

NBER WORKING PAPER SERIES

GLOBAL V. LOCAL METHODS IN THE ANALYSIS OF OPEN-ECONOMY MODELS
WITH INCOMPLETE MARKETS

Oliver de Groot
Ceyhun Bora Durdu
Enrique G. Mendoza

Working Paper 26426
<http://www.nber.org/papers/w26426>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2019

We thank Tom Holden, Matteo Iacoviello, and Rob Vigfusson for helpful discussions and comments, and Holt Dwyer, Alex Martin and Sergio Villalvazo for excellent research assistance. We are also grateful for comments by conference and seminar participants at the 2018 SED Meeting, the 2018 North American Meetings of the Econometric Society, the Bank of Finland, George Washington University, and the 2018 Central Bank Macroeconomic Modeling Workshop. The views expressed herein are those of the authors and should not be attributed to the National Bureau of Economic Research, the Board of Governors of the Federal Reserve System or its staff.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w26426.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Oliver de Groot, Ceyhun Bora Durdu, and Enrique G. Mendoza. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Global v. Local Methods in the Analysis of Open-Economy Models with Incomplete Markets
Oliver de Groot, Ceyhun Bora Durdu, and Enrique G. Mendoza
NBER Working Paper No. 26426
November 2019
JEL No. E44,E47,F41,F44,F47

ABSTRACT

Global and local methods are widely used in international macroeconomics to analyze incomplete-markets models. We study solutions for an endowment economy, an RBC model and a Sudden Stops model with an occasionally binding credit constraint. First-order, second-order, risky steady state (RSS), and DynareOBC solutions are compared v. fixed-point-iteration global solutions in the time and frequency domains. The solutions differ in key respects, including measures of precautionary savings, cyclical moments, impulse response functions, financial premia and macro responses to credit constraints, and periodograms of consumption, foreign assets and net exports. The global method is easy to implement and fast albeit slower than local methods, except DynareOBC which is of comparable speed. These findings favor global methods except when prevented by the curse of dimensionality and urge caution when using local methods. Of the latter, first-order solutions are preferable because results are very similar to second-order and RSS methods.

Oliver de Groot
University of Liverpool
Management School
Chatham Street,
Liverpool, L69 7ZH
United Kingdom
oliverdegroot@gmail.com

Enrique G. Mendoza
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104
and NBER
egme@sas.upenn.edu

Ceyhun Bora Durdu
Federal Reserve Board
20th Street and Constitution Avenue N.W.
Washington, DC 20551
bora.durdu@frb.gov

1 Introduction

All major strands of the International Macroeconomics literature study topics in which incomplete asset markets play a key role (e.g., business cycles in emerging economies, sovereign default, Sudden Stops, global imbalances, nominal rigidities, macroprudential regulation, unconventional monetary policies, currency carry trade, etc). Since the dynamics of external wealth or net foreign assets (NFA) generally lack analytic solutions under incomplete markets, researchers rely on numerical methods to study the implications of theoretical models. Choosing the appropriate numerical methods is difficult, however, because deterministic models yield stationary equilibria dependent on initial conditions and, under uncertainty, market incompleteness makes the evolution of wealth state-contingent and influenced by precautionary savings. Certainty equivalence fails and the NFA position grows infinitely large when setting the rate of interest equal to the rate of time preference.

The literature follows two approaches to address these issues. The first approach, based on the influential work by Schmitt-Grohé and Uribe (2003), modifies the models by introducing “stationarity-inducing” assumptions that yield a well-defined deterministic steady state for NFA, independent of initial conditions, and implements log-linear or first-order (1OA) approximations around that steady state, recovering certainty equivalence. Schmitt-Grohé and Uribe proposed introducing one of three assumptions: a debt-elastic interest rate (DEIR) function by which the real interest rate rises as NFA falls below its steady state, quadratic costs that make NFA costly to deviate from steady state, or a subjective discount factor that varies with individual or aggregate consumption.¹ Important innovations were made subsequently to improve local methods, including second- (2OA) and higher-order approximation methods (e.g. Devereux and Sutherland (2010), Fernandez-Villaverde et al. (2011)), the risky steady state (RSS) method proposed by Coeurdacier et al. (2011), and the algorithms for solving models with occasionally binding constraints (e.g. OccBin by Guerrieri and Iacoviello (2015), DynareOBC by Holden (2016a)).

The second approach uses global methods to solve directly for the nonlinear decision rules and long-run distribution of wealth of the models in their original form. These methods are analogous to those used in other fields in Macroeconomics, particularly to solve models of heterogeneous agents with incomplete markets. Open-economy applications of these methods date back to Mendoza’s (1991) RBC model of a small open economy, and there are now many applications in quantitative studies of sovereign default (e.g., Aguiar and Gopinath (2006), Arellano (2008)), emerging markets

¹Schmitt-Grohé and Uribe showed that the business cycle moments of an RBC small-open-economy model solved using any of these assumptions are similar, and impulse response functions to a TFP shock are virtually identical.

business cycles (e.g., Neumeyer and Perri (2005), Mendoza (1995), Uribe and Yue (2006)), global imbalances (e.g. Mendoza et al. (2009)), Sudden Stops (e.g., Mendoza and Smith (2006), Durdu et al (2009), Mendoza (2010)), and macro-financial regulation (e.g., Bianchi (2011), Benigno et al. (2016), Bianchi et al. (2012), Schmitt-Grohé and Uribe (2017)).

Tables 1 and 2 document the usage of local and global methods for solving open-economy models with incomplete markets in research articles and policy institutions. Table 1 includes 61 papers. The list is meant to be illustrative, because constructing a complete list of the articles in the literature that use these methods is beyond the scope of this paper. It includes the 50 most cited papers in *Google Scholar* that cite Schmitt-Grohé and Uribe (2003), excluding textbooks and review articles. It also includes all quantitative papers in the references of this paper that are not in that top-50 list, and several well-known papers going back to the early 1990s when the first numerical solutions of open-economy models with incomplete markets were produced. Table 2 lists the models used in eight policy institutions, using information obtained from publicly available documents.

Table 1 shows that both local and global methods are widely used in research. IOA is the most common local method, and from the assumptions to induce stationarity, DEIR is the most common. Table 2 shows that all eight policy models use the IOA method. Five of these models use DEIR. Of all solutions that use DEIR, the majority set the value of the debt elasticity parameter ψ to an arbitrary small number, with the aim of preventing the DEIR function from playing a role other than inducing stationarity, since ideally this function should be an endogenous object.² The ψ values range from 0.00001 to 0.01, and the most common is 0.001, which is the value proposed by Schmitt-Grohé and Uribe.³ In other cases, the value of ψ is set by calibration or obtained via estimation, although for some policy models the value of ψ and/or the method followed to set it are not reported. In calibrated cases (three research papers and three policy models), ψ is in the 0.01-0.1 range, and in estimated cases (four research papers and one policy model), the point estimates or the medians of posterior distributions in Bayesian estimation are in the 0.00014-2.8 range.

While global methods have the advantage that the models are solved in their original form and capture the dynamics of wealth accurately, they suffer from the traditional curse-of-dimensionality problem: They are impractical in large models because the methods become exponentially inefficient

²Garcia-Cicco et al. (2010) explain that, following Schmitt-Grohé and Uribe (2003), the standard practice is to set ψ to a small value because the DEIR function aims to obtain independence of the deterministic steady state from initial conditions without affecting cyclical dynamics. They also studied a model in which ψ represents financial frictions, and in this case they estimated ψ using Bayesian methods.

³Note, however, that the DEIR functional forms are not always the same, so ψ values are not directly comparable. When relevant for our quantitative analysis, we control for this by making comparisons in terms of the elasticity of the interest rate with respect to percent deviations of NFA from steady state.

as the number of endogenous state variables rises. The local methods have the advantage that they can be applied in large-scale models, but have the shortcoming that they need the stationarity-inducing assumptions that are not part of the original models, and more importantly, as we show in this paper, these extra assumptions may not be innocuous for the numerical results. Hence, these tradeoffs pose two key questions: Under what conditions are local solutions better approximations to the “exact” solutions obtained with global methods? When those conditions are satisfied, which features of the global solutions are approximated accurately and which are not?

This paper provides answers to these questions by comparing local and global solutions. For the local methods, we evaluate 1OA, 2OA, RSS, and DynareOBC solutions. For the global methods, we use the fixed-point iteration (*FiPit*) algorithm proposed by Mendoza and Villalvazo (2019), which applies fixed-point iteration to solve for NFA (asset prices) in the Euler equation of bonds (capital).⁴ For the models we study here, this algorithm is fast and easy to implement in Matlab. The local methods (except DynareOBC) are faster but they yield larger Euler equation errors. DynareOBC and *FiPit* are of comparable speed.

With the global method, the existence of a well-defined stochastic steady state follows from the same condition as in the Bewley-Aiyagari-Hugget class of heterogeneous-agents models: the rate of time preference must be lower than the interest rate.⁵ In multicountry models this is a general equilibrium result, because if the rate of interest equals the rate of time preference, all countries would desire an infinitely large stock of NFA for self-insurance, which is inconsistent with world general equilibrium (see Mendoza et al. (2009)). Hence, assuming an interest rate lower than the rate of time preference in small-open-economy models is an *implication* of the assumption that the interest rate is a world-determined price. With the local methods, the DEIR function is constructed so that at a pre-determined steady state the rate of interest equals the rate of time preference.

We compare local v. global solutions for three widely-used small open economy models: An endowment economy model, a real business cycle (RBC) model, and a model of Sudden Stops (SS) with an occasionally-binding credit constraint linked to market prices. In each case, we start with “*baseline calibrations*” in which the local methods use DEIR with $\psi = 0.001$ and the center of approximation of the 1OA, 2OA and DynareOBC methods is the deterministic steady state, and

⁴We also solved the models using the canonical value function iteration method and obtained very similar results (see Appendix D.1). This method is slow in models with more than one endogenous state variable and often requires efficient equilibria, but does not require differentiability or convexity of the optimization problems, which *FiPit* requires. Moreover, *FiPit* does not use the contraction mapping property so its convergence is not guaranteed.

⁵Preferences with endogenous discounting proposed by Uzawa (1968) and Epstein (1983) also support a well-defined limiting distribution of NFA. In Appendix Section D.1, we study global solutions obtained with endogenous discounting. Qualitatively, the comparison v. local methods yields similar findings as those reported in the paper.

RSS is centered at its risky steady state. Then we consider “*targeted calibrations*,” for which ψ is calibrated to match the first-order autocorrelation of NFA in the global solutions.⁶ For RSS and DynareOBC, we also solve variants without DEIR in which the rate of interest is lower than the rate of time preference, so that credit constraints bind at the deterministic steady state. We conduct both time-domain comparisons of statistical moments and dynamics in response to shocks, and frequency-domain comparisons based on spectral density analysis. In addition, for the SS model we compare multipliers, financial premia and macro responses when the credit constraint binds.

The results show important differences between local and global solutions, which are due to differences in the decision rule driving NFA, the key endogenous state variable in open economy models. In the global solutions, the NFA decision rule is solved for “exactly,” and the model’s equilibrium stochastic processes are jointly determined by this decision rule (and the one for capital in the RBC and SS models) and the stochastic process of exogenous shocks. In the local solutions (other than DynareOBC), the NFA decision rule is a first- or second-order approximation around the deterministic steady state or the risky steady state for the RSS method. Hence, differences across the solution methods can be summarized by comparing NFA decision rules, particularly their first-order autocorrelations and their implications for two key moments: (a) the mean of NFA, which is an indicator of precautionary savings and the critical variable in research topics like global imbalances or reserves accumulation; and (b) the autocorrelation of net exports, since net exports are also a key variable in the analysis of open-economy models (see Mendoza (1991), Aguiar and Gopinath (2006), and Garcia Cicco et al. (2010)).

Global methods pin down the “true” NFA autocorrelation, while local solutions yield approximations that differ depending on the solution method and the criteria used to set ψ and the center of approximation. Both global and local methods produce a near-unit-root NFA process. In all calibrated exercises, the NFA autocorrelations exceed 0.95. The autocorrelations differ slightly, however, and because they are near-unit-roots these small differences have major implications.

The largest difference between global and local solutions is in the long-run average of NFA. Small differences in the autocorrelation and intercepts of the NFA decision rules cause large differences in mean NFA, and the local methods can yield NFA averages above or below the global solution. The local methods also yield very different results for the effects of parameter changes that alter precautionary savings. For instance, the global solution predicts large increases in mean NFA with

⁶We also studied an alternative in which the center of approximation is the average NFA of the global solutions, but targeting the autocorrelation produces a closer match to the global solutions.

higher variability of shocks, lower rate of time preference or higher coefficient of relative risk aversion (CRRA). In contrast, the 1OA method preserves certainty equivalence, and hence keeps mean NFA equal to the deterministic stationary equilibrium implied by the DEIR function. The 2OA and RSS methods produce long-run averages of NFA that are much larger or smaller than the mean of the global methods, depending on the model, the parameter change considered, and whether we use baseline or targeted calibrations. For example, for the endowment economy and variations in income volatility, the 2OA method overestimates average NFA while the RSS underestimates it using baseline calibrations, while for targeted calibrations both methods underestimate average NFA significantly. Similarly, DynareOBC calibrated to a steady state in which the credit constraint binds (does not bind) yields mean NFA significantly below (above) the global solution. Thus, a major drawback of the local methods is that they do poorly at quantifying precautionary savings.

Small differences in the autocorrelation of NFA also yield large differences in the autocorrelations of the net exports-GDP ratio (NX), because NX is a quasi first-difference of a near-unit-root NFA process, as shown in Section 2. For instance, in the endowment model with the baseline calibration, the global solution predicts that raising the persistence of income from near 0 to 0.8 makes the autocorrelation of NFA rise from 0.83 to 0.99 but that of NX varies from -0.09 to 0.77. In contrast, the 2OA and RSS solutions predict that the autocorrelation of NFA always exceeds 0.99 while that of NX varies from 0.24 to 0.95. For a given autocorrelation of income in the 0-0.8 range, the local solutions always overestimate the autocorrelations of NFA and NX.

The local methods do better at matching the moments of the global solution with the targeted calibrations (i.e. with ψ set to match the autocorrelation of NFA in the global solution). Averages of NFA get closer to the global solution. The autocorrelations of NFA and NX are similar when the persistence of income is such that the autocorrelation of NFA in the global solution is greater or equal to that obtained with the baseline calibration (0.977), but for less persistent income, and hence lower autocorrelations of NFA in the global solution, again the local solutions overestimate the autocorrelations of NFA and NX. Moreover, the targeted calibration approach has the drawback that it requires obtaining first the global solution so as to find the first-order autocorrelation of NFA to calibrate ψ , and doing this again for any parametric change that alters the NFA autocorrelation.

Comparing alternative local methods also yields important findings. 1OA, 2OA and RSS solutions yield similar second- and higher-order moments and similar impulse response functions for all endogenous variables in the endowment and RBC models. We use analytic solutions of the endowment model NFA decision rules to show that this is due to two features of the results: First, the

coefficient on lagged NFA is nearly the same in the RSS and 2OA solutions when ψ is small (less than 0.1), unless the deterministic and risky steady states of NFA differ by a very large margin (at least 40 percentage points of GDP). Second, the coefficients in the square and interaction terms of the 2OA decision rules are very small. As a result, the three methods yield different first moments because of their different centers of approximation (and different intercept for 1OA), but the second- and higher-order moments and impulse response functions are about the same.

Targeted calibrations require increasing ψ from the common calibration setting of 0.001 to values of 0.0469 (0.0063) and 0.0469 (0.0045) for the 2OA and RSS methods applied to the endowment (RBC) economy respectively. These variations imply large increases in the elasticity of the DEIR function, by factors of 4.5 to 25, which make this function play a more significant role than when $\psi = 0.001$. In particular, the required ψ values effectively make deviations of NFA from steady state too costly, and as a result even the first moments of 2OA and RSS are similar and they are both also close to the 1OA solution (i.e. certainty equivalence approximately holds). In these cases, 1OA should be the preferable local method, but it also means that precautionary savings are disregarded entirely. On the other hand, keeping $\psi = 0.001$ implies that predictions about some key moments, like the autocorrelation of NX, deviate significantly from the “true” solutions.

The global and local solutions also show important differences in the moments that measure the variability, persistence and output-correlation of consumption in the three models. In contrast, and in line with the findings of Schmitt-Grohé and Uribe (2003), the moments for output and investment in the RBC and Sudden Stops models are similar. This is an implication of the small excess return on capital typical of RBC models, which keeps the capital decision rule close to the one consistent with Fisherian separation of investment from saving due to arbitrage of asset returns.

The frequency-domain analysis shows that local and global methods yield equilibrium stochastic processes that display different behavior at most frequencies, and not just in business cycle moments and long-run averages. Non-parametric periodograms for the endowment economy with the baseline calibrations show that local methods overestimate significantly the contribution of low-frequency movements to the variability of NFA and NX, which is consistent with the result indicating that the local methods overestimate the autocorrelation of NFA. As a result, the local methods also overestimate (underestimate) the contribution of low-frequency (high-frequency) movements to consumption fluctuations. The targeted calibrations perform much better at approximating the spectral density of NFA and NX, but for consumption they still underestimate the contribution of high-frequency fluctuations by a large margin. Similar results are obtained for

NFA, NX and consumption in the RBC model, but in this case the local and global methods yield similar periodograms for capital, investment and output, in line with the finding that the time-series moments for these variables are similar with both methods.

Standard 2OA and RSS methods cannot be used for solving the SS model, because they cannot handle the occasionally binding constraint. We used instead Holden’s (2016a) DynareOBC toolkit. DynareOBC uses also a local approximation but introduces news shocks which hit every time the constraint is violated and push the relevant variables back to the constraint. To ensure the solution is consistent with rational expectations, these news shocks are constructed as if they were expected by agents along a perfect-foresight path and so are akin to “endogenous news shocks.” This method, when solved in first order and without integrating over future uncertainty, ignores precautionary savings and the *risk* of alternative future paths in which the constraint may or may not bind, and hence it also ignores the equity risk premium. We examined DynareOBC solutions with and without the constraint binding at the deterministic steady state.

The results show that our findings from the endowment and RBC model extend to the SS model. Local and global solutions yield large differences in the amount of precautionary savings induced by the credit constraint, business cycle moments, the probability of hitting the constraint, impulse responses, and spectral densities. Moreover, the near-unit-root nature of NFA increases DynareOBC execution time considerably, because it requires multiple, long perfect-foresight paths to form the news shocks realizations needed to implement the constraint, and long time-series simulations to attain convergence of long-run moments. DynareOBC also underestimates significantly the tightness of the credit constraint and its effects on financial premia and macro responses. In particular, it yields mean shadow interest premia when the constraint binds of 0.13 to 0.8 percent v. 2.6 percent in the global solution and average equity premia of 0.1 to 0.64 percent v. 2.2 percent. In turn, lower equity returns imply higher equity prices and investment when the constraint binds, and therefore higher borrowing capacity. As a result, DynareOBC with the constraint binding at steady state yields weaker Sudden Stop macro responses when the credit constraint binds, and DynareOBC with the constraint not binding at steady state cannot capture Sudden Stop responses in macro variables.

This paper is related to the broader literature comparing solution methods of nonlinear rational expectations models. The comprehensive study by Taylor and Uhlig (1990) showed that several methods that approximate global solutions of a canonical RBC model yield different results for simulated sample paths, density functions and cyclical moments. Dou et al. (2019) compared 1OA, 2OA and OccBin local solutions v. global solutions of a DSGE model with financial frictions and

found that the local solutions approximate poorly the model’s nonlinear dynamics and yield biased impulse responses. Rabitsch et al. (2015) compared the DS local method proposed by Devereux and Sutherland (2010) for solving portfolio allocations in a two-country, incomplete-markets model v. a global method. In the DS method, non-stationarity of the NFA position remains an issue, but given NFA it yields an accurate portfolio structure. They found that the DS method is accurate only with particular calibrations and with symmetric countries with long-run NFA set to 0. With asymmetric countries, and using endogenous discounting to induce stationarity, the DS method performs poorly unless the center of approximation matches the global solution, and more so if NFA decision rules are nonlinear. Our work differs from these studies in three key respects: We study 1OA, 2OA, RSS and DynareOBC methods using the dominant DEIR approach to induce stationarity; we compare results in both the time and frequency domains; and we consider endowment, RBC and SS models, and for the latter we compare the global and DynareOBC solutions.

It is also worth noting that Holden (2016b) compared DynareOBC v. global solutions of a small open economy model with endowment income (i.e. NFA as the single endogenous state) and subject to a constant debt limit and two non-negativity constraints as occasionally binding constraints. Our results differ in that we solve an SS model with two endogenous states (capital and NFA) and a collateral constraint that depends on both states and on asset prices. Holden found that DynareOBC is not very far from the global solution, while we found that the two solutions differ sharply.

The rest of the paper is organized as follows. Section 2 compares the endowment model solutions, providing analytic and quantitative results. Section 3 compares the RBC model solutions. Section 4 compares the global and DynareOBC solutions of the SS model. Section 5 provides conclusions.

2 Endowment economy

2.1 Model structure and equilibrium

Consider first a small open economy model with stochastic endowment income. We use this setup to derive analytic results and characterize NFA dynamics under incomplete markets. The economy is inhabited by a representative agent with preferences given by:

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t) \right\}, \quad u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}. \quad (1)$$

where $\beta \in (0, 1)$ is the subjective discount factor, c_t is consumption and σ is the CRRA coefficient.

The economy's resource constraint is:

$$c_t = e^{z_t} \bar{y} - A + b_t - qb_{t+1}, \quad (2)$$

where $e^{z_t} \bar{y}$ is stochastic income or GDP that fluctuates around a mean \bar{y} with shocks z_t of exponential term e^{z_t} , b_t denotes the NFA position in one-period, non-state-contingent discount bonds traded in a frictionless global credit market at an exogenous price $q = \frac{1}{1+r}$, where r is the world real interest rate, and A is a constant that represents investment and government expenditures (which is introduced so that the model can be calibrated to average consumption-GDP ratios observed in the data).⁷ Income shocks follow an AR(1) process: $z_t = \rho_z z_{t-1} + \sigma_z \varepsilon_t^z$ where ε_t^z is i.i.d.

The agent chooses the optimal sequences of bonds and consumption so as to maximize (1) subject to (2). This optimization problem is analogous to the one solved by a single individual in heterogeneous-agent models of precautionary savings (e.g., Bewley (1977), Aiyagari (1994) and Hugget (1993)). As in those models, the Inada condition of CRRA utility implies that the marginal utility of consumption goes to infinity as consumption goes to zero from above. This implies that the small open economy faces Aiyagari's Natural Debt Limit (NDL), by which the NFA position never exceeds the annuity value of the worst realization of net income $b_{t+1} \geq -\min(e^{z_t} \bar{y} - A)/r$, otherwise agents would be exposed to the possibility of nonpositive consumption with positive probability. We can also impose a tighter ad-hoc debt limit φ , following Aiyagari (1994), such that $b_{t+1} \geq \varphi \geq -\min(e^{z_t} \bar{y} - A)/r$. As we show later, this is useful for model calibration.

Using the resource constraint, we can express the Euler equation for bonds as follows:

$$u_c(e^{z_t} \bar{y} - A + b_t - qb_{t+1}) = (1+r)\beta E_t [u_c(e^{z_{t+1}} \bar{y} - A + b_{t+1} - qb_{t+2})] + \mu_t, \quad (3)$$

where $u_c(t)$ is the marginal utility of c_t and μ_t is the Lagrange multiplier of the debt limit.

A competitive equilibrium for this economy is defined by stochastic sequences $[c_t, b_{t+1}]_{t=0}^{\infty}$ that satisfy equation (3) for all t with $b_{t+1} \geq \varphi$. As we discuss next, the incompleteness of asset markets plays a critical role in determining key features of this equilibrium.

Consider first the equilibrium under complete markets of contingent claims. If income shocks are idiosyncratic to the small open economy, the economy diversifies away all the risk of its endowment fluctuations. As a result, the equilibrium features a constant consumption stream and the economy's wealth position is time- and state-invariant. The solution would be the same as in a perfect-foresight

⁷We assume that r is constant, but study later the implications of allowing it to be stochastic.

model with $\beta(1+r) = 1$ and wealth (the present value of income plus the initial NFA holdings) scaled to represent the same wealth as in the complete markets economy.

Under incomplete markets, the equilibrium differs significantly, because wealth becomes state-contingent and consumption cannot attain a perfectly smooth path. In particular, as shown in Chapter 18 of Ljungqvist and Sargent (2012), the Euler equation (3) implies that $M_t \equiv (1+r)^t \beta^t u_c(t)$ forms a supermartingale, which must converge almost surely to a non-negative random variable as an implication of the Supermartingale Convergence Theorem. If $\beta(1+r) \geq 1$, this convergence implies that consumption and NFA diverge to infinity because marginal utility converges to zero almost surely, which is the cause of the non-stationarity problem that led to the use of the DEIR function in local solution methods. The economy builds an infinitely large stock of precautionary savings so that self-insurance can sustain a consumption path that satisfies $u_c(t) \geq \beta(1+r)E_t[u_c(t+1)]$ and still leads to convergence in M_t . Hence, in this case the model lacks a “nice” stochastic stationary state with a well-defined limiting distribution of NFA. In contrast, if $\beta(1+r) < 1$, the economy attains a well-defined stochastic steady state with finite long-run averages of assets and consumption, and the rest of the moments of the model’s endogenous variables are also well-defined. A crude intuition is that the opposing forces of the incentive to save for self-insurance and the pro-borrowing incentive implied by $\beta(1+r) < 1$ make NFA fluctuate within an ergodic set. When NFA falls too much the precautionary savings force prevails and when NFA rises too much the pro-borrowing force prevails.

It is important to note that $\beta(1+r) < 1$ is also a general equilibrium outcome in multi-country models with incomplete markets, because otherwise all countries would want an infinite amount of NFA, which is inconsistent with world market clearing in the market of risk-free assets (see Mendoza et al. (2009)). Thus, in small-open-economy models $\beta(1+r) < 1$ is not an arbitrary assumption but an implication of the assumption that the real interest rate is a world-determined price. This also implies that the issues raised here are not particular to small open economy models, they are relevant for multi-country models and closed-economy macro models with incomplete markets.

2.2 Global methods

Global methods solve for the competitive equilibrium in recursive form. Since the equilibrium is efficient in this economy, we can represent it as the solution to this dynamic programming problem:

$$V(b, z) = \max_{c, b'} \left\{ \frac{c^{1-\sigma}}{1-\sigma} + \beta \sum_{z'} \pi(z', z) V(b', z') \right\}, \quad (4)$$

$$s.t. \quad c = e^z \bar{y} - A + b - qb', \quad b' \geq \varphi.$$

The AR(1) process of income is approximated as a discrete Markov chain with transition probability matrix $\pi(z', z)$. The solution to the Bellman equation is characterized by a decision rule $b'(b, z)$ and the associated value function $V(b, z)$. This decision rule and the Markov process of the shocks induce a joint ergodic (unconditional) distribution of NFA and income $\lambda(b, z)$.

We solve for $b'(b, z)$ over a discrete state space of (b, z) pairs using the *FiPIt* method.⁸ This method solves the recursive equilibrium conditions using a fixed-point iteration algorithm (see Mendoza and Villavazo (2019) for details). For this simple model, the *FiPIt* method iterates on the following representation of the Euler equation:

$$c_{j+1}(b, z) = \left\{ \beta R \sum_{z'} \pi(z', z) \left[\left(c_j(\hat{b}'_j(b, z), z') \right)^{-\sigma} \right] \right\}^{-\frac{1}{\sigma}} \quad (5)$$

Given a conjecture of the decision rule $\hat{b}'_j(b, z)$ in iteration j , the associated consumption function is $c_j(b, z) = e^z \bar{y} - A + b - q\hat{b}'_j(b, z)$. This consumption function is interpolated over its first argument in order to determine $c_j(\hat{b}'_j(b, z), z')$, so that the above equation solves directly for a new consumption function $c_{j+1}(b, z)$. Using the resource constraint, the new consumption function yields a new decision rule for bonds $b'_{j+1}(b, z)$, which is re-set to $b'_{j+1}(b, z) = \varphi$ if $b'_{j+1}(b, z) \leq \varphi$. Then the decision rule conjecture is updated to $\hat{b}'_{j+1}(b, z)$ as a convex combination of $\hat{b}'_j(b, z)$ and $b'_{j+1}(b, z)$, and the process is repeated until $b'_{j+1}(b, z) = \hat{b}'_j(b, z)$ up to a convergence criterion.

FiPIt uses a standard fixed-point iteration approach to solve transcendental equations so that the Euler equation solves directly instead of using a non-linear solver. Mendoza and Villavazo (2019) provide a detailed comparison of this algorithm with other Euler equation methods, including time iteration and endogenous grids methods. They show that this method is fast and easy to implement and optimize in Matlab for models with one or two endogenous state variables, including problems with occasionally binding constraints. In contrast, time-iteration with two endogenous states requires solving two nonlinear simultaneous equations and the endogenous grids method requires complex interpolation techniques because of the irregularity of the endogenous grids.

The global method solves the model in its “original form” without additional assumptions to impose stationarity. If $\beta(1+r) = 1$, the true solution is that NFA diverges to infinity, which is unpleasant but *is* the equilibrium outcome. $\beta(1+r) < 1$ is, however, the relevant case, because

⁸We show in Appendix D.1 that solving by value function iteration yields nearly identical results.

as noted above it is implied by world general equilibrium. Note also that with $\beta(1+r) < 1$ the *deterministic* stationary state converges to the natural or ad-hoc debt limit, with consumption falling at a gross rate of $(\beta(1+r))^{1/\sigma}$. Hence, without quantitative analysis, we can predict already that the long-run average of NFA in the stochastic, incomplete-markets model will differ significantly from the deterministic steady state and that the difference is due to precautionary savings.

2.3 Local methods

The local methods solve for the competitive equilibrium using a local approximation of the optimality conditions (equations (2) and (3)) either around the deterministic steady state (b^{dss}) for the 1OA and 2OA methods or the risky steady state (b^{rss}) for the RSS solutions. Also, since the model in its original form yields a deterministic stationary state that is either dependent on initial conditions (if $\beta(1+r) = 1$) or the debt limit (if $\beta(1+r) < 1$), 1OA and 2OA methods require one of the stationarity-inducing assumptions. As noted in the Introduction, the most widely used of these assumptions is to introduce the DEIR function, which generally adopts the following form:

$$1 + r_t = 1 + r + \psi \left[e^{b^{dss} - b_{t+1}} - 1 \right], \quad (6)$$

where ψ determines the elasticity of r_t with respect to NFA. For small deviations of NFA from b^{dss} , the elasticity of r_t with respect to percent deviations of b_{t+1} from b^{dss} is given by $\eta^r \equiv -\psi b^{dss}$.

We implement the 1OA and 2OA methods following Scmitt-Grohe and Uribe (2004) and the RSS method following Coeurdacier et al. (2011).⁹ 1OA and 2OA solve for local approximations around b^{dss} in standard fashion, by solving a first- or second-order approximation of the decision rules jointly with approximations of the same order of the model's optimality conditions. In contrast, RSS uses b^{rss} as center of approximation and assumes $\beta(1+r) < 1$ (see the Appendix for details). The risky steady state aims to take into account future risk, so that the center of approximation may approximate better the precautionary savings motive. The value of b^{rss} is obtained from a second-order approximation to the conditional expectation of the steady-state Euler equation, which is solved jointly with the coefficients of a first-order approximation of the decision rules using a conditional second-order approximation of the full equilibrium conditions' Jacobian, which requires the third derivatives of the equilibrium conditions. As explained by de Groot (2014), this second part of the solution is crucial to obtain stationary NFA dynamics in the RSS solution. We also

⁹Using Andreasen et al. (2018) to solve using 2OA and 3OA yields similar findings. Kliem and Uhlig (2016) also proposed a RSS method, in the context of an asset pricing model.

consider a variant of the RSS method in which b^{rss} is computed in the same way, but is then used together with the DEIR function and standard first-order approximations of the decision rules and equilibrium conditions to obtain stationary dynamics. We refer to the original RSS method as “full RSS” and the alternative with the DEIR function as “partial RSS.” Partial RSS is faster and, for other than first-order moments, it yields similar results as full RSS in our quantitative applications.

2.4 Calibration

We use the same baseline calibration as in Durdu et al. (2009), which was based on annual data for Mexico. Table 3 lists the calibration parameters. The Table shows a subset of parameters common to local and global solutions, as well as the parameters that are particular to each, including for the local methods both *baseline* and *targeted* calibrations of ψ .

The common baseline calibration parameters are set as follows. The CRRA coefficient is set to $\sigma = 2$, which is a standard value. Mean income is normalized to $\bar{y} = 1$. The steady state real interest rate is set to 5.9 percent, which is the average of Uribe and Yue’s (2006) real interest rate including the EMBI spread for Mexico. The target average ratio of net foreign assets to GDP is set to -44 percent, which is the average of Mexico’s NFA-GDP ratio over the period 1985-2004 in the database constructed by Lane and Milesi-Ferretti (2006), and the target consumption-GDP ratio is 0.692, which is the average ratio in Mexican data for the 1965-2005 period. These parameter values and the steady-state resource constraint imply that $A = \bar{y} + rb - c = 0.282$. The standard deviation and first-order autocorrelation of the endowment process are set to estimates obtained using the HP-detrended cyclical component of GDP computed also from the Mexican data. This yields $\sigma_z = 0.0327$ and $\rho_z = 0.597$. When using local methods, these parameters define the AR(1) process of income shocks that is part of the equilibrium conditions. In the global solution, we approximate this process as a five-point Markov chain using Tauchen and Hussey (1991)’s quadrature method.

The baseline calibration for the global solutions requires also calibrating the values of φ and β . As explained in Durdu et al. (2009), these parameters are set to values such that the model matches the -0.44 long run average of the NFA-GDP ratio observed in Mexican data together with Mexico’s cyclical variability of private consumption of 3.28 percent over the 1965-2005 period. This implies $\varphi = -0.51$ and $\beta = 0.94$. Two parameters are required to identify the calibration, because while the average NFA-GDP ratio can be matched by simply adjusting φ , this can result in stochastic steady states in which the distribution of bond holdings is clustered near the debt limit and consumption fluctuates too much, or has a high variance and consumption fluctuates too little.

In the baseline calibration for the local methods, we set the deterministic steady state of NFA to the same value of φ in the calibration of the global solutions, hence $b^{dss} = -0.51$. This is done so that both solutions have the same b^{dss} (recall that the deterministic steady state of the global solution is φ because $\beta(1+r) < 1$). The value of β for the local methods (except in the full RSS solution) is set to $\frac{1}{1+r}$, because the functional form of the DEIR function implies that the rate of interest equals the rate of time preference in the deterministic steady state. Note, however, that the discount factors of the global and local calibrations differ only slightly (0.944 v. 0.94). The *baseline* value of ψ is the commonly-used value of 0.001, taken from Schmitt-Grohé and Uribe (2003). In the *targeted* calibrations, we set ψ to values so that the solution for a given local method matches the autocorrelation of NFA obtained with the global solution. This yields $\psi = 0.0469$ for both the 2OA and RSS. We do this because, as we show below, approximating closely the autocorrelation of the equilibrium law of motion of NFA gives the local methods the best chance to match the global solution. Recall also that for a given value of ψ the elasticity of debt with respect to percent deviations of NFA from steady state varies with b^{dss} , since $\eta^r \equiv -\psi b^{dss}$ for small deviations.

2.5 Comparison of quantitative results

We compare global and local solutions along five dimensions. First, the NFA decision rule and its implications for net exports. Second, the full set of long-run moments. Third, impulse response functions for income shocks. Fourth, spectral density functions. Fifth, an environment with interest-rate shocks, in addition to income shocks.

a) Decision rules of NFA and Net Exports

We start by examining differences in the first-order autocorrelations of NFA and net exports (nx). Given that the trade balance is a quasi first-difference of NFA, since $nx_t = qb_{t+1} - b_t$, and assuming that b follows an AR(1) process with autocorrelation coefficient ρ_b , we show in Appendix Section A that the first-order autocorrelation of net exports (ρ_{nx}) is:¹⁰

$$\rho_{nx} = \frac{q^2 \rho_b + \rho_b - q - q\rho_b^2}{1 + q^2 - 2q\rho_b}. \quad (7)$$

Hence, if ρ_b is close to 1, as is typically the case in models with non-state-contingent assets, small differences in ρ_b induce large differences in ρ_{nx} . Thus, small errors in the local solution for ρ_b can result in large errors in their solution for ρ_{nx} . We show below that this is indeed the case here.

¹⁰NFA is an AR(1) process in the IOA and RSS solutions. In the 2OA solution it includes squared and interaction terms in b_{t-1} and y_t , but these terms are quantitatively negligible in all of our experiments.

The global methods produce “exact solutions” for ρ_b and ρ_{nx} that are computed using $\lambda(b, z)$, $\pi(z', z)$, $b'(b, z)$ and the definition of net exports. The local methods produce solutions implied by the local decision rules. The 2OA decision rule of NFA can be expressed as:

$$\tilde{b}_{t+1} = h_b \tilde{b}_t + h_y \tilde{y}_t + \frac{1}{2} \left(h_{bb} \tilde{b}_t^2 + h_{yy} \tilde{y}_t^2 \right) + h_{by} \tilde{b}_t \tilde{y}_t + \frac{1}{2} (h_{\sigma_z \sigma_z} \sigma_z^2) \quad (8)$$

where $\tilde{b}_t \equiv b_t - b^{dss}$ and $\tilde{y}_t \equiv y_t - \bar{y}$. The 1OA and RSS methods have similar expressions, except that they only have the first two terms in the right-hand side, because they both use a first-order approximation of the decision rules. The RSS method also differs in that it uses b^{rss} instead of b^{dss} .

The key coefficient to analyze is h_b , because it is the main determinant of ρ_b . This is the case even for the 2OA solutions because in all of our quantitative applications h_{bb} , h_{yy} and h_{by} are negligibly small.¹¹ The term $h_{\sigma_z \sigma_z} \sigma_z^2$ is also important, because it isolates the effect of income variability on mean NFA. It is an estimate of the amount of precautionary savings that the 2OA solution captures. Moreover, since this is the only term that is *quantitatively* relevant of those that distinguish 2OA from 1OA solutions, and both of these solutions have the same h_b term (as shown by Schmitt-Grohé and Uribe (2004)), these results also imply that the 2OA and 1OA solutions should be very similar, except in their first moments.

For the RSS method, de Groot (2014) showed that income volatility matters for determining b^{rss} because the coefficient of variation of consumption (relative to its risky steady state) is constant, at a level that depends on β , r and σ .¹² Intuitively, this captures precautionary savings because, if income variability rises and the shares of income allocated to savings v. consumption remain unchanged, the volatility of consumption would increase. But by increasing NFA relative to income, more of the disposable income comes from interest income, so that the coefficient of variation of consumption can remain constant. Since the RSS decision rule follows from a first-order approximation, however, the ρ_b value will differ from that implied by the 1OA and 2OA solutions only to the extent that b^{dss} and b^{rss} differ, and as we document below, this requires larger differences than those obtained under the baseline and targeted calibrations. Hence, 1OA, 2OA and partial RSS solutions are likely to be very similar, except for differences in first moments.

We now study how ψ and the center approximation affect the determination of h_b in the local decision rules. For tractability, assume log utility and i.i.d income. Under these assumptions, we

¹¹For instance, as we document in Section B.3 of the Appendix, for the 2OA solution with baseline calibration: $h_{bb} = 0.004$, $h_{by} = 0.005$, and $h_{yy} = 0.00005$.

¹²Corollary 5 in de Groot (2014) shows that $\frac{\text{var}(c)}{(c^{rss})^2} = \frac{2}{\sigma(1+\sigma)} \frac{1-\beta R}{\beta R}$.

show in Appendix B.3 that h_b in all the local methods is given by the following expression:¹³

$$h_b(\psi, b^*) = \frac{R + e^{b^*\psi}(1 - b^*\psi + \psi) - \sqrt{R^2 + 2e^{b^*\psi}(b^*\psi + \psi - 1)R + e^{2b^*\psi}(1 - b^*\psi + \psi)^2}}{2e^{b^*\psi}}, \quad (9)$$

where $R \equiv 1 + r$ and $b^* = b^{dss}$ for 1OA and 2OA or b^{rss} for RSS. Since the terms h_{bb} , h_{yy} and h_{by} in (8) are quantitatively irrelevant, it follows that $\rho_b(\psi, b^*) \approx h_b(\psi, b^*)$ for 1OA, 2OA and RSS methods. Hence, eq. (9) can be used to analyze how ψ and b^* affect the autocorrelation of the equilibrium stochastic process of NFA produced by the local methods. Moreover, the above expression also implies that the value of h_b obtained with 1OA and 2OA differs from the RSS solution only to the extent that b^{dss} and b^{rss} differ.

Equation (9) formalizes the argument that setting the value of ψ imposes implicitly the equilibrium autocorrelation of NFA, and in particular choosing very low values of ψ implies values of ρ_b close to 1. For given R and b^* , ψ determines ρ_b and in fact, as the numerical results reported below show, ρ_b is decreasing (increasing) in ψ for relatively low (high) ψ . For the 1OA and 2OA methods, eq. (9) describes fully the relationship between ψ and ρ_b , because $b^* = b^{dss}$ and b^{dss} is exogenous, but for the RSS method we need to consider that $b^* = b^{rss}$ and b^{rss} is solved together with the coefficients of the decision rules for \tilde{b}_{t+1} and \tilde{c}_t , which also depend on ψ .

Equation (9) can also be used to illustrate the non-stationarity of the local solutions if a stationarity-inducing transformation is not used. If $\psi = 0$, the solution of $\rho_b(\psi, b^*)$ has two roots, R or 1, so NFA is non-stationary. In contrast (and assuming $b^* = 0$ for simplicity), if $\psi > 0$ the smaller of the two roots that solve $\rho_b(\psi, 0)$ is less than unitary, and thus yields a stable solution.¹⁴

We study numerically how variations in ψ and b^* alter $\rho_b(\psi, b^*)$. To this end, we set $R = 1.059$ as in the baseline calibration and solve for $\rho_b(\psi, b^*)$ or a set of values of ψ and b^* . The results are summarized in Figure 1. The Figure plots $\rho_b(\psi, b^*)$ for ψ in the interval $[0, 0.9]$ and three values of b^* : 0, -0.41 (b^{rss} in the baseline calibration) and -0.51 (b^{dss} in the baseline calibration).

Figure 1 yields a key finding: The value of ρ_b is nearly identical across 2OA and RSS for any $0 \leq \psi \leq 0.1$, which is an interval that includes the baseline and targeted calibration values and also all the values of η^r implied by the ψ values used in the papers listed in Table 1.¹⁵ Hence, for

¹³This result applies for both full and partial RSS.

¹⁴Schmitt-Grohé and Uribe (2003) obtained similar results by deriving the analytic solution of the decision rule of NFA for an endowment economy with log utility and using the stationarity-inducing assumption by which the rate of time preference depends on aggregate consumption.

¹⁵The highest ψ in the literature is $\psi = 2.8$ in Garcia-Cicco et al. (2010), and with their value of $b^{dss} = -0.007$ yields $\eta^r = -0.0196$. For the interval $0 \leq \psi \leq 0.1$ with our $b^{dss} = -0.51$ we obtain an interval of elasticities $0 \geq \eta^r \geq -0.051$. In our baseline and targeted calibrations, $\psi = 0.001, 0.0468$ which implies $\eta^r = -0.0051, -0.0239$.

the values of ψ used in the literature, the choice of approximating around b^{dss} v. b^{rss} and solving with 1OA, 2OA or partial RSS does not make a difference. The two steady-state estimates would have to differ much more than what the baseline calibration and small variations around it would predict. We start to notice a non-negligible difference only if b^{dss} is more than forty percentage points of GDP below b^{rss} . Moreover, since in the baseline and targeted calibrations it is also the case that the terms related to the quadratic and interaction terms of the 2OA decision rule of b are nearly zero, it follows that we can expect the 2OA and RSS solutions to produce similar second and higher-order moments for all endogenous variables, as the results reported below will confirm.¹⁶

Figure 1 also shows that ρ_b switches from decreasing to increasing in ψ at a sufficiently high value of ψ (i.e. $\psi \geq 0.5$). However, $\psi \geq 0.5$ would imply much larger values of η^r than those used in the literature, including estimated models.

The above findings indicate that the implications of ρ_b for ρ_{nx} derived in condition (7) by *assuming* that NFA follows an AR(1) process actually apply to the *equilibrium* stochastic processes produced by the local methods. The DEIR function with very small ψ imposes values of ρ_b near 1, and small differences between these imposed values and the exact solutions of the global solution result in large differences in the predicted values of ρ_{nx} , as we document next.

Table 4 illustrates the relationship between ρ_b and ρ_{nx} in the global (GLB) and local solutions. To induce changes in ρ_b , we vary ρ_z from 0 to 0.8. Panel i) shows the GLB results using the baseline calibration. Panel ii) shows the 2OA and partial RSS results for their baseline calibration with $\psi = 0.001$. Panel iii) shows local solutions for targeted calibrations with ψ set to match $\rho_b = 0.996$ (the value in the GLB solution shown in Table 5) which implied $\psi = 0.0469$ for the 2OA and partial RSS solutions. Panel iv) shows an additional scenario in which, for each value of ρ_b obtained with the GLB solution as ρ_z changes, we re-calibrate the value of ψ in the local solutions so as to match that value of ρ_b (the corresponding values of ψ are also show in this panel).

The first result evident in Table 4 is that 2OA and partial RSS results are always very similar, indicating that the gap between b^{dss} and b^{rss} and the quadratic and interaction terms in the 2OA decision rules are not large enough to alter significantly the values of ρ_b and ρ_{nx} of the two methods under any of the combinations of ρ_z and ψ considered. Recall also that for the same reasons 1OA and 2OA solutions are nearly the same.

Panel i) shows that as ρ_z rises from 0 to 0.8, the GLB solution indicates that the true value of

¹⁶The analytic solution for $h_b(\psi, b^*)$ is strictly valid only for log utility and i.i.d. shocks, but these implications of the analysis still hold quantitatively in the solutions with AR(1) shocks.

ρ_b rises from 0.82 to 0.99 and ρ_{nx} rises from almost -0.1 to 0.85. Thus, as predicted by eq. (7), small variations in ρ_b near 1 cause large changes in ρ_{nx} . In contrast, Panel ii) shows that with $\psi = 0.001$, the local solutions yield values of ρ_b that always exceed 0.99, which in turn yield values of ρ_{nx} ranging from 0.3 to 0.95. The errors relative to the global solutions are large. For $\rho_z = 0$, the true values of ρ_b and ρ_{nx} are 0.83 and -0.1 respectively, while the 2OA and RSS methods yield $\rho_b = 0.99$ and values of ρ_{nx} of 0.27 and 0.24 respectively. For the calibrated value of $\rho_z = 0.597$ estimated from the Mexican data, the global solution yields $\rho_b = 0.977$ and $\rho_{nx} = 0.54$, while the 2OA and RSS methods yield $\rho_b = 0.999$ and $\rho_{nx} = 0.82$ (see Table 5).¹⁷ Thus, these results show that in order for the local methods to yield an accurate approximation to the model's true value of ρ_{nx} to contrast with the data, they first need to produce a value of ρ_b *very* close to the true value.

Panel iii) shows that the local solutions perform better using the targeted calibrations with $\psi = 0.0469$. In this case, GLB and local solutions yield the same ρ_b for the calibrated value of ρ_z (0.597) by construction. For lower values of ρ_z , the local solutions overestimate slightly the values of ρ_b and ρ_{nx} relative to the GLB solutions. Panel iv) shows that, if we re-calibrate ψ as we change ρ_z so that the local solutions match the ρ_b of the GLB solutions in each column of the Table, the local methods do a good job at matching the global solutions. This is true by construction for ρ_b , but the solutions of ρ_{nx} are also close. In these results, however, the value of ψ had to be increased as ρ_z falls, up to a value of 0.185 for $\rho_z = 0$. The required values of ψ range from 0.027 to 0.185. These are significantly larger than the ideal value of 0.001 that keeps the DEIR inessential, and effectively they make deviations of NFA from its steady state very costly. Moreover, knowing the true value of ρ_b that ψ needs to target requires first solving the model with global methods.

b) Long-run moments and precautionary savings

Table 5 shows long-run moments produced by the global and local solutions, including for the latter results from baseline and targeted calibrations and for partial and full RSS.¹⁸ The local solutions under the baseline calibration do poorly at matching the GLB moments. The global solution yields an average NFA position of -41.3 percent of GDP, nearly 10 percentage points higher than the -51 percent at the deterministic steady state (which is the average for the 1OA solution because of certainty equivalence). The 2OA and partial RSS methods yield long-run averages of -28.2 and -45.1 percent of GDP respectively. The former (latter) overestimates (underestimates) precautionary savings by about 13 (4) percentage points. The full RSS solution yields a much lower

¹⁷The 1OA and 2OA solutions are nearly identical, and hence we omit the former from the Table.

¹⁸1OA solutions are not shown because they are nearly the same as the 2OA solutions, except for the averages.

average NFA position of nearly -1120 percent of GDP. This is because this method has the same $\beta R < 1$ of the global method, but it lacks the occasionally binding debt limit φ that allows the global method to simultaneously match the mean of NFA and the variability of consumption observed in the data.¹⁹ Hence, the full RSS generates much higher debt and yields a consumption process with a much lower mean and higher variability than in both the GLB solution and the Mexican data.

Local solutions with baseline calibrations also do poorly at matching the rest of the moments shown in the Table. The three local methods overestimate significantly the variability and persistence of c , nx and b . The correlations with GDP show the opposite pattern. For these higher-order moments, full and partial RSS generate similar values, so the large differences in means do not translate into large differences in higher-order moments. Since partial RSS is much faster and is much closer to the true mean of NFA, we conclude that partial RSS is better than full RSS.

The local methods again perform better at approximating the global results if we switch to targeted calibrations. The major exception is that they do *worse* at capturing the effect of precautionary savings, with both RSS and 2OA solutions yielding an average of b/y of nearly -0.51. These averages are very close to the deterministic steady state, and this occurs because (as explained below) higher ψ is akin to a higher cost of moving b away from its steady state.²⁰ The local solutions with the targeted calibrations also have the drawback that they still overestimate the autocorrelation of consumption, as they did under the baseline calibration. For the rest of the moments, the targeted calibrations deliver a better approximation to the GLB solution than the baseline calibrations, but again they require knowledge of the global solution to determine the target value of ψ and the implied ψ values are much larger than 0.001.

The results in Table 4 showing that the 2OA and partial RSS solutions yield similar ρ_b and ρ_{nx} extends to nearly all the moments shown in Table 5 under both baseline and targeted calibrations. The only exception is the mean of b/y under the baseline calibration, which is -0.28 with 2OA v. -0.45 with RSS, but under the targeted calibration even this moment is nearly the same across 2OA and RSS. This is further evidence indicating that the different centers of approximation in these solutions and the extra terms in the 2OA decision rules have negligible quantitative effects.

Figure 2 provides further evidence of the inaccuracy of the local methods at accounting for precautionary savings. Panels (a) and (b) show how average NFA changes with σ_z and ρ_z under the

¹⁹Full RSS also does not face the NDL, since it yields stationary dynamics around b^{rss} , which makes the effects of the Inada condition on $u'(c)$ irrelevant and the multiplier on the NDL is ignored.

²⁰We could target ψ to match the mean of NFA in the GLB solution (-0.41) instead, but then the local methods do poorly at matching the GLB value of ρ_{nx} . Using this approach, $\rho_{nx} = 0.74$ and 0.88 for the 2OA and RSS solutions respectively, whereas $\rho_{nx} = 0.543$ in the global solution.

baseline calibration. Panels (c) and (d) do the same but for the targeted calibrations. Recall that for 1OA solutions, certainty equivalence implies that mean of NFA is always -0.51 (the deterministic steady state), regardless of the variability and persistence of income and the value of ψ .

These plots yield two important findings. First, the local methods cannot approximate accurately the long-run averages of NFA produced by the GLB solutions in general, and hence they yield incorrect measures of precautionary savings. The continuous, blue curves for the GLB solutions show that the model embodies a strong precautionary savings motive. Panels (a) and (c) ((b) and (d)) indicate that increasing the standard deviation (autocorrelation) of income from 1 to 8 percent (0 to 0.8) increases the long-run average of NFA from -0.5 to near zero (-0.47 to -0.34). In contrast, the local solutions with the baseline calibrations show that the 2OA method overestimates the increase in precautionary savings significantly, with a gap that widens as the variability or persistence of income increase, while the partial RSS method mostly underestimates average NFA, although with a smaller error in absolute value than the 2OA solution. Note also that for high enough ρ_z in Panel (b), partial RSS actually predicts slightly higher mean NFA than GLB.

The second finding is that the local methods with targeted calibrations are worse than the baseline calibrations at capturing precautionary savings. Panels c) and d) show that the average of NFA in the targeted calibrations increases very slightly above b^{dss} as the variability and persistence of income rise. Using $\psi = 0.1$ and 0.2 , the result is even stronger, with mean NFA becoming invariant to income variability and persistence. Moreover, since the terms depending on the quadratic and interaction terms of the 2OA solutions continue to be quantitatively irrelevant, this result suggests that unless ψ is kept very low, the 1OA solutions are nearly the same as the 2OA and RSS solutions in all dimensions, even long-run averages. In addition, the 2OA and RSS solutions also become nearly identical, since b^{rss} becomes very similar to b^{dss} . Thus, while calibrating ψ to match the persistence of NFA in the GLB solution improves the local methods' ability to match second- and higher-order moments, it also removes precautionary savings almost entirely and renders the 2OA and RSS solutions quantitatively consistent with certainty equivalence!

The fact that precautionary savings nearly vanish from the 2OA and partial RSS solutions as ψ rises implies that the terms driving the deviations of the unconditional averages of b from b^{dss} are vanishing too. To shed light on why this happens, we use again the decision rules for log utility and i.i.d. shocks (together with the quantitative result that the quadratic and interaction terms of the

2OA solutions are negligible) to express the unconditional means of the local solutions as:

$$E[b]^{2OA} = b^{dss} + \sigma_z^2 \frac{h_{\sigma_z \sigma_z}}{2(1 - h_b)}, \quad E[b]^{RSS} = b^{dss} + \sigma_z^2 \frac{g_y^2}{2\psi}, \quad (10)$$

where g_y is the coefficient of the linearized consumption decision rule on income. We showed earlier that h_b is decreasing in ψ for $\psi < 0.5$. Hence, the denominator in the right-hand-side of both of the above expressions rises with ψ , which brings the unconditional average closer to b^{dss} . The coefficients $h_{\sigma_z \sigma_z}$ and g_y also depend on ψ , so there are also effects of ψ on mean NFA operating through these coefficients, but the quantitative evidence indicating that for $\psi \geq 0.047$ the mean NFA is about the same as b^{dss} indicates that these effects are either working in the same direction as those operating via the denominators of the expressions, or they are more than offset by them.

The economic intuition for the above result follows from the argument by Schmitt-Grohé and Uribe (2003) showing that the DEIR function is similar to a setup without DEIR but where agents incur a quadratic cost $(\tilde{\psi}/2)(b_{t+1} - b^{dss})^2$ for deviating from steady-state bond holdings. The log-linearized Euler equations of the two formulations become equivalent if we set $\tilde{\psi} = \psi/(1+r)$. Hence, the model with DEIR can be re-interpreted as a model in which agents are penalized for deviating from b^{dss} , and the cost increases with ψ .²¹ Moreover, the cost has variable and fixed components, since it can be decomposed into these two terms: $(\tilde{\psi}/2)(b_{t+1} - 2b^{dss})b_{t+1}$ and $(\tilde{\psi}/2)b^{dss}$. If the fixed cost is larger than the benefit derived from precautionary savings, it would be suboptimal to allow the long-run average of bonds to deviate from b^{dss} . Thus, local solutions using targeted calibrations have the shortcoming that while they can approximate better the true values of second- and higher-order moments by setting ψ high enough to match the true value of ρ_b , it only takes a modest increase in ψ to make precautionary savings nearly vanish and to make 1OA, 2OA and RSS solutions about the same in all dimensions.

The last panel in Table 5 shows performance metrics of the algorithms. All solutions are quite fast. The GLB solution takes 36 seconds but the local solutions are faster at about 1/4th of that time (see the Table footnote for details on the software and hardware used in all the computations). On the other hand, the GLB solution is significantly more accurate, as indicated by the smaller maximum error in the Euler equation and the mean and maximum percent differences in decision rules for NFA and consumption of the local solutions v. the GLB solution. The NFA local decision rules show average (maximum) differences ranging from 7.5 to 12 (12.3 to 37.8) percent.

²¹With DEIR, for $b_{t+1} < b^{dss}$ ($b_{t+1} > b^{dss}$) agents pay more (get less) for borrowing (saving) more.

c) Impulse response functions

Figure 3 provides impulse response functions to a one-standard-deviation income shock. Income dynamics are shown in Panel g). For the local solutions, Panels a), c), and e) (b), d), and f)) show results for the baseline (targeted) calibration. In each plot, the GLB solution is shown in blue. Consumption and output are in percent deviations from their long-run means, while NFA/GDP and nx /GDP are in differences relative to their long-run means (since ratios are already in percent).

The plots show nearly identical impulse responses for 1OA, 2OA and RSS in each calibration case, in line with the results that the h_b coefficients of the bond decision rules are similar and the quadratic and interaction terms of the 2OA solutions are very small. On the other hand, the local impulse responses with the baseline calibration differ sharply from the GLB ones. GLB predicts a weaker decline in NFA (i.e. less borrowing) and much faster mean reversion than the local solutions. Accordingly, consumption falls nearly twice as much on impact in the GLB solution, and continues to decline before recovering, displaying also faster mean reversion. These differences in consumption responses also reflect smaller trade deficits on impact and in the first periods of transition and a faster recovery into trade surpluses with the GLB solution. Local solutions with targeted calibrations yield impulse responses that approximate better the GLB solutions, but still show noticeable discrepancies. In particular, the local solutions now overestimate the fall in consumption on impact.

d) Spectral density analysis

We compare next nonparametric sample periodograms of the various solutions. The goal is to determine whether in addition to the differences in time-series properties, the methods differ in their predictions about the relevance of fluctuations at different frequencies for overall variability.

Figure 4 shows periodograms for simulated data of b , c and nx corresponding to a multivariate spectrum using a Bartlett window with the smoothing parameter set to 100.²² These periodograms are computed based on long time-series simulations including 4500 periods. The y-axis shows the population spectrum, the x-axis shows the frequency in years, and the vertical lines isolate the business cycle frequency. The panels on the left (right) are for the baseline (targeted) calibrations. As in the previous charts, the plots for the GLB solution are identical in both sets of plots, because the global solution has a single calibration. In addition, as with the time-series results, the spectral density functions are nearly identical for 2OA and RSS methods, because the local decision rules have similar h_b terms and the quadratic and interaction terms of the 2OA solutions are irrelevant.

²²We follow Hamilton(1994) in setting the value of the smoothing parameter. The results for parametric estimates of the spectral densities are generally smoother but show similar patterns as those of the nonparametric estimates.

All the periodograms are generally downward sloping because the equilibrium stochastic processes are similar to AR(1) processes. Hence, the contribution of lower frequencies to the variances of the variables exceeds that of business cycle and lower frequencies. The results show, however, that the local methods under the baseline calibration overestimate the contribution of low frequency movements to the total variance of all three series, which is consistent with their slower mean-reversion and higher values of ρ_b relative to the GLB solution. Moreover, while the contribution of fluctuations at the business cycle frequency or higher for the variability of b is slightly higher with the local solutions than in the GLB solution, for nx the local methods overestimate it and for c they underestimate it. In particular, the local methods underestimate significantly the fraction of consumption fluctuations explained by business cycle and higher frequencies and under-predict significantly the contribution of low frequencies.

For targeted calibrations, the periodograms of b are nearly the same in the GLB and local solutions almost by construction, because the targeted calibrations are built to match the AR(1) coefficient of the GLB solution. However, the periodograms of c and nx for the local solutions still differ sharply from the GLB ones. They still underestimate significantly the contribution of consumption fluctuations at business cycle and higher frequencies to overall consumption variance.

e) Interest-rate shocks

We examine next the effects of adding interest-rate shocks. We do this for two reasons. First, because Courdacier et al. (2011) and de Groot (2013) showed that interest-rate shocks play an important role in the quantitative performance of the RSS method, and second to facilitate comparisons of the endowment model results with those of the RBC model, which also has interest-rate shocks. The gross real interest rate is now defined as: $R_t = e^{z_t^R} \bar{R}$, where z_t^R is a shock with exponential support and \bar{R} is the mean interest rate.

The two shocks follow a diagonal VAR representation:

$$\begin{bmatrix} z_t \\ z_t^R \end{bmatrix} = \begin{bmatrix} \rho_z & 0 \\ 0 & \rho_{z^R} \end{bmatrix} \cdot \begin{bmatrix} z_{t-1} \\ z_{t-1}^R \end{bmatrix} + \begin{bmatrix} \varepsilon_t^z \\ \varepsilon_t^R \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma_{\varepsilon^z}^2 & \sigma_{\varepsilon^z, \varepsilon^R} \\ \sigma_{\varepsilon^z, \varepsilon^R} & \sigma_{\varepsilon^R}^2 \end{bmatrix}. \quad (11)$$

where Σ is the variance-covariance matrix of the innovations. Note that if the variance of interest-rate shocks is zero, the endowment shock process is identical to the one used earlier in this Section.

The value of ρ_z is the same 0.597 of the original calibration. We minimize the size of the state space in the GLB solution by using a Markov chain for two shocks defined by the Simple Persistence

Rule, which imposes the same autocorrelation on both shocks. Hence, we set $\rho_{z^R} = 0.597$.²³ The value of $\sigma_{\varepsilon^z}^2$ is set at 0.00069, so that given ρ_z we obtain a standard deviation of the endowment income of $\sigma_z = \sqrt{\sigma_{\varepsilon^z}^2 / (1 - \rho_z^2)} = 0.0327$, which is the value from the Mexican data in the original calibration. For the terms that involve the interest-rate process, we solve the model with values of $\sigma_{\varepsilon^R}^2$ and $\sigma_{\varepsilon^z, \varepsilon^R}$ such that σ_{z^R} takes values ranging from 0 to 2.5 percent and the correlation between endowment income and the interest rate is $\rho_{z, z^R} = -0.669$, which matches the correlation of the interest rate with TFP in Mendoza (2010), and is also the value we use in the calibration of the RBC model in the next Section. The values of $\sigma_{\varepsilon^z, \varepsilon^R}$ and $\sigma_{\varepsilon^R}^2$ change as we change σ_{z^R} , and they are given by: $\sigma_{\varepsilon^z, \varepsilon^R} = (1 - \rho_z \rho_{z^R}) \rho_{z, z^R} \sigma_z \sigma_{z^R}$ and $\sigma_{\varepsilon^R}^2 = \sigma_{z^R}^2 / (1 - \rho_{z^R}^2)$.

The Simple Persistence Markov chain is defined by a set of pairs of realizations of the shocks (z, z^R) and a matrix π of transition probabilities of moving from any realization pair to any other pair in one period. Each shock has two realizations equal to plus/minus one-standard deviation of each shock ($z_1 = -z_2 = 0.0327, z_1^R = -z_2^R = \sigma_{z^R}$, with σ_{z^R} ranging from 0 to 2.5 percent). The Simple Persistence rule produces a 4x4 matrix π with elements defined by a formula such that the standard deviations of the shocks match the realization values, and the unconditional correlation and autocorrelations of the shocks match the values set in the calibration.

With interest-rate shocks, a well-defined limiting distribution of NFA requires $\beta \bar{R} < 1$, otherwise $\beta^t \prod_{j=1}^t R_j$ diverges to infinity (see Chamberlain and Wilson (2000)). In addition, there are realizations, and hence histories of realizations, with R_t lower (higher) than \bar{R} for many periods, which imply much weaker (stronger) precautionary savings incentives than with a constant interest rate set at \bar{R} . This is easier to see for histories of realizations with $\beta R_t > 1$, because they produce equilibrium subsequences where b_{t+1} can grow very large, since there is no pro-borrowing effect due to $\beta R_t < 1$ offsetting the precautionary savings incentive.²⁴ At some point, each of these histories shifts to histories of interest-rate realizations with sufficiently low R_t to induce mean-reversion in NFA. Note also that the NDL is now computed with the highest realization of $R_t - 1$, so it will be tighter than when computed with $\bar{R} - 1$. These effects are at work only in the GLB solution, not in the local ones, because they result from expectations of histories of future shocks that take the economy far from its long-run average and the deterministic steady state.

When using the DEIR function, the interest rate is modified so that interest-rate shocks hit only

²³This is reasonable because in the Mexican data ρ_z and ρ_{z^R} are 0.537 and 0.572, respectively (see Mendoza (2010)).

²⁴Reducing \bar{R} keeping σ_{z^R} constant accentuates these effects, because histories with even larger gaps between β and R_t are possible and with higher probability.

the world interest rate component, not the premium driven by steady state deviations of b_{t+1} :

$$1 + r_t = e^{z_t^r} \bar{R} + \psi \left[e^{b^{dss} - b_{t+1}} - 1 \right] \quad (12)$$

Table 6 shows some of the key moments produced by the different solution methods under baseline and targeted calibrations as σ_{zR} rises from 0 to 2.5 percent. The baseline and targeted calibrations shown in Table 3 are left unchanged. For the GLB solution, we show results with both the calibrated ad-hoc debt limit ($\varphi = -0.51$) and the NDL, with the aim of comparing the roles that the ad-hoc debt limit and interest-rate shocks play in inducing higher mean NFA outcomes, and with the similar effect of interest-rate shocks in local solutions.

Comparing the GLB solution with the solutions under the baseline calibration for partial RSS and 2OA, we find that increasing σ_{zR} from 0 to 2.5 percent has a much stronger effect on mean NFA in the local solutions than in the GLB solution. In the former, $E(b/y)$ increases by 140 (109) percentage points of GDP for the partial RSS (2OA) solution, and actually turns from negative to positive, while in GLB it increases by just about 3 percentage points of GDP. Notice also that the ability of the partial RSS v. the 2OA solutions to generate precautionary savings changes as the variability of the interest rate rises. With low or no interest-rate variability, 2OA generates significantly more precautionary savings ($E(b/y) = -0.285$ v. -0.451), but for interest-rate variability of 2.5 percent the opposite is true ($E(b/y) = 0.806$ v. 0.942). Larger interest-rate shocks also alter the result that the baseline RSS and 2OA solutions have similar second- and higher-order moments.

The above findings suggest that interest-rate shocks in the partial RSS solutions with baseline ψ could be helpful for matching mean NFA, playing the role of NDL in the global solution calibration. This strategy fails, however, because consumption fluctuates too much in all the scenarios for partial RSS and 2OA solutions. All the local solutions shown in Table 6 overestimate the variability of consumption in the GLB solutions by ratios ranging from 1.04 (for partial RSS with targeted ψ and $\sigma_{zR} = 0.5\%$) to 4.01 (for partial RSS with baseline ψ and $\sigma_{zR} = 2.5\%$).

Comparing now GLB solutions v. local solutions with targeted calibrations, we find that the adjustment-cost-like effect of higher ψ bringing mean NFA close to the deterministic steady state still dominates. The local solutions yield small increases in $E(b/y)$ of about 1.8 percentage points (with $\sigma_{zR} < 1.5\%$ there is almost no change) and second- and higher-order moments for RSS and 2OA are again very similar. Hence, the result that higher ψ values neutralize precautionary savings and yield very similar 1OA, 2OA and RSS solutions is robust to adding interest-rate shocks.

Table 6 also shows that full and partial (baseline) RSS solutions do not yield similar second- and higher-order moments once interest-rate shocks are added. The full RSS solution generates sharply higher variability in consumption and NFA, higher autocorrelations in net exports, and very low values of $E(b/y)$. In fact, full RSS is closer to the GLB solution that replaces the ad-hoc debt limit with the NDL than to the baseline or targeted partial RSS solutions. The full RSS solution and the GLB solution with the NDL have, however, the major shortcoming that they produce unreasonably large net debt positions of 3 to 11 times the income of the economy! Moreover, for net interest rates just a notch below the rate of time preference (0.0638), the full RSS method always generates much lower values of $E(b/y)$ than the global solutions with either ad-hoc or natural debt limits. At low interest rates, the RSS solution violates the NDL very often (e.g., for $\bar{R} = 1.01$, NDL is -68.44 while the *average* NFA of the full RSS solution is -69.62). Hence, full RSS performs poorly in general at approximating accurately the long-run average of NFA, even if we remove the ad-hoc debt limit from the GLB solution.

Summing up, the endowment model analysis yields four key findings: First, local solutions for baseline calibrations perform poorly in several key dimensions, such as the mean of net foreign assets, the cyclical moments of consumption and net exports, the impulse response functions to income shocks, and the characteristics of spectral density functions. Second, targeted calibrations perform better, but they are still unable to match some key features of the global solutions, and in order to construct them one needs the exact value of ρ_b from the global solution to calibrate ψ . Moreover, the implied values of ψ (0.0469 for both the 2OA and RSS solutions) are much higher than the 0.001 baseline value used to keep the DEIR function inessential and imply much higher interest-rate elasticities. These high ψ values also make moving NFA from its steady state too costly, thereby causing precautionary savings to vanish. Third, the targeted calibrations also underperform significantly at capturing the effects of parameter variations, particularly those that affect the precautionary savings motive. In order to keep targeted calibrations close to the global solutions, ψ needs to be recalibrated to match the autocorrelation of NFA of each global solution, which makes targeted calibrations impractical. Fourth, all the local methods (1OA, 2OA and partial RSS) yield similar solutions, because second-order terms of the decision rules, other than the variance term, are quantitatively small and the deterministic and risky steady states are not sufficiently different. For $\psi = 0.001$ only the first moments differ, while for $\psi = 0.047$ or higher even the first moments are similar. These results are largely robust to the addition of interest-rate shocks, except that local solutions with $\psi = 0.001$ can generate sizable precautionary savings effects

as interest-rate variability rises. Moreover, we also found that for the baseline ψ value the full and partial RSS solutions yield similar second- and higher-order moments without interest-rate shocks, but not when these shocks are present, and the full RSS always yields excessively large average net debt-income ratios (larger than GLB solutions constrained only by the natural debt limit).

3 Real business cycle model

We compare next local and global solutions for a workhorse RBC model based on those proposed by Neumeyer and Perri (2005), Uribe and Yue (2006) and Mendoza (2010).

3.1 Model structure and equilibrium

As in Mendoza (2010), we characterize the model's competitive equilibrium as the solution to a representative firm-household problem that is akin to a planner's problem, except that the wage rate w_t entering in the calculation of working capital is taken as given by the representative agent and set to satisfy the equilibrium condition that equates w_t with the marginal disutility of labor. The economy produces gross output using a Cobb-Douglas technology that requires capital, k_t , labor, L_t , and imported inputs, v_t :

$$\exp(\varepsilon_t^A)F(k_t, L_t, v_t) = \exp(\varepsilon_t^A)k_t^\alpha L_t^\gamma v_t^\eta, \quad 0 \leq \alpha, \gamma, \eta \leq 1, \quad \alpha + \gamma + \eta = 1, \quad \varepsilon_t^A > 0. \quad (13)$$

Gross output is a tradable good sold at a world-determined price which is the numeraire and is assumed to be constant and equal to 1. The price of imported inputs is also determined in world markets, with the relative price of these inputs in terms of tradable goods given by $p_t = p \exp(\varepsilon_t^P)$, where p is the mean price and ε_t^P is a terms-of-trade shock. There are also TFP shocks, ε_t^A , and interest rate shocks ε_t^R .

A standard working capital constraint requires a fraction ϕ of the cost of L_t and v_t to be paid in advance of sales. Working capital loans are obtained from foreign lenders at the beginning of each period and repaid at the end, so that the financing cost of these loans is the net interest rate $R_t - 1$.

Capital is costly to adjust, with adjustment costs per unit of net investment $(k_{t+1} - k_t)$ determined by the function $\Psi\left(\frac{k_{t+1} - k_t}{k_t}\right) = \frac{a}{2} \left(\frac{k_{t+1} - k_t}{k_t}\right)$, with $a \geq 0$. This functional form satisfies Hayashi's conditions so that the average and marginal values of Tobin's Q are equal at equilibrium.

The representative firm-household chooses $[c_t, L_t, i_t, v_t, b_{t+1}, k_{t+1}]_{t=0}^{\infty}$ so as to maximize:

$$E_0 \left\{ \sum_{t=0}^{\infty} \beta^t \frac{\left(c_t - \frac{L_t^\omega}{\omega} \right)^{1-\sigma}}{1-\sigma} \right\}, \quad (14)$$

subject to this budget constraint:

$$c_t(1 + \tau) + i_t = \exp(\varepsilon_t^A) F(k_t, L_t, v_t) - p_t v_t - \phi(R_t - 1)(w_t L_t + p_t v_t) - q_t b_{t+1} + b_t, \quad (15)$$

with w_t set to satisfy the labor supply condition $w_t = \bar{L}_t^{\omega-1}$, and \bar{L}_t denoting the aggregate labor allocation taken as given by the agent. The only financial asset available is a non-state-contingent discount bond traded in world markets at price q_t . The left-hand-side of the resource constraint is the sum of consumption, inclusive of an ad-valorem tax τ which will be used to calibrate the ratio of government expenditures to GDP, plus gross investment, i_t , where $i_t = \delta k_t + (k_{t+1} - k_t) \left[1 + \Psi \left(\frac{k_{t+1} - k_t}{k_t} \right) \right]$ and δ denotes the depreciation rate of physical capital. The right-hand-side equals total supply, which consists of GDP (gross output minus the cost of intermediate goods, $\exp(\varepsilon_t^A) F(k_t, L - t, v_t) - p_t v_t$) net of foreign interest payments on working capital loans ($\phi(R_t - 1)(w_t L_t + p_t v_t)$) minus (plus) net resources lent (borrowed) abroad ($q_t b_{t+1} - b_t$). The trade balance is therefore $q_t b_{t+1} - b_t + \phi(R_t - 1)(w_t L_t + p_t v_t) = GDP_t - c_t(1 + \tau) - i_t$

The competitive equilibrium is defined by stochastic sequences of allocations $[c_t, L_t, k_{t+1}, b_{t+1}, v_t, i_t]_0^{\infty}$ and prices $[w_t]_0^{\infty}$ such that (a) the representative firm-household solves its optimization problem taking as given the wage rate and the initial conditions (k_0, b_0) , (b) wages satisfy $w_t = \bar{L}_t^{\omega-1}$, and (c) the labor market clears: $\bar{L}_t = L_t$.

We solve this model using the same methods as in the previous Section, extended to include the capital stock as a second endogenous state. For the global solution, we use *FiPIt* with a state space consisting of grids of k and b with 30 and 80 nodes respectively. The algorithm iterates to convergence on the decision rule for bonds and the pricing function for capital. Mendoza and Villalvazo (2019) provide full details, Matlab codes and an Appendix describing how to use them.

For the local solutions, the DEIR function is now defined as:

$$1 + r_t = \exp(\varepsilon_t^R) \bar{R} + \psi \left[e^{\frac{b^{dss}}{y^{dss}} - \frac{b_{t+1}}{y^{dss}}} - 1 \right]. \quad (16)$$

Note that r_t depends on the gap between b_{t+1}/y^{dss} and b^{dss}/y^{dss} . The elasticity of the interest rate

with respect to percent deviations of b_{t+1} from b^{dss} is $\eta^r \approx \psi b^{dss}/y^{dss}$ for small deviations. This facilitates comparisons across calibrations of GLB and local solutions, since output is no longer equal to 1 at steady state and NFA is calibrated to match the NFA-GDP ratio in the data.

3.2 Calibration

Table 7 shows the calibration parameters for the RBC model, most of which were taken from Mendoza's (2010) calibration to Mexican data. The main difference is that we calibrate φ and β in the GLB solution following the same strategy as in the endowment model, by targeting those two parameters to approximate the mean NFA and the standard deviation of consumption in the data. Note, however, that since output is endogenous, steady state GDP cannot be normalized to 1. Hence, we searched over values of β and the lower bound of the NFA grid that yield model solutions close to the target data moments, and then expressed φ as a ratio of that lower bound to y^{dss} . Setting $\beta = 0.92$ and $\varphi = -0.758$ (implied by a lower bound of NFA of -300 and $y^{dss} = 396$) we obtained a mean NFA-GDP ratio of -0.372 (v. -0.44 in the data) and a variability ratio of consumption to GDP of 1.29 (v. 1.25 in the data). For the local solutions, we also proceed as in the case of the endowment model, by setting b^{dss}/y^{dss} to be the same as in the GLB solution and studying baseline and targeted calibrations of ψ . The baseline value is again $\psi = 0.001$ and the targeted values are $\psi = 0.0109$ for the 2OA solution and $\psi = 0.008$ for the RSS solution.

The stochastic process of the shocks is also taken from Mendoza (2010). We characterize the joint process of the three shocks as the following diagonal VAR system:

$$\begin{bmatrix} \varepsilon_t^A \\ \varepsilon_t^R \\ \varepsilon_t^p \end{bmatrix} = \begin{bmatrix} \rho^A & 0 & 0 \\ 0 & \rho^R & 0 \\ 0 & 0 & \rho^p \end{bmatrix} \cdot \begin{bmatrix} \varepsilon_{t-1}^A \\ \varepsilon_{t-1}^R \\ \varepsilon_{t-1}^p \end{bmatrix} + \begin{bmatrix} u_t^A \\ u_t^R \\ u_t^p \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \sigma_{u^A}^2 & \sigma_{u^A, u^R} & 0 \\ \sigma_{u^A, u^R} & \sigma_{u^R}^2 & 0 \\ 0 & 0 & \sigma_{u^p}^2 \end{bmatrix}. \quad (17)$$

In line with the estimates in Mendoza (2010), the above specification imposes the conditions that the co-movement between TFP and interest-rate shocks is driven only by the covariance of their innovations, and the price shocks are independent of the other two shocks. Mendoza reports estimates of the standard deviations of the shocks of $\sigma_{\varepsilon^A} = 0.013$, $\sigma_{\varepsilon^R} = 0.0196$ and $\sigma_{\varepsilon^p} = 0.0335$ respectively. The first-order autocorrelation coefficients are $\rho^A = 0.537$, $\rho^R = 0.572$ and $\rho^p = 0.737$, and the correlation between TFP and R is $\rho_{\varepsilon^A, \varepsilon^R} = -0.669$. Since, as we explain below, the discrete approximation to this VAR system in the global solution requires $\rho^A = \rho^R$, we set the common autocorrelation to the average of the two data estimates, and hence $\rho^A = \rho^R = 0.555$. We impose

the same condition on the VAR representation of the shocks used for the local solutions. Given all these estimates, the elements of the Σ matrix are given by: $\sigma_{u^A}^2 = 1.0273e - 04$, $\sigma_{u^A, u^R} = -0.0047$, $\sigma_{u^R}^2 = 2.4387e - 04$, and $\sigma_{u^P}^2 = 5.1097e - 04$.

In the global solution, the shocks are approximated using symmetric two-point Markov processes defined with the Simple Persistence Rule. These processes consist of a set \mathbb{E} of all combinations of realizations of the shocks $\varepsilon_t = (\varepsilon_t^A, \varepsilon_t^R, \varepsilon_t^P)$, and a matrix π of transition probabilities of moving from ε_t to ε_{t+1} . Each shock has two realizations equal to +/- one-standard-deviation of their corresponding data counterparts: $\varepsilon_1^A = -\varepsilon_2^A = 0.0134$, $\varepsilon_1^R = \varepsilon_2^R = 0.0196$, $\varepsilon_1^P = -\varepsilon_2^P = 0.0335$, so \mathbb{E} contains 8 triples. The Simple Persistence Rule produces an 8x8 matrix π which yields variances, correlations and autocorrelations for all the shocks that match those in the data, except that the procedure requires shocks that are correlated (i.e. ε^A and ε^R) to have the same autocorrelation. As noted above, we set $\rho^A = \rho^R = 0.555$. This restriction is immaterial, because the two shocks have very similar autocorrelation coefficients in the data ($\rho^A = 0.537$, $\rho^R = 0.572$).

3.3 Comparison of quantitative results

a) Long-run moments and performance metrics

Table 8 presents unconditional moments of GLB, 2OA and RSS solutions. 1OA results are omitted because, as with the endowment model, second- and higher-order moments are nearly identical to those obtained with 2OA.²⁵ First, we highlight briefly the differences between the RBC and endowment results in the GLB solutions: The RBC model predicts higher variability of consumption relative to GDP and countercyclical net exports, both of which bring the model closer the data. These changes are due to the presence of the working capital constraint and capital accumulation. The former amplifies the effects of TFP and input price shocks, and induces higher imports of inputs during expansions in response to the countercyclical interest-rate shocks. Capital accumulation also generates an incentive to increase imports and run external deficits during expansions, because of the positive autocorrelation of the three shocks present in the model: “Good times,” in which TFP is high and both input prices and the interest rate are low, have positive persistence, which makes it optimal to borrow from abroad to finance investment due to the expectation that favorable realizations of these shocks will continue to be observed in the near future. The countercyclical net exports due to these effects contributes to the excess variability in consumption relative to GDP.

²⁵As we document in Appendix C.4.1, the first-order coefficients of the decision rules for 1OA and 2OA solutions are identical, and those for higher-order terms in the 2OA solution, except the variance, are negligible.

Compare next the RSS and 2OA solutions under the baseline calibration. The moments for consumption, net exports and NFA differ slightly between these two solutions in the endowment model, but in the RBC model the differences are larger. This is, however, consistent with the arguments presented earlier, because in the RBC model b^{rss} and b^{dss} differ sharply (36 v. -76 in percent of GDP), while in the endowment model the difference was too small to matter. The first-order coefficients of the decision rules (reported in Appendix C.4.1) are again similar for 2OA and RSS, and the second-order coefficients of the 2OA solution yield again negligible effects, but the large difference between b^{rss} and b^{dss} yields larger differences in long-run moments. This is particularly the case for the means of the ratios of net exports and NFA to GDP, which are -4.2 and 73.2 percent respectively in the 2OA solution v. -18.5 and 255.9 percent in the RSS solution.

Comparing local v. global solutions under the baseline calibration, the performance of the former at approximating the GLB solution for the average NFA-GDP ratio worsens markedly in the RBC model v. the endowment model. In particular, while for the endowment model the 2OA and RSS methods produced mean ratios of -0.28 and -0.45, relative to -0.41 in the GLB solution, in the RBC model they produce *positive* ratios of 0.73 and 2.56 respectively (i.e., the economy is a net lender) relative to -0.38 in the GLB solution. Hence, the precautionary savings motive is sharply overstated by the local solutions. This is partly because the RBC model includes interest-rate shocks, and we documented earlier that when these shocks are included 2OA and RSS solutions overestimate significantly the mean NFA position even in the endowment model.²⁶ These findings are also in line with results reported by de Groot (2014), showing large, positive mean NFA-GDP ratios of 3.6 and 41 in the two stable equilibria produced by the full RSS method for an endowment economy.²⁷

For second- and higher-order moments, the results are largely in line with what we observed in the endowment model. In particular, the local solutions overestimate again the persistence of the balance of trade. The GLB solution generates an autocorrelation of net exports around 0.71 whereas both local methods generate values around 0.85. This occurs again because NFA is a near-unit root process and small differences in the autocorrelation of NFA (0.996 in GLB v. 0.999 in 2OA and 0.998 in RSS) imply large differences in the autocorrelation of net exports. Moreover, and also in line with the endowment model results, the local solutions overestimate significantly the variability of consumption, net exports and NFA relative to GDP.

²⁶In the endowment model with $\sigma_{zR} = 2.5\%$, the 2OA (RSS) method produced a mean NFA-GDP ratio of 0.806 (0.942), v. -0.38 in the GLB solution.

²⁷Interestingly, de Groot's analysis showing spurious multiplicity of the full RSS solution shows an additional weakness of this method, namely that it can produce two stable solutions whereas the exact global solution is unique.

Despite the differences in the moments for consumption, net exports and NFA, the cyclical moments for investment, capital, imported inputs, labor and output are similar across the solutions. For investment and the capital stock, this occurs because, as shown in Mendoza (1991), the Fisherian separation of investment from savings and consumption decisions that holds strictly under perfect foresight, holds approximately in the RBC model. Intuitively, the RBC model is in the wide class of models consistent with negligible equity premia, and in the limit with zero premium Fisherian separation holds exactly. In addition, the GHH structure of preferences prevents consumption and savings from affecting labor supply, and hence output and all factors of production. The near-Fisherian separation property is verified in the negligible coefficients of the capital decision rules on lagged NFA in the 2OA and RSS solutions and the near-zero numerical derivatives of the decision rule for $k'(b, k, \varepsilon)$ with respect to b in the global solutions (the largest of which was 0.0064).

Consider next the local solutions with targeted calibrations, for which matching the GLB value of ρ_b required $\psi = 0.0109$ and $\psi = 0.008$ in the 2OA and RSS solutions, respectively (see Table 7). These are considerably smaller than the value needed for the targeted calibrations of the endowment economy (0.0469 for both 2OA and RSS). These differences, together with the different NFA-GDP ratios in the deterministic steady states of the endowment and RBC models, imply values of η^r of 0.0083 and 0.0061 for the 2OA and RSS solutions of the RBC model respectively, lower by a factor of 3 than the 0.0239 for the endowment model solutions. This is the case mainly because the GLB solution of ρ_b is higher in the RBC than in the endowment model (0.996 v. 0.977).

The lower ψ values under the targeted calibrations of the RBC model v. the endowment model also imply that the mean of NFA can now rise above the deterministic steady state by non-trivial margins, because the implicit cost of deviating from b^{dss} is smaller. This is even more the case for the targeted RSS solution, which has a lower ψ than the 2OA solution and thus allows $E(b/y)$ to rise by more (-0.397 v. -0.62 in the 2OA solution and -0.758 in the deterministic steady state).

The gap between b^{rss} and b^{dss} in the RSS and 2OA solutions narrows in the targeted calibrations relative to the baseline calibrations: b^{rss} is now -0.591, compared with -0.758 for b^{dss} . With this smaller difference, we recover the result that the decision rules for RSS and 2OA yield similar second- and higher-order moments and similar impulse response functions, as shown below.

As in the endowment model, targeted calibrations generally yield moments closer to the GLB solution than baseline calibrations. It is still the case, however, that in order to target the calibration of the local solutions we need to know the GLB solution for ρ_b . The targeted RSS solution performs markedly better than the 2OA solution in that it yields a mean NFA-GDP ratio much closer to

the GLB solution. 2OA yields a ratio of -0.62, nearly 24 percentage points of GDP lower than the true value (-0.38), whereas the RSS solution yields about -0.4, just 1.3 percentage points below the true value. The targeted RBC calibrations, however, do not get as close to the GLB solution moments as in the case of the endowment model, even with the RSS solution: The variability of the NFA-GDP ratio is roughly half of what the GLB solution yields and its correlation with GDP is 3.5 times bigger. The leverage ratio is also much less variable and has a much lower correlation with GDP. Fisherian separation continues to approximately hold, so moments for output, investment, and factors of production are similar in the targeted calibrations and the GLB solution.

In terms of execution times, the local solutions are still faster than the GLB solution but with a much smaller margin than in the endowment model solutions. The local solutions take about 2/3rds of the time taken up by the GLB solution (which takes 61 seconds), instead of 1/4th for the endowment model. The 1OA solution has a similar execution time, and as explained earlier yields similar second- and higher-order moments. The accuracy of local method results show similar limitations as in the endowment model: the local solutions yield significantly larger Euler equation errors (for capital and NFA with RSS and for capital with 2OA), the average (maximum) differences in the decision rules for k and c are in the 1.7-1.9 (5.1-6.6) percent range, and those for NFA are much larger at above 8 (50) percent for the average (maximum) respectively. These large differences occur at the debt limit φ , because the local methods do not handle it as occasionally binding.

We conducted a robustness analysis of the results reported here by altering the values of some of the model's key parameters. The details are provided in Section D of the Appendix. We examined scenarios increasing the variability of TFP, input price and interest-rate shocks one at a time, as well as increasing the coefficient of relative risk aversion, the correlation between interest rate and TFP shocks, and the subjective discount factor. As in the case of the endowment model, local solutions with a fixed value of ψ calibrated to match ρ_b in the baseline GLB solution are not useful for analyzing the effects of any of these parameter changes, because they yield solutions that differ sharply from the GLB solutions for the same parameter variations. In particular, the local solutions continue to perform poorly at capturing precautionary savings effects (i.e. the true solution for the mean of the NFA-GDP ratio differs sharply from what the local solutions yield).²⁸ In addition, the local solutions underestimate the variability of NFA and net exports, overestimate (underestimate) the correlations of NFA and consumption (net exports) with GDP, and underestimate the autocor-

²⁸The only exception was the experiment doubling the variability of input price shocks. This has minor effects because imported inputs are only 10 percent of gross output, so that their share in GDP net of working capital is 11 percent ($0.1/(1 - 0.1) = 0.11$). Thus, rising $\sigma_{\epsilon P}$ from 3.4 to 6.8 percent increases net income variability just a notch.

relation of net exports. The local solutions are closer to the GLB solutions if we re-calibrate ψ to target the new value of ρ_b from the GLB solution for each new parameterization, but this implies having obtained the GLB solution first and in addition the long-run moments are not as close to those of the GLB solution as in the case of the endowment model.

b) Impulse response functions

Figures 5 and 6 show impulse response functions to a negative one-standard-deviation TFP shock for the GLB and local solutions (baseline and targeted).²⁹ We plot only the first 100 periods to highlight the differences across the solutions. All impulse responses return to zero in about 500 periods. As in the endowment model, impulse responses for 1OA and 2OA are nearly identical under baseline and targeted calibrations. This occurs because again the first-order coefficients of decision rules are identical, and the second-order terms (other than the variance terms) are quantitatively irrelevant. Hence we omit the plots for the 1OA results.

For the baseline calibration, RSS yields markedly different responses for NFA-GDP ratio (Panel a.), consumption (Panel b.) and the net exports-GDP ratio (Panel c.) than 2OA. This is because, as noted earlier, the gap between b^{rss} and b^{dss} is large enough to affect the results. However, since near-Fisherian separation still holds, the other variables (capital, investment, labor, imported inputs and GDP) display similar responses in the two local solutions.

2OA and RSS baseline impulse responses differ sharply from those of the GLB solution. In particular, RSS overestimates the initial rise in the NFA-GDP ratio while 2OA underestimates it (see panel a.). In fact, RSS yields above-average NFA-GDP ratios for the first 17 periods, while in both GLB and 2OA the NFA position is always below average. After the 15th period, the two local solutions predict larger mean deviations of the NFA-GDP ratio than GLB. The local solutions remain uniformly above the GLB solution until mean reversion is attained. These differences in NFA are reflected in differences in consumption and net exports (see panels b. and c.). Initially, the mean deviation of consumption under the GLB solution is lower (higher) than in the RSS (2OA) solution. After the 30th period, the GLB solution yields significantly smaller mean deviations of consumption than the two local solutions and the opposite is observed for the net exports-GDP ratio. Differences in investment, output and factors of production are less noticeable because near-Fisherian separation holds, but still capital falls slightly more initially in the GLB than in the local solutions, and then between periods 15 and 80 the GLB solution rises slightly above the local ones. Since there is no wealth effect on labor supply, these differences in capital stock dynamics translate

²⁹We also computed impulse response functions for interest-rate and input price shocks in Appendix D.

into qualitatively similar but quantitatively smaller differences in labor, imported inputs and GDP.

Under the targeted calibrations (Figure 6), the gap between b^{rss} and b^{dss} becomes again too small to make a difference for the 2OA and RSS impulse responses. Hence, our findings for the endowment and RBC models indicate that, if the choice is limited to local methods, a 1OA solution is simpler and nearly identical to 2OA and RSS solutions. Relative to the GLB solution, the targeted local solutions still fail to match important features of the GLB impulse responses. Initial differences are smaller than with the baseline solutions, and now the NFA-GDP ratio always shows negative deviations from its mean in all three solutions. Beyond the 15th to 20th period, however, NFA-GDP, consumption and the net exports-GDP ratio in the targeted solutions differ sharply from the GLB solution, with similar qualitative features as with respect to the baseline solutions, and in some cases with even larger quantitative differences. The reason for this is that, even though the targeted calibrations force the same ρ_b across GLB and local solutions, the required higher values of ψ imply that NFA has much less variability than in the GLB solution (see Table 8). The higher volatility with similar persistence in the GLB solution yield an impulse response for NFA-GDP that rises more initially and then drops more before returning to zero in the long run. In contrast, in the local solutions the high ψ values make large deviations of NFA-GDP from its mean too costly, and hence NFA-GDP never falls more than about 2 percentage points below its mean (v. about 7 percentage points in the GLB solution).

c) Spectral density functions

Figures 7 and 8 show nonparametric periodograms for the key macroeconomic aggregates of the RBC model produced by the global and local methods. As in the case of the endowment model, all the periodograms are downward sloping, indicating that lower frequencies contribute more to the variability of the simulated data than business cycle and higher frequencies. In contrast with the endowment economy, however, the periodograms for NFA, consumption and net exports produced by the 2OA and RSS solutions under the baseline calibration are different, because in the RBC baseline calibration the gap between b^{rss} and b^{dss} is large enough for 2OA and RSS results to differ. The other periodograms for 2OA and RSS are similar because of the near-Fisherian-separation property noted earlier. Relative to the GLB solution, 2OA and RSS periodograms show differences that are less stark than for the endowment model, but RSS still overstates the contribution of business cycle and higher frequencies to the variability of NFA, and 2OA and RSS still overstate the contribution of very low frequencies to the variability of NFA, consumption and net exports, as well as the contribution of business cycle and higher frequencies to the variability of GDP.

For the targeted calibrations, the 2OA and RSS periodograms are nearly identical, reflecting the result that in this case the gap between b^{rss} and b^{dss} is too small to affect the results. Relative to the GLB solution, both RSS and 2OA yield periodograms that approximate their GLB counterparts better than under the baseline calibration, in line with what we found for the endowment model. The local methods underestimate slightly overall NFA and GDP variability. The periodograms for investment and factors of production are very similar to those under both the GLB solution and the local baseline calibrations, again because of the near-Fisherian-separation property.

In summary, the analysis of the RBC model yields several key results in line with those obtained for the endowment model: Local methods do poorly at quantifying the effects of precautionary savings. Local methods with baseline calibrations yield very different results than the global solution for consumption, net exports and NFA. Targeted calibrations perform better but in order to target the value of ψ it is necessary to solve the model globally to find the exact solution for ρ_b , and this needs to be re-done for any parameter variation. 1OA and 2OA solutions yield nearly identical results (other than first moments), because they have identical first-order terms and the second-order terms of the 2OA solution (other than the variance term) are quantitatively irrelevant.

The RBC results differ from the endowment model results in that 2OA and RSS solutions with baseline calibrations differ significantly, because b^{rss} and b^{dss} differ enough to yield non-negligible differences in first-order coefficients of the decision rules. In the targeted calibrations, however, b^{rss} and b^{dss} are close again, and hence 2OA and RSS solutions are very similar. Thus, with targeted calibrations, 1OA, 2OA and RSS solutions differ only in their first moments, while higher-order moments, impulse responses and spectral density functions are nearly identical. This makes 1OA the preferable local method if first moments are not being studied. A second important difference relative to the endowment model results is that the targeted local solutions are less accurate at approximating the GLB solution results for NFA, consumption and net exports. This is because the required ψ values make fluctuations in NFA costly and this reduces NFA variability to about half of what the GLB solution yields. Results for investment, output and factor allocations are similar across local and GLB solutions because Fisherian separation of savings and investment nearly holds.

4 Sudden Stops Model

This Section compares local v. global solutions for the SS model proposed by Mendoza (2010), which adds to the RBC model an occasionally binding credit constraint. Debt and working capital

financing cannot exceed a fraction κ of the market value of capital:

$$q_t^b b_{t+1} - \phi R_t (w_t L_t + p_t v_t) \geq -\kappa q_t k_{t+1}. \quad (18)$$

4.1 Solution Methods

For the GLB solution, we use again the *FiPit* algorithm. In each iteration, the algorithm assumes first that the constraint does not bind, solves for allocations and prices using the model’s optimality conditions and then evaluates the constraint. If it binds the results are discarded and new allocations and prices are solved for with the constraint holding with equality. The algorithm iterates over three recursive functions of the state variables: the NFA decision rule, the price of capital, and the Lagrange multiplier on the borrowing constraint (see Mendoza and Villalvazo (2019) for details).

For the local solution, Appendix Sections C.4 and C.5 describe the model formulation and the DynareOBC method. This method treats the occasionally binding constraint as a source of endogenous news about the future along perfect-foresight paths. If the constraint is (is not) binding at the deterministic steady state, it uses news shocks to solve for unconstrained (constrained) periods along those paths by representing the solution as a mixed integer linear programming problem. For instance, if the constraint does not bind at steady state, when agents anticipate that the constraint will bind at some date $t + j$ conditional on the date- t state variables and the deterministic evolution of the exogenous shocks, this provides “news” that bond holdings will follow a path in which they are higher than otherwise. This approach is akin to assuming that agents live in a world without the constraint, but whenever they are on a path that would lead them to borrow above what the constraint allows, a series of news shocks hit that makes them borrow only what is allowed and moderate their borrowing before that happens.³⁰

The main output of DynareOBC is a time-series simulation constructed by stitching together the date- t values of perfect-foresight equilibrium paths conditional on (k_t, b_t) and the date- t realizations of the exogenous shocks. Each path is obtained using an extended path algorithm that traces equilibrium dynamics up to T periods ahead of t , with the shocks following their deterministic VAR dynamics. The extended path can be obtained using first- or higher-order approximations, but we report here results based on the former.³¹ The path computed for a given starting date t

³⁰The model with the constraint is approximately equivalent to the same model without the constraint but with sequences of news shocks chosen to yield the same equilibrium as the model with the constraint. This equivalence holds exactly if the model is linear and shock variances are zero, such that any shocks that occur are truly “unexpected.”

³¹DynareOBC is equivalent to the Guerrieri-Iacoviello OccBin method when using a first-order approximation, with the difference that it is guaranteed to converge in finite time. Holden (2016b) showed that a second-order

determines the values of the endogenous states (k_{t+1}, b_{t+1}) and these together with the realizations of the shocks at t and the optimality conditions determine the date- t values of all the endogenous variables. The rest of the path is discarded and the process is repeated at $t + 1$ to generate the values of the time-series simulation for that period. The efficiency of this method hinges on three factors: (a) T needs to be large enough so that for $t > T$ no further news shocks are needed (if the constraint does not bind at the deterministic steady state, T needs to be large enough so that the constraint never binds again, and if it binds at steady state T needs to be large enough so that it always binds), (b) for each path requiring news shocks, the algorithm needs to find the sequence of news shocks that supports the correct equilibrium path, and (c) the time-series simulation needs to be long enough for long-run moments of the endogenous variables to converge. The algorithm is less efficient in models with very persistent dynamics, which require a large T and a long simulation length, and models in which the news shocks are needed frequently.

Figure 9 illustrates the DynareOBC method in a simple example based on the endowment model of Section 2 with its ad-hoc debt limit and using the DEIR function so that this limit does not bind at the deterministic steady state (see Appendix C.5.1 for details). Panels (a) and (b) show the solutions of consumption and bond holdings for $t=100$ to 250 (black, solid curves) and eleven of the perfect-foresight paths (red, dashed curves) that generated them, with the corresponding date- t solution marked with a red circle. In panels (b) and (d), the green, dashed lines indicate the ad-hoc debt limit. The constraint never binds in seven of the perfect-foresight paths shown in Panel (b) and in four it does. Hence, endogenous news shocks are needed only in the latter.

Panels (c) and (d) isolate periods $t=140$ to 180 and show the extended path that generates the equilibrium values at $t=141$ (red, dashed curve). DynareOBC computes a sequence of news shocks that sustains this path as an equilibrium. The comparable path of NFA in the solution without credit constraint is also provided in Panel (d) (black, dashed curve). Panel (d) shows that the constraint first becomes binding along the perfect-foresight path at $t=144$. It also shows that, relative to the model without a credit constraint, agents choose higher bond holdings (less debt) earlier, in anticipation of the constraint becoming binding with perfect foresight (i.e. the red, dashed curve is above the black, dashed curve at $t=142,143$). Income is rising gradually on its path back to its deterministic steady state, so that the constraint continues to bind for several periods, until

approximation integrating over future uncertainty can be used to approximate precautionary savings in models with simple constraints, but this method is significantly slower and for the SS model produced results that deviate sharply from the GLB and first-order DynareOBC solutions. In particular, investment and the net exports-GDP ratio had negative serial autocorrelation and NFA had near-zero autocorrelation.

income is sufficiently high for NFA to also start rising back towards its steady state (after $t=170$).

For comparing v. the GLB solution, it is critical to note that first-order DynareOBC ignores the *risk* of hitting the constraint and fluctuating across states where it binds or not. At each date t , it does not consider the histories of future shocks and associated equilibrium allocations and prices that can occur, it only considers the news-shock-adjusted perfect-foresight path computed for t and the date- t shock realizations. Hence, wealth and precautionary-saving effects of the occasionally binding constraint are ignored, and forward-looking objects like asset prices and excess returns also abstract from them. This is important in SS models, because a financial crisis with a deep recession and collapsing prices occurs when the constraint binds, and the risk of these events strengthens precautionary savings and alters asset prices even in “good times” (see Mendoza (2010), Durdu et al. (2009)). In contrast, in each DynareOBC perfect-foresight path, if the constraint binds (does not bind) at the deterministic steady state, agents anticipate reaching a long-run equilibrium in which the constraint binds (does not bind) regardless of whether it binds or not at t . They anticipate deterministically hitting (escaping) the constraint at some date $t + j$ if unconstrained (constrained) at t and adjust their optimal plans before $t + j$ accordingly, but these plans do not factor in the *risk* of the constraint becoming binding or non-binding, and how this risk affects incentives to self-insure, asset prices and the equity premium.

4.2 Calibration

The calibration is very similar to that of the RBC model (see Table 7). The value of κ is set to 0.2, as in Mendoza (2010). For the GLB solution, we increase the lower bound of the bonds grid to -200 (i.e. a tighter ad-hoc debt limit), which is about about half of steady state GDP. We do this because strong precautionary savings due to the credit constraint imply that NFA-GDP ratios below -0.192 are never observed in the ergodic distribution of the SS model, and hence we can set φ to a higher value so as to use fewer nodes in the bonds grid to make the algorithm more efficient.

Before discussing the DynareOBC calibration, it is important to note that if the constraint binds at the deterministic steady state, the steady-state equilibrium is well defined and there is no need to induce stationarity. This follows from the steady-state bonds Euler equation:

$$1 = \beta R + \frac{\mu_{ss}}{u'(c_{ss})}, \quad (19)$$

where μ is the Lagrange multiplier on the borrowing constraint. Note that $\beta R < 1 \iff \mu_{ss} > 0$,

hence having the constraint bind at steady state requires $\beta R < 1$ and viceversa. This Euler equation is solved together with the other steady-state equilibrium conditions (particularly the steady-state resource and credit constraints) to solve for the steady-state equilibrium, including c_{ss} and μ_{ss} .

We study DynareOBC solutions with both $\mu_{ss} > 0$ and $\mu_{ss} = 0$. The first is labeled “DynareOBC- $\beta R < 1$ ” and the second “DynareOBC-DEIR,” because in the latter case the DEIR function induces stationarity. For the DynareOBC- $\beta R < 1$ case, β is the same as in the GLB solution, which ensures that the deterministic steady state is the same. Hence, in this case DynareOBC and GLB use identical calibrations. For the DynareOBC-DEIR case, the target NFA position in the DEIR function is set so that the deterministic steady state of the NFA-GDP ratio matches the long-run average obtained in the GLB solution, with $\beta = 1/R$ (which is required for $\mu_{ss} = 0$) and $\psi = 0.001$ (which is the inessential value). The rationale for looking at this scenario is that in the GLB solution the constraint binds rarely and mean NFA is much higher than at the deterministic steady state. Hence, a local approximation around an unconstrained steady state would be more in line with the unconstrained nature of the GLB long-run equilibrium solution.

4.3 Comparison of quantitative results

a) Long-run moments, impulse response functions & performance metrics

Table 9 shows that several moments of the DynareOBC solutions differ from their GLB counterparts, with smaller differences for supply-side variables as was the case for the RBC model. Long-run averages show the largest differences, particularly for c , nx/y , b/y and the leverage ratio. The large differences in $E[b/y]$ indicate that the local solutions continue to perform poorly in this dimension. $E[b/y] = 1.5$ percent in the GLB solution, relative to -37.2 percent in the RBC model, indicating that the credit constraint strengthens precautionary savings sharply. DynareOBC- $\beta R < 1$ (DynareOBC-DEIR) underestimates (overestimates) this result significantly, yielding $E[b/y] = -10$ (20.6) percent. Differences this large have important implications for key research questions. For example, quantifying the optimal amount of foreign reserves to manage the risk of Sudden Stops requires determining first how the economy responds to this risk without policy intervention (e.g. Durdu et al. (2009)). DynareOBC would yield findings that deviate sharply from the correct result.

Several of the moments that are underestimated in DynareOBC- $\beta R < 1$ v. the GLB solution tend to be overestimated in the DynareOBC-DEIR solution. Relative to the GLB solution, DynareOBC- $\beta R < 1$ (DynareOBC-DEIR) yields markedly lower (higher) variability in NFA, net exports, consumption, and leverage, lower (higher) correlations of GDP with investment, net exports,

and intermediate goods, and lower (higher) persistence in consumption and net exports. Interestingly, DynareOBC-DEIR, which was calibrated so that b^{dss}/y^{dss} equals $E[b/y]$ in the GLB solution also yields a value of $\rho(b)$ close to the GLB solution (0.995 v. 0.99). Hence, it is in line with the notion of the “targeted calibrations” of the previous Sections even though we kept $\psi = 0.001$.

Local and global solutions also differ sharply in that the constraint binds much more often in the former than in the latter (51.8 and 71.1 percent in the two local solutions v. 2.6 percent in the GLB solution). This is due in part to the fact that agents do not respond to the *risk* of the constraint becoming binding in the first-order DynareOBC solution. However, recall that agents do anticipate the constraint becoming binding in the perfect-foresight paths for which this happens and borrow less before hitting the constraint. Hence, although the constraint binds nearly half the time with DynareOBC, it does so with very small multipliers. Later in this Section we study in more detail the extent to which the local and global solutions differ when the constraint binds.

The long-run averages of NFA in the DynareOBC solutions are higher than the corresponding deterministic steady states, indicating that certainty equivalence does not hold, even though the perfect-foresight paths are first-order approximations: b^{dss}/y^{dss} is set at -0.192 (0.015) in DynareOBC- $\beta R < 1$ (DynareOBC-DEIR) while $E[b/y] = -0.1$ (0.206). This is not, however, due to precautionary savings, since DynareOBC does not take into account the risk of future shocks and the constraint becoming binding. Instead, this occurs because the constraint induces asymmetric (nonlinear) responses to shocks even without risk. This feature of DynareOBC can be illustrated using again the endowment model example of Figure 9 (see Appendix Section C.5.1 for details). A negative endowment shock that causes the constraint to bind along the perfect-foresight path determining the date- t value of the solution reduces bond holdings by less than the increase in bond holdings in response to a positive shock of the same size in absolute value. As a result, upward movements in b when positive shocks hit are larger than downward movements when negative shocks hit if the economy is near or at a point where the constraint binds. Moreover, b cannot move below the lower bound set by the constraint but it can wander off to high values after sequences of positive shocks (see again Figure 9). Hence, the outcome is a DynareOBC time-series with observations “biased” above the deterministic steady state, which implies a mean above b^{dss}/y^{dss} .³²

The GLB solution features a similar asymmetry but in addition it takes into account precautionary savings effects due to the risk of future shocks and the constraint becoming binding. It does not follow, however, that the DynareOBC solutions must always yield mean bond positions lower than

³²Recall that the constraint in this example is a fixed debt limit while in the SS model it depends on $q_t k_{t+1}$.

the GLB solution. DynareOBC- $\beta R < 1$ (DynareOBC-DEIR) yields a significantly lower (higher) mean of NFA. Both of these results abstract from precautionary savings, but in the DynareOBC-DEIR solution b^{dss}/y^{dss} is set equal to the value of $E[b/y]$ in the GLB solution (0.015) and the constraint does not bind at steady state. Hence, the DynareOBC solution is “biased” above 0.015 and therefore it must yield a long-run average higher than that value.

If one had to choose between Dynare- $\beta R < 1$ and Dynare-DEIR, the former is preferable. Both yield long-run moments that differ from the GLB solution, as shown in Table 9, but as we show later in this Section, Dynare- $\beta R < 1$ does better at approximating the effects of the collateral constraint. Moreover, Dynare- $\beta R < 1$ uses the exact same calibration as the GLB solution and does not require a stationarity-inducing transformation. In contrast, Dynare-DEIR requires setting a value for b^{dss}/y^{dss} . We used the average of the GLB solution, but this resulted in a much higher mean for the NFA-GDP ratio and requires prior knowledge of the mean in the GLB solution.

We compare next performance metrics.³³ The main result is that the speed advantage of the local methods shrinks considerably, particularly for DynareOBC- $\beta R < 1$. Speed ratios relative to the GLB solution rise to 0.901 for DynareOBC- $\beta R < 1$ and 0.70 for DynareOBC-DEIR (the GLB solution takes 268 seconds). Relative to the local RBC solutions using standard 10A and 20A algorithms, the DynareOBC execution times are considerable higher at 244 and 188 seconds for DynareOBC- $\beta R < 1$ and DynareOBC-DEIR, respectively. This is due to the three determinants of the efficiency of DynareOBC noted earlier together with the fact that NFA follows a near-unit-root process. Each extended path required at least 60 periods and the time-series simulations needed 100,000 periods to converge to invariant moments, particularly those for b/y , nx/y and the standard deviation of consumption.³⁴ DynareOBC- $\beta R < 1$ is slower than DynareOBC-DEIR because it required more searches to construct the news shocks sequences that implement the constraint.

The speed comparisons of DynareOBC v. *FiPit* need to be pondered carefully. On one hand, *FiPit* suffers from the standard curse of dimensionality of GLB methods related to the number of state variables, and more so if the model specification requires using a root-finder when the constraint binds.³⁵ But once the decision rules are solved for, generating stochastic time-series simulations is very fast. On the other hand, the number of state variables is much less of an issue

³³See footnote to Table 5 for the hardware and software used to record the performance metrics.

³⁴Intuitively, consider that the estimators of the mean and autocorrelation of an AR(1) process are consistent but biased in finite samples. The bias is higher the higher the true autocorrelation but it falls as the sample size rises. For a near-unit-root process, the sample needs to be quite large to make the estimation bias negligible.

³⁵As Mendoza and Villalvazo (2019) explain, this is not needed for several specifications of credit constraints. This is the case in the SS model without working capital in the constraint, which reduces the *FiPit* run time by 57 percent.

for DynareOBC, but execution time rises with the required length of perfect-foresight paths, the length of the time-series simulation needed for convergence of long-run moments, and the iterations required to compute the news-shocks sequences that implement the constraint. As shown in Table 5 of Appendix D, DynareOBC- $\beta R < 1$ is much slower than *FiPIt* if the simulation length rises to 150,000 periods (350 seconds v. 268 seconds) or with fewer exogenous shocks so that the curse of dimensionality is less severe (with TFP shocks only, *FiPIt* solves in 42 seconds v. 230 seconds with DynareOBC). Relaxing the credit constraint by setting κ to 0.3 also results in *FiPIt* solving much faster than DynareOBC (137 v. 228 seconds).³⁶ Moreover, solving in second- or higher order and/or adding the option to integrate over future uncertainty also slows down DynareOBC considerably. Holden (2016b) noted that solving a model similar to the endowment model of Section 2 using first-order DynareOBC integrating over 45 periods of future uncertainty with the highest accuracy required 2,855 seconds (with different hardware than we used).

In terms of accuracy, *FiPIt* produces again very accurate GLB results, as indicated by the small maximum errors of the bonds and capital Euler equations. The accuracy of the DynareOBC solutions cannot be assessed this way because the solution method produces a time-series simulation, instead of decision rules. Hence, we follow Holden (2016b) and evaluate their accuracy by constructing consumption time-series simulations of the GLB solution for the same initial conditions and sequence of shocks as in each of the two DynareOBC solutions, and computing the maximum absolute values of the differences across them. The maximum differences in log base 10 are 1.292 percent with DynareOBC- $\beta R < 1$, and 1.342 percent with DynareOBC-DEIR. These are significantly larger than Holden’s estimates with for the endowment model.

Figure 10 shows impulse responses to a one-standard deviation negative TFP shock. The impulse responses for the GLB solution are conditional on starting at the long-run averages of capital and NFA, and those for DynareOBC solutions on starting at the deterministic steady states (which for DynareOBC-DEIR are the same as the GLB averages). The GLB impulse responses are very similar to those of the RBC model, because the constraint binds very infrequently in the left tail of the ergodic distribution. Hence, in the “correct” SS model solution, the responses of macro variables triggered by shocks of standard magnitudes starting off from long-run means are nearly unaffected by the credit friction. The means of variables other than the supply-side variables change

³⁶Using DynareOBC also poses logistical hurdles. Since it relies on Dynare, updates to Dynare can make older versions of DynareOBC inoperable, and some versions of Dynare operate only in certain operating systems and software environments. For instance, the DynareOBC toolbox we used operates with Dynare 4.4.3 and only with Matlab2016a. Dynare 4.4.3 operates with the Ubuntu 14.04 Linux operating system but not with Ubuntu 18.04.

across the RBC and SS GLB solutions, but the deviations from means triggered by identical shocks are nearly unchanged. In contrast, the DynareOBC impulse responses are very different from the local-methods impulse responses of the RBC model, even those for the 1OA solution.

DynareOBC impulse responses of several variables differ sharply from their GLB counterparts. For DynareOBC- $\beta R < 1$, Panel (a) shows that NFA/GDP hardly moves, and Panel (b) shows that nx/y moves into a surplus on impact, because of reduced demand for imported inputs. This occurs because the constraint binds at date 0 and the TFP shock tightens the constraint more. For DynareOBC-DEIR, NFA/GDP declines, offsetting the fall in imported inputs to yield an almost unchanged trade balance. In contrast, in the GLB solution net exports jump on impact nearly twice as much as under DynareOBC- $\beta R < 1$ and NFA/GDP rises gradually to peak roughly 150 basis points above its mean, and after that it falls slowly to a trough 400 basis points below its mean before gradually reverting to its mean. The responses of capital are also very different. Both DynareOBC solutions yield a decline on impact, while in the GLB solution capital is nearly unchanged. Then capital declines slightly and starts recovering in the two local solutions, while in the GLB solution it falls by nearly three times as much reaching nearly 1.5 percent below average before starting to recover. Qualitatively, the responses of consumption, investment, labor, inputs and GDP are similar in all three solutions, but the declines on impact are significantly larger in the GLB solution.

b) Spectral density functions

Figure 11 shows the nonparametric periodograms for the DynareOBC and GLB solutions. As in the endowment and RBC models, since all of the variables follow AR(1)-like processes, the periodograms are generally downward sloping, indicating that low frequencies account for a larger fraction of the variance of the variables than business cycle and higher frequencies. Moreover, the periodograms for the GLB solution are very similar to those for the RBC GLB solution, in line with the finding that the GLB impulse responses of the RBC and SS models are similar because the credit constraint binds infrequently. The periodograms of the DynareOBC solution differ from those of the 2OA and RSS solutions of the RBC model, so the local methods fail to match the property of the GLB solutions that spectral densities of the RBC and SS models are similar.

The DynareOBC periodograms for NFA, consumption, net exports, investment and labor differ sharply from the GLB results. Consumption has the highest variance in the GLB solution (121.4), followed by DynareOBC-DEIR (119.2) and DynareOBC- $\beta R < 1$ with a much lower variance (65.8). In contrast, the autocorrelation of consumption is highest in DynareOBC-DEIR (0.91) and about the same in GLB and DynareOBC- $\beta R < 1$ (0.83). As a result, the consumption periodograms

for the latter two have the same intercept but the one for GLB is uniformly higher otherwise, while the periodogram for DynareOBC-DEIR has the highest intercept but is generally below the periodogram for GLB. The DynareOBC solutions assign significantly less consumption variability to business cycle and lower frequencies than the GLB solution. Net exports also show higher persistence in the DynareOBC-DEIR solution while DynareOBC- $\beta R < 1$ and GLB have similar persistence, and opposite from what we observe for consumption, the GLB solution has less overall variance and less variability at all frequencies. Investment has higher variance and persistence in the GLB than in the local solutions, and it has uniformly higher variability at all frequencies.

c) Sudden Stops and Risk Effects

We compare next the results that DynareOBC and GLB yield for the effects of the credit constraint, particularly for financial premia and for sudden stop responses of macro variables when the constraint binds. The financial premia include the shadow interest rate premium (SIP), the equity premium (EP), its components due to unpledgeable capital ($(1 - \kappa)SIP$) and risk premium (RP), and the Sharpe ratio (S). As we explain below, these premia depend on both the tightness of the constraint and its risk effects. For macro aggregates, we compare deviations from long-run averages in consumption, the net exports-GDP ratio, investment, GDP, labor and imported inputs.

SIP is the amount by which the subjective rate at which agents are willing to trade current for future consumption (i.e. the intertemporal marginal rate of substitution $u'(t)/[\beta E_t(u'(t+1))]$) exceeds R_t . The bonds' Euler equation implies that $SIP_t = \mu_t(1 + \tau)/[\beta E_t(u'(t+1))]$. Using the same Euler equation, the denominator can be rewritten without the conditional expectation:

$$SIP_t = \frac{R_t \mu_t (1 + \tau)}{u'(t) - \mu_t (1 + \tau)}. \quad (20)$$

Since $SIP_t = 0$ if $\mu_t = 0$, this shadow interest premium is relevant only when the constraint binds. We can also infer that SIP_t rises as the constraint becomes more binding, because μ_t rises and $E_t(u'(t+1))$ falls, since the constraint forces agents to defer consumption.

The equity premium is defined as $EP_t \equiv E_t[R_{t+1}^q] - R_t$, where $R_{t+1}^q \equiv (d_{t+1} + q_{t+1})/q_t$ is the return on equity and d_{t+1} is the dividend payment, where $d_t \equiv \exp(\epsilon_t^A) F_k(t) - \delta + \frac{a}{2} \frac{(k_{t+1} - k_t)^2}{k_t^2}$. Using the Euler equations for bonds and capital it follows that:

$$EP_t = (1 - \kappa)SIP_t + RP_t, \quad RP_t \equiv -\frac{COV_t[u'(t+1), R_{t+1}^q]}{E_t[u'(t+1)]}. \quad (21)$$

EP_t has two components: RP_t is the standard risk premium driven by the conditional covariance

of marginal utility and equity returns, and the term $(1 - \kappa)SIP_t$ is the fraction of SIP_t pertaining to the share of new capital that cannot be pledged as collateral $1 - \kappa$. EP_t rises when $\mu_t > 0$ for two reasons: First, SIP_t rises, as explained above. Second, RP_t rises, because $COV_t[u'(t+1), R_{t+1}^q]$ becomes more negative as consumption is harder to smooth and $E_t[u'(t+1)]$ falls as the credit constraint forces consumption into the future. Thus, EP responds both to the tightness of the constraint itself via SIP_t and to the larger *risk* premium that the constraint induces. The Sharpe ratio measures the compensation for risk-taking, defined as the excess returns obtained per unit of variability in returns ($S_t = E[EP]/\sigma(R^q)$). We compute the Sharpe ratio following the standard practice of using moments that are not conditional on date t .

For the GLB solution, the financial premia are computed for each triple (b, k, ε) in the state space by applying the above formulae using the Markov process of ε and the recursive optimal decision rules (see Appendix Section B.5 for details). Conditional and unconditional averages are then computed using conditional and unconditional distributions of (b, k, ε) of the GLB solution.

For the DynareOBC solution, SIP_t is computed using the time-series simulation produced by Holden's algorithm. The equity premium is then generated as $EP_t = (1 - \kappa)SIP_t$ because $RP_t = 0$ by construction, since each date- t solution is determined by a perfect-foresight path along which agents base decisions on expectations such that $COV_t[u'(t+1), R_{t+1}^q] = 0$. Computing the unconditional covariance for the simulated dataset also produces very small values for $COV[u'(\cdot), R^q]$.

Table 10 reports quintile distributions of μ conditional on $\mu > 0$, the associated within-quintile averages of the financial and macro variables, their overall means and medians, and the Sharpe ratios.³⁷ The Table shows key differences across GLB and local solutions. The magnitudes of μ are very small in all five quintiles of both solutions, but this is because μ is in units of marginal utility with *CRRA* preferences and $\sigma = 2$. For instance, evaluated at the unconditional means of c and L , marginal utility is about 2.05E-05 (-4.688 in log base 10). Hence, small μ values do not imply that the constraint is negligibly binding or that model solutions are similar, as we document below.

Table 10 shows that the local solutions yield significantly smaller results than the GLB solution for the credit constraint multipliers, the financial premia and the macro responses when the constraint binds. The differences grow larger for higher μ (i.e. in the fourth and fifth quintiles), and they are larger compared with the local solution that has the unconstrained deterministic steady state (DynareOBC-DEIR) v. the one with the constrained steady state (DynareOBC- $\beta R < 1$).

³⁷Variables are assigned into quintiles according to the quintile distribution of μ : if a given μ_i belongs to a particular quintile of μ , then the corresponding values of the financial and macro variables are assigned to that same quintile.

For financial premia, GLB yields overall means of 2.6, 2.2, 2.1 and 0.1 percent for SIP , EP , $(1 - \kappa)SIP$ and RP , respectively, while DynareOBC- $\beta R < 1$ (DynareOBC-DEIR) yields 0.8, 0.6, 0.6 and 0 (0.13, 0.10, 0.10 and 0) percent, respectively. In the GLB solution, RP is about 0.1 percent on average in each of the five quintiles of μ , but EP still increases sharply with μ because $(1 - \kappa)SIP$ rises sharply. In the fifth quintile, GLB yields averages of 6.6, 5.4, and 5.3 percent for SIP , EP , and $(1 - \kappa)SIP$, respectively, while DynareOBC- $\beta R < 1$ (DynareOBC-DEIR) yields 3.3, 2.7 and 2.7 (0.64, 0.51 and 0.51) percent, respectively. The local solutions underestimate SIP , and hence EP , significantly. They also miss the risk premium, but this accounts for a small fraction of the differences in EP . GLB yields a Sharpe ratio of 1.16, nearly 5 and 30 times the DynareOBC- $\beta R < 1$ and DynareOBC-DEIR results, respectively. Since RP is small in the GLB solution and zero in the local solutions, these differences in S are mostly explained by the large gap in SIP .

Large differences in SIP and EP are important because the model predicts that they cause very different responses of macro variables when the constraint binds. To explain how financial premia affect macro responses, we follow Mendoza and Smith (2005) in using the definition of expected returns $E[R_{t+1}^q] \equiv E_t[(d_{t+1} + q_{t+1})/q_t]$ to express the forward solution for the price of capital as:

$$q_t = E_t \left(\sum_{i=1}^{\infty} \left[\prod_{j=0}^i \frac{1}{E_t(R_{t+1+j}^q)} \right] d_{t+1+i} \right) \quad (22)$$

Since eq. (21) implies that $E_t[R_{t+1}^q] = (1 - \kappa)SIP_t + RP_t + R_t$, lower SIP_t in the DynareOBC solutions implies higher equity prices when $\mu_t > 0$, which in turn imply weaker Fisherian deflation effects of the binding credit constraint. Moreover, since equity prices and investment are monotonic functions of each other due to the Tobin Q nature of the investment setup, k_{t+1} should be higher, and so should be borrowing capacity ($\kappa q_t k_{t+1}$), which is critical for determining allocations when $\mu_t > 0$. This also affects the future stream of dividends, which causes further feedback effects into equity prices and borrowing capacity.

The differences in macro responses when $\mu > 0$ across the GLB and local solutions reported in Table 10 reflect the above arguments. In the GLB solution, the responses are in line with the standard features of Sudden Stops (i.e. large recessions and sharp reversals in the external accounts). The mean percent declines (relative to long-run averages) are -3.6 in c , -4.1 in i , -1.0 in GDP , -0.7 in L , and -1.8 in v while NX/GDP rises 2.6 percentage points above its long-run average. The responses are generally larger when the constraint binds more, reaching means of -4.9 for c and -13.5 for i with a trade balance reversal of 5.1 percentage points for the top quintile of μ . The

DynareOBC- $\beta R < 1$ solution underestimates the mean responses of consumption and net exports (-1.9 v. -3.6 for c and 1.2 v. 2.6 for NX/GDP) and overestimates those for L , v and GDP . It also fails to match the property that the responses should be larger when the constraint binds more, as it yields the largest responses in the third quintile of μ . The DynareOBC-DEIR performs even more poorly, producing positive mean responses for c and i and a negative mean response for NX/Y (i.e. on average, when the constraint binds consumption and investment rise above their long-run means and the trade balance falls below its long-run mean). Moreover, these counterfactual responses grow larger when the constraint binds more, in the fourth and fifth quintiles of μ .

Dynare- $\beta R < 1$ and DynareOBC-DEIR yield very different results because in the former (latter) the constraint binds (does not bind) at the deterministic steady state, and in DynareOBC-DEIR the steady state was set at the mean of NFA of the GLB solution. The constraint binds very often in DynareOBC-DEIR (71 percent frequency) but in all but the first quintile of μ it yields below average net exports and above average consumption and investment. This failure to produce Sudden Stops when the constraint binds is an important shortcoming of this solution vis-a-vis both the GLB and Dynare- $\beta R < 1$ solutions.

In summary, the results documented in this Section are consistent with those of the previous Sections indicating that local solutions fail to capture important features of the global solution in terms of long-run moments, impulse-response functions and spectral density functions. In addition, the results show that local solutions fail to match important properties of the GLB solution when the constraint binds. This is due to the inability of the first-order DynareOBC algorithm to capture risk effects and to approximate accurately the frequency, magnitude and effects of the occasionally binding credit constraint. DynareOBC yields large differences in the amount of precautionary savings caused by the constraint, the probability of hitting the constraint, the tightness of the constraint when it binds and the associated financial premia and macro responses.

5 Conclusions

In this paper, we compared global and local solutions of open-economy models with incomplete markets in the time and frequency domains and found major differences. This was the case for endowment, RBC and Sudden Stops (i.e. credit constrained) models of a small open economy. The local solutions were produced using 1OA, 2OA, RSS and DynareOBC methods while the global solutions were generated using an Euler-equation fixed-point iteration method (*FiFit*). Most local

solution methods need a stationarity-inducing transformation, for which we chose the widely-used DEIR function that makes the world real interest rate a decreasing function of the NFA position. We considered the standard “inessential” approach to set the DEIR elasticity such that the interest rate remains close to the world interest rate as well as a variation in which the elasticity is calibrated to match the autocorrelation of NFA in the global solution.

The main limitation of the local methods is their inability to approximate accurately the effects of precautionary savings on NFA, net exports and consumption obtained with the global solution, even using higher-order methods such as 2OA and RSS. On the other hand, *FiPIt* is a fast global method but it is slower than standard local methods and is hampered by the curse of dimensionality. For the Sudden Stops model with the occasionally binding credit constraint, however, *FiPIt* is significantly more accurate than DynareOBC and of comparable speed, although the curse of dimensionality remains an issue. First-order DynareOBC has two key disadvantages: it underestimates the tightness of the credit constraint and its effects on financial premia and macro responses to a binding constraint, and it does not capture *risk* effects of the credit constraint and their implications for precautionary savings and the determination of forward-looking variables like asset prices.

The results show that NFA is a near-unit-root process in all three model applications. Hence, small errors in calculating its “true” autocorrelation coefficient lead the local methods to produce non-trivial errors in other key items, including the long-run averages of NFA, consumption and the balance of trade, as well as various features of business cycle moments, impulse responses and spectral densities. Local solutions with the DEIR elasticity targeted to match the NFA autocorrelation of the global solution perform better, but imply elasticities akin to imposing large costs in moving NFA from its steady state and require knowing the global solution. Interestingly, 1OA, 2OA, and RSS local methods produce very similar second- and higher-order moments, impulse responses and periodograms. This is because they yield decision rules that differ mainly in their intercepts, but otherwise have similar first-order terms and negligible higher-order terms. Hence, if one is restricted to choose from local methods only, and as long as first moments are not central to the question under study, the 1OA method is the best alternative.

Overall, these findings argue in favor of using global methods except when the curse of dimensionality makes them unfeasible. The results are robust to several parameter modifications. Including setting the DEIR elasticity to an inessential low value v. targeting the global solution, replacing the DEIR function with the assumption that the rate of interest is lower than the rate of time preference, introducing different shocks, and changing the variability and persistence of shocks.

References

- [1] Aiyagari, S. Rao, 1994. "Uninsured idiosyncratic risk and aggregate saving," *Quarterly Journal of Economics* 109 (No. 3), 17-31.
- [2] Aguiar, Mark and Gita Gopinath, 2007. "Emerging Market Business Cycles: The Cycle is the Trend," *Journal of Political Economy*, Vol. 115(1), pp. 69-102.
- [3] Andreasen, Martin M., Jesús Fernández-Villaverde and Juan F. Rubio-Ramírez, (2018). "The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications", *Review of Economic Studies* Volume 85(1), pp. 1-49.
- [4] Hamilton, James D., 1994. *Time Series Analysis*, Princeton: Princeton University Press.
- [5] Bewley, T., 1977. "The permanent income hypothesis: A theoretical formulation," *Journal of Economic Theory*, 16(2), 252-92.
- [6] Bianchi, J., Boz, E., Mendoza, E., 2012. "Macroprudential Policy in a Fisherian Model of Financial Innovation," *IMF Econ Rev*, 60: 223.
- [7] Coeurdacier, N., H. Rey, and P. Winant, 2011, "The Risky Steady State," *American Economic Review* 101(3), 398-401.
- [8] de Groot, O., 2013. "Computing the risky steady state of DSGE models," *Economics Letters* 120, 566-569.
- [9] de Groot, O., 2014. "The risky steady state and multiple (spurious) equilibria: The case of the small open economy model," Unpublished manuscript.
- [10] Devereux, M. and A. Sutherland, 2010. "Country Portfolio Dynamics," *Journal of Economic Dynamics and Control*, vol 34(7), pp. 1325-1342, July.
- [11] Durdu, C. Bora, Enrique G. Mendoza and Marco E. Terrones, 2009. "Precautionary Demand for Foreign Assets in Sudden Stop Economies: An Assessment of the New Mercantilism," *Journal of Development Economics*, Vol. 89 pp. 194-209, July.
- [12] Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramírez, M. Uribe, 2011. "Risk Matters: The Real Effects of Volatility Shocks," *American Ec. Review*, 101(6), pp. 2530-61.
- [13] Garcia-Cicco, Javier, Roberto Pancrazzi and Martin Uribe, 2010, "Real Business Cycles in Emerging Countries?," *American Economic Review*, December, vol. 100(5), pp. 2510-31.
- [14] Guerrieri, Luca & Matteo Iacoviello, 2015. "OccBin: A Toolkit to Solve Models with Occasionally Binding Constraints Easily," *Journal of Monetary Economics*, March, vol. 70, pp. 22-38.
- [15] Holden, Thomas, 2016a. "Existence and uniqueness of solutions to dynamic models with occasionally binding constraints," Unpublished manuscript
- [16] Holden, Thomas, 2016b. "Computation of solutions to dynamic models with occasionally binding constraints," Unpublished manuscript
- [17] Hugget, Mark, 1993, "The risk free rate in heterogenous-agent, incomplete-insurance economies," *Journal of Economic Dynamics and Control* 17 (No. 5-6), 953-969.

- [18] Kliem, M. and H. Uhlig, 2016, “Bayesian estimation of a dynamic stochastic general equilibrium model with asset prices,” *Quantitative Economics*, 7: 257-287.
- [19] Ljungqvist, L. & T. J. Sargent. 2012. Recursive Macroeconomics Theory. MIT Press.
- [20] Epstein, L. G. (1983). “Stationary Cardinal Utility and Optimal Growth under Uncertainty.” *Journal of Economic Theory*, v. 31, pp. 133-152.
- [21] Mendoza, Enrique G., 1991. “Real Business Cycles in a Small Open Economy,” *American Economic Review*, Vol. 81, pp. 797-818, September.
- [22] Mendoza, Enrique G., 1995. “The Terms of Trade, the Real Exchange Rate, and Economic Fluctuations” *International Economic Review*, Vol. 36, pp. 101-137, September.
- [23] Mendoza, Enrique, 2010. “Sudden Stops, Financial Crises and Leverage,” *American Economic Review*, Vol. 100(5), pp. 1941-1966.
- [24] Mendoza, Enrique, Vincenzo Quadrini & Jose Victor Rios-Rull, 2009. “Financial Integration, Financial Development & Global Imbalances,” *Journal of Pol. Economy*, 117(3), pp. 371-416.
- [25] Mendoza, Enrique and Sergio Villalvazo, 2019, “*FiPit*: A Simple, Fast Global Method for Solving Models with Two Endogenous States & Occasionally Binding Constraints,” mimeo, Department of Economics, University of Pennsylvania.
- [26] Neumeyer, Pablo A. and Fabrizio Perri, 2005, “Business Cycles in Emerging Economies: the Role of Interest Rates,” *Journal of Monetary Economics*, Vol. 52, pp. 345-380.
- [27] Rabitsch, Katrin, Serhiy Stepanchuk & Viktor Tsyrennikov, 2015, “International Portfolios: A Comparison of Solution Methods,” *Journal of International Economics*, Vol. 97(2), pp. 404-422.
- [28] Schmitt-Grohé, Stephanie and Martin Uribe, 2003, “Closing Small Open Economy Models,” *Journal of International Economics*, Vol. 61, pp. 163-185.
- [29] Schmitt-Grohé, Stephanie & Martin Uribe, 2004, “Solving Dynamic General Equilibrium Models using a Second-Order Approximation to the policy function,” *Journal of Econ. Dynamics & Control*, Vol. 28, pp. 755-775.
- [30] Taylor, John B. and Harald Uhlig, 1990, “Solving Nonlinear Stochastic Growth Models: A Comparison of Alternative Solution Methods,” *Journal of Business and Economic Statistics*, Vol. 8, No. 1, pp. 1-17.
- [31] Uribe, Martin and Vivian Z. Yue, 2006. “Country Spreads and Emerging Countries: Who Drives Whom?,” *Journal of International Economics*, Vol. 69, pp. 6-36.
- [32] Uzawa, H., 1968. “Time Preference, the consumption function and optimum asset holdings,” J.N. Wolfe (Ed.), *Value, Capital & Growth*, Univ. of Edinburgh Press, Edinburgh, pp. 485-504.
- [33] Wei Dou, Winston, Xiang Fang, Andrew W. Lo, and Harald Uhlig, 2019, “Macro-Finance Models with Nonlinear Dynamics,” Univ. of Chicago, Becker Friedman Institute Working Paper, forthcoming.

Table 1: Methods Used to Solve Open-Economy Incomplete Markets Models

Authors	Year	Publication	Type of model	Solution method	Stationarity assumption	ψ
Adolfson et al.	2007	JIE	SOE	1OA	DEIR	.145(e)
Aguiar and Gopinath	2007	JPE	SOE	1OA	DEIR	.001(s)
Angeloni and Ehrmann	2007	BEJ Macro	$N = 12$	1OA	DEIR	.1(c)
Arellano	2008	AER	SOE	GLB		
Arellano and Mendoza	2002	NBER	SOE	GLB		
Baxter and Crucini	1995	IER	$N = 2$	1OA	AHC	
Bengui et al.	2012	JME	$N = 2$	GLB		
Benigno and Thoenissen	2008	JIMF	$N = 2$	1OA	AHC	
Benigno et al.	2016	JME	SOE	GLB		
Bergin	2006	JIMF	$N = 2$	1OA	DEIR	.00384(e)
Bianchi	2011	AER	SOE	GLB		
Bianchi and Mendoza	2018	JPE	SOE	GLB		
Bianchi et al.	2012	IMFER	SOE	GLB		
Bianchi et al.	2016	JIE	SOE	GLB		
Bodenstein	2011	JIE	$N = 2$	1OA	ED	
Bodenstein et al.	2011	JIE	$N = 2$	1OA	DEIR	0.0001(s)
Boz et al.	2011	JME	SOE	1OA	DEIR	.001(s)
Boz and Mendoza	2014	JME	SOE	GLB		
Buch et al.	2005	JIMF	$N = 2$	1OA	AHC	
Cavallo and Ghironi	2002	JME	$N = 2$	1OA	OLG	
Coeurdacier et al.	2011	AER P&P	SOE	RSS		
Correia et al.	1995	EER	SOE	1OA	AHC	
Corsetti et al.	2008	RESTUD	$N = 2$	1OA	ED	
Cuadra and Sapriza	2008	JIE	SOE	GLB		
Devereux et al.	2006	EJ	SOE	1OA	AHC	
Devereux and Sutherland	2010	JMCB	$N = 2$	1OA	ED	
Devereux and Sutherland	2011	JEEA	$N = 2$	1OA	ED	
Durdu et al.	2009	JDE	SOE	GLB		
Durdu and Mendoza	2006	JIE	SOE	GLB		
Enders et al.	2011	JIE	$N = 2$	1OA	ED	
Engel and Wang	2011	JIE	$N = 2$	1OA	AHC	

Continued on next page

Table 1 – continued from previous page

Authors	Year	Journal	Type of model	Solution method	Stationarity assumption	ψ
Fernandez and Chang	2013	IER	SOE	1OA	DEIR	.001 (s)
Fernandez-Villaverde et al.	2011	AER	SOE	3OA	AHC	
Fogli and Perri	2006	NBER	$N = 2$	GLB		
Garcia-Cicco et al.	2010	AER	SOE	1OA	DEIR	.001 (s), 2.8(e)
Gertler et al.	2007	JMCB	SOE	1OA	DEIR	0.0001(s)
Ghironi	2006	JIE	$N = 2$	1OA	OLG	
Ghironi and Melitz	2005	QJE	$N = 2$	1OA	AHC	
Hatchondo and Martinez	2009	JIE	SOE	GLB		
Heathcote and Perri	2002	JME	$N = 2$	1OA	AHC	
Heathcote and Perri	2013	JPE	$N = 2$	2OA, 3OA		
Jaimovich and Rebelo	2008	JMCB	SOE	1OA	DEIR	.00001(s)
Justiniano and Preston	2010	JIE	SOE	1OA	DEIR	.01(c)
Lubik and Schorfheide	2005	NBER Macro	$N = 2$	1OA	CM	
Mendoza	1991	AER	SOE	GLB		
Mendoza	1992	IMFSP	SOE	GLB		
Mendoza	1995	IER	SOE	GLB		
Mendoza	2010	AER	SOE	GLB		
Mendoza and Smith	2006	JIE	SOE	GLB		
Mendoza et al.	2009	JPE	$N = 2, 3$	GLB		
Mendoza and Yue	2012	QJE	SOE	GLB		
Monacelli	2005	JMCB	SOE	1OA	CM	
Nason and Rogers	2006	JIE	SOE	1OA	DEIR	.00014, .007(e)
Neumeyer and Perri	2005	JME	SOE	1OA	AHC	
Rabanal and Tuesta	2010	JEDC	$N = 2$	1OA	AHC	
Raffo	2008	JIE	$N = 2$	1OA	AHC	
Rebelo and Vegh	1995	RESTUD	SOE	1OA	AHC	
Smets and Wouters	2002	JME	SOE	1OA	OLG	
Uribe and Yue	2006	JIE	SOE	1OA	AHC	

Note: SOE denotes a small open economy model. $N =$ denotes a multicountry model with N countries. 1OA, 2OA and 3OA are the first-, second- and third-order approximation methods respectively, RSS is the risky steady state method, and GLB indicates models solved with global methods (including models with standard preferences and $\beta R < 1$, endogenous discounting, or overlapping generations). The approaches used to induce stationarity when using local methods are the debt-elastic interest rate (DEIR), asset holding costs (AHC), endogenous discounting (ED), overlapping generations (OLG) and complete markets (CM). For cases using DEIR, (s), (c) and (e) denote whether the debt-elasticity parameter ψ was chosen to be small, estimated, or calibrated respectively.

Table 2: Solution Methods Used in Policy Models

Institution	Model name	Type of model	Solution method	Stationarity assumption	ψ
Bank of Canada	GEM	$N = 5$	1OA	DEIR	n.d.
Bank of England	COMPASS	SOE	1OA	ED	
ECB	NAWM	$N = 2$ and SOE	1OA	DEIR	.01(s)
European Commission	QUEST	SOE	1OA	DEIR	.02(e)
Federal Reserve Board	SIGMA	$N = 2$	1OA	PAC	
IMF	GIMF	$N \geq 2$	1OA	OLG	
Norges Bank	NEMO	SOE	1OA	DEIR	n.d.
Riksbank	RAMSES	SOE	1OA	DEIR	.01(c)

Note: See note to Table 1 for details on abbreviations. n.d. denotes that there is no public document disclosing what value was used.

Table 3: Calibration of the Endowment Economy Model

Notation	Parameter/Variable	Value
1. Common parameters		
σ	Coefficient of relative risk aversion	2.0
y	Mean endowment income	1.00
A	Absorption constant	0.28
R	Gross world interest rate	1.059
σ_z	Standard deviation of income (percent)	3.27
ρ_z	Autocorrelation of income	0.597
2. Global solution parameters		
β	Discount factor	0.940
φ	Ad-hoc debt limit	-0.51
3. Local solution parameters		
<i>Common parameters</i>		
β	Discount factor	0.944
\bar{b}	Deterministic steady state value of NFA	-0.51
<i>Baseline calibration</i>		
ψ	Inessential DEIR coefficient	0.001
<i>Targeted calibration</i>		
ψ	DEIR coefficient for 2OA	0.0469
ψ	DEIR coefficient for RSS	0.0469

Note: 2OA and RSS denote the second-order and risky-steady state solutions, respectively.

Table 4: Autocorrelations of Net Exports, NFA, and Income: Endowment Economy Model

	ρ^ϵ									
	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	
<i>i) GLB</i>										
ρ_{NFA}	0.827	0.866	0.899	0.926	0.947	0.964	0.977	0.987	0.993	
ρ_{NX}	-0.088	0.010	0.11	0.213	0.321	0.432	0.547	0.661	0.768	
<i>ii) Baseline (inessential ψ)</i>										
2OA										
ρ_{NFA}	0.995	0.996	0.996	0.997	0.998	0.998	0.999	0.999	0.999	
ρ_{NX}	0.238	0.348	0.455	0.556	0.652	0.741	0.821	0.89	0.946	
RSS										
ρ_{NFA}	0.995	0.996	0.997	0.997	0.998	0.998	0.999	0.999	1.000	
ρ_{NX}	0.239	0.35	0.457	0.559	0.655	0.745	0.826	0.896	0.952	
<i>iii) Targeted (calibrated ψ to match baseline GLB ρ^b)</i>										
2OA										
ρ_{NFA}	0.912	0.928	0.941	0.952	0.961	0.97	0.977	0.984	0.990	
ρ_{NX}	-0.01	0.089	0.188	0.287	0.387	0.486	0.586	0.685	0.786	
RSS										
ρ_{NFA}	0.912	0.928	0.941	0.952	0.961	0.97	0.977	0.984	0.99	
ρ_{NX}	-0.01	0.089	0.188	0.287	0.386	0.485	0.585	0.684	0.784	
<i>iv) Targeted for all ρ^ϵ (calibrated ψ to match each GLB ρ^b)</i>										
2OA										
ψ	0.185	0.158	0.13	0.106	0.083	0.064	0.046	0.034	0.027	
ρ_{NFA}	0.827	0.866	0.899	0.926	0.947	0.964	0.977	0.987	0.993	
ρ_{NX}	-0.029	0.068	0.166	0.267	0.37	0.476	0.586	0.698	0.807	
RSS										
ψ	0.185	0.158	0.13	0.106	0.083	0.064	0.046	0.034	0.027	
ρ_{NFA}	0.827	0.866	0.899	0.926	0.947	0.964	0.977	0.987	0.993	
ρ_{NX}	-0.030	0.067	0.166	0.266	0.369	0.475	0.585	0.696	0.804	

Note: GLB, 2OA and RSS denote the global, second-order and risky-steady state solutions, respectively.

Table 5: Long-run Moments: Endowment Economy Model

	Baseline Calibration			Targeted Calibration		
	GLB	2OA	RSS		2OA	RSS
		DEIR	$\beta R < 1$	DEIR	DEIR	DEIR
$\psi =$	<i>na</i>	0.001	<i>na</i>	0.001	0.0469	0.0469
<i>Averages</i>						
$E(c)$	0.694	0.701	0.093	0.692	0.689	0.689
$E(nx/y)$	0.022	0.015	0.625	0.025	0.028	0.028
$E(b/y)$	-0.413	-0.282	-11.210	-0.451	-0.502	-0.506
<i>Variability relative to variability of income</i>						
$\sigma(c)/\sigma(y)$	0.992	1.594	1.161	1.617	1.001	0.997
$\sigma(nx)/\sigma(y)$	0.660	1.327	1.202	1.346	0.730	0.730
$\sigma(nx/y)/\sigma(y)$	0.643	1.311	1.161	1.331	0.709	0.709
$\sigma(b)/\sigma(y)$	7.461	62.327	1.706	40.078	6.647	6.576
$\sigma(b/y)/\sigma(y)$	7.735	61.989	1.892	40.213	7.174	7.118
<i>Income correlations</i>						
$\rho(y, c)$	0.755	0.202	0.188	0.197	0.684	0.684
$\rho(y, nx)$	0.729	0.572	0.312	0.567	0.705	0.708
$\rho(y, nx/y)$	0.704	0.572	0.006	0.567	0.705	0.708
$\rho(y, b)$	0.449	0.128	0.070	0.124	0.489	0.488
$\rho(y, b/y)$	0.549	0.156	0.445	0.149	5.593	0.592
<i>First-order autocorrelations</i>						
$\rho(c)$	0.840	0.995	0.996	0.995	0.929	0.929
$\rho(nx)$	0.543	0.819	0.934	0.823	0.583	0.582
$\rho(nx/y)$	0.551	0.826	0.995	0.830	0.591	0.590
$\rho(b)$	0.977	0.999	0.999	0.999	0.977	0.977
$\rho(b/y)$	0.961	0.998	0.953	0.998	0.958	0.959
<i>Performance metrics</i>						
Execution time (secs.)	35.9	8.5	n.a.	9.9	8.5	9.8
ratio rel. to GLB	1.0	0.237	n.a.	0.276	0.237	0.273
Max. Abs. Euler eq. errors	9.60E-05	1.10E-03	4.45E-03	1.00E-03	2.60E-03	2.50E-03
Decision rule diff b		0.120 (0.248)	n.a.	0.099 (0.378)	0.086 (0.127)	0.075 (0.123)
Decision rule diff c		0.025 (0.049)	n.a.	0.028 (0.055)	0.019 (0.037)	0.020 (0.039)

Note: GLB, 2OA and RSS refer to the global, second-order and risky-steady state solutions, respectively. $\sigma(\cdot)$ denotes the coefficient of variation for variables in levels and the standard deviation for variables in ratios (nx/y , b/y and the leverage ratio $lev/rat.$). The results were obtained using Matlab 2016a in a Linux cluster with 128gb of RAM, two 10-core Intel(R) Xeon(R) CPU E5-2690 v2 @ 3.00GHz processors, and a Samsung SSD 840 512GB hard drive. We used the same software and hardware for all the computations reported in the paper. Execution times are not reported for the RSS solution with $\beta R < 1$ because we used analytic solutions. The Euler equation errors and decision rule differences are computed for all (b, z) pairs in the state space of the GLB solution. The decision rule differences for the local solutions in the last two rows are differences between the local and GLB solutions in percent of the latter. We report the mean and maximum (the maximum in brackets) differences conditional on bond values that have positive probability in the ergodic distribution of the GLB solution.

Table 6: Endowment Economy Model with Income and Interest-Rate Shocks

	Std Dev of Int Rate (percent)					
	0.0	0.5	1.0	1.5	2.0	2.5
<i>Global calibrated</i>						
$E(b/y)$	-0.413	-0.411	-0.408	-0.403	-0.396	-0.384
$\sigma(c)/\sigma(y)$	0.992	0.977	1.009	1.058	1.126	1.214
$\sigma(b)/\sigma(y)$	7.461	7.169	7.465	8.009	8.874	10.311
$\rho(y, nx)$	0.729	0.681	0.617	0.527	0.415	0.298
$\rho(nx)$	0.543	0.540	0.542	0.546	0.551	0.559
$\rho(b)$	0.977	0.973	0.975	0.976	0.978	0.981
<i>Global with NDL</i>						
$E(b/y)$	-10.778	-9.249	-7.445	-5.991	-4.875	-3.956
$\sigma(c)/\sigma(y)$	9.747	7.375	6.962	6.189	5.563	4.906
$\sigma(b)/\sigma(y)$	1.682	2.418	4.232	5.771	7.194	8.374
$\rho(y, nx)$	0.684	0.457	0.343	0.308	0.297	0.301
$\rho(nx)$	0.858	0.880	0.924	0.931	0.927	0.914
$\rho(b)$	0.999	0.998	0.998	0.998	0.998	0.997
<i>Full RSS w. $\beta\bar{R} < 1$</i>						
$E(b/y)$	-11.21	-9.098	-7.182	-5.577	-4.226	-3.075
$\sigma(c)/\sigma(y)$	12.484	11.171	9.672	8.209	6.745	5.322
$\sigma(b)/\sigma(y)$	19.067	38.394	49.967	53.952	52.038	45.600
$\rho(y, nx)$	0.315	0.077	0.011	-0.021	-0.044	-0.066
$\rho(nx)$	0.933	0.987	0.993	0.994	0.992	0.986
$\rho(b)$	0.999	0.999	0.999	0.999	0.999	0.999
<i>Partial RSS w. baseline ψ</i>						
$E(b/y)$	-0.451	-0.426	-0.279	-0.018	0.381	0.942
$\sigma(c)/\sigma(y)$	1.617	1.645	1.773	2.085	2.894	4.969
$\sigma(b)/\sigma(y)$	40.078	43.072	71.486	1327.807	94.562	71.228
$\rho(y, nx)$	0.567	0.560	0.531	0.469	0.357	0.217
$\rho(nx)$	0.823	0.823	0.830	0.856	0.910	0.965
$\rho(b)$	0.999	0.999	0.999	0.999	0.999	0.999
<i>2OA w. baseline ψ</i>						
$E(b/y)$	-0.285	-0.319	-0.179	0.056	0.384	0.806
$\sigma(c)/\sigma(y)$	1.594	1.612	1.664	1.747	1.857	1.990
$\sigma(b)/\sigma(y)$	62.327	55.583	100.480	313.282	47.421	23.101
$\rho(y, nx)$	0.572	0.568	0.555	0.536	0.512	0.485
$\rho(nx)$	0.819	0.816	0.809	0.798	0.785	0.771
$\rho(b)$	0.999	0.999	0.999	0.999	0.999	0.999
<i>Partial RSS w. targeted ψ</i>						
$E(b/y)$	-0.506	-0.507	-0.505	-0.501	-0.495	-0.487
$\sigma(c)/\sigma(y)$	0.997	1.016	1.068	1.150	1.254	1.375
$\sigma(b)/\sigma(y)$	6.576	6.571	6.657	6.805	7.022	7.315
$\rho(y, nx)$	0.708	0.695	0.663	0.619	0.571	0.523
$\rho(nx)$	0.582	0.580	0.576	0.570	0.564	0.559
$\rho(b)$	0.977	0.977	0.977	0.977	0.977	0.977
<i>2OA w. targeted ψ</i>						
$E(b/y)$	-0.502	-0.505	-0.502	-0.498	-0.492	-0.484
$\sigma(c)/\sigma(y)$	1.001	1.020	1.073	1.157	1.264	1.391
$\sigma(b)/\sigma(y)$	6.647	6.612	6.694	6.833	7.030	7.287
$\rho(y, nx)$	0.705	0.693	0.660	0.615	0.564	0.514
$\rho(nx)$	0.583	0.581	0.577	0.572	0.566	0.561
$\rho(b)$	0.977	0.977	0.977	0.977	0.977	0.977

Note: The variability and persistence of endowment shocks are kept as in Table 3. The correlation between endowment and interest-rate shocks is set to -0.669 , for all columns with the exception of the first column for which the correlation is set to 0. GLB, 2OA and RSS refer to the global, second-order and risky-steady state solutions, respectively.

Table 7: Calibration of the RBC & Sudden Stops Models

Notation	Parameter/Variable	Value
1. Common parameters		
σ	Coefficient of relative risk aversion	2.0
R	Gross world interest rate	1.0857
α	Labor share in gross output	0.592
γ	Capital share in gross output	0.306
η	Imported inputs share in gross output	0.102
δ	Depreciation rate of capital	0.088
ω	Labor exponent in the utility function	1.846
ϕ	Working capital constraint coefficient	0.2579
a	Investment adjustment cost parameter	2.75
τ	Consumption tax	0.168
κ	Collateral constraint coefficient	0.20
y^{dss}	GDP at the deterministic steady state	396
2. RBC global solution parameters		
β	Discount factor	0.920
φ	Ad-hoc debt limit as a share of y^{dss}	-0.758
3. RBC local solution parameters		
<i>Common Parameters</i>		
β	Discount factor	0.9211
b^{dss}/y^{dss}	NFA/GDP at the deterministic steady state	-0.758
<i>Baseline Calibration</i>		
ψ	Inessential DEIR coefficient	0.001
<i>Targeted Calibration</i>		
ψ	DEIR coefficient for 2OA	0.0109
ψ	DEIR coefficient for RSS	0.008
4. Sudden Stops global solution parameters		
β	Discount factor	0.920
φ	Ad-hoc debt limit as a share of y^{dss}	-0.505
b^{dss}/y^{dss}	NFA/GDP at the deterministic steady state	-0.192
5. Sudden Stops local solution parameters		
<i>DynareOBC with $\beta R < 1$</i>		
β	Discount factor	0.920
b^{dss}/y^{dss}	NFA/GDP at the deterministic steady state	-0.192
<i>DynareOBC with DEIR</i>		
β	Discount factor	0.9211
ψ	Inessential DEIR coefficient	0.001
b^{dss}/y^{dss}	NFA/GDP at the deterministic steady state	0.015

Note: 2OA and RSS denote the second-order and risky-steady state solutions, respectively. For the Sudden Stops model, the GLB solution has two credit constraints, namely φ and the collateral constraint. Credit is constrained at the deterministic steady state, since $\beta R < 1$, but φ is set low enough so that the collateral constraint binds first.

Table 8: Long-run Moments: RBC model

	GLB	Baseline Calibration		Targeted Calibration	
		2OA	RSS	2OA	RSS
$\psi =$	na	0.001	0.001	0.0109	0.008
<i>Averages</i>					
$E(y)$	393.847	396.990	396.190	397.050	397.210
$E(c)$	264.021	294.900	342.850	259.180	265.420
$E(I)$	67.53	68.008	67.747	68.035	68.063
$E(nx/y)$	0.045	-0.042	-0.185	0.065	0.046
$E(b/y)$	-0.372	0.732	2.559	-0.620	-0.397
$E(lev.rat.)$	-0.286	-0.237	-1.100	0.400	0.295
$E(v)$	42.649	42.938	42.852	42.946	42.975
$E(L)$	18.433	18.519	18.499	18.521	18.528
<i>Variability relative to variability of GDP</i>					
$\sigma(y)$	0.040	0.039	0.039	0.041	0.040
$\sigma(c)/\sigma(y)$	1.291	1.910	1.412	1.268	1.212
$\sigma(I)/\sigma(y)$	3.386	3.467	3.493	3.320	3.388
$\sigma(nx/y)/\sigma(y)$	0.885	1.293	1.212	0.712	0.731
$\sigma(b/y)/\sigma(y)$	7.589	13.824	12.909	3.758	4.269
$\sigma(lev.rat.)/\sigma(y)$	3.614	6.549	6.084	1.849	2.053
$\sigma(v)/\sigma(y)$	1.481	1.496	1.504	1.463	1.482
$\sigma(L)/\sigma(y)$	0.596	0.600	0.600	0.596	0.598
<i>Correlations with GDP</i>					
$\rho(y, c)$	0.773	0.631	0.509	0.929	0.904
$\rho(y, I)$	0.640	0.632	0.628	0.661	0.648
$\rho(y, nx/y)$	-0.227	-0.278	0.026	-0.476	-0.381
$\rho(y, b/y)$	0.090	0.200	-0.160	0.511	0.343
$\rho(y, lev.rat.)$	0.112	-0.206	0.150	-0.532	-0.366
$\rho(y, v)$	0.834	0.831	0.830	0.839	0.835
$\rho(y, L)$	0.995	0.995	0.995	0.995	0.995
<i>First-order autocorrelations</i>					
$\rho(y)$	0.830	0.824	0.820	0.841	0.853
$\rho(b)$	0.996	0.999	0.998	0.996	0.996
$\rho(c)$	0.885	0.940	0.918	0.873	0.862
$\rho(I)$	0.516	0.511	0.509	0.519	0.513
$\rho(nx/y)$	0.711	0.850	0.843	0.555	0.563
$\rho(lev.rat.)$	0.997	0.999	0.998	0.991	0.995
$\rho(v)$	0.780	0.777	0.774	0.788	0.782
$\rho(L)$	0.808	0.808	0.799	0.819	0.810
<i>Performance metrics</i>					
Time in sec.	61.0	37.8	40.6	37.9	39.6
ratio rel. to GLB	1.0	0.620	0.666	0.621	0.649
Max. Abs. b Euler eq. error	1.17E-07	1.33E-07	6.21E-04	1.43E-07	1.13E-03
Max. Abs. k Euler eq. error	3.84E-16	4.52E-07	8.92E-05	3.95E-07	6.59E-05
Decision rule diff b		0.089 (0.546)	0.089 (0.546)	0.081 (0.505)	0.081 (0.505)
Decision rule diff k		0.017 (0.051)	0.017 (0.051)	0.017 (0.049)	0.017 (0.049)
Decision rule diff c		0.019 (0.066)	0.019 (0.066)	0.019 (0.077)	0.019 (0.077)

See Note to Table 5.

Table 9: Long-run Moments: Sudden Stops model

	GLB	DynareOBC- $\beta R < 1$	DynareOBC-DEIR
<i>Averages</i>			
$E(y)$	393.619	391.390	395.230
$E(c)$	274.123	269.610	279.970
$E(I)$	67.481	66.714	67.897
$E(nx/y)$	0.015	0.025	0.000
$E(b/y)$	0.015	-0.100	0.206
$E(lev.rat.)$	-0.102	-0.157	-0.011
$E(v)$	42.617	42.263	42.712
$E(L)$	18.426	18.364	18.469
<i>Variability relative to variability of GDP</i>			
$\sigma(y)$	0.039	0.032	0.032
$\sigma(c)/\sigma(y)$	1.023	0.937	1.207
$\sigma(I)/\sigma(y)$	3.383	3.492	3.777
$\sigma(nx/y)/\sigma(y)$	0.746	0.927	1.262
$\sigma(b/y)/\sigma(y)$	4.980	3.703	9.595
$\sigma(lev.rat.)/\sigma(y)$	2.340	1.705	4.498
$\sigma(v)/\sigma(y)$	1.495	1.632	1.612
$\sigma(L)/\sigma(y)$	0.599	0.571	0.569
<i>Correlations with GDP</i>			
$\rho(y, c)$	0.842	0.823	0.557
$\rho(y, I)$	0.641	0.309	0.224
$\rho(y, nx/y)$	-0.117	0.176	0.223
$\rho(y, b/y)$	-0.120	0.027	-0.054
$\rho(y, lev.rat.)$	-0.111	0.008	-0.056
$\rho(y, v)$	0.832	0.777	0.775
$\rho(y, L)$	0.994	0.987	0.986
<i>First-order autocorrelations</i>			
$\rho(y)$	0.825	0.752	0.754
$\rho(b)$	0.990	0.980	0.995
$\rho(c)$	0.829	0.826	0.910
$\rho(I)$	0.500	0.470	0.502
$\rho(nx/y)$	0.601	0.456	0.651
$\rho(lev.rat.)$	0.992	0.988	0.996
$\rho(v)$	0.777	0.753	0.756
$\rho(L)$	0.801	0.761	0.774
Prob. ($\mu > 0$)	2.58	51.80	71.06
<i>Performance metrics</i>			
Time in sec.	268.0	243.5	187.4
Max. Abs. b Euler eq. error	2.62E-04	na	na
Max. Abs. k Euler eq. error	4.25E-16	na	na


Note: See Note to Table 5.

Table 10: Credit Constraint Multiplier, Macro & Financial Variables Conditional on $\mu > 0$

$\log(\mu)$		Financial Premia			Macro variables						
upper limit	mean	SIP	EP	$(1-\kappa)SIP$	RP	c	NX/GDP	i	L	v	
		means in each quintile of μ			means of deviations from long-run averages in each quintile of μ						
Panel a. GLB											
<i>Quintiles of μ</i>											
First	-6.563	0.32	0.37	0.26	0.10	-2.76	1.98	-1.76	-0.60	-0.26	0.35
Second	-6.320	1.07	0.96	0.85	0.11	-2.17	1.37	2.70	0.08	0.15	-1.25
Third	-6.088	1.82	1.56	1.46	0.10	-3.80	2.30	-3.00	-1.35	-0.82	-1.29
Fourth	-5.843	2.98	2.48	2.38	0.09	-4.72	2.58	-5.46	-2.26	-1.42	-3.35
Fifth	-3.374	6.59	5.38	5.27	0.10	-4.86	5.10	-13.45	-1.21	-1.37	-2.98
<i>Overall mean</i>	-7.627	2.59	2.17	2.07	0.10	-3.64	2.64	-4.05	-1.04	-0.73	-1.78
<i>Overall median</i>	-6.198	1.79	1.52	1.43	0.11	-3.22	1.60	-1.64	-1.02	-0.57	-2.15
<i>Ex-post Sharpe ratio = 1.16</i>											
Panel b. DynareOBC-BetaR < 1											
<i>Quintiles of μ</i>											
First	-9.000	0.00	0.00	0.00	0.00	-0.78	0.73	-3.17	-0.40	-0.30	-0.53
Second	-8.523	0.01	0.01	0.01	0.00	-1.74	1.38	-6.36	-1.10	-0.77	-1.58
Third	-8.155	0.02	0.01	0.01	0.00	-3.23	2.33	-11.28	-2.24	-1.52	-3.13
Fourth	-6.295	0.66	0.53	0.53	0.00	-2.02	1.24	-4.04	-1.13	-0.68	-1.54
Fifth	-5.523	3.32	2.65	2.65	0.00	-1.85	0.38	-1.35	-1.38	-0.81	-1.92
<i>Overall mean</i>	-6.623	0.800	0.64	0.64	0.00	-1.92	1.21	-5.24	-1.25	-0.82	-1.74
<i>Overall median</i>	-8.337	0.015	0.01	0.01	0.00	-1.87	0.93	-4.71	-1.23	-0.81	-1.74
<i>Ex-post Sharpe ratio = 0.25</i>											
Panel c. DynareOBC DEIR											
<i>Quintiles of μ</i>											
First	-15.380	0.00	0.00	0.00	0.00	-1.09	1.05	-0.27	0.12	0.06	0.20
Second	-15.064	0.00	0.00	0.00	0.00	-0.14	-0.13	0.62	-0.13	-0.05	-0.22
Third	-14.856	0.00	0.00	0.00	0.00	1.10	-1.44	1.20	-0.30	-0.13	-0.36
Fourth	-14.659	0.00	0.00	0.00	0.00	2.67	-3.09	1.68	-0.55	-0.25	-0.68
Fifth	-5.603	0.64	0.51	0.51	0.00	2.22	-3.20	1.83	-0.99	-0.47	-1.34
<i>Overall mean</i>	-7.422	0.13	0.10	0.10	0.00	0.95	-1.36	1.01	-0.37	-0.16	-0.48
<i>Overall median</i>	-14.954	0.00	0.00	0.00	0.00	0.89	-1.13	0.96	-0.37	-0.18	-0.48
<i>Ex-post Sharpe ratio = 0.04</i>											

Note: SIP is the Shadow interest premium, EP is the equity premium, and RP is the risk premium component of EP . The quintile distribution of μ is conditional on $\mu > 0$. Means for other variables are computed using the distribution of μ within each quantile and the overall distribution of μ conditional on $\mu > 0$. $\log(\mu)$ is the base-10 logarithm of the multiplier on the credit constraint. The moments for GLB are computed using the recursive equilibrium decision rules and ergodic distribution of the model. For EP , we compute the equity premium conditional on all date-t states (b, k, s) and then average across them using the ergodic distribution. The Sharpe ratio is computed using the conditional ergodic distribution for $\mu > 0$. The moments for DynareOBC are ex-post moments, computed using time-series simulation output of DynareOBC. RP is set to zero because the covariance between future equity returns and marginal utility is zero along the perfect-foreflight paths that determine each date-t solution in DynareOBC.


Figure 1: The first-order coefficient of NFA decision rules as the elasticity of the DEIR function varies




Note: This figure shows how the first-order coefficient of the NFA decision rules, $\rho_b(\psi, b^*)$, varies with ψ for three values of b^* : -0.51 (deterministic steady state), -0.41 (risky steady state) and zero.

Figure 2: Average NFA in the endowment economy as the variability and persistence of output rise

Baseline Calibration ($\psi = 0.001$)




Targeted Calibration ($\psi = 0.0469$)




Note: GLB refers to global solution, 2OA refers to second-order solution, RSS refers to risky-steady state solution.

Figure 3: Endowment Model Impulse Response Functions to a Negative Income Shock




Note: GLB, 1OA, 2OA and RSS denote global, first-order, second-order and risky-steady state solutions, respectively. GLB impulse responses are forecast functions of the equilibrium Markov processes of the endogenous variables with initial conditions set to $E[b]$ and a value of z equal to a one-standard-deviation shock.

Figure 4: Spectral Density Functions in the Endowment Economy Model




Note: These graphs show parametric estimates of spectral density functions. GLB, 2OA and RSS denote the global, second-order and risky-steady state solution, respectively.

Figure 5: RBC Impulse Response Functions to a Negative TFP shock: Baseline Calibration




Note: GLB, 1OA, 2OA, RSS refer to the global, first-order, second-order and risky-steady state solution, respectively. GLB impulse responses are forecast functions of the equilibrium Markov processes of the endogenous variables with initial conditions set to $E[b]$, $E[k]$ and a value of TFP equal to a one-standard-deviation shock. Variables are plotted as percent deviations from long-run means, with the exception of NFA and net exports, which are plotted as differences relative to their long-run means (since these variables are measured as GDP ratios, and hence are already in percent).

Figure 6: RBC Impulse Response Functions to a Negative TFP Shock: Targeted Calibration




Note: GLB, 1OA, 2OA, RSS refer to the global, first-order, second-order and risky-steady state solution, respectively.

Figure 7: Spectral Density Functions for the RBC model: Baseline Calibration



Note: These graphs show parametric estimates of spectral density functions. GLB, 2OA, and RSS refer to the global, second-order and risky-steady state solution, respectively.

Figure 8: Spectral Density Functions for the RBC model: Targeted Calibrations



Note: These graphs show parametric estimates of spectral density functions. GLB, 2OA, and RSS refer to the global, second-order and risky-steady state solution, respectively.

Figure 9: Perfect foresight paths for DynareOBC




Figure 10: Sudden Stops Model: Impulse Response Functions to a Negative TFP Shock




Figure 11: Spectral Density Functions for the Sudden Stops Model

