A Global Dynamic Nonlinear Solution Framework and the Repeated Transition Method[1]

Online Appendix

A Proof of the sufficient statistic proposition

Proposition 1 (The qualification for the sufficient statistic)**.**
Let Assumptions 1–4 in the main text (stability, uniqueness, Harris recurrence, finite exogenous support) hold. For a sufficiently large T, if there exists a time series of a variable $\{e_t\}_{t=0}^{T}$ such that for each time partition $\mathcal{T}_S = \{t|S_t = S\}$, for all S and for all (a, s) in the support of Φ_t,

$$(i) \quad e_{\tau_0} < e_{\tau_1} \iff V_{\tau_0}(a,s) < V_{\tau_1}(a,s) \text{ for any } \tau_0, \tau_1 \in \mathcal{T}_S$$

or

$$(ii) \quad e_{\tau_0} < e_{\tau_1} \iff V_{\tau_0}(a,s) > V_{\tau_1}(a,s) \text{ for any } \tau_0, \tau_1 \in \mathcal{T}_S,$$

then within each $\mathcal{T}_S$, periods sharing the same (e_t, S_t) yield the same continuation objects: for any $\widetilde{\tau}$ with $e_{\widetilde{\tau}} = e_t$ and $S_{\widetilde{\tau}} = S_t$, $V_{\widetilde{\tau}}(a,s) = V_t(a,s)$ for all (a,s) in the support of Φ_t. Consequently, whenever an exact match is available, the RTM's constructed conditional expectations coincide with the true ones. The finite-T case with only bracketing candidates is handled by the Interpolation Error Bound (Proposition 2 in the main text; proof below).

Proof.
Suppose $\widetilde{\tau} \in \mathcal{T}_S$ satisfies $e_{\widetilde{\tau}} = e_t$. By the contrapositive of condition (i) (or (ii)): if $V_{\widetilde{\tau}}(a,s) \neq V_t(a,s)$ for some (a,s) in the support of Φ_t, then $e_{\widetilde{\tau}} \neq e_t$, contradicting the hypothesis. Therefore $V_{\widetilde{\tau}}(a,s) = V_t(a,s)$ for all (a,s) in the support of Φ_t, and the RTM's constructed conditional expectations coincide with the true ones whenever an exact match is available. The finite-T case with only bracketing candidates is handled by the Interpolation Error Bound (Proposition 2 in the main text; proof below).

■

Proposition 2 (Interpolation Error Bound)**.**
Let Assumptions 1–4 in the main text hold. Suppose:

[1]Hanbaek Lee, University of Cambridge, Email: hl610@cam.ac.uk

(i) For each S, the conditional expectation $\mathcal{E}(\mathbf{x}; e, S) \equiv \mathbb{E}\big[V(\mathbf{x}'; \Phi', S') \,\big|\, e, S\big]$ is Lipschitz continuous with constant L_S uniformly over the individual state $\mathbf{x}$.

(ii) The converged simulation path visits each state S at times $\mathcal{T}_S$, with mesh width $h_S \equiv \max_i(e_{(i+1)} - e_{(i)})$, where $e_{(1)} \le \cdots \le e_{(|\mathcal{T}_S|)}$ are the sorted values of $\{e_t : t \in \mathcal{T}_S\}$.

Then the pointwise interpolation error satisfies $\max_t \big|\mathcal{E}^{\mathrm{RTM}}(\mathbf{x}; e_t, S_t) - \mathcal{E}(\mathbf{x}; e_t, S_t)\big| \le \frac{L}{2} \cdot h$, where $L = \max_S L_S$ and $h = \max_S h_S$.

Proof.
Fix a query point (e^*, S) and individual state $\mathbf{x}$. The RTM identifies two consecutive iso-shock candidates $e_\ell \le e^* \le e_u$ with $S_\ell = S_u = S$ and $e_u - e_\ell \le h_S$. The interpolation weight is $\omega = (e^* - e_\ell)/(e_u - e_\ell)$, so

$$\mathcal{E}^{\mathrm{RTM}}(\mathbf{x}; e^*, S) = (1-\omega)\,\mathcal{E}(\mathbf{x}; e_\ell, S) + \omega\,\mathcal{E}(\mathbf{x}; e_u, S).$$

By the triangle inequality and Lipschitz continuity:

$$\begin{aligned}\big|\mathcal{E}^{\mathrm{RTM}} - \mathcal{E}(\mathbf{x}; e^*, S)\big| &\le (1-\omega)\,L_S(e^* - e_\ell) + \omega\,L_S(e_u - e^*) \\ &= L_S\big[(1-\omega)(e^* - e_\ell) + \omega(e_u - e^*)\big].\end{aligned}$$

Substituting $\omega = (e^* - e_\ell)/(e_u - e_\ell)$:

$$= \frac{2\,L_S\,(e^* - e_\ell)(e_u - e^*)}{e_u - e_\ell}.$$

By the AM–GM inequality, $(e^* - e_\ell)(e_u - e^*) \le (e_u - e_\ell)^2/4$, so

$$\big|\mathcal{E}^{\mathrm{RTM}} - \mathcal{E}(\mathbf{x}; e^*, S)\big| \le \frac{L_S}{2}(e_u - e_\ell) \le \frac{L_S}{2}\,h_S \le \frac{L}{2}\,h.$$

■

Proposition 3 (Consistency)**.**
Suppose the conditions of Proposition 1 and the Interpolation Error Bound hold, with bounded returns and compact state spaces. Suppose additionally that (i) the Bellman operator $\mathcal{T}$ is a contraction on the space of bounded continuous functions with modulus $\beta < 1$; (ii) there exists a class $\mathcal{F}$ of bounded continuous functions, invariant under both $\mathcal{T}$ and $\widetilde{\mathcal{T}}_T$ and containing their fixed points, such that the map $e \mapsto \beta\,\mathbb{E}[f(\cdot) \mid e, S]$ is Lipschitz with a uniform constant L over $f \in \mathcal{F}$; and (iii) $h_T \to 0$ almost surely as $T \to \infty$. If the RTM iterations converge to a fixed point V^{RTM} of the approximate operator $\widetilde{\mathcal{T}}_T$, then the fixed point is close to the true

value function:

$$\|V_T^{\mathrm{RTM}} - V^*\|_\infty \leq \frac{L \cdot h_T}{2(1-\beta)} \xrightarrow{a.s.} 0 \quad \text{as } T \to \infty,$$

where V^* is the unique fixed point of $\mathcal{T}$.

Proof.

The proof uses the standard fixed-point perturbation inequality for contractions: if $\mathcal{T}$ and $\widetilde{\mathcal{T}}$ are contractions on $(\mathcal{F}, \|\cdot\|_\infty)$ with modulus $\beta < 1$ and fixed points V^* and V^{RTM}, and if $\sup_{f \in \mathcal{F}} \|\widetilde{\mathcal{T}} f - \mathcal{T} f\|_\infty \leq \varepsilon$, then $\|V^{\mathrm{RTM}} - V^*\|_\infty \leq \varepsilon/(1-\beta)$.

We apply this with $\mathcal{T}$ the Bellman operator and $\widetilde{\mathcal{T}}_T$ the RTM's approximate operator (which replaces exact continuation values with interpolated values on the simulated grid of length T). At convergence, the RTM's simulation path is fixed, so $\widetilde{\mathcal{T}}_T$ is a well-defined operator on the space of bounded continuous functions. Since $\widetilde{\mathcal{T}}_T$ composes discounting (by $\beta < 1$), maximization over actions (non-expansive), and linear interpolation of continuation values (non-expansive, since the weights ω depend only on the sufficient statistic e, not on the individual state $\mathbf{x}$), $\widetilde{\mathcal{T}}_T$ is a contraction with the same modulus β. By the Banach fixed-point theorem, it has a unique fixed point $V^{\mathrm{RTM}} = \widetilde{\mathcal{T}}_T V^{\mathrm{RTM}}$.

Step 1: Bound the operator approximation error. The only difference between $\widetilde{\mathcal{T}}_T$ and $\mathcal{T}$ is that the continuation value at aggregate state e_{t+1} is replaced by an interpolation between the two nearest grid points. Extending the argument of the Interpolation Error Bound (Proposition 2 in the main text) to any $f \in \mathcal{F}$ in place of the equilibrium conditional expectation: by the uniform Lipschitz assumption (condition (ii) of the proposition), the map $e \mapsto \beta\, \mathbb{E}[f(\cdot, s')|e, S]$ is Lipschitz with constant L uniformly over $f \in \mathcal{F}$, and the approximate operator satisfies

$$\|\widetilde{\mathcal{T}}_T f - \mathcal{T} f\|_\infty \leq \frac{L \cdot h_T}{2}$$

for all $f \in \mathcal{F}$, where h_T is the mesh width. Since $\mathcal{F}$ is invariant under both operators and contains both fixed points V^* and V_T^{RTM}, the bound applies at the fixed points.

Step 2: Apply the perturbation inequality. With $\varepsilon_T = L h_T/2$ from Step 1,

$$\|V^{\mathrm{RTM}} - V^*\|_\infty \leq \frac{\varepsilon_T}{1-\beta} = \frac{L \cdot h_T}{2(1-\beta)}.$$

Step 3: Convergence. By hypothesis (iii), $h_T \to 0$ almost surely as $T \to \infty$ (sufficient conditions—compact support of e conditional on S with density bounded away from zero, together with Harris recurrence—are discussed in Section 2.5 of the main text). Therefore $\|V_T^{\mathrm{RTM}} - V^*\|_\infty \to 0$ almost surely. ■

B Implementation details

B.1 Implementation

The method starts by simulating a single path of exogenous aggregate shocks for a long-enough period T, $\mathbb{S} = \{S_t\}_{t=0}^{T}$, using the aggregate transition matrix Π^S. A time partition $\mathcal{T}(S)$ is defined as a group of periods that share the same aggregate exogenous state realization as follows.

$$\mathcal{T}_S := \{\tau | S_\tau = S\} \subseteq \{0, 1, \ldots, T\} \text{ for } S \in \{B, G\}.$$

The pseudo algorithm of the RTM is as follows:

Step 1. Guess on the paths of the value functions, the endogenous states, and the prices:[2]

$$\{V_t^{(n)}, \Phi_t^{(n)}, p_t^{(n)}\}_{t=0}^{T}.$$

Step 2. Solve the model backward from the terminal period T in the following sub-steps. The explanation is based on an arbitrary period t. For an illustrative purpose, I assume $S_t = G$ and $S_{t+1} = G$:

2-a. Find $\widetilde{t}+1$ where the endogenous aggregate allocation is closest to the one in period $t+1$, but the shock realization is different from period $t+1$ ($S_{\widetilde{t}+1} = B$):[3]

$$\widetilde{t}+1 = \arg \inf_{\tau \in \mathcal{T}_B} ||\Phi_\tau^{(n)} - \Phi_{t+1}^{(n)}||_\infty.$$

2-b. Compute the expected future value function as follows:

$$\mathbb{E}_t \widetilde{V}_{t+1} = \pi_{G,G} V_{t+1}^{(n)} + \pi_{G,B} V_{\widetilde{t}+1}^{(n)}.$$

2-c. Using $\mathbb{E}_t \widetilde{V}_{t+1}$ and $p_t^{(n)}$, solve the individual agent's problem at the period t. Then, I obtain the solution $\{V_t^*, a_{t+1}^*\}$. The individual problem is solved via policy function iteration on the Euler equation with linear interpolation on the individual endogenous state grid; Appendix C.4 confirms that cubic spline interpolation produces virtually identical results for one-dimensional grids. The forward simulation step distributes household mass to bracketing grid points via linear interpolation.

[2] In practice, I use the stationary equilibrium allocations for all periods as the initial guess.

[3] Such $\widetilde{t}+1$ might not be unique. However, any of such $\widetilde{t}+1$ is equally good to be used in the next step.

After taking these sub-steps for all t, $\{V_t^*, a_{t+1}^*\}_{t=0}^T$ are available.

Step 3. Using $\{a_{t+1}^*\}_{t=0}^T$, simulate forward the time series of the endogenous states $\{\Phi_t^*\}_{t=0}^T$ starting from $\Phi_0^* = \Phi_0^{(n)}$.[4]

Step 4. Using $\{\Phi_t^*\}_{t=0}^T$, all the aggregate allocations over the whole path such as $\{K_t^*\}_{t=0}^T$ can be obtained. Using the market-clearing condition, compute the time series of the market-clearing price. If the model features a non-trivial market clearing condition, compute the time series of the implied prices $\{p_t^*\}_{t=0}^T$.[5]

Step 5. Check the distance between the implied prices and the guessed prices.

$$\sup_{T_{\text{burn}} \leq t \leq T - T_{\text{burn}}} ||p_t^* - p_t^{(n)}||_\infty < tol$$

Note that the distance is measured after excluding the burn-in periods at the beginning and the end of the simulated path. This is an adjustment to handle a potential bias from the imperfect guesses on the terminal period's value function $V_T^{(n)}$ and the initial period's endogenous state $\Phi_0^{(n)}$. The convergence criterion can be augmented by including the distance in other allocations, such as value functions or endogenous states.

Boundary conditions. At the terminal period T, the backward pass computes continuation values by treating the unrealized future outcomes as usual—that is, the matching step identifies counterfactual candidates from the simulated path and forms conditional expectations in the standard way. At period 1, the forward pass does not update the initial distribution $\Phi_0^{(n)}$; it is held fixed at the steady-state distribution (or, optionally, at the terminal period's distribution $\Phi_T^{(n-1)}$ from the previous iteration, which can economize on burn-in by starting the simulation closer to the ergodic set). In

[4] In this step, if the endogenous state is a distribution, I use the non-stochastic simulation method (Young, 2010).

[5] It is worth noting that the prices here are not the market-clearing prices that are determined from the interactions between demand and supply. Rather, they are the prices implied by the market-clearing condition given either demand or supply fixed at the n^{th} iteration:

$$p_t^* = \arg_{\widetilde{p}}\{Q^D(p_t^{(n)}, X_t, X_{t+1}) - Q^S(\widetilde{p}, X_t, X_{t+1}) = 0\} \text{ or}$$
$$p_t^* = \arg_{\widetilde{p}}\{Q^D(\widetilde{p}, X_t, X_{t+1}) - Q^S(p_t^{(n)}, X_t, X_{t+1}) = 0\}.$$

In the computation method used in Krusell and Smith (1997), a market-clearing price needs to be computed in an additional loop due to the non-trivial market-clearing condition. The implied price cannot replace the market-clearing price in this method, as the misspecified price prediction rule can lead to a divergent law of motion of the aggregate allocation. In contrast, due to the missing market clearing step, the RTM significantly saves computation time.

all cases, sufficiently long burn-in periods at both ends of the simulated path ensure that these boundary choices do not affect the measured equilibrium dynamics in the interior. A practical diagnostic for burn-in adequacy is to verify that the convergence error $\|p_t^* - p_t^{(n)}\|$ does not systematically increase near the burn-in boundaries; if it does, the burn-in should be extended.

If the distance is smaller than the tolerance level, the algorithm is converged. A practical note on early convergence: if the algorithm converges unusually quickly (e.g., within a few dozen iterations), users should verify that this reflects genuine solution accuracy rather than insufficient exploration of the ergodic set. Fast convergence arises naturally when aggregate dynamics are approximately linear (e.g., the Krusell and Smith 1998 model); in such cases, early convergence is benign. When the model features genuinely nonlinear dynamics, fast convergence may instead indicate that the initial distribution is a special case that does not allow the simulated path to explore the full ergodic set. Users can diagnose this by checking the linearity of key allocations in the sufficient statistic (Section 2.5 of the main text) and by verifying robustness across different random seeds (Section 4.3 of the main text). The stationary equilibrium provides a robust default initial guess.

Otherwise, I make the following updates on the guess:[6]

$$p_t^{(n+1)} = p_t^{(n)}\psi_1 + p_t^*(1-\psi_1)$$

$$V_t^{(n+1)} = V_t^{(n)}\psi_2 + V_t^*(1-\psi_2)$$

$$\Phi_t^{(n+1)} = \Phi_t^{(n)}\psi_3 + \Phi_t^*(1-\psi_3)$$

for all $t \in \{0, 1, \ldots, T\}$. With the updated guess $\{V_t^{(n+1)}, \Phi_t^{(n+1)}, p_t^{(n+1)}\}_{t=0}^T$, I go back to Step 1.

(ψ_1, ψ_2, ψ_3) are the parameters of convergence speed in the algorithm. If ψ_i is high, then the algorithm conservatively updates the guess, leaving the algorithm to converge slowly. If the equilibrium dynamics are almost linear, as in Krusell and Smith (1998), uniformly setting ψ_i at around 0.9 guarantees convergence at a fairly high convergence speed. However, if a model is highly nonlinear, the convergence speed needs to be controlled to be substantially

[6] In highly nonlinear aggregate dynamics, I have found that the log-convex combination updating rule marginally dominates the standard convex combination updating rule in terms of convergence speed. The log-convex combination rule is as follows:

$$\log(p_t^{(n+1)}) = \log(p_t^{(n)})\psi_1 + \log(p_t^*)(1-\psi_1).$$

slower than the one in the linear models. This is because the nonlinearity can lead to a sudden jump in the realized allocations during the iteration if a new guess is too dramatically changed from the last guess. A heterogeneous updating rule $\psi_i \neq \psi_j$ $(i \neq j)$ is also helpful in cases where the dynamics of certain allocations are particularly more nonlinear than the others.

Depending on the models, the guessed bundle can include the paths of the optimal decision rules on top of the basic three allocations. For example, in an application for a global nonlinear solution for the canonical New Keynesian model with Rotemberg price adjustment cost, the paths of labor demand and consumption are all included in the bundle of guess and simultaneously updated with the inter-temporal policy functions, endogenous states, and prices.[7]

As can be seen from the convergence criterion in Step 5, the algorithm stops when the predicted allocation paths (nth iteration) are close enough to the realized allocation paths (with asterisks). Therefore, once convergence is achieved, the predicted and realized aggregate paths coincide within the specified tolerance. If this internal consistency is summarized using R^2 or mean-squared error, as in Krusell and Smith (1998), the RTM's predicted and realized paths approach $R^2 = 1$ and a negligible MSE. This is a convergence diagnostic rather than an external accuracy test, which is why Section 4 of the main text also reports Euler equation errors and off-equilibrium validation.

B.2 The required length of simulation path

The required simulation length depends on two factors: the exogenous state process and the model's nonlinearity. The simulation must visit each exogenous state frequently enough for accurate interpolation. With two aggregate states and moderate persistence ($\mathbb{P}(S' = G|S = B) = 0.125$ as in Krusell and Smith 1998), 1,000 periods suffice. With a 7-point Tauchen grid covering three standard deviations, at least 3,000 periods are needed so that tail states receive at least 30 visits. Greater shock persistence requires proportionally longer simulations. Nonlinearity also matters: more observations per exogenous state provide denser interpolation nodes. Section 4.3 of the main text confirms that $T = 2{,}000$ provides adequate coverage for the HANK model with six combined exogenous states.

[7] Sample codes are available at https://sites.google.com/view/hanbaeklee/computation-lab.

B.3 Implementation of the sufficient statistic approach

In the algorithm explained in the previous section, Step 2-a is the most demanding step for heterogeneous-agent models, as it needs to find a period $\widetilde{t}+1$ that is closest to period $t+1$ in terms of distribution. Therefore, the similarity of the distributions across the periods needs to be measured, which is computationally costly.

However, if there is a sufficient statistic that can perfectly represent a period's endogenous aggregate state, the computational efficiency can be substantially improved. This enables one to locate the target period $\widetilde{t}+1$ by only comparing the distance between these sufficient statistics instead of the distributions. For example, in Krusell and Smith (1998), if the aggregate capital is the sufficient statistic, Step 2-a becomes easier as follows:

$$\widetilde{t}+1 = \arg \inf_{\tau \in \mathcal{T}_B} ||K_{\tau}^{(n)} - K_{t+1}^{(n)}||_{\infty}.$$

As the algorithm relies on the ergodicity, a sufficiently long period of simulation is needed for accurate computation. However, in practice, the simulation ends in finite periods. Therefore, two periods with arbitrarily close sufficient statistics might not exist. For this hurdle, the following adjusted versions of Step 2-a and Step 2-b interpolate between the two closest matching periods to improve the accuracy of the solution:

2-a′. Find $\widetilde{t}^{up}+1$ where the sufficient statistic of the endogenous aggregate state is closest to the one in period $t+1$ from above, but the shock realization is different from period $t+1$:

$$\widetilde{t}^{up}+1 = \arg \inf_{\tau \in \mathcal{T}_B \text{ s.t. } e_{\tau}^{(n)} \geq e_{t+1}^{(n)}} ||e_{\tau}^{(n)} - e_{t+1}^{(n)}||_{\infty},$$

where e_{τ} denotes the sufficient statistic of the endogenous aggregate state in period τ. Similarly, find $\widetilde{t}^{dn}+1$ where the sufficient statistic of the endogenous aggregate state is closest to the one in period $t+1$ from below, but the shock realization is different from period $t+1$:

$$\widetilde{t}^{dn}+1 = \arg \inf_{\tau \in \mathcal{T}_B \text{ s.t. } e_{\tau}^{(n)} < e_{t+1}^{(n)}} ||e_{\tau}^{(n)} - e_{t+1}^{(n)}||_{\infty}.$$

Then, I have $e_{\widetilde{t}^{up}+1}^{(n)}$ and $e_{\widetilde{t}^{dn}+1}^{(n)}$ that are closest to $e_{t+1}^{(n)}$ from above and below, respectively. Using these two, I compute the weight ω to be used in the convex combination

of value functions in the next step:

$$\omega = \frac{e^{(n)}_{t+1} - e^{(n)}_{\widetilde{t}^{dn}+1}}{e^{(n)}_{\widetilde{t}^{up}+1} - e^{(n)}_{\widetilde{t}^{dn}+1}}.$$

2-b′. Compute the expected future value function as follows:

$$\mathbb{E}_t \widetilde{V}_{t+1} = \pi_{G,G} V^{(n)}_{t+1} + \pi_{G,B} \left(\omega V^{(n)}_{\widetilde{t}^{up}+1} + (1-\omega) V^{(n)}_{\widetilde{t}^{dn}+1} \right).$$

Step 2-a' and Step 2-b' construct a synthetic counterfactual conditional value function by the convex combination of the two value functions that are for the most similar periods to period $t+1$. These adjusted steps help accurately solve the problem in relatively short periods of simulation. For example, the model in Krusell and Smith (1998) can be accurately solved using $T = 1{,}000$ periods of simulation (with 100 burn-in periods at the beginning and the end of the simulated path).

The step of interpolation after finding the closest periods in terms of sufficient statistics can be understood as a piecewise interpolation, in contrast to the unconditional linear interpolation used in the state space-based approach based on the regression coefficients.

C Computational analysis

C.1 Computational scaling

This subsection documents how the RTM's per-iteration computation time scales with model dimensions. The analysis uses the heterogeneous-household portfolio-choice model of Krusell and Smith (1997) with endogenous labor supply, two assets, and five idiosyncratic income states. All computations use a MacBook Pro with M3 Pro chip running MATLAB R2023b.

Table C.1 reports results for sweeps over the individual-state grid size (Panel A) and aggregate shock states (Panel B). The backward pass scales sublinearly in grid size (3.6× for 8× grid) and linearly in shock states (2.5× for 2.5× states). The forward pass scales super-linearly in grid size (14.8× for 8× grid) due to sparse matrix operations on the distribution. Final convergence errors are comparable across configurations.

C.2 Warm-start initialization

The RTM's standard initialization sets all aggregate paths to their stationary-equilibrium values ("cold start"). A natural alternative is to initialize using aggregate paths predicted

Table C.1: RTM computational scaling

Panel A: Wealth grid size (2 shock states, $T = 1{,}000$)					
Grid points	Total (s)	BW/iter (s)	FW/iter (s)	Final MSE	SS K
50	172.9	0.52	0.34	1.55×10^{-6}	12.800
100	265.4	0.66	0.67	1.62×10^{-6}	12.780
200	601.8	1.07	1.94	1.73×10^{-6}	12.760
400	1412.3	1.88	5.17	1.63×10^{-6}	12.760
Panel B: Aggregate shock states (200 grid points, $T = 1{,}000$)					
Shock states	Total (s)	BW/iter (s)	FW/iter (s)	Final MSE	
2	494.8	1.03	1.44	1.73×10^{-6}	
3	601.3	1.56	1.44	2.52×10^{-6}	
5	897.9	2.58	1.91	4.24×10^{-6}	

Notes: Each configuration runs for 200 RTM iterations. BW and FW denote per-iteration backward and forward pass times. Final MSE is the convergence error after 200 iterations.

by a law of motion (LoM) fitted to a previously converged solution ("warm start"). This subsection evaluates whether warm-start initialization improves convergence.

I fit log-linear LoMs from the converged baseline solution:

$$\log K_{t+1} = \beta_0 + \beta_1 \log K_t + \beta_2 \log A_t + \beta_3 \log K_t \cdot \log A_t, \tag{1}$$

and analogously for aggregate labor supply L_t and bond price q_t. The fitted LoMs achieve R^2 values of 0.9999 for K, 0.9994 for L, and 0.583 for q. The low R^2 for the bond price reflects the inherently nonlinear bond pricing relationship, which a log-linear specification cannot fully capture. This itself illustrates why global nonlinear methods are needed.

For the warm start, I forward-simulate these LoMs from the stationary equilibrium to generate initial aggregate paths. Both cold and warm starts use identical individual policy initializations. The warm start reduces the initial MSE by a factor of 7.6 (4.3×10^{-2} versus 5.7×10^{-3}). The critical difference emerges in later iterations: the cold start stalls at approximately 2.2×10^{-7}, oscillating without further improvement, while the warm start continues to decrease smoothly to 3.0×10^{-8}. This divergence arises because the bond price channel converges slowly under conservative damping ($\psi_q = 0.999$); when aggregate paths are already close to the true solution, bond price updates remain coherent.

C.3 Full-distribution and moment-based matching

The RTM identifies target periods in the backward pass by matching on a finite-dimensional sufficient statistic, typically aggregate capital K_t. As established in the sufficient statistic

theory section, this is justified when the sufficient statistic indexes the continuation objects relevant for matching along the ergodic path. This subsection provides a direct numerical test by replacing the K-based matching with matching on the full cross-sectional distribution Φ_t.

Specifically, in the modified backward pass, instead of bracketing the target period by finding two iso-shock candidates with the nearest K values, I find two candidates with the smallest sup-norm distributional distance:

$$d(t, \tilde{t}) = \max_{(a,s)} |\Phi_t(a, s) - \Phi_{\tilde{t}}(a, s)| . \tag{2}$$

Let $\tilde{t}_1$ and $\tilde{t}_2$ denote the two nearest neighbors with distances d_1 and d_2. The interpolation weights are computed via inverse-distance weighting: $w_1 = d_2/(d_1 + d_2)$. All other components (forward simulation, convergence criterion, damping weights, shock path) remain identical to the baseline K-only implementation.

The most informative comparison is the iteration-by-iteration convergence trajectory: if K is a valid sufficient statistic, both matching methods should select the same candidate periods and converge at the same rate. Table C.2 confirms this. The two methods produce nearly identical MSE through the entire convergence phase, and the converged solutions yield a maximum capital deviation of only 0.24% of steady state.

Table C.2: Convergence trajectory: full-distribution matching versus K-only matching

Iteration	Distribution matching (MSE)	K-only matching (MSE)
1	4.34×10^{-2}	4.34×10^{-2}
21	9.60×10^{-4}	9.77×10^{-4}
41	1.21×10^{-4}	1.21×10^{-4}
61	3.02×10^{-5}	3.05×10^{-5}
101	3.90×10^{-6}	3.94×10^{-6}
201	1.44×10^{-6}	1.73×10^{-6}

Notes: The first 5 iterations of the distribution matching use K-based matching as a fallback (distributions are identical at initialization). Both methods use the same shock realization, damping weights, and initialization.

The two solutions track closely: the maximum capital deviation is 0.24% of steady state, and the near-identical convergence trajectories in Table C.2 confirm that K-matching and distributional matching select the same candidates.

HANK model: moment-based matching For the HANK model, I replace sup-norm matching with moment-based matching. Variant 1 uses K only; Variant 2 adds B; Variant 3 adds $\text{Var}(k)$, $\text{Var}(b)$, $\text{Cov}(k, b)$. Variants 2–3 are warm-started from Variant 1 and run for

20 additional iterations.

Table C.3: Moment-based matching: two-asset HANK model

Matching variables	K (%)	L (%)	q (%)	π (pp)	RP (pp)
K	—	—	—	—	—
K, B	0.005	0.002	0.000	0.000	0.003
K, B, Var(k), Var(b), Cov(k, b)	0.007	0.002	0.000	0.000	0.003

Notes: Each entry reports $\max_t |x_t^{\text{variant}} - x_t^{\text{baseline}}|/x_{ss}$ as a percentage. The baseline is the K-only solution (Variant 1, 200 iterations from steady state). Variants 2 and 3 are warm-started from Variant 1 and run for 20 additional iterations. RP denotes the annualized risk premium.

Enriching matching beyond K produces negligible changes (Table C.3): maximum capital deviation of 0.01%, with all other variables showing near-zero deviations. Adding second moments beyond (K, B) produces no additional change, confirming K-sufficiency.

C.4 Interpolation scheme robustness

Table C.4 compares three interpolation schemes: the baseline two-point linear (A), inverse-distance weighted k-NN (B), and local linear regression (C, KS1998 only).

All schemes produce very similar results: LoM R^2 and aggregate paths are nearly indistinguishable, while EE means remain in the same accuracy range. The RTM's accuracy derives primarily from the simulated grid quality (determined by T) rather than the interpolation rule.

C.5 Mesh convergence and accuracy scaling with T

The convergence rate discussion in Section 2.5 of the main text predicts that the mean mesh width decreases at rate $O(1/T)$, since $\bar{h} = O(N_S/T)$ when the ergodic density is bounded below. I verify this prediction by solving both models at multiple values of T, holding the random seed fixed so that the first $T_{\min}$ periods of the shock path are identical across runs.

Table C.5 confirms $\bar{h} = O(1/T)$: for KS1998, $\bar{h}$ falls from 0.015 ($T = 500$) to 0.001 ($T = 5{,}000$), tracking the theoretical rate precisely. For HANK, from 0.003 ($T = 2{,}000$) to 0.001 ($T = 5{,}000$). Euler equation errors are bounded across all T values, fluctuating within a narrow range as different simulation lengths visit different ergodic samples.

Formal error bound estimates The bound $L \cdot h/2$ (Proposition 2) has two components: the Lipschitz constant L and mesh width h. For KS1998 ($T = 5{,}000$), $h = 0.089$ (0.25% of K_{ss}) with $L_{99} = 10.4$. For HANK ($T = 2{,}000$), $h = 0.028$ (0.26% of K_{ss}) with larger

Table C.4: Interpolation scheme robustness

(a) Panel A: Krusell–Smith (1998), $T = 2{,}000$

Scheme	Iters	Time (s)	LoM R^2	Mean K	EE mean
A: Baseline (2-pt linear)	1,045	793	0.9999	36.232	−4.124
B: k-NN IDW ($k = 4$)	961	772	0.9999	36.232	−3.623
C: Local linear ($k = 4$)	1,001	2,975	0.9999	36.234	−3.651
A + cubic spline (individual)	989	755	0.9999	36.222	−4.124

(b) Panel B: Two-Asset HANK, $T = 2{,}000$

Scheme	Iters	Time (s)	LoM R^2	Mean K	EE mean
A: Baseline (2-pt linear)	677	9,160	0.9999	10.930	−4.825
B: k-NN IDW ($k = 4$)	668	8,867	0.9999	10.931	−4.578
C: Local linear ($k = 4$)	660	8,934	0.9999	10.932	−4.389

Notes: EE mean is $\log_{10}$ of the distribution-weighted absolute Euler equation error. Scheme A is the baseline two-point linear interpolation between the two bracketing iso-shock candidates. Scheme B uses inverse-distance weighted interpolation over $k = 4$ nearest neighbors. Scheme C uses local linear projection $\hat{V}(K) = \hat{\alpha} + \hat{\beta}K$ over $k = 4$ neighbors. "A + cubic spline (individual)" uses Scheme A for aggregate matching but cubic spline interpolation on the individual endogenous state grid; for the HANK model, cubic spline interpolation on the non-uniform two-dimensional individual state grid (k, b) introduces numerical instability due to oscillatory behavior near grid boundaries, and is therefore infeasible. All schemes use the same random seed ($s = 100$) and simulation length.

$L_{99} \sim 10^5$ due to the capital Euler equation's interaction of returns, adjustment costs, and marginal utility. The actual EE error ($10^{-4.82}$) is orders of magnitude below the formal upper bound. The required T scales linearly with N_S: both models achieve h/K_{ss} well below 1%.

C.6 Failure case: Calvo pricing with interest rate smoothing

Section 2.7 of the main text discusses the RTM's potential failure when multiple slow-moving endogenous aggregate states interact. This subsection provides quantitative evidence using a representative-agent New Keynesian model with Calvo pricing.

Model The model is a standard representative-agent New Keynesian economy with Calvo pricing ($\theta = 0.66$), log utility ($\sigma = 1$), Frisch elasticity of 1, elasticity of substitution $\varepsilon = 6$, and a Taylor rule responding to inflation ($\phi_\pi = 1.5$) and the output gap ($\phi_y = 0.125$). Aggregate shocks include TFP ($\rho = 0.95$, $\sigma = 0.007$, 3 grid points), preference ($\rho = 0.90$, $\sigma = 0.015$, 3 grid points), and monetary policy (iid, $\sigma = 0.0025$, 3 grid points), for a total of 27 combined exogenous states. Under Calvo pricing, price dispersion Δ_t evolves according

Table C.5: Accuracy scaling with simulation length T

(a) Panel A: Krusell–Smith (1998)

T	Iters	Time (s)	LoM R^2	EE mean	$\bar{h}$
500	876	136	0.9999	−4.305	0.015
1,000	815	244	0.9999	−3.998	0.008
2,000	737	501	0.9999	−3.742	0.004
3,000	813	1,078	0.9999	−3.698	0.002
5,000	818	2,148	0.9999	−3.697	0.001

(b) Panel B: Two-Asset HANK

T	Iters	Time (s)	LoM R^2	EE mean	$\bar{h}$
2,000	677	9,160	0.9999	−4.825	0.003
3,000	494	10,100	0.9999	−4.811	0.002
5,000	481	16,200	0.9999	−4.667	0.001

Notes: For each simulation length T, the model is solved from scratch using the same random seed ($s = 100$). EE mean is $\log_{10}$ of the distribution-weighted absolute Euler equation error. $\bar{h}$ is the mean mesh width (average gap between adjacent iso-shock candidates), averaged across aggregate shock states. All KS1998 runs use burn-in = 100; HANK runs use burn-in = 500.

to

$$\Delta_t = (1-\theta)(1+\pi_t^*)^{-\varepsilon}(1+\pi_t)^{\varepsilon} + \theta(1+\pi_t)^{\varepsilon}\Delta_{t-1},$$

making it a slow-moving endogenous aggregate state. The sufficient statistic for the matching step is Δ_{t-1}. The simulation length is $T = 2{,}000$ with a burn-in of 500 periods, and the damping weight is 0.99 for both specifications.

I compare two specifications that differ only in the Taylor rule:

- *Non-persistent*: $1 + i_t = \frac{1}{\beta}(1+\pi_t)^{\phi_\pi}(Y_t/Y_t^f)^{\phi_y} e^{\varepsilon_t^m}$
- *Persistent*: $1 + i_t = (1+i_{t-1})^{\rho_r}\left[\frac{1}{\beta}(1+\pi_t)^{\phi_\pi}(Y_t/Y_t^f)^{\phi_y}\right]^{1-\rho_r} e^{\varepsilon_t^m}$ with $\rho_r = 0.85$

The persistent specification makes i_{t-1} a second slow-moving endogenous state alongside Δ_{t-1}. Both specifications use Δ_{t-1} as the sufficient statistic for the matching step.

Results Table C.6 reports the MSE trajectory (reported at selected iterations). The non-persistent specification converges: MSE falls from 6.79×10^{-6} to 3.81×10^{-8} over approximately 1,300 iterations, with non-monotonic oscillations in the early phase (rising from 1.23×10^{-6} at iteration 100 to 6.57×10^{-6} at iteration 250 before resuming decline). The persistent specification diverges: MSE briefly falls from 6.79×10^{-6} to 4.40×10^{-6} over 59 iterations, then grows exponentially to 1.36×10^{-2} by iteration 750.

Table C.6: RTM convergence: non-persistent vs. persistent Taylor rule

Iteration	Non-persistent MSE	Persistent MSE
1	6.79×10^{-6}	6.79×10^{-6}
50	1.76×10^{-6}	4.45×10^{-6}
100	1.23×10^{-6}	4.87×10^{-6}
150	2.91×10^{-6}	7.04×10^{-6}
200	5.15×10^{-6}	1.57×10^{-5}
250	6.57×10^{-6}	3.88×10^{-5}
350	6.13×10^{-6}	1.47×10^{-4}
500	2.90×10^{-6}	1.02×10^{-3}
750	5.68×10^{-7}	1.36×10^{-2}
1000	1.24×10^{-7}	—
1500	4.65×10^{-8}	—
2000	1.72×10^{-7}	—
2500	9.22×10^{-7}	—

Notes: MSE is the mean squared error between predicted and realized equilibrium paths across all variables. Both specifications use identical parameters, shocks, simulation length ($T = 2{,}000$), damping weight (0.99), and sufficient statistic (Δ_{t-1}). "—" indicates the algorithm was terminated after divergence.

Diagnosis The key difference is the number of slow-moving endogenous aggregate states. With the contemporaneous Taylor rule, Δ_{t-1} is the sole slow-moving state and serves as a valid sufficient statistic: it indexes the continuation values monotonically along the simulated path. Interest rate smoothing introduces a second slow-moving state i_{t-1} that independently affects continuation values through future monetary policy. When matching uses Δ_{t-1} alone, two matched periods may have different i_{t-1} values and therefore different continuation values, violating the sufficient statistic condition. As iterations progress, errors in the unmatched state i_{t-1} feed back into inflation dynamics through the persistent Taylor rule and amplify, producing the explosive pattern in Table C.6. Consistent with this diagnosis, the Spearman rank correlation between Δ_{t-1} and the forward-looking continuation value proxies is moderate ($|\rho| \approx 0.35$) rather than near unity, reflecting the weak monotonicity inherent in models with multiple slow-moving endogenous states. As noted in Section 2.7 of the main text, it is unclear whether this failure reflects a limitation of the algorithm or the non-existence of an ergodic recursive competitive equilibrium under this particular configuration.

Practical guidance Users should monitor MSE for initial decline followed by sustained reversal, which signals multiple interacting slow-moving states. Remedies include reformulating the model (e.g., Rotemberg instead of Calvo, contemporaneous instead of persistent

Taylor rules) or extending the RTM to multi-dimensional endogenous matching. Exploratory experiments with two-dimensional matching on (Δ_{t-1}, i_{t-1})—using Delaunay triangulation, kernel regression, and various update frequencies—confirm that the pair is the correct matching space (achieving transient MSE improvements of two orders of magnitude), but do not yet deliver stable convergence. Whether the remaining convergence difficulty reflects a fundamental limitation of multi-dimensional interpolation or the non-existence of an ergodic equilibrium under this configuration remains an open question.

D Performance comparison with existing methods

This section benchmarks the RTM against linearization, OccBin (Guerrieri and Iacoviello, 2015), and GDSGE (Cao et al., 2023) using an RBC model with irreversible investment (McGrattan, 1996; Christiano and Fisher, 2000). The RTM also converges approximately ten times faster than the Krusell and Smith (1997) algorithm for the Khan and Thomas (2008) heterogeneous-firm model, by avoiding nested market-clearing loops. All computations use a MacBook Pro with M3 Pro chip.

Consider a representative household solving:

$$V(a; X) = \max_{c,a'} \frac{c^{1-\sigma}}{1-\sigma} + \beta \mathbb{E} V(a'; X') \tag{3}$$

$$\text{s.t. } c + a' - (1-\delta)a = Aa^{\alpha} \tag{4}$$

$$a' - (1-\delta)a \geq \phi I_{ss} \tag{5}$$

where V is the value function of a household. The value function's arguments are wealth a and the aggregate state X. c is consumption and σ is the risk-aversion parameter. I_{ss} is the steady-state investment level. ϕ is the parameter for the degree of the irreversibility. δ is the depreciation rate, and α is the capital share in the production function $F(a; A) := Aa^{\alpha}$. Apostrophes denote next-period variables. The aggregate state X is as follows:

$$X = [K, A]. \tag{6}$$

K is the aggregate capital stock, satisfying $a = K$ in equilibrium, as the capital market clears. A is TFP that follows the log AR(1) process:

$$\log(A') = \rho \log(A) + \sigma\epsilon, \quad \epsilon \sim N(0, 1). \tag{7}$$

The model features highly nonlinear aggregate dynamics due to the occasionally bind-

ing irreversibility constraint for capital investment. Therefore, besides the macroeconomic implications, the model serves as an ideal testing ground for the accuracy of the different methods for the nonlinear solutions. For precise comparison, I generate a single TFP path using the Tauchen method (7 grid points covering three standard deviations) and apply this path to all solution methods. The simulation runs for 5,000 periods with 500 burn-in periods. Each method exhibits a trade-off between accuracy and computational efficiency depending on convergence criteria. In this comparison, I tune the repeated transition method to stop after around 90 seconds, matching the speed of the GDSGE toolkit.

The first four rows of Table D.1 compare solution accuracy across methods, with columns presenting results for the RTM, GDSGE, OccBin, and linearized solutions, respectively. I evaluate accuracy using two criteria: prediction error ($Error_t$) and Euler equation error (EE_t) following Judd (1992).[8] The prediction error measures the discrepancy between predicted and realized equilibrium paths and is defined as:

$$Error_t = K_t^{(n)} - K_t^* \tag{8}$$

where $\{K_t^{(n)}\}_{t=1}^T$ represents the capital stock sequence from each solution method, and K_t^* is the implied capital path assuming agents expect $\{K_t^{(n)}\}_{t=1}^T$. The RTM constructs the period-specific expected future allocations by properly combining allocations in the predicted path (the previous iteration). Then, it provides the realized allocations implied by the prediction path. Convergence requires these predicted and realized paths to coincide. This allows the RTM to serve as a diagnostic tool for other solution methods—by feeding their simulated paths as predicted paths into the RTM algorithm, the RTM can evaluate their prediction accuracy.

Across all four error metrics in Table D.1, the RTM shows the lowest values among the comparison methods.

D.1 Confidence intervals for state-dependent GIRFs

Figure D.1 replicates the main-text MP GIRF figure with 95% bootstrap confidence intervals around each regime's mean. For each conditioning period t_0 in the tight or loose group, the GIRF is computed by averaging 2,000 paired stochastic simulations, yielding a per-conditioning-period mean response. The bootstrap (10,000 resamples of conditioning periods within each regime) then constructs a confidence interval for the regime mean at every horizon. Non-overlapping bands between the tight and loose regimes indicate statistically significant state dependence.

[8] The Euler equation error specification follows Guerrieri and Iacoviello (2015).

Table D.1: Comparison across the solution methods

	RTM	GDSGE	OccBin	Linear
Accuracy				
$\max(\lvert Error_t \rvert)$ (% of steady-state K)	0.003	0.735	1.317	2.019
$\sqrt{\text{mean}(Error_t^2)}$ (% of steady-state K)	0.001	0.025	0.217	0.559
$\max(\lvert EE_t \rvert)$ (% of contemp. C_t)	0.014	0.057	2.854	2.323
$\sqrt{\text{mean}(EE_t^2)}$ (% of contemp. C_t)	0.001	0.059	0.775	0.707
Business cycle stat.				
mean(I)	0.363	0.363	0.365	0.363
mean(C)	1.165	1.166	1.164	1.160
vol(I)	0.022	0.022	0.023	0.023
vol(C)	0.053	0.052	0.052	0.052
skewness(I)	1.388	1.320	1.307	1.407
skewness(C)	-0.210	-0.213	-0.322	-0.095
kurtosis(I)	4.514	4.578	4.513	4.255
kurtosis(C)	2.728	2.546	2.858	2.796

Notes: The upper part of the table compares the accuracy of different computation methods based on four criteria: 1) maximum absolute prediction error, 2) square root of mean squared prediction error, 3) maximum absolute Euler equation error, 4) square root of mean squared Euler equation error. The bottom part of the table compares the computed equilibrium's business cycle statistics.

For the monetary policy shock (Figure D.1), the bands are clearly separated at impact for all six variables. A bootstrap test for the difference in means confirms significance at the 95% level: the bootstrap CI for the tight-minus-loose difference excludes zero for output ($[-0.124, -0.082]$ percentage points), consumption ($[+0.007, +0.029]$), inflation ($[-0.112, -0.065]$ pp ann.), bond price ($[+0.028, +0.044]$), dividends ($[+0.602, +0.973]$), and risk premium ($[+0.120, +0.192]$ pp ann.). The state-dependent patterns documented in the main text are thus statistically significant despite substantial within-regime heterogeneity across conditioning periods.

Table D.1's lower panel reports business cycle statistics across solutions. While lower-order moments show negligible differences, higher-order moments (skewness and kurtosis) reveal significant variations, reflecting each method's treatment of the occasionally binding constraint.

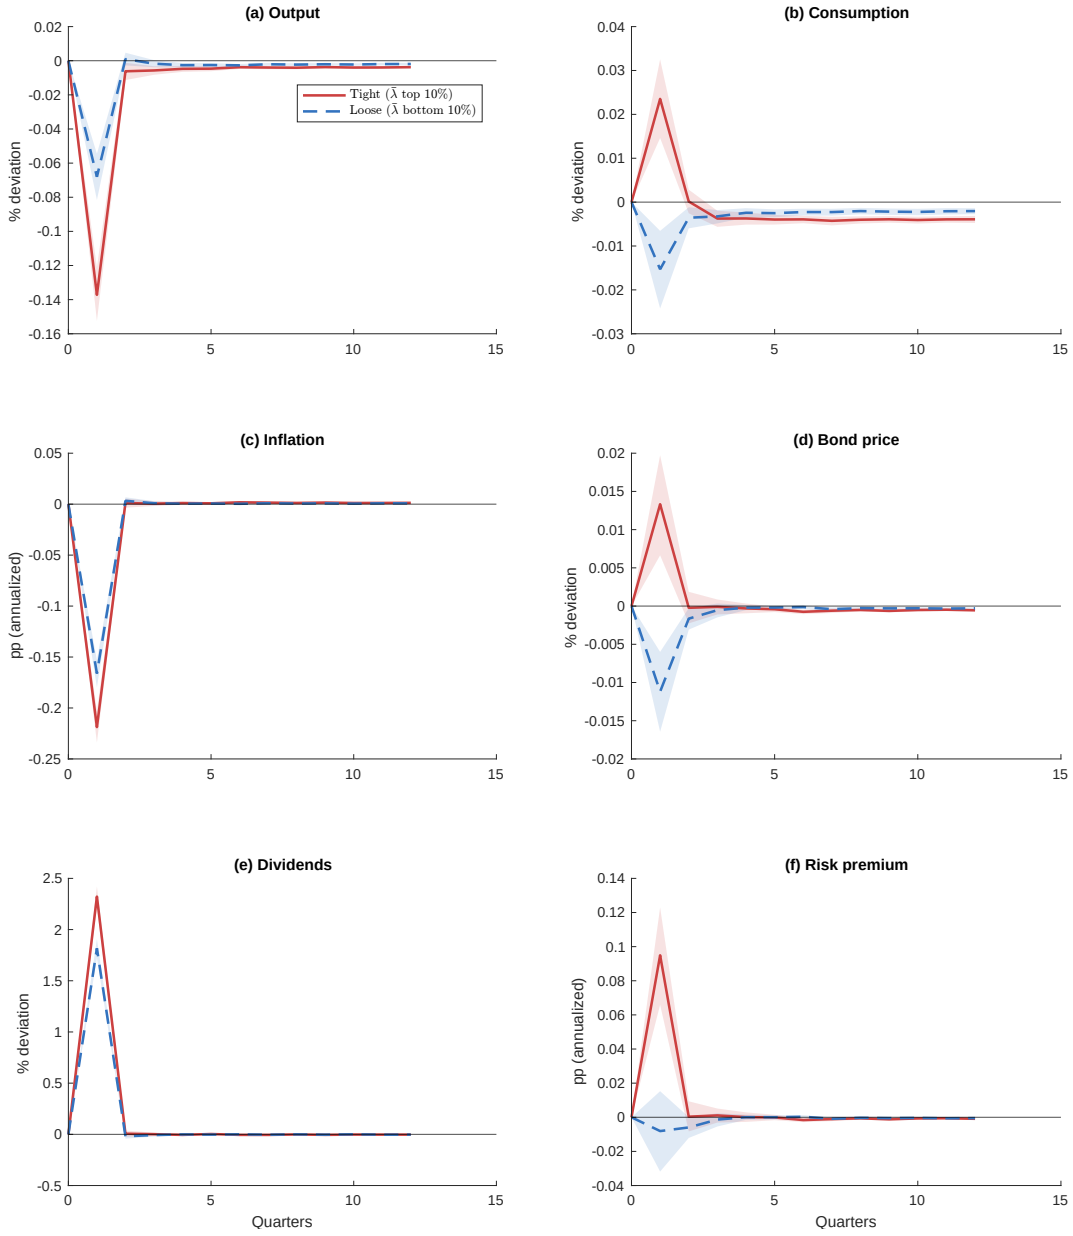

Figure D.1: State-dependent GIRFs to a contractionary MP shock (with 95% bootstrap CI)

Notes: As in the main-text MP GIRF figure (without the linearized comparison). Red solid lines: tight constraint ($\bar{\lambda}$ top 10%); blue dashed lines: loose constraint ($\bar{\lambda}$ bottom 10%). Shaded bands indicate 95% bootstrap confidence intervals for each regime's mean GIRF (10,000 resamples of conditioning periods). Non-overlapping bands indicate statistically significant state dependence (see text for numerical values).

E Generalized transition function (GTF)

In this section, I define a generalized transition function (GTF), which is a sub-path of the recursive competitive equilibrium (RCE), and the RTM solution directly yields the GTF. The GTF nests generalized impulse response functions (GIRF) (Koop et al., 1996; Andreasen

et al., 2017) and stochastic growth paths (Justiniano and Primiceri, 2008; Hansen et al., 2008).[9]

Definition 1 (Generalized transition function (GTF))**.**
Given an aggregate state realization (Φ_0, S_0) *in the RCE, the generalized transition function* g_j *of the variable of interest* v *is defined as follows:*

$$g_j(v; \Phi_0, S_0) = \int v(x; \Phi_j, S_j) d\Phi_j, \quad S_j \sim \Gamma^j(S_j; S_0), \quad j \geq 1 \tag{9}$$

where S_j *is a random variable of the exogenous aggregate state, which follows a* j*-length Markov chain* Γ^j *from the initial realization of* S_0*, and* $\Phi_j = \Phi_j(S_1, \ldots, S_j; \Phi_0)$ *is the distribution at horizon* j *determined by the equilibrium law of motion and the realized shock sequence. Since* S_j *is random,* g_j *is itself a random variable for each* j*.*

Notably, Φ_0 is assumed to be a RCE allocation. However, a realistic analysis of the stochastic growth or structural transformation might need to consider Φ_0 off the RCE. Cao (2020) establishes the existence of dynamic stochastic competitive equilibrium (DSCE) paths starting from arbitrary initial points under mild assumptions that encompass variants of Krusell and Smith (1998). However, characterizing the initial DSCE path from a non-equilibrium initial condition requires non-trivial computational steps because the path dynamics cannot be reduced to a function of the current aggregate state alone.[10] In computing the generalized transition function (GTF) from an initial state off the RCE, I assume that agents immediately adopt the policy function associated with the closest endogenous equilibrium state that shares the same exogenous state. This assumption is justified by the observation that the RTM's simulated path, under Harris recurrence, visits a dense set of endogenous states in the support of the stationary distribution. The matching step therefore involves interpolation rather than extrapolation for states of economic interest.[11]

Definition 2 (An extended generalized transition function)**.**
Given an aggregate state realization (Φ_0^{off}, S_0) *off the RCE, the extended generalized transition function* g_j *of the variable of interest* v *is defined as follows:*

$$g_j(v; \Phi_0^{off}, S_0) = g_j(v; \Phi_0, S_0) = \int v(x; \Phi_j, S_j) d\Phi_j, \quad S_j \sim \Gamma^j(S_j; S_0), \quad j \geq 1 \tag{10}$$

[9] The definition assumes a heterogeneous-agent setup with micro-level v for comprehensive explanation, but a representative-agent counterpart is consistently defined by omitting the integration.

[10] A state admitting such a reduction would already belong to the set of RCE allocations given the uniqueness of the RCE.

[11] Cao (2020) provides theoretical justification under transversality conditions. Since the ergodic set covers a broad range of distributions, the initial jump is typically small relative to the scale of the transition being analyzed. At states far from the ergodic set, the matching step may involve extrapolation, reducing accuracy.

where $\Phi_0 = \arg \inf_{\widetilde{\Phi} \in \Omega^{RCE}(S_0)} |e(\Phi_0^{off}) - e(\widetilde{\Phi})|$ *is the nearest equilibrium distribution with the same initial exogenous state* S_0*, measured in terms of the sufficient statistic e (Section 2.5 of the main text), and* S_j *is a random variable of the exogenous aggregate state, which follows a* j*-length Markov chain* Γ^j *from the initial realization of* S_0*.*[12]

For brevity, I denote g_j as a transition function at horizon j, so each g_j is also a random variable. Notably, $\mathcal{G} := \{g_1, g_2, \dots\}$ is a sub-path of the RCE. Therefore, $\mathcal{G}$ is readily computable from the global solution computed by the RTM. Specifically, for each (Φ_j, S_j), v can be obtained by identifying the period sharing the closest aggregate states in the global solution path and then interpolating the functions.

A certain magnitude of exogenous shocks is often considered in impulse response analyses. In the GTF, one can flexibly consider an arbitrary magnitude of the initial exogenous shock by fixing the initial exogenous state realization S_1. Still, this response path is a subpath of the RCE. I formally define the refined generalized impulse response function as follows:

Definition 3 (A refined generalized impulse response function)**.**
Given an aggregate state realization (Φ_0, S_0) *in the RCE, the refined generalized impulse response* g_j^{girf} *of the variable of interest* v *to an exogenous state realization* S_1 *in the following period is as follows:*

$$g_1^{girf}(v; \Phi_0, S_0, S_1) = \int v(x; \Phi_1, S_1) d\Phi_1 \tag{11}$$

$$g_j^{girf}(v; \Phi_0, S_0, S_1) = \int v(x; \Phi_j, S_j) d\Phi_j, \quad S_j \sim \Gamma^{j-1}(S_j; S_1), \quad j > 1 \tag{12}$$

where S_j *is a random variable of the exogenous aggregate state, which follows a* $(j-1)$*-step Markov chain from the fixed initial shock realization* S_1*.*

Given Line (11), the first component g_1^{girf} is deterministic, as the magnitude of the shock $|S_1 - S_0|$ is set by a researcher from choosing S_1. The average path with the 95% confidence intervals can be characterized based on simulated shock paths.[13] Notably, the average GTF path differs from the GTF path under the average exogenous process because the GTF can be nonlinear. It also differs from the perfect-foresight growth path studied in Section F.1. The state dependence in the response path could be flexibly analyzed through varying Φ_0 and S_0 with $|S_1 - S_0|$ fixed.

[12] $\Omega^{RCE}(S_0)$ denotes the set of all distributions realized along the recursive competitive equilibrium path when the exogenous state equals S_0. The projection uses the sufficient statistic rather than a norm on distributions, consistent with the RTM's matching step.

[13] Because the RTM solution already contains a dense collection of equilibrium sub-paths on the simulated ergodic support, computing GTFs for many shock paths incurs minimal extra computation cost.

F Additional applications

The subsequent sections present two applications based on heterogeneous household business cycle models. These applications address important macroeconomic questions that have remained unexplored due to computational barriers.

F.1 Application I: A heterogeneous-household model with endogenous labor supply, investment irreversibility, and fiscal spending shock

The first application is a heterogeneous-household business cycle model with endogenous labor supply, investment irreversibility, and both aggregate TFP and fiscal spending shocks. A unit continuum of ex-ante identical households face uninsurable idiosyncratic labor productivity shocks and make endogenous labor supply decisions n. The recursive formulation is:

$$V(a,z;X) = \max_{c,n,a'} \ \log(c) - \frac{\eta}{1+\frac{1}{\chi}} n^{1+\frac{1}{\chi}} + \beta \mathbb{E} V(a',z';X') \tag{13}$$

$$\text{s.t.} \quad c + a' = (1+r(X))a + w(X)zn - T(X) \tag{14}$$

$$a' - (1-\delta)a \geq \phi I_{ss} \tag{15}$$

$$\Phi' = \Gamma_X(X), \quad S' \sim \pi(S'|S), \quad z' \sim \pi(z'|z) \tag{16}$$

where r and w are competitively determined capital rent and wage; I_{ss} is steady-state investment; T is lump-sum tax; χ is the Frisch elasticity; η is labor disutility; and ϕ governs saving irreversibility. A CRS Cobb-Douglas production sector operates as $\max_{K,L} AK^{\alpha}L^{1-\alpha} - w(X)L - (r(X)+\delta)K$. The aggregate state is $X = \{\Phi, A, G\}$, where A and G follow log AR(1) processes. Government balances its budget by lump-sum taxation: $T(X) = G$.[14] Parameters are in Appendix H.

The recursive competitive equilibrium clears labor and capital markets:

$$\text{(Labor)} \quad L(X) = \int zn(a,z;X)d\Phi, \qquad \text{(Capital)} \quad K(X) = \int ad\Phi. \tag{17}$$

The labor market clearing is non-trivial as wages determine individual labor supply, requiring aggregation. The monotonicity condition in Proposition 1 is verified via Spearman rank correlation tests, with minimum coefficients exceeding $1-10^{-15}$ across all state combinations.

[14] Uniform lump-sum transfers are not neutral in heterogeneous-agent models when households have heterogeneous marginal propensities to consume. An alternative is transfers proportional to productivity. The uniform specification is adopted here for simplicity, as the focus is on demonstrating the solution method.

The solution reveals highly nonlinear capital dynamics driven by the occasionally binding constraint: while predicted and realized RTM paths coincide, a log-linear law of motion deviates significantly ($R^2 = 0.920$). Comparing the heterogeneous-household model (HH) with a representative-household variant (RH) that eliminates labor productivity heterogeneity, HH exhibits 5% lower output volatility and 18% lower investment volatility (Table F.1), while heterogeneous-firm models show the opposite pattern. When the irreversibility constraint is removed, HH and RH paths coincide with log-linear predictions, confirming that the constraint drives the nonlinearity.

Table F.1: Business cycle statistics: heterogeneous (HH) vs. representative (RH)

	Heterogeneous	Representative
Mean		
Output	0.512	0.508
Consumption	0.288	0.287
Investment	0.122	0.120
Volatility		
log(*Output*)	0.042	0.044
log(*Consumption*)	0.057	0.058
log(*Investment*)	0.049	0.060
Skewness		
log(*Output*)	-0.026	0.024
log(*Consumption*)	-0.363	-0.349
log(*Investment*)	0.937	0.874

Notes: Business cycle statistics for the heterogeneous agent model (first column) and the representative counterpart (second column). Full time-series figures are in the supplementary material.

F.1.1 State-dependent fiscal multipliers

The occasionally binding constraint creates state-dependent MPCs: constrained households exhibit near-unity consumption responses to negative income shocks. With $\phi = 0.975$ as in Guerrieri and Iacoviello (2015), the model produces 33.9% hand-to-mouth households at steady state, consistent with Kaplan and Violante (2014), with 21.6% holding above-average wealth ("wealthy hand-to-mouth").

A positive fiscal shock generates heterogeneous output responses depending on the share of constrained households:

$$[\textit{Mostly constrained}]: \quad \underbrace{C_t}_{\downarrow\downarrow\downarrow} + \underbrace{I_t}_{\downarrow} + \underbrace{G_t}_{\uparrow\uparrow} = Y_t \ \uparrow\uparrow \quad (\because \text{Large wealth effect}) \tag{18}$$

$$[\textit{Mostly unconstrained}]: \quad \underbrace{C_t}_{\downarrow\downarrow} + \underbrace{I_t}_{\downarrow\downarrow} + \underbrace{G_t}_{\uparrow\uparrow} = Y_t \ \uparrow \quad (\because \text{Small wealth effect}) \tag{19}$$

I estimate state dependence using simulated data:

$$Y_t = \beta_0 + \beta_1 G_t + \beta_2 G_t \times \Lambda_t + \beta_3 \log(K_t) + \beta_4 \log(A_t) + \epsilon_t, \tag{20}$$

where $\Lambda_t := \int \lambda(a, z; X_t) d\Phi_t$ captures constraint intensity. Table F.2 reports results across HH, RH, and GHH specifications. The baseline yields a fiscal multiplier of approximately 0.8 over two years, consistent with Ramey (2020). Including the interaction term reveals that state dependence accounts for 23.5% of output variation, with prediction errors below 10^{-6}. The negative correlation (-0.788) between Y_t and Λ_t implies counter-cyclical fiscal effectiveness. The GHH specification confirms the mechanism operates through state-dependent MPCs and the wealth effect.

Table F.2: State-dependent fiscal spending multipliers

	Dependent variable: Y_t (\$)			
	Hetero. (HH)		Rep. (RH)	GHH
G_t (\$)	0.402	0.182	0.206	0.000
	(0.005)	(0.002)	(0.000)	(0.001)
G_t (\$) $\times \Lambda_t$		0.533	0.534	0.000
		(0.003)	(0.002)	(0.000)
$\log(K_t)$	0.143	0.105	0.104	0.496
	(0.002)	(0.001)	(0.000)	(0.000)
$\log(A_t)$	0.463	0.598	0.591	1.504
	(0.001)	(0.001)	(0.001)	(0.000)
Constant	Yes	Yes	Yes	Yes
Observations	3,000	3,000	3,000	3,000
R^2	0.992	0.999	0.999	0.999
Adjusted R^2	0.992	0.999	0.999	0.999

Notes: Regression results based on specification (20). Columns: HH without interaction, HH with interaction, RH, and RH with GHH utility.

F.1.2 The GTFs: an endogenous disaster

The model's nonlinear nature implies that identical shock paths produce different outcomes depending on the hand-to-mouth share. Comparing states with high (43.7%) versus low (34%) hand-to-mouth shares at identical exogenous conditions, the high hand-to-mouth state amplifies output declines by 28% and consumption declines by 25% following a negative two-standard-deviation TFP shock. The amplification operates through excess consumption sensitivity from high cross-sectional MPCs, resonating with Petrosky-Nadeau et al. (2018)'s concept of *endogenous disasters*. Full GIRF figures with 95% confidence intervals are in the supplementary material.

F.2 Application II: A heterogeneous-household model with portfolio choice and endogenous labor supply

The second application features a heterogeneous-household model with portfolio choice between risky capital and risk-free bonds. A unit continuum of households solve:

$$V(a,b,z;X) = \max_{c,n,a',b'} \frac{c^{1-\sigma}}{1-\sigma} - \frac{\eta}{1+\frac{1}{\chi}} n^{1+\frac{1}{\chi}} + \beta \mathbb{E} V(a',b',z';X') \tag{21}$$

$$\text{s.t. } c + a' + q^b(X)b' = a(1+r(X)) + b + zw(X)n \tag{22}$$

$$a' \geq 0, \quad b' \geq \underline{b} \tag{23}$$

$$\Phi' = \Gamma_X(X), \quad S' \sim \pi(S'|S), \quad z' \sim \pi(z'|z) \tag{24}$$

where a is the risky asset earning capital rent $r(X)$; b is risk-free bond holding with price $q^b(X)$; σ is risk aversion; other notation follows Section F.1. The aggregate state is $X := \{\Phi, A\}$, where A follows the two-state Markov chain of Krusell and Smith (1997). The bond market clears at zero net supply: $\int b'(a,b,z;X)d\Phi(X) = 0$. Parameters are in Appendix H.

The model combines two inter-temporal assets, two occasionally binding constraints, and two non-trivial market clearing conditions (labor and bonds). The zero net bond supply creates a unique computational challenge: the clearing condition $q^b(X) \times 0 = 0$ is non-invertible. I resolve this by introducing a dummy bond term $\overline{B} > 0$ to enable relative price updates:

$$q_t^{b*} = \frac{Y_t^{(n)} - C_t^* - I_t^* + B_t^* + q_t^{b(n)}\overline{B}}{\overline{B}}, \tag{25}$$

where asterisks denote aggregations given the nth iteration's guessed prices.[15] Through iteration, $\{q_t^{b*}\}$ converges to market-clearing levels and $B_t^* \to 0$. Computational details are in Appendix B.

F.2.1 Nonlinear bond price and heterogeneous portfolio adjustment over the business cycle

The bond price dynamics are particularly nonlinear: a log-linear prediction achieves only $R^2 \approx 0.50$, yet the sufficient statistic K maintains strict monotonicity (Spearman coefficient indistinguishable from unity). The heterogeneous agent model predicts pro-cyclical bond prices ($\text{corr}(Y_t, q_t^b) = 0.64$), while the representative agent predicts counter-cyclical prices ($\text{corr}(Y_t, q_t^b) = -0.43$), with volatility more than 9× greater—reflecting heterogeneous hedging motives absent in representative agent models.

The global solution reveals heterogeneous portfolio adjustment over the business cycle.

[15] Appendix G provides details. Conservative updating (weight = 0.999) is needed due to high sensitivity.

Wealth-rich and middle-class households maintain stable portfolio weights (middle class at ∼110%, leveraging in debt markets), while wealth-poor households dramatically rebalance counter-cyclically. Richer households hold larger risky asset positions (consistent with Calvet et al. 2009), while wealth-poor households finance theirs through aggressive leverage near the borrowing limit. Full time-series figures are in the supplementary material.

F.2.2 State-dependent GTFs of the risk premium

Figure F.1 shows state-dependent impulse responses of the risk premium to ±2% TFP shocks, differentiated by pre-shock portfolio composition. Following a negative shock (panel (a)), economies with high risky asset exposure see the risk premium spike as constrained households sell off risky assets, while low-exposure economies experience a decline (flight to safety). The pattern reverses for positive shocks (panel (b)). Responses are highly asymmetric and persistent, driven by wealth-poor households whose portfolio composition determines risk premium state dependence (see supplementary material for detailed portfolio dynamics).

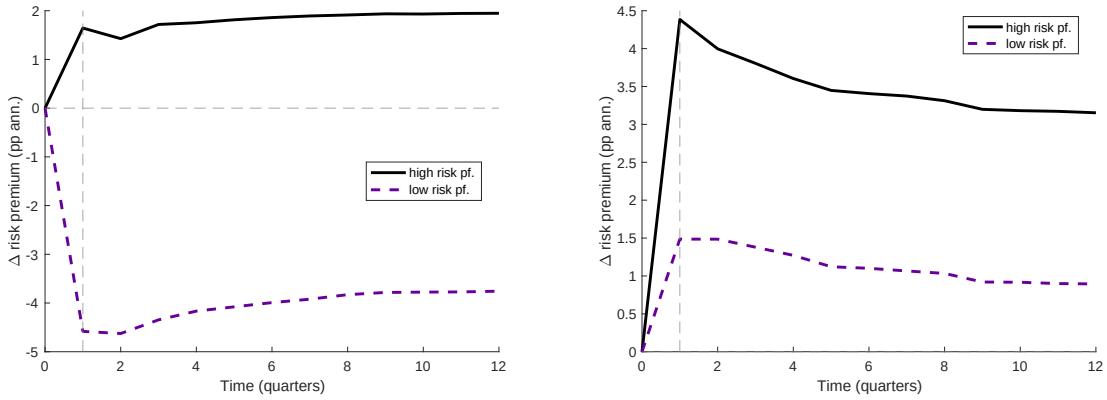

(a) Negative TFP shock (b) Positive TFP shock

Figure F.1: State-dependent risk premium dynamics

Notes: GTFs of the risk premium for ±2% TFP shocks, shown as deviation from the pre-shock level (annualized percentage points). Each line is the mean path across 1,000 stochastic forward simulations from a single conditioning period, differentiated by pre-shock risky asset share. Solid: highest pre-shock risky asset share. Dashed: lowest pre-shock risky asset share.

G Bond market clearing via fictitious supply

The net zero bond supply condition in the extended Krusell and Smith (1997) model leads to a non-invertible identity:

$$q^b(X)B' = B \iff q^b(X) \times 0 = 0. \tag{26}$$

Therefore, it is demanding to specify the implied price q_t^{b*} in the bond market. To overcome this issue, I introduce a dummy bond $\overline{B} > 0$ such that

$$q_t^{b*}\overline{B} := Y_t^{(n)} - C_t^* - I_t^* + B_t^* + q_t^{b(n)}\overline{B} \implies q_t^{b*} = \frac{Y_t^{(n)} - C_t^* - I_t^* + B_t^* + q_t^{b(n)}\overline{B}}{\overline{B}}. \quad (27)$$

In equilibrium, the national accounting identity must hold. During the iterative solution process, any discrepancy between the implied output Y_t^* and the n^{th}-guess $Y_t^{(n)}$ provides valuable information for updating the bond price $q_t^{b(n)}$ to satisfy this identity given the n^{th}-guess $q_t^{b(n)}$. [16] This price adjustment mechanism requires a dummy bond as an anchor point for relative price updates. Without this reference asset (i.e., if the dummy bond value were zero), there would be no basis for calibrating the bond price adjustments needed to satisfy market clearing conditions.

This approach can be generally used to pin down the implied price level required for non-trivial market clearing conditions under net zero supply. By contrast, state-space-based approaches directly solve for the market clearing price (q_t^b) at the exact point where supply and demand curves intersect at zero, eliminating the need for this dummy bond.

H Parameter tables

Tables H.1 and H.2 report parameter values for the two additional applications in Appendix F. Application I uses annual frequency; Application II uses quarterly frequency. The irreversibility parameter in Application I follows Guerrieri and Iacoviello (2015); borrowing limits and TFP processes in Application II follow Krusell and Smith (1997).

References

Andreasen, M. M., J. Fernández-Villaverde, and J. F. Rubio-Ramírez (2017, 06). The pruned state-space system for non-linear dsge models: Theory and empirical applications. *The Review of Economic Studies 85*(1), 1–49.

Calvet, L. E., J. Y. Campbell, and P. Sodini (2009, 02). Fight or flight? portfolio rebalancing by individual investors*. *The Quarterly Journal of Economics 124*(1), 301–348.

Cao, D. (2020). Recursive equilibrium in krusell and smith (1998). *Journal of Economic Theory 186*, 104978.

Cao, D., W. Luo, and G. Nie (2023, January). Global DSGE Models. *Review of Economic Dynamics 51*, 199–225.

[16] The information gain does not necessarily have to come from the output. In an unreported result, I found the difference in consumption also provides identifying variation for the implied bond price.

Parameter	Description	Value
α	capital share	0.330
β	discount factor	0.960
δ	depreciation	0.100
ϕ	borrowing constraint parameter	0.975
χ	Frisch elasticity	1.000
η	labor disutility	10.850
ρ	idiosyncratic productivity persistence	0.800
σ	idiosyncratic productivity volatility	0.014
$\overline{G}$	steady-state government demand	$0.200 \times Y^{ss}$
ρ_A	aggregate TFP persistence	0.900
σ_A	aggregate TFP shock volatility	0.013
ρ_G	government demand persistence	0.800
σ_G	government demand shock volatility	$0.01 \times Y^{ss}$

Table H.1: Parameters for Application I (KS1998 + fiscal shocks, annual)

Parameter	Description	Value
α	capital share	0.360
β	discount factor	0.990
δ	depreciation	0.025
σ	risk aversion (CRRA)	1.000
χ	Frisch elasticity	1.000
η	labor disutility	8.000
ρ	idiosyncratic productivity persistence	0.900
σ_z	idiosyncratic productivity volatility	0.050
$\underline{b}$	borrowing limit for the risk-free asset	-2.400
$\overline{B}$	dummy bond term (only for computation)	20.000
ρ_A	aggregate TFP persistence	0.859
σ_A	aggregate TFP volatility	0.014

Table H.2: Parameters for Application II (KS1997 + portfolio choice, quarterly)

Christiano, L. J. and J. D. Fisher (2000). Algorithms for solving dynamic models with occasionally binding constraints. *Journal of Economic Dynamics and Control 24*(8), 1179–1232.

Guerrieri, L. and M. Iacoviello (2015, March). OccBin: A toolkit for solving dynamic models with occasionally binding constraints easily. *Journal of Monetary Economics 70*, 22–38.

Hansen, L., J. Heaton, and N. Li (2008). Consumption strikes back? measuring long-run risk. *Journal of Political Economy 116*(2), 260–302.

Judd, K. L. (1992, December). Projection methods for solving aggregate growth models. *Journal of Economic Theory 58*(2), 410–452.

Justiniano, A. and G. E. Primiceri (2008, June). The time-varying volatility of macroeconomic fluctuations. *American Economic Review 98*(3), 604–41.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Khan, A. and J. K. Thomas (2008, March). Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics. *Econometrica 76*(2), 395–436.

Koop, G., M. Pesaran, and S. M. Potter (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics 74*(1), 119–147.

Krusell, P. and A. A. Smith, Jr. (1997, June). Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns. *Macroeconomic Dynamics 1*(2), 387–422.

Krusell, P. and A. A. Smith, Jr. (1998, October). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy 106*(5), 867–896.

McGrattan, E. R. (1996). Solving the stochastic growth model with a finite element method. *Journal of Economic Dynamics and Control 20*(1), 19–42.

Petrosky-Nadeau, N., L. Zhang, and L.-A. Kuehn (2018, August). Endogenous disasters. *American Economic Review 108*(8), 2212–45.

Ramey, V. A. (2020). The macroeconomic consequences of infrastructure investment. NBER Working Paper.

Young, E. R. (2010, January). Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control 34*(1), 36–41.