BANK OF ENGLAND

Working Paper No. 484

GDP-linked bonds and sovereign default

David Barr, Oliver Bush and Alex Pienkowski

January 2014

BANK OF ENGLAND

Working Paper No. 484
GDP-linked bonds and sovereign default

David Barr,[1] Oliver Bush[2] and Alex Pienkowski[3]

Abstract

Using a calibrated model of endogenous sovereign default, we explore how GDP-linked bonds can raise the maximum sustainable debt level of a government, and substantially reduce the incidence of default. The model explores both the costs (in particular the GDP risk premium) and the benefits of issuing GDP-linked bonds. It concludes that significant welfare gains can be achieved by indexing debt to GDP.

Key words: Fiscal policy, contingent pricing, debt management, sovereign debt, sovereign default.

JEL classification: E95, G16, H63.

(1) Bank of England. Email: david.barr@bankofengland.co.uk
(2) Bank of England. Email: oliver.bush@bankofengland.co.uk
(3) Bank of England.

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or the Financial Policy Committee. We thank participants and discussant Enrique Kawamura at the IEA roundtable on sovereign debt crises, participants at seminars at Universidad Nacional de Córdoba and the University of Warwick, Martin Brooke, Andy Haldane, Glenn Hoggarth and Rhiannon Sowerbutts for helpful comments. This paper was finalised on 13 December 2013.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Summary

This paper investigates the conditions under which GDP-linked bonds help to protect governments (or 'sovereigns') from unexpected poor growth outturns that might otherwise push them into a debt crisis. This is because the return on these bonds varies in proportion to the country's GDP – when growth is weak, the debt servicing cost and repayment amount automatically declines; and when growth is strong, the return on the bond increases. This helps to stabilise a sovereign's debt to GDP ratio and makes it less likely that a deep recession will trigger a debt crisis and cause a default. GDP-linked bonds, therefore, can be viewed as a form of 'recession insurance' for sovereigns. While all countries might experience some benefit from the use of GDP-linked debt, economies with higher GDP growth volatility (such as emerging market economies) or countries where monetary policy is constrained (such as those in a monetary union) are likely to benefit most.

We use a calibrated model of sovereign default based on work by International Monetary Fund authors, which delivers a calculation of the maximum level of debt that a sovereign is likely to be able to sustain before it risks facing a crisis. This model is estimated for a 'representative' sovereign in two scenarios – i) when all debt is issued as conventional bonds; and, ii) when all debt is in the form of GDP-linked bonds.

Given the simplicity of the model, these debt thresholds should not be interpreted as hard limits. In fact, historical experience suggests that many countries can exceed these levels without facing repayment problems. Instead, the focus of this paper is to consider how GDP-linked bonds can help to reduce the risks to a sovereign within this simple framework. This implies more attention should be focused on the amount GDP-linked bonds can potentially increase debt limits, rather than the absolute value of the debt limits themselves. Under the simplest model set-up we find that GDP-linked bonds have a substantial impact on a sovereign's debt limit – raising it by around 100% of GDP.

This analysis abstracts from a number of important considerations, so the baseline model is then made more realistic with two innovations. First, investors are now assumed to be risk-averse and require an additional premium to hold risky assets. This means that when the return on the asset is uncertain – either due to a risk of payment default, or in the case of GDP-linked bonds, because future growth outturns are uncertain – investors will charge a higher interest rate on debt. Second, it is assumed that when a sovereign changes its fiscal policy stance in order to try and stabilise debt, this has an impact on growth. For instance, when a sovereign increases its primary balance this will drag down on GDP growth. When these two modifications are included in the model, the additional 'fiscal space' derived from the introduction of GDP-linked bonds is around 45% of GDP.

The final section of analysis considers the welfare implications of issuing GDP-linked bonds. Sovereign defaults have the potential to damage the domestic economy significantly, so reducing the incidence of this will improve welfare. A stable and predictable fiscal policy is also desirable, as taxpayers are not faced with unexpected and erratic changes in tax. GDP-linked bonds help both to reduce the incidence of sovereign default and to stabilise fiscal

policy. But on average taxpayers will have to pay higher interest payments on GDP-linked bonds (at least at low and moderate debt levels) compared to conventional bonds, which will lower taxpayer welfare. On balance, however, we conclude that GDP-linked bonds may provide a substantial net benefit in welfare terms – in our calibration this is equivalent to consumption equal to between 1% and 9% of GDP in perpetuity.

In summary, GDP-linked bonds have the potential to reduce the incidence of costly sovereign default and allow fiscal policy to be more stable and predictable. The welfare gains from this outweigh any additional costs associated with issuing such debt, especially for sovereigns with volatile GDP. GDP-linked bonds also have the potential to improve the functioning of the international monetary and financial system, by encouraging greater country self-insurance, and reducing the reliance on large-scale official sector support programmes to resolve crises.

SECTION 1: INTRODUCTION

In this paper we explore the ways in which GDP-linked bonds can stabilise sovereign debt dynamics and reduce the probability of default. GDP-linked bonds provide cash payments that vary positively with the level of GDP, thereby helping to stabilise the debt-to-GDP ratio.

GDP-linked bonds belong to a wider class of state-contingent assets. In the Arrow-Debreu framework, a complete set of state-contingent assets allows agents to achieve a Pareto efficient allocation of risk. Since real world financial markets are far from 'complete', increases in the number of state-contingent assets have the potential to increase welfare.

To a certain extent, payments on conventional sovereign bonds are state-contingent, since the sovereign always has the option of reneging on its commitments. Therefore losses in adverse states of the world can be passed on to investors, which is why sovereign bonds carry a credit spread.

Loss-sharing through reneging on debt is inefficient for three reasons however. First, the degree of 'risk-sharing' is typically confined to tail-risk events. This means that there are probably many states of the world in which it would be efficient to share sovereign risk, but these opportunities are not realised.

Second, there can be significant deadweight costs to both debtors and creditors associated with sovereign debt crises and default. For example, De Paoli, Hoggarth and Saporta (2006) find that sovereign defaults are often associated with banking sector and currency crises, which can significantly exacerbate the cost of default. Across their sample of defaults, the mean output loss following a sovereign default is 15 per cent of GDP. Further, Benjamin and Wright (2009) show that, on average, debtor nations exit their period of default with an increased debt-to-GDP ratio, despite having imposed haircuts on creditors. Levy Yeyati and Panizza (2011) show that large output contractions often occur prior to default, suggesting either that there are costs associated with the anticipation of default or that defaults do not generate significant deadweight costs.

Third, there is a significant amount of uncertainty associated with sovereign default. Debt restructuring negotiations are often protracted, and losses upon default can vary significantly in size. Much of this uncertainty occurs because there is no pre-defined framework or convention for negotiating sovereign debt write-downs. And this uncertainty can increase the cost of sovereign borrowing. GDP-linked bonds can mitigate each of these inefficiencies.

GDP-linked bonds also have the potential to improve the functioning of the international monetary and financial system (IMFS), by encouraging greater country self-insurance. At present, the predominant policy response to sovereign debt crises is large scale official sector liquidity support from the International Monetary Fund and other official sector financing arrangements. While there are strong arguments for providing this public insurance, there are also a number of potential costs. First, this support can promote both creditor and debtor moral hazard, and incentivise short-term lending (as these creditors have a higher likelihood of being repaid by the official sector). Second, large official sector loans can in principle jeopardise tax-payer resources if they are not repaid in full. Third, negotiating write-downs with private sector creditors to achieve sustainable sovereign debt levels can become harder when a significant proportion of the sovereign's debt is held by the official sector, which has de-facto senior creditor status (Brooke, Mendes, Pienkowski and Santor, 2013).

By promoting greater country self-insurance, through the use of state-contingent assets including GDP-linked bonds, these costs may be significantly reduced. Indeed they could potentially improve the overall functioning of IMFS by both reducing the incidence and the impact of sovereign debt crises.

An alternative policy response might be to pursue a statutory mechanism to improve the orderliness and predictability of sovereign debt restructurings. The most high profile example of this is the Sovereign Debt Restructuring Mechanism (SDRM) proposed a decade ago (Krueger, 2002). No consensus was reached on this by the international community, and collective action clauses (CACs) were instead promoted. This suggests that contractual, market-based, responses to inefficiencies in the sovereign debt restructuring framework may be preferable to international policymakers.

The primary focus of the paper is to sketch out a framework which allows for endogenous sovereign default and can capture the benefits of GDP-linked bonds in lowering this probability. In our model, we find that GDP-linked bonds raise the maximum sustainable debt-to-GDP ratio considerably and that both the probability of default and fiscal policy volatility fall. Furthermore, in our simple experiment, the welfare gains of switching to GDP-linked bonds are significant, especially if there are large costs of default.

The rest of the paper is organised as follows. Section 2 summarises some relevant literature on GDP-linked bonds and sovereign debt crises. Section 3 sets up the basic framework of our model. Section 4 introduces GDP-linked bonds and explores how they can reduce the incidence of sovereign default under the basic set-up. Sections 5 and 6 extend the model by introducing risk-averse investors and feedback from fiscal consolidation to growth. Section 7 explores the welfare implications of GDP-linked bonds. Section 8 concludes.

SECTION 2: REVIEW OF THE LITERATURE

Following the debt crises of the 1980s, several papers (see for example, Krugman (1988) and Froot, Scharfstein and Stein (1989)) suggested that sovereigns should link debt cash-flows to macroeconomic variables such as commodity prices, exports and real GDP. This would promote greater risk-sharing with creditors and reduce the incidence of costly sovereign defaults. In general, these early papers favoured indexing against variables that are out of the control of governments, such as global commodity prices or world interest rates, to avoid moral hazard complications.

The first economist to focus on the merits of GDP-linked bonds was Robert Shiller. Shiller (1993, 2003) proposed introducing 'Trills' – perpetual claims on one-trillionth of a country's GDP – that would allow households and companies to take an 'equity stake' in a country's economic fortunes, and would, therefore, support risk diversification and hedging. Barro (1995) argued that an optimal debt management strategy from the sovereign's perspective would be to issue sovereign bonds linked to government expenditure in order to smooth taxes through time. However, given the clear moral hazard risks associated with this strategy, the next best alternative is to link bonds to GDP. Despite the fact that GDP can be influenced by the government, the incentive to lower output in order to reduce debt payments is much smaller.

Borensztein and Mauro (2004) identify four major benefits to the sovereign from issuing GDP-linked bonds. First, during an economic downturn the likelihood of sovereign default, with all the associated spillover costs to other sectors and countries, will decline. Second, the sovereign is less likely to be forced to undertake damaging pro-cyclical fiscal policies. Third, the sovereign is able to implement a smoother intertemporal tax path, which reduces uncertainty over consumption and investment decisions. Fourth, the higher interest payments associated with higher growth can provide a natural break on unsustainable business cycle booms, reducing the likelihood that governments will over-spend.

Borensztein and Mauro (2004) estimate the premium required to compensate investors for taking on GDP volatility risk using a simple CAPM. They show that the premium required by an investor with a broad portfolio of US equities to hold an Argentine GDP-linked bond would only be around 100bps above the yield on Argentine vanilla bonds. Similar calculations made by Kamstra and Shiller (2009) suggest that US Trills would have a yield of 150bps over short-term US government debt.

Chamon and Mauro (2005) introduce default risk into their pricing model. They develop a Monte Carlo framework that simulates the evolution of the debt-to-GDP ratio under vanilla and GDP-linked bonds. Simulations are calibrated against historical data for a number of emerging markets. The market price of vanilla bonds is used to predict the ratio at which default occurs. The relative frequency of defaults under vanilla and GDP-linked bonds is used to price default risk. They find that GDP-linked bonds reduce the probability of default and therefore reduce the required return on the remaining conventional bonds.

Ruban, Poon and Vonatsos (2008) develop this approach further by using a structural model of default based on a sovereign's ability to pay. Default occurs in this model when a sovereign's 'assets' fall below a fixed level of *dollar-denominated* debt due to exogenous shocks. The dollar value of these assets is defined as a function of the country's potential output denominated in *local currency*. It is, therefore subject to growth and exchange rate shocks. They calibrate their model using data for Brazil, Mexico and Argentina, and test a number of contract designs. They find that GDP-linked bonds generate a significant reduction in the default probability only if the debt contract is linked to the *US dollar* value of a country's GDP.

Our paper differs from this approach in two ways. First, we do not complicate the model by using foreign currency denominated debt. This allows us to focus solely on the risk-sharing benefits of GDP-linked bonds. Second, we develop a more sophisticated model of sovereign default, which calculates a sovereign's 'debt limit' endogenously, and allows liabilities to change through time. The origins of this model are discussed next.

In principle, assessing the solvency of a sovereign is similar to that of any private institution. An assessment of current and future assets and liabilities is made in an environment where both sides of the balance sheet are subject to uncertain shocks. While the liabilities side of a sovereign's balance sheet is similar to that of a corporate or bank, the asset side is particularly hard to estimate. This is because much of the sovereign's 'assets' are a function of expected primary balances, which are very hard to predict.

How can one assess the potential size of future primary balances? Bohn (1998, 2005) argues that if it can be shown that the primary balance reacts in a positive way to debt, this provides evidence that a sovereign is solvent. He demonstrates that when the business cycle and temporary expenditure shocks (e.g. wars) are controlled for, the US primary balance responds positively to debt. This, he argues, provides evidence that US sovereign debt does not follow a random walk but reverts to some steady-state level.

Abiad and Ostry (2005) use a similar methodology to enhance the IMF's debt sustainability analysis for a wide panel of emerging markets. In addition to the level of debt, they consider a range of factors that might also influence the potential size of the primary balance, such as commodity prices and the quality of institutions. Developing this approach further, Mendoza and Ostry (2007) find a non-linear relationship between the primary balance and debt, and that, at high debt levels, the responsiveness of the primary balance to increases in debt declines, i.e. the sovereign becomes less 'fiscally responsible'. This makes intuitive sense, as the primary balance cannot grow indefinitely to stabilise debt without the country suffering from some kind of 'fiscal fatigue'. This is likely to result from political constraints from undertaking severe austerity measures.

This concept of 'fiscal fatigue' is formalised in a model by Ghosh, Kim, Mendoza, Ostry and Qureshi (2011) and built upon by Miller and Zhang (2013). We use this model as the foundation of our analysis of GDP-linked bonds. Ghosh *et al* model the determination of a country's debt limit and estimate this for a number of advanced economies. A primary balance reaction function is estimated as a cubic function of debt, country fixed effects and a variety of other factors. Crucially the cubic function allows for fiscal fatigue, as the primary balance no longer responds positively, and may even decline, as debt levels increase. Consequently there may come a point at which the primary balance is no longer large enough to counteract the upward pressure on debt from the interest rate – growth differential and debt becomes explosive. This model is presented in more detail in Section three.

This approach to modelling sovereign debt dynamics does not consider the *strategic* motive for default; it focuses on the question of *when* and under *what conditions* default occurs rather than on *why* it occurs. The literature associated with the motives for sovereign default, beginning with Eaton and Gersovitz (1981), is older and deeper than this empirical approach and would provide an alternative way to assess GDP-linked bonds.

Sandleris, Sapriza and Taddei (2008) develop a dynamic stochastic general equilibrium (DSGE) model, based on Aguiar and Gopinath (2006) and Arellano (2008), with the addition of GDP-linked bonds. Calibrated for Argentina, Sandleris *et al* show that GDP-linked bonds can improve welfare by an amount equivalent to 0.5 per cent of average aggregate consumption, and that sustainable debt levels increase significantly. However, the model is complex and the conclusions are sensitive to changes in parameter values. It is difficult, therefore, to extract generalised results for a wider set of countries. In our paper, we aim to complement this more rigorous technique by providing more generalised and intuitive results.

SECTION 3: MODEL FOUNDATIONS

Our model is based on the paper by Ghosh, Kim, Mendoza, Ostry and Qureshi (2011). By combining the concept of 'fiscal fatigue' with a standard debt dynamics equation, this model provides estimates of the debt limit of a number of advanced economies. This debt limit is then used to model sovereign default, with both conventional and GDP-linked bonds.

Sovereign debt dynamics

The evolution of the sovereign's debt stock is governed by the standard government budget constraint:

$$d_{t+1} - d_t = \frac{(r_t - g_{t+1})d_t}{(1+g_{t+1})} - s_{t+1} + \varepsilon_{t+1} \qquad (1)$$

where d_t is the debt level as a proportion of GDP; r_t is the real interest rate on debt in period t, due in period t+1; g_{t+1} is the real GDP growth rate from period t to period t+1; s_{t+1} is the primary balance as a proportion of GDP in period t+1; and ε_{t+1} is a shock to the debt-to-GDP ratio in period t+1. All debt has a maturity of one period, so r_t applies to the entire debt stock in each period. It is important to note that our model does not allow for nominal shocks.

The endogenous interest rate, r_t, is made up of a risk-free rate, r^*, which is exogenous and constant, and a credit spread that compensates investors for expected losses due to sovereign default. Initially, creditors are assumed to be risk-neutral, so require no compensation for the uncertainty associated with future states of the world. The market interest rate satisfies the following equilibrium condition:

$$(1 - p_{t+1}).(1 + r_t) + p_{t+1}.\theta.(1 + r^*) = (1 + r^*) \qquad (2)$$

where p_{t+1} is the probability of default in the next period and $\theta(1 + r^*)$ is the recovery value on a bond in the event of default. If the probability of default is zero, i.e. $p_{t+1} = 0$, the interest rate will equal the risk-free rate, $r_t = r^*$. Re-arranging yields the following expression for the risky interest rate on government debt:

$$(1 + r_t) = (1 + r^*).\frac{(1 - p_{t+1}.\theta)}{1 - p_{t+1}} \qquad (3)$$

Unlike Ghosh *et al*, we assume GDP growth to have a stochastic element around a steady-state value:

$$g_t = g^* + u_t \qquad u_t \sim iid \qquad (4)$$

where g^* is trend real GDP growth and u_t is a stochastic shock. In Section 6, we endogenise growth to react to changes in fiscal policy.

The 'ε_t' term in (1) represents shocks to the debt level that are not caused by the primary balance or interest payments (let's call them 'debt shocks'). For example, this shock could capture circumstances when a sovereign unexpectedly takes on previously off balance sheet (contingent) liabilities. Evidence (from Campos, Jaimovich and Panizza (2006) and Abbas, Belhocine, El Ganainy and Horton (2011)) suggests that historically these shocks have been an important determinant of a sovereign's debt dynamics. We assume that this shock is independent of the growth shock, u_t.

Ghosh *et al* hypothesise that, up to a certain debt-to-GDP ratio, the sovereign reacts to increases in debt by raising the primary balance in order to stabilise the debt-GDP ratio. Beyond this point, fiscal fatigue sets in, perhaps because of political resistance or because of institutional constraints. In support for this model, find that this fiscal reaction function is cubic.

The primary balance fiscal reaction function, s_{t+1}, in our model is slightly different –

$$s_{t+1} = \min(\alpha + \beta(r_t - g^*)d_t, \gamma) \quad (5)$$

Here the primary balance can react to *all* non-stochastic factors that affect debt dynamics – the debt level as a proportion of GDP (d_t), the interest rate (r_t) and expected/trend GDP growth (g^*) – rather than allowing the government to react only to changes in debt. This reaction function suggests that the relationship between debt and the primary balance will curve upwards over some range as the sovereign reacts both to a higher debt-to-GDP ratio and to the higher spread associated with it. This is perhaps a more realistic characterisation of how a government may respond to changes in its debt dynamics and the upward curve is also consistent with other studies of the relationship (e.g. IMF (2003)).

There are no exogenous shocks to the primary balance, so the government can always achieve its desired level. There are also no 'automatic stabilisers' so the primary balance does not react directly to shocks to growth.

At some point, however, a maximum primary balance to GDP ratio is reached. This limit, γ, is the maximum primary balance that a sovereign can achieve before 'fiscal fatigue' prevents further consolidation. The 'α' intercept in (5) captures other factors that might influence the size of the primary balance, such as the age dependency ratio or the degree of political stability, and is a constant in our model. This term should be negative to yield a positive steady-state debt-to-GDP ratio. The β coefficient determines the vigour with which the sovereign reacts to any shocks which push debt-to-GDP away from its steady-state.

Combining equations (1), (3), (4) and (5) gives a fuller picture of how government debt evolves:

$$d_{t+1} - d_t = \frac{(r^* - g^* - u_{t+1})d_t}{(1 + g^* + u_{t+1})} + (1 - \theta).(1 + r^*).d_t.\left[\frac{p_{t+1}}{1 + p_{t+1}}\right] - \min(\alpha + \beta(r_t - g^*)d_t, \gamma) + \varepsilon_{t+1} \quad (6)$$

The process is determined by: i) two stochastic elements, the shocks to GDP growth (u_t) and the debt shock (ε_t); ii) one behavioural component, the fiscal reaction function, $\min(\alpha + \beta(r_t - g^*)d_t, \gamma)$; and iii) the endogenous probability of default in the next period, p_{t+1}. The rest of the parameters are exogenous and constant.

To illustrate the basic structure of the model assume that there are no stochastic shocks ($u_t = 0, \varepsilon_t = 0$). Without this stochastic element, the probability of default must be zero, so p_{t+1} can also be set to zero ('default' will be defined more precisely in the next section). The debt evolution equation can now be broken into two components: the effect that the growth-interest rate differential has on debt, $(r^* - g^*).d_t$ and the primary balance, $\alpha + \beta\big((r^* - g^*)d_t\big)$.

These opposing factors can be illustrated using the phase diagram (Figure 1) with the debt level multiplied by the growth-interest differential on the horizontal axis, and the primary balance on the vertical axis. There are two steady-state values of debt, d^* and $\bar{d}$. If debt is below $\bar{d}$, then debt will move towards d^* where it will stabilise. If debt exceeds $\bar{d}$ it will grow at an explosive rate, as the primary balance implied by the fiscal reaction function will not be sufficient to counteract the effects from the growth-interest rate differential. We assume that markets will not finance an exploding debt-to-GDP ratio, so above $\bar{d}$, the sovereign is forced to default on its debt. Therefore, $\bar{d}$ is defined as a country's debt limit. One of the objectives of this paper is to consider how GDP-linked bonds affect this limit. First, however, the debt limit with the stochastic elements of equation (6) needs to be determined.

Deriving the debt limit

The re-introduction of the stochastic elements into (6) creates a risk that one or more negative shocks may force the sovereign to breach its debt limit, $\bar{d}$. This implies a non-zero probability of default, which means that creditors require a credit spread[1] as compensation for holding this sovereign debt. The debt limit is a function of the market interest rate and vice versa, so the debt limit and the probability of default are jointly determined.

The precise details of how the debt limit is derived can be found in Ghosh *et al*, but this section will summarise the broad methodology. The limit is defined as the highest level of d_t that the government can sustain at finite interest rates while satisfying its budget constraint and the interest rate equilibrium condition.

The probability of default is simply the probability that in the next period the debt ratio will breach the limit:

$$p_{t+1} = Pr\left[d_{t+1} > \bar{d}\right] \qquad (7)$$

This is equal to the probability that the two exogenous shocks ϵ_t and u_t are large enough to counter the combination of the other factors influencing debt, and push the sovereign above its debt limit. To solve for the debt limit, we search numerically for the highest d_t with a probability of default below 1, where $\bar{d}$ and p_{t+1} satisfy:

$$p_{t+1} = Pr\left[\frac{(r^* - g^* - u_{t+1})}{(1+g^* + u_{t+1})}\bar{d} + (1-\theta)(1+r^*)\bar{d}.\left[\frac{p_{t+1}}{1+p_{t+1}}\right] - \min\left(\alpha + \beta(r_t - g^*)\bar{d}, \gamma\right) + \varepsilon_{t+1} > 0\right] \qquad (8)$$

Calibrating the model

Table 1 summarises the key parameter values in the model. We have deviated from Ghosh *et al* on most, but not all, of the parameter values as well as on the functional form of the fiscal reaction function. First, we work with a representative country, rather than performing country specific experiments. This allows us to focus on the debt dynamics of a typical advanced economy, rather than considering country-specific circumstances.

Second, although we use the same difference between r^* and g^*, we calculate g^* using long-run data from Schularick and Taylor (2012), covering 11 advanced economies over the 1870-2008 period (we exclude the world wars).

[1] The credit spread is equal to $(1-\theta).(1+r^*).\left[\frac{p_{t+1}}{1+p_{t+1}}\right]$, which is embedded in equation 6.

Third, we assume different probability distributions for our shocks. Ghosh *et al* assume that 'spending shocks' have a triangular distribution with finite supports. This model takes a more agnostic view on the appropriate distribution of shocks. Haldane (2012) shows that over a long horizon, real GDP growth has very significant 'fat-tails', which are difficult to model using standard probability distribution functions. Instead, we use data from Schularick and Taylor (2012) to simulate shocks to GDP growth. Random draws from this data series are used in the model, using a 'bootstrap' technique[2]. While the variance of these shocks is relatively modest, the kurtosis and skew is large, as shown in Figure 2.

Table 1: Model calibration

r^*	0.034	Risk-free interest rate
g^*	0.021	Trend GDP growth rate
u_t	0.038	Standard deviation of GDP growth shocks
ε_t	0.030	Standard deviation of debt shocks
θ	0.7	Recovery rate
α	-0.03	Intercept on fiscal reaction function
β	3.5	Responsiveness of the primary balance to changes in $(r_t - g^*)d_t$
γ	0.1	Maximum primary balance as a proportion of GDP

A similar technique is used for the 'debt shocks', ε_t. This is calculated by finding the change in country debt over each year and then subtracting the fiscal balance, to find what Campos *et al* (2006) call the 'unexplained' component of the evolution of sovereign debt. We use OECD from 1950 to 2010 (where available) for the same sample of countries as Schularick and Taylor (2012). Again, a 'bootstrap' technique is used to draw random samples from this distribution for the model simulations.

Fourth, our recovery rate, θ, is 0.7, compared to 0.9 in Ghosh *et al*. This seemed to us somewhat more conservative and, given our wider shock distributions, we need a lower recovery rate to ensure that debt-to-GDP always falls below the debt limit following a default. This parameter is broadly consistent with the findings of Cruces and Trebesch

[2] Thousands of draws are randomly taken, with replacement, from the data set.

(2011). They estimate average net present value losses of around 35% (0.65 recovery rate), for a sample of 180 sovereign debt restructurings.

Fifth, given our altered fiscal reaction function, we need to find parameters α, β and γ. α we set at -0.03 so the steady-state debt-to-GDP ratio is positive. γ is set to 0.1 as it is very unusual for sovereigns to run primary balance to GDP ratios higher than this (once other variables are controlled for). Periods of primary balances averaging close to 0.1 have been seen, for instance in the United Kingdom after both world wars. Finally β is set so that the elasticity of a change in the primary balance to a change in debt to the GDP ratio is equal to 0.045 while the credit spread is zero. This is the elasticity suggested by IMF (2003).

Our preference would be to directly estimate this fiscal reaction function in a similar way to Ghosh *et al*, but data limitations prevented this. In particular, with the sample of advanced economies used (Schularick and Taylor (2012)) it is difficult to find reliable data on long-term interest rates and GDP growth expectations. Furthermore, we suspect that over parts of this period interest rates were not always determined by market forces, but instead subject to 'financial repression' (Reinhart and Sbrancia (2011)).

Model results

Given the simplicity of the model, the debt limits derived in this paper should be viewed as indicative. Attention should focus on the ability of GDP-linked bonds (explored in the next section) to increase the debt limit of a sovereign, rather than the specific values of the limit alone.

Solving (8) with the parameter values in Table 1 generates a debt limit of 93 per cent of GDP. Beyond this limit, we assume fiscal fatigue prevents the sovereign from increasing the primary balance enough to counteract the effect of the growth-interest rate differential. Debt will then become explosive, and the sovereign will not be able to issue new debt at any finite interest rate. Of course, this looks like a rather low estimate of advanced economies' capacity to bear debt, which no doubt reflects failings in our model, which is not surprising given how simple it is.

The debt dynamics below the debt limit are also of interest. The steady-state level of debt (d^*) is 78 per cent of GDP. Although at high debt levels the primary balance is larger than the growth-interest rate differential, which will act to reduce debt over time, debt could still increase if there were sufficiently large negative shocks to growth and/or the debt level. Thus there is a risk that debt may exceed the limit, leading to default. Investors recognise this risk and demand a credit spread, which further worsens the sovereign's debt dynamics. Under this set-up the sovereign will default around once every two hundred years (as we move to a richer framework this default rate will increase).

Sensitivity analysis

How do these parameter values affect the results of the model, and which are key determinants of the debt limit? Table 2 shows to the size of the change in each parameter value required to *increase* the debt limit by 10pp of GDP. The effect on the steady-state debt level and the incidence of default are also presented.

Table 2: Sensitivity of parameter values

Parameter	Baseline parameter value	Change in parameter	Effect on debt limit ($\bar{d}$)	Effect on steady-state debt level (d^*)	Effect on average default rate (per 100 years)
$r^* - g^*$	0.013	-0.08	0.1	0.14	-0.5
$\sigma(u_t)$	0.038	-0.008	0.1	0.07	-0.04
$\sigma(\varepsilon_t)$	0.03	-0.013	0.1	0.07	-0.05
θ	0.7	0.07	0.1	0.07	0
α	-0.03	0.01	No Change[3]	-0.19	-0.4
β	3.5	1	No Change[4]	-0.18	-0.4
γ	0.1	0.009	0.1	0.06	-0.1

The difference between the growth and interest rate ($r^* - g^*$) determines the model results – the actual growth and interest rates do not matter independently for the behaviour of the debt-to-GDP ratio. If this differential is reduced, the slope of the red line in Figure 1 declines, which increases the debt limit. The equilibrium debt level also increases and the net result is a large decline in the incidence of sovereign default.

The model results are highly sensitive to the variance of the growth shock. A reduction of less than 0.01 of the standard deviation of these shocks will increase the debt limit by 10pp

[3] See below for an explanation.

and the steady-state debt level by 7pp. This effect is almost twice as large as the effect from increasing the standard deviation of the debt shock by an equivalent amount.

Increasing the recovery rate also has a positive effect on the debt limit and steady-state debt level. The reason for this is that as debt approaches the debt limit, the credit spread required to compensate investors for expected losses declines (see (3)). The sovereign can therefore sustain higher probabilities of default without debt becoming explosive.

In this set-up, the debt limit lies beyond the point at which the primary balance has reached its maximum value of 10per cent of GDP. This means that a modest changes to the intercept (α) and slope (β) terms on the fiscal reaction function have no effect on the debt limit (consider Figure 1). However, in both cases an increase in these parameters values does decrease the steady-state debt level. And because this steady-state debt level is now further away from the unchanged debt limit, the incidence of defaults in our simulations decrease. Not surprisingly, if the maximum capacity of the primary balance (γ) increases, then the debt limit rises and the incidence of default declines.

In the next section GDP linked bonds are introduced to this set-up and we consider how these might affect the debt limit and debt dynamics of the sovereign.

Figure 1: Debt dynamic phase diagram

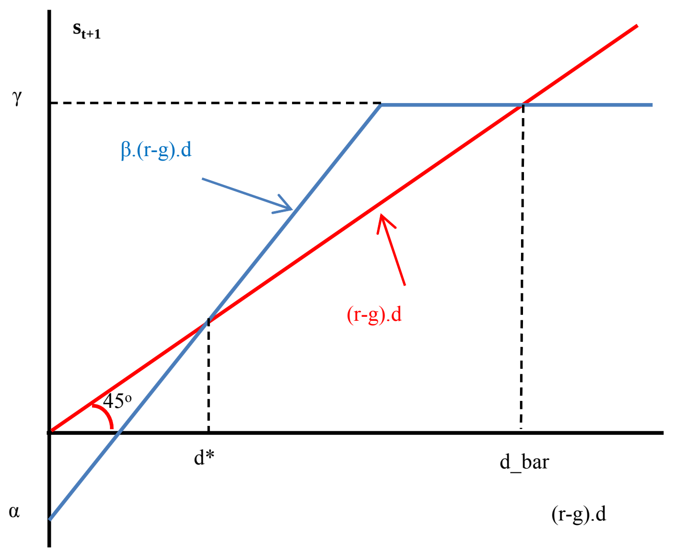

Figure 2: Real GDP growth outturns 1870-2008

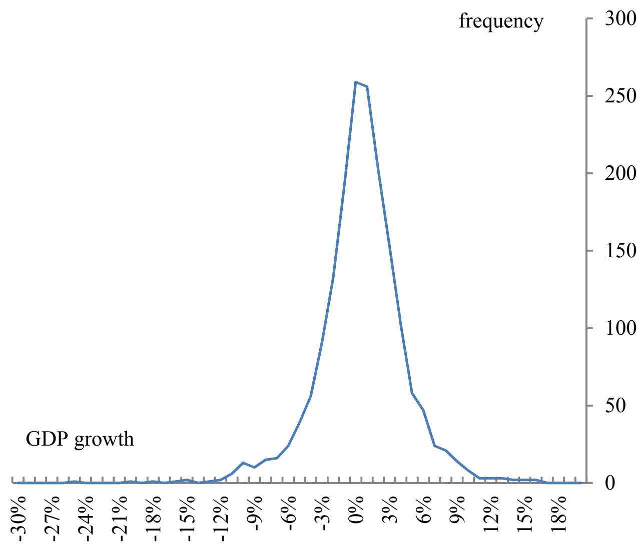

Source: Schularick and Taylor (2012)

SECTION 4: INTRODUCING GDP-LINKED BONDS

GDP-linked bonds in this model are similar to UK inflation-linked bonds. The debt's redemption value is linked to the level of GDP, which means that if a sovereign issues only GDP-linked bonds, its entire debt stock will adjust in proportion to GDP. The interest rate on the bond is defined as a fixed percentage of this principal, so it too adjusts with GDP. The redemption value of a single GDP-linked bond, b_t^{gdp}, will change through time as follows:

$$b_{t+1}^{gdp} = b_t^{gdp}.(1 + g_{t+1}) = b_t^{gdp}.(1 + g^* + u_{t+1}) \qquad (9)$$

Aggregating debt, the government budget constraint with GDP-linked bonds becomes:

$$d_{t+1}^{gdp} - d_t^{gdp} = \frac{r_t^{gdp}.d_t^{gdp}}{(1+g^*)} - s_{t+1} + \varepsilon_{t+1} \qquad (10)$$

Shocks to GDP growth no longer enter into this debt dynamics equation. While a negative shock to growth increases the debt-to-GDP ratio, the fall in the redemption value of GDP-linked bonds will reduce it by the same amount. Since this ratio is now immune to GDP shocks, GDP-linked bonds reduce the volatility of the sovereign's debt-to-GDP ratio, *ceteris paribus*.

In the absence of credit risk and risk aversion, the equilibrium interest rate on GDP-linked bonds is $r_t^{gdp} \approx r^* - g^*$ (both of these conditions will be relaxed in subsequent sections). This demonstrates that the difference between the long-run safe interest rate and trend growth is important for sovereign debt dynamics even when all debt is GDP linked. GDP-linked bonds protect a sovereign from unexpected temporary falls in growth relative to the interest rate, but do not protect it against permanent declines in trend growth. (However, if GDP-linked bonds had a longer maturity, then the sovereign would be better hedged against lower trend growth. Perpetual GDP-linked bonds (Shiller's Trills) would hedge the *entire* existing debt stock against negative shocks to trend growth.)

Adding GDP-linked bonds to the model

Using (10) instead of (1) to determine sovereign debt dynamics generates the following debt evolution equation:

$$d_{t+1} - d_t = \frac{(r^*-g^*)}{(1+g^*)}.d_t + (1-\theta).(1+r^*).d_t.\left[\frac{p_{t+1}}{1+p_{t+1}}\right] - \min\left(\alpha + \beta\left(r_t^{gdp} - g^*\right)d_t, \gamma\right) + \varepsilon_{t+1}$$
(11)

The major difference from (6) is the removal of the GDP growth shock, u_t. This reduces the probability of default at any given debt level which in turn lowers the credit spread demanded by investors. As a result, the required return on GDP-linked bonds is lower than for conventional bonds, which, from the logic above (and in Figure 1), implies an increase in the debt limit.

Model results

If the sovereign's debt stock is made up of GDP linked bonds only, the indicative debt limit rises to 195 per cent of GDP – a substantial rise of over 100 percentage points of GDP. Furthermore, the frequency of default is almost zero. On top of this, the steady-state debt-to-GDP ratio is higher with GDP-linked bonds (93 per cent compared to 78 per cent of GDP). This should not be a big surprise after Table 2: lowering the variance of the GDP shock resulted in exactly the same movements.

In a world of risk-neutral investors and costly defaults, GDP-linked bonds would be pretty attractive. In practice, however, one of the main challenges to introducing GDP-linked bonds is that investors will demand a risk premium to hold this debt which the sovereign may judge to be too expensive. In the next section, risk-averse investors are introduced.

SECTION 5: ADDING RISK-AVERSE INVESTORS

Adding risk-averse investors affects two parts of the model. First, creditors demand a premium for holding bonds with default risk. This increases the credit spread associated with any particular expected probability of default. Second, holders of GDP-linked bonds demand a risk premium for holding bonds that have an uncertain return due to GDP volatility. Each affect is explored in turn.

Default risk premia

In the previous set-up, risk-neutral investors derive the same utility from holding risk-free or risky bonds, as long as the expected return is equal. Therefore, rather than model utility directly, the expected returns on the bonds are set equal to each other to fulfil the equilibrium condition (2).

When investors are risk-averse, a utility function must be specified in order to derive the equilibrium condition under which the expected *utility* of holding each bond is equal. We use the familiar constant relative risk aversion (CRRA) utility function:

$$U = \frac{V^{1-\tau}-1}{1-\tau} \qquad (11)$$

where τ is the coefficient of relative risk aversion and V is the value of an asset. We assume that these agents have no other source of income or wealth aside from a portfolio of risk-free and GDP-linked bonds. The utility function leads to the following condition for the return on risky bonds (see appendix for full derivation):

$$(1+r_t) = (1+r^*).\left[\frac{1-p_{t+1}.\theta^{1-\tau}}{1-p_{t+1}}\right]^{\frac{1}{1-\tau}} \qquad (12)$$

Compare this with equation (3) for risk-neutral investors: for any risk-averse investor ($\tau > 0$), the credit spread will be higher for a given probability of default and recovery rate. This will increase the interest rate on both conventional and GDP-linked bonds, as both are subject to credit risk.

GDP risk premia

Ignoring default risk for now, the required return on a GDP-linked bond with risk-*neutral* investors is (see the appendix for derivation):

$$(1+r_t) = \frac{(1+r^*)}{(1+g^*)} \qquad (13)$$

This shows that as expected/trend GDP growth increases, investors demand a lower interest rate on GDP-linked bonds. This is because they expect to receive a capital gain as growth increases the value of their bonds. The risk-neutrality assumption means that investors are interested only in expected returns, and not the volatility of returns.

With risk-*averse* investors, the variance of the growth outturns does matter. Investors demand a risk premium to compensate them for this (see the appendix for derivation):

$$(1 + r_t) = \frac{(1+r^*)}{(1+g^*-\vartheta)} \quad (14)$$

where ϑ is the risk premium on GDP volatility. Comparing (13) and (14), it is clear that the required return on GDP-linked bonds with is higher risk-averse investors.

In a richer setting, the GDP risk premium would depend on the covariance of local GDP with the world market return. This could in principle be negative (and as Borensztein and Mauro (2004) show, this is the case for several countries). Were this the case, the GDP risk premium could be negative. But for an average small open economy, one might expect a covariance of around one and our assumption that there are no diversification benefits would be reasonable.

Model results

Using a coefficient of relative risk aversion of 4 (at the upper end of what is normally assumed in similar macroeconomic models), the model yields an indicative debt limit for conventional bonds of 63 per cent of GDP, significantly lower than the risk-neutral case (93 per cent of GDP). The introduction of risk-averse investors also substantially lowers the debt limit for GDP-linked bonds (from 195 per cent to 114 per cent of GDP).

Part of this reduction in debt limit is because creditors demand a higher credit spread for any given probability of default. This raises the interest rate on government debt, thus lowering the debt limit.

On GDP-linked bonds, creditors also demand a risk premium for taking on GDP volatility risk. This risk premium applies to GDP-linked bonds regardless of the level of debt or the probability of default. Simulating the volatility of returns due to the real GDP shock using the data from Schularick and Taylor (2012) produces a GDP risk premium of around 0.35 per cent. This increases the effective interest rate on GDP-linked bonds, steepening the slope of the interest – growth differential line in Figure 1.

However, just as it is difficult to account for the size of the equity risk premium with standard levels of risk aversion, it may be the case that the actual level of compensation required to hold GDP-linked bonds would be higher than suggested by this model. To test this, Table 3 shows the coefficient of relative risk aversion used to calculate the GDP risk premium (the level of risk aversion used for credit risk is kept constant at 4) gradually raised until the debt limit falls below the value of conventional bonds.

Table 3: Debt limits under various coefficients of relative risk aversion

Coefficient of relative risk aversion (τ)	Implied risk premium	Debt limit
4	0.35%	114%
8	0.7%	99%
12	1.0%	93%
15	1.5%	84%
20	2.5%	75%
24	3.7%	63%

Kamstra and Shiller (2009) used a CAPM to estimate the risk premium over conventional bonds to be 1.5 per cent. This implies a debt limit of 84 per cent, still well above that of conventional bonds. Only a coefficient of relative risk aversion of around 24 (implying a risk premium of over 3.7 per cent) would imply a lower debt limit than for conventional bonds.

Therefore, as long as the risk premium on GDP volatility is lower than around 3.5 per cent, GDP-linked bonds will have a higher debt limit and sovereigns will experience a lower incidence of default than with conventional bonds.

In some circumstances, GDP-linked bonds can both reduce the risk of sovereign default *and* be cheaper to issue than conventional bonds. The reason for this is that when debt approaches the debt limit, the credit spread on conventional bonds can *exceed* the GDP risk premium.

Figure 3 shows the expected return (equivalent to the sovereign's cost of borrowing) on conventional and GDP-linked bonds as a function of the debt level, again assuming our model-based GDP risk premium of 0.35 per cent. When debt is below 45 per cent of GDP, it is cheaper to issue conventional bonds, as the GDP risk premium is higher than the credit spread on conventional debt. But as debt approaches the debt limit (from 45% to 63% of

GDP), the credit spread on conventional bonds dominates the GDP risk premium. Of course, debt can only be issued at a finite interest rate above 63 per cent of GDP if the debt stock is made up of GDP-linked bonds.

Figure 3: Expected return on conventional and GDP-linked bonds

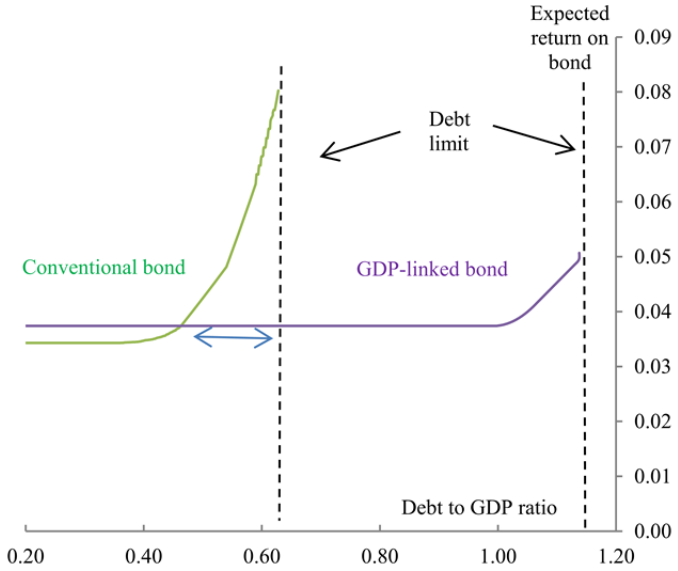

SECTION 6: FISCAL CONSOLIDATION AND GROWTH

One of the major benefits of GDP linked bonds identified by Borensztein and Mauro (2004) is that they can reduce the need for governments to undertake pro-cyclical fiscal policy during recessions. In our set-up, negative shocks to GDP growth worsen the debt-to-GDP ratio, as the denominator of this ratio falls. In our set-up, we assume that the government reacts to this shock by tightening fiscal policy in order to stabilise debt dynamics. In this section, the model incorporates feedback from fiscal consolidation to growth, through standard fiscal multiplier effects. This means that growth will weaken further, worsening debt dynamics by even more.

GDP-linked bonds can break this cycle, because sovereigns do not need to undertake fiscal consolidation to stabilise debt. Instead, debt is stabilised because the face value of GDP-linked bonds is reduced.

The fiscal feedback loop

Output in the economy is modelled through a very simple function:

$$Y_t = A_t\left(1 - \frac{S_t}{Y_t}\right) = A_t(1 - s_t) \qquad (15)$$

where Y_t is the level of total output and A_t is total factor productivity (TFP). Total factor productivity, A_t, is subject to the same growth shocks as previously defined, hence $A_{t+1} = A_t(1 + g^* + u_t)$.

This highly stylised characterisation of output is a simple way to generate a relationship between fiscal consolidation and output. Over the long run, this might be viewed as a supply-side relationship – a higher primary balance implies higher taxes, a lower incentive to work, and lower output. In the extreme, as the primary balance as a proportion of GDP approaches 100 per cent, the incentive to work is so low that output tends to zero. Of course this cannot occur in our model as the maximum feasible primary balance to GDP ratio is 0.1.

Over the short term, (15) might be viewed as a demand-side relationship. A higher primary balance implies higher taxes, which reduces consumption and therefore output. (15) is equivalent to assuming a fiscal multiplier of 1. The size of this multiplier can be varied to generate qualitatively similar results.

Total debt can be defined as a proportion of TFP, $\frac{D_t}{A_t}$. This is the state variable. Once this is defined, both output (Y_t) and the primary balance as a proportion of output (s_t) can be solved for simultaneously[4]. And this relationship, along with the model previously defined, can be used to estimate the debt limit with endogenous output.

Perhaps surprisingly, the maximum sustainable debt-to-GDP ratio is somewhat higher with a GDP feedback loop than without (70 per cent compared to 63 per cent for the model with risk aversion and exogenous output). This, however, reflects a lower denominator as the maximum sustainable debt *level* in the first period is lower with the GDP feedback loop. The steady-state debt-to-GDP ratio is 50 per cent and, reflecting the accelerator effect of the feedback loop, the frequency of default is markedly higher. On average, the sovereign defaults almost once every fifty years (a default rate of 1.8 per cent).

Introducing GDP-linked bonds

This model with endogenous output requires a change in the way GDP-linked bonds are priced. Rather than indexing GDP-linked bonds to the level of output, they are now indexed to the level of TFP, A_t. The reason for this is that in this simple two period set-up, it is impossible for investors to derive the interest rate on GDP-linked bonds[5]. Given the complications associated with this, it is much simpler to assume that GDP-linked bonds are instead indexed to TFP. (This modification is designed to simplify the model calculations, rather than a proposal to link GDP-linked bond returns to TFP -which is difficult to measure accurately - instead of GDP).

This means, however, that our state-contingent bonds lose some of their stabilisation role. They protect the sovereign from any shocks to TFP, and therefore break the feedback loop from these shocks to additional consolidation and even weaker growth. However, shocks to debt (ε_t) will still induce an increase in the primary balance which will reduce growth. But in this circumstance, the value of GDP-linked bonds will be unchanged, and will not provide any protection when this affects output. The subsequent analysis therefore underestimates the full beneficial effects of GDP-linked bonds.

The debt limit with GDP-linked bonds in this model is 115 per cent of GDP, similar as a ratio to GDP to the previous model but as above this reflects a lower denominator rather than a higher numerator. The fall in default risk compared to the case with conventional bonds is

[4] We exploit the fact that $\frac{D_t}{A_t}$ and $\frac{D_t}{Y_t}$ cannot diverge (i.e. the debt to TFP ratio cannot explode without the debt to output ratio exploding and vice versa) because $\frac{Y_t}{A_t}$ is bounded.

[5] This is because the price of a GDP-linked bond at t depends on expectations for GDP $t+1$, which in turn depend on expectations for the primary balance at $t+1$and thus expectations for the GDP at $t+2$ and so on.

particularly large – from 1.8 per cent per annum to just 0.15 per cent per annum (i.e. twice every three millennia as opposed to once every 55 years).

SECTION 7: WELFARE

Let us recap on what we have shown so far. In all of our different model versions, GDP-linked bonds would significantly increase the maximum sustainable debt-to-GDP ratio and lower the frequency of sovereign default. Once investor risk aversion is accounted for, GDP-linked bonds are more expensive than conventional bonds for low to medium debt-to-GDP ratios. But as debt approaches its limit, GDP-linked bonds become cheaper to issue than conventional bonds. And clearly beyond the debt limit for conventional bonds, the sovereign has a choice of indexing or default (at least in our simple model). So GDP-linked bonds provide insurance against default and the relative price of this insurance depends on how distant sovereign default appears to be. Table A1 summarises this information.

Thus far, we have ignored the plight of taxpayers. We have not written a general equilibrium model of sovereign default and the economy, so we have no internally consistent way of making welfare judgements. Nevertheless, taxpayers should be front and centre of policymakers' attention. So what are the implications of the sovereign debt indexation choice for them? The benefit we have considered – a reduced likelihood of default – certainly seems important. If De Paoli *et al* (2006) are to be believed, the output cost of default can be very large and persistent, even if they are rare. But the potential cost – a higher interest burden in normal times to compensate investors for taking on GDP volatility – also seems relevant if it results in higher taxes or lower (useful) government spending.

There is, however, one more obvious potential benefit of GDP-linked bonds. Investors demand a risk premium for holding aggregate risk because they dislike it. But it is by no means obvious that a sovereign is better placed to bear this risk. Ultimately this risk is borne by current and future taxpayers (assuming the sovereign does fulfil its obligations). In theory, the sovereign could use its ability to borrow and lend to smooth taxes through the economic cycle, rather than tightening policy when transitory shocks emerge. But recent experience shows the difficulties faced by policymakers in judging transitory from permanent shocks.

In our model, GDP, or productivity in the last version, follows a random walk with drift. Should a negative shock occur, which in turn raises the debt-to-GDP ratio above its steady-state, the policymaker will react by tightening policy. Figures 4 and 5 show this in action in simulations of the model with risk aversion and exogenous output. With GDP indexation (Figure 5), the elasticity of the change in primary balance with respect to output growth is much smaller than with conventional bonds (Figure 4) as the government is hedged against movements in GDP.

Figure 4: Output growth and the change in the primary balance, conventional bonds

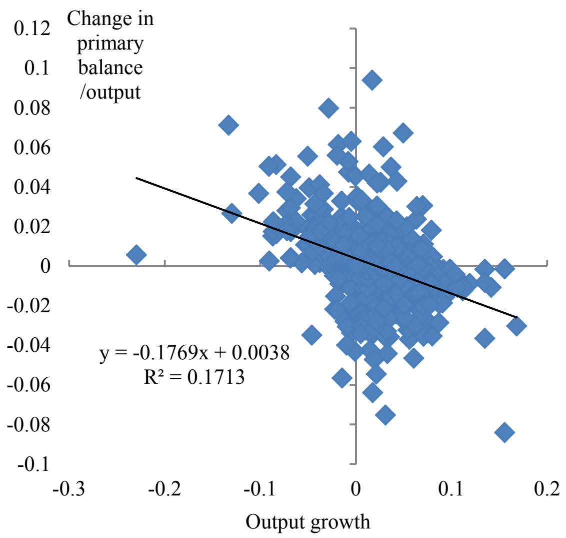

Figure 5: Output growth and the change in the primary balance, GDP bonds

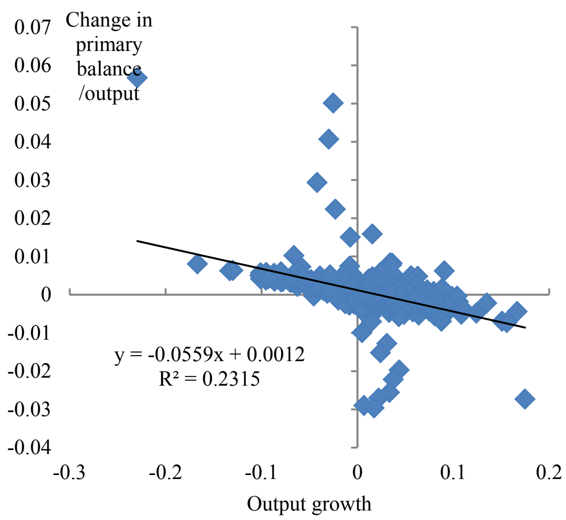

As a result, the standard deviation of the primary balance to GDP ratio is less than half as large with GDP-linked bonds. This volatility of taxes or government spending is likely to have a deleterious impact on welfare (and in our GDP feedback model, it raises the standard deviation of GDP growth by 20 per cent).

How should we weigh these pros and cons? One option is to assume that consumer-taxpayers are unable to borrow and lend abroad, while the government only borrows from foreigners. In this world, the government's primary balance would be equal to the country's trade balance. This is course a very restrictive assumption, but makes a welfare proxy much easier. Assuming away investment, consumption is simply equal to output minus the primary balance.

Table A1 shows the mean and standard deviation of the growth rate of our consumption proxy across four different model set-ups. The first three have been introduced in Sections 3 to 6. The final one is identical to the third, except a default penalty of 10 per cent of output is assumed. This output is lost and never regained (although in the year of default output tends to fall only very slightly as the primary balance falls sharply, offsetting most of the effect of the default cost). This output loss is within the range suggested by De Paoli *et al* (2006) and the judgement that output costs are permanent is informed by Cerra and Saxena (2008)'s finding that after major financial and political crises countries do not on average return to their previous output trend[6].

The figures in Table A1 are calculated by running a series of very long simulations (exploiting the law of motion from the debt accumulation equation and taking random draws from the shock distributions) and averaging across to obtain the summary statistics.

Mean consumption growth is 2.1 per cent, the growth rate of output/productivity, in every version of the model except the one incorporating default costs. In this version, default on conventional bonds is sufficiently frequent (once every 55 years) for the average growth rate to fall to 1.9 per cent.

The standard deviation of consumption growth shows much more variation. It is always larger for conventional bonds, reflecting the insurance benefits of GDP-linked bonds. The gap between the standard deviations for conventional and GDP-linked bonds jumps from around 40 to 60 basis points in the first two versions of the model to 110 to 120 basis points in the model with endogenous feedback as uninsured risk is exacerbated by the fiscal feedback loop.

Finally, the consumption share of output is lower for GDP-linked bonds in steady-state in all our model versions. There are two reasons for this. First, the cost of issuing GDP-linked bonds is, on average, higher. Second, the steady-state debt-to-GDP ratio is higher for sovereigns that issue GDP-linked bonds. But the difference between interest burdens is never very large: it is at most 0.5 per cent of GDP per annum.

Bringing these different factors into a utility function allows us to investigate their net effect. We use the same utility function for consumers as investors: CRRA with a coefficient of relative risk aversion equal to 4. As above, we compute expected utilities for the different model versions by running many series of long simulations and taking the mean[7].

Our baseline results, using a discount factor of 0.99, show that the gains from switching to GDP-linked bonds are equal to:

- Certainty equivalent consumption equal to **one per cent of first period GDP** in perpetuity in the version with risk-averse investors and exogenous output.

[6] A simple application of Cerra and Saxena's framework to sovereign debt crises and inflation crises (jumps to a high rate of inflation) suggests permanent output costs are also associated with these events. Results available on request.
[7] We use the same starting point for each type of bond (the steady-state debt to GDP ratio under conventional debt).

- Certainty equivalent consumption equal to **two per cent of first period GDP** in perpetuity in the version with risk-averse investors and endogenous output.
- Certainty equivalent consumption equal to **nine per cent of first period GDP** in perpetuity in the version with risk-averse investors, endogenous output and a output cost of default.

The significantly larger gain from switching in a world with large default costs should be no surprise as the costs are sufficient to reduce the average growth rate of consumption. It is also the case that economies with less volatile GDP (for instance, those with an independent central bank) will experience smaller welfare gains.

The gains from switching to GDP-linked bonds are decreasing in the discount factor, reflecting the fact that a government making such a choice would borrow temporarily as it adjusts towards the higher steady-state debt-to-GDP ratio. We also used a discount factor of 1 and obtained qualitatively and quantitatively similar results (except for the gains in the version with default costs which were much larger).

Our results suggest much larger potential welfare gains from GDP-linked bonds than Sandleris *et al* (2008), who found a gain of 0.5% in certainty equivalent consumption. We believe the bulk of this difference is accounted for by our assumption of a permanent output cost of default as opposed to their assumption of a temporary output cost.

There are of course considerable uncertainties around our results. First, we have constrained the behaviour of consumers unreasonably. They could, for example, change their risk profiles to counteract the stabilising impact of GDP-linked bonds. But they might equally embrace a new risk-sharing instrument, perhaps encouraged by regulators who appreciate the benefits for financial stability.

Second, the behaviour of the fiscal policymaker is very crudely modelled and, even if an approximate guide to past behaviour, may be a very poor guide to policy once GDP-linked bonds are introduced. It is possible to imagine imprudent policymakers taking full advantage of the higher debt limits associated with GDP-linked bonds, risking much larger sovereign defaults than we have seen hitherto.

Third, our modelling approached bond pricing in a very simple manner which may under-price GDP risk and takes no account of the liquidity premia associated with issuing a new type of instrument.

Nevertheless, our simple model and welfare experiment captured what we believe to be the main cost and benefit of issuing GDP-linked bonds. The results suggest that there may be a strong case for using such instruments in a more default-prone world. Indeed, in our analysis, countries which are close to the brink may even find it cheaper to issue GDP-linked bonds.

SECTION 8: CONCLUSIONS AND SUGGESTIONS FOR FURTHER RESEARCH

In this paper we have adapted the model of Ghosh *et al* to allow a comparison of the costs and benefits of conventional and GDP-linked bonds. *Ceteris paribus* GDP-linked bonds raise the maximum sustainable level debt of a sovereign and reduce the probability of default. A default, in the conventional sense of reneging on an existing bond contract, is replaced by a contracted decline in the redemption value of the bond. The avoidance of the significant deadweight costs associated with sovereign debt crises and default could lead to material increases in welfare. And greater risk-sharing across more states of the world is Pareto improving. Greater use of GDP-linked bonds would also move the sovereigns away from a reliance on public insurance towards greater self-insurance.

While all countries might experience some benefit from the use of GDP-linked debt, economies with higher GDP growth volatility (such as EMEs) or countries where monetary policy is constrained (such as those in a monetary union) are likely to benefit most.

Our model differs from that of Ghosh *et al* in a number of respects. In particular, our baseline model includes GDP growth that follows a stochastic process; and a primary fiscal reaction function that reacts to all factors that influence sovereign debt dynamics, rather than just the debt level. We then add GDP-linked bonds, risk-averse investors, exogenous shocks to GDP growth and feedback effects from the primary balance to GDP. In each variant of the model we find that GDP-linked bonds raise the sovereign's debt limit, and reduce the incidence of default.

The introduction of endogenously determined GDP allows us to examine the effect of pro-cyclical fiscal adjustment as negative growth shocks increase the debt-to-GDP ratio and the government raises the primary surplus ratio in line with its fiscal reaction function. With GDP-linked bonds these growth shocks have no direct effect on the debt ratio, which breaks this cycle.

While GDP-linked bonds increase the debt limit, investors demand a premium for providing this GDP-volatility insurance. As the debt-to-GDP ratio increases however, this cost gets overturned as the lower debt limit associated with conventional bonds leads to their default premium increasing, while that on GDP-linked bonds remains unchanged, with the result than conventional bonds carry the greater servicing cost. As to the whether or not this cost outweighs the benefits of being able to tolerate larger shocks to the debt ratio without defaulting depends in part on the sovereign's preferences and on the costs of default. We attempt to capture this using a simple time-separable social welfare function combined with an arbitrary default cost.

Our experiment suggests that the gains from issuing GDP-linked bonds – lower consumption volatility and lower default frequency – outweigh the higher expected interest burden associated with them. The gain is large in magnitude for the version of the model with the default penalty.

Countries that issue debt in local currency (and are outside of a monetary union) also have the ability to inflate away the real value of debt obligations. Although we have not explicitly included this option into our model, we would argue that this is analogous to a sovereign default, with similar costs to creditors and the wider economy[8].

Future work would be merited to: (i) investigate the determinants of the optimal degree of GDP indexation; (ii) allow for a richer description of sovereign and consumer behaviour, perhaps in a DSGE framework so long that it could remain tractable; (iii) focus on the interaction between sovereign debt structure and the ability to run countercyclical fiscal policy.

Finally there may be scope for further assessment of other costs and obstacles to GDP-linked bond issuance. Griffith-Jones and Sharma (2009) report a much wider range of investor concerns than are covered in this paper. Although some concerns such as how to price the bonds and uncertainty about GDP data are easily surmountable in our view, the costs associated with starting and ensuring liquidity in a new market may be more teething. They suggest both a role for international financial institutions in market-making and a co-ordinated issuance as solutions. We believe further analysis of these problems and potential solutions such as these would be valuable.

[8] This is supported by footnote 6 above.

BIBLIOGRAPHY

Abbas, S, Belhocine, N, El Ganainy, A, and Horton, M (2011), 'Historical patterns of public debt – evidence from a new database', *IMF Economic Review*, Vol. 59, Issue 4.

Abiad, A And Ostry, J (2005),'Primary surpluses and sustainable debt levels in emerging market countries', *IMF Policy Discussion Paper PDP/05/6*.

Aguiar, M and Gopinath, G (2006), 'Default able debt, interest rates and the current account', *Journal of International Economics*.

Arellano, C (2008), 'Default risk and income fluctuations in emerging markets', *The American Economic Review*.

Barro, R (1995), 'Optimal debt management', *NBER Working Paper No. 5327*.

Benjamin, D and Wright, M (2009), 'Recovery before redemption: a theory of delays in sovereign debt renegotiations', available at SSRN: http://ssrn.com/abstract=1392539.

Bohn, H (1998), 'The behaviour of US public debt and deficits', *The Quarterly Journal of Economics*.

Bohn, H (2005), 'The sustainability of fiscal policy in the United States', *CESifo Working Paper No. 1446*.

Borensztein, E and Mauro, P (2004), 'GDP-indexed bonds', *Economic Policy*.

Brooke, M., Mendes, R., Pienkowski, A. and Santor, E. (2013) ,'Sovereign Default and State-Contingent Debt' Bank of England Financial Stability Paper 27

Campos, C, Jaimovich, D and Panizza, U (2006), 'The unexplained part of public debt', *Inter-American Development Bank Working Paper No. 554*.

Cerra, V and Saxena, S (2008), 'Growth dynamics: the myth of economic recovery', *American Economic Review*, Vol. 98, No. 1, pages 439–57.

Chamon, M and Mauro, P (2005), 'Pricing growth-indexed bonds', *IMF Working Paper No. WP/05/216*.

Cruces, J and Trebesch, C (2011), 'Sovereign defaults: the price of haircuts', *CESifo Working Paper No. 3604*.

De Paoli, B, Hoggarth, G and Saporta, V (2006), 'Costs of sovereign default', *Bank of England Financial Stability Paper No. 1*.

Eaton, J, and Gersovitz, M (1981), 'Debt with potential repudiation: theory and empirical analysis', *Review of Economic Studies*.

Froot, K, Scharfstein, D and Stein, J (1989), 'LDC debt: forgiveness, indexation, and investment incentives', *The Journal of Finance*.

Ghosh, A, Kim, J, Mendoza, E, Ostry, J and Qureshi, M (2011), 'Fiscal fatigue, fiscal space and debt sustainability in advanced economies', *NBER Working Paper No. 16782*.

Griffith-Jones, S and Sharma, K (2009), 'GDP-indexed bonds: making it happen', Innovative Financing for Development: 79.

Haldane, A (2012), 'Tails of the unexpected', speech given at 'The Credit Crisis Five Years On: Unpacking the Crisis', conference held at the University of Edinburgh Business School, 8-9 June, available at: www.bankofengland.co.uk/publications/Documents/speeches/2012/speech582.pdf.

International Monetary Fund (2003), 'Public debt in emerging markets: is it too high?' *World Economic Outlook*, September 2003.

Kamstra, M and Shiller, R (2009), 'The case for trills: giving the people and their pension funds a stake in the wealth of the nation', *Cowles Foundation Discussion Paper No. 1717*.

Krueger, A (2002), 'A new approach to sovereign debt restructuring', available at: www.imf.org/external/pubs/ft/exrp/sdrm/eng/sdrm.pdf.

Krugman, P (1988), 'Financing vs. forgiving a debt overhang', *NBER Working Paper No. 2486*.

Levy Yeyati, E and Panizza, U (2011), 'The elusive costs of sovereign defaults', *Journal of Development Economics*.

Mendoza, E and Ostry, J (2007), 'International evidence on fiscal solvency: is fiscal policy 'responsible?', *NBER Working Paper No. 12947*.

Miller, M and Zhang, L (2013), 'Fiscal consolidation: Dr Pangloss meets Mr Keynes', *mimeo*.

Reinhart, C and Sbrancia, M (2011), 'The liquidation of government debt', *NBER Working Paper No. 16893*.

Ruban, O, Poon, S, and Vonatsos, K (2008), 'GDP linked bonds: contract design and pricing', available at SSRN: http://ssrn.com/abstract=966436.

Sandleris, G, Sapriza, H and Taddei, F (2008), 'Indexed sovereign debt: an applied framework', *Collegio Carlo Alberto Working Paper No. 104*.

Schularick, M and Taylor A (2012), 'Credit booms gone bust: monetary policy, leverage cycles, and financial crises, 1870–2008', *American Economic Review*, No. 102, pages 1029–61.

Shiller, R (1993), 'Macro markets: creating institutions for managing society's largest economic risks', Oxford University Press.

Shiller, R (2003), 'The new financial order: risk in the 21st century', Princeton University Press.

APPENDIX

Interest rate with risk-averse investors

Default risk premia

The no arbitrage condition with risk-neutral investors means that the expected returns from investing in both risk-free and risky bond must be equal (2).

When investors are risk-averse however, expected returns, consumption and portfolio shares are jointly determined, usually in the familiar expected-utility maximisation framework. Since we want to use this framework to determine expected returns we have to supply some information about portfolio shares or consumption, and because we are not working with a general equilibrium model, we make an arbitrary assumption about portfolio shares i.e. that at the margin investors are indifferent between adding an extra bond or safe asset to their portfolio. [9] This requires that the expected *utility* from holding an extra unit of either asset be equal. For investors with CRRA utility (11) this gives us

$$U[(1+r^*)] = (1-p_{t+1}).U[(1+r_t] + p_{t+1}.U[\theta(1+r^*)]$$

$$(1+r^*)^{1-\tau} = (1-p_{t+1})(1+r_t)^{1-\tau} + p_{t+1}.\theta^{1-\tau}(1+r^*)^{1-\tau}$$

$$(1+r_t) = (1+r^*).\left[\frac{1-p_{t+1}.\theta^{1-\tau}}{1-p_{t+1}}\right]^{\frac{1}{1-\tau}} \quad \text{(A1)}$$

which is equation (12) in the body of the paper.

GDP risk premia

The risk premium on GDP volatility associated with GDP-linked bonds is derived in a similar way. In a risk-neutral set-up (assuming no default risk for simplicity), equating the expected return on risk-free and risky bonds gives

$$(1+r^*) = E_t(1+r_t).(1+g^*+u_{t+1})$$

$$(1+r^*) = (1+r_t)E_t(1+g^*+u_{t+1}).$$

$$(1+r^*) = (1+r_t).(1+g^*)$$

$$(1+r_t) = \frac{(1+r^*)}{(1+g^*)} \quad \text{(A2)}$$

[9] The use of any arbitrary assumption to close the model is not ideal but it does serve to give a sense of the way in which risk premia and default probabilities are related. We intend to develop this research using a general equilibrium model in a future paper.

where r_t is the market interest rate on a GDP-linked bond. For a risk-averse investor the equality of expected returns leads to:

$$U\{1+r^*\} = E(U\{(1+r_t).(1+g^*+u_{t+1})\})$$

$$\frac{(1+r^*)^{1-\tau}}{1-\tau} = \frac{(1+r_t)^{1-\tau}}{1-\tau}.E_t[\frac{(1+g^*+u_{t+1})^{1-\tau}}{1-\tau}]$$

$$(1+r^*)^{1-\tau} = (1+r_t)^{1-\tau}E_t[(1+g^*+u_{t+1})^{1-\tau}].$$

Here, the variable under the expectations operator is based on a non-standard distribution and therefore it difficult to represent this as an analytical solution. Instead, we will solve for this numerically, and re-named is as 'G'-

$$(1+r^*)^{1-\tau} = (1+r_t)^{1-\tau}.E_t[G]$$

$$(1+r_t) = \frac{(1+r^*)}{E(G)^{1/_{1-\tau}}} = \frac{(1+r^*)}{(1+g^*-\vartheta)} \qquad \text{(A3)}$$

where equation (A3) is the market interest rate on a GDP-linked bond for risk-averse investors. Because investors are risk-averse, it is always the case that $E(G)^{1/_{1-\tau}} < (1+g^*)$ and therefore the interest rate on a GDP-linked bonds with risk-averse investors (A3) is higher than that of a risk-neutral investor (A2). Another representation is to explicitly include the risk premium, ϑ, in the interest rate term.

Table A1: Summary of key model outputs

Model features	Conventional bonds					GDP-linked bonds				
	Steady-state debt level (d^*)	Debt limit ($\bar{d}$)	Frequency of default (per 100 years)	Mean consumption growth	Standard deviation of consumption growth	Steady-state debt level (d^*)	Debt limit ($\bar{d}$)	Frequency of default (per 100 years)	Mean consumption growth	Standard deviation of consumption growth
Basic set-up	78%	93%	0.5	2.1%	4.2%	93%	195%	≈0	2.1%	3.8%
Risk-averse investors	52%	63%	1	2.1%	4.7%	76%	114%	0.1	2.1%	4.1%
Risk-averse investors, endogenous output	50%	70%	1.8	2.1%	5.5%	76%	115%	0.2	2.1%	4.4%
Risk-averse investors, endogenous output, default penalty	50%	70%	1.8	1.9%	5.6%	76%	115%	0.2	2.1%	4.4%