WP/18/82

IMF Working Paper

Interest-Growth Differentials and Debt Limits
in Advanced Economies

by Philip Barrett

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

WP/18/82

IMF Working Paper

Interest-Growth Differentials and Debt Limits in Advanced Economies

by Philip Barrett

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

I N T E R N A T I O N A L M O N E T A R Y F U N D

IMF Working Paper

Fiscal Affairs Department

Interest-Growth Differentials and Debt Limits in Advanced Economies

Prepared by Philip Barrett¹

Authorized for distribution by Catherine Pattillo

April 2018

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

Abstract

Do persistently low nominal interest rates mean that governments can safely borrow more? To address this question, I extend the model of Ghosh et al. [2013] to allow for persistent stochastic changes in nominal interest and growth rates. The key model parameter is the long-run difference between nominal interest and growth rates; if negative, maximum sustainable debts (debt limits) are unbounded. I show how both VAR- and spectral-based methods produce negative point estimates of this long-run differential, but cannot reject positive values at standard significance levels. I calibrate the model to the UK using positive but statistically plausible average interest-growth differentials. This produces debt limits which increase by only around 5% GDP as interest rates fall after 2008. In contrast, only a tiny change in the long-run average interest-growth differential – from the 95th to the 97.5th percentile of the distribution – is required to move average debt limits by the same amount.

JEL Classification Numbers: E62, H62, H63

Keywords: Fiscal policy, Government Debt, Uncertainty

Author's E-Mail Address: pbarrett@imf.org

¹ I am grateful for helpful comments from Tamon Asonuma, Olivier Blanchard, Luc Eyraud, Vitor Gaspar, Jun Kim, Divya Kirti, Catherine Pattillo, Abdel Senhadji, Leonardo Martinez, Raphael Espinoza, Mark Watson, Philippe Wingender, and participants of the FAD seminar series and 2017 MFM Summer Session.

1 Introduction

Government debt in advanced economies is very high. In many countries debt relative to GDP has risen steadily since 1980 (Figure 1), and is now at or near its highest in the post-war period. During the same time period interest rates have declined, including quite dramatically during the global financial crisis of 2008 (Figure 2), lowering the cost of borrowing for the private and public sectors. This raises two related questions: quite how much public debt such countries can sustain? And has the maximum sustainable debt level increased as interest rates have fallen?

Figure 1: General government gross debt to GDP ratio, percent, in G7 economies 1980-2016 (annual data)

Figure 2: Central Bank interest rates in G7 economies 1980-2016 (annual data)

This paper is an attempt to address these issues, and argues that the critical factor determining a country's maximum sustainable debt level is the difference between its future nominal interest and growth rates. This interest-growth differential determines the rate at which a country's public debt rises relative to its output, sometimes termed the "natural" debt dynamics. A higher interest-growth differential means that a country must raise larger surpluses in order to stabilize its debt-GDP ratio¹.

This issue is directly relevant for policymakers today, who face interest-growth differentials that are very low by the standards of the last 20 years. This paper maps those latest changes into their consequences for debt sustainability. In addition, understanding how changing interest-growth differentials affect debt sustainability also matters for our interpretation of the past. For

1. Or, equivalently, that a country can only service a lower debt given its surpluses.

example, the worldwide rise in the differential during 1980s and early 1990s will have reduced sustainable public debt levels. The analysis in this paper tells us by how much².

In what follows, I examine historical interest-growth differentials in Advanced Economies since 1960, and show that these have exhibited three notable characteristics. First, average risk-free interest-growth differentials are usually less than zero. Since 1880, the average annual interest-growth differential in six advanced economies has been -1.7 percent, and -0.8 percent since 1960. Second, that there have been persistent medium-run deviations from this average level. For example, during the era of relative peace following 1960, decadal average interest-growth differentials have varied from around minus five percentage points per year in the 1960s and 70s, to nearly two percentage points in the 1980s and 90s, before dropping back below zero more recently. Third, the short-term volatility of interest-growth differentials in these countries is very high, frequently changing by several percentage points per year.

These three aspects of the data have contrasting implications for maximum debt levels. Low average differentials mean that maximum sustainable debt levels are likely to be high, consistent with the lack of debt distress in these countries. Persistent medium-run movements suggest that maximum sustainable debt levels may vary substantially over time. And because large high-frequency volatility induces unpredictable variation in the debt-GDP ratio, this may increase default risk overall, pulling down maximum sustainable debt levels.

While all three of these properties have important implications for sustainable debt levels, the first (the long-run differential) is paramount. Point estimates of the long run average interest-growth differential in advanced economies are frequently negative. If true, the consequences are rather unpalatable: unless governments can commit to infinitely large deficits, they can issue as much debt as they like without becoming insolvent³. Given this, it seems wise to test statistically whether the long-term interest-growth differential may, in fact, be positive.

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. I apply this to two datasets on interest and growth rates in five advanced economies, one annual since 1880 and one quarterly since 1956. For both datasets, the point estimates from the VAR are negative⁴, but the associated likelihood ratio test cannot reject the possibility that the long-

2. The short answer: not by much, and certainly by less than our margin of error on these estimates.

3. This does not mean that high debt levels are *necessarily* safe. Even sovereigns with sufficient future income to repay their current debts can be exposed to self-fulfilling debt crises if their gross financing needs are large, as first discussed in Cole and Kehoe (2000).

4. 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. 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 the “equity premium puzzle”.

run interest-growth differential is positive. This upper bound is typically in the order of zero to one percent per year. 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 unlikely, but still consistent with the data.

I then use a model to map the estimated process for interest and growth rates into conservative-but-plausible debt limits. 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. The particular framework used to map the statistical behavior of interest and growth rates into debt limits is based on Ghosh, Kim, Mendoza, Ostry, and Qureshi (2013). In common with that paper, the model used here takes as given the estimated function relating the average surplus-GDP ratio to the aggregate state of the economy, including outstanding debts. Deficits are funded by issuing noncontingent defaultable bonds to risk-neutral investors. The model is subject to correlated aggregate shocks to the nominal risk-free rate (which determines investors' cost of funds) and the growth of nominal GDP (which determines the value of future surpluses), as well as idiosyncratic shocks to governments' primary balances.

A debt limit in this model is the debt level at which the government loses access to debt markets. Investors price government debt at a premium over the risk free rate, reflecting default risk. As public debt increases, the probability that realized future surpluses will be high enough to repay these debts falls. The default premium charged on the debt reflects this probability. But a higher default premium also makes repayment more onerous, and so more default more likely. If debt is high enough, though, any finite default premium will cause such high debt service obligations that default is almost a certainty. As a result, investors refuse to lend to the government at any finite interest rate. In other words, government lose market access. The debt level at which this occurs is the model's debt limit.

To illustrate the key aspects of the relationship between interest-growth differentials and debt limits, I use the model to generate conservative estimated historical debt limits for the UK. 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 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 2.5% critical value for the VAR-based estimates (instead of the 5% value) leads to maximum sustainable debt-GDP ratios which are about 5pp lower. More intuitively, the increase in fiscal space due to the large declines in interest rates since 2008 would be entirely offset by a fractionally more conservative approach to uncertainty from just one key parameter: the long-run interest-growth differential.

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

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 the same model as this paper, but extended for stochastic variation in interest and growth rates⁵. Ostry and Kim (Forthcoming) and Kim and Asonuma (Forthcoming) develop versions of this model with uncertainty over growth, but without persistent shocks or uncertainty over the risk-free interest rate.

The second related literature is that on measuring 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, but limited to 1991-2008, Escolano (2010) finds interest-growth differentials are typically small and positive⁶. More recently, Mehrotra (2017) recovers negative average differentials, using data from advanced economies in the post-war period. The fundamental difference between those papers and this one is that I go beyond point estimates and formulate statistical tests of the long-run interest-growth differential. 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.⁷ In contrast, I take the behavior of interest-growth differentials as a given, purely

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

6. Although with considerable heterogeneity; many countries are estimated to have a negative differential.

7. And a yet further group of papers attempts to explain why real rates have fallen during the same period,

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).

Last, while the model in this paper is one of capacity to repay rather than strategic default, the stochastic driver of debt distress shares much in common with models of strategic sovereign default. 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.

The paper proceeds as follows. Section 2 motivates the paper, highlighting the sensitivity of debt limits to interest-growth differentials in a very simple model of fixed debt limits, as well as providing a brief overview of the data. Section 3 introduces a larger model of dynamic debt limits driven by shocks to interest and growth rates, and highlights how the model predictions are crucially dependent on the sign of the long-run differential. Section 4 then estimates processes for the interest and growth rate shocks which drive the model, and tests statistically the sign of the differential. Section 5 applies the model to UK data and discusses the sensitivity of the resulting debt limits to short- and long-run changes in the interest-growth differential, and assumptions such as shock volatility and debt maturity. Section 6 concludes.

2 Motivation

This section outlines a very simple model of debt sustainability, emphasizing how basic notions of sustainable debt levels are sensitive to the interest-growth differential. Then follows a brief overview of the empirical properties of interest-growth differentials in advanced economies as well as a discussion of their likely implications for sustainable debt levels viewed through the lens of the simple model.

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)

2.1 A motivating model

Consider an economy with constant nominal growth rate $1 + g$ and risk-free nominal interest rate $1 + r$. The government chooses its *expected* primary surplus s_t^e from within the finite range $[\underline{s}, \bar{s}]$, but the realized surplus is stochastic. It is given by $s_t = s_t^e + u_t$, where u_t is a mean-zero random variable with support $[-\bar{u}, \bar{u}]$ ⁸.

The government finances deficits by issuing one-period zero coupon nominal bonds. The price of these bonds is q_t , and their face value is b_t . The government budget constraint is therefore:

$$q_t b_t + s_t = \frac{b_{t-1}}{(1 + g)}$$

In order to identify the maximum debt could be sustained, the government is assumed to default on its outstanding debt if and only if meeting the government budget constraint would mean that future debt was unbounded, and that it will always raise expected primary surpluses $s_t^e = \bar{s}$ to avoid a default.

For simplicity, assume that investors are infinitely risk averse; they assign zero value to a bond with any non-zero default risk. Then the bond price is:

$$q_t = \begin{cases} 1/(1 + r) & \text{If the probability of default is zero} \\ 0 & \text{Otherwise} \end{cases}$$

Then two results follow (see Appendix A.1 for proofs):

Proposition 1

1. *If $r > g$ then there exists $\bar{b}$, the maximum sustainable debt. This is the largest b^* such that the government defaults whenever $b_{t-1} > b^*$. This is given by:*

$$\bar{b} = \frac{(\bar{s} - \bar{u})(1 + g)(1 + r)}{r - g}$$

2. *If $r < g$ then there exists no such $\bar{b}$. The government never defaults, no matter the debt level.*

From this simple model we learn two key lessons about the relationship between the interest-growth differential and the maximum debt a government can sustain. First, that when it exists,

8. An upper bound on the primary surplus function arises because, at the very least, the amount of resources that the government can expropriate are finite. Although in practice, political economy reason, or bounds imposed by a tax Laffer curve may be more important. The lower bound simply says that the government cannot run an infinite deficit.

the debt limit is enormously sensitive to the interest-growth differential, particularly when $r - g$ is close to zero⁹.

Second, that the sign of this differential is critical for even the existence of a sustainable debt limit. If the difference between the risk-free rate and the nominal growth rate is negative, then there is *no* level of debt which is unsustainable, so long as deficits are bounded.

More sophisticated models will produce different formulae for the maximum sustainable debt. But they will retain these key properties. For example, if r and g are stochastic, or the expected surplus is a bounded function of the debt, the path for debt will diverge to infinity with probability one so long as the long-run expectation of $r - g$ is negative. Proposition 3 proves this in a more general setting.

2.2 Data on interest-growth differentials

Figure 3 shows two measures of the annual risk-free interest-growth differential for the USA. Data are taken from Jordà, Schularick, and Taylor (2017). The two interest rates used are a short-term rate (equal to the Federal Funds Rate in the post-war period) and a long-term rate.

Figure 3: US interest-growth differentials 1880-2016

9. The risk preference assumption here is not essential for this result; it simply gives an analytic solution for $\bar{b}$. If investors are infinitely risk-loving, then $\bar{b} = (\bar{s} + \bar{u})(1 + g)(1 + r)/(r - g)$. And if investors have some intermediate degree of risk tolerance, then the debt limit is somewhere between these two extremes.

Three features stand out. First, that negative values are far from rare, and in fact are quite common. Using the short-term rate, the average interest-growth differential was -1.9 percent during 1880-2016, and -1.6 percent since 1950. So recent declines in the interest-growth differential since 2008 are not so much a rare event as a return to historically common levels. Second, the interest-growth differential departs substantially from this average for long periods of time. The decadal average is -2.5 percent per year in the 60s and 70s, but 2.5 in the 80s. Third, that at short frequencies the differential is very volatile. In the USA it moves by over 2.5 percentage points in more than one-fifth of years. This can cause the debt-GDP ratio to jump considerably; a 2.5 percentage point fall in the annual growth rate¹⁰ causes the debt-GDP ratio to increase mechanically by 2.5 percentage points.

These patterns are not unique to the United States. Figure 4 shows the five-year mean interest-growth differentials using short term rates in seven advanced economies (data come from Jordà, Schularick, and Taylor (2017)). All exhibit a similar pattern: interest-growth differentials are typically negative, but deviate persistently from their long-run average. They are also volatile at short horizons, much as in the US. In addition, these facts hold true even when we restrict our attention only to periods of (relative) peace, such as from 1960 onward.

Figure 4: Five-year average annualized interest-growth differentials for a sample of advanced economies: 1960-2016

10. Changes in the interest rate have similar effects in the long run, but the timing of the impact is complicated by debt maturity.

Viewed through the lens of the simple motivating model, these three aspects of interest-growth differentials in advanced economies are highly relevant for sustainable debt levels. Low average differentials may help to explain why high debt levels can be sustained, even in the face of continued deficits. Persistent changes in differentials raise the possibility that major episodes in economic history, such as the Volcker disinflation (where interest rates in the US rose far above growth rates), may lead to debt limits that change significantly over time. And the high short term volatility will increase default risk, pulling down on debt limits. The next section outlines a model which can quantify these effects.

3 Model

This section outlines a model which can systematically connect changes in interest-growth differentials to maximum sustainable debt limits. The model does this by providing a government with a stream of future surpluses against which they can borrow to fund current deficits and debt repayments. The government's ability to borrow is therefore impacted by two sources of risk.

First, the growth rate of surpluses is tied to that of nominal GDP, which is stochastic. This affects the resources available to the government to repay current obligations. Second, the rate at which investors discount future payoffs, i.e. the risk-free nominal rate, is also stochastic. This affects the ability of the government to roll over future debts. The key summary statistic of the government's debt sustainability is the difference between these two variables, the interest-growth differential.

The model can be thought of as a version of that of (Ghosh et al. 2013) with two extensions and one restriction. The first extensions is that the interest and growth rates are stochastic, depending on a finite-state Markov process. The second extension is that the surplus process is also permitted to covary with interest and growth rates, capturing the notion that fiscal policy may be pro- or counter-cyclical. The restriction is that the surplus function is bounded, meaning that the government is unable to run arbitrarily large surpluses or deficits.

3.1 Environment

I model the government as taking as given nominal risk-free interest rates and nominal GDP growth rates. Each period, an exogenous state x_t is drawn from a N -state discrete Markov process with transition matrix M . The risk-free and nominal growth rates are functions of the realized states, so $R_t = R(x_t)$ and $G_t = G(x_t)$. This allows for arbitrary cross- and auto-

correlation of these variables. I assume that x_t is ergodic¹¹

The government raises surpluses according to a surplus rule. Then surpluses depend on outstanding debt d_{t-1} , the aggregate state x_t and a mean-zero shock:

$$s_t = s(d_{t-1}, x_t) + e_t, \quad e_t \sim F(\cdot), \quad \mathbb{E}e_t = 0 \quad (1)$$

I assume that $s(d, x)$ is weakly increasing in d and bounded. So the government is able to raise a larger surplus to repay a larger debt, but that this ability is limited. (Ghosh et al. 2013) term this effect “fiscal fatigue”¹². I also assume that $F(\cdot)$ has finite support $[-\bar{\epsilon}, \bar{\epsilon}]$.

The government has access to a non-contingent one-period bond. The government’s end-of-period debt to GDP ratio is denoted d_t , which is valued at market prices. When debt is repaid, this evolves according to the government budget constraint:

$$\frac{d_{t-1}}{G(x_t)q_{t-1}} = s_t + d_t \quad (2)$$

This equation says that maturing debt due in period t is paid by surpluses plus new debt issuance¹³. This is expressed in terms of period t GDP by dividing by the GDP growth rate G_t . The government is prevented from accumulating assets, so $d_t \geq 0$.

Investors are risk-neutral¹⁴ and their opportunity cost of funds in period t is given by the risk-free rate R_t . In the event of default, investors reclaim a share ϕ of outstanding liabilities. If investors in period t expect that the government will default in the next period with probability p_t , then the debt price satisfies the risk-neutral pricing condition¹⁵.

$$q_t = \frac{1}{R(x_t)} [p_t \phi q_t + (1 - p_t)] \quad (3)$$

11. This would hold, for example, if x_t were irreducible (i.e. all states are accessible from all other states) and has at least one aperiodic state (i.e. at least one element of the diagonal of M is non-zero).

12. In fact, they go further, assuming that $\lim_{d \rightarrow \infty} s(d, x) = -\infty$. When the interest-growth differential is negative this will produce finite debt limits. But, as we saw in the motivating example of section 2.1, and prove in Proposition 3, a bounded-surplus model will not. In other words, the (Ghosh et al. 2013) model will produce finite debt limits when interest-growth differentials are negative *only* because the government can create infinitely large deficits

13. This equation is still valid despite quantitative easing programs, such as those pursued in recent years in Advanced Economies, so long as the government commits to repay obligations held by the central bank, i.e. there is no money financing of the deficit.

14. Risk aversion is not hard to introduce in this model. Simply replacing the probability of default, p_t , with p_t^θ for some $\theta > 1$ would produce generate bonds which trade at a discount to their risk-free price. The motivating example of Section 2 effectively considered $\theta = \infty$.

15. Note that in equation (3) the repayment in default is proportional to the price in the preceding period. This is an important technical assumption, as it prevents the debt price from being bounded below by ϕ . This means that recovery-in-default is proportionate to the loan’s principal, rather than the principal plus interest.

The model is closed with an equilibrium condition connecting the probability of default p_t to the government's default rule. This is most easily expressed in the recursive formulation of the next section.

3.2 Recursive equilibrium

To solve the model numerically, I express it in recursive form. Intuitively, equilibrium consists of three state-dependent objects: debt prices that are consistent with investors' beliefs about the probability of default; debt thresholds which are the maximum debt the government can sell at those prices; and investor's beliefs about default probabilities p_t equal to the probability of default given debt thresholds. Given that fiscal policy follows a pre-determined rule (much like monetary policy is assumed to when following a Taylor rule, for example), there is no notion of strategic default here.

The ability of a government to repay its future obligations can be summarized by two state variables: x , the aggregate state; and d , the market value of outstanding government debt. We will express the equilibrium recursively in terms of these states. For the purposes of exposition, we will start by assuming that the debt price and default threshold period $t + 1$ are known. Then we write down equations for period t defining the debt price, default probability and default threshold consistent with each other and the assumed rules for period $t + 1$. Equilibrium occurs when the period t debt price and default threshold rules thus defined equal those assumed to hold in period $t + 1$.

Formally, denote by p , investors' period t beliefs of the probability of default in period $t + 1$. Then the period debt price in period t is a function of this, $\hat{q}(x, p)$ which satisfies:¹⁶

$$\begin{aligned}\hat{q}(x, p) &= \frac{1}{R(x)} [p\phi\hat{q}(x, p) + (1 - p)] \\ \Rightarrow \hat{q}(x, p) &= \frac{1 - p}{R(x) - p\phi}\end{aligned}\tag{4}$$

Let $D(x_{t+1})$ be the debt price and default rule in period $t + 1$. This means that $D(\cdot)$ is a vector of debt levels such that $a + t = 0$ if $d_{t+1} > D(x_{t+1})$. Given the period t price function $\hat{q}(x, p)$ and a default rule $D(\cdot)$, the probability of default in period t by a government with d units

16. Note that when $p = 1$, the only solution to equation (4) is that $\hat{q}(x, p) = 0$. This is a consequence of the way that recovery rates are parameterized; they are expressed proportionate to the purchase price of the debt. Looking forward, this will be important when we come to look for equilibrium, which requires that expected default probabilities equal realized default probabilities. In this framework we can always guarantee that $p = 1$ satisfies this restriction, as a zero price will mean that next-period default is inevitable.

of outstanding debt is given by:

$$Z(x, d, p, D) = \sum_{x'} \pi(x'|x) F(H(x, x', d, p, D(x'))) \quad (5)$$

$$\text{Where:} \quad H(x, x', d, p, D(x')) = \frac{d}{G(x')\hat{q}(x, p)} - D(x') - s(x', d)$$

So $H(\cdot)$ is the minimum surplus shock required in state x' to push the period $t + 1$ debt level over the default threshold $D(x')$. And $Z(\cdot)$ is the sum of the probabilities of default in each state, weighted by the state-transition probabilities $\pi(x'|s)$. Note the timing here: uncertainty is only over the growth rate $G(x')$. In contrast, the risk-free interest rate $R(x)$, which affects the price $\hat{q}(x, p)$, is pre-determined.

In equilibrium we require that investor beliefs are rational. That is, the realized probability of default defined by $Z(\cdot)$ equals investors perceived probability of default p . Following Ghosh et al. (2013), we define $p^*(x, d, D)$ be the smallest fixed point of:

$$p = Z(x, d, p, D) \quad (6)$$

Such a $p^*(x, d)$ always exists because there is always a fixed point at $p = 1$ for any x and any $d > 0$. This follows from the set-up of the recovery ratio in equation (4). As noted before, $\hat{q}(x, p) = 0$ if $p = 1$. Therefore $H(x, x', d, p) = \infty$, and so $Z(x, d, p) = 1$. More intuitively, because recovery scales with the present value of the loan (as opposed to its face value), the expected recovery in default is zero if investors think default is assured. The bond price therefore also falls to zero, requiring infinite repayments on any non-zero debt, and so guaranteeing default.

The period t default boundary is a vector $\hat{D}$ such that:

$$\exists \quad p^*(x, d, D) < 1 \quad \forall x, \forall d \leq \hat{D}(x) \quad (7)$$

$$\nexists \quad p^*(x, d, D) < 1 \quad \forall x, \forall d > \hat{D}(x) \quad (8)$$

So $\hat{D}(\cdot)$ is the threshold above which the government defaults for sure in each state. A recursive equilibrium is when the period $t + 1$ default thresholds and prices equal those in period t . Formally:

Definition 2 (Recursive equilibrium) *A recursive equilibrium is given by functions $D(x')$, $\hat{q}(x, p)$, $p^*(x, d, D)$, $\hat{D}(x)$ such that:*

1. $\hat{q}(x, p)$ satisfies equation (4) for all (x, p) .
2. For all (x, d) , $p^*(x, d, D)$ is the smallest p satisfying equation (6).

3. Given $p^*(x, d, D)$, $\hat{D}(x)$ satisfies equations (7) and (8).

4. Current and future default thresholds are equal:

$$D(x) = \hat{D}(x) \quad \forall x$$

5. If $\{d_t\}_{t=0}^\infty$ is the sequences of debt levels generated by these rules, then:

$$\lim_{T \rightarrow \infty} \mathbb{P}(d_T > d) = 1 \quad \forall d > 0 \quad \Rightarrow \quad d_0 > D(x_0)$$

This definition requires that three key objects must be mutually consistent at every debt level. Prices must be consistent with investors' beliefs about the probability of default. Investors' beliefs must be consistent with the true probability of default given prices and the government's default rule. And the default rule must be the highest debt that the government can raise at the consistent prices. The fourth condition requires that the default rule is stable. And the final requirement of equilibrium says that if the government debt grows without bound, then the government must default. This is the stochastic analogue of the idea that the debt limit is the “point of no return”: that if the government debt were to grow without limit, then it must be above the debt limit¹⁷.

3.3 Properties of equilibrium

We now show two key properties of equilibrium:

Proposition 3 *There exists an equilibrium with $D(x) = \infty \quad \forall x$ if and only if $\mathbb{E} \log(G_t/R_t) > 1$.*

Because $\mathbb{E} \log(G_t/R_t) \simeq \mathbb{E}(g_t - r_t)$, where g_t and r_t are the net interest and growth rates, then this proposition says: debt limits are infinite if and only if on average over the long run, the nominal growth rate is larger than the *risk-free* interest rate. This generalizes the existence result for the simple example in Proposition 1, which also depended only on the risk-free differential. A formal proof is in Appendix A.2. But the intuition is that for an equilibrium when $D(x) = \infty \quad \forall x$, the risk premium is always zero, and so the path of debt depends only on the risk-free interest-growth differential. If this is positive, then the debt level will diverge to infinity with probability one. But because d_0 is finite, then $d_0 < D(x_0)$, violating the fourth equilibrium condition. This is the crucial theoretical justification for focusing on the risk-free interest-growth differential in the

17. Or, as Ghosh et al. 2013 say, “beyond $\bar{d}$ [the debt limit], ... debt grows continuously and the government necessarily defaults, which is the definition of the debt limit.”

empirical work that follows: it is the long run value of the risk-free differential which determines whether debt limits are finite.

Proposition 4 *At an equilibrium, a tangency condition holds:*

$$\frac{\partial}{\partial p} Z \left(x, \hat{D}(x), p^*(x, \hat{D}(x)), \hat{D} \right) = 1 \quad (9)$$

The proof of this condition follows (Ghosh et al. 2013), although extended to a multi-state case, and is included in Appendix A.3. This tangency condition reduces computation of equilibrium to the solution of a set of simultaneous equations. With N states of the world, the N tangency conditions and N fixed-point requirements provide $2N$ equations, which can be solved for $2N$ unknowns: the default probabilities and thresholds for each state. This ease of solution is the modeling payoff for assuming one-period debt. There is no analogous condition to equations 9 in a model with long-maturity debt. The cost, though, is that the model will over-state the government's exposure to risk from fluctuations in interest-growth differentials. Section 5.4 attempts to quantify and correct for this.

In Section 5 we solve this model numerically, calibrated to the UK. But first we address the key empirical issue for this model: what is the process defining $R(x)$ and $G(x)$?

4 Empirical properties of interest and growth rates

This section has two purposes. First, to estimate a dynamic stochastic process for $R(\cdot)$ and $G(\cdot)$ that can be fed into the model to produce predicted debt limits. Second, 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.

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 this section proceeds first by estimating 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. Mehrotra (2017) presents a complementary approach, using a probit regression to infer the probability that

future interest rates will be positive at a given horizon, and identifies the determinants of this event.

I find that the evidence on the long run interest-growth differential is mixed. 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 values¹⁸. But because this is a much more general approach, these confidence sets are wider.

4.1 Data on interest and growth rates

I use two separate datasets, each covering a set of major Advanced Economies¹⁹, 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 data²⁰. For all countries the data covers 1880-2015. 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.

Country	Observations	Years
USA	134	1880-2015
UK	134	1880-2015
France	114	1880-1913, 1922-1938, 1951-2015
Germany	128	1880-1921, 1925-1944, 1950-2015

Table 1: Sample periods for interest and growth rates: Annual data

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, 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

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

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.

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).

Data are for nominal interest and growth rates simply because debt is typically issued on nominal terms. And hence debt sustainability is driven by nominal shocks as much as real ones. For example, if inflation is higher than expected, the debt-GDP ratio falls, as nominal growth outstrips the nominal interest rate. So in order to capture the stochastic changes in interest and growth rates that are relevant for debt sustainability, I estimate stochastic processes for nominal variables²¹.

I focus policy interest rates in the empirical analysis for three reasons. First, for consistency with the model. The conceptual framework that produces debt limits described in section 3 is driven by the risk-free rate. Second, the key determinant of debt sustainability is the future interest-growth differential. A measure such as the effective rate is inherently backward-looking, as it is lagging combination of past interest rates, which (depending on the sample) may be systematically above or below future interest rates. Third, for reasons of data availability; short-term policy rates are more consistently defined, and available at higher frequencies over longer periods than other interest rates. Nevertheless, in Appendix B.2, I re-run the analysis using interest rates on long-term government bonds and effective interest rates, and conclude that the main findings are much the same.

4.2 Describing the dynamic behavior of interest and growth rates

To get a broad sense of the data, Table 2 presents the results of a simple one-lag VAR for the risk free rate and nominal growth rate, estimated on annual data during 1880-2015. All rates are annualized. Later, I use longer-lagged VARs to capture the dynamic behavior of interest and growth rates in a more sophisticated fashion, as well as for inference. But Table 2 helps us get a sense of the data initially. Note that the second row of the coefficient table is the long run mean, which is a nonlinear combination of the parameters²².

Three points are worth noting about the dynamic behavior of nominal interest and growth in all seven countries during this period. First, that in both datasets, interest rates are highly

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.

22. 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 $\bar{y} = (I - \sum_{i=1}^n A_i)^{-1}a$

persistent but past nominal growth appears to hold little predictive power for interest rates²³. Second, that nominal GDP growth is highly volatile, annual standard deviations of 6 percentage points are common, and has a lower persistence than interest rates. Third, that the innovations to the two series covary positively. The correlation is around 0.3 for the USA and similar in other countries.

	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)
LR 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)
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)
Log likelihood	687.5		610		542.1		596.1	
Observations	134		134		114		128	

Table 2: 1-lag VAR for sample of countries. Annual data 1880-2015. Robust likelihood-based standard errors in parentheses.

Although the model depends on interest and growth rates separately, their average difference is a crucial summary statistic for the existence of finite debt limits, as we saw in the simple motivating model. Even in the simple specification reported in Table 2 the long run interest-growth differential is negative, at, between -1.6 and -2.1 for all countries. However, this is imprecisely

23. See, for example, the coefficient of 0.03 on lagged growth for the US interest rate regression

estimated. For the US, the standard error of the interest-growth differential is about 1.4²⁴. A simple Wald test of the hypothesis that the long-run interest-growth differential is less than zero is has a p-value of 0.06. So even in this very simple statistical model we cannot reject the hypothesis that the interest-growth differential is zero at standard significance levels, such as 5%. The next subsection explores this idea more rigorously.

4.3 VAR-based inference on the long run interest-growth differential

While the results from the simple specification in Table 2 are interesting, it is purely illustrative. In this section, I test more seriously the idea that the interest-growth differential is negative.

I start by estimating a more flexible empirical specification, in this case a long-lagged VAR. Then, I employ a likelihood ratio test to answer the question: what is the largest long run interest-growth differential which is consistent with the data? Measuring the statistical uncertainty of this parameter allows us to make plausibly conservative estimates of the debt limit. Conservative because they are based on interest-growth differentials higher than are implied by point estimates; plausible because the long-run interest-growth differentials cannot be rejected by the data.

Tables 3 and 4 show the results from four-lag VARs on annual and quarterly data for the risk-free interest and nominal growth rates²⁵. Again, the second row shows the long-run means implied by the dynamic model. The extended lag structure is not so straightforward to interpret, in that the auto and cross-correlations cannot be read off so easily from the parameter estimates. But two points obviously carry over from the simpler setting. First, while the extra dynamic information leads to slightly different point estimates compared to the 1-lag specification, differences in estimates of the means are negative for all countries. Second, innovations to interest and growth are still positively correlated, and growth is still much noisier.

We now use these estimates to test statistically what is the long-run mean of the interest-growth differential. Formally, let θ be the estimated interest-growth differential for some VAR.

24. This estimate comes from the delta method. Alternatively, the standard error of the interest-growth differential can be thought of as the standard error of the mean for the interest rate, plus that of the growth rate, less twice the covariance, which is positive.

25. 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.

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

Table 3: Multi-lag VAR for sample of countries. Annual data 1880-2015. Robust likelihood-based standard errors in parentheses.

	USA		France		United Kingdom		Canada	
	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>								
Constant	1.74	-0.28	1.35	-0.01	1.33	0.04	1.57	-0.15
	(0.565)	(0.124)	(0.58)	(0.122)	(0.681)	(0.113)	(0.635)	(0.118)
LR Mean	6.02	4.4	6.79	4.94	7.14	6.01	6.75	5.39
	(0.827)	(1.403)	(1.489)	(1.639)	(1.675)	(1.693)	(1.178)	(1.785)
Growth (-1)	0.35	0.07	0.13	0.03	0.1	0	0.63	0.06
	(0.068)	(0.015)	(0.067)	(0.014)	(0.067)	(0.011)	(0.072)	(0.013)
Int. rate (-1)	0.32	1.14	0.67	0.98	0.91	1.12	0.07	1.15
	(0.312)	(0.067)	(0.322)	(0.067)	(0.403)	(0.067)	(0.404)	(0.073)
Growth (-2)	0.21	0.03	0.36	0.03	0.26	0.02	-0.05	0
	(0.071)	(0.015)	(0.067)	(0.014)	(0.064)	(0.011)	(0.087)	(0.016)
Int. rate (-2)	-1.62	-0.49	-0.75	-0.01	-1.55	-0.18	-0.51	-0.37
	(0.459)	(0.099)	(0.449)	(0.093)	(0.602)	(0.101)	(0.61)	(0.11)
Growth (-3)	0.04	0	0.14	0.01	0.29	0.03	0.1	0
	(0.073)	(0.016)	(0.068)	(0.014)	(0.065)	(0.011)	(0.087)	(0.016)
Int. rate (-3)	1.54	0.49	0.2	-0.09	1.02	-0.09	-0.05	0.23
	(0.469)	(0.101)	(0.451)	(0.094)	(0.609)	(0.102)	(0.61)	(0.11)
Growth (-4)	0.09	-0.01	0.12	-0.02	0.03	0	0.07	0
	(0.069)	(0.015)	(0.068)	(0.014)	(0.068)	(0.011)	(0.076)	(0.014)
Int. rate (-4)	-0.19	-0.21	-0.05	0.06	-0.23	0.08	0.5	-0.07
	(0.309)	(0.067)	(0.315)	(0.066)	(0.397)	(0.066)	(0.394)	(0.071)
<i>Innov. covar.</i>								
Growth	10.83		19.1		23.04		15.55	
	(1.024)		(1.806)		(2.177)		(1.594)	
Int. rate	0.53	0.5	0.43	0.83	0.49	0.64	0.26	0.5
	(0.16)	(0.048)	(0.268)	(0.078)	(0.26)	(0.061)	(0.203)	(0.051)
Log likelihood	880		1012.3		1004.3		800.9	
Observations	243		243		243		210	

Table 4: 4-lag VAR for sample of countries. Annualized quarterly data 1956Q1-2016Q4 (from 1960Q1 for Canada). Robust likelihood-based standard errors in parentheses.

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: 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 the largest interest-growth differential still consistent with the data.

Figure 5 represents this exercise graphically for the USA. Each panel of Figure 5 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 estimation reported in Table 4; the point estimates are by construction, impossible to reject against the alternative of themselves. 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 5a, this crosses the 95% critical values for the χ^2_1 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 “soak up” 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 5 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

(a) Annual data 1880-2015

(b) Quarterly data 1956-2016

Figure 5: 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.

level. In sufficiently 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, as here, the conditional statistics can be quite misleading; θ_0 can be very large and positive (larger than the largest observed value even) but not rejected. In Figure 5a (and even more so in Figures 6a and 6b) 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 6 and 7 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)²⁶.

4.4 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%

26. Table 7 (in Appendix B.1) 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 6: Test statistics for hypothesis test of long run means. Annual data 1880-2015.

(a) United Kingdom

(b) France

(c) Canada

Figure 7: Test statistics for hypothesis test of long run means. Annualized quarterly data 1956Q1-2016Q4 (from 1960Q1 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 4), 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.²⁷

Table 5 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 data²⁸. 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

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

28. Appendix B.3 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.

data admits a moving average representation. The second lines estimates the degree of partial integration of the time series by maximum likelihood and by Bayesian estimation²⁹ respectively. The third line, labeled “Bayes superset” expands the Bayesian estimates to guarantee that the 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 5: 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 5 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 5 and 7, 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

29. Assuming a flat prior over the degree of partial integration in $[-1, 1]$

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 fourth line of Table 5 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 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.

To conclude this section, Figure 8 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 8 is clear. Point estimates of the long-run interest-growth differential are negative. This is robust across countries, periods, and estimation methods. This represents a very serious challenge to models of debt sustainability; if true it means that debt limits are not finite. However, upper bounds of confidence sets for this average are 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 B.2 shows that these basic findings are also robust to using alternative interest rate measures. So conservative estimates of sustainable debt levels should a) feature long-run differentials that are somewhere in this range, and b) explain clearly the sensitivity of the results to the assumed long-run differential. This is the exercise that we pursue in the next section.

5 Estimated debt limits

I now use the statistical processes for interest and growth differentials estimated in section 4 in the model developed in section 3. I focus on the UK 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 UK does have a large share of index-linked bonds relative to other countries, they are a recent phenomenon by historical standards (first issued in 1981) and even now only constitute around 25% of outstanding liabilities.

Figure 8: Comparison of point estimates and upper boundaries for interest-growth differentials in five advanced economies. Data annual 1880-2015, except where otherwise noted.

5.1 Estimating the surplus function

The average surplus function $s(x', d)$ is a critical determinant of realized debt limits. Any empirically plausible average surplus function must satisfy four properties: 1) so that average debt levels are positive, average balances must be negative when debt is low; 2) so that debt limits are positive, the government must be about to raise surpluses high enough to pay the (average) risk-free interest on the debt for at least *some* level of debt; 3) so that debt limits are finite, surpluses must be bounded; and 4) the correlation of the surplus ratio with the state must realistically reflect the cyclical nature of surpluses. To meet these four requirements, I parameterize the surplus function as a simple bounded quadratic function of debt:

$$s(x', d) = \alpha + \beta(d - \gamma)^2 I_{d < \gamma} + \delta(R(x') - 1) + \eta(G(x') - 1) \quad (10)$$

Where $\alpha, \beta, \gamma, \delta, \eta$ are constants. When $\alpha, \gamma > 0$ and $\beta < -\alpha/\gamma^2$, this satisfies all four requirements stipulated above. It also embodies the concept of “fiscal fatigue” that Ghosh et al. (2013) discuss - this is the notion that governments can raise surpluses in order to pay large debts, but that their ability to do so is eventually limited³⁰. In this case, the maximum expected surplus that the government can raise, even when debt is high, is α . Figure 9 illustrates the shape of this function when $R = G = 1$.

The relationship described in equation (10) is not meant to be a causal description of the government’s response to debt, interest and growth. Rather, it merely captures the usual correlation of surpluses with the other variables in the model. As such, we can simply fit the functional form to the data, without any other controls³¹.

Table 6 shows the parameters for this surplus function, estimated using generalized least squares on UK data from 1880 to 2015 (data is taken from Mauro et al. (2015)). Several points stand out from the estimation. First, that for high enough debt levels, surpluses are positive (as $\alpha > 0$). Second, that because $\beta < 0$, then surpluses are low when debt is low. Third, that surpluses are positively conditionally correlated with nominal rates (as $\delta > 0$), and negatively with growth rates (as $\eta < 0$). So the fiscal policy rule used here will (all else equal) produce higher surplus ratios in response to a decline in growth, i.e. fiscal policy is counter-cyclical. In spite of this, the sum $\delta + \eta$ is positive, so the surplus function responds positively to the interest-growth differential, over and above the dynamic response acting through movements in the debt stock.

30. 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.

31. Attempts to control for “omitted variables” throw away meaningful variation in surpluses that is correlated with the states of the model.

Figure 9: Stylized surplus function

This is an important determinant of the model-predicted debt limits. Because surpluses increase when debt dynamics deteriorate, this acts to mitigate the increase in the debt in states of the world when the interest-growth differential is high. Appendix C contains further discussion of the estimation of the surplus function.

Parameter	α	β	γ	δ	η
Estimate	3.36	-0.0067	63.6	0.99	-0.83

Table 6: Estimated surplus parameters for the UK, annual data 1880-2015

5.2 The dynamic behavior of debt limits

I now solve the model, and compute the associated estimated debt limits. The recovery rate on debt in default, ϕ , is set to 90%, a standard value for advanced economies (see Benjamin and Wright (2009)).

The discrete-state Markov process for interest and growth rates mimics a one-lag VAR with

the long-run interest-growth differential restricted to be at the 95% critical value for the VAR³². This means that the long-run differential is small and positive; around 0.16pp per year. The VAR for interest and growth rates is then discretized into 61 nodes states, producing 3721 possible levels of the realized interest-growth differential³³. Appendix D includes further details of the discretization process.

Figure 10 shows the impulse responses of the model. In blue are the responses to a one-standard-deviation increase in the interest rate innovation. Due to the high persistence of nominal interest rates, this is a long-duration shock, with a half-life of around five years. Because higher interest rates have little effect on future nominal growth rates (see the blue line in panel 10b). As a result, the interest-growth differential is substantially and persistently higher (panel 10c), causing a decline in the sustainable debt-GDP ratio of around 0.6pp on impact. So the model predicts an interest-elasticity of the maximum debt-GDP ratio of around -0.5. That is, for every 1pp increase in the nominal interest the debt limit falls by around 0.5% of GDP³⁴.

In red are the impulse responses for a one-standard-deviation shock to the innovation for nominal growth. While growth is much more volatile (reflected in the fact that the innovation is much larger), it is much less persistent, with a half-life of around two years. Higher growth induces a small interest rate response in the next period. This is consistent with policy responding sluggishly to an increase in nominal growth. Because interest rates are so persistent, this effect accumulates until interest growth differential actually rises in the long term (see the red line in panel 10a). As a result, the debt limit rises a little on impact in response to a positive growth shock, even though the initial interest-growth differential declines sharply. This occurs because the debt limit is forward looking (see equation (5)). And so debt issued today is sustainable only if it can be rolled over repeatedly in future. And so if higher growth is associated with higher future interest-growth differentials, positive growth shocks may permit only small increases in the amount of debt a government can sustain. As a result, the growth-elasticity of the debt limit is only around 0.3 - lower than for interest rates.

From a policy perspective, these results caution against expecting large variations in sustainable debt levels due to changes in interest-growth differentials, particularly when growth fluctuations are the source of the shock. As these are short-lived, the impact on sustainable

32. When conducting inference in section 4 we 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 61 nodes for interest and growth rates would mean solving the model for $61^4 \simeq 13,800,000$ values of x_t !

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 $61^2 = 3721$ possible interest-growth differentials.

34. Computed as the ratio of the initial changes in the debt limit and interest rate: $0.63/1.23 \simeq 0.5$

(a) Nominal interest rate

(b) Nominal growth rate

(c) Interest-growth differential

(d) Debt limit (percent of GDP)

Figure 10: Model impulse responses to one-standard-deviation shocks

debt levels is typically small, and the sign of the response is probably highly dependent on the associated interest rate response. However, interest rate movements do have a more noticeable impact on debt sustainability, as they are much more persistent.

5.3 Implied historical debt limits

We now use the model to produce estimated debt limit. Figure 11 shows in black the estimated debt limit for the model-implied annual debt limit for the same calibration as in the section 5.2, expressed as a share of GDP³⁵. This typically averages around 90-95% for the period in question, which seems a little on the low side. After all, the UK's public debt has exceeded this level during this period (most notably following world war II) without defaulting. This tells us something about the shortcomings of the model, either in its structure or parameterization. The factor that most likely affects this is the maturity of the debt. With multi-period debt, the government need repay only a fraction each period, reducing its gross financing needs.

For now, though, we defer discussion of the likely impact of including long-maturity debt to section 5.4, and start by highlighting key some results evident in Figure 11. First, that the variation in the debt limit is typically quite small. The entire range of $D(x)$ across all states is only around 13 percent of GDP. Even large and persistent changes, such as the decline in the interest-growth differential since 2008, have resulted in small changes in the debt limit over time, increasing only around 5pp during 2008-15.

In contrast, small changes to the long-run interest-growth differential are required to have similar effects on the level of the debt limit. The blue and green lines show two experiments which increase the long-run interest-growth differential by 0.5pp. In blue is shown the impact of shifting the assumed mean interest rate up by 0.5pp. And in green we see the effect of permanently lower growth rates, again by 0.5pp. These are very small changes relative both to the intertemporal variance of interest and growth rates over time, and relative to our degree of statistical uncertainty over the long-run interest-growth differential. For example, this is change is equivalent to moving from the 95% to 97.5% critical values for the VAR test, or moving about a quarter of the way to the 95% critical value of the Müller-Watson Bayes estimate. The estimated debt limits in these examples are around 5pp lower than in the baseline.

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.

35. The year-to-year fluctuations are rather noisy, so to make medium-term trends more apparent, this is shown as a five-year rolling average. The overall range of the series changes little, though.

Figure 11: Estimated debt limits for the UK, 1880-2015 (rolling five-year average)

Figure 12: Dependence of average debt limits on variance of interest-growth differentials

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.

The red line in Figure 11 performs one last experiment with this parameterization of the model. This line contemplates a small reduction in the degree of variation in interest and growth rates, reducing the standard deviation of both by 5%, while keeping cross- and auto-correlations the same. This causes debt limits to increase by around 5pp. This highlights a key mechanism of the model, and explains why the model-implied debt limits are low, at least relative to Ghosh et al. (2013). Fluctuations in interest and growth rates represent a source of risk to the government. If growth rates are lower than expected, the debt-GDP ratio rises. And if interest rates are higher than expected, then the cost of rolling over a given debt increases. This risk represents a cost to investors; unexpected movements in the interest-growth differential increase the probability of default. Investors require compensation for bearing this risk. This pushes down debt prices in all states of the world, and makes debt in general much less affordable. So in equilibrium, the maximum amount of debt a government can maintain must decline.

This channel is the key reason why the debt limits in Figure 10d are so much lower than those of Ghosh et al. (2013), who omit this mechanism. They assume that interest-growth differentials are constant, and generate debt limits for the UK in the range of 165-185% GDP. Figure 12, illustrates this point, plotting the the average model-implied debt limit produced by reducing the standard deviations of interest and growth rates yet further. This shows how the current model

can recover average debt limits in excess of 150% GDP by depressing the variance of interest and growth rates below empirically plausible levels. But this is not a terribly satisfactory way to produce more realistic-seeming debt limits. So instead, we return to the issue of debt maturity.

5.4 Debt maturity and shock volatility

It is possible to write down a version of the model solved here with long-maturity debt. However, the equilibrium conditions of such a model are very challenging to solve, and we do not attempt that here³⁶. So to attempt to approximate the extent to which the debt maturity might affect the level of the debt limit, Figure 13 presents the debt limits for alternate parameterizations of the model. In each case, we convert the annual stochastic processes for interest, growth rates, and surplus shocks into their equivalent averages over longer periods.

This is, of course, an imperfect representation of long-duration debt, as it assumes that the debt is still rolled over all-at-once, just at longer intervals. However, this will capture a critical aspect of longer-maturity debt; that by issuing longer debt, the relevant interest-growth differential is the average over the lifetime of the debt, effectively reducing the volatility of the interest-growth rate to which the government is exposed. This attempts to correct the major shortcoming of the model: that one-period debt means that the government is much more exposed to fluctuations in the interest-growth differential that is likely to be realistic. As we saw in the foregoing discussion on exactly this topic, the level of the debt limit is highly sensitive to this volatility.

The main quantitative result of Figure 13 is that increasing the period length of the model, the maximum sustainable debt limit increases back to levels that are significantly above current debt levels. The x-axis values shown in this figure are also quantitatively reasonable. Ellison and Scott (2017) show that before 1920 almost all UK debt was consols (and therefore of infinite maturity), but that since around 1960 the maturity of debt has been more evenly spread, with a median maturity since 1960 close to 8-10 years. Figure 14 shows the elasticity of the debt limit to shocks to the interest-growth differential, as a function of the period length. This represents the sensitivity of the debt limit to short-term fluctuations in the interest-growth differential.

Figure 14 preserves, and even strengthens, the central conclusion from the one-year calibration: that the debt limit does fall when interest-growth differentials rise, but not by much. As the period length increases, the magnitude of this response actually declines, as at longer frequencies, shocks are less persistent. So changes in the interest-growth differential in the current period have less

36. The reason is that the debt price in the current period depends on the likely resale price of the debt in the next period. As a result, one can no longer solve for the probability of default at the debt limit along, but instead have to solve for default probability at all debt levels simultaneously. Including this extension is an interesting topic for future research.

impact of debt sustainability going forward. In contrast, the impact of changes to the long-run interest-growth differential will be very similar, as the period length has no effect on the long-run properties of the model.

Figure 13: Dependence of average debt limits on period length

Figure 14: Elasticity of debt limit to interest-growth differential

We can summarize this section on estimated historic debt limits in six points. First, that debt limits change little as interest-growth differentials fluctuate, typically by only a few percentage points even in response to large shocks. Second, that similar movements in the average debt limit can be generated by changes to the long-run interest-growth differential which are much smaller, and also small relative to the statistical uncertainty over this parameter. Third, that the levels of the debt limits in an annual calibration are lower than can be easily reconciled with the debt levels seen in the data. Fourth, that past work which had generated higher debt limits relied crucially on suppressing variation in interest-growth differentials. Fifth, that the model can produce debt limits which are more plausible if instead the period of the model is lengthened to match empirically reasonable debt maturities. And sixth, that changing the period length of the model does not affect the first two points: that short-term fluctuations in debt limits are much smaller than our uncertainty over them.

6 Conclusion

The policy question which motivates this paper is simple: do declines in nominal interest rates following the global financial crisis mean that governments can safely borrow more? This question is not only interesting in its own right, but is important to understand in environment when interest rates are likely to rise once more in future.

The short answer to that question is: “yes, but probably only by a few percentage points”. But more interesting than this final answer are the other issues that we have been forced to address along the way.

We 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. We then extended the approach of Ghosh et al. (2013) to build a framework in which we can assess quantitatively how changes in this variable might impact the maximum debt a government can borrow. We showed that in such a model, maximum debt limits exist only if the interest-growth differential is positive on average in the long run.

So we set out to measure the long run interest-growth differential in a sample of advanced economies and found that while point estimates are indeed negative, a variety of statistical techniques cannot reject the possibility that this differential is small and positive. We 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 too.

Bearing this in mind, we calibrated the model to the historical data for the UK 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 very sensitive to the choice of the long-run differential, even when restricted to the set of those which are statistically plausible.

The level of the debt limits produced in this model were lower than could be easily reconciled with the data; around 90-95%. However, this was due entirely to the feature we wanted to capture: variability of interest-growth differentials. As such, this represented a challenge as much to past work ignoring this effect as to the model presented here.

As a way to address this shortcoming, though, we approximated the inclusion of a key missing feature from the model: long-maturity debt. We did so by extending the length of the period to something closer to the average debt maturity of the UK. This recovered average debt levels which seem more reasonable, in the order of 140% of GDP, but without undermining the earlier about parameter sensitivity.

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, solvency considerations alone cannot provide a motive for finite debt limits. Yet default risk is on-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. As such, the conclusions of this paper give added weight to work, such as Aguiar et al. (2017), that seek to examine the interaction of solvency with Cole and Kehoe (2000) style rollover crises.

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.
- 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.”
- 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.
- 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.

- Ellison, Martin, and Andrew Scott. 2017. “Managing the UK National Debt 1694-2017.”
- 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, Leonardo Martinez, and Francisco Roch. 2017. “Robust Fiscal Rules.” *Working paper*.
- 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, Noémie, 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

A.1 The motivating model

Proposition 1

1. If $r > g$ then there exists $\bar{b}$, the maximum sustainable debt. This is the largest b^* such that the government defaults whenever $b_{t-1} > b^*$. This is given by:

$$\bar{b} = \frac{(\bar{s} - \bar{u})(1 + g)(1 + r)}{r - g}$$

2. If $r < g$ then there exists no such $\bar{b}$. The government never defaults, no matter the debt level.

Proof.

1. Assume $r > g$. Then for any $\bar{s} - \bar{u} > 0$, there exists $b^* > 0$ such that risk-free interest payments can be met even with the lowest possible surplus shock. That is, b^* satisfies:

$$\frac{1}{1 + r}b^* + \bar{s} - \bar{u} = \frac{1}{1 + g}b^* \quad (11)$$

Then for any $b_{t-1} \leq b^*$, the government can choose an expected surplus $s_t^e = \bar{s}$ to guarantee that $b_t \leq b_{t-1}$ for any $u_t \in [-\bar{u}, \bar{u}]$ if the debt price is $1/(1+r)$. The debt is therefore bounded with probability one, so the government will repay for certain. As a result, investors will be willing to buy debt at $q_t = 1/(1 + r)$, verifying the assumed debt price.

But for any $b_{t-1} > b^*$, there is some risk that $b_t > b_{t-1}$, even if $q_t = 1/(1 + r)$, its maximum possible value. Therefore $b_t > b^*$ as well, and so the debt grows without bound. The government therefore must default for sure no matter the debt price. And so the only debt price satisfying the investors' pricing requirements is $q_t = 0$.

Thus b^* separates the debt level into two regions, $q_t > 0$ when $b_{t-1} \leq b^*$, and $q_t = 0$ when $b_{t-1} > b^*$. Rearranging equation (11) gives:

$$b^* = \frac{(\bar{s} - \bar{u})(1 + g)(1 + r)}{r - g}$$

2. Now let $r < g$. There still exists $b^* > 0$ as defined in equation (11). Then for any $b_{t-1} > b^*$, $b_t < b_{t-1}$ (because now $r < g$). So for any $s_t \geq \bar{s} - \bar{u}$, the government can meet its budget constraint in each period with a declining debt if $q_t = 1/(1+r)$. In this case, debt is therefore bounded with probability one and the government never defaults. And so $q_t = 1/(1+r)$ does indeed solve investors' debt pricing equation. As this holds for *any* $b_{t-1} > b^*$, then there exists no upper bound on the sustainable debt.

■

A.2 Debt limits with negative average interest-growth differentials

Proposition 3 *There exists an equilibrium with $D(x) = \infty \ \forall x$ if and only if $\mathbb{E} \log(G_t/R_t) > 1$.*

Proof. If $D(x) = \infty \ \forall x$, then conditions 1-4 of the definition of equilibrium are obviously satisfied with the probability of default equal to zero everywhere. The only outstanding issue is whether debt will diverge to infinity for sure when d_0 is finite.

Iterating forward on equation (2) gives:

$$d_T = \frac{d_0}{Q_T^1} - \sum_{r=1}^T \frac{s_r}{Q_r^t} \quad (12)$$

$$\text{Where: } Q_1^T = \prod_{s=1}^T \frac{G_s q_{s-1}}{(1 - p_s)}$$

Where $p_t = p^*(x_t, d_t, D)$. In an equilibrium with $D(x) = \infty \ \forall x$, we must have that $p_t = 0$, $q_t = 1/R_t$ for all finite d_t . And so:

$$Q_1^T = \prod_{s=1}^T \frac{G_s}{R_{s-1}}$$

Then:

$$\begin{aligned} \frac{1}{T} \log Q_1^T &= \frac{1}{T} \sum_{s=1}^T (\log G_s - \log R_{s-1}) \\ &= \left(\frac{T-1}{T} \right) \frac{1}{T-1} \sum_{s=1}^{T-1} \log \frac{G_s}{R_s} + \frac{1}{T} (\log G_T - \log R_{t-1}) \\ &\rightarrow 1 \times \mathbb{E} \log(G_s/R_s) + 0 \quad \text{with probability 1 as } T \rightarrow \infty \end{aligned}$$

Where the last line follows from the ergodicity of the Markov process x_t . So by the continuous

mapping theorem:

$$Q_1^T \rightarrow (e^{\mathbb{E} \log(G_s/R_s)})^T \quad \text{with probability 1 as } T \rightarrow \infty$$

$$= \begin{cases} \infty & \text{if } \mathbb{E} \log(G_s/R_s) > 1 \\ 0 & \text{if } \mathbb{E} \log(G_s/R_s) < 1 \end{cases}$$

Case 1: $\mathbb{E} \log(G_s/R_s) > 1$: If the interest-growth differential is negative, then equation (12) becomes:

$$\mathbb{P} \left(d_T = - \sum_{r=1}^T \frac{s_r}{Q_r^t} \right) \rightarrow 1 \quad \text{as } T \rightarrow \infty$$

Then so long as deficits are not infinite (which holds because we assume that the surplus function is bounded), then the right-hand side is finite. So d_0 will never diverge to infinity, and this rule is an equilibrium.

Case 2: $\mathbb{E} \log(G_s/R_s) < 1$: If the interest-growth differential is positive, then equation 12 becomes unbounded with probability one unless $d_{t-1} \leq \sum_{r=t}^{\infty} Q_r^t s_r$. But then this boundary *is* a debt limit: debt higher than this will eventually become unbounded with probability one. This contradicts the assumption that $D(x) = \infty \forall x$. Therefore there is no equilibrium here.

■

A.3 Tangency condition

Proposition 4 *At an equilibrium, a tangency condition holds:*

$$\frac{\partial}{\partial p} Z \left(x, D(x), p^*(x, D(x)), \hat{D} \right) = 1$$

Proof. Assume, for a contradiction, that the condition fails, so for some x :

$$\frac{\partial}{\partial p} Z \left(x, D(x), p^*(x, D(x)), \hat{D} \right) \neq 1$$

Step 1 We first show that

$$\frac{\partial}{\partial p} Z \left(x, D(x), p^*(x, D(x)), \hat{D} \right) < 1$$

Imagine not. Then because $Z(\cdot)$ and continuous, there must be some range of p for $p < p^*(x, D(x))$ where $Z < p$. Let $\tilde{p}$ be such a point in this range. Because $Z(\cdot)$ is non-decreasing, the image of $Z(\cdot)$ restricted to $[0, \tilde{p}]$ is a subset of $[0, \tilde{p}]$. Therefore there must be a fixed point in $[0, \tilde{p}]$ (possibly

at $p = 0$). Denote this fixed point by $\hat{p}$. So $\hat{p} < \tilde{p} < p^*(x, D(x))$. But $p^*(x, D(x))$ is defined as the smallest fixed point of $Z(x, D(x), p, \hat{D})$, so

Step 2 We now show that there exists $\bar{p} \in (p^*(x, D(x)), 1)$ such that $Z(x, D(x), p, \hat{D}) < p \ \forall p \in (p^*(x, D(x)), \bar{p})$.

This follows because $\frac{\partial}{\partial p} Z(x, D(x), p, \hat{D}) \rightarrow 0$ as $p \rightarrow 1$. By continuity there must exist some $\tilde{p} < 1$ such that $Z(x, D(x), \tilde{p}, \hat{D}) > \tilde{p}$. By continuity again, there therefore exists at least one fixed point in $(p^*(x, D(x)), \tilde{p})$. Let $\bar{p}$ be the smallest such fixed point. Then $\bar{p}$ satisfies the claim of this step.

Step 3 We now complete the proof. Clearly, $\frac{\partial}{\partial d} Z(x, d, p, \hat{D}) > 0$; higher debt means a higher probability of crossing the default threshold in all states. Then for $d = (1 + \delta)D(x)$ with $\delta > 0$, the line $Z(x, D(x), p, \hat{D})$ increases everywhere, including for $p \in (p^*(x, D(x)), 1)$. $Z(x, D(x), p^*(x, D(x)), \hat{D}) > p^*(x, D(x))$. But for small enough δ , there exists $\underline{p} \in [p^*(x, D(x)), \bar{p}]$ such that $Z(x, D(x), \underline{p}, \hat{D}) < \underline{p}$. So by continuity, there exists a fixed point $p^{**} \in (p^*(x, D(x)), \underline{p})$. Because $\bar{p} < 1$, then $p^{**} < 1$ as well. So for $d = (1 + \delta)D(x) > D(x)$, there exists a fixed point for the default probability which is strictly less than one. In other words, $D(x)$ is not the largest debt level with a non-zero probability of repayment, i.e. not a debt limit. So by contradiction, the proposition holds.

■

B Further empirical work

B.1 Restricted VAR estimation

Table 7 shows the VAR estimates restricted such that the mean difference is at the 5% critical value for each country. Comparing to Table 4 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.

B.2 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).

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

be at the 5% critical value under the unconditional LR test. Annual data 1880-2015. Robust Table 7: Multi-lag restricted VAR for sample of countries. Lon-run means are restricted to likelihood-based standard errors in parentheses

Table 8 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 periods where the risk-free rate has been surprisingly high (such as 1980-2000).

	<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 8: Average spread over risk-free rate

Figure 15 shows the results from the 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.

B.3 Application of the spectral method to quarterly data

Table 9 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 4). 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 9. The USA and UK have an estimated degree of integration of unity (at least,

(a) VAR estimates

(b) I(0) spectral estimates

Figure 15: Long-run interest-growth differential using alternative interest rates. Annual data, 1880-2015

the estimation process does not consider higher-order integration). Indeed, only Germany has a stationary representation during this period³⁷. As a result, the confidence sets grow without 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 stationarity³⁸, 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 9: 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).

C Estimating the surplus function

Figure 16 shows the two estimated surplus functions. In black, a fifth-order polynomial, in blue the bounded quadratic functional form featured in the text. Both cases also include linear terms in the interest and growth rates, which are held at their sample means in the curves shown in Figure 16. Note how the qualitative form of the high-order polynomial - sharply increasing for low lagged debt levels, broadly flat for high debt levels - except without the sharp divergence outside the range of estimates.

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

38. Except for Canada; see the fractional integration line of Table 5

Figure 16: Surplus function.

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 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 5 are shown in Figure 17, 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 to 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

Figure 17: Discretized values for nominal interest and growth rates with $D = 20, R = 3$. Points shown in proportion to their ergodic frequency.

typically cannot match the conditional moments without error using no-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 10 compares the estimates from a one-lag VAR based on the UK 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.

	Data		Simulation	
	Growth	Int. rate	Growth	Int. rate
<i>Coefficients</i>				
Constant	1.43	0.27	1.81	0.85
LR Mean	5.02	5.18	5.01	5.17
Growth (-1)	0.67	0.05	0.64	0.03
Int. rate (-1)	0.04	0.9	0	0.81
<i>Innov. covar.</i>				
Growth	22.05		24.02	
Int. rate	1.43	1.52	2.12	2.86

Table 10: Estimated restricted VAR based on data, and VAR estimated from discretized approximation simulated 100,000 periods

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 18 shows the fitted discretized series in comparison to the UK data. The discretization handles the volatility in growth 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.

Figure 18: Time series of nominal interest and growth rates: data in black, discretized process in blue.