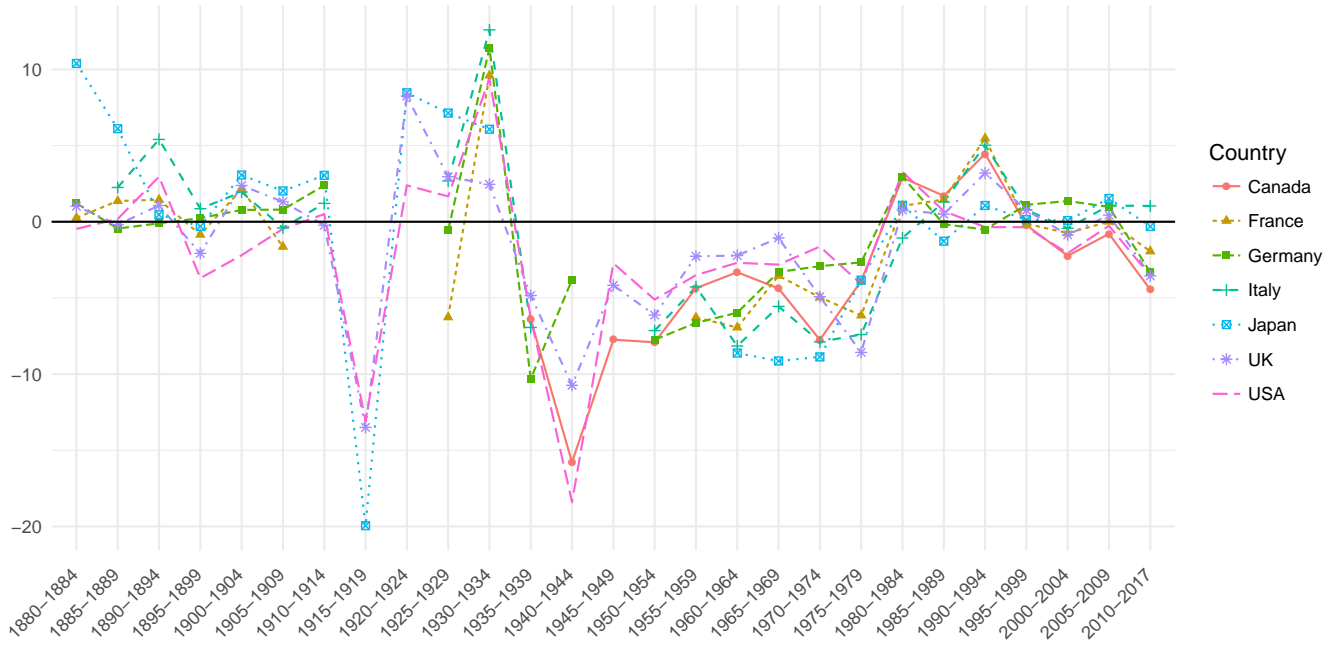


Figure 5: Five-year average annualized interest-growth differentials for a sample of advanced economies: 1960-2016

of fiscal sustainability.

3.2 VAR based tests of the long-run interest-growth differential

Table 2 presents summary statistics for the historical sample, including the interest-growth differential for each country. Given the broad applicability of the central limit theorem, the sample average of the interest-growth differential is an unbiased estimator of the long-run average differential. This is negative for all countries in both samples. If observations were uncorrelated, the sample standard deviation divided by the square root of the number of observations would give the standard error of the estimate, and we would with high confidence reject the possibility that interest growth differentials are positive in the long run (except for Germany). But the observations clearly are auto-correlated, so we instead need to formulate a test based on an estimated dynamic process for interest-growth differentials.²³

A natural starting point is to estimate a simple vector autoregression (VAR) for interest and growth rates. The long-run mean is then just a combination of the estimated parameters.²⁴ Table 3 presents the results of this exercise using annual data, estimated via maximum likelihood. The

23. In small samples, the time series dynamics may also prove informative about the point estimates of the long-run mean, as this is the point to which a stable time series will drift towards.

24. For any VAR $y_t = a + \sum_{i=1}^n A_i y_{t-i}$, the long-run mean of y_t is given by $\bar{y} = (I - \sum_{i=1}^n A_i)^{-1} a$

Country	n	Interest Rate		Growth Rate		Interest-Growth Differential			
		Mean	SD	Mean	SD	Mean	SD	Fraction positive	SD/ \sqrt{n}
<i>Annual data</i>									
USA	134	3.96	2.84	6.06	7.46	-1.96	7.59	0.34	0.66
UK	134	4.25	3.33	5.86	6.37	-1.50	5.81	0.38	0.50
France	114	4.39	3.17	6.24	7.96	-1.71	7.24	0.39	0.68
Germany	128	4.17	2.16	7.98	120.60	-3.88	120.12	0.41	10.62
<i>Quarterly data</i>									
USA	228	5.03	3.49	6.34	3.87	-1.31	4.09	0.34	0.27
UK	228	6.75	3.86	7.75	6.21	-1.00	5.48	0.43	0.36
France	228	5.73	3.86	6.99	5.82	-1.26	5.19	0.34	0.34
Canada	223	5.97	3.66	7.13	5.27	-1.16	5.49	0.37	0.37

Table 2: Summary statistics for interest and growth rates

point estimates broadly agree with the sample means in Table 2, around -2 for the USA and slightly closer to zero for the others. As expected, the standard errors are larger than those reported in the final column of Table 3 but are still sufficiently small that a Wald test would – in most cases – come close to rejecting the hypothesis that the long-run average interest-growth differential is positive. For example, the z -statistic for the USA is $-2.15/1.1 = -1.96$, almost exactly the 2.5% critical value.

This simple approach is deficient in two obvious ways. First, given that the estimates are potentially contaminated by model mis-specification error, one lag alone may not be appropriate. Second, a Wald test is relatively restricted and is based on a linear approximation to the likelihood function, thus permitting only one dimension of the parameter space to vary at a time. The more general, but more computationally intensive, likelihood ratio test relaxes this assumption, allowing all model parameters to vary.

I therefore extend my analysis by performing a likelihood ratio test on a longer-lagged VAR.²⁵ If θ represent Formally, if θ be the estimated interest-growth differential for some VAR. Then the hypothesis that we wish to test is:

Null hypothesis The long-run mean is θ_0

Alternative hypothesis The long-run mean is $\theta \neq \theta_0$

Where θ_0 is some arbitrary number. Typically, one specifies that $\theta_0 = 0$, therefore testing if the estimated coefficient is statistically distinguishable from zero. But here we want to ask:

²⁵. Indeed even a more sophisticated linear model may still be mis-specified, so in the next section I relax the model specification yet further.

	USA		United Kingdom		France		Germany	
	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>								
Constant	4.79	0.51	1.67	0.16	1.47	0.29	3.59	0.86
	(1.106)	(0.255)	(0.689)	(0.178)	(0.98)	(0.204)	(1.148)	(0.295)
Growth (-1)	0.43	0.03	0.65	0.06	0.52	0.02	0.65	0.04
	(0.081)	(0.019)	(0.071)	(0.019)	(0.088)	(0.019)	(0.071)	(0.019)
Int. rate (-1)	-0.34	0.82	0.07	0.88	0.32	0.9	-0.37	0.73
	(0.214)	(0.049)	(0.136)	(0.035)	(0.202)	(0.042)	(0.239)	(0.062)
<i>Innov. covar.</i>								
Growth	44.71		21.77		33.87		27.39	
	(5.618)		(2.74)		(4.678)		(3.594)	
Int. rate	3.29	2.41	1.45	1.49	2.15	1.6	0.91	1.88
	(0.968)	(0.304)	(0.523)	(0.187)	(0.749)	(0.22)	(0.673)	(0.248)
<i>Long-run</i>								
Mean	6.11	3.95	5.6	3.97	5.7	3.97	5.88	4.07
	(0.901)	(0.737)	(1.341)	(1.23)	(1.76)	(1.365)	(1.195)	(0.458)
Difference	-2.15		-1.63		-1.72		-1.81	
	(1.1)		(0.911)		(1.053)		(1.209)	
Log likelihood	687.5		610		542.1		596.1	
Observations	134		134		114		128	

Table 3: 1-lag VAR for sample of countries. Annual data 1880-2013. Robust likelihood-based standard errors in parentheses.

what is the largest interest-growth differential which cannot be rejected by the data? In other words, we wish to find the critical value of θ_0 at which the null hypothesis is rejected for a given significance level. This is then the largest interest-growth differential still consistent with the data.

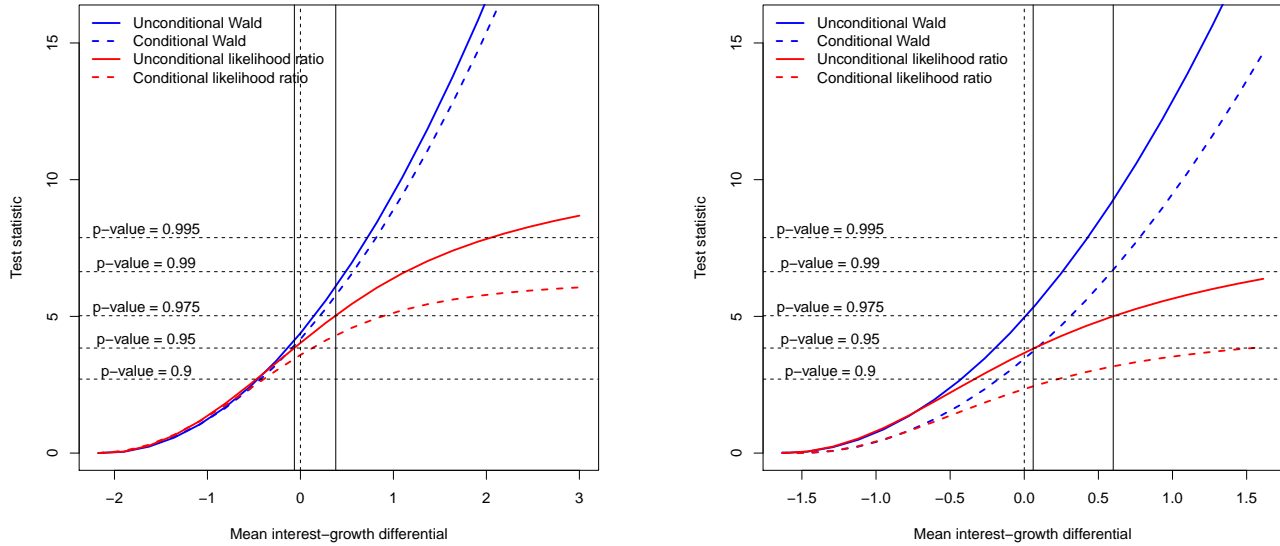
Figure 6 represents this exercise graphically for the USA using a four-lag VAR.²⁶ Each panel of Figure 6 shows four different test statistics for the hypothesis described above, with the null value θ_0 plotted on the x-axis. So when the x axis is zero, for example, the four lines show four different test statistics for the hypothesis test that the long run average interest-growth differential is zero. The test statistics are minimized when θ_0 is the value implied by the unrestricted estimation²⁷. As the point estimates are negative, so are the values at which these minima are attained. The solid red line is the preferred test statistic, and shows the likelihood ratio test statistic using the unconditional likelihood. For annual data, shown in Figure 6a, this crosses the 95% critical values for the χ_1^2 test at -0.06. So using this test we can conclude that the data allows us (just) to reject the hypothesis that the quarterly interest-growth differential is positive at the 5% significance level. But at any more restrictive significance levels, this test will fail. Furthermore, using the quarterly post-war data, we cannot reject the hypothesis that the interest-growth differential is positive; the 5% critical value is 0.06.

The solid blue line is the Wald test for the same models and hypotheses. While this rejects a larger set of hypotheses, there are valid grounds to prefer the likelihood ratio test. While simpler to calculate, the Wald test is only a local approximation to the likelihood ratio test, and for large deviations from the null (as shown here) can be inaccurate. Intuitively, the Wald test evaluates the hypothesis that that long-run mean has changed with all other model parameters held fixed. In contrast, the likelihood ratio test re-estimates the full model subject to the constraint that the mean interest-growth differential is that proposed in the null. As a result, the likelihood ratio test allows other parameters to change to fit the variation in the data. Given we are interested in testing what statistical model can best fit the data, this seems a more appropriate test.

Figure 6 also shows, in the broken lines, the test statistics for the conditional likelihood. The difference between the conditional and unconditional likelihoods is that the former uses only information about changes in the time series, whereas the latter uses information about the level. Specifically, the unconditional likelihood also includes a term for the initial level of the series, under the assumption that it is drawn from the long-run stationary distribution. In sufficiently

26. Both the Akaike information and likelihood ratio tests for lag length suggest that all countries in the sample are best represented by a four-lag VAR when run on quarterly data. For annual data, the optimal lag length varies, but is typically 3 or 4. For simplicity, I therefore only show results with four lags. Using other lag lengths does not materially affect the results.

27. The results of the longer-lagged VARs are reported in the Appendix in Tables 9 and 10. The unrestricted estimates of the interest-growth differentials are very similar to the one-lag case.



(a) Annual data 1880-2015

(b) Quarterly data 1956-2016

Figure 6: Test statistics and critical values for an unrestricted four-lag VAR for the USA. Vertical lines show the 5% critical values of the unconditional LR test.

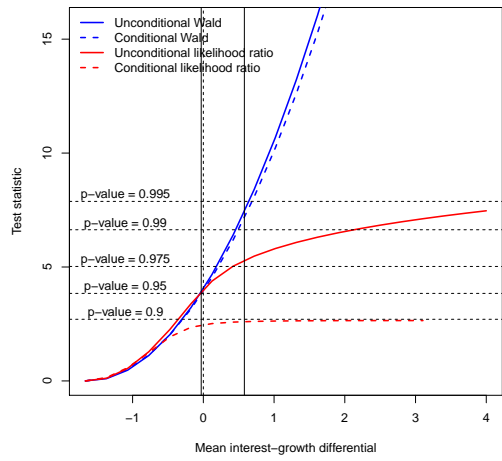
long samples, these are the same, as the long-run level of the series is just a function of past innovations. As a result, so the computational simplicity of the conditional likelihood means that it is usually preferred. But when sample is small, or the data is highly persistent, there can be meaningful information in the initial level of the data. Ignoring this can produce conditional statistics that are quite misleading, as most obviously in the quarterly sample. In particular, θ_0 can be very large and positive but not rejected. In Figure 6a (and even more so in Figures 7a and 7b) this effect shows up as the leveling out of the conditional likelihood test statistic in historical data (the broken red line) for the mean differential greater than about 2.

Figures 7 and 8 show the equivalent tests for the other four countries in the sample. These are very similar, suggesting that we cannot reject the possibility that in the long run, risk free rates will exceed nominal growth rates. At the 5% confidence level, conservative estimates of this average vary from around 0 to 0.4 percentage points per year for the UK, France, and Germany, and nearer 0.8 for Canada (due in part to the shorter sample).²⁸

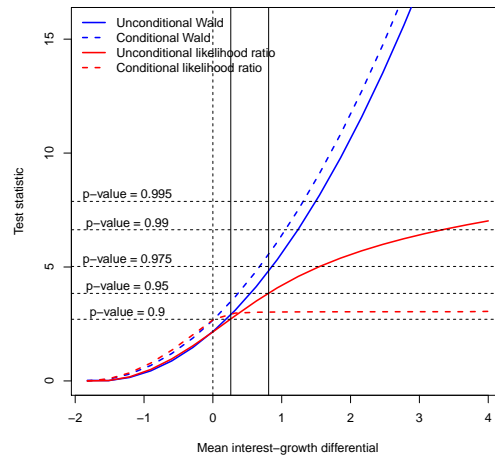
3.3 Spectral inference on the long run interest-growth differential

Müller and Watson (2016) develop a spectral method for dealing with uncertainty surrounding long-run predictions. These approach is designed to answer questions such as: what is a 90%

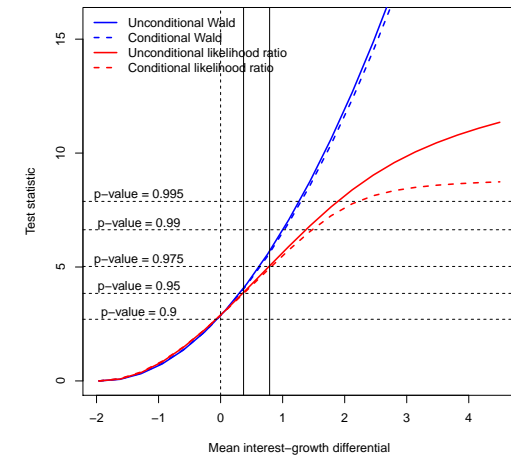
28. Table 10 (in Appendix C.2) reports the parameter estimates from the restricted VARs, and concludes that the dynamics of the statistical process are little-changed by the level shift imposed by the restriction.



(a) United Kingdom

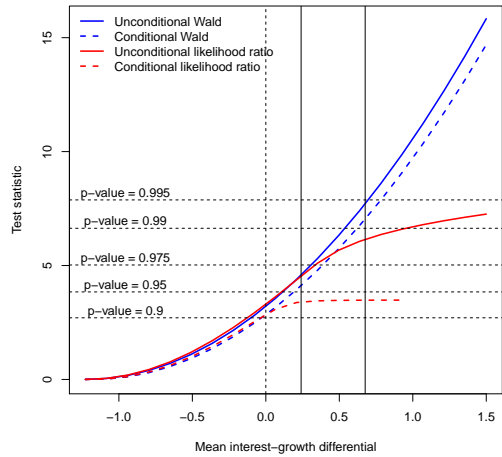


(b) France

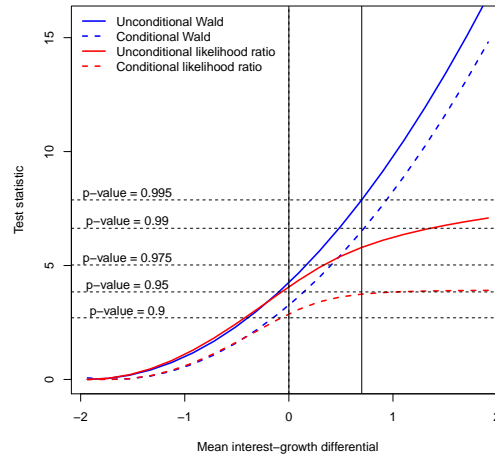


(c) Germany

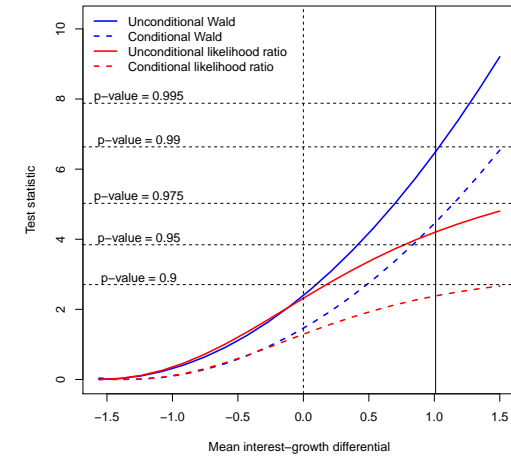
Figure 7: Test statistics for hypothesis test of long run means. Annual data 1880-2015.



(a) United Kingdom



(b) France



(c) Canada

Figure 8: Test statistics for hypothesis test of long run means. Annualized quarterly data 1960Q1-2016Q4 (from 1961Q4 for Canada).

confidence interval for GDP growth (or any other univariate time series) over the next 50 or 100 years?

The advantage of approaching this problem from a spectral perspective is that it avoids the need to assume or fit specific functional forms for the data generating process. Instead, the data are decomposed into movements at different frequencies, and standard errors of the future mean are computed using the low-frequency movements. The intuition behind this calculation is simple: because it is by definition an average, a central limit theorem applies to the forecast of the average level of a time series over the future. The central insight of Müller and Watson (2016) is to show that the variance of this limiting distribution can be expressed in terms of the spectral weights near zero, i.e. the low-frequency movements of the time series.

This focus on the spectral properties of the series mean that the method is robust to a wide set of assumptions about the long-run movements of a series, such as intermittent regime changes, or fractional integration. Beyond even this, it is particularly relevant in the current setting. Low-frequency variation not only seems characteristic of the time series for interest growth differentials (see Figure 5), but is also the type of variation which is hard to estimate using a VAR. This is because VAR estimation infers the long-run properties of a time series from their estimated innovations, which for a n -lag VAR depend on only n consecutive periods. Of course, fitting longer-lagged VARs to the data will go some way to capturing low-frequency variability, but this solution is hampered by the ever-expanding number of coefficients which would need to be estimated.

In contrast, the Müller-Watson method decomposes a time series into fluctuations at different frequencies, and focuses specifically on the slowest-moving components. This includes information on the relationship between data points at distant horizons. So this captures the very slow-moving components of the data much more as effectively than a VAR. Furthermore, determining the degree of partial integration of the series, and hence its stationarity, is an important step in the Müller-Watson method. As such, this method also provides a partial cross-check on the stationary assumptions embedded in the VAR analysis.²⁹

Table 4 presents the mid-point and upper bound of three different Müller-Watson spectral estimates for the average interest-growth differential in the 100 years following the end of the sample using annual data³⁰. The estimates are increasingly general, varying in their assumptions about how to fit a spectral density to the data. The first estimate, the $I(0)$ estimate, assumes that the (log) spectral density is uniform, and would be accurate if the data were integrated of order zero. This is the most restrictive assumption, and is equivalent to assuming that the data

29. For more on the relationship between partial integration and stationarity see Parke (1999)

30. Appendix C.4 discusses the application of this approach to the quarterly data, and concludes that the sample is too short to usefully apply this method to quarterly data.

admits a moving average representation. The second line allows for estimation of the degree of partial integration of the time series by Bayesian estimation assuming a flat prior over the degree of partial integration in $[-1, 1]$. The third line, labeled “Bayes superset” expands the Bayesian estimates to guarantee that they have frequentist coverage. That is, the $(1 - \alpha)\%$ confidence set includes the true average $(1 - \alpha)\%$ of the time, under the true data distribution. The last line of the table reports the maximum likelihood estimate of the degree of fractional integration, which we discuss further below.

	USA		UK		France		Germany	
	Mid.	95% c.v.	Mid.	95% c.v.	Mid.	95% c.v.	Mid.	95% c.v.
I(0)	-1.90	0.90	-1.50	1.10	-2.00	0.90	-3.40	2.50
Bayes	-1.90	1.60	-1.30	2.20	-1.30	4.40	-2.90	5.10
Bayes superset	-1.50	3.50	-0.90	5.10	-0.60	10.10	-0.80	13.80
<i>Memo</i> : 4-lag VAR	-2.12	-0.07	-1.63	-0.03	-1.65	0.26	-1.92	0.27
ML fractional integration	-0.10		0.00		0.30		0.00	

Table 4: Midpoint and 95% upper critical value of Mueller-Watson prediction ranges for average interest-growth differential of next 100 years. Annual data 1880-2015.

The results in Table 4 are broadly consistent with the VAR analysis. For most specifications, the midpoint of the confidence interval are around -1 to -2 percent for the I(0) and Bayes measures, in line with the the point estimates from the VAR setup. The upper boundaries of the confidence set, comparable to the 95% critical values marked in Figures 6 and 8, vary considerably across the the measures. For the US, UK and France, the 95% critical value for the I(0) estimators are typically one percentage points higher than the VAR estimates (nearer two percentage points for Germany). This is not surprising - the spectral method is less restrictive, so should result in wider confidence sets. However, as the spectral estimates become increasingly general, moving from I(0) to Bayes, and then to the Bayes superset, the the upper critical values become ever larger. This occurs because the estimation is increasingly sensitive to outliers. France and Germany in particular experienced some rather extreme fluctuations in their nominal growth rates either side of World War II. This drives the very large upper bounds on the confidence sets, particularly for the the Bayes superset.

So how should we interpret the difference between the Bayesian and I(0) estimators? The source of the difference is the differing functional forms fitted to the sample spectral densities. In the I(0) case, this spectral density is forced to be flat. But in the Bayesian setting, the

functional form for the spectral density is estimated from a much more flexible family. This estimated Bayesian spectral density depends on two factors: the prior, and the likelihood. The fifth line of Table 4 helps separate out the contributions of each of these factors. It shows the maximum likelihood estimate of the degree of fractional integration. In all cases this is close to zero, and always less than 0.5 in absolute magnitude (the threshold between stationarity and non-stationarity). This means that the likelihood contribution to the Bayesian estimates (that is, the part informed by data rather than imposed) would select a model very close to $I(0)$. So the difference between the Bayesian and $I(0)$ estimates is essentially due to the flat prior imposed on the family of spectral densities initially. In this circumstance, both estimates reflect data-consistent assumption about the degree of integration, just that the $I(0)$ estimate does so more parsimoniously. And so it seems reasonable to prefer the estimates coming from this method over the Bayesian ones.

3.4 Discussion

To conclude this section, Figure 9 summarizes the empirical work in terms of the main estimates for the long run mean interest-growth differential, and the 95% upper bound. As a result of the foregoing discussion, only the $I(0)$ spectral estimates are shown.

Overall, the message of Figure 9 is clear. Point estimates of the long-run interest-growth differential are negative. This is robust across countries, periods, and estimation methods. The upper bounds of confidence sets for this average are usually positive. For countries with long, unbroken datasets and few extreme events (UK, USA, France) we can be more precise: both VAR-based and spectral estimates agree that the largest plausible value for the interest-growth differential over the long run is somewhere between 0 and 2 percent per year. Appendix C.3 shows that these basic findings are also robust to using alternative interest rate measures.

This finding represents a very serious challenge to current models of debt sustainability. In light of these empirical findings, the most likely implication of the vast majority of models of debt sustainability is that advanced economy governments can borrow unlimited amounts without fear of default. This seems unlikely, and suggests that these models are wrong in some fundamental way. Other mechanisms instead play an important role in determining maximum sustainable debt levels. These might include equilibrium selection criteria (in the spirit of Cole and Kehoe (2000)) or fiscal policies which affect long-run interest-growth differentials (as in Blanchard (2019)).

And yet, what if one is not ready to entirely reject the standard, solvency-based models? After all, there is at least some uncertainty over the sign of the long-run interest-growth differential. And contemporary solvency-driven models of debt sustainability remain a central part of the policy discussion. So in the next section, I ask is there some way usefully retain these for policy

advice? The answer – overwhelmingly – is no.

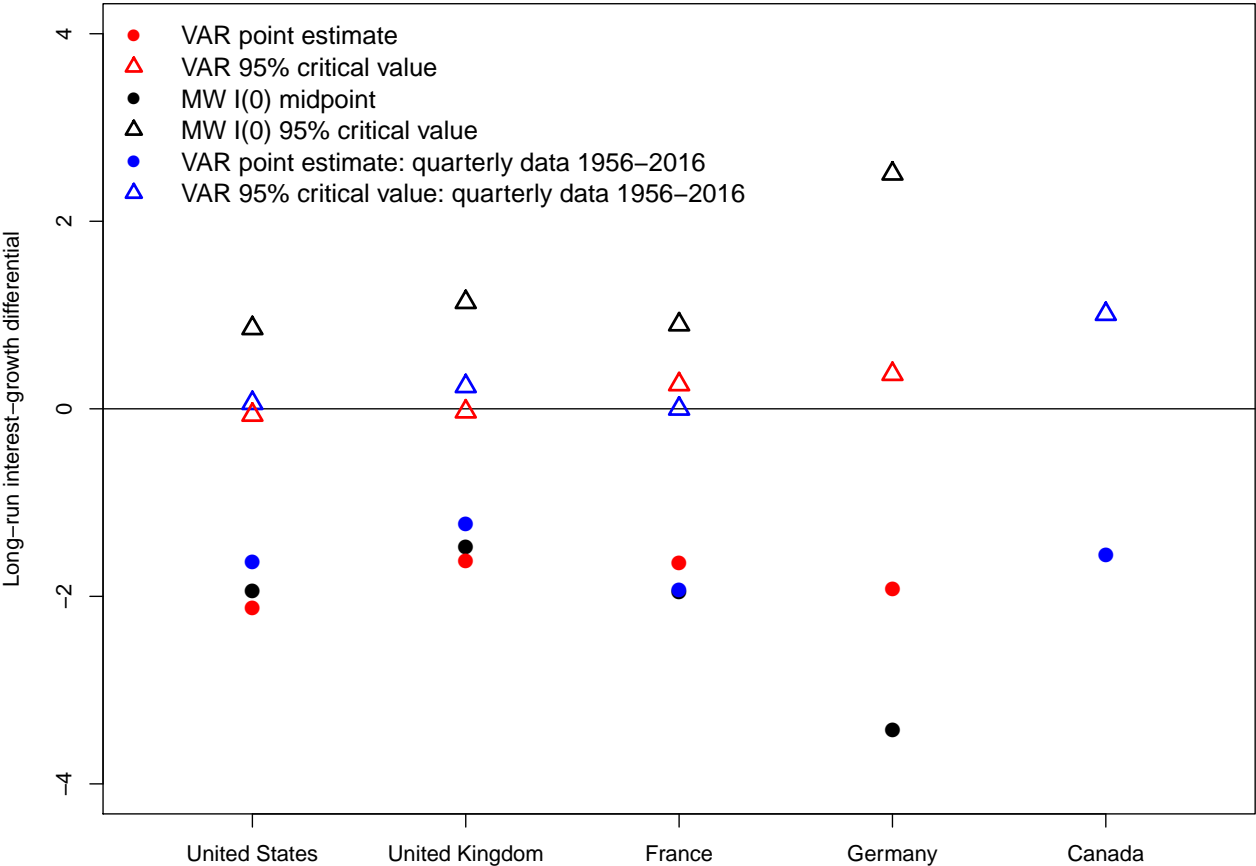


Figure 9: Comparison of point estimates and upper boundaries for interest-growth differentials in five advanced economies. Data annual 1880-2015, except where otherwise noted.

4 Estimated debt limits

The lesson of the preceding section is that, while point estimates of long-run interest-growth differentials are negative, there is considerable uncertainty about these estimates and that they might in fact be positive. The aim of this section is to ask: what are the quantitative implications of this uncertainty if we are not yet ready to reject the standard set of models of debt sustainability?

I answer this question by exploring the solution of the model in Section 2 under a range of statistical processes for the interest-growth differential. These statistical processes vary by their long-run mean. The range of long-run means that I consider is not arbitrary, though. Instead, it is described by the upper percentiles of the confidence intervals estimated in Section 3. Thus, this section is essentially a sensitivity test of the key parameter in a wide class of models, with range of experimentation informed by statistical evidence.

I focus on the USA as an example, as Jordà, Schularick, and Taylor (2017) include a long, unbroken time-series of annual data on surpluses, debt, interest and growth rates. While the USA does issue some index-linked bonds, they are a relatively recent phenomenon by historical standards (first issued in 1997) and only constitute around 9% of outstanding liabilities.³¹ I also allow explicitly for long-term debt (i.e. $\lambda < 1$).

4.1 Quantitative implications

I calculate a discrete-state Markov process for interest and growth rates which mimics a one-lag VAR for US risk-free interest and growth rates in the post-war period, 1948-2016.³² The VAR for interest and growth rates is then discretized into 13 nodes, producing 169 possible levels of the realized interest-growth differential.³³ This discretization locates nodes along the axis of (unconditional) correlation of interest and growth rates, allowing a relatively sparse discretization to cover the data (see Figure 18 in Appendix D).

The matching realized risk-free interest rates and growth rates to their nearest values in the discretized values $R(x)$ and $G(x)$ space implies a sequence of past values of the state variable x

31. Some \$1.5trn of \$16.1trn at end-2018.

32. When conducting inference in section 3 I focused on the results from a four-lag VAR. As the size of the state space for x_t is the number of nodes to the power of the number of lags, I use a one-lag VAR for reasons of computational simplicity. Using a four-lag VAR with 13 nodes for interest and growth rates would mean solving the model for $13^4 \simeq 28,561$ values of x .

33. This occurs due to a timing effect. The relevant interest-growth differential for the debt limit in period t is the difference between the nominal interest rate in period t and the nominal growth rate realized in period $t + 1$. So for every one of the 61 interest rates between this period and the next, there are 61 possible growth rates over the same period. And therefore $13^2 = 169$ possible interest-growth differentials.

during the sample period. Figure 11 shows the resulting fitted values for the risk-free interest-growth differential, $R(x_t) - G(x_t)$. Appendix D includes further details of the discretization process and tests of its accuracy.

I assume the the surplus function is a spline, cubic for $b < b_2$ and constant for $b \geq b_2$:

$$s(x, b) = \begin{cases} f(b) + \eta(x) & \text{if } b < b_2 \\ s_2 + \eta(x) & \text{if } b \geq b_2 \end{cases} \quad (8)$$

Where $\eta(x)$ is a state-specific shift which introduces uncertainty over the realized surplus. The function $f(b)$ is a cubic, parameterized by: the intercept s_1 ; the peak (b_1, s_1) ; and the constant threshold (b_2, s_2) . Thus, it satisfies:

$$\begin{aligned} f(b_1) &= s_1 & f'(b_1) &= 0 \\ f(b_2) &= s_2 & f'(b_2) &= 0 \\ f(0) &= s_0 & s_1 &\geq \max(s_0, s_2) \end{aligned}$$

This is relatively flexible functional form satisfying Assumption 2, and is illustrated in Figure 10.

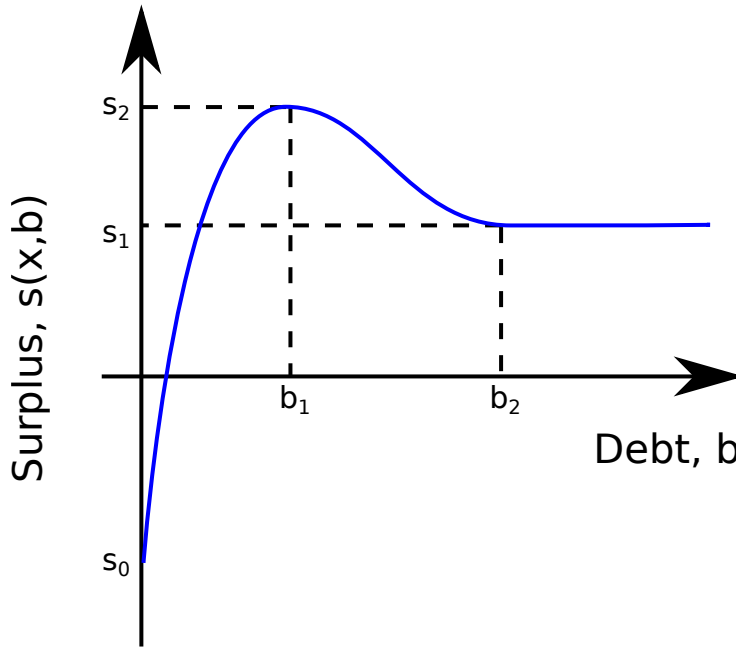


Figure 10: Estimated surplus function (with $\eta(x) = 0$)

I estimate the parameters of the surplus function using nonlinear least squares to match

primary surpluses in the postwar period. The estimates for the surplus function are shown in Table 5 and the fitted values for both the interest-growth differential and the primary surplus in Figure 11. Three points stand out from these. First, that the fitted values broadly pick up the trends in surpluses and interest-growth differentials over the past half-century. Interest-growth differentials are typically negative prior to 1980, when they turn positive before slowly declining to negative levels after the 2008 global financial crisis. Likewise with the primary balance, which is positive through the 50s and 60s, negative until the 1990s, then positive until large negative shocks following the 2001 and 2008 recessions.

Second, that the data suggest that a strictly concave surplus function is appropriate. By selecting $s_2 = s_1$, the data reject the idea surpluses decline at high debt levels (i.e. a strict version of the “fiscal fatigue” hypothesis of Ghosh et al. (2013)).³⁴ The absence of standard errors in Table 5 follows from this. Because b_2 has no effect on the surplus function when $s_1 = s_2$, and because the constraint $s_2 \leq s_1$ holds with equality, these parameters are not identified by the data: so long as these constraints remain, other values would not affect the model fit.

Third, while the standard errors suggest that the coefficient of the surplus function are relatively precisely estimated, these should not be interpreted as causal. The aim here is simply to estimate a surplus function which broadly captures the reduced-form surplus process for the United States.

	Value	Standard error
s_0	-3.2	0.2
s_1	5.0	0.7
b_1	252.7	33.0
s_2	5.0	
b_2	254.8	

Table 5: Surplus function estimation results

The surplus level shock $\eta(x)$ is chosen to equalize the conditional averages of the estimated surplus function and the data in each state. That is, that $\sum_{t|x_t=x} (s(x_t, b_t) - s_t) = 0 \forall x$. The key unconditional cyclical correlations are shown in Figure 12. This shows that fiscal policy is counter-cyclical, running larger deficits when growth is low. The correlation with interest rates is weaker but still positive, likely reflecting the fact low interest rates are correlated with low growth and make debt both more affordable.

Table 6 provides further model validation and explanation. The first two lines compare the

34. Fournier and Fall (2017) also estimate fiscal reaction functions, using a piecewise-linear approach, and solve for cases where the long-run interest-growth differential is negative. As in Ghosh et al. (2013), though, this relies on assuming that the government can commit to infinitely large deficits.

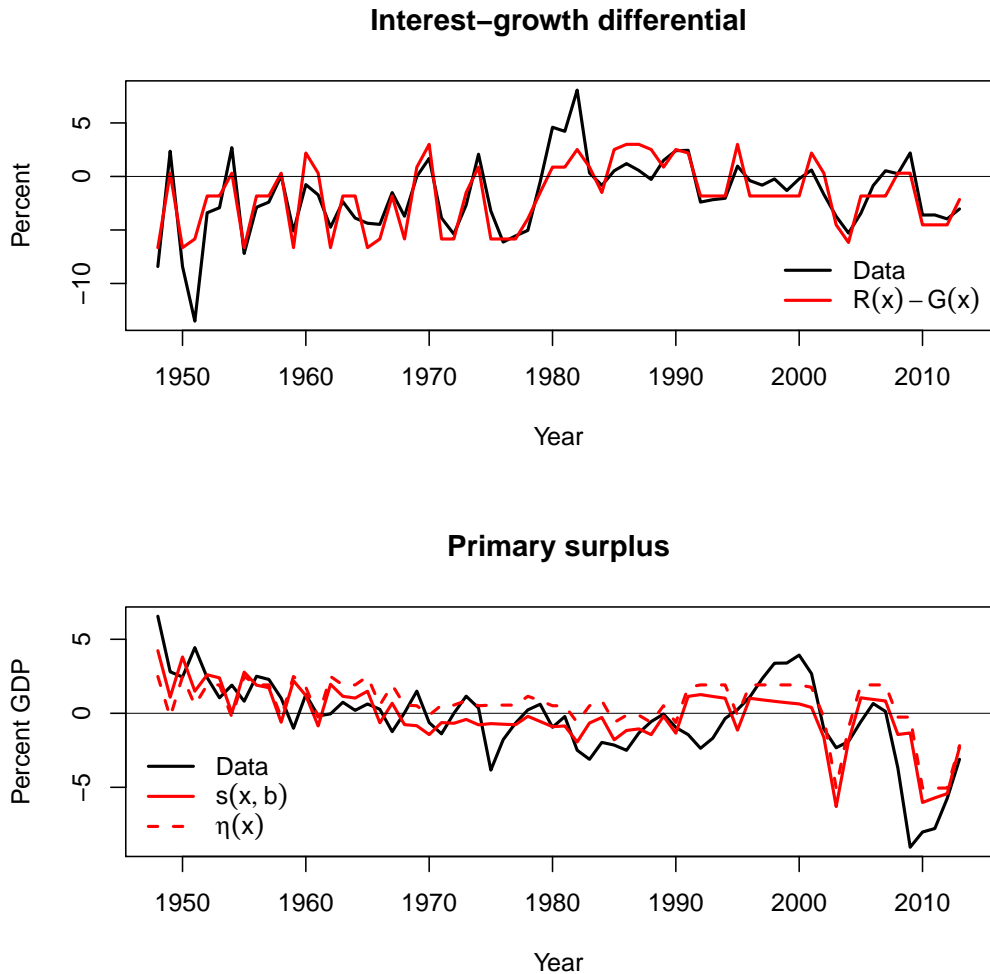


Figure 11: Fitted interest-growth differentials and surpluses)

correlation of fitted surplus to that in the data. These are not an explicit target of the estimation. Yet they are very close, providing an additional external cross-check on the model. The negative correlation of surpluses with the interest-growth differential highlights the trade-off between cyclicalty and sustainability. Counter-cyclical policy means a fiscal expansion when growth is low and hence future surpluses are lower, so debt sustainability more at risk.

The third line shows the correlation stemming purely from the additive surplus shocks. This allow a decomposition of the driver of the surplus response. That due solely to the state-contingent shock $\eta(x)$ has little correlation with the interest-growth differential. The negative correlation in the model is therefore driven by the debt response $f(b)$. This might seem somewhat curious; after all, one would expect that all else equal higher interest-growth differentials would produce increasing debt and thus larger surpluses via $f(b)$. Yet this is largely a function of the initial

conditions: with a high debt level early in the sample, the surplus function does much of the work of reducing debt, driving the negative correlation between fitted surpluses and interest-growth differentials. Over a longer simulation, the correlations are in line with one's intuition: that there is a negative correlation between surpluses and interest-growth differentials but that this is driven by cyclical shocks (i.e. the $\eta(x)$) rather than the endogenous response to debt levels).³⁵

	Risk-free interest rate	Growth rate	Interest-growth differential
Surplus, data	0.10	0.29	-0.18
Surplus, model	0.11	0.26	-0.14
$\eta(x)$, model	0.28	0.35	-0.06

Table 6: Correlations of government surpluses, data and fitted, USA 1947-2016

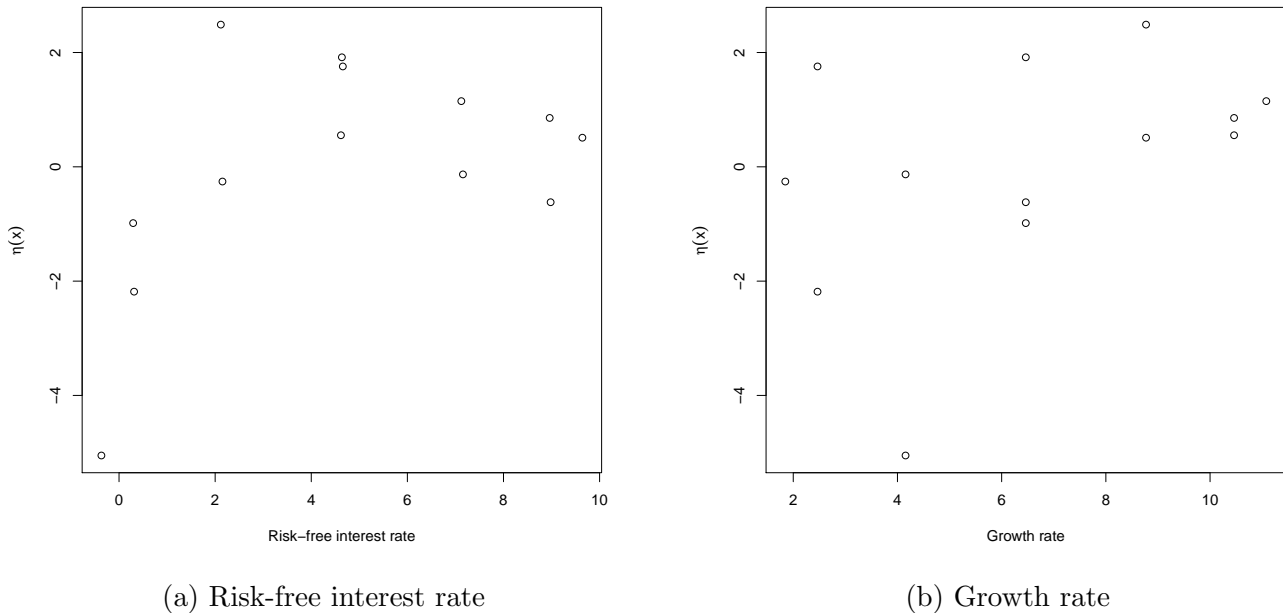


Figure 12: Cyclical correlations of the surplus shock $\eta(x)$.

Estimating a surplus function directly is a modeling choice. The obvious alternative is to estimate the parameters of some optimal policy model. Yet this has few advantages over the current approach at the cost of being more opaque. The paucity of advantages arises because the deep parameters ultimately map into a state-dependent rule for surpluses. If that reflects the data then this can be recovered by direct estimation. If not, then the model does not fit the data well.

³⁵. Over a 100,000 period simulation, the correlation of interest-growth differentials with surpluses is -0.29 and with $\eta(x)$ is -0.24 .

Optimal policy methods are also more opaque than direct estimation. Debt limits are crucially dependent on the out-of-equilibrium surplus rule – investors price debt by contemplating the hypothetical of what the government *would* do if debt *were* at a certain level. With optimal policy models the government can often guarantee repayment by promising potentially implausible surpluses.³⁶ With direct estimation of the surplus function, this relationship is entirely transparent and it is entirely evident from the estimated parameters.

Three model parameters remain: the fraction of debt repayable each period, λ ; and the repayment fraction ψ . I set $\lambda = 0.2$, in line with average US debt maturity. The U.S. Treasury publishes historical average debt maturity by month from 2000. The average maturity is almost exactly five years. The equivalent model statistic is the average duration of the bond coupon payments, $1/\lambda$, thus I set $\lambda = 0.2$. In line with most of the literature on sovereign default, I assume that investors are risk-neutral, in common with much of the literature, although I experiment with this in the various robustness exercises [TBC]. Finally, I set $\psi = 0.9$, a standard value for advanced economies (see Benjamin and Wright (2009)). Table 7 summarizes these calibrated parameters.

Parameter	Interpretation	Value	Reason
λ	Debt share repaid each period	0.2	Average U.S. debt maturity of five years
ψ	Repayment share in default	0.9	Benjamin and Wright (2009)

Table 7: Calibrated parameters

4.2 Implied historical debt limits

I now use the model to produce estimated historical debt limits. There is an obvious tension in doing this, as the debt limit under the stochastic process estimated for interest and growth rates is trivial. The long-run average risk-free interest-growth differential is negative – around -1.8% per year – so Theorem 4 says that the government can borrow as much as it likes without losing market access. That is, the debt limit is infinite.

Yet there is considerable uncertainty over the true value of the long-run differential, as discussed at length in Section 3. So in order to compute conservative-but-plausible debt limits I therefore aim off the estimated process for interest and growth rates, increasing the interest rate

³⁶ For example, in the canonical strategic default model of Arellano (2008) the government can potentially confiscate 100% of output via a lump-sum tax.

in all states of the world. This yields an average interest-growth differential is small and positive; in the baseline solution the long-run average interest-growth differential is 2% higher than the point estimate, at 0.2%. This is in line with the 5% critical value for the likelihood-based tests of Section 3.

Figure 13 shows the implied historical debt limits from this baseline solution. Two points are evident. First, that implied debt limits as much higher than the debt ratio in the data. This is despite forcing interest-growth differentials to be at the outer limits of what is empirically plausible. Second, that the estimated debt limits barely vary at all over time. Even as decadal average interest-growth differentials swung from negative and positive and back again, the corresponding theoretical debt limit varied by only a little more than 10% (around 25% of GDP). And despite the sharp decline in policy rates during and after the 2008 recession, the estimated debt limit moved by barely 5 % of GDP during this period. In otehr words, short-term fluctuations in the interest-growth differential have little quantitative impact on sustainable debt levels.

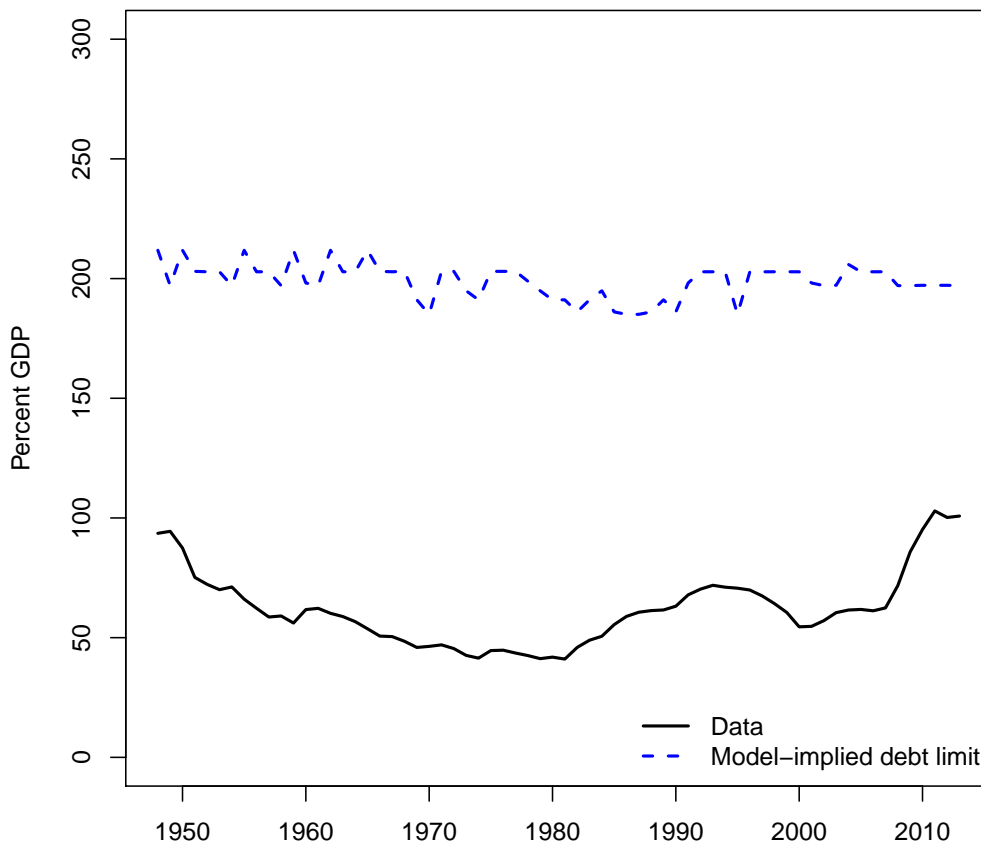


Figure 13: Historical debt and model-implied limits

The impact of short-term movements in the interest-growth differential on maximum sustainable debt levels is highlighted further in Figure 14, which plots the cross-sectional variation of the debt limit with the interest-growth differential. It shows that, on average,³⁷ a one percent increase in the current interest-growth differential is associated with a two percentage point reduction in the affordable debt-GDP ratio. To put this into context, the US debt ratio increased by 40 percentage points between 2007 and 2011, implying that an offsetting increase in the debt limit would have required a huge fall in the interest-growth differential – around 20 percentage points.

The mechanism that determines the debt limit is the loss of market access. This is illustrated in Figure 15, which shows the model-implied spread on government debt over the risk-free rate. This is almost zero until around 180% of GDP, as the risk of default on even a five-year bond is minimal. It then rises with the likelihood of exceeding the debt limit in the next period. In states where the government is certain to exceed the debt limit in the next period, the spread diverges to infinity.

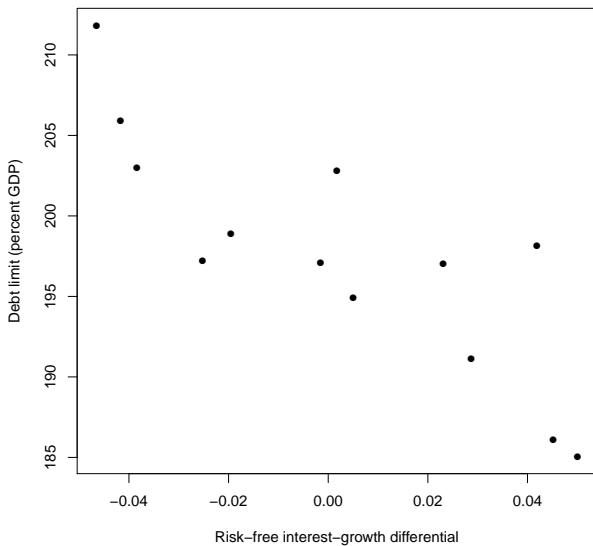


Figure 14: Estimated debt limits: cross section.

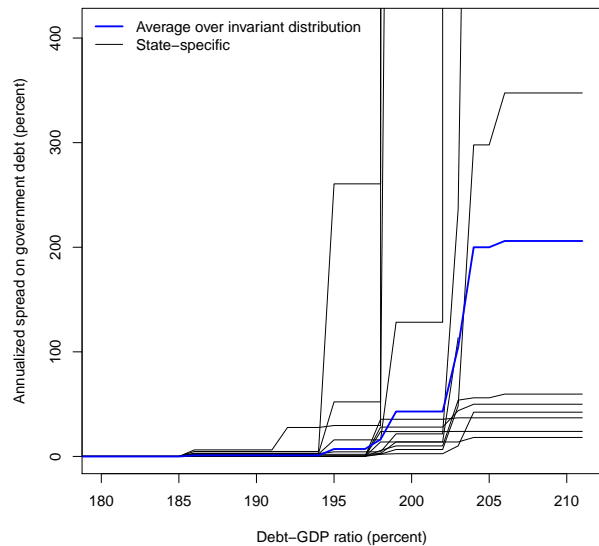


Figure 15: Model-implied spread on government debt over risk-free rate.

37. Over a long period of time, as opposed to the specific realization of history in Figure 13.

4.3 Sensitivity to the long-run interest-growth differential

To give some context to the changes in historical debt limits, Figure 16 plots the average debt limit as a function of the average long-run interest-growth differential. Here, the range of the x axis spans is equivalent to a one percentage point change in the one-sided significance test, increasing from the 97.5th to the 98.5th percentile of the confidence set around the point estimate of the long-run differential.³⁸

In statistical terms, then, this is a tiny perturbation in the long-run differential. Yet the corresponding change in the debt limits is huge. The average debt limit ratio drops by around 35 percentage points. This is far larger than the change in the debt limit associated with a relatively large economic shock, such as the 2008 recession.

From a policy perspective, this means that although debt limits may have indeed risen on account of the reduction in interest rates since the global financial crisis, it is hard to advocate that this has resulted in large increases in fiscal space. Moreover, these movements are very small relative to the impact of our uncertainty over the long-run level of the interest-growth differential. The largest fluctuations that we see year-to-year are comparable in magnitude only to changes in the long-run differential which are (in a statistical sense) tiny.

5 Conclusions

The policy question which motivates this paper is simple: do declines in nominal mean that governments can safely borrow more? This question is of critical importance in a world where fiscal policy seems to be expected to bear more of the burden of cyclical stabilization.

The short answer to that question is: yes, but probably only by a few percentage points. In arriving at this answer, we touched on several other interesting issues along the way.

This paper started out by noting that the key determinant of debt sustainability is not the nominal interest rate alone, but its level relative to the nominal growth rate. I then built a framework in which nests a large class of commonly-used models, which can then be applied to assess quantitatively how changes in interest-growth differentials might impact the maximum debt a government can borrow. I showed that in such models, finite maximum debt limits exist only if the interest-growth differential is positive on average in the long run.

Naturally, the next step was to measure the long run interest-growth differential in a sample of advanced economies. That analysis showed that while point estimates are indeed negative,

38. The point estimate of the long-run mean is -1.8%, and the standard error is around 0.9. The upper bound of the 95% confidence set (i.e. the 97.5th percentile of a symmetric interval) is therefore at $-1.8 + 1.96 \times 0.9 \simeq 0$. The percentile of the upper range of the x axis is given by $\phi((0.2 - -1.8)/0.9) \simeq 0.985$.

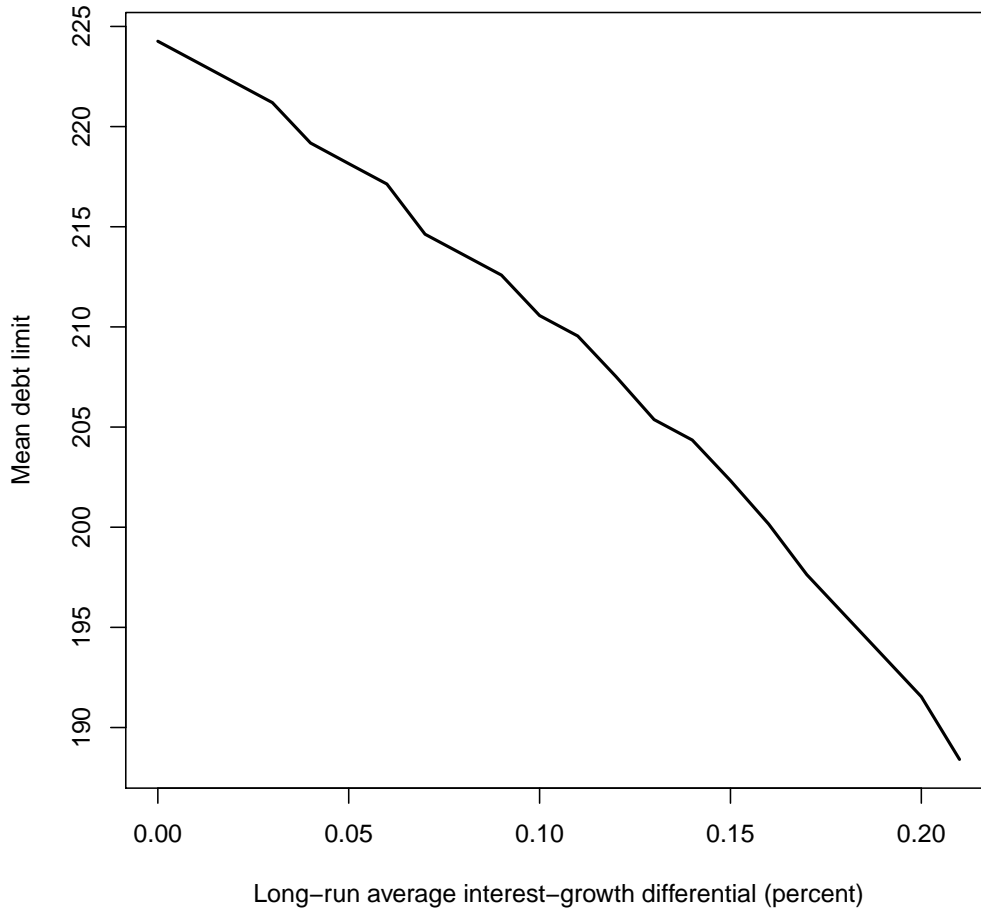


Figure 16: Sensitivity of debt limits to long-run interest-growth differentials

a variety of statistical techniques cannot reject the possibility that this differential is small and positive. I therefore concluded that in order to be conservative in our predictions of maximum sustainable debt levels, models of debt sustainability should feature interest-growth differentials which are small and positive.

Bearing this in mind, I calibrated the model to the historical data for the USA using a long-run interest-growth differential which was small and positive, yet statistically difficult to reject. This produced debt limits which vary little over time, but are enormously sensitive to the choice of the long-run differential, even when restricted to the set of those which are statistically plausible.

The latter part of this paper has focused on understanding debt limit in the presence of small, positive long run interest-growth differentials. While possible, it is much more likely that long run differentials are in fact negative. So it is worth also asking what this would mean for sustainable debt levels if true. Clearly, this would mean that solvency considerations alone cannot provide a

motive for finite debt limits. Yet default risk is non-zero, as even a cursory look at CDS markets will confirm. To reconcile this with negative interest-growth differentials, we must therefore look for other mechanisms to explain default risk. For example, models where even a solvent government can face a “run” on its debt, akin to those suffered by banks. The canonical example is Cole and Kehoe (2000), but more recent work by Aguiar et al. (2017) seek to examine the interaction of solvency with such rollover crises. These models do map into the basic framework in this paper, with self-fulfilling crises interpreted as a prohibition on universal repayment. The conclusions of this paper give added weight to this line of research.

References

- Aguiar, Mark, Satyajit Chatterjee, Harold Cole, and Zachary Stangebye. 2017. *Self-Fulfilling Debt Crises, Revisited: The Art of the Desperate Deal*. Technical report. National Bureau of Economic Research.
- Aguiar, Mark, and Gita Gopinath. 2006. “Defaultable debt, interest rates and the current account.” *Journal of international Economics* 69 (1): 64–83.
- Arellano, Cristina. 2008. “Default risk and income fluctuations in emerging economies.” *The American Economic Review*: 690–712.
- Auclert, Adrien, and Matthew Rognlie. 2016. “Unique equilibrium in the Eaton–Gersovitz model of sovereign debt.” *Journal of Monetary Economics* 84:134–146.
- Ball, Laurence, Douglas W Elmendorf, and N Gregory Mankiw. 1998. “The Deficit Gamble.” *Journal of Money, Credit and Banking*: 699–720.
- Benjamin, David, and Mark LJ Wright. 2009. “Recovery before redemption: A theory of delays in sovereign debt renegotiations.”
- Bi, Huixin, and Eric M Leeper. 2013. *Analyzing fiscal sustainability*. Technical report. Bank of Canada Working Paper.
- Blanchard, Olivier. 2019. “Public Debt and Low Interest Rates.” *AEA Presidential Address*.
- Blanchard, Olivier, and Philippe Weil. 2001. “Dynamic efficiency, the riskless rate, and debt Ponzi games under uncertainty.” *Advances in Macroeconomics* 1 (2).
- Bohn, Henning. 1998. “The behavior of US public debt and deficits.” *the Quarterly Journal of economics* 113 (3): 949–963.
- Borio, Claudio, Piti Disyatat, Mikael Juselius, and Phurichai Rungcharoenkitkul. 2017. “Why so low for so long? A long-term view of real interest rates.” *BIS Working Papers*, no. 685.
- Caballero, Ricardo J, Emmanuel Farhi, and Pierre-Olivier Gourinchas. 2008. “An equilibrium model of” global imbalances” and low interest rates.” *American economic review* 98 (1): 358–93.
- Carvalho, Carlos, Andrea Ferrero, and Fernanda Nechio. 2016. “Demographics and real interest rates: inspecting the mechanism.” *European Economic Review* 88:208–226.
- Chatterjee, Satyajit, and Burcu Eyigungor. 2012. “Maturity, indebtedness, and default risk.” *American Economic Review* 102 (6): 2674–99.
- Cole, Harold L, and Timothy J Kehoe. 2000. “Self-fulfilling debt crises.” *The Review of Economic Studies* 67 (1): 91–116.
- Collard, Fabrice, Michel Habib, and Jean-Charles Rochet. 2015. “Sovereign debt sustainability in advanced economies.” *Journal of the European economic association* 13 (3): 381–420.

- D’Erasmus, Pablo, Enrique G Mendoza, and Jing Zhang. 2016. “What is a sustainable public debt?” *Handbook of Macroeconomics* 2:2493–2597.
- Escolano, Julio. 2010. *A practical guide to public debt dynamics, fiscal sustainability, and cyclical adjustment of budgetary aggregates*. International Monetary Fund.
- Fournier, Jean-Marc, and Falilou Fall. 2017. “Limits to Government Debt Sustainability in OECD Countries.” *Economic Modelling* 66:30–41.
- Ghosh, Atish R, Jun I Kim, Enrique G Mendoza, Jonathan D Ostry, and Mahvash S Qureshi. 2013. “Fiscal fatigue, fiscal space and debt sustainability in advanced economies.” *The Economic Journal* 123 (566): F4–F30.
- Hatchondo, Juan Carlos, and Leonardo Martinez. 2009. “Long-duration bonds and sovereign defaults.” *Journal of international Economics* 79 (1): 117–125.
- Hatchondo, Juan Carlos, Leonardo Martinez, and Francisco Roch. 2017. “Robust Fiscal Rules.” *Working paper*.
- HC Deb. 4 September 2019. “Spending Round 2019.” *House of Commons Hansard* 664:177.
- IMF. 2017. *Fiscal Monitor: Achieving More with Less*. Washington, DC: International Monetary Fund, April.
- Jordà, Òscar, Moritz Schularick, and Alan M Taylor. 2017. “Macrofinancial history and the new business cycle facts.” *NBER Macroeconomics Annual* 31 (1): 213–263.
- Kim, Jun I, and Tamon Asonuma. Forthcoming. “Growth Uncertainty and Fiscal Space.” *IMF Working Paper*.
- Lisack, Noëmie, Rana Sajedi, and Gregory Thwaites. 2017. “Demographic Trends and the Real Interest Rate.” *Bank of England Working Paper*, no. 701.
- Lunsford, Kurt Graden, et al. 2017. “Productivity Growth and Real Interest Rates in the Long Run.” *Economic Commentary*, no. November.
- Mauro, Paolo, Rafael Romeu, Ariel Binder, and Asad Zaman. 2015. “A modern history of fiscal prudence and profligacy.” *Journal of Monetary Economics* 76:55–70.
- Mehrotra, Neil. 2017. “Debt sustainability in a low interest rate world.” *Hutchins Center Working Papers*.
- Müller, Ulrich K, and Mark W Watson. 2016. “Measuring uncertainty about long-run predictions.” *Review of Economic Studies* 83 (4): 1711–1740.
- Ostry, Jonathan, and Jun I Kim. Forthcoming. “Boosting Fiscal Space: The Roles of GDP-linked Debt and Maturity Structure.” *IMF Working Paper*.
- Parke, William R. 1999. “What is fractional integration?” *The review of economics and statistics* 81 (4): 632–638.
- Piketty, Thomas. 2014. *Capital in the twenty-first century*. Harvard University Press.

- Rachel, Lukasz, and Thomas Smith. 2015. "Secular drivers of the global real interest rate."
- Rouwenhorst, K.G. 1995. "Asset Pricing Implications of Equilibrium Business Cycle Models." Chap. 10 in *Frontiers of business cycle research*, edited by Thomas F Cooley, 294–330. Princeton, N.J.: Princeton University Press.
- Sajedi, Rana, and Gregory Thwaites. 2016. "Why Are Real Interest Rates So Low? The Role of the Relative Price of Investment Goods." *IMF Economic Review* 64 (4): 635–659.
- Tourre, Fabrice. 2016. "A Macro-Finance Approach to Sovereign Debt Spreads and Returns."
- Turner, David, and Francesca Spinelli. 2011. "Explaining the interest-rate-growth differential underlying government debt dynamics."
- Woo, Jaejoon, Anna Shabunina, and Julio Escolano. 2017. "The Puzzle of Persistently Negative Interest Rate-Growth Differentials: Financial Repression or Income Catch-Up?" *Fiscal Studies* 38 (2): 179–217.

A Proofs

It is useful to define set set of functions $\mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$, as:

Definition 7 $\mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$ is the set of functions which maps $\mathcal{X} \times \mathbb{R} \times \{0, 1\} \rightarrow \mathbb{R}$ and are continuous on their second input.

A.1 Investor rationality

The fixed-point problem in equation (4) can be expressed as a contraction.

Definition 8 Given d, H :

$$J_{|d,H} : \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\}) \rightarrow \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$$

$$q(x, b', \delta) \mapsto \frac{\lambda(1 - p(x, b', \delta)) + (1 - \lambda) \sum_{x'} \pi(x'|x, b', \delta)(1 - \lambda)q(x', H(x', b', d(x', b'), d(x', b')), d(x', b'))}{R(x)(1 - \psi p(x, b', \delta))}$$

Proposition 2 $J_{d,H}(q)$ is a contraction if $\min_x R(x) > 1 - \lambda$

Proof. By Blackwell's sufficient conditions. *Monotonicity:* If $\tilde{q} \geq q \forall (x, b', \delta)$ then $J_{|d,H}(\tilde{q}) \geq J_{|d,H}(q) \forall (x, b', \delta)$. *Discounting:*

$$J(q + a) = J(q) + \frac{(1 - \lambda)(1 - p(x, b', \delta))a}{R(x)(1 - \psi p(x, b', \delta))}$$

$$< J(q) + a \quad \Leftrightarrow \quad \min_x R(x) > 1 - \lambda$$

■

$\min_x R(x) > 1 - \lambda$ is to be understood in the remaining proofs.

Proposition 3 For any $d(x, b)$ and $H(x, b, \delta)$ weakly increasing in b for all (x, b, δ) :

$$\frac{\partial q}{\partial b'} < 0 \forall (x, b', \delta) \Rightarrow \frac{\partial J|(d, H)(q)}{\partial b'} < 0 \forall (x, b', \delta)$$

Proof.

$$J_{|d,H} = \frac{\lambda(1 - p(x, b', \delta))}{R(x)(1 - \psi p(x, b', \delta))} + \frac{(1 - \lambda) \sum_{x'} \pi(x'|x, b', \delta)(1 - \lambda)q(x', H(x', b', d(x', b'), d(x', b')), d(x', b'))}{R(x)(1 - \psi p(x, b', \delta))}$$

Then because $\psi < 1$ the first term is weakly increasing in b' . And because q is weakly decreasing in b and H is weakly increasing in b' then so too is $J(q)$. ■

Proof. (Of Proposition 1)

1. Follows from $J_{|d,H}$ a contraction
2. Follows because $J_{|d,H}$ a contraction on the set of weakly decreasing functions.

■

Note that given rules for default and debt issuance, we can compute prices relatively straightforwardly. As $J_{|d,H}$ is a contraction, iterative application converges at rate $\frac{1-\lambda}{\min_x R(x)}$.

A.2 Period feasibility

Again, period feasibility is a fixed point problem in that the debt issuance rule H must raise funds to cover gross financing needs given future prices that depend on H and d . Thus, we can define a mapping from the space of functions $\mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$ into itself, the fixed point of which satisfies period feasibility given $d(x, b)$

Definition 9 *Given a default rule $d(x, b)$:*

$$\chi_{|d} : \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\}) \rightarrow \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$$

$$H(x, b, \delta) \mapsto (1 - d(x, b))(1 - \lambda)b/G(x) + \frac{\max\left(\frac{\lambda b}{G(x)}(1 - d(x, b)) - s(x, b), 0\right)}{q_{|d,H}(x, H(x, b, \delta), d(x, b))}$$

The following proposition is then obvious:

Proposition 4 *$d(x, b)$ is period feasible if and only if there exists $H \in \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$ such that:*

1. H is a fixed point of $\chi_{|d}$: $H = \chi_{|d}(H)$
2. $H(x, d, \delta) < \infty \quad \forall \delta = d(x, b)$

Note also that $\chi_{|d}$ preserves monotonicity of H in the following sense:

Proposition 5 *If $s(x, 0) = 0$ and H is weakly increasing in b , then $\chi_{|d}(H)$ is also weakly increasing in b for any threshold rule $d(x, b)$.*

Proof. If H is weakly increasing in b , then by Proposition 3, $q_{|d,H}$ is weakly decreasing in b . We can now divide the proof into two cases. First, if $b < b_s(x)$ then because $s(x, 0) = 0$ then the mean value theorem says that $\lambda b/G(x) < s(x, b)$ and so $H = (1 - \lambda)b/G(x)$, which is strictly increasing. Second, if $b \geq b_s(x)$ then $s'(x, b) < \lambda/G(x)$ which also implies that H is at least weakly increasing. ■

In contrast to $J_{|d,H}$ (the mapping for debt prices), $\chi_{|d}$ is not a contraction. Until further notice, assume that $s(x, 0) = 0$. For a threshold default rule given by \bar{b} , I denote $J_{|d_0(\bar{b}),H}$, $\chi_{|d_0(\bar{b})}$ by $J_{|\bar{b},H}$, $\chi_{|\bar{b}}$.

If d is a threshold default rule, then the set on which $\chi_{|d}(H)$ is finite *defines* a new threshold rule. In other words, $\chi_{|d}$ preserves threshold default rules.

Proposition 6 *Let $s(x, 0) = 0$, then for any $d(x, b)$ and H increasing in b there exists $\beta_{|d,H}(x) \geq b_s(x) > 0$ such that $\forall x$:*

$$\chi_{|d}(H)(x, b, 0) < \infty \quad \Leftrightarrow \quad b \leq \beta_{|d,H}(x)$$

Furthermore, if $d(x, b)$ is a finite threshold rule then $\beta_{|d,H}(x) < \infty$

Proof. Clearly $\chi_{|d}(H)(x, b, 0) = 0 \forall b < b_s(x)$, $\forall x$. And as $\chi_{|d}$ is weakly increasing in b , then if there exists some b for which $\chi_{|d}(x, b, 0)$ is infinite then so too must be $\chi_{|d}$ for all larger b . That is, there must be a threshold for b as described. And if $d(x, b)$ is finite then for sufficiently large b , $J(q)(x, b', \delta) = 0 \forall (x, b', \delta)$, i.e. we can guarantee that $\chi_{|d}$ is eventually infinite. ■

We can combine $\beta_{\bar{b},H}$ and $\chi_{|\bar{b}}$ into one mapping:

Definition 10 $\Theta(\bar{b}, H)$ *defines the default thresholds and debt issuance rules in period t when \bar{b} and H are the default and debt issuance rules in period $t + 1$.*

$$\Theta : \mathbb{R}^m \times \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\}) \rightarrow \mathbb{R}^m \times \mathcal{C}(\mathcal{X}, \mathbb{R}, \{0, 1\})$$

$$\begin{pmatrix} \bar{b} \\ H \end{pmatrix} \mapsto \begin{pmatrix} \min(\bar{b}(x), \beta_{|\bar{b},H}(x)) \\ \chi_{|\bar{b}}(H) \end{pmatrix}$$

And we denote the elements of the output by $\Theta_1(\bar{b}, H)$, $\Theta_2(\bar{b}, H)$ respectively.

Note that if H is finite whenever $\delta = d_0(\bar{b})$, then so too is $\chi_{|\bar{b}}(H)$ whenever $\delta = d_0(\Theta_1(\bar{b}, H))$. As a result, we have that:

Proposition 7 (Period feasibility) *A threshold default rule $d_0(\bar{b})$ is period-feasible with corresponding debt issuance function H if and only if (\bar{b}, H) is a fixed point of Θ .*

Proof. If (\bar{b}, H) is a fixed point of Θ then:

$$\bar{b}(x) = \min(\bar{b}(x), \beta_{|\bar{b}, H}(x)) \quad H = \chi_{|\bar{b}}(H)$$

Therefore equation (5) is satisfied whenever $\delta = d_0(\bar{b})$

If \bar{b} is period feasible, then $H = \chi_{|\bar{b}}(H)$. Moreover, H is finite on $b \leq \beta_{|\bar{b}, H}(x)$ and $\beta_{|\bar{b}, H}(x) \leq \bar{b}(x)$. ■

Moreover, repeated application of Θ guarantees the existence of a period feasible default rule and corresponding issuance rule (and gives a convergent computational method for finding it).

Proposition 8 (Existence of period-feasible threshold default rules) *For any $(\bar{b}_0, H_0) \in C(\mathcal{X}, \mathbb{R}, \{0, 1\})$ such that $\bar{b}(x) > b_s(x) \forall x$ and H is weakly increasing in b , the sequence $\{\bar{b}_n, H_n\}$ defined by*

$$\bar{b}_{n+1}, H_{n+1} = \Theta(\bar{b}_n, H_n)$$

converges to $(\bar{b}_\infty, H_\infty)$, a fixed point of Θ with H_∞ weakly increasing in b .

Proof. By definition, we have that $\bar{b}_{n+1} \leq \bar{b}_n$. And by Proposition 6, $\bar{b}_n > b_s(x) \forall x$. Thus, the sequence \bar{b}_n is weakly decreasing and bounded below by zero. It therefore converges to a fixed point, $\bar{b}_\infty > 0$. Now assume for the purpose of a contradiction that H_n does not converge, even when $\bar{b} = \bar{b}_\infty$. Then $H_{n+1}(x, b', \delta) \leq H_n(x, b', \delta) \forall (x, b', \delta)$. If not, $\beta_{|\bar{b}_\infty, H_n}(x) < \bar{b}_\infty(x)$ for some x implying that \bar{b}_∞ is not a fixed point. And so for \bar{b}_n sufficiently close to \bar{b}_∞ , H_n is a weakly decreasing and bounded below by risk-free continuation debt issuance. Thus, H_n also converges to a fixed point. Weak monotonicity of H_∞ follows from the monotonicity of H_0 and Proposition 5. ■

We can now put these parts together to prove Theorem 1.

Proof. (Of Theorem 1) Under the assumptions of the theorem, repeated application of Θ generates a sequence $(\bar{b}_0, \bar{b}_1, \bar{b}_2, \dots)$ convergent to a fixed point of Θ . Thus a fixed point exists with H weakly increasing in b . So by Proposition 7 this defines a period-feasible default rule. ■

Proof. (Of Theorem 2) We proceed by showing that for \bar{b}_0 sufficiently large there exists $\hat{n} \leq N$ such that $\bar{b}_{\hat{n}}(x) < \bar{b}(x) \forall x$. This proves the claim because there can be no fixed point in a region where repeated application of Θ is strictly monotone. Throughout the proof we focus on the bounding case, where the price of debt is given by $q^{rf}(x)$. This is because for any d, H , $q_{d, H}(x, b', \delta) \geq q^{rf}(x)$ and thus $\beta_{|d, H}(x)$ is bounded above by the equivalent when prices are given

by $q^{rf}(x)$ (i.e. the maximum period-feasible debt is always weakly less than that under risk-free prices).

Let the default rule be initially given by a threshold rule $d(\bar{b}_0)$. Then:

$$\begin{aligned} q_{|d,H}(x, b') &\leq q^{rf}(x) \sum_{x': b' \leq b(x')} \pi(x'|x) \\ &\leq q^{rf}(x) 1_{b' \leq \bar{B}_0} \end{aligned}$$

Where $\bar{B}_0 = \max_x \bar{b}_0(x)$. Then $\chi_{|d,H}(c, b', \delta) = b'$ is infinite if:

$$(1 - \lambda) \frac{b}{G(x)} + \frac{1}{q^{rf}(x)} \left[\frac{\lambda b}{G(x)} - s(x, b) \right] > \bar{B}_0$$

$$\text{That is: } b > \frac{G(x)q^{rf}(x)\bar{B}_0}{\lambda + (1 - \lambda)q^{rf}(x)} + \frac{s(x, b)}{\lambda + (1 - \lambda)q^{rf}(x)}$$

Then for sufficiently large $\bar{b}_0(x)$ we can guarantee that this b is strictly less than $\bar{b}_0(x)$ if:

$$\frac{G(x)q^{rf}(x)}{\lambda + (1 - \lambda)q^{rf}(x)} < 1 \quad (9)$$

Which by equation (6) must be true for some x . We denote the set of such x by X_1 . Then, for $x \in X_1$, $\beta_{d,H}(x) = \tilde{b}_1(x) < \bar{b}_0(x)$. The next iteration of the threshold function is then:

$$\beta(x) = \bar{b}_1(x) \leq \begin{cases} \bar{b}_0(x) & R(x) \leq G(x) \\ \tilde{b}_1(x) & R(x) > G(x) \end{cases}$$

For $n = 2$, if $b' \in \left(\max_{x \in X_1} \tilde{b}_1(x), \bar{b}_0(x) \right)$ then the expected default probability is the probability that $x \in X_1$ in the next period, i.e. that:

$$p(x) = \sum_{x' \in X_1} \pi(x'|x, 0)$$

And thus for b' in this region:

$$q_{|d,H}(x, b') \leq q^{rf}(x) \sum_{x' \in X_1} \pi(x'|x, b', 0)$$

Applying the same reasoning, we can derive a new condition which guarantees that $b_2(x) < b_1(x)$.

That is, the new version of Equation (10):

$$\frac{p(x)G(x)q^{rf}(x)}{\lambda + (1 - \lambda)q^{rf}(x)} < 1 \quad (10)$$

But by equation (6) we know that there must exist some x such that 1) $x \notin X_1$ and 2) $p(x)G(x)q^{rf}(x) < \lambda + (1 - \lambda)q^{rf}(x)$. In other words, we can guarantee that we have strict monotonicity for at least one value of x where the interest-growth differential is negative.

To extend this to all values of x , we simply reapply the same logic: equation (6) guarantees that under the next iteration of Θ there must be at least one more level of x for which Θ is strictly monotone. Thus, there must exist some $\hat{n} \leq N$ such that for sufficiently large $\bar{b}_0, \bar{b}_{\hat{n}} < \bar{b}_0(x) \forall x$.

■

A.3 Long-run stability

Proof. (Of Theorem 3) We use the fact that a Markov chain with at least one recurrent state has a long-run invariant distribution. The process for x_t is irreducible, and under a finite threshold default rule, b_t is bounded for $t \geq 1$. Therefore, there must be some recurrent state (x, b) , and so there exists a well-defined long-run distribution for $\{(x_t, b_t)\}_{t=0}^{\infty}$. The long-run distribution for $\{b_t\}_{t=0}^{\infty}$ is then just the corresponding marginal distribution. ■

Proof. (Of Theorem 4) Case II: Assume $\mathbb{E} \log (R(x)/G(x)) < 0$. Then, under the universal-repayment rule, $q(x, b') = q^{rf}(x) \forall (x, b')$. As $s(x, b)$ is bounded, then for any arbitrarily small $\epsilon > 0$, there exists \hat{b} such that $s(x, b)/b < \epsilon \forall b \geq \hat{b}$. Then for $b > \hat{b}$, evolution of (log) debt is given by:

$$\begin{aligned} \log b_{t+1} &\geq \log \left[\frac{b_t}{G(x_t)} \left((1 - \lambda) + \frac{\lambda}{q(x_t)} - \epsilon \right) \right] \\ &\geq \log b_t + \log \left((1 - \lambda)q^{rf}(x_t) + \lambda \right) - \log (q^{rf}(x_t)) - \log (G(x_t)) \\ \Rightarrow \quad b_{t+n} &= \log b_t + n \times \left[\frac{1}{n} \sum_{k=0}^{n-1} \log \left((1 - \lambda)q^{rf}(x_{t+k}) + \lambda \right) - \frac{1}{n} \sum_{k=0}^{n-1} \log (q^{rf}(x_{t+k})) - \frac{1}{n} \sum_{k=0}^{n-1} \log (G(x_{t+k})) \right] \\ &= \log b_t + nJ_n \end{aligned} \quad (11)$$

But $q^{rf}(x) = \frac{1}{R(x)} [\lambda + (1 - \lambda) \sum_{x' \in \mathcal{X}} \pi(x', |x, 0, \delta) q^{rf}(x')]$. Substituting into the equation for

$\log b_{t+1}$ above, and applying the pointwise ergodic theorem for a Markov chain:

$$J_n \xrightarrow{\mathbb{P}} \mathbb{E} \log R(x) - \mathbb{E} \log G(x) = \mathbb{E} \log \left(\frac{R(x)}{G(x)} \right)$$

And so the drift term in equation (11) is positive for large enough b (i.e. small enough ϵ). As is well-known, this process has no well-defined long-run distribution; any probability distribution will “drift off” towards infinity.

Case II: Assume $\mathbb{E} \log (R(x)/G(x)) < 0$. Then by a similar argument, $\log b_t$ can be bounded above (this time using the fact that the surplus process is bounded below) by a random walk with negative drift. Thus b_t is bounded above with probability one. Because $b_t > 0$ then there must exist a recurrent state (x, b) , and thus an invariant long-run distribution for b_t . ■

When we have departures from risk-neutrality, the limit of J_n becomes:

$$J_n \xrightarrow{\mathbb{P}} \mathbb{E} \log (\lambda + (1 - \lambda)q^{rf}(x)) - \mathbb{E} \log \left(\lambda + (1 - \lambda) \sum_{x'} \tilde{\pi}(x'|x)q^{rf}(x') \right) + \mathbb{E} \log \left(\frac{R(x)}{G(x)} \right)$$

Where expectation is over the stationary distribution of x . Taking a first-order approximation:

$$\begin{aligned} J_n &\xrightarrow{\mathbb{P}} \mathbb{E} (\lambda + (1 - \lambda)q^{rf}(x) - 1) - \mathbb{E} \left(\lambda + (1 - \lambda) \sum_{x'} \tilde{\pi}(x'|x)q^{rf}(x') - 1 \right) + \mathbb{E} \log \left(\frac{R(x)}{G(x)} \right) \\ &= (1 - \lambda)\mathbb{E} \left(q^{rf}(x) - \sum_{x'} \tilde{\pi}(x'|x)q^{rf}(x) \right) + \rho \\ &= (1 - \lambda)\mathbf{v}'(I - \tilde{\Pi})\mathbf{q}^{rf} + \rho \\ &= (1 - \lambda)\mathbf{v}'(\Pi - \tilde{\Pi})\mathbf{q}^{rf} + \rho \end{aligned}$$

Where the last line follows from $\mathbf{v}' = \mathbf{v}'\Pi$. This is equation (7).

B Examples

The framework described above nests several well-known models of debt sustainability, including both optimal default models and models of affordability. Here, I show how the canonical exponents of both genres, Arellano (2008) and Ghosh et al. (2013) can be recast in these terms. Moreover, that the solutions to these models all satisfy period feasibility and long-run stability. I also show that this framework can also nest a simple general equilibrium model with a tax Laffer curve.

A final example I discuss is that of Blanchard (2019). This cannot be mapped into this setting;

examining why not is instructive.

Arellano (2008)

This is a model of optimal sovereign default. The government faces a state-dependent fluctuating income $y(x)$, and can borrow from international investors to smooth these fluctuations, paying for household consumption c . Debt is one-period, so $\lambda = 1$. Investors are risk-neutral and have a constant outside option with guaranteed return R . Upon default, the country is in autarky – the government is excluded from borrowing, and is re-admitted to international capital markets in subsequent periods with constant per-period probability γ . For simplicity I assume that y can take two levels, high and low, denoted y_h & y_l with i.i.d. probability p that income is high. Growth, therefore, can take three levels, $G_h = y_h/y_l$, $G_m = 1$, and $G_l = y_l/y_h$.

The stochastic setting in this model can be expressed in a eight-state Markov chain, with $x \in \{x_{hh}, x_{hl}, x_{lh}, x_{ll}, x_{hh}^a, x_{hl}^a, x_{lh}^a, x_{ll}^a\}$. These four states are the four combinations of income levels in and out of autarky. The growth function is then:

$$G(x) = \begin{cases} 1 & x \in \{x_{ll}, x_{hh}, x_{ll}^a, x_{hh}^a\} \\ y_h/y_l & x \in \{x^{lh}, x_{lh}^a\} \\ y_l/y_h & x \in \{x^{hl}, x_{hl}^a\} \end{cases}$$

The transition probabilities are then given by:

$$\Pi_{d=0} = \begin{bmatrix} p & 1-p & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & p & 1-p & 0 & 0 & 0 & 0 \\ p & 1-p & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & p & 1-p & 0 & 0 & 0 & 0 \\ (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) & 0 & 0 \\ 0 & 0 & (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) \\ (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) & 0 & 0 \\ 0 & 0 & (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) \end{bmatrix}$$

$$\Pi_{d=1} = \begin{bmatrix} 0 & 0 & 0 & 0 & p & 1-p & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & p & 1-p \\ 0 & 0 & 0 & 0 & p & 1-p & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & p & 1-p \\ (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) & 0 & 0 \\ 0 & 0 & (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) \\ (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) & 0 & 0 \\ 0 & 0 & (1-\gamma)p & (1-\gamma)(1-p) & 0 & 0 & \gamma p & \gamma(1-p) \end{bmatrix}$$

These matrices define the transition function $\pi(x'|x, \delta)$.

If r is the risk-free return of investors' outside option, interest and risk-aversion functions are:

$$R(x) = 1 + r \quad \forall x \quad \phi(x) = \begin{cases} 1 & x \in \{x_l, x_h\} \\ \infty & x \in \{x_l^a, x_h^a\} \end{cases}$$

Then if $c(x, b)$ represents the government's optimal expenditure function (which in this model equals consumption), the surplus ratio is then:

$$s(x, b) = 1 - \frac{c(x, b)}{y(x)}$$

Which satisfies Assumption 2 if $c(x, b)$ is decreasing and convex in b , which holds if households' *marginal* utility is convex. Thus, the framework can represent the solution to the Arellano model. And so all the results described here, including those about the long-run level of the interest-growth differential, still hold in this setting.

Because a budget constraint in the form of equation (5) holds, period feasibility is satisfied.

And because the optimal default rule is a finite threshold rule, then long run sability also holds.

Ghosh et al. (2013)

In this setting, there is no uncertainty over the risk-free rate, growth, or investor risk preferences, but there is uncertainty over surpluses. Debt is again short term, so $\lambda = 1$. The government's surplus is a given function of debt, $f(b)$, subject to i.i.d. shocks each period, which is calibrated to a triangle distribution. This can be represented arbitrarily closes using a suitable discretization $\{x_1, \dots, x_N\}$, with probability density $p(x_n)$. Then the surplus function is given by

$$s(x, b) = f(b) + x$$

Ghosh et al. (2013) focus on surplus functions capturing “fiscal fatigue”. That is, the idea that the government cannot continue to raise ever higher surpluses as debt levels increase. The limit in condition 2 of Assumption 2 reflects this property.

However, the surplus function that Ghosh et al. use violates Assumption 2 in way that is material to implications for maximum sustainable debt levels. They assume a cubic function for $f(b)$, with the result that the the deficit is not bounded and instead becomes arbitrarily large as b grows. This violation of Assumption 2 means that Theorem 4 does not hold. For a high enough debt level, the deficit will always exceed interest payments, even if the interest-growth differential is negative. Other than this difference, the Ghosh et al. model maps directly into the framework of Section 2.

A simple optimal tax model

A drawback of both the Arellano and Ghosh et al. models is that they feature no meaningful trade-off between tax revenues and welfare. In Arellano, the government can raise lump-sum taxes up to 100% of output. In Ghosh et al. the surplus function is taken as empirical rather than micro-founded. Here I present a simple optimal taxation model with a distorting tax and show that it produces a surplus function which satisfies Assumption 2. And so there is both an increasing marginal welfare cost of taxes and a Laffer curve which limits total tax revenues.

As the focus here is on establishing concavity of the surplus function given the state x . I therefore omit an explicit description of the stochastic process for x but instead assume that it is chosen so that Assumption 1 holds.

Households have period preferences over consumption c , labor n and government spending g .

$$\mathcal{U}(c, n, g) = \frac{c^{1-\sigma}}{1-\sigma} - \chi \frac{n^{1+\eta}}{1+\eta} + \theta \log g$$

Households have no access to savings technology, and so maximize this function subject to

$$c = w(x)n(1 - \tau(x))$$

where x is an aggregate state and $w(x)$ and $\tau(x)$ the pre-tax wage and tax rate in state x . Wages are the marginal productivity of the linear aggregate technology function, $A(x)$. Thus, $w(x) = A(x)$ and the household's optimal consumption and labor supply choices are given by:

$$c(x) = A(x)\kappa(x)(1 - \tau(x))^{1+\alpha} \quad n(x) = \kappa(x)(1 - \tau(x))^\alpha$$

Where:

$$\kappa(x) = \left(\frac{A^{1-\sigma}}{\chi} \right)^{\frac{1}{\eta}} \quad \alpha = \frac{1-\sigma}{\eta+\sigma}$$

Then tax revenue is given by:

$$Rev(\tau, x) = A(x)\kappa(x)\tau(x)(1 - \tau(x))^\alpha$$

Which is maximized at $\tau = 1/(1 + \alpha)$. Then indirect welfare of a tax rate τ is given by:

$$\mathcal{V}(\tau, x) = (A(x)\kappa(x))^{1-\sigma} \frac{(1-\tau)^{(1+\alpha)(1-\sigma)}}{1-\sigma} - \chi \kappa(x)^{1+\eta} \frac{(1-\tau)^{\alpha(1+\eta)}}{1+\eta}$$

Note that $(1 + \alpha)(1 - \sigma) = \alpha(1 + \eta)$ and so $\mathcal{V}(\tau, x)$ is concave and decreasing in τ if $\sigma > 1$. I assume that the government sets policy each period to maximize household welfare less a penalty on the deficit. This penalty increases with the outstanding debt, in order to produce stable debt levels. That is, the government solves:

$$\max_{\tau, g} \mathcal{V}(\tau, x) + \theta \log g - \gamma (Rev(x, \tau) - g) b$$

The first order conditions are:

$$g = \frac{\gamma}{\theta b} \quad \mathcal{V}'(\tau, x) = \gamma b \frac{dRev(\tau, x)}{d\tau}$$

So g is clearly decreasing and convex in b . And because $\mathcal{V}'(\tau, x)$ is decreasing and concave if $\sigma > 1$ (this follows directly from twice differentiating the definition), and $Rev(\tau, x)$ is increasing and concave then the optimal revenue is increasing and concave in b . Thus, the optimal surplus is increasing and concave in b , thus satisfying Assumptions 2.

Blanchard (2019)

This setting cannot be mapped into the framework described in section 2. This is an overlapping generations model, so the long-run demand for assets is not infinitely elastic. As a result, government policy can affect the risk-free rate over the long run. In particular, if the government issues enough debt this crowds out private capital, pushing up the risk-free rate. Assumption 1 therefore fails, and so Theorem 4 fails.

C Further empirical work

C.1 Measures of the risk premium

Table 8 shows the variances and covariances of consumption growth and inflation since 1950 for seven advanced economies. These are very small, and implied risk premia are similarly tiny unless investors have very high risk aversion.

C.2 Results from long-lagged VAR estimation

Table 9 reports the results from the long-lagged VAR estimation. On their own, the coefficients are not terribly interesting. However, the comparison of the long-run estimates to the Table 3 is illuminating. While the point estimates of the long-run means of growth and interest rates individually are a little different to those in the one-lag results, the differences – both point estimates and standard errors – are very similar. In other words, confidence intervals for the long-run interest-growth differential seem to be robust to model specification.

To illustrate how the restricted estimation process works, Table 10 shows the VAR estimates restricted such that the mean difference is at the 5% critical value for each country. Comparing to Table 9 we see that the dynamics coefficients are little changed; only the intercept term has changed. In other words, the dynamics of interest and growth rates seem, from a statistical perspective, to be pinned down fairly well. But the relative levels are not, hence the uncertainty over the interest-growth differentials.

Country	Period	$Var_t g_t^c$	$Var_t \pi_t$	$Cov_t(g_t^c, \pi_t)$	Risk prem. ($\sigma = 1$)	Risk prem. ($\sigma = 10$)
Canada	1950-69	0.04	0.05	-0.02	0.02	1.62
	1970-89	0.04	0.08	-0.03	0.03	1.93
	1990-	0.02	0.01	-0.00	0.01	0.98
France	1950-69	0.03	0.19	-0.01	0.10	1.34
	1970-89	0.02	0.12	-0.01	0.06	0.91
	1990-	0.01	0.01	-0.00	0.01	0.67
Germany	1950-69	0.05	0.03	-0.01	0.03	2.67
	1970-89	0.03	0.04	-0.01	0.03	1.43
	1990-	0.02	0.01	0.00	0.02	0.81
Italy	1950-69	0.04	0.08	-0.01	0.05	1.75
	1970-89	0.03	0.25	-0.02	0.12	1.48
	1990-	0.05	0.02	0.00	0.04	2.43
Japan	1950-69	0.06	0.14	0.03	0.13	3.48
	1970-89	0.05	0.27	-0.05	0.12	2.38
	1990-	0.01	0.01	0.00	0.01	0.59
UK	1950-69	0.03	0.05	-0.03	0.01	1.03
	1970-89	0.07	0.24	-0.10	0.06	2.58
	1990-	0.04	0.02	-0.02	0.01	1.86
USA	1950-69	0.04	0.04	-0.01	0.03	1.74
	1970-89	0.04	0.09	-0.05	0.02	1.47
	1990-	0.02	0.01	0.00	0.02	1.03

Table 8: G7 Consumption growth and inflation variances (logs, multiplied by 100).

C.3 Alternative interest rate measures

As a robustness check on the empirical results, here I test for the sign of the long-run interest-growth differential using two alternative measures of the interest rate: the interest rate on long-term government debt, and the effective interest rate on the stock of government debt. In both cases, the data are drawn from Mauro et al. (2015).

Table 11 introduces the new interest rate series, presenting their spreads relative to the risk-free rate. The long-term rate is, on average, higher than the risk-free rate as term- and risk-premia are both likely to be positive. Within-country spread variation correlates negatively with the current risk-free rate, reflecting mean-regression of the risk-free rate. For example, spreads over risk-free rates are high after 2000, a period when risk-free rates were low. In contrast, effective interest rates are not uniformly higher than the risk-free rate on average. Interpreting within-country movements is also hard, as the effective rate is the interest rate paid on the outstanding stock of debt, which is likely to be a time-varying combination of different maturities. However, as the backward-looking nature of this measure means that the spread is typically lower during

	USA		United Kingdom		France		Germany	
	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>								
Constant	4.4 (1.388)	0.31 (0.304)	1.53 (0.749)	0.05 (0.185)	1.46 (1.097)	0.24 (0.217)	2.94 (1.425)	0.71 (0.373)
Growth (-1)	0.49 (0.1)	0.01 (0.022)	0.64 (0.097)	0.06 (0.025)	0.43 (0.107)	0 (0.023)	0.76 (0.097)	0.06 (0.026)
Int. rate (-1)	-0.91 (0.452)	0.96 (0.101)	-0.4 (0.381)	0.96 (0.097)	0.51 (0.515)	1.08 (0.108)	-0.5 (0.372)	0.79 (0.098)
Growth (-2)	0.01 (0.108)	0.02 (0.024)	0.17 (0.114)	-0.01 (0.029)	0.18 (0.116)	0.02 (0.024)	-0.26 (0.121)	-0.05 (0.032)
Int. rate (-2)	0.91 (0.622)	-0.43 (0.139)	0.29 (0.537)	-0.33 (0.137)	-0.18 (0.728)	-0.47 (0.154)	0.35 (0.48)	-0.14 (0.126)
Growth (-3)	-0.08 (0.106)	0.02 (0.024)	-0.14 (0.113)	-0.01 (0.029)	0.12 (0.115)	0.02 (0.024)	0.28 (0.122)	0.02 (0.032)
Int. rate (-3)	-0.53 (0.622)	0.37 (0.139)	0.35 (0.537)	0.16 (0.137)	-0.07 (0.726)	0.32 (0.153)	-0.25 (0.475)	0.11 (0.125)
Growth (-4)	-0.04 (0.094)	-0.01 (0.021)	-0.07 (0.097)	0.02 (0.025)	-0.17 (0.106)	-0.01 (0.022)	-0.14 (0.098)	0.03 (0.026)
Int. rate (-4)	0.35 (0.449)	-0.05 (0.1)	-0.08 (0.366)	0.09 (0.093)	0.01 (0.502)	-0.03 (0.106)	0.2 (0.371)	-0.01 (0.098)
<i>Innov. covar.</i>								
Growth	43.56 (5.788)		20.94 (2.774)		32.22 (4.769)		25.63 (3.565)	
Int. rate	3.69 (0.972)	2.17 (0.287)	1.55 (0.52)	1.36 (0.18)	2.04 (0.739)	1.44 (0.211)	0.81 (0.666)	1.77 (0.246)
<i>Long-run</i>								
Mean	5.96 (0.853)	3.84 (0.942)	5.49 (1.598)	3.87 (1.654)	5.67 (1.943)	4.02 (1.542)	5.96 (1.231)	4.04 (0.543)
Difference	-2.12 (1.013)		-1.63 (0.809)		-1.65 (1.119)		-1.92 (1.137)	
Log likelihood	660.3		586.5		519.2		574.2	
Observations	134		134		114		128	

Table 9: Multi-lag VAR for sample of countries. Annual data 1880-2013. Robust likelihood-based standard errors in parentheses.

	USA		United Kingdom		France		Germany	
	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>								
Constant	4.37 (1.394)	0.32 (0.271)	1.3 (0.578)	-0.09 (0.188)	0.17 (11.488)	-0.06 (1.119)	1.65 (1.255)	0.85 (0.301)
Growth (-1)	0.49 (0.134)	0.01 (0.019)	0.65 (0.084)	0.06 (0.024)	0.44 (0.32)	0.01 (0.124)	0.79 (0.174)	0.06 (0.032)
Int. rate (-1)	-0.89 (0.441)	0.99 (0.127)	-0.35 (0.444)	0.98 (0.099)	0.57 (1.784)	1.1 (0.369)	-0.54 (0.357)	0.8 (0.115)
Growth (-2)	0.01 (0.164)	0.02 (0.019)	0.17 (0.161)	-0.01 (0.034)	0.18 (0.161)	0.02 (0.06)	-0.25 (0.173)	-0.05 (0.028)
Int. rate (-2)	0.91 (0.684)	-0.43 (0.126)	0.3 (0.521)	-0.32 (0.18)	-0.15 (0.538)	-0.46 (0.243)	0.38 (0.492)	-0.14 (0.159)
Growth (-3)	-0.08 (0.142)	0.02 (0.019)	-0.14 (0.129)	-0.01 (0.03)	0.12 (0.192)	0.02 (0.072)	0.3 (0.166)	0.01 (0.032)
Int. rate (-3)	-0.52 (0.657)	0.39 (0.124)	0.37 (0.496)	0.17 (0.198)	-0.06 (0.737)	0.32 (0.266)	-0.26 (0.45)	0.11 (0.188)
Growth (-4)	-0.05 (0.11)	-0.02 (0.016)	-0.08 (0.163)	0.02 (0.028)	-0.18 (0.374)	-0.02 (0.109)	-0.12 (0.097)	0.03 (0.022)
Int. rate (-4)	0.35 (0.368)	-0.04 (0.115)	-0.04 (0.26)	0.11 (0.102)	0.14 (1.044)	-0.01 (0.197)	0.27 (0.299)	-0.02 (0.104)
<i>Innov. covar.</i>								
Growth	43.83 (7.945)		21.14 (6.755)		32.82 (15.554)		26.47 (7.024)	
Int. rate	3.75 (1.065)	2.24 (0.346)	1.61 (0.681)	1.39 (0.256)	2.17 (1.242)	1.47 (0.37)	0.74 (0.914)	1.8 (0.37)
<i>Long-run</i>								
Mean	5.66 (0.88)	5.6 (1.806)	10.24 (2.538)	10.21 (3.836)	-4.29 (43.88)	-4.03 (4.384)	3.6 (1.827)	3.97 (0.608)
Difference	-0.06 (2.031)		-0.03 (3.623)		0.26 (47.241)		0.37 (1.675)	
Log likelihood	662.2		588.9		521.7		576.2	
Observations	134		134		114		128	

Table 10: Multi-lag restricted VAR for sample of countries. Lon-run means are restricted to be at the 5% critical value under the unconditional LR test. Annual data 1880-2013. Robust likelihood-based standard errors in parentheses

	<i>Long-term interest rates</i>				<i>Effective interest rates</i>			
	USA	UK	France	Germany	USA	UK	France	Germany
1880-2013	0.55	0.90	1.16	1.28	-0.88	0.35	-0.46	0.12
1880-1900	-0.12	0.20	0.91	0.64	-0.36	1.30	0.79	1.10
1900-1920	-0.01	-0.15	0.53	0.53	-1.90	0.19	0.13	0.46
1920-1940	0.30	1.54	1.20	1.27	0.55	2.00	-0.48	-4.50
1940-1960	0.83	1.75	2.61	1.90	0.19	0.26	-0.70	-2.10
1960-1980	0.30	1.78	1.08	2.33	-3.10	-1.80	-2.40	0.87
1980-2000	1.16	-0.05	0.79	1.18	-2.20	-0.40	-1.30	0.85
2000-2013	1.91	1.43	1.73	1.37	1.90	1.20	1.70	1.70

Table 11: Average spread over risk-free rate

periods where the risk-free rate has been surprisingly high (such as 1980-2000).

Tables 12 and 13 repeat the simple one-lag VAR exercise, using the long-term and effective rates on government debt respectively. With long-term rates, interest-growth differentials are typically a little larger, notably so in the case of France, suggesting that default and term premia have an important impact on the results of the estimation. In the case of effective rates, surprisingly high inflation in the 1970s and 1980s results in somewhat lower estimated for most countries, most obviously the USA.

Finally, Figure 17 shows the summary results of the hypothesis tests using the long-lagged VAR- and spectral-based estimates for the long-run interest-growth differential using these measures. Estimates using long-term rates are, unsurprisingly, higher than the risk-free rate (particularly for France). But the overall story is qualitatively similar: point estimates are universally negative, and upper limits of the confidence sets are usually small and positive.

C.4 Application of the spectral method to quarterly data

Table 14 reports the spectral results for the quarterly dataset. In contrast to the annual data, these estimates diverge wildly from those implied by VARs. This is due to the failure of stationarity in the sample period. In most of the sample countries, interest-growth differentials have been steadily declining since the early 1980s (see Figure 5). In the spectral analysis, this long, slow movement is interpreted as sufficiently persistent as to affect the long-run properties of the time series. This is reflected in the estimated degree of fractional integration, reported in the last line of Table 14. The USA and UK have an estimated degree of integration of unity (at least, the estimation process does not consider higher-order integration). Indeed, only Germany has a stationary representation during this period.³⁹ As a result, the confidence sets grow without

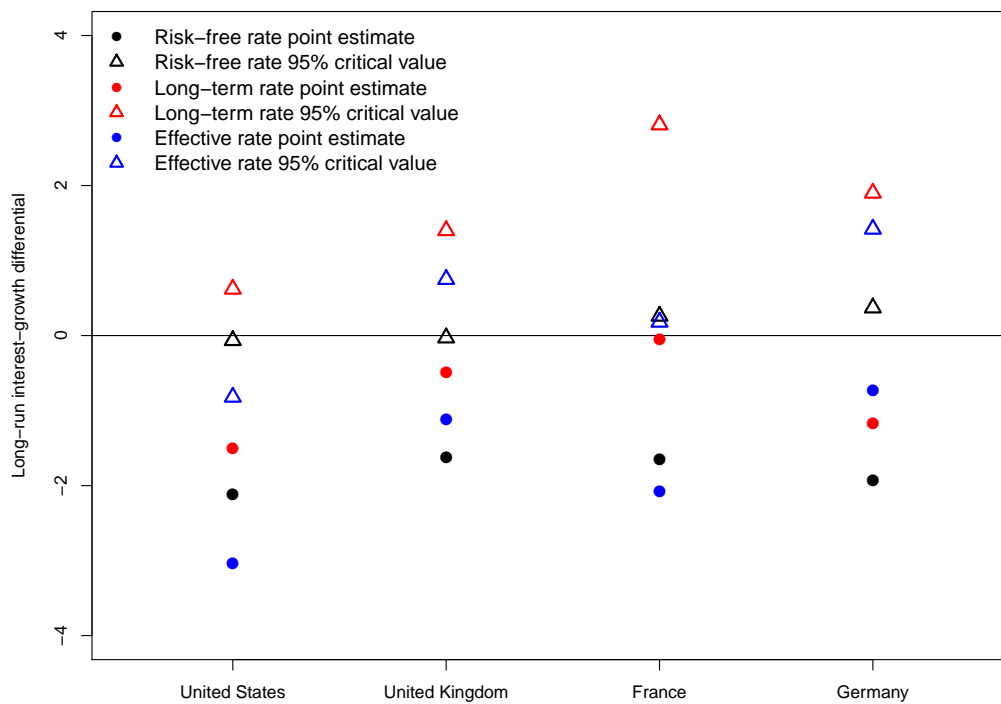
³⁹. Fractional integration of $\frac{1}{2}$ is the cutoff between stationarity and non-stationarity

	USA		United Kingdom		France		Germany	
	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>								
Constant	3.26 (1.359)	0.12 (0.112)	0.44 (0.778)	0.1 (0.095)	-0.19 (1.39)	0.23 (0.157)	1.74 (1.703)	0.36 (0.211)
Growth (-1)	0.41 (0.081)	0.01 (0.007)	0.58 (0.073)	0.03 (0.01)	0.38 (0.092)	0.02 (0.011)	0.63 (0.072)	0.01 (0.009)
Int. rate (-1)	0.05 (0.257)	0.96 (0.022)	0.37 (0.152)	0.95 (0.019)	0.67 (0.249)	0.94 (0.029)	0.11 (0.297)	0.91 (0.037)
<i>Innov. covar.</i>								
Growth	45.55 (5.735)		20.85 (2.622)		42.38 (5.776)		34.39 (4.503)	
Int. rate	0.67 (0.371)	0.37 (0.047)	0.21 (0.25)	0.38 (0.048)	0.93 (0.526)	0.67 (0.092)	0.03 (0.415)	0.59 (0.077)
<i>Long-run</i>								
Mean	5.93 (1.06)	4.18 (1.224)	5.03 (2.13)	4.5 (1.78)	4.96 (2.268)	4.89 (1.602)	6.08 (1.538)	5.04 (0.824)
Difference		-1.75 (1.28)		-0.53 (0.892)		-0.07 (1.078)		-1.04 (1.411)
Log likelihood	569.4		519.8		513.7		543.6	
Observations	134		134		115		128	

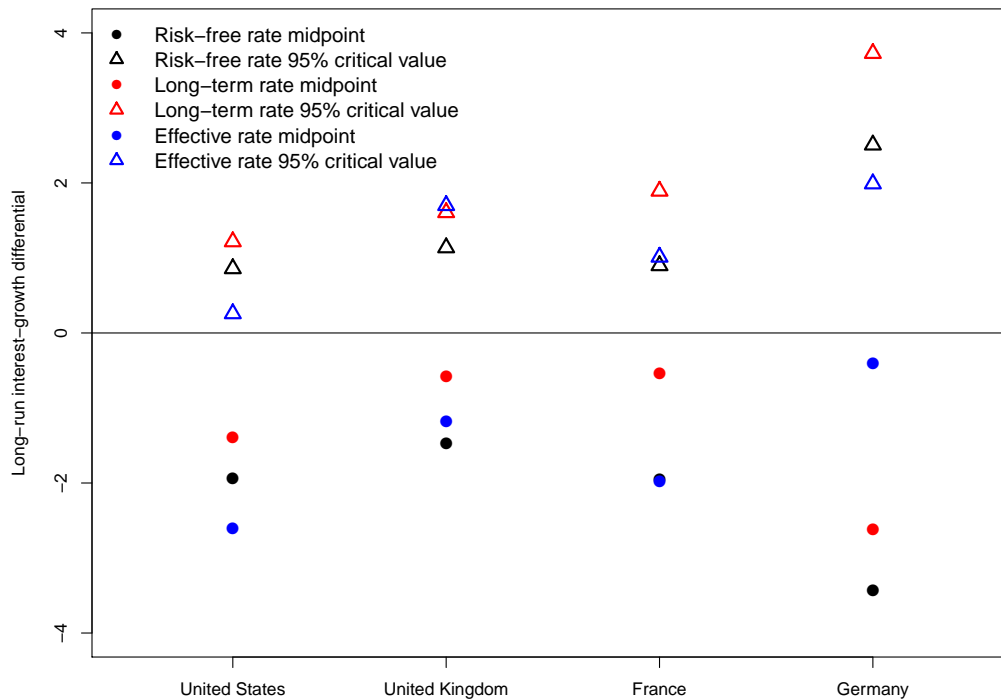
Table 12: 1-lag VAR for sample of countries. Annual data 1880-2013 using long-term government bond yields. Robust likelihood-based standard errors in parentheses.

	USA		United Kingdom		France		Germany	
	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>								
Constant	6.04 (1.973)	0.17 (0.104)	0.51 (0.948)	0.24 (0.146)	3 (1.63)	0.21 (0.141)	1.16 (1.239)	0.37 (0.196)
Growth (-1)	0.4 (0.082)	0 (0.005)	0.63 (0.068)	0.02 (0.011)	0.5 (0.089)	0.01 (0.008)	0.62 (0.079)	0 (0.013)
Int. rate (-1)	-0.72 (0.532)	0.95 (0.028)	0.34 (0.188)	0.93 (0.029)	0.03 (0.366)	0.93 (0.032)	0.15 (0.248)	0.91 (0.04)
<i>Innov. covar.</i>								
Growth	45.58 (5.755)		21.57 (2.745)		33.69 (4.837)		17.99 (2.571)	
Int. rate	0.15 (0.229)	0.14 (0.018)	-0.39 (0.316)	0.57 (0.072)	0.81 (0.324)	0.28 (0.041)	-0.22 (0.306)	0.51 (0.072)
<i>Long-run</i>								
Mean	6.11 (1.146)	3.27 (0.542)	5.58 (1.603)	4.49 (1.056)	6.21 (1.174)	3.98 (0.78)	4.69 (1.144)	4.22 (0.761)
Difference	-2.84 (1.487)		-1.09 (1.118)		-2.23 (1.117)		-0.47 (1.252)	
Log likelihood	498.7		540.5		413.9		418.6	
Observations	132		132		105		106	

Table 13: 1-lag VAR for sample of countries. Annual data 1880-2013 using effective rates on government debt. Robust likelihood-based standard errors in parentheses.



(a) VAR estimates (four lags)



(b) I(0) spectral estimates

Figure 17: Long-run average interest-growth differential using alternative interest rates. Annual data, 1880-2015

bound over the future, and are very sensitive to the particular type of estimator used.

So how should we handle this apparent lack of stationarity in the quarterly data, and what does it mean for the earlier VAR estimates using this data? The most reasonable explanation is simply that the dataset is too short to allow the data to determine the degree of fractional integration in the series, and so we should not put too much weight on these results. The advantage of the spectral approach is that it allows us to consider a more diverse set of statistical processes. But this comes at the cost of having to allow for non-stationary. As the longer annual data series implies stationarity,⁴⁰ the apparent non-stationarity in the quarterly dataset tells us that the interest-growth differential must be a product of the short dataset, rather than anything more fundamental.

	USA		UK		France		Germany	
	Mid.	95% c.v.	Mid.	95% c.v.	Mid.	95% c.v.	Mid.	95% c.v.
I(0)	3.40	5.20	4.90	6.80	4.00	6.00	4.40	6.50
Bayes	-0.70	7.80	-0.70	8.20	0.00	9.10	1.20	10.60
Bayes superset	-1.10	12.20	-1.00	12.90	-0.50	14.00	0.60	15.90
ML fractional integration	1.00	1.00	1.00	1.00	0.90	0.90	0.80	0.80

Table 14: Midpoint and 95% upper critical value of Mueller-Watson prediction ranges for average interest-growth differential of next 100 years. Annualized quarterly data 1956Q1-2016Q4 (from 1961Q1 for Canada).

D Discretizing the growth-interest rate process

The model uses a discrete-state approach. To convert the estimated processes for interest and growth rates, which are continuous, into a form suitable for the model, I develop a method for discretizing multi-dimensional processes accurately and quickly. This consists of a matrix of two-dimensional nodes, each with an associated interest and growth rate (the $R(x)$ and $G(x)$), and a transition matrix governing the probability of transferring between each pair of nodes (the \mathbf{M}).

First, I pick R and D , the number of distances and directions. These generate $RD + 1$ discrete nodes, each with an associated interest and growth rate. The method works by splitting a circle by D equally-spaced rays. Then, nodes are placed along each ray using the Gauss-Hermite quadrature formula for a normal with the conditional density of the long-run density along the

40. Except for Canada; see the fractional integration line of Table 4

ray in question. Intuitively, this spaces the nodes such that they are distributed most densely where the long-run probability of the joint process for interest and growth rates is highest.

The generating formula for the node with index (m, d) is:

$$\mathbf{x}(r, d) = \mu + \sqrt{2r} \mathbf{A} \begin{bmatrix} \cos\left(\frac{d2\pi}{D+1}\right) \\ \sin\left(\frac{d2\pi}{D+1}\right) \end{bmatrix}$$

Where: μ = Long run mean

Σ = Long run variance

$$\mathbf{A}\mathbf{A}^T = \Sigma$$

This gives MD nodes. The remaining node is generated from the same formula using $r = d = 0$, i.e. the long-run mean.

The discretized nodes used in Section 4 are shown in Figure 18, with size in proportion to their long-run probability. The Figure clearly illustrates how nodes are placed along rays, and are more tightly grouped where the process has greater probability mass. The positive correlation of the points comes from the positive correlation of interest and growth rates; there is little incremental value adding nodes at which interest rates are high and growth rates low (or vice versa).

The second step is to select the matrix of transition probabilities, \mathbf{M} . In the spirit of Rouwenhorst (1995), I select the probability of row i minimize the errors on the conditional mean, variance and skew of the Markov process in state i , relative to the continuous process we wish to approximate. At first sight, it might seem as if this can be done exactly. With only three linear restrictions (mean, variance, and skew), then if there are at least 4 states, there exists a probability vector for state i which both sums to one and also matches exactly the three conditional moments. However, the exact solution(s) often entail negative probabilities on some nodes, so we typically cannot match the conditional moments without error using non-negative probabilities. So for each i , I minimize a weighted combination of the target moments.

To test the discretization algorithm, I simulate the fitted Markov chain for 100,000 periods and then estimate a VAR from the simulation. Table ?? compares the estimates from a one-lag VAR based on the US historical data (labeled “Data”), to a a VAR estimated from the simulated Markov process generated via the discretization algorithm, labeled “Simulation”. The coefficients throughout are almost identical. In other words, the discretized process fed into the model well-represents the VAR estimated from the data.

To produce the figures used in the text, I then generate a historical time series for model states by selecting in each time period the node closest to the data. Figure 19 shows the fitted discretized series in comparison to the US data. The discretization handles the volatility in growth

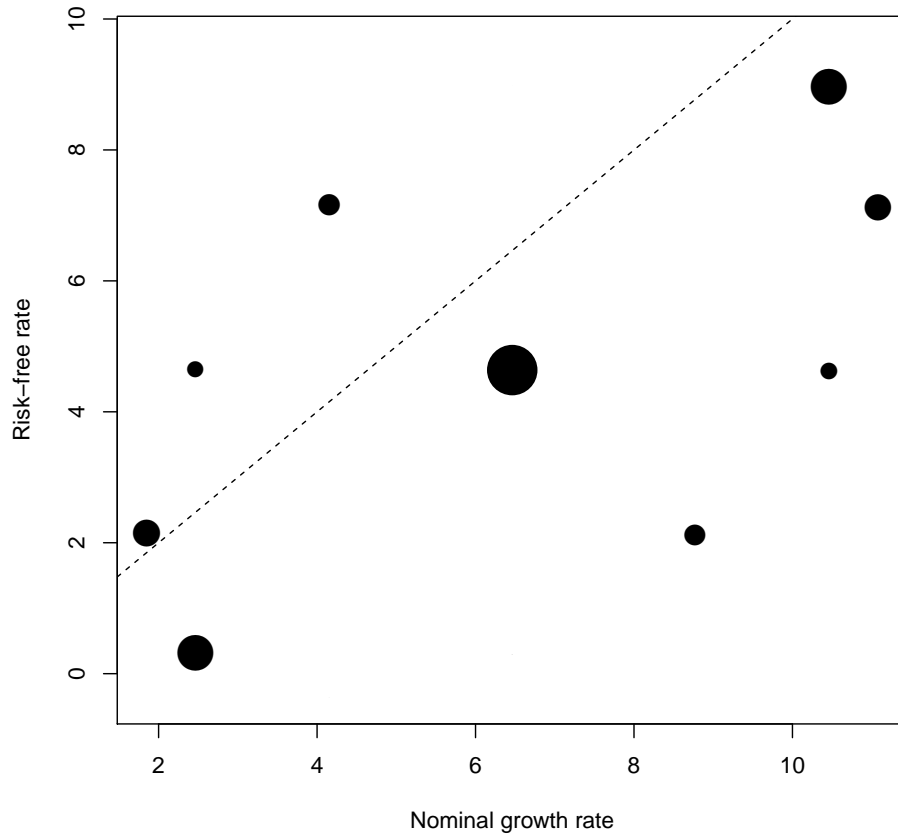


Figure 18: Discretized values for nominal interest and growth rates with $D = 20, R = 3$. Points shown in proportion to their ergodic frequency.

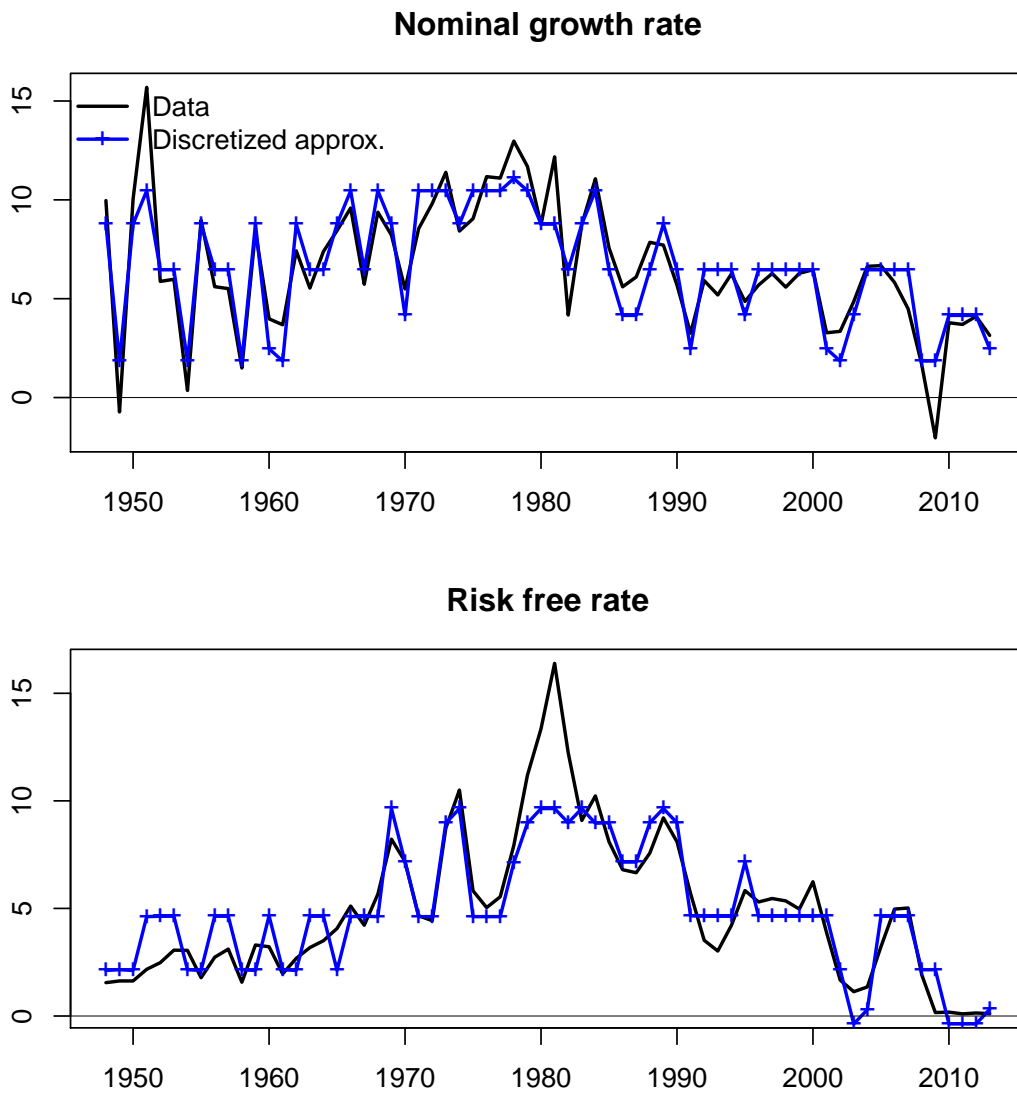


Figure 19: Time series of nominal interest and growth rates: data in black, discretized process in blue.

	Data		Simulation	
	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>				
Constant	3.98	-0.48	2.65	0.01
LR Mean	6.46	4.64	6.46	4.63
Growth (-1)	0.28	0.24	0.46	0.28
Int. rate (-1)	0.14	0.77	0.19	0.61
<i>Innov. covar.</i>				
Growth	9.08		8.19	
Int. rate	2.19	2.34	1.85	1.79

Table 15: Estimated VAR based on data, and VAR estimated from discretized approximation simulated 100,000 periods

rates well, and also covers the large range of both series adequately. However, in some places the discretization inevitably has some limitations, struggling in particular during World War II, where financial repression meant that the usual correlation between interest and growth rates was broken.