DE GRUYTER BE J. Macroecon. 2014; aop

Advances

Alexis Anagnostopoulos* and Xin Tang

Evaluating linear approximations in a two-country model with occasionally binding borrowing constraints

Abstract: Under a linear approximation, a standard two-country business cycles model with incomplete markets delivers consumption and debt dynamics that are non-stationary (unit root) and a bond price that is independent of the wealth distribution. We argue that these two features are due to the local nature of the approximation and we show that they survive even when second or third order local approximations are used. However, these features disappear when debt limits and the associated precautionary motives are taken into account by a standard, global solution method. We subsequently investigate whether this qualitative difference has significant quantitative implications regarding the linear solution as an approximation to the model's equilibrium dynamics. Policy function differences between the local and global solutions can be large and remain significant even in the case of debt limits as loose as the natural debt limit. These differences can lead to significant discrepancies in implied simulated second moments. In a benchmark calibration, the cross-country correlation of consumption is 0.61 under linearization, but only 0.38 when a policy iteration algorithm is used.

Keywords: borrowing constraints; incomplete markets; linearization; perturbation methods.

JEL classification: E21; F44; C63.

DOI 10.1515/bejm-2013-0175

***Corresponding author: Alexis Anagnostopoulos,** Stony Brook University – Economics, 100 Nicolls Road, Stony Brook, NY 11794-4384, USA,
e-mail: alexis.anagnostopoulos@stonybrook.edu
Xin Tang: Stony Brook University – Economics, Stony Brook, NY, USA

1 Introduction

The two-country real business cycles model of Backus, Kehoe, and Kydland (1992), and its extension to incomplete financial markets as in Baxter (1995), has played the role of a workhorse model in open economy macroeconomics for more than two decades. Given the amount of research still being produced in this context, it is important to have a good understanding of the equilibrium dynamics of consumption and debt arising from commonly used solution methods. Local linear approximations are often used, partly because of their ease of implementation. Under such an approximation, consumption and debt dynamics exhibit a unit root which raises some concerns, both theoretically and computationally. In this paper, it is argued that the equilibrium of this model features stationary, mean-reverting stochastic processes for consumption and debt and a solution method that does not use *only* local information will deliver this. Whether this qualitative difference leads to significant quantitative differences, and along which dimensions, is the main object of analysis of this article. It is shown that approximation errors resulting from a linear approximation can be significant. These errors can accumulate in a simulation resulting in significant inaccuracy of the linear approximation with respect to statistics such as cross-country consumption correlations, which are often used as a metric to evaluate the model against the data.

Stochastic dynamic general equilibrium models characterize equilibria through a collection of non-linear stochastic difference equations that are necessary conditions for equilibrium. Most often, that system of difference equations for optimal policy functions presents a non-trivial problem that does not admit an analytical solution and these models are therefore commonly analyzed using numerical solution methods. A common method consists in obtaining a local linear approximation of the system in a neighborhood of the deterministic steady state, a procedure described, for example, in King, Plosser, and Rebelo (1988). Equilibrium laws of motion for all variables can then be computed using the methods presented by Blanchard and Kahn (1980), Uhlig (1999) or Christiano (2002) amongst others. In a special issue of the *Journal of Business and Economic Statistics* in 1990, a number of researchers combined in comparing a variety of numerical methods and their accuracy in the context of the stochastic growth model. Christiano (1990) and Mcgrattan (1990), in particular, provided linear methods and found that the accuracy was acceptable. Being local approximations, these methods rely on the assumption that the economy fluctuates in a region around the steady state of the corresponding deterministic economy. In one sense, finding accuracy to be acceptable validates the above assumption, in the context of the particular model used.[1]

The same linear methods have also been used by Baxter (1995), Baxter and Crucini (1995) and Kollman (1996) amongst others, to study international business cycles with incomplete markets. The model considered in these studies consists of two countries, each populated by a representative agent maximizing utility in the face of idiosyncratic as well as aggregate uncertainty but with limited ability to insure against idiosyncratic risk through financial markets. It is found that these ingredients imply unit root dynamics in bonds. From a theoretical perspective, this finding is disconcerting. It implies that debt can rise to arbitrarily high levels and eventually surpass any debt level, including the natural debt limit. This raises the question of whether the equilibrium computed is consistent with the model's underlying assumptions. On the other hand, the computed equilibrium is only intended as an approximation to the true equilibrium. If the approximation is good enough for the intended use of this equilibrium, the theoretical concern could be seen as second order. It is the objective of the present paper to evaluate the quality of approximation along several dimensions. This is important because linear approximations are very appealing due to their low computational cost but also due to the fact that they deliver closed-form solutions that lay the model's mechanisms bare.

There are good a priori reasons to suspect that the linearization approach might be worse in this setup than in a standard stochastic growth model. Local linear methods are based on the assumption that the economy fluctuates around the deterministic steady state. If the equilibrium law of motion has a unit root (or close to it), this assumption is no longer valid since, even if the economy begins at steady state, it will eventually drift away. In addition, equilibrium existence necessitates the imposition of debt limits and one can expect policy functions to exhibit significant nonlinearity close to those limits. When the debt process is highly persistent, these are points that will potentially lie inside the stationary distribution. Although there is a substantial literature on comparing different computational methods for incomplete markets economies[2], this literature typically focuses on closed economy models with a continuum of agents. In these comparisons, linearization methods are not considered, mainly because they are typically not used by researchers working with such models. Our evaluation of linearization methods has added value in the context of a two-country business

1 In a different context, namely the Diamond-Mortensen-Pissarides model, Petrosky-Nadeau and Zhang (2013) have recently argued that perturbation methods can be inaccurate.

2 See the January 2010 special issue of the *Journal of Economic Dynamics and Control* and the references therein.

cycle model with incomplete markets because such methods are often used within that literature. We find the following results.

First, the unit root result under linearization is reproduced in a model that is stripped down to the bare essentials. Relative to the international business cycles literature motivating this paper, this means abstracting from production and assuming exogenous income processes. The model is, thus, similar to those used by Telmer (1993), Heaton and Lucas (1996), Marcet and Singleton (1999) and den Haan (2001) to study asset prices under incomplete markets in a closed economy. Simplifying the model in this way clarifies the crucial elements responsible for the unit root result, namely two agents facing uninsurable idiosyncratic risk and trading a risk free bond.[3] A second, closely related feature of the equilibrium under linearization is that bond prices only respond to aggregate income changes. In particular, changes in the distribution of wealth leaves bond prices unaffected.

Second, an accurate, non-linear global solution to the model is computed using a standard policy iteration algorithm and compared to the linear approximation. Linearization ignores debt limits and delivers a certainty equivalent solution whereas the policy iteration solution incorporates the effects of risk and the possibility of future binding constraints, even at points in the state space that are far from the debt limits. The resulting precautionary motives reduce the willingness of agents to accumulate debt and lead to a stationary, albeit highly persistent, equilibrium debt process.[4] In addition, the distribution of wealth does affect the bond price. The reason has to do with the concavity of the consumption function [see Carroll and Kimball (1996)] and has been eloquently explained in den Haan (2001).

Third, we argue that these two aspects of the equilibrium cannot be captured by considering higher order perturbation methods. We prove this analytically for the case of a second order approximation and confirm it numerically for a third order approximation. This suggests that the absence of distributional effects in the equilibrium bond price and the non-stationarity of debt are not due to the linearity of the approximation but rather result from the *local* nature of the approximation. Local approximation methods necessarily ignore the effects of debt limits and these have effects on equilibrium behavior even if they do not

3 Indeed, it is straightforward to show that this result carries over to many incomplete markets settings including models with two goods, models with equities as well as bonds as in Marcet and Singleton (1999) but also optimal (Ramsey) taxation models like the one analyzed in Marcet and Scott (2009). The presence of a risk free bond is sufficient.

4 Precautionary motives are well understood in the consumption/savings literature, see for example, Zeldes (1989), Kimball (1990), Deaton (1991) or Carroll and Kimball (1996, 2001). This is simply an application to a two-agent, general equilibrium model with an international flavor.

bind in equilibrium, as in the case of natural debt limits. For this reason, we conjecture that perturbation methods of even higher order would still miss these effects.

One implication of this difference arising from solution methods, is that equilibrium paths for consumption and debt can be substantially different depending on the solution method. We illustrate this using long simulations produced under the different solutions. Consumption paths implied by perturbation methods eventually imply negative (and arbitrarily large so) consumption, a feature absent from the global solution method because, as consumption falls, marginal utility rises fast and prevents consumption from falling to zero.

Fourth, we provide a quantitative analysis of the performance of the linearization method along several dimensions. Here, we focus mainly on the case of *ad hoc* debt limits which delivers debt levels lying within the empirically observed range.[5] We find significant distance between the two computed solutions. As an indication, under the benchmark calibration the maximum distance between consumption policy functions is more than 7%, even when we ignore points of the state space where the limits bind. In turn, this can lead to significant differences in the implied impulse response functions and cross-country consumption correlations, two statistics that are commonly used when confronting this model with data. With respect to the former, although the initial impact is similar, the long run effects can be substantially different due to the non-stationary nature of the linearly approximated equilibrium. Consumption correlations can diverge significantly suggesting that the quantitative success of an international business cycle model can often be sensitive to approximation error. In a benchmark experiment, the linear solution yields a consumption correlation of 0.61, whereas the policy iteration solution implies a value of 0.38. In Appendix C we show that the policy iteration solution is indeed a very accurate approximation by computing both static and dynamic Euler errors.

Finally, we conduct extensive sensitivity analysis with respect to all of the model's parameters. We find that the linear approximation deteriorates as the variance of the exogenous income process or the subjective discount factor increases. The persistence of the exogenous income process has a non monotonic effect on the accuracy of the linear approximation. The accuracy deteriorates with persistence at low levels of persistence but improves as income shocks become close to permanent. We also investigate the performance of linear approximations as the borrowing limit is relaxed. Accuracy improves with looser limits, but errors persist even when the limit is chosen to be the natural debt limit. This

5 Imposing only the natural debt limit leads to excessive variation in debt levels in the model, a feature that is inconsistent with the evolution of net foreign asset positions in practice.

provides support for our claim that these methods incur approximation errors not simply because they cannot deal with binding limits *per se*, but rather because they do not incorporate the precautionary effects arising from the *possibility of future binding limits*.

The rest of the paper is organized as follows. Section 2 presents the model, Section 3 discusses the solution obtained using first, second and third order perturbation methods and discusses potential pitfalls of these methods that are specific to this model. Section 4 presents and discusses the main numerical results and Section 5 concludes.

2 The model

The model used is a standard two-agent general equilibrium model with incomplete markets as in Telmer (1993). It is interpreted and calibrated as a two-country model, where markets are complete within each country, so that we can assume a representative agent in each country, but incomplete across countries. It is thus a simplified version of Baxter (1995), Baxter and Crucini (1995) or Kollman (1996), the simplification being that production is assumed exogenous here, whereas it is endogenous in those studies.

The representative agent in each country $i=1, 2$ aims to maximize the expected discounted sum of period utility from consumption c_{it}

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_{it}) \tag{1}$$

where $\beta \in (0, 1)$ is the subjective discount factor. The period utility function $u(\cdot)$ is strictly increasing, strictly concave and satisfies $\lim_{c\to 0} u'(c)=\infty$. Income in country i, period t is given by an exogenous stochastic endowment y_{it} which follows an AR(1) process

$$y_{it}-\bar{y}=\rho(y_{i,t-1}-\bar{y})+\varepsilon_{it} \quad i=1,\ 2 \tag{2}$$

where $\bar{y}$ is mean income, ρ is the persistence of income and $\varepsilon_{it}\sim N(0, \sigma_{\varepsilon}^2)$. Asset markets are restricted by assuming that countries can only trade one-period risk free bonds b_{it} which are bought at price p_t and pay one unit of consumption in the following period. The budget constraint of each country is thus

$$\begin{aligned} &c_{it}+p_t b_{it}=y_{it}+b_{i,t-1} \\ &b_{i,-1} \text{ given} \end{aligned} \tag{3}$$

In addition, exogenous upper limits on debt are imposed for each i

$$b_{it} \geq -K \tag{4}$$

with K a positive parameter. Clearly, if equilibrium is to exist, these limits have to be chosen so that they are at least as strict as the natural borrowing limit. On the other hand, if one wanted to avoid introducing additional frictions to the model on top of the missing asset markets, one would want to make sure these limits never bind in equilibrium. Anticipating the results of the next section, it should be noted here that this is impossible under a linearized solution. It is also important to clarify that, even if limits never bind in equilibrium, they may affect the properties of the equilibrium. Under rational expectations, agents anticipate that there is a point above which their debt cannot grow any further and take this into account in their decisions, more so the closer they are to the limit. Indeed that is the whole purpose of using debt limits to ensure equilibrium existence, otherwise they would be redundant.[6] In the computations of the following sections, we will choose an ad hoc level for the limit K and investigate the effects of varying the parameter K all the way up to the natural debt limit.

The model is closed by assuming that bond markets clear

$$b_{1t} + b_{2t} = 0 \ \forall t \tag{5}$$

and goods markets clear

$$c_{1t} + c_{2t} = y_{1t} + y_{2t} \ \forall t \tag{6}$$

Letting $\beta^t \lambda_{it}$ be the multiplier on the debt limit at t for country i, the equilibrium is characterized by

$$p_t u_c(c_{it}) - \lambda_{it} = \beta \mathbb{E}_t u_c(c_{i,t+1}) \quad i=1,\ 2 \tag{7}$$

$$\lambda_{it}(b_{it} + K) = 0 \quad i=1,\ 2 \tag{8}$$

together with budget constraints (3), market clearing conditions (5) and (6) as well as a transversality condition. Equation (7) is the usual Euler equation stating that the marginal benefit from using debt to increase consumption at t must be greater or equal to the expected marginal loss at t+1 arising from the additional

6 For a thorough analysis of debt limits under complete markets see Ljungqvist and Sargent (2004). Levine and Zame (1996) and Magill and Quinzii (1994) discuss those in the case of market incompleteness.

debt. It will be exactly equal whenever $\lambda_{it}=0$, i.e., whenever i is unconstrained. Equation (8) is the complementary slackness condition.

We compute the equilibrium of this model using alternative solution methods. Perturbation methods of first, second and third order are discussed in the following section. A policy iteration algorithm along the lines of Coleman (1990) is used in subsequent sections as an example of a global approximation method. The policy iteration algorithm is described in Appendix B.

3 Perturbation methods

It is well known that, in principle, the presence of occasionally binding inequality constraints makes the linearization approach inappropriate for the solution of this model. However, many researchers have used this approach by simply ignoring these inequality constraints.[7] Since one of the objectives of this paper is to quantitatively evaluate these approximations, we obtain a linear solution ignoring the debt limits.

Formally, the linearization method is simply a special case of the more general approximation method known as perturbation.[8] The state variables in this model include the exogenous endowments y_{1t} and y_{2t} as well as the distribution of assets. In this two-country model, the debt level of country 1 is a sufficient statistic for the whole distribution of assets. As a result, $b_{1,t-1}$ is used as the third state variable. The objective is to obtain approximations for the equilibrium policy functions for bonds and consumption of country 1, b_{1t} and c_{1t}, as well as for the equilibrium bond price p_t. Consumption and bonds in country 2 can always be inferred from market clearing and are therefore omitted. We use bars to denote the non-stochastic steady state values of the endogenous variables and a circumflex to denote the deviation of a variable from its steady state. Using this notation, we express the equilibrium policies in terms of deviations from steady state as follows

$$\hat{b}_{1t}=B(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}; \sigma_\varepsilon)$$
$$\hat{c}_{1t}=C(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}; \sigma_\varepsilon)$$
$$\hat{p}_t=P(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}; \sigma_\varepsilon)$$

7 The justification for doing so is that one can choose the exogenous limit K to be large enough so that it never binds in equilibrium, in which case a linear approximation could be a good first approximation to the equilibrium of this model.

8 See Judd (1998) and the references therein. For a simple introduction to the practical application of these methods, we refer the reader to Schmitt-Grohé and Uribe (2004).

Note that we have made explicit the dependence of these policy functions on the standard deviation of the exogenous endowments σ_ε, which will play the role of the perturbation parameter. We obtain approximations of these policy functions in the neighborhood of the non-stochastic steady state, i.e., near the point $(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}; \sigma_\varepsilon)=(b_{1,-1}, 0, 0; 0)$. At the deterministic steady state, the absence of income variation implies no bond trade, which in turn means that bonds remain at their initial level $b_{1,-1}$. A Taylor approximation of these policy functions is characterized by the derivatives of each policy function evaluated at the deterministic steady state. For a first order (linear) approximation of $B(\cdot)$, this means that we need to compute four numbers B_b, B_{y1}, B_{y2}, B_σ where the subscript denotes the variable with respect to which derivatives are taken and it is understood that these are evaluated at $(b_{1,-1}, 0, 0; 0)$. Approximating the functions $C(\cdot)$ and $P(\cdot)$ in a similar fashion, a first order approximation boils down to the computation of 12 numbers. For a second order approximation, one would also require second order derivatives such as $B_{b,b}$, $B_{b,y1}$, $B_{b,y2}$, $B_{b,\sigma}$ and so on, in total an additional 48 numbers. These are obtained by implicitly differentiating the equilibrium conditions and subsequently solving the resulting system of simultaneous equations. Applying this method to our economy, the following results can be proved.[9] We first analyze the first order perturbation.

Result 1 *(Linear Approximation):* $B_\sigma=C_\sigma=P_\sigma=0$, $B_b=1$, $P_b=0$, $P_{y1}=P_{y2}$.

As expected, in the linear approximation all the σ_ε coefficients are zero, indicating a certainty equivalent solution. More interestingly, the equilibrium law of motion for bonds has a unit root and the equilibrium bond price is independent of the wealth distribution, i.e., it depends only on current aggregate income $y_{1t}+y_{2t}$. This result does not depend on specific utility assumptions, in fact it is true even under heterogeneous utilities. In addition, even though our proof assumes an AR(1) processes for the endowments as in (2), it is easy to extend it to a wide range of processes for these endowments. Finally, note that the result is independent of the value of $b_{1,-1}$. The unit root in bonds and the absence of bonds from the price equation are closely related. In a similar setting with exogenous prices (a small open economy), Schmitt-Grohé and Uribe (2003) also note the presence of a unit root in bonds and show that a stationary law of motion for bonds can be obtained by exogenously introducing debt-elastic interest rates. In the following sections, we show that interest rates are debt-elastic in general equilibrium and bonds follow a stationary law of motion.

9 The proofs of these results involve straightforward but tedious algebra. They are omitted from the paper but available in an online appendix.

In order to discuss the intuition obtained from the linear approximation, we now make some simplifying assumptions.[10] First, we assume the utility function is the same across countries and of the constant relative risk aversion (CRRA) form

$$u(c)=\begin{cases}\dfrac{c^{1-\gamma}-1}{1-\gamma} & \text{for } \gamma\neq 1\\ \log(c) & \text{for } \gamma=1\end{cases} \tag{9}$$

Second, we assume $b_{i,-1}=0$, which ensures countries are ex ante identical. The latter simplification is standard in the literature. It is also consistent with the stochastic version of the model in the following sense: because of the presence of precautionary motives, the mean of the stationary distribution for bonds is equal to zero for any value of $b_{i,-1}$.[11] With these assumptions, the first order approximation is given by

$$\hat{b}_{1,t}=\hat{b}_{1,t-1}+\frac{1-\rho}{2(1-\beta\rho)}(\hat{y}_{1,t}-\hat{y}_{2,t}) \tag{10}$$

$$\hat{p}_t=\frac{\gamma(1-\rho)}{2\bar{y}}\beta(\hat{y}_{1,t}+\hat{y}_{2,t}) \tag{11}$$

$$\hat{c}_{1t}=(1-\beta)\hat{b}_{1,t-1}+\left(1-\beta\frac{1-\rho}{2(1-\beta\rho)}\right)\hat{y}_{1,t}+\beta\frac{1-\rho}{2(1-\beta\rho)}\hat{y}_{2,t} \tag{12}$$

$$\hat{c}_{2t}=-(1-\beta)\hat{b}_{1,t-1}+\left(1-\beta\frac{1-\rho}{2(1-\beta\rho)}\right)\hat{y}_{2,t}+\beta\frac{1-\rho}{2(1-\beta\rho)}\hat{y}_{1,t} \tag{13}$$

$$\hat{y}_{it}=\rho\hat{y}_{i,t-1}+\hat{\varepsilon}_{it}\ \text{for } i=1,\ 2 \tag{14}$$

Borrowing is decreasing in relative income so that the country with the higher relative income will be a lender. Consumption is increasing in both home and foreign incomes and increasing in assets inherited from the previous period. The slope of consumption with respect to one's own income is higher than with respect to the other's income which implies that risk sharing is imperfect. In addition, as income becomes more persistent, there is less risk sharing. In the limit as $\rho\rightarrow 1$, there is no bond trade and countries live in autarky. This is the essence

10 Appendix A presents and discusses the more general case.
11 For this reason, local approximation around a point other than zero results in worse accuracy. We thank a referee for drawing attention to this point.

of Baxter (1995)'s result that permanent income shocks can lead to a significant reduction in consumption correlations.

Now let us turn to the second order perturbation. Assuming homogeneous utilities of the CRRA form the following result can be proved:

Result 2 *(Quadratic Approximation):* $B_{b,b}=P_{b,b}=P_{b,y2}=P_{b,y1}=C_{b,b}=0$, $P_{y1,y2}=P_{y1,y1}=P_{y2,y2}$ *and* $B_{b,y2}=B_{b,y1}$.

In a second order approximation, the coefficients on the first order terms are the same as in the linear approximation, but now there are also second order terms. As usual, there are no terms involving σ_ε in the second order approximation, but there is a term involving σ_ε^2 correcting for risk as long as the third derivative of the utility is not zero. The bond price function is still independent of the distribution of wealth. Compared to the linear bond price there is only one additional term involving $(\hat{y}_{1,t}+\hat{y}_{2,t})^2$ and the risk correction involving σ_ε^2. The bond policy function still implies a non-stationary law of motion for bonds. An interaction term involving $(\hat{y}_{1,t}+\hat{y}_{2,t})\hat{b}_{1,t-1}$ is added which makes this a stochastic unit root process.[12] That is, the AR(1) coefficient on bonds is now stochastic with a mean of one.

If we assume again that $b_{i,-1}=0$, the second order approximated policies can be written in terms of the first order policies $P^{LIN}(\cdot)$, $B^{LIN}(\cdot)$ and $C^{LIN}(\cdot)$ as follows

$$\hat{p}_t=P^{LIN}(\hat{b}_{1,t-1},\hat{y}_{1t},\hat{y}_{2t})+\frac{1}{2}P_{y,y}(\hat{y}_{1t}+\hat{y}_{2t})^2+\frac{1}{2}P_{\sigma,\sigma}\sigma_\varepsilon^2 \tag{15}$$

$$\hat{b}_{1t}=B^{LIN}(\hat{b}_{1,t-1},\hat{y}_{1t},\hat{y}_{2t})+B_{b,y}\hat{b}_{1t-1}(\hat{y}_{1t}+\hat{y}_{2t})+\frac{1}{2}B_{y,y}(\hat{y}_{1t}^2-\hat{y}_{2t}^2) \tag{16}$$

$$\hat{c}_{1t}=C^{LIN}(\hat{b}_{1,t-1},\hat{y}_{1t},\hat{y}_{2t})+C_{b,y}\hat{b}_{1t-1}(\hat{y}_{1t}+\hat{y}_{2t})+\frac{1}{2}C_{y,y}(\hat{y}_{1t}^2-\hat{y}_{2t}^2) \tag{17}$$

where $P_{y,y}\equiv P_{y1,y1}=P_{y2,y2}=P_{y1,y2}$, $B_{b,y}\equiv B_{b,y1}=B_{b,y2}$, $C_{b,y}\equiv C_{b,y1}=C_{b,y2}$, $B_{y,y}\equiv B_{y1,y1}=-B_{y2,y2}$ and $C_{y,y}\equiv C_{y1,y1}=-C_{y2,y2}$. These illustrate that for purely redistributive shocks where $\hat{y}_{1t}+\hat{y}_{2t}=0$, the consumption and bond policies coincide with the first order approximations and the price function only differs according to a constant correction term related to σ_ε^2.

Although we have no analytical results for the case of a third order approximation, numerical results obtained using Dynare indicate that the two main

12 See Granger and Swanson (1997) and the references therein.

results are also true in that case. That is, even in a third order approximation, debt follows a non-stationary stochastic process and the bond price is independent of the distribution of wealth. We conjecture that this is the case for a local approximation of any order. The following section provides a global approximation method and argues that these two results are not true in this model. The reason is that the behavior of rational, forward-looking agents is affected by the possibility of future binding constraints. Put differently, the behavior of consumption and debt can be affected by the presence of limits even when debt is far from the constraint. This mechanism is naturally captured by a global approximation method but cannot be captured by an approximation method using only local information away from the limits. We conclude that these two properties are a result of the local nature of the approximation.

Local linear approximations tend to work well for some economic models (e.g., stochastic growth model) because the economy being approximated fluctuates around the steady state without drifting too far away and the policy functions in the neighborhood of the steady state are close to linear. Higher order perturbation methods can take into account non-linearity close to steady state but they are still local in nature. The reason why local methods might not work as well in this economy is that they imply non-stationary debt and consumption dynamics, which means that the economy can drift arbitrarily far away from the point of approximation. Debt limits have been ignored based on the supposition that they can be chosen loose enough so that they never come into play. The non-stationarity of debt implies instead that any debt limit, even the natural debt limit, will eventually be violated.[13] These considerations generate some concern regarding the quality of the local linear approximation. On the other hand, it could be that the global, non-linear solution, which takes into account the effects of risk and limits, does imply a stationary distribution for bonds with a small variance around zero as well as policy functions that are indeed approximately linear in the stationary distribution. The extent to which this is true is the subject of investigation in the following section.

4 Numerical results

Under a benchmark calibration, we compute a solution to the model based on a global approximation method and compare it to the linear solution, highlighting the fact that both slopes and levels of the policy functions can differ, even in a

13 This is confirmed numerically in Section 4.4.

neighborhood of the steady state. Subsequently, we compute a measure of the distance between the two solutions and investigate the magnitude of this distance across a range of parameterizations. Finally, we argue that statistics based on simulated time series can markedly differ across solution methods and provide an example of this, namely cross-country consumption correlations.

4.1 Benchmark parametrization

For the benchmark parametrization, we consider a quarterly version of the model and choose the discount factor β to be 0.99, implying an average annual interest rate close to 4%. Period utility is given in (9) and the coefficient of relative risk aversion is chosen to be $\gamma=1$. To obtain the income process, an AR(1) process is fitted to detrended, quarterly US Real GDP data for the period from 1947 up to 2009. The estimates are, approximately, $\rho=0.99$ for the persistence parameter and $\sigma_\varepsilon=0.01$ for the standard deviation of the innovation. Finally, we choose $K=1$ which, given that $\bar{y}=1$, implies a borrowing limit equal to 100% of mean GDP. To put this limit into perspective, Lane and Milesi-Ferretti (2007) provide estimates of net foreign asset positions (NFA) as a percentage of GDP for 145 countries and report the vast majority of countries within the 100% of GDP limit and even the outliers within a 250% upper bound.[14] Imposing $K=1$ ensures debt in the model remains within the empirically relevant range. This is not simply because the model economy hits the limits and cannot diverge further. An important mechanism is the possibility of future binding constraints which keeps debt from hitting the limits in the first place. To the extent that there are real world costs of allowing NFA to drift to levels that are too high, the imposition of ad hoc limits is intended to crudely (exogenously) capture these potential costs. The last part of Section 4.3 discusses several departures from this benchmark calibration including alternative values for the debt limit K. In Section 4.4, we also consider the case of natural debt limits, in which case the model implies excessive variation in NFA.

The combination of low variability in the exogenous shocks, logarithmic utility and limited variability in debt when $K=1$, imply that the additional terms in the quadratic policies (15)–(17) are either zero or very small. As a result, the second order approximation does not offer a significant improvement upon the first order approximation. For this reason, we focus only on the first order

14 Reading from their plots and following their categorizations, all industrial countries fall below 100%, all countries in Asia, Latin America, Emerging Europe, the Middle East and the Commonwealth of Independent states fall below 150% and only a handful of African countries lie in the neighborhood of 200%.

approximation in Sections 4.2 and 4.3 and return to the second order approximation in Section 4.4, where we allow debt to vary more and risk aversion to differ from 1.

4.2 Qualitative comparison

4.2.1 Policy functions

Figures 1–3 plot the policy functions for bonds, $B(b, y_1, y_2)$ and country 1 consumption, $C(b, y_1, y_2)$, as functions of b for specific values of y_1 and y_2. Country 2's policy functions are completely symmetric so they are omitted from the discussion that follows. Solid lines correspond to the policy iteration solution (PI) and dashed lines correspond to the linear solution. Since mean income $\bar{y}=1$, the scale on these graphs can be interpreted as fractions of mean income. In choosing the values of y_1 and y_2, we distinguish between non-redistributive (NR) and purely redistributive (R) realizations of (y_1, y_2). For each case, we plot three policies corresponding to three realizations. For the NR case, these are (y_{min}, y_{min}), $(\bar{y}, \bar{y})$ and (y_{max}, y_{max}) and for the R case, these are (y_{max}, y_{min}), (y_+, y_-) and (y_{min}, y_{max}). Here, y_{min} (y_{max}) denotes the smallest (largest) value for y dictated by the discretization used in the numerical solution and y_+ (y_-) denotes the value of y that is right above (below) $\bar{y}$ in the discretization.

The case of purely aggregate shock realization (NR) is depicted in Figures 1 and 2. The linear bond policy function is a straight line with a slope equal to one going through the origin indicating that no exchange of assets takes place under the linear rule when shock realizations are non-redistributive. This is true regardless of the level of the aggregate endowment, i.e., the linear bond policy functions coincide for all three realizations. The PI policy functions also go through the origin implying no asset trade if initial debt is 0. However, for all realizations of a non-redistributive shock, the policy functions have a slope less than one. These policies lie above the 45 degree line for $b<0$ and below for $b>0$, that is debt is mean reverting. This is visually hard to notice in the top panel of Figure 1 because the effect is quantitatively small, so Figure 2 illustrates this by focusing on a smaller range of b values. In addition, the policies do not coincide for different levels of the aggregate shock. For higher income realizations, the policy lies further above the 45 degree line, i.e., it prescribes a more aggressive repayment of debt. Note, however, that even in the case of the lowest level of income realization some debt repayment still takes place. The smaller slope of the bond policy function (compared to the linear case) translates to a larger slope for the consumption function, as is evident in the bottom panel of Figure 1

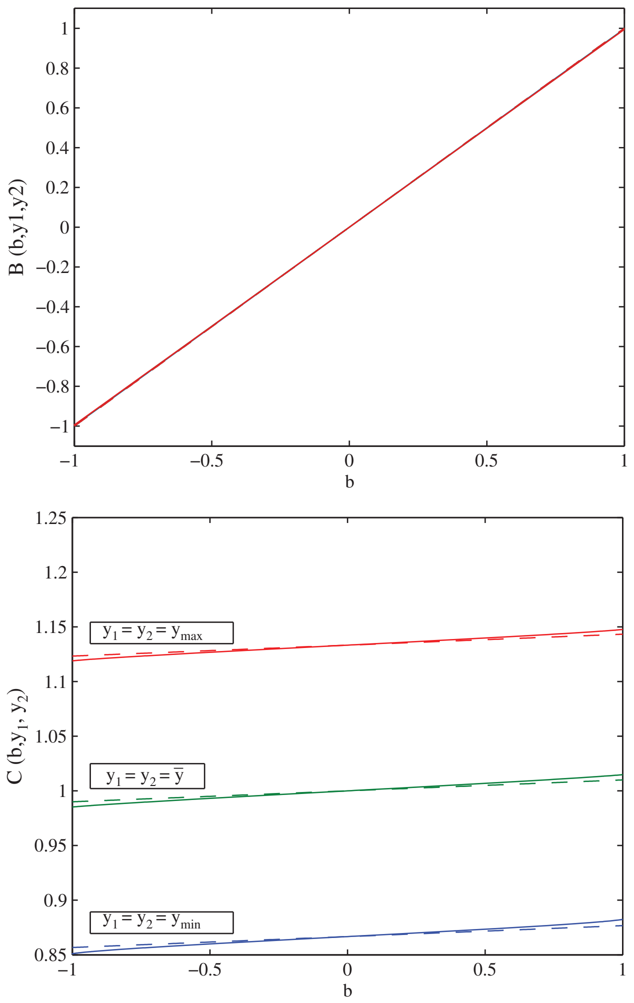

Figure 1 Bond and consumption policy functions for purely aggregate (NR) shocks. Dashed line is for linearized solution and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the ratio of current bond and consumption choices to mean income. The top panel reports the policies for bond and the bottom panel reports those for consumption. For the bond policies, linear and non-linear functions for all three realizations of (y_1, y_2) described in the text are visually indistinguishable. For the consumption policies, labels indicate the three realizations of (y_1, y_2) described in the text.

which depicts the consumption policy functions. That is, consumption responds more aggressively to changes in wealth than what the linearization procedure would suggest.

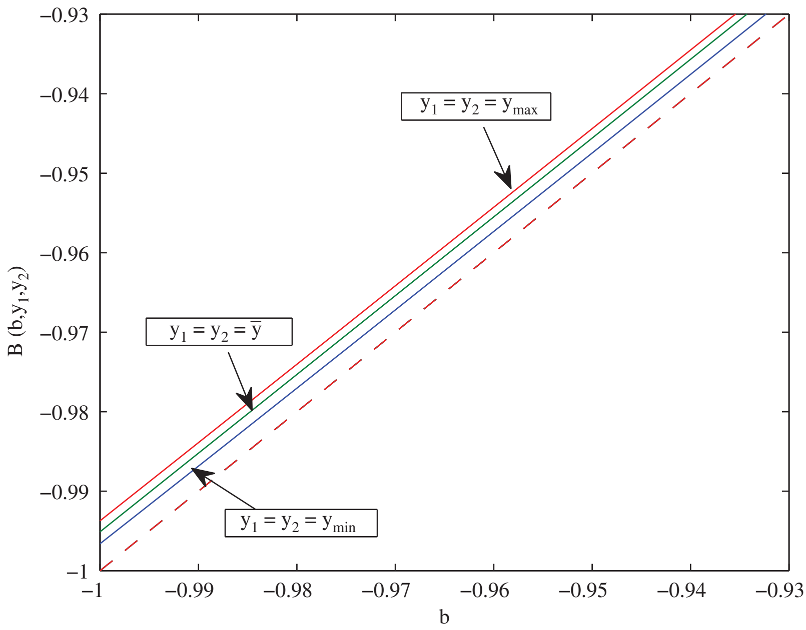

Figure 2 Bond policy functions for purely aggregate (NR) shocks, limited range of b. Dashed line is for linearized solution and solid line for policy iteration solution. The linear policy is the same for all three realizations of (y_1, y_2) described in the text. The non-linear policies are different and less steep than the linear one.

In Figure 3, policy functions for purely redistributive (R) shocks are plotted. Redistributive income shocks do lead to bond trade under linearization, but the linear policy rules still have a slope equal to one. As in the NR case, the PI bond policy functions are less steep than under linearization. In addition, they are closer to the 45 degree line implying less willingness to engage in bond trade for consumption smoothing purposes. Consumption policy functions are more steep, i.e., the marginal propensity to consume is everywhere higher than under certainty equivalence and they prescribe more variation of consumption in response to (redistributive) income shocks. The crucial difference with the NR case is that these effects are now more pronounced and the quantitative difference from the linearized solution is much larger. Importantly, even when debt is 0 and relative income shocks are moderate, linearization underestimates the response of consumption to income by a non-negligible amount. To put it differently, the linearization method predicts significantly different consumption allocations even in the neighborhood of the steady state. The differences

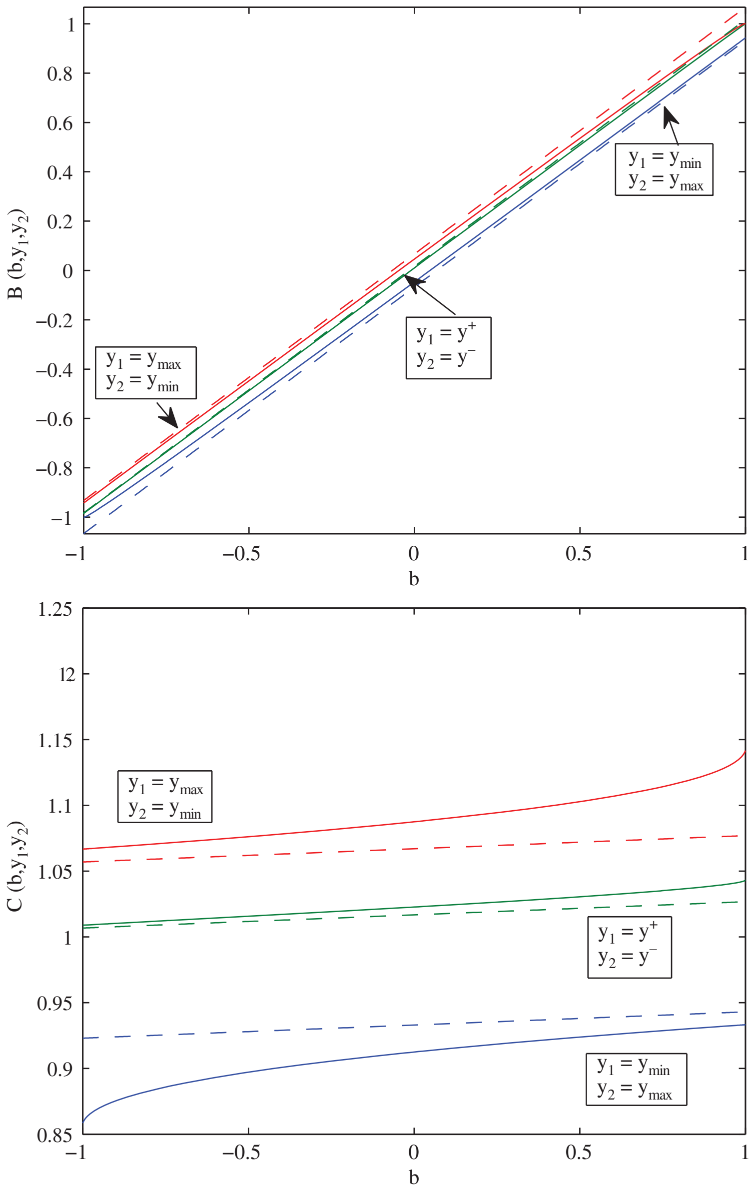

Figure 3 Bond and consumption policy functions for purely redistributive (R) shocks. Dashed line is for linearized solution and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the ratio of current bond and consumption choices to mean income. The top panel reports the policies for bond and the bottom panel reports those for consumption. Labels indicate the three realizations of (y_1, y_2) described in the text.

become larger as one moves away from zero debt and become significant close to the limits.

These departures from the certainty equivalent policy functions obtained under linearization can be understood by first considering the departures that would arise in a partial equilibrium context, i.e., fixing bond prices, and subsequently considering the additional effects arising from general equilibrium bond price variation. In a partial equilibrium context, the consumption-savings literature has attributed these departures from certainty equivalence to precautionary motives arising from two interrelated sources: convex marginal utility and the possibility of future binding constraints. As shown numerically in Zeldes (1989) and proved analytically in Carroll and Kimball (1996), uncertainty combined with convex marginal utility leads to a concave consumption function, with the marginal propensity to consume decreasing in wealth but always higher than the certainty case. Deaton (1991) and Carroll and Kimball (2001) argue convincingly that the combination of uncertainty and borrowing limits has a similar effect.[15] This higher marginal propensity to consume out of wealth is immediately apparent in the bottom panels of Figures 1 and 3. Next, we turn to the behavior of equilibrium bond prices to explain the additional general equilibrium effects.

Figures 4 and 5 plot the equilibrium bond price as a function of inherited debt for the NR and R cases, respectively. Equation (11) shows that, under linearization, the bond price is independent of bond holdings and depends only on aggregate income. As a result the price as a function of bonds is flat and, whenever aggregate income is equal to is mean, the bond price is simply equal to β. In the PI solution, precautionary motives introduce additional effects on the equilibrium bond price, both regarding its average level and regarding its response to the state variables. The level effect is easier to understand in the case of no aggregate uncertainty. In a general equilibrium model with no aggregate uncertainty, Huggett (1993) and Aiyagari (1994) have shown that the equilibrium risk free rate (the inverse of the bond price) will be lower when markets are incomplete than under complete markets. The reason is that precautionary motives push the demand for assets upward and, in equilibrium, the interest rate has to fall to clear the bond market. This explains why, on average, the bond price is higher than in the certainty equivalent case. Note, however, that the presence of aggregate shocks means that this is true only on average, but not at all points in the state space. In particular, when both agents receive a low income realization as in the

15 There can be complex interactions between these two components when they are added simultaneously. These are analyzed in detail in Carroll and Kimball (2001). An early attempt to decompose the effects of convex marginal utility and borrowing constraints in a finite horizon, partial equilibrium model can be found in Xu (1995).

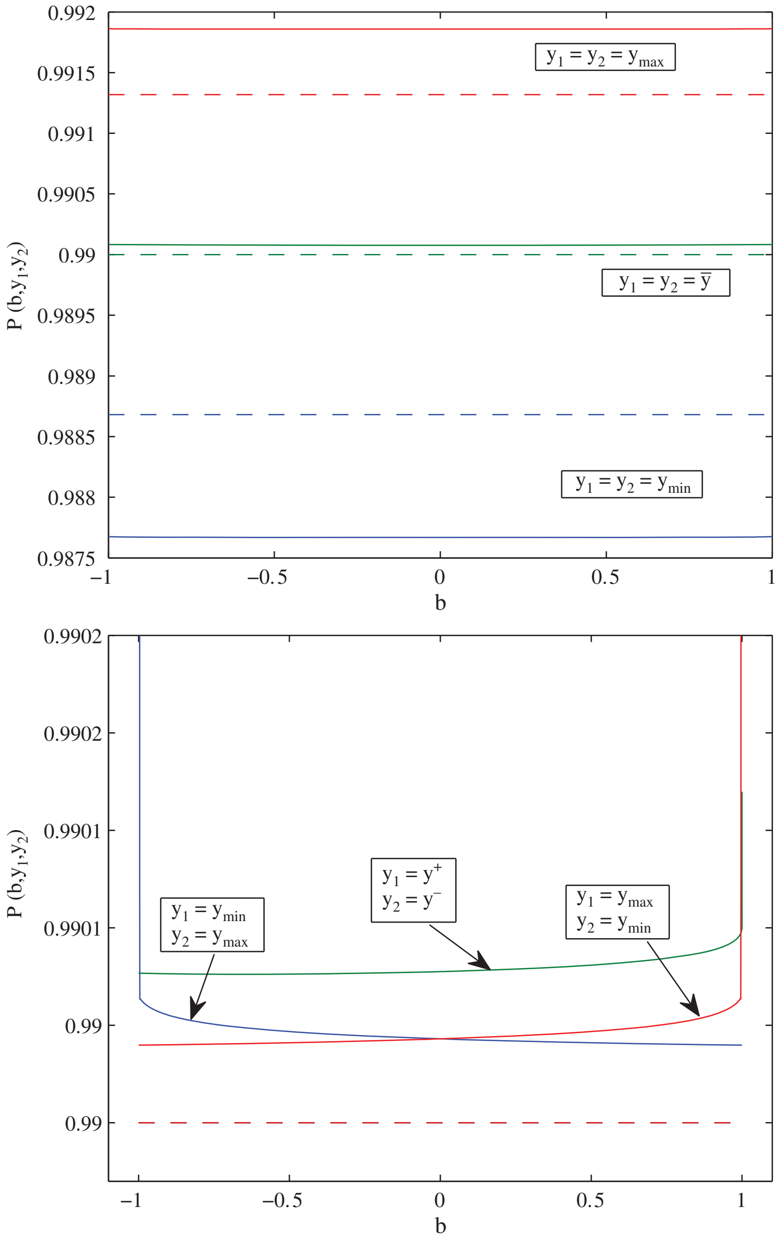

Figure 4 Bond prices.
Dashed line is for linearized solution and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the price level. The top panel reports that for the purely aggregate (NR) shocks and the bottom panel reports that of the purely redistributive (R) shocks.

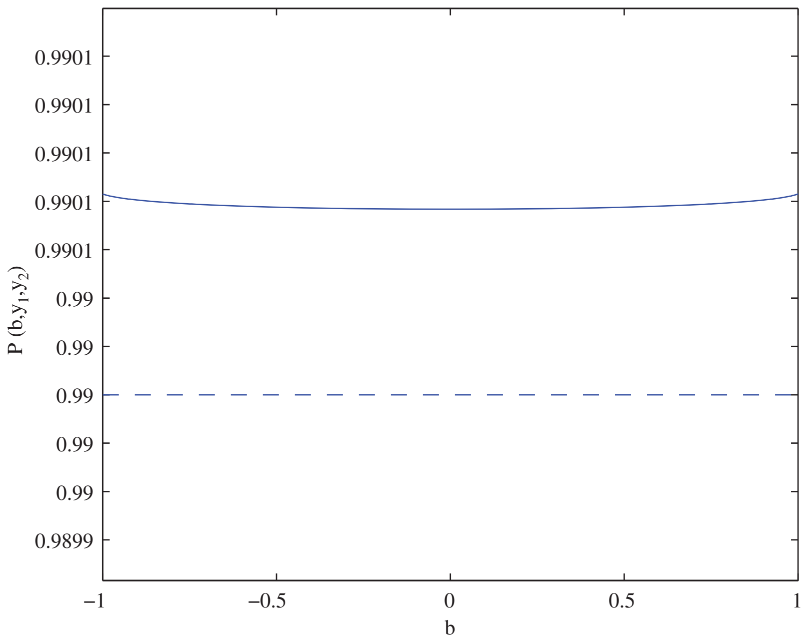

Figure 5 Bond price for purely aggregate (NR) shocks.
Dashed line is for linearized solution and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the price level. Only the case of mean income shown to illustrate the U-shape.

bottom lines of the top panel of Figure 4, the demand for assets falls, pushing the bond price below β. In the PI solution, the drop in the bond price is larger. This just reflects the fact that the variability of the asset prices in response to aggregate shocks is higher than in the linear solution.

Contrary to the linearization case, in the PI solution the bond price is debt-elastic, i.e., it varies with the level of debt b. More precisely, the bond price responds to wealth dispersion and the price functions have an U-shape.[16] As explained in den Haan (2001), the concavity of the (partial equilibrium) consumption function is once again the reason for this. Focusing on Figure 5 where income realizations are equal to the mean for both countries, when country 1 is a lender ($b>0$) it is wealthier than country 2. Holding prices fixed, concavity implies that a marginal increase in b will induce an increase in country 1's consumption demand that is smaller than the decrease in country 2's consumption demand.

16 With NR shocks this effect is small so it is hard to detect visually in the top panel of Figure 4. Figure 5 which focuses on the case of $y_1=y_2=\bar{y}$ makes this more clear.

As a result, country 1 wants to save a larger fraction of this additional wealth than what country 2 is willing to borrow. In order for the bond market to clear, a higher bond price is required. Thus the price is increasing in b for $b>0$. The same argument implies that the price must be decreasing in b for $b<0$. Close to the debt limits, this heterogeneity in the response of desired saving to marginal wealth changes is exacerbated by the unwillingness of the borrower to increase debt up to the limit and the bond price rises even further. At $b=0$, agents are equally wealthy and marginal changes in b have no effect on price. That is why a linear approximation around the steady state $b=0$, yields a bond price that is independent of bond holdings. A similar situation arises in the case of redistributive shocks (R) shown in the bottom panel of Figure 4, the only difference being that the point where agents are equally wealthy (and thus the slope is zero) is not exactly at $b=0$. The reason is that when the representative agent in country 1 (for example) has higher relative income then at $b=0$ this agent is wealthier. Overall wealth, i.e., taking both income and financial wealth into account, is equalized when that agent is also indebted ($b<0$). At the points in the state space where the debt constraint binds, the borrower cannot increase debt at all in response to a negative income realization. In order to induce the lender to hold the same level of assets as before, the bond price has to increase significantly. Note that, as den Haan (2001) points out, the variation in prices resulting from wealth dispersion provides some additional consumption smoothing possibilities since borrowing becomes cheaper exactly when the borrower needs it the most. Quantitatively, the effect of wealth dispersion on bond prices is larger in the case of redistributive shocks, close to the debt limits and, especially, at the debt limits. However, for this benchmark parameterization, this effect is quantitatively small.

Given this general equilibrium variation in bond prices, we now return to complete the discussion on the consumption functions. Bond price variation has implications for the slope of the equilibrium consumption functions shown in bottom panels of Figures 1 and 3 and how these differ from the linear case. Recall that, with exogenously fixed prices, the marginal propensity to consume out of wealth is everywhere higher due to precautionary motives in the face of uncertainty. At points in the state space where the bond price is higher than under linearization, this effect is mitigated. This occurs at points where there is wealth dispersion due to non-zero b or due to redistributive income realizations. It also occurs at points where income realizations are high for both countries. On the other hand, when income realizations are relatively low for both countries the bond price falls below the linear case. In that scenario, the precautionary motive effect on the slope of the consumption function is exacerbated. Quantitatively, the effect of non-redistributive shocks on prices is larger than the wealth dispersion effect.

It should come as no surprise that the local linear approximation, which ignores borrowing limits by construction, will do a poor job close to those limits. The main point to take from this qualitative comparison is that the linear approximation can mistake both the level and slope of the equilibrium bond and consumption policy functions even close to the steady state, i.e., very far from the borrowing limits. The next section illustrates this point further by presenting impulse response functions.

4.2.2 Impulse responses

In the business cycle literature, impulse responses are often used to elucidate the economic mechanisms operative in the model. Baxter (1995) reports impulse responses in a two-country international business cycles model and finds that temporary shocks have permanent effects when markets are incomplete. It might appear that a *local* linear approximation should be enough to obtain a reasonable approximation for the impulse response functions since, after all, impulse responses are inherently local statistics. The economy starts at the steady state and then a single, one-standard-deviation innovation is introduced so that, in principle, endogenous variables never stray too far away from the steady state. This section is intended to show that such an approximation can lead to misleading results even when one only focuses on the behavior close to the steady state, where the approximation is expected to be accurate.

Figure 6 presents impulse responses for c_{1t}, c_{2t}, b_{1t}, and p_t to a positive, one standard deviation (1%) innovation in country 1's income. The units are in percentage deviation of the variable from its steady state, except for b_{1t} which is reported as a fraction of mean income $\overline{y}$. The solid line represents the response in the non-linear equilibrium and the dashed line the response under linearization.

The impulse responses in the two cases share some qualitative properties. The persistent increase in country 1's income leads to an increase in consumption that is smaller than the income increase and a gradual accumulation of assets by country 1. Partial, but not perfect, risk sharing takes place in the sense that country 2's consumption also temporarily increases as a result of the increase in aggregate income, but by less than country 1's consumption increase. As the income shock dies out, consumption in both countries slowly decreases and assets of country 1 keep accumulating. However, the long run response is qualitatively different across the solutions and the short run response is quantitatively also different. The linear equilibrium implies asset accumulation by country 1 for as long as its relative income is higher and results in a permanently higher wealth level for country 1. This also translates to a permanently higher (lower)

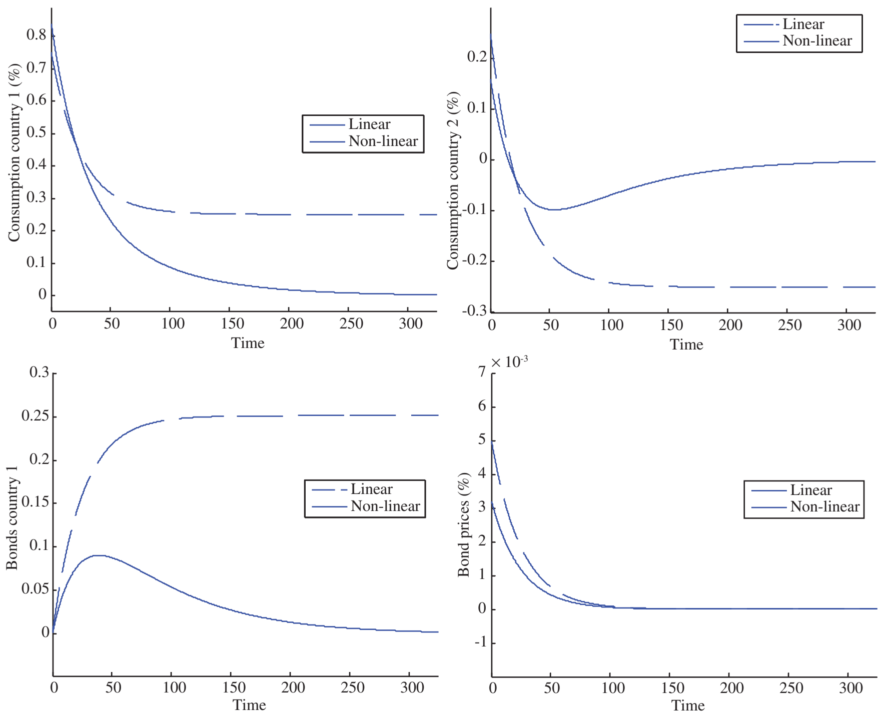

Figure 6 Impulse responses to a 1% increase in y_1.
Clockwise from the top-left panel are, respectively, impulse responses of c_{1t}, c_{2t}, p_t and b_{1t}.

consumption level for country 1 (country 2). By contrast, in the non-linear solution, asset accumulation stops earlier and is eventually reversed. The intuition is that country 2 has incentives to start running down its debt in normal times due to the precautionary motive. As its debt increases, the probability of being constrained in the future increases and this induces a precautionary reduction in debt. The end result is that debt returns to its pre-innovation level and the same is true for consumption. From a quantitative perspective, the initial increase in country 1 consumption is higher (0.84% vs. 0.75%) and the initial increase in country 2 consumption is lower (0.16% vs. 0.25%) in the non-linear equilibrium. That is, the non-linear solution implies stronger sensitivity of consumption to income variations and lower co-movement of consumption across countries. The initial differences are relatively small but they accumulate to be much larger in the long run. It is important to note that limits never even come close to be binding (debt in country 2 remains below 10% of GDP in the non-linear case and 25% of GDP in the linear case, the limit is 100%). Mean reversion results from the possibility of constraints binding in the distant future. Bond prices show very similar behavior across solution methods since they do not inherit the unit root dynamics of debt under linearization.

The qualitative differences observed in the policy functions as well as the impulse responses are present under any choice of parameters. On the other hand, whether the differences are quantitatively significant will depend on the parameterization. In the following section we provide a quantitative comparison across a range of parameterizations. We use two insights from the qualitative analysis to guide this. First, approximation errors under linearization might be small at a given point, but accumulate to significant size in a simulation. Second, linear policy functions for assets and consumption give the certainty equivalent solution, meaning that the standard deviation of innovations, σ_ε, does not affect those functions. Since uncertainty will matter for the non-linear solution, this indicates that this parameter can matter significantly for any differences.

4.3 Quantitative comparison

4.3.1 Policy function differences

We use a simple comparison of the differences in allocations and prices between the linear and PI solutions across all points in the state space to gauge the success of the linear approximation. Since the non-linear solution is also an approximation to the true solution, this measure provides information about the success of linearization only to the extent that the PI solution is an accurate approximation to the true solution. In this section, we assume this to be the case. In Appendix C, we show that this is indeed the case by looking at the Euler errors.

For each variable $x=c_1$, c_2, b_1, p and at each point in the state space, we compute the absolute difference between the value x^{PI} given by the policy iteration algorithm and the value x^L obtained using the linear policy functions. We express this difference as a percentage of x^{PI} for $x=c_1$, c_2, p and as a fraction of mean income for bonds b, and denote them by Δx

$$\begin{aligned} &\Delta x \equiv 100\left|\frac{x^L - x^{PI}}{x^{PI}}\right| \text{ for } x=c_1,\, c_2,\, p \\ &\Delta b \equiv 100\,|b^L - b^{PI}| \end{aligned} \tag{18}$$

For each variable x, the left panel of Table 1 reports the maximum and weighted average value of Δx across all points in the state space. The latter measure uses the probability of different points in the state space implied by the stationary distribution of the PI solution, to weigh the values of Δx accordingly.

Table 1 Policy function differences: linear vs. policy iteration solutions.

	Whole state space		Binding points excluded	
	Max	**Mean**	**Max**	**Mean**
Δc	7.51%	1.52%	7.24%	0.828%
Δb	6.70%	1.55%	6.30%	0.836%
Δp	0.182%	0.031%	0.103%	0.018%

The values for Δx are similar for allocations of consumption and debt and an order of magnitude smaller for bond prices. Focusing on consumption differences, the maximum difference is more than 7% of consumption and the average is equal to 1.52% of consumption. The maximum values occur exactly at the limits. These differences are large, but in one sense this is to be expected since the linear approximation assumes no limits. We are also interested in how large these differences are away from the limits. The right panel of Table 1 reports the same statistics but instead of using all points in the state space, it excludes points where the limits bind. To obtain the weighted average in this case, we truncate the distribution accordingly. Away from the limits, differences look smaller as expected, but still very large. The maximum is still above 7% of consumption, the average is about 0.83%. We conclude that, at least for the benchmark calibration, linearization can lead to a significantly different solution even at points where the limits don't bind. This is especially true for allocations, but less so for prices. In the last part of this section, we look at how these differences vary with model parameters, focusing on consumption allocations only.

4.3.2 Cross-country consumption correlations

The well-known consumption correlation puzzle refers to the inability of a standard two-country business cycle model to produce consumption correlations that are consistent with those observed in the data. Complete markets models produce consumption correlations that are very high and, in particular, higher than income correlations, a result that is contrary to the empirical observation. This is a reflection of risk sharing and has been found to be robust to significant reductions in the menu of assets that can be traded. Using a linearization method in a model where only a risk free bond is traded and production is endogenous and subject to TFP shocks, Baxter (1995) finds that introducing market incompleteness has a small to moderate effect on consumption correlations except when exogenous shocks are permanent (ρ=1).

Judd (1992) points out that even if a linear approximation yields a reasonable approximation for allocations, that does not necessarily imply a good approximation of second moments such as consumption correlations. In this section, we compare the cross-country consumption correlations implied by this incomplete markets model under the two solution methods across a range of values for the income persistence parameter ρ.

Implied consumption correlations are computed as follows. For each solution, the economy is simulated for 50 periods, the HP filter (with a smoothing parameter 1600) is then applied to the series of consumptions and the correlation is computed. This is repeated 20,000 times and the average of the correlation over these 20,000 replications is obtained. We keep the length of each simulation short in order to avoid periods of binding constraints and thus give the linear solution method the best chance to match the correlations of the non-linear method. We have also experimented with longer simulations, in which case limits sometimes bind for the non-linear simulation and, as a result, consumption correlations are lower in the non-linear case (but the same for the linear).

In the model under consideration, the cross-country consumption correlation under complete markets is exactly 1. For levels of persistence <0.9, the incomplete markets model produces correlations that are almost equal to 1 regardless of the solution method. Thus, the main qualitative idea in Baxter (1995) is confirmed even when using a non-linear method; incomplete markets alone cannot change this correlation for low levels of income persistence.[17] We focus on the range $\rho \in [0.9, 0.995]$ and present the implied correlations in Table 2.

Income correlation is exogenous in this paper and equal to 0. All correlations shown in Table 2 are positive and therefore higher than income correlation. For

Table 2 Consumption correlations.

ρ	Consumption correlation		
	Linear	**PI**	**PI (σ_ε adjusted)**
0.9	0.98	0.97	0.83
0.95	0.94	0.87	0.71
0.98	0.80	0.59	0.52
0.99	0.61	0.38	0.38
0.995	0.46	0.30	0.33

17 This is conditional on modelling market incompleteness as an exogenous restriction on the menu of assets traded. Kehoe and Perri (2002) show that endogeneizing market incompleteness through imperfect enforceability can go a long way towards explaining the data in this dimension.

the benchmark calibration, the consumption correlation is 0.38 as opposed to 0.61 under linearization. Put differently, the linear approximation leads to a significant overestimation of this correlation. The reason is straightforward. Consumption correlations depend on the extent to which idiosyncratic income fluctuations can be insured against. In turn, this depends on the ability to use the risk free bond available in this model to insulate consumption from income shocks. The linear approximation overestimates this ability because it ignores the precautionary motives that limit agents' willingness to accumulate debt. In particular, debt constraints do not only restrict debt accumulation at the constraint but also away from it. Indeed, due to our use of short run simulations, limits never become binding for the non-linear economy in this experiment. Thus, we conclude that the *possibility of future binding constraints* introduces limitations in debt accumulation even away from the constraints and this alone implies lower consumption correlations.

Changing persistence while keeping the income innovation variance σ_ε^2 fixed also introduces changes in the unconditional variance of income $\sigma_y^2=\sigma_\varepsilon^2/(1-\rho^2)$. More specifically, as persistence ρ is increased, income variability also increases. An alternative experiment would be to adjust σ_ε together with ρ in such a way as to maintain the unconditional variance of the income process constant. The result of this experiment is shown in the last column of Table 2, labeled "σ_ε adjusted." At values of ρ below the benchmark 0.99, the standard deviation of innovations σ_ε has to be increased in order to maintain the unconditional variance of income constant. This, in turn, implies more uncertainty and lower consumption correlations than when σ_ε is kept fixed.

To summarize, consumption correlations are decreasing in ρ and also decreasing in σ_ε (for fixed ρ) as expected. The linear approximation overestimates these correlations, especially when persistence ρ is high. If longer simulations were used, introducing periods of binding constraints, this overestimation would be even more severe.

4.3.3 Sensitivity to parameter choices

Although the main qualitative aspects of a linear approximation and its differences from the non-linear solution are true regardless of the parameterization, the quantitative importance of such differences will depend on the specific calibration. In this section, we consider a wide range of values for the model's parameters and aim to provide some guidelines as to which of those parameters are most important.

Table 3 reports the maximum and mean value of Δc as defined in equation (18), i.e., the percentage difference in consumption policy functions between the linear and policy iteration solutions, for a range of values for the preference parameters γ and β, the debt limit K and the exogenous income process parameters ρ and σ_ε. In each of those cases, we start from the benchmark calibration and change only one of the parameters as indicated on Table 3. Table 4 provides the same information but excluding all points in the state space where the limits bind.

As expected, the differences become smaller as the borrowing limit is relaxed (higher K). In the benchmark where the limit is 100% of GDP, the maximum consumption deviation is approximately 7% and the average is 1.5% (or 0.8% if we exclude binding points). Increasing the limit to 400% does not significantly reduce the maximum deviation, but it does reduce the frequency with which the model visits points close to or at the limit in the stationary distribution. As a result, the average deviation drops to 0.9% (or 0.5% excluding binding points). The maximum does not drop significantly even in the case of unrealistically loose limits of 1000% of GDP. The average keeps dropping as the possibility of binding constraints falls towards zero, but even at this level we find differences of 0.5% (0.3% excluding binding points). The case of natural debt limits, which are even looser, is discussed in the following section.

Table 3 Maximum and mean values of Δc under alternative parameterizations (whole state space).

$K/\bar{y}$	**0.5**	**1**	**2**	**3**	**4**	**10**
Max	7.58	7.51	7.39	7.28	7.18	6.73
Mean	1.72	1.52	1.26	1.07	0.94	0.54
σ_ε	**0.005**	**0.01**	**0.015**	**0.02**	**0.025**	**0.03**
Max	3.48	7.51	12.22	17.79	24.48	32.66
Mean	0.63	1.52	2.49	3.50	4.58	5.72
ρ	**0**	**0.5**	**0.9**	**0.95**	**0.99**	**0.995**
Max	2.68	3.31	5.37	6.09	7.51	5.92
Mean	0.004	0.013	0.23	0.54	1.52	1.38
$\rho \& \sigma_\varepsilon$	**0**	**0.5**	**0.9**	**0.95**	**0.99**	**0.995**
Max	34.52	25.22	19.10	14.84	7.51	3.98
Mean	0.26	0.57	1.53	1.77	1.52	0.92
β	**0.9**	**0.93**	**0.95**	**0.97**	**0.99**	**0.995**
Max	1.23	1.76	2.40	3.68	7.51	10.06
Mean	0.17	0.25	0.36	0.62	1.52	2.18
γ	**0.5**	**1**	**2**	**3**	**5**	**10**
Max	7.62	7.51	7.29	7.08	6.66	5.63
Mean	1.53	1.52	1.50	1.48	1.41	1.21

Table 4 Maximum and mean values of Δc under alternative parameterizations (excluding binding points).

$K/\bar{y}$	0.5	1	2	3	4	10
Max	7.32	7.24	7.14	7.04	6.95	6.52
Mean	0.90	0.83	0.70	0.61	0.54	0.34
σ_ε	**0.005**	**0.01**	**0.015**	**0.02**	**0.025**	**0.03**
Max	3.37	7.24	11.76	17.13	23.57	31.40
Mean	0.35	0.83	1.33	1.85	2.40	3.01
ρ	**0**	**0.5**	**0.9**	**0.95**	**0.99**	**0.995**
Max	0.49	1.63	4.73	5.67	7.24	5.75
Mean	0.004	0.012	0.20	0.42	0.83	0.65
$\rho \& \sigma_\varepsilon$	**0**	**0.5**	**0.9**	**0.95**	**0.99**	**0.995**
Max	3.77	11.48	16.51	13.83	7.24	3.86
Mean	0.26	0.54	1.22	1.28	0.83	0.44
β	**0.9**	**0.93**	**0.95**	**0.97**	**0.99**	**0.995**
Max	1.01	1.53	2.13	3.41	7.24	9.78
Mean	0.086	0.11	0.15	0.26	0.83	1.33
γ	**0.5**	**1**	**2**	**3**	**5**	**10**
Max	7.34	7.24	7.06	6.84	6.45	5.45
Mean	0.83	0.83	0.82	0.81	0.78	0.71

The difference between the linear and the non-linear solutions increases monotonically in σ_ε. This monotonicity arises from the combination of two separate effects that the variance of the exogenous income process has on bond trade. First, when σ_ε is higher, more bond trade is required to achieve the same level of consumption smoothing. Second, higher uncertainty increases the precautionary motive making agents less willing to accumulate significant debt. This leads to a frequency of binding debt limits in equilibrium that increases with σ_ε for low values but eventually starts decreasing as the precautionary motive becomes stronger at high levels of σ_ε. This indicates that for low σ_ε, the increase in Δc is mainly due to the increasing possibility of binding debt limits which is ignored in the linear solution, but for high σ_ε this is mainly due to the effects of risk on prudence. Notice that the difference becomes extremely large for high levels of σ_ε. The fact that in the benchmark case, where σ_ε=0.01, we find relatively moderate deviations is a direct result of a relatively small variance in the calibrated GDP process.

For a given level of σ_ε, increasing the persistence ρ of the income process has a non-monotonic effect on the difference between linear and non-linear solutions. Starting from ρ=0 and increasing ρ gradually, initially increases the difference but eventually, above ρ=0.99, that difference decreases. This is true for the maximum over the state space as well as for the mean. The non-monotonicity reflects several counteracting effects of persistence on equilibrium allocations. For very low levels of persistence, the likelihood of receiving several bad income realizations in a row is very low, which makes it very unlikely that there will be a substantial amount of debt accumulation. As a result, debt limits almost never bind in equilibrium. Thus, when $\rho=0$, the difference between linear and non-linear solutions simply reflects the effect of risk on precautionary motives, which is absent in the linear solution. As persistence increases, the likelihood of hitting the limit increases and feeds back into equilibrium consumption behavior in the non-linear case even away from the limits. As a result consumption allocations start diverging from the linear case. However, as ρ tends to one, income shocks become almost permanent. It is well known that, at the limit, there is no bond trade and the two countries live in autarky, since there is no scope in consumption smoothing in the face of permanent shocks. Given the fact that this is true for both the linear and non-linear equilibrium, for very high persistence the difference between the two solutions starts falling.

The row labeled "$\rho\&\sigma_\varepsilon$" presents the alternative experiment where, as ρ increases, the unconditional income variance is kept fixed by appropriately adjusting σ_ε. Thus, reading from left to right, the value of ρ increases and, simultaneously, the value of σ_ε decreases. The maximum difference, which occurs exactly at the debt limit, decreases with ρ. This simply reflects the higher conditional variance of income since, at the limit, all of the shock is absorbed by consumption in the non-linear case but not so in the linear case where there are no limits. Comparing to the case where only ρ is adjusted, and focusing on the iid case, notice that the average difference is larger here (0.26% vs. 0.004%). With iid shocks, the limits play almost no role, but in the current experiment the amount of risk is substantially larger since σ_ε is much larger to maintain the same level of σ_y. The mean difference increases with ρ due to the higher possibility of binding constraints and despite the simultaneously decreasing σ_ε. As before, there is a non-monotonicity that kicks in at high values of ρ where bond trade becomes less useful and limits start binding less. For $\rho=0.995$, the mean difference is lower than in the benchmark. Comparing again with the case of only ρ changing, the difference is now smaller (0.92% vs. 1.38%) since the amount of risk is now smaller.

Turning to preference parameters, the differences computed are increasing in the discount factor β. Recall that, as β tends to 1, the linear consumption functions look like the complete markets solution in the sense that they become functions of aggregate income only and the effects of idiosyncratic income or wealth dispersion vanish. In the presence of borrowing limits, this level of risk sharing is not feasible since it would imply excessive variation of assets and lead to bond limits becoming binding. This aspect is reflected in the non-linear solution but not in the linear one, which is why the difference becomes larger as β increases. To put it differently, the possibility of future binding limits has a stronger effect on behavior when the future is discounted by less. Regarding the risk aversion parameter γ, we find little variation in these differences as γ is changed. Recall that the linear consumption policy functions are independent of γ, so another way to put this is that the non-linear consumption functions do not change a lot with γ. With higher risk aversion, agents would like to use more bond trade to insulate consumption from shocks, but the debt limits significantly restrict their ability to do so. In fact, stronger precautionary motives would imply less willingness to build up on debt and get close to the constraints. As a result, consumption is only slightly better smoothed when risk aversion is higher. Since the linear solution allows for unlimited use of bonds to smooth consumption and achieves better smoothing, this explains why the consumption policies are closer to the linear ones with higher γ. In fact, most of the extra smoothing comes from bond price variation, not reported here in the interest of brevity. Bond prices vary more when risk aversion is high. This variability is higher in the non-linear case and more sensitive to the value of γ, so bond prices provide better insurance than in the linear case. Indeed, the differences in bond prices between the two solutions are much more sensitive to γ than the differences in allocations and can become very significant with high risk aversion.

To summarize, the linear approximation deviates significantly from the non-linear solution even away from the limits. The deviation in consumption allocations increases in income variance and in the patience parameter, decreases in the size of the limit and decreases slightly in risk aversion. Higher risk aversion increases instead the deviation in equilibrium bond prices. Finally, the effect of income persistence is non-monotonic, with the deviation in consumption allocations increasing in income persistence for moderate levels of ρ but decreasing for very high levels of ρ.

4.4 The case of natural debt limits

Since perturbation methods ignore debt limits, it seems reasonable that they will provide a better approximation for the model with very loose debt limits. The loosest possible debt limit is the natural debt limit, defined as the level of debt that can be repaid under any possible future realization of the income shocks without letting consumption become negative. Computing an economy under the natural debt limit is a non-trivial task in our setup because of the presence of aggregate uncertainty. The presence of aggregate uncertainty implies that bond prices fluctuate stochastically even in the long run, so we do not have an explicit formula for the natural debt limit. Instead, we approximate the limits numerically using a guess-and-verify approach. With the additional assumption of a CRRA utility which satisfies Inada conditions, strictly positive consumption is always optimally chosen, meaning that the natural debt limit should, in theory, not bind in equilibrium. We obtain a solution of the economy with a natural debt limit by computing economies with ad hoc limits which we progressively relax. For our benchmark calibration, we have computed an economy with K=80.7 implying a level of debt that is more than 8000% of mean GDP. With this value of K, the equilibrium prescribes a value of consumption equal to 0.001 when inherited debt is at the limit. Increasing K further causes a breakdown in the computational algorithm because consumption becomes zero or negative at some point in the state space. We conclude that this value of K is a reasonable approximation for the natural debt limit.

Figure 7 presents long run simulations produced using the policy iteration solution under the natural debt limit as well as simulations using the first and second order perturbation solutions. The bold line corresponds to the PI simulation, the thinner, solid line corresponds to the linear and the dashed line corresponds to the second order approximation. The figure illustrates the fact that, for the perturbation methods, debt is non-stationary and this non-stationarity is inherited by consumption. Debt exceeds the natural debt limit and can grow to arbitrarily large levels. As a result, consumption becomes negative and can grow to arbitrarily large and negative levels. This is not true when the global solution method is used. Even though debt and consumption are very persistent, they remain bounded. The debt limit never binds, the highest debt level in a simulation of 500,000 periods is 55, which is still far from the debt limit imposed and consumption remains bounded above zero.[18] One implication is that in long simulations, consumption paths under the perturbation method solutions ultimately diverge away from those implied by the global solution method. Another implication is that, even under the global solution, debt can become unrealistically large. This is the reason why we chose to focus on stricter limits in the preceding sections. However, the case of the natural debt limit offers some additional insights into the properties of equilibrium consumption and bond prices.

Figures 8 and 9 present policy functions for bond prices and consumption respectively. Solid lines correspond to the global solution, dashed lines correspond to the first order and dotted lines correspond to the second order approximation. The policy functions for the perturbation methods are not affected by the choice of debt limits and are therefore identical to the ones previously analyzed. Regarding the policy functions obtained using policy iteration, Figure 8 confirms

18 The figure only shows the initial part of the simulation for expositional clarity.

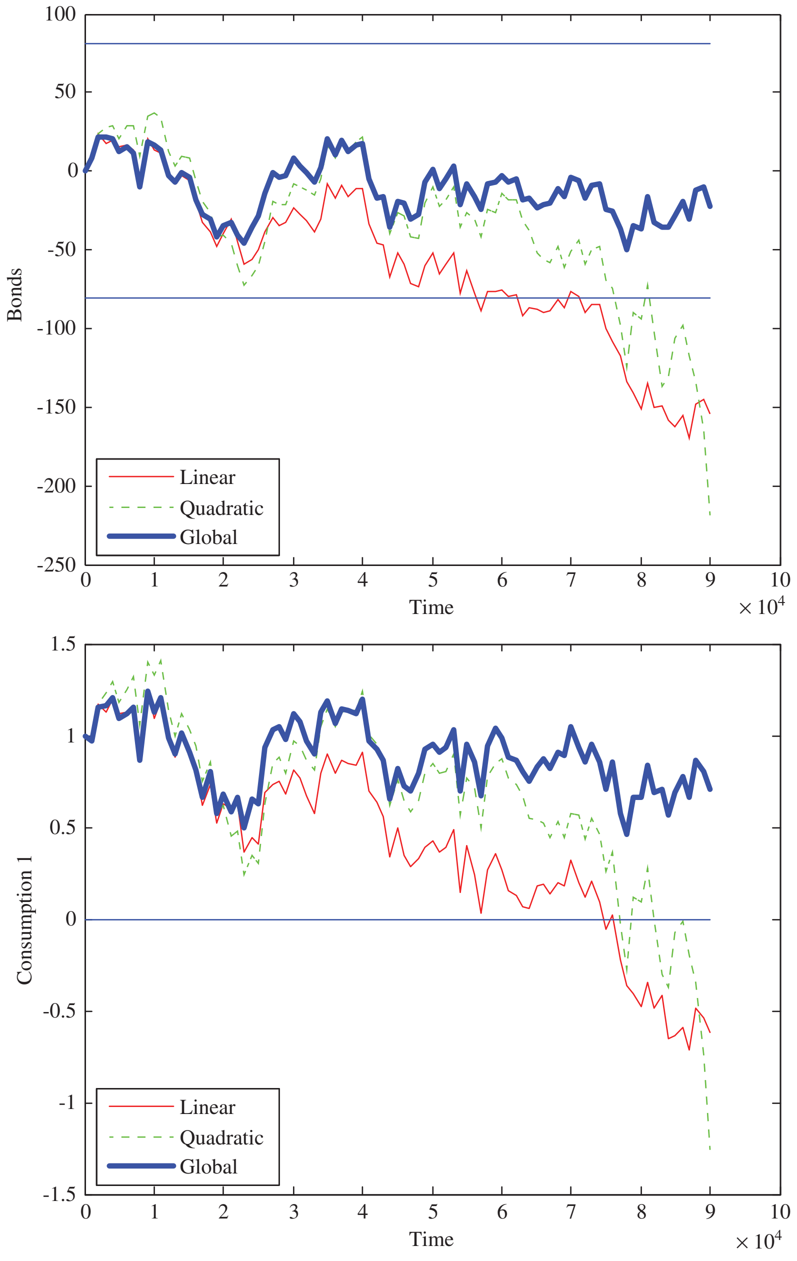

Figure 7 Long simulation for b_{1t} and c_{1t} under three solution methods.
The solid line represents the linear approximation, the dashed line represents the quadratic approximation and the thicker, bold line represents the policy iteration method. The top panel reports the simulation for b_{1t} and the bottom panel reports that of c_{1t}.

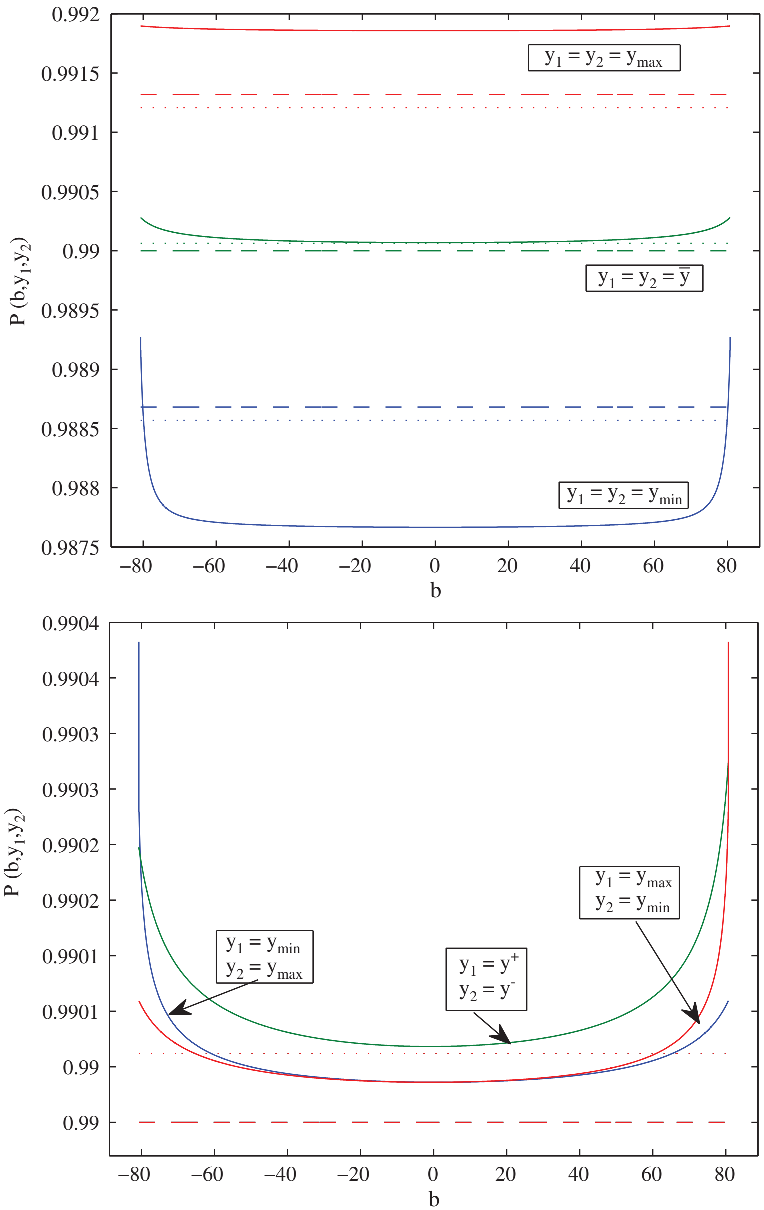

Figure 8 Bond prices under natural debt limits.
Dashed line is for first order, dotted line for second order and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the price level. The top panel reports that for the purely aggregate (NR) shocks and the bottom panel reports that of the purely redistributive (R) shocks.

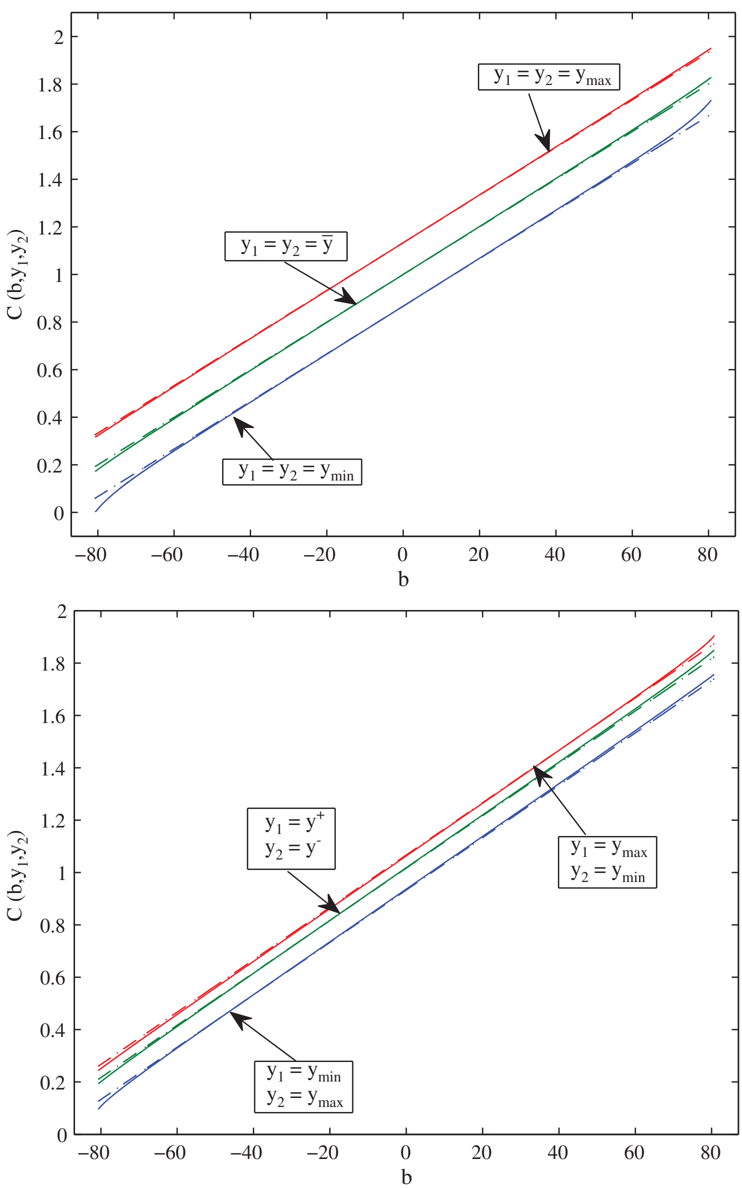

Figure 9 Consumption policies under natural debt limits.
Dashed line is for first order, dotted line for second order and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the ratio of current consumption choice to mean income. The top panel reports that for the purely aggregate (NR) shocks and the bottom panel reports that of the purely redistributive (R) shocks. Labels indicate the three realizations of (y_1, y_2) described in the text.

the U-shape in bond prices. The effect is now more pronounced as the level of debt, and hence wealth inequality, can be much larger. The figure also illustrates the correction for risk that the second order approximation incorporates as well as the fact that bond prices are independent of bond levels even in the second order perturbation method. The consumption policy functions in Figure 9 are steeper than the linear ones. For the cases of purely aggregate and purely redistributive shocks, the consumption policies under a second order approximation coincide with the first order ones. This follows from noticing that the additional terms introduced by the second order approximation in equations (16) and (17) are zero. For purely redistributive shocks both of these terms are zero regardless of the values of parameters. For purely aggregate shocks, the second term is again zero. The first term is also zero, but only because $\gamma=1$ implies $C_{b,y}=0$. With logarithmic utility, substitution and wealth effects cancel out and the consumption function slope under a second order approximation coincides with the one under linearization. This observation also indicates that the case $\gamma=1$ could be misleading regarding the accuracy of the linear approximation. Indeed, Figure 10 presents consumption policy functions under purely aggregate shocks for a model with $\gamma=3$. Now, the price differences translate to differences in the slope of consumption. The consumption function is steeper for low aggregate shocks (price) and flatter for high aggregate shocks (price). Under the linear approximation, the slope of the consumption function is invariant to the state of the economy. The second order approximation, on the other hand, incorporates this effect and, as a result, stays much closer to the policies implied by the global solution.

Table 5 reports policy function differences between the policy iteration solution and the perturbation methods. Comparing to Table 1, these differences are smaller when the debt limits are loose. However, the differences are still substantial. Focusing on consumption the maximum difference is close to 6% and the average difference 0.2%.[19] Notice also that moving from a first order to a second order approximation does not significantly reduce these differences. Consistent with the preceding discussion, with $\gamma=3$ the differences are larger and the second order approximation yields a more substantial improvement. Table 6 in Appendix C.1 shows Euler errors and confirms that these differences are not due to inaccuracy in the policy iteration method. Instead, it is the perturbation methods that remain relatively inaccurate even under a natural debt limit.

The case of natural debt limits provides some support for the use of linear approximations with regard to cross-country correlations. Even though the

19 Computing Δc as in equation (18) leads to extremely large differences, because c^{PI} is very close to zero at some points in the state space. In this section, we choose to report consumption differences as percentages of steady state consumption, i.e., we divide by $\overline{c}=1$ instead of c^{PI}.

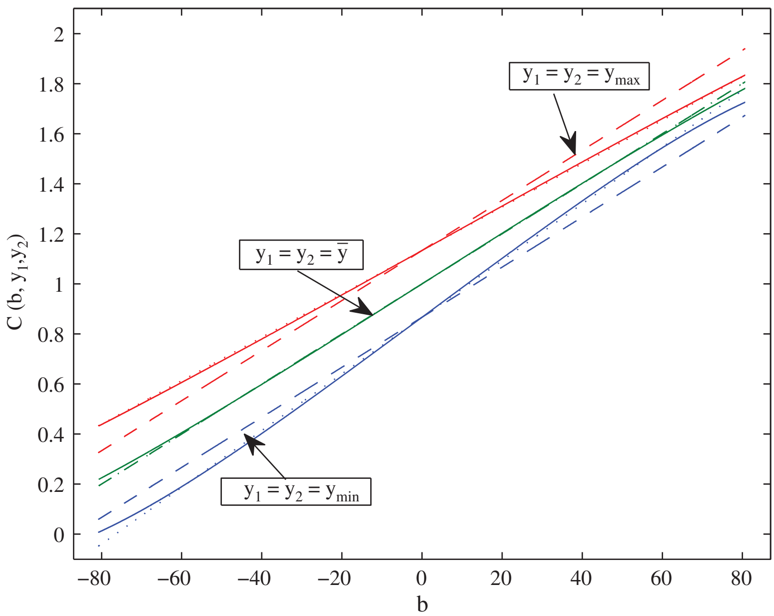

Figure 10 Consumption policies under natural debt limits and γ=3 , purely aggregate (NR) shocks. Dashed line is for first order, dotted line for second order and solid line for policy iteration solution. The *x*-axis shows the ratio of inherited bonds to mean income, the *y*-axis shows the ratio of current consumption choice to mean income. Labels indicate the three realizations of (y_1, y_2) described in the text.

linear solution implies non-stationary consumption paths, HP filtering removes this non-stationarity. Once the series are HP filtered, consumption correlations under the linear approximation provide a good approximation for the consumption correlation obtained using the policy iteration solution when γ=1. Increasing the coefficient of relative risk aversion has no effect on the linear consumption

Table 5 Policy function differences between perturbation and policy iteration solutions, with *K*=80.7.

	Case: γ=1				Case: γ=3			
	First order		Second order		First order		Second order	
	Max	Mean	Max	Mean	Max	Mean	Max	Mean
Δc	5.87%	0.211%	5.87%	0.207%	10.60%	0.466%	5.781%	0.194%
Δb	21.69%	0.845%	31.07%	0.482%	134.43%	0.970%	153.20%	0.762%
Δp	0.317%	0.018%	0.324%	0.017%	1.96%	0.064%	1.96%	0.053%

functions, but it does affect the true consumption functions. The cross-country correlation of consumption is 0.60 under linearization for any value of γ. The corresponding numbers under policy iteration are found to be 0.60, 0.58, 0.56 and 0.55 for the cases of a risk aversion of 1, 3, 5, and 10, respectively. Thus, the approximation deteriorates with risk aversion but remains reasonably close to the true consumption correlations. Interestingly, a second order approximation makes matters worse. Consumption and debt follow a stochastic unit root under the second order approximation and we find, numerically, that HP filtering does not render those processes stationary.

5 Conclusion

In summary, this paper contributes to three distinct strands of economic literature. First, it provides an application of the concept of precautionary savings and associated departures from certainty equivalence to an international business cycles setting. The result is that debt exhibits mean reverting dynamics and, as a direct consequence, consumption is more responsive to own income than under certainty equivalence. Second, it contributes to the discussion on the consumption correlation puzzle by showing how market incompleteness can provide a significant reduction in consumption correlations without resorting to endogenous market incompleteness or unrealistically stringent borrowing constraints. It has been shown that a simple model of exogenous market incompleteness could generate low cross-country consumption correlations as long the variation in net foreign asset positions is reasonably restricted. Third, it contributes to the computational methods literature by comparing perturbation methods of first and second order to a global approximation method. In this sense, the paper can be thought of as an analysis of the approximation errors incurred when using local approximation methods in this setting. It has been shown that important, qualitative features of this model can be obscured by the use of perturbation methods.

We have focused mainly on exogenous borrowing limits that are tighter than the natural borrowing limit. The reason is that the model with natural debt limits implies unrealistically large variation in net foreign asset positions. Tighter limits ensure the model predicts net foreign asset positions within the observed range for developed economies. Importantly, we have focused most of our analysis on behavior far from the limits. We have argued behavior can be significantly affected by the possibility of future binding limits even when those limits are far from binding. These effects are not captured by perturbation methods, partly because those build upon the solution to a deterministic model where limits are irrelevant.

Acknowledgments: We have greatly benefited from suggestions, discussions and comments by Wouter den Haan, Albert Marcet and Morten Ravn. We also wish to acknowledge helpful comments made by Eva Carceles-Poveda, Tom Muench, Tom Sargent, Andrew Scott, Kjetil Storesletten and participants at the X Workshop on Dynamic Marcoeconomics in Vigo, Spain.

Alexis Anagnostopoulos wishes to acknowledge financial support from the European Commission, MAPMU network under contract HPRN-CT-2002-00237.

Appendix A Perturbation methods

In this appendix, we provide the first order approximation of the policy functions in the more general case where the point of approximation is chosen to be $\bar{b}_1 \neq 0$ and the utility function is left unspecified. All derivations are omitted but available upon request.

The first order approximation is given by

$$\begin{aligned}
\hat{b}_{1t} &= B^{LIN}(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}) = \hat{b}_{1,t-1} + B_{y_1}\hat{y}_{1t} + B_{y_2}\hat{y}_{2t} \\
\hat{p}_t &= P^{LIN}(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}) = P_y(\hat{y}_{1t} + \hat{y}_{2t}) \\
\hat{c}_{1t} &= C^{LIN}(\hat{b}_{1,t-1}, \hat{y}_{1t}, \hat{y}_{2t}) = (1-\beta)\hat{b}_{1,t-1} + C_{y_1}\hat{y}_{1t} + C_{y_2}\hat{y}_{2t}
\end{aligned}$$

where B^{LIN}, C^{LIN} and P^{LIN} represent the policy functions for country 1's bonds, consumption and for equilibrium prices, respectively. The coefficients B_{y1}, B_{y2}, P_y, C_{y1} and C_{y2} represent the derivative of the corresponding true function with respect to the variable indicated in the subscript, evaluated at the steady state. These derivatives are given in terms of the parameters in what follows

$$P_y \equiv P_{y_1} = P_{y_2} = -\beta(1-\rho)\frac{\dfrac{u_{cc}(\bar{c}_1)}{u_c(\bar{c}_1)}\dfrac{u_{cc}(\bar{c}_2)}{u_c(\bar{c}_2)}}{\dfrac{u_{cc}(\bar{c}_2)}{u_c(\bar{c}_2)} + \dfrac{u_{cc}(\bar{c}_1)}{u_c(\bar{c}_1)}}$$

$$B_{y_1} = \frac{1-\rho}{1-\beta\rho}\left[\frac{\dfrac{u_{cc}(\bar{c}_1)}{u_c(\bar{c}_1)}}{\dfrac{u_{cc}(\bar{c}_2)}{u_c(\bar{c}_2)} + \dfrac{u_{cc}(\bar{c}_1)}{u_c(\bar{c}_1)}} - \bar{b}_1 P_y\right]$$

$$B_{y_2}=-\frac{1-\rho}{1-\beta\rho}\left[\frac{\frac{u_{cc}(\bar{c}_2)}{u_c(\bar{c}_2)}}{\frac{u_{cc}(\bar{c}_2)}{u_c(\bar{c}_2)}+\frac{u_{cc}(\bar{c}_1)}{u_c(\bar{c}_1)}}+\bar{b}_1P_y\right]$$

$$C_{y_1}=1+\beta\frac{1-\rho}{1-\beta\rho}\frac{(1-\beta)\bar{b}_1-\frac{u_c(\bar{c}_2)}{u_{cc}(\bar{c}_2)}}{\frac{u_c(\bar{c}_1)}{u_{cc}(\bar{c}_1)}+\frac{u_c(\bar{c}_2)}{u_{cc}(\bar{c}_2)}}$$

$$C_{y_2}=\beta\frac{1-\rho}{1-\beta\rho}\frac{\frac{u_c(\bar{c}_1)}{u_{cc}(\bar{c}_1)}+(1-\beta)\bar{b}_1}{\frac{u_c(\bar{c}_1)}{u_{cc}(\bar{c}_1)}+\frac{u_c(\bar{c}_2)}{u_{cc}(\bar{c}_2)}}$$

Note that the bond law of motion has a unit root regardless of the assumptions on utility and of the steady state bond level $\bar{b}_1$. Similarly, the bond price is only a function of aggregate income shocks. These two main properties discussed in the text are also true in the more general setup. If we assume the same CRRA utility for the two countries, then we can make some further points regarding the effects of $\bar{b}_1\neq 0$ on the consumption savings decisions. Whereas with $\bar{b}_1=0$, bonds respond only to redistributive shocks, with $\bar{b}_1>0$ bonds also respond to purely aggregate shocks. The reason is that aggregate shocks affect bond prices and this in turn has redistributive effects when $\bar{b}_1>0$. A purely aggregate positive income shock effectively reduces interest rates, redistributing from the saver to the borrower. As a result, the saver (country 1 in this scenario with $\bar{b}_1>0$) reduces bonds and the borrower repays some debt. The effect on consumption depends on the relative strength of wealth and substitution effects. It can be shown that consumption is unaffected when $\gamma=1$ in the CRRA case, but it increases (decreases) when $\gamma<1$ ($\gamma>1$) because the substitution (wealth) effect dominates.

Appendix B Policy iteration method

This appendix provides a brief description of the policy iteration algorithm used and the parameter choices made. The equilibrium conditions that need to be satisfied at every possible state of the economy $s_t=(y_{1t}, y_{2t}, b_{1t-1})$ are

$$p(s_t)c_1(s_t)^{-\gamma}-\lambda_1(s_t)=\beta E_t[c_1(s_{t+1})^{-\gamma}] \tag{B1}$$

$$p(s_t)c_2(s_t)^{-\gamma}-\lambda_2(s_t)=\beta E_t[c_2(s_{t+1})^{-\gamma}] \tag{B2}$$

$$c_1(s_t)+p(s_t)b_1(s_t)=y_1(s_t)+b_1(s_{t-1}) \tag{B3}$$

$$c_1(s_t)+c_2(s_t)=y_1(s_t)+y_2(s_t) \tag{B4}$$

$$\lambda_1(s_t)(b_1(s_t)+K)=0 \tag{B5}$$

$$\lambda_2(s_t)(-b_1(s_t)+K)=0 \tag{B6}$$

Note that the bond market clearing condition has already been used to substitute out $b_2(s_t)=-b_1(s_t)$ and the budget constraint for country 2 is omitted since it is automatically satisfied by Walras' law. All variables are expressed as functions of the state vector s_t. As discussed in Ljungqvist and Sargent (2004), and contrary to the complete markets case, the equilibrium is not recursive in the exogenous state variables (y_{1t}, y_{2t}). That is, the wealth redistribution needs to be added as a state variable given the incompleteness of financial markets. Because the current model consists of only two agents, adding the variable b_1 to the state vector is enough to keep track of the wealth distribution. Note that exogenous income is trivially dependent on the state vector as $y_1(s_t)=y_{1t}$ and $y_2(s_t)=y_{2t}$.

The first step is to discretize the state space. The AR(1) process given in (2) is approximated with a Markov chain with $M=9$ states for each shock, implying a total of $M^2=81$ states for the two shocks together. The discretization procedure follows the method of Tauchen and Hussey (1991) adjusted as suggested in Flodén (2008) to improve its performance for processes with high level of persistence. The resulting Markov chain has values for persistence and variance that match the corresponding ρ and σ_ε^2 values of the AR(1) process extremely well for values of $\rho\leq 0.95$. As ρ rises above 0.95, the Tauchen and Hussey approximation starts deteriorating even with the Floden adjustment. As pointed out in Floden, the older method of Tauchen (1986) tends to be more robust. Thus, for $\rho>0.95$ this method is used instead. This method uses nodes equally spaced between $\pm\alpha\sigma_y$ ln M, where σ_y is the standard deviation of y, M is the number of nodes and $\alpha=1.2$. For each $\rho>0.95$, we check the resulting Markov chain's implied values for ρ, σ_ε and σ_y and compare them to the targeted values, i.e., the values assumed in the AR(1) process. We adjust α as ρ gets closer to one to obtain the best fit.

The endogenous state variable (bond) lies in an interval $[-K, K]$ where $K=1$ is the debt limit. This interval is discretized using $N=301$ points with higher concentration of points closer to the limits (points are distributed according to a quadratic function). Although 301 points is not a very large number, Euler errors end up being relatively small. This is to a large extent due to the use of interpolation for the space in between points. In the case of loose limits (K=80.7) we use N=3000 to accommodate the fact that the state space is much wider.

To summarize up to now, we have obtained a collection of M^2N points in the state space. From here on we drop the time subscripts and denote the current state at a point i in the grid by $s_i=(b_{1j}, y_{1k}, y_{2r})$, $i\in\{(j, k, r):j=1, \dots, N, k=1, \dots, M, r=1, \dots, M\}$. Also let s and s' denote a generic value for the current and future state vectors, respectively. Given guesses $B_1^0(s)$, $C_1^0(s)$ and $C_2^0(s)$ for the bond and consumption policy functions and using the Markov chain it is straightforward to compute conditional expectations of marginal utility $\phi_1(s)\equiv E[C_1^0(s')^{-\gamma}|s]$ and $\phi_2(s)\equiv E[C_2^0(s')^{-\gamma}|s]$ for any point s in the discretized state space. Note that $B_1^0(s)$ is not necessarily on the grid so interpolation is used to compute $C_1^0(s')$ and $C_2^0(s')$. Armed with the conditional expectations, we can solve for the implied bond and consumption functions $B_1^1(s)$, $C_1^1(s)$ and $C_2^1(s)$ [as well as price and multiplier functions $P(s)$, $\Lambda_1(s)$ and $\Lambda_2(s)$] at any state s using (B1)–(B6) as follows:

1. Assume limits do not bind at $s_i=(b_{1j}, y_{1k}, y_{2r})$ so that $\Lambda_1(s_i)=\Lambda_2(s_i)=0$. Using (B1)–(B4), straightforward algebra gives

$$C_1^1(s_i)=\frac{\phi_2(s_i)^{\frac{1}{\gamma}}}{\phi_2(s_i)^{\frac{1}{\gamma}}+\phi_1(s_i)^{\frac{1}{\gamma}}}(y_{1k}+y_{2r})$$

$$C_2^1(s_i)=\frac{\phi_1(s_i)^{\frac{1}{\gamma}}}{\phi_2(s_i)^{\frac{1}{\gamma}}+\phi_1(s_i)^{\frac{1}{\gamma}}}(y_{1k}+y_{2r})$$

$$P(s_i)=\beta\frac{\phi_1(s_i)}{C_1^1(s_i)^{-\gamma}}=\beta\frac{\phi_2(s_i)}{C_2^1(s_i)^{-\gamma}}$$

$$B_1^1(s_i)=\frac{y_1(s_i)+b_{1j}-C_1^1(s_i)}{P(s_i)}$$

2. If $B_1(s_i)<-K$, then set $B_1^1(s_i)=-K$, $\Lambda_2(s_i)=0$ and solve (B1)–(B4) for $C_1^1(s_i)$, $C_2^1(s_i)$, $P(s_i)$ and $\Lambda_1(s_i)$.
3. If $B_1(s_i)>K$, then set $B_1^1(s_i)=K$, $\Lambda_1(s_i)=0$ and solve (B1)–(B4) for $C_1^1(s_i)$, $C_2^1(s_i)$, $P(s_i)$ and $\Lambda_2(s_i)$.

It is not possible to obtain closed form expressions in cases 2 and 3, so a nonlinear equation solver is used to obtain a solution numerically. When this procedure is finished for all points s in the state space, we have implied policy functions $B_1^1(s)$, $C_1^1(s)$ and $C_2^1(s)$ which we use to update the guesses $B_1^0(s)$, $C_1^0(s)$ and $C_2^0(s)$. This process is repeated until the implied functions are close enough to the guesses. Specifically, the stopping criterion requires the maximum (over all points in the state space and across all the policy functions) of the absolute distance to be $<10^{-11}$. As Rendahl (2014) points out, there is no theoretical guarantee of convergence for the policy iteration algorithm. In practice, with the linearization solution as an initial guess, we always obtain convergence in our experiments.

Appendix C Euler errors

This appendix provides accuracy measures for our policy iteration method as well as for the perturbation methods. We show that Euler errors are small for our global solution which supports the claim in the main text that the policy function differences between the global solution and the perturbation methods are not due to inaccuracy in the global method.

C.1 Static Euler errors

Euler errors, as suggested by Judd (1992) and described in more detail in den Haan (2010), are a widely used measure of the accuracy of numerical solutions of dynamic models. The policy iteration algorithm delivers policy functions which satisfy the Euler equations (up to the tolerance level used for algorithm convergence, in our case 10^{-11}) at the points in the state space chosen during the discretization procedure. The idea is to check the size of the approximation error at other points in the state space. We look at the midpoints lying in between discretization points and calculate the Euler equation errors at these points. Specifically, at each point we use our (linearly interpolated) policy functions to obtain "actual" values for current consumption, bonds and prices. We then compare these to the "implied" values obtained by directly solving the non-linear equations (3)–(8) for c_{it}, b_{1t}, p_t with the conditional expectations calculated using our policy functions.[20] Santos (2000) shows that these Euler errors can be used to evaluate the accuracy of the solution since their size is closely related to the distance of the approximated policy functions to the true solution.[21] For each variable $x=c_1$, c_2, b_1, p, we compute the absolute difference between actual values x^A and implied

20 We take the discrete Markov chain as a model primitive so that conditional expectations are exactly computed.

21 Note, however, that Santos (2000) relies on an interority assumption that is not satisfied here.

values x^I at each point considered. We express this difference as a percentage of x^I for $x=c_1$, c_2, p and as a fraction of mean income for bonds b, and denote them by Δx

$$\Delta x \equiv 100\left|\frac{x^A - x^I}{x^I}\right| \text{ for } x=c_1,\ c_2,\ p$$
$$\Delta b \equiv 100\,|b^A - b^I| \tag{C7}$$

Tables 6 and 7 reports the Euler errors for the benchmark calibration.

The maximum and the average values of these are reported in Table 6. Euler errors are of similar order of magnitude for the three variables, so we focus on consumption. At most points in the state space Euler errors are very small, as indicated by the average being 1.8×10^{-5}%. Euler errors are larger close to the limits, with the maximum at 0.0077% of consumption. Overall, we find the errors to be reasonably small and, importantly, several orders of magnitude smaller than the differences between linear and policy iteration solutions reported in the previous section. This provides some confidence that the differences reported in section 4.3 of the main text are not simply due to large approximation errors in the policy iteration solution.

We have also computed Euler errors for the linear solution which we present in Table 7. Focusing on consumption, maximum Euler errors are almost 4% of consumption and average errors are 0.53%. These are several orders of magnitude

Table 6 Static Euler errors, policy iteration solution.

	Whole state space		Binding points excluded	
	Max	**Mean**	**Max**	**Mean**
Δc	0.0077%	1.82×10^{-5}%	0.0026%	1.90×10^{-5}%
Δb	0.0129%	2.11×10^{-5}%	0.0024%	1.89×10^{-5}%
Δp	0.0063%	0.30×10^{-5}%	0.90×10^{-5}%	4.90×10^{-8}%

Table 7 Static Euler errors, linear solution.

	Whole state space		Binding points excluded	
	Max	**Mean**	**Max**	**Mean**
Δc	3.86%	0.530%	0.041%	0.009%
Δb	6.69%	1.042%	0.237%	0.082%
Δp	3.13%	0.502%	0.102%	0.037%

higher than the policy iteration errors. The maximum error is of similar size to the maximum policy difference, but the average errors appear to be smaller than what the policy differences would suggest. Indeed, when we exclude the limit points, the maximum falls to 0.04% of consumption and the average is approximately 0.01% of consumption. Although still significantly larger than the PI solution errors, these could be argued to be reasonably small for a first approximation and they paint a very different picture than what policy differences would suggest. We interpret this as evidence of the importance of the expectations of future binding constraints, even when constraints do not currently bind. Whereas the global approximation method builds this in the conditional expectations, the linear solution does not. When computing standard Euler errors for the linear solution, we compute conditional expectations using the linear policy, i.e., ignoring the possibility of future binding constraints. Under this assumption, current actual decisions are not too far away from the implied ones. But the global solution is actually very different exactly because these expectations are very different. Thus, focusing on Euler errors can be misleading as a measure of the quality of approximation obtained with linear methods for a model with occasionally binding constraints, even if one focuses on points that are away from the limits.

The Euler errors of the solutions under the natural debt limit are presented in Table 8.

C.2 Dynamic Euler errors

den Haan (2010) suggests a second accuracy measure which he terms dynamic Euler errors, as opposed to the static Euler errors computed in the previous section. The idea is to evaluate whether small static Euler errors can accumulate to large errors in a simulation. This is especially relevant for our case, since we know that the linear solution yields non-stationary policy functions so one of the main concerns about its accuracy is exactly this accumulation of errors in a simulation.

Table 8 Static Euler errors, with K=80.7 and γ=1.

	Policy iteration		First order		Second order	
	Max	Mean	Max	Mean	Max	Mean
Δc	0.293%	1.76×10^{-7}%	3.50%	0.014%	5.43%	0.012%
Δb	0.060%	2.54×10^{-7}%	17.20%	0.794%	10.75%	0.377%
Δp	0.0007%	$2.44\times10-9$%	0.131%	0.018%	0.141%	0.0173%

To produce a simulated series, we first draw a sequence of realizations for the exogenous income shocks and initialize the series at the steady state values for the state variables. For the "actual" series, we use directly the policy functions in each period to compute values for bonds, consumptions and prices given the exogenous shock. The bond choice at t is then the state variable for t+1 and a simulation is produced recursively. For the "implied" series, we compute current values by evaluating expectations according to the policy function but directly solving the non-linear equilibrium conditions for current bonds, consumptions and prices. This implied value for the bond choice at t is then used as the state variable for t+1 and the simulation is thus produced recursively again. In order to focus on the possibility of future binding constraints we produce many short run simulations as opposed to one very long one. This is intended to give the linear solution the best chance to perform well. In a very long simulation, the unit root in the linear bond law of motion would imply bonds that drift arbitrarily far away from the steady state and, in particular, above the limits. We choose the simulation length to be 50 periods and compute average errors over 20,000 repetitions. The draws of the exogenous shocks are kept the same in the simulations using the linear and global solution to make those directly comparable.

We define Δx as in (C7) and report these dynamic Euler errors for the policy iteration method in the top panel of Table 9. It is important to note that the limits never bind in any of the simulations. This is partly due to the fact that the debt policy function is mean-reverting but also due to the fact that we are looking at short run simulations starting at zero bonds. The errors are thus small because the static Euler errors away from the limits are small and there does not seem to be substantial accumulation of errors in a 50 period simulation. The maximum consumption error occurs when bonds come close to the limit and is 0.0022%. The economy spends most of these simulations far from the limits so the average is very small at 3.74×10^{-6}%. The corresponding statistics for the

Table 9 Dynamic Euler errors.

	Max	Mean
Policy iteration solution		
Δc	0.0022%	3.74×10^{-6}%
Δb	0.0028%	7.08×10^{-5}%
Δp	8.24×10^{-7}%	9.75×10^{-9}%
Linear solution		
Δc	3.86%	0.033%
Δb	140.25%	0.50%
Δp	3.13%	0.041%

linear solution are shown in the bottom panel of Table 9. Despite the short simulation length, in the linear equilibrium debt sometimes reaches the debt limits. In the implied series, this limit is imposed but in the actual series it is not, since the linear policies do not assume such limits. As a result, the maximum error in bonds can be extremely high, indeed it can be arbitrarily high if we increase the simulation length.[22] For our short simulation experiment, this only happens for <2% of the periods. For consumption, the maximum error is more than 3% of consumption and the average is 0.033%. Consumption errors are significantly smaller than errors in bonds because of a moderating effect from prices. In the implied allocations, when bonds cannot adjust due to the limits, prices increase significantly, sometimes rising even above one. This provides significant consumption smoothing for the borrowing constrained country, keeping the consumption allocations from diverging even further away from the ones in the actual simulation.

To sum up, Euler errors can accumulate significantly for bonds as they drift above the limits, but this does not necessarily translate to large accumulation of errors for consumption, partly because prices adjust to provide some consumption smoothing. Before the limits are reached, actual and implied series stay very close even for the linear solution, indicating small dynamic Euler errors as long as the limits don't bind in the simulation. Similarly to the previous section, we note that the use of the linear policies to evaluate expectations essentially shuts down any effects from the possibility of future binding constraints. As a result, small dynamic Euler errors for the linear solution does not necessarily imply that the dynamic properties of the linear economy match well with those of the nonlinear economy with constraints. This point is illustrated by the cross-country consumption correlations reported in the main text.

References

Aiyagari, S. R. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *The Quarterly Journal of Economics* 109: 659–684.

Backus, D. K., P. J. Kehoe, and F. E. Kydland. 1992. "International Real Business Cycles." *Journal of Political Economy* 100: 745–775.

Baxter, M. 1995. "International Trade and Business Cycles." In *Handbook of International Economics*, vol. 3, chap. 35, 1801–1864, 1 ed. edited by G. M. Grossman and K. Rogoff. Amsterdam, The Netherlands: Elsevier/North-Holland.

22 We have also tried the alternative of imposing limits ex post for the actual series. In that case, we find very similar numbers for consumption and price errors to those reported below, but bond errors are now significantly smaller and similar in magnitude to the consumption errors.

Baxter, M., and M. J. Crucini. 1995. “Business Cycles and the Asset Structure of Foreign Trade.” *International Economic Review* 36: 821–854.

Blanchard, O. J., and C. M. Kahn. 1980. “The Solution of Linear Difference Models under Rational Expectations.” *Econometrica* 48: 1305–1311.

Carroll, C. D., and M. S. Kimball. 1996. “On the Concavity of Consumption Function.” *Econometrica* 64: 981–992.

Carroll, C. D., and M. S. Kimball. 2001. “ Liquidity Constraints and Precautionary Saving.” NBER Working Paper 8496.

Christiano, L. J. 1990. “Solving the Stochastic Growth Model by Linear-Quadratic Approximation and by Value-Function Iterations.” *Journal of Business and Economic Statistics* 8: 23–26.

Christiano, L. J. 2002. “Solving Dynamic Equilibrium Models by a Method of Undetermined Coefficients.” *Computational Economics* 20: 21–55.

Coleman, W. J. I. 1990. “Solving the Stochastic Growth Model by Policy-Function Iteration.” *Journal of Business & Economic Statistics* 8: 27–29.

Deaton, A. 1991. “Saving and Liquidity Constraints.” *Econometrica* 59: 1221–1248.

den Haan, W. J. 2001. “The Importance of the Number of Different Agents in a Heterogeneous Asset-Pricing Model.” *Journal of Economic Dynamics and Control* 25: 721–746.

den Haan, W. J. 2010. “Comparison of Solutions to the Incomplete Markets Model with Aggregate Uncertainty.” *Journal of Economic Dynamics and Control* 34: 4–27.

Flodén, M. 2008. “A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1) Process.” *Economic Letters* 99: 516–520.

Granger, C. W. J., and N. R. Swanson. 1997. “An Introduction to Stochastic Unit-Root Process.” *Journal of Econometrics* 80: 35–62.

Heaton, J., and D. J. Lucas. 1996. “Evaluating the Effects of Incomplete Markets on Risk Sharing and Assets Pricing.” *Journal of Political Economy* 104: 443–487.

Huggett, M. 1993. “The Risk-free Rate in Heterogeneous-agent Incomplete-insurance Economies.” *Journal of Economic Dynamics and Control* 17: 953–969.

Judd, K. L. 1992. “Projection Methods for Solving Aggregate Growth Models.” *Journal of Economic Theory* 58: 410–452.

Judd, K. L. 1998. *Numerical Methods in Economics*. Cambridge, Massachusetts and London, England: MIT Press.

Kehoe, P. J., and F. Perri. 2002. “International Business Cycles with Endogenous Incomplete Markets.” *Econometrica* 70: 907–928.

Kimball, M. S. 1990. “Precautionary Saving in the Small and in the Large.” *Econometrica* 58: 53–73.

King, R. G., C. I. Plosser, and S. T. Rebelo. 1988. “Production, Growth and Business Cycles: 1. The Basic Neoclassical Model.” *Journal of Monetary Economics* 21: 195–232.

Kollman, R. 1996. “Incomplete Asset Markets and the Cross-Country Consumption Correlation Puzzle.” *Journal of Economic Dynamics and Control* 20: 945–961.

Lane, P. R., and G. M. Milesi-Ferretti. 2007. “The External Wealth of Nations mark II: Revised and Extended Estimates of Foreign Assets and Liabilities, 1970–2004.” *Journal of international Economics* 73: 223–250.

Levine, D. K., and W. R. Zame. 1996. “Debt Constraints and Equilibrium in Infinite Horizon Economies with Incomplete Markets.” *Journal of Mathematical Economics* 26: 103–131.

Ljungqvist, L., and T. J. Sargent. 2004. *Recursive Macroeconomic Theory*, 2nd ed. Cambridge, Massachusetts and London, England: MIT Press.

Magill, M., and M. Quinzii. 1994. "Infinite Horizon Incomplete Markets." *Econometrica* 62(4): 853–880.

Marcet, A., and A. Scott. 2009. "Debt and Deficit Fluctuations and the Structure of Bond Markets." *Journal of Economic Theory* 144: 473–501.

Marcet, A., and K. J. Singleton. 1999. "Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints." *Macroeconomic Dynamics* 3: 243–277.

Mcgrattan, E. R. 1990. "Solving the Stochastic Growth Model by Linear-Quadratic Approximation." *Journal of Business and Economic Statistics* 8: 41–44.

Petrosky-Nadeau, N., and L. Zhang. 2013. "Solving the DMP Model Accurately." *NBER Working Paper*.

Rendahl, P. 2014. "Inequality Constraints and Euler Equation Based Solution Methods." *Economic Journal*. DOI: 10.1111/ecoj.12115.

Santos, M. S. 2000. "Accuracy of Numerical Solutions Using the Euler Equation Residuals." *Econometrica* 68: 1377–1402.

Schmitt-Grohé, S., and M. Uribe. 2003. "Closing Small Open Economy Models." *Journal of International Economics* 61: 163–185.

Schmitt-Grohé, S., and M. Uribe. 2004. "Solving Dynamic General Equilibrium Models using a Second-Order Approximation to the Policy Function." *Journal of Economic Dynamics and Control* 28: 755–775.

Tauchen, G. 1986. "Finite State Markov-chain Approximations to Univariate and Vector Autoregressions." *Economic Letters* 20: 177–181.

Tauchen, G., and R. Hussey. 1991. "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models." *Econometrica* 59: 371–396.

Telmer, C. I. 1993. "Assets-Pricing Puzzles and Incomplete Markets." *The Journal of Finance* 48: 1803–1832.

Uhlig, H. 1999. "A Toolkit for Analyzing Non-linear Dynamic Stochastic Models Easily." In *Computationaly Methods for the Study of Dynamic Economies*, chap. 3, 1 ed., edited by R. Marimon and A. Scott, Oxford, UK: Oxford University Press.

Xu, X. 1995. "Precautionary Saving and Liquidity Constraints: A Decomposition." *International Economic Review* 36: 675–690.

Zeldes, S. P. 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *The Quarterly Journal of Economics* 104: 275–298.