A Global Dynamic Nonlinear Solution Framework and the Repeated Transition Method*

Hanbaek Lee†

University of Cambridge

April 20, 2026

Abstract

This paper develops the repeated transition method (RTM), a global nonlinear solution framework for dynamic stochastic general equilibrium models. The method computes state-contingent expectations by identifying periods with matching aggregate states on a simulated equilibrium path, without assuming perfect foresight or parameterizing laws of motion. I establish conditions under which a scalar sufficient statistic enables accurate solutions with high-dimensional aggregate states. As a leading application, I globally solve a two-asset heterogeneous-agent New Keynesian model featuring two-dimensional aggregate uncertainty (TFP, monetary policy), portfolio choice, occasionally binding constraints, non-trivial market clearing, frictional capital adjustment, and endogenous labor supply. The global solution reveals that the interplay of binding constraints and portfolio choice generates sizable, state-dependent risk premium dynamics, and that the asset distribution qualitatively alters monetary policy transmission.

Keywords: Global nonlinear solution, repeated transition method, heterogeneous agents, HANK, risk premium, occasionally binding constraints.

JEL codes: C63, E32, E21, E44

*I am extremely grateful to Jesús Fernández-Villaverde, Dirk Krueger, Andrew Abel, and Frank Schorfheide for their invaluable guidance and support. I also thank Adrien Auclert, Andrea Ajello, Fernando Alvarez, David Argente, Kosuke Aoki, Isaac Baley, Gadi Barlevy, Florin Bilbiie, Timo Boppart, Christopher Carroll, Vasco Carvalho, Martin Eichenbaum, Martin Ellison, David Evans, Wouter Den Haan, Xiang Fang, Joel Flynn, James Graham, Sebastian Graves, Matthias Hansel, Zhen Huo, In Hwan Jo, Miles Kimball, Robert Kirkby, Sagiri Kitao, Dongya Koh, Felix Kubler, Eunseong Ma, Albert Marcet, Yusuf Mercan, Giuseppe Moscarini, Makoto Nirei, Alessandro Peri, Bruce Preston, José-Víctor Ríos-Rull, Andreas Schaab, Alp Simsek, John Stachurski, Kjetil Storesletten, Ludwig Straub, Stephen Terry, Christian Wolf, Yucheng Yang, Donghai Zhang, and seminar participants at the KEA, SCE, RES Symposium, Keio University, University of Colorado Boulder, Midwest Macro, University of Tokyo, Osaka University ISER, University of Sydney, Monash University, University of Melbourne, Queen Mary University of London, University of Cambridge, Tinbergen Institute, University of Zurich, Korea University, CEA, HKUST/Jinan Macro Workshop, NUS/Dynare Conference, Yale Cowles Foundation Macroeconomics Conference, Alicante Macro Workshop, SITE 2025, Bank of England, and ANU $Tv = v$ reading group for their insightful comments and discussions. All errors are my own. This research was supported by the University of Cambridge, Faculty of Economics, Keynes Fund.

†Faculty of Economics, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom. Email: hl610@cam.ac.uk

1 Introduction

Modern macroeconomic analysis faces a critical challenge: real-world economies exhibit rich heterogeneity and complex nonlinearities that our analytical tools struggle to capture fully. This methodological gap has limited our capacity to answer fundamental questions in economics including how growth interacts with uncertainty, what drives recessions, and when a policy can effectively intervene.

This paper contributes to closing this gap by developing the *repeated transition method* (RTM), a global nonlinear solution framework. The core idea is intuitive: simulate a long stochastic path of the economy, then construct each period's continuation value by finding other periods on the same path where the economy reached a similar endogenous state but experienced a different exogenous shock. By iterating between backward optimization (given the matched continuation values) and forward simulation (given the updated policies), the predicted and realized aggregate paths converge to a self-consistent equilibrium. This approach delivers four key advantages: (i) it accurately computes conditional expectations under aggregate uncertainty, without assuming perfect foresight; (ii) it avoids specification error in models with nonlinear dynamics, without parameterizing laws of motion; (iii) it efficiently handles non-trivial market clearing conditions by tracking implied prices along the sequence, without solving period-by-period fixed-point problems; (iv) it accurately solves models with occasionally binding constraints, without relying on piecewise linear approximations.

The key contributions of this paper are twofold. First, it develops the RTM algorithm and establishes its theoretical foundations. The method requires four properties of the model economy: stability, uniqueness, and Harris recurrence of the recursive competitive equilibrium, along with a finite exogenous state space. I provide a formal proposition showing that when value functions are strictly monotone in a scalar summary of the aggregate state, matching on that scalar recovers nearly identical continuation objects to those from direct distributional matching along the ergodic path. This monotonicity condition is empirically testable through a distribution-weighted Spearman rank correlation. This result enables the RTM to solve models with rich heterogeneity—including joint distributions over multiple individual state variables—using a single aggregate variable for matching. I also establish consistency: under contraction of the Bellman operator, the approximation error vanishes at rate $O(N_S \log T/T)$ in the simulation length T and the number of exogenous states N_S. Because the converged path encodes the full stochastic equilibrium, any sub-path defines a *generalized transition function* (GTF) that nests generalized impulse response functions and stochastic growth paths. This yields state-contingent policy analysis at no additional computational cost. The

leading application exploits GTFs to characterize nonlinear aggregate dynamics and monetary policy transmission.

Second, I globally solve a two-asset heterogeneous-agent New Keynesian (HANK) model featuring two-dimensional aggregate uncertainty (TFP, monetary policy), endogenous portfolio choice between capital and bonds, multiple occasionally binding constraints, and non-trivial market clearing conditions in both the labor and bond markets. The global nonlinear solution reveals qualitatively different economic conclusions from linearized methods. Several aggregate variables involved in non-trivial market clearing—the bond price, wage, and labor—exhibit meaningful residual variation under linear projections on aggregate capital (R^2 of 0.929, 0.934, and 0.307, respectively), with negligible improvement from adding a quadratic term. This non-approximability reflects a complex dependence on the distributional state that resists simple parametric approximation. The interplay of binding constraints and endogenous portfolio choice produces a sizable and state-dependent risk premium that averages approximately 4.1 percentage points annualized, fluctuating between 3.8 and 4.3 percentage points across TFP states. Using the GTF to condition on the aggregate constraint multiplier $\bar{\lambda}$ *within each TFP state*, I find that monetary policy transmission is strongly state-dependent: the risk premium responds sharply to a contractionary MP shock under tight constraints but is essentially unchanged under loose constraints, and output contracts more strongly under tight than under loose constraints, even after controlling for TFP. Bond price and consumption responses also differ across regimes in direction—an additional feature the global solution captures that linearization cannot. These results illustrate how the asset distribution—a key dimension of heterogeneity—can quantitatively alter aggregate dynamics and policy transmission, mechanisms that are inherently absent from first-order perturbation solutions.

The RTM provides a unified framework for globally solving a broad class of dynamic stochastic models, from simple representative-agent RBC economies to rich heterogeneous-agent models with multiple aggregate shocks and frictions. The supplementary material provides sample MATLAB codes for models with heterogeneous households and firms, occasionally binding constraints, frictional labor markets (DMP), sticky prices (New Keynesian), multiple aggregate shocks including uncertainty shocks, and multi-dimensional endogenous aggregate states.[1]

A battery of accuracy tests confirms the solution's reliability. Euler equation errors for the HANK model average $10^{-4.82}$ in consumption-equivalent units. The solution is robust across eight independent random seeds, with coefficient of variation (CV) below 0.35% for

[1]All sample codes are available at https://sites.google.com/view/hanbaeklee/computation-lab. Concurrently, Ferreira et al. (2025) applies the RTM to study business cycle implications of corporate cash holdings, Lee et al. (2024) to DMP models with endogenous separation, and Lee and Nomura (2024) to nonlinear inflation dynamics at the zero lower bound.

all equilibrium aggregates, and across simulation lengths with CV at most 0.10%. External benchmarks against Den Haan (2010) for the Krusell and Smith (1998) model and against GDSGE/OccBin for an RBC model with irreversible investment confirm accuracy relative to established methods. An off-equilibrium forward simulation test further verifies that the converged policy functions generalize beyond the equilibrium path, extending the Den Haan (2010) protocol to the HANK model.

Despite this accuracy and versatility, the method has limitations: it is less reliable for states not visited along the simulated path and is slower than linearization. It can also fail to converge when multiple slow-moving endogenous aggregate states interact—for example, Calvo pricing with interest rate smoothing, where both price dispersion and the lagged interest rate are slow-moving (see Section 2.7). These limitations are discussed in Sections 2.5, 2.7, and 5.

Related literature The RTM builds on the simulation-based approach of Krusell and Smith (1997, 1998), who use parameterized laws of motion to approximate aggregate dynamics. While powerful for models with approximately linear aggregate dynamics, their approach faces challenges with nonlinear dynamics due to potential misspecification of the law of motion. The RTM avoids this by computing conditional expectations directly from realized equilibrium outcomes, without functional approximation.[2]

The method is most closely related to Reiter (2010), who also computes conditional expectations without a full law of motion. Reiter (2010) defines value functions over a pre-specified grid of aggregate moments and constructs a proxy distribution at each grid point to recover equilibrium allocations. This requires the researcher to select which moments to track and how to map from moments back to full distributions—choices that grow increasingly burdensome as heterogeneity increases. The RTM eliminates the proxy-distribution construction step by using the converged equilibrium path itself as an endogenous computational grid, with non-parametric interpolation between matched periods. This grid automatically concentrates on the ergodic set. When a sufficient statistic is used, it serves only to index matches on the simulated path, rather than to parameterize the aggregate state space; Proposition 1 provides a formal monotonicity criterion for when the matching problem can be reduced to one dimension, complementing moment-based dimension-reduction approaches. For models with approximately linear aggregate dynamics, Reiter's approach and the RTM achieve comparable accuracy (Online Appendix D). The non-parametric structure of the RTM is

[2]The method of parameterized expectations (Marcet, 1988; Den Haan and Marcet, 1990) and the explicit aggregation approach of Den Haan (2010) address nonlinearity through richer basis functions but remain anchored to functional approximation.

especially natural for settings with occasionally binding constraints, endogenous portfolio choice, and non-trivial market clearing, as illustrated by the HANK application. For the same reason, the RTM is distinguished from the broader class of moment-based state-space methods (Den Haan, 1996, 1997; Reiter, 2001; Algan et al., 2008, 2010) and global functional approximation approaches (Cao et al., 2023; Elenev et al., 2021; Brumm et al., 2022). In particular, adaptive sparse grid methods enable global solutions in high-dimensional state spaces (Schaab and Zhang, 2022).

The RTM differs fundamentally from the sequence-space Jacobian (SSJ) approach of Auclert et al. (2021), which achieves remarkable speed and enables rapid likelihood-based estimation by linearizing around perfect-foresight transition paths. The RTM instead handles aggregate uncertainty directly and captures global nonlinearities, including risk premia and occasionally binding constraints, at the cost of longer computation times. The two approaches are complementary: SSJ excels at estimation and linear dynamics, while the RTM is designed for global nonlinear analysis. The RTM is also distinguished from other perfect-foresight methods (Fair and Taylor, 1983; Juillard, 1996; Boppart et al., 2018) and from perturbation-based approaches, including higher-order methods such as Taylor projection (Reiter, 2009; Ahn et al., 2018; Winberry, 2018; Bayer and Luetticke, 2020; Levintal, 2018), which extend the reach of local approximations but remain anchored to a reference state in the state space.

The RTM shares computational philosophy with Judd et al. (2011) and Maliar et al. (2011), who concentrate computation on the realized ergodic set. These methods still employ parametric approximation—polynomial projection—on the simulated points. The RTM eliminates this step: the simulated equilibrium path serves simultaneously as the computational grid and the source of continuation values, with non-parametric interpolation between matched periods replacing functional approximation. This unified structure—where the equilibrium path is both the object being solved for and the platform for constructing expectations—is what enables the RTM to handle complex models without specifying functional forms for either laws of motion or value functions.

Deep learning methods offer another approach to high-dimensional nonlinear problems, approximating value or policy functions via neural networks trained to minimize equilibrium residuals (Duarte, 2018; Azinovic et al., 2022; Maliar et al., 2021; Fernández-Villaverde, 2025). These methods achieve scalability in high-dimensional state spaces but require careful network architecture design, hyperparameter tuning, and convergence verification. The RTM provides a complementary simulation-based alternative that requires no neural network architecture design or extensive hyperparameter optimization; in turn, deep learning could enhance the RTM's interpolation step when the sufficient statistic is multi-dimensional, replacing piecewise linear interpolation with learned mappings from aggregate states to continuation values.

The HANK application connects to the large and growing literature on heterogeneous-agent models with nominal rigidities (Kaplan et al., 2018; Luetticke, 2021; Bayer and Luetticke, 2020; Auclert et al., 2024). Most existing HANK solutions rely on linearization. The RTM enables fully nonlinear analysis, revealing features such as the endogenous risk premium and state-dependent constraint dynamics that are absent in linearized solutions. The finding of a nonlinear, state-dependent risk premium further relates to the heterogeneous-agent asset pricing literature (Krusell and Smith, 1997; Heaton and Lucas, 2000; Gomes and Michaelides, 2007; Bayer et al., 2019), where portfolio choice under incomplete markets generates risk premia through market incompleteness. The RTM captures this mechanism in a model with nominal rigidities and monetary policy, extending the analysis beyond real models.

Roadmap Section 2 explains the repeated transition method, including the algorithm, sufficient statistic theory, and treatment of non-trivial market clearing conditions. Section 3 presents the leading application to a two-asset HANK model. Section 4 evaluates accuracy and robustness. Section 5 concludes. Appendices provide proofs, implementation details, performance comparisons, extensions of the generalized transition function framework, and additional applications.

2 Repeated transition method

2.1 A generic model framework

This section introduces a generic model framework that encompasses a broad class of dynamic stochastic general equilibrium (DSGE) models. The framework accommodates both heterogeneous-agent and representative-agent specifications. For clarity of exposition, the presentation uses a parsimonious benchmark; the method extends to richer environments—including heterogeneous firms, multiple aggregate shocks, frictional labor markets, and sticky prices—as summarized at the end of this section, with sample codes in the supplementary material. I denote the individual state as x and the aggregate state as X. The individual state x consists of the endogenous individual state a and the exogenous individual state s (the idiosyncratic shocks). The aggregate state X consists of the endogenous aggregate state Φ and the exogenous aggregate state S (the aggregate shocks). The endogenous aggregate state Φ takes different forms depending on the model class. In heterogeneous-agent models, Φ includes the distribution of individual states x, and may additionally include aggregate variables such as a smoothed policy instrument. In representative-agent models, it captures

the relevant aggregate allocations and prices.

$$\text{(Individual state)} \qquad x = (a, s) \tag{1}$$

$$\text{(Aggregate state)} \qquad X = (\Phi, S) \tag{2}$$

The idiosyncratic and aggregate shock processes are assumed to follow Markov processes with transition matrices Π^s and Π^S, respectively. The value function is denoted as V. Following standard notation, variables with apostrophes indicate future period values. The objective function of an economic agent consists of the contemporaneous part $f(y, a', x; X)$ and the expected future value. The agent maximizes the objective function by choosing (y, a'), where y is a vector of control variables that affects only the contemporaneous period. Then, the recursive formulation of an agent's problem is as follows:[3]

$$V(x; X) = \max_{y,a'} \; f(y, a', x; X) + \mathbb{E}m(X, X')V(a', s'; X') \tag{3}$$

$$\text{s.t.} \qquad (y, x') \in \mathcal{B}(x; X, X', q), \quad \Phi' = F(X) \tag{4}$$

where $m(X, X')$ is the stochastic discount factor; $q(X, X')$ is a price bundle; $\mathcal{B}(x; X, X', q)$ is the budget constraint; $F(X)$ is the law of motion known to an individual agent.[4] For notational convenience, I combine the price bundle (m, q) into p. The following market clearing condition pins down the price p:[5]

$$\text{[Market clearing]}: \quad p(X, X') = \arg_{\widetilde{p}}\{Q^D(\widetilde{p}, X, X') - Q^S(\widetilde{p}, X, X') = 0\}, \tag{5}$$

where Q^D and Q^S are the functions of demand and supply, which are endogenously determined by the model. The market clearing condition can be replaced by frictional price determination. For example, the method applies seamlessly to frictional labor markets with bilaterally determined wages, as in Diamond-Mortensen-Pissarides (DMP) models; applications are available in the supplementary material. For expositional clarity, I consider a simple case where the exogenous aggregate state S can take two possible values $\{G, B\}$ with a 2×2 transition matrix Π^S.[6] Based on this setup, Definition 1 defines the recursive competitive

[3] By expressing the model directly in recursive form, I implicitly assume the standard regularity conditions—namely, Assumptions 4.1 and 4.2 in Stokey et al. (1989).

[4] The stochastic discount factor can be a constant, for example β, as in a canonical dynamic household's problems. In a dynamic firm problem, the stochastic discount factor needs to be included.

[5] Any period-specific fixed point problem can be considered in the method, such as the externality effect as a function of endogenous allocations or non-trivial market clearing conditions. For brevity, I only include the non-trivial market clearing condition.

[6] The method's applicability is not limited to a certain number of grid points for the aggregate shocks. Moreover, multiple aggregate shocks can be considered an exogenous state.

equilibrium (RCE).

Definition 1 (The recursive competitive equilibrium)**.**
$(V, g_y, g_a) : \mathcal{X}_x \times \mathcal{X} \to \mathbb{R}$, $F : \mathcal{X} \to \Omega$, *and* $(Q^S, Q^D, p) : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ *are the recursive competitive equilibrium if* (i) (V, g_y, g_a) *solves the agent's problem, where* V *is the value function,* g_y *is the policy function for control variables, and* g_a *is the policy function for the endogenous state;* (ii) Q^D, Q^S *represent the demands and supplies of inputs and outputs;* (iii) p *satisfies the market clearing condition; and* (iv) F *is consistent with the realized future distribution implied by* g_a*.*[7]

In the following sections, I explain the method based on the recursive form in Equation (3) for completeness of exposition. However, the method applies seamlessly when the value function is replaced by first-order derivatives or marginal value functions. In the supplementary material, I provide multiple applications where the expected marginal value function is computed instead of the expected value function.

The method operates on a simulated equilibrium path. Therefore, although the converged equilibrium allocations are fully describable in recursive form, I denote equilibrium objects in sequential notation, such as $\{V_t\}$, for clarity. Given a realized state $\{x_t, X_t\}$ for an individual (or representative) agent at period t, the sequential expression V_t and the recursive form $V(\cdot; X_t)$ are used interchangeably.

2.2 Assumptions

This section states the key conditions a model must satisfy for the method to apply.

Assumption 1 (Stability)**.** *The recursive competitive equilibrium is stable: small perturbations to the aggregate path decay over time, so that the backward–forward iteration converges to the equilibrium path.*

Assumption 2 (Uniqueness)**.** *The recursive competitive equilibrium is unique: given the exogenous shock sequence, there is a single equilibrium path of endogenous aggregate states and prices.*

Assumption 3 (Harris recurrence)**.** *The Markov process over endogenous aggregate states is Harris recurrent: regardless of the initial condition, the process visits any set of positive measure under the invariant distribution infinitely often, with probability one.*[8]

[7] $\mathcal{X}_x$, $\mathcal{X}$ are the σ-algebras generated from all possible individual and aggregate states. Ω is a set of all distributions Φ.

[8] Harris recurrence is implied by standard regularity conditions: 1) the Feller property, 2) a Lyapunov drift condition (boundedness), 3) irreducibility, and 4) minorization. These are typically satisfied by dynamic

Assumption 4 (Finite exogenous support)**.** *The exogenous aggregate state S takes values in a finite set $\{S_1, \ldots, S_{N_S}\}$ with a known transition matrix Π^S.*

Harris recurrence (Assumption 3) is essential: without it, certain equilibrium allocations may appear only once and never reoccur in the future, making it impossible to implement the method using repeated transitions. Combined with aperiodicity (i.e., the absence of deterministic cycles), Assumptions 1–4 guarantee the existence of a unique stationary distribution over the endogenous state variables—a distribution over distributions in heterogeneous-agent settings—from which ergodicity follows. Assumption 4 ensures that for each possible future shock realization, the matching step can identify candidate periods that experienced the same exogenous state; any model with a finite Markov chain for aggregate shocks (e.g., $S \in \{B, G\}$) trivially satisfies this assumption. In what follows, I focus exclusively on models that satisfy Assumptions 1–4.

To ensure a well-defined equilibrium, I assume there is no redundancy in the representation of the aggregate state X. Specifically, I require a one-to-one mapping between the economy's fundamental state and the aggregate state X, formalized as:

$$V(x; X) = V(x; X') \text{ for all } x \quad \Longleftrightarrow \quad X = X'. \tag{6}$$

This condition rules out redundant state variables that could artificially generate equilibrium multiplicity through superfluous expansions of the state space. Equivalently, X is the natural minimal state variable (Cao, 2020).

2.3 The methodology

The method's key innovation lies in computing conditional expectations directly from realized equilibrium allocations of previous iterations. In particular, the RTM exploits Harris recurrence of DSGE models' recursive competitive equilibria: a simulated path of a stationary aggregate shock process revisits any set of positive measure under the invariant distribution infinitely often with probability one, populating the ergodic set of aggregate states as T grows. This enables state-contingent future allocations to be approximated using realized equilibrium outcomes on the simulated path. The interpolation error is quantified in Proposition 2 under Lipschitz continuity of the conditional-expectation map. Full consistency (Proposition 3) additionally requires the mesh width h_T (the maximum gap between adjacent matched

macroeconomic models with smooth objective functions and bounded returns—see Assumption 4.2 in Stokey et al. (1989). For formal statements and proofs, see Chapter 4 of Stokey et al. (1989) and Chapter 9 of Meyn and Tweedie (1993).

sufficient statistics) to shrink to zero, assumed explicitly under coverage conditions given in Section 2.5 rather than implied by Harris recurrence alone. Stability and uniqueness of the RCE then ensure that the conditional expectation function converges to its true counterpart as iterations proceed. This approach eliminates the need to specify parametric laws of motion—the method requires only a metric to assess similarity between aggregate states across periods.

To illustrate, consider an agent's infinite-horizon problem with two aggregate states: G (Good) and B (Bad). To solve the agent's problem at period t, a researcher needs to construct the expected value (marginal value) function of period $t+1$. For each possible future state $S_{t+1} \in \{G, B\}$, the RTM identifies a period on the previous iteration's allocation path where the endogenous state most closely matches that of period $t+1$. The expected future value function is then constructed by combining the time-specific value functions from these identified periods. Finite-sample approximation error in this step is controlled by Proposition 2, with the mesh-width condition $h_T \to 0$ stated explicitly in Proposition 3. Consequently, the expected future value (marginal value) function at each period can be accurately constructed by combining these realized outcomes from the simulation path.[9]

The method iteratively updates the predicted allocation path using the realized path from the previous iteration until convergence.[10] Each iteration transmits information from the entire realized allocation sequence to the next. By avoiding parametric specifications of transition dynamics, the RTM can capture equilibrium dynamics in highly nonlinear models without functional-form error.

The basic structure of the methodology is as follows: Suppose T periods of aggregate exogenous states $\{S_t\}_{t=0}^{T}$ are simulated, and the path is long enough that the realized states form a dense grid on the ergodic support.[11] The solution process starts by conjecturing three time series: 1) value functions, $\{V_t^{(0)}\}_{t=0}^{T}$, 2) endogenous states $\{\Phi_t^{(0)}\}_{t=0}^{T}$, and 3) prices $\{p_t^{(0)}\}_{t=0}^{T}$. Using these guesses, I solve the allocations backward from the terminal period T to obtain the implied value function solution $\{V_t^*\}_{t=0}^{T}$, and simulate the economy forward using the solution. The forward simulation generates the endogenous states $\{\Phi_t^*\}_{t=0}^{T}$ and implied prices $\{p_t^*\}_{t=0}^{T}$ from the market clearing conditions. Here p_t^* is the price implied by the market clearing condition, not the market clearing price itself. This distinction is discussed

[9]The method's name—repeated transition method—reflects its key feature of utilizing *repeated transitions* between the same (similar) endogenous states with different exogenous states.

[10]The terminology "predicted" means predicted from the perspective of a researcher outside the model. It is equivalent to the nth guess for the allocation paths.

[11]In theory, an infinitely long simulation needs to be considered, but for illustrative purposes, I consider a T-period long simulation. Later in the application, a long-enough finite simulation is used as an approximation for the infinitely long ergodic path. The RTM's reliance on a single, sufficiently long simulated path of aggregate shocks connects to recent work by Kahou et al. (2021).

in detail in Section 2.6. The guess is then updated through convex combinations of prior guesses and realized allocations to form $\{V_t^{(1)}, \Phi_t^{(1)}, p_t^{(1)}\}_{t=0}^T$. While this broad approach shares similarities with perfect-foresight methods (Fair and Taylor, 1983), it differs fundamentally in the backward solution step due to its treatment of conditional expectations.

To clarify this point, consider period t in iteration $n+1$, after solving backward from T to $t+1$. Suppose the exogenous state at period $t+1$ is G ($S_{t+1} = G$). To solve an agent's problem at t, one needs to construct an expected future value function $\mathbb{E}_t \widetilde{V}_{t+1}$.[12] This presents a challenge: while $V_{t+1}^*(\cdot, S = G)$ is available from the backward solution, $V_{t+1}^*(\cdot, S = B)$ is not, as only one exogenous state realizes in each period. I define this unobserved $V_{t+1}^*(\cdot, S = B)$ as a *counterfactual* conditional value function.

The standard state space-based approach addresses this challenge by replacing time indices with endogenous and exogenous aggregate states, interpolating endogenous states through an assumed law of motion. The solution's accuracy thus critically depends on correctly specifying this law of motion. However, verifying the specification's accuracy is impossible before solving the equilibrium. An incorrect specification requires restarting the solution process with a new guess, presenting two fundamental challenges: determining which statistics to include and selecting appropriate functional forms. This problem cannot be easily resolved unless the aggregate dynamics are known to be log-linear, as in Krusell and Smith (1998).

The RTM takes a fundamentally different approach. Instead of specifying a law of motion, it obtains the counterfactual conditional value function from another period $\widetilde{t}+1$ where the endogenous aggregate state matches that of period $t+1$ but the counterfactual exogenous state is realized:

$$\Phi_{\widetilde{t}+1}^{(n)} = \Phi_{t+1}^{(n)} \tag{7}$$

$$S_{\widetilde{t}+1} = B \neq G = S_{t+1}. \tag{8}$$

Under these conditions, all aggregate states in period $\widetilde{t}+1$ match those of the *counterfactual state* in period $t+1$, implying

$$V_{\widetilde{t}+1}^{(n)}(\cdot, S = B) = V_{t+1}^{(n)}(\cdot, S = B). \tag{9}$$

Importantly, $V_{\widetilde{t}+1}^{(n)}(\cdot, S = B)$ is the observed *factual* conditional value function available in the nth iteration. With both $V_{t+1}^*(\cdot, S = G)$ and $V_{t+1}^{(n)}(\cdot, S = B)(= V_{\widetilde{t}+1}^{(n)}(\cdot, S = B))$ available from iteration n, the expected future value function $\mathbb{E}_t \widetilde{V}_{t+1}$ can be consistently computed.

[12] The method can potentially accommodate various expectation formations beyond rational expectations. The conditional expectation computation step can be adjusted to any well-defined expectation structure.

This approach extends naturally to finer discretizations of the aggregate shock process.[13] As T grows, the realized iso-shock candidates provide increasingly fine coverage of the ergodic support; for finite T, the matching step uses bracketing and interpolation between the two nearest candidates, with approximation error controlled by the mesh width h_T (Proposition 2). Figure 1 illustrates the step for computing the conditional expectation in period t.

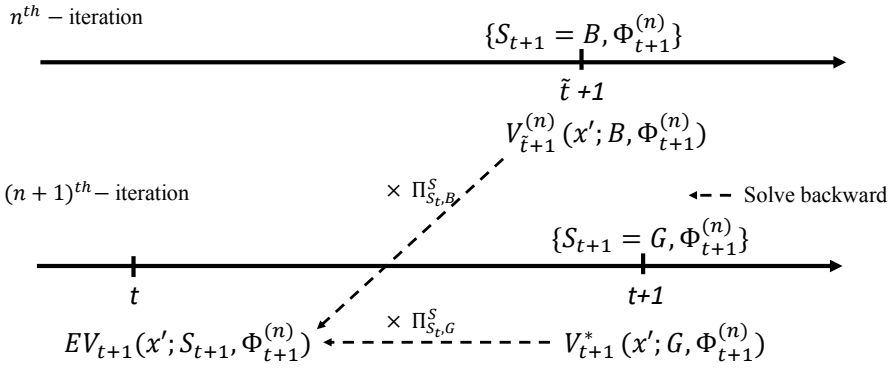

Figure 1: The computing step for conditional expectation based on the RTM

Notes: The conditional expectation in period t can be computed by $\mathbb{E}V_{t+1}(x'; S_{t+1}, \Phi_{t+1}^{(n)}) = \Pi_{S_t,B}^S \times V_{\tilde{t}+1}^{(n)}(x'; B, \Phi_{t+1}^{(n)}) + \Pi_{S_t,G}^S \times V_{t+1}^*(x'; G, \Phi_{t+1}^{(n)})$.

This approach eliminates the need to specify a law of motion for computing expected future value functions. Instead, the critical step becomes identifying period $\widetilde{t}+1$ that replicates the counterfactual conditions of period $t+1$. This identification relies on tracking the sequence of endogenous aggregate states $\{\Phi_t^{(n)}\}_{t=0}^T$, which serves as the key criterion for locating appropriate matching periods:

$$\widetilde{t}+1 = \arg\inf_{\tau} \|\Phi_{\tau}^{(n)} - \Phi_{t+1}^{(n)}\| \tag{10}$$

Addressing the curse of dimensionality The RTM addresses the curse of dimensionality in two respects. First, it replaces the infinite-dimensional law of motion with a scalar distance comparison: to locate matching periods, the algorithm requires only a distance metric over endogenous aggregate states that induces a total ordering. This reduction is inherent to the method—any implementation must rank periods by similarity—though it requires pairwise comparisons across the T simulated periods. This cost is independent of the sufficient statistic result developed in Section 2.5, which provides formal conditions under which a particular scalar choice is without loss. Second, the RTM indexes value and policy functions by time period rather than by aggregate state, so the dimensionality of the function domain does not

[13] Most applications in Online Appendix F employ finer grids than two grid points for the exogenous aggregate states.

grow with the aggregate state space. In contrast, state-space-based approaches require value and policy functions whose dimensions grow with the aggregate state.

Table 1: Repeated Transition Method (RTM) Algorithm

The RTM algorithm

Step 1 (Initialization) Simulate a long path of aggregate shocks $\{S_t\}_{t=0}^{T}$. Conjecture initial sequences of value functions $\{V_t^{(0)}\}_{t=0}^{T}$, aggregate states $\{\Phi_t^{(0)}\}_{t=0}^{T}$, and prices $\{p_t^{(0)}\}_{t=0}^{T}$.

Step 2 (Backward solution) Starting from the terminal period T, solve agents' problems backward using expectations based on $\{V_t^{(n)}\}_{t=0}^{T}$. For each exogenous state in $t+1$:

(a) Find periods τ such that $\Phi_\tau^{(n)} \approx \Phi_{t+1}^{(n)}$ and collect realized state-contingent value functions.

(b) Use these values to construct expectations $\mathbb{E}_t[V_{t+1}]$.

Then, obtain the optimal value function V_t^* and the policy function g_t^{a*}.

Step 3 (Forward simulation) Simulate the model forward using $\{g_t^{a*}\}_{t=0}^{T}$ to generate $\{\Phi_t^*\}_{t=0}^{T}$ and implied prices $\{p_t^*\}_{t=0}^{T}$.

Step 4 (Update) Update the guessed sequences via convex combination:

$$V_t^{(n+1)} = \xi V_t^{(n)} + (1-\xi)V_t^*, \quad \text{similarly for } \Phi_t,\ p_t.$$

Step 5 (Iteration) Repeat Steps 2–4 until convergence:

$$\sup_t \left\| p_t^{(n+1)} - p_t^{(n)} \right\| < \varepsilon.$$

In Appendix B of the online appendix, I elaborate on the detailed steps to implement the RTM and the required length of the simulated path. Because the RTM extracts information from the entire simulated path, it does not necessarily require longer simulations than traditional methods in low-dimensional shock processes. For example, as in Krusell and Smith (1998), if only two aggregate states are realized based on a symmetric transition probability of a moderate level (e.g., $\mathbb{P}(S' = G|S = B) = 0.125$), $T = 1{,}000$ simulated periods suffice for the solution to remain stable under further lengthening. Section 4.3 confirms this formally: Table 6 reports that equilibrium aggregates remain stable (CV at most 0.10%) across simulation lengths $T \in \{2{,}000, 3{,}000, 5{,}000\}$.

The RTM's outer loop is structurally analogous to Krusell and Smith (1998): both iterate until the aggregate objects used to form expectations are consistent with the outcomes implied by the resulting individual decisions. The key difference is that Krusell–Smith iterates on law-of-motion coefficients, while the RTM iterates on the entire sequence of aggregate variables, enabling nonlinear dynamics without functional-form assumptions.

2.4 The algorithm: An illustrative real business cycle example

To fix ideas, I first explain the RTM through a canonical real business cycle (RBC) model with occasionally binding investment irreversibility (McGrattan, 1996; Christiano and Fisher, 2000).[14] A representative household solves:

$$V(a;X) = \max_{c,a'} \frac{c^{1-\sigma}}{1-\sigma} + \beta \mathbb{E} V(a';X') \tag{11}$$

$$\text{s.t. } c + a' - (1-\delta)a = Aa^{\alpha}, \quad a' - (1-\delta)a \geq \phi I_{ss} \tag{12}$$

where a is capital, A is TFP following a Markov chain, and the irreversibility constraint $a' - (1-\delta)a \geq \phi I_{ss}$ prevents investment from falling below a fraction ϕ of its steady-state level.

The RTM proceeds as follows:

Step 0: Initialization Generate a long (T periods) simulated path of exogenous shocks $\{S_t\}_{t=1}^{T}$ from the Markov transition matrix. Initialize the aggregate capital path $\{K_t^{(0)}\}_{t=1}^{T}$ to the stationary-equilibrium level and the value function $V_t^{(0)}$ to the stationary-equilibrium value function, replicated at each t.

Step 1: Backward solution (from $t = T$ to $t = 1$). At each period t:

Step 1a: Matching. For each possible future shock $S' \in \{G, B\}$, find the period $\widetilde{t}(S')$ on the *previous iteration's* path where the endogenous state most closely matches $K_{t+1}^{(n)}$:

$$\widetilde{t}(S') = \arg \inf_{\tau : S_\tau = S'} \left| K_\tau^{(n)} - K_{t+1}^{(n)} \right| \tag{13}$$

This matching step is the key innovation of the RTM. It identifies periods where the economy was in a similar endogenous state but experienced a different exogenous shock, enabling construction of state-contingent expectations without parameterizing the law of motion. In practice, I find two bracketing candidates for each S' and interpolate between them using linear interpolation weights.[15][16]

[14] This model serves as a useful starting point because it features genuinely nonlinear aggregate dynamics while remaining simple enough for a clear exposition. The method applies identically to heterogeneous-agent models; the only difference is that the endogenous aggregate state becomes the distribution Φ rather than a scalar.

[15] For the realized successor shock S_{t+1}, the continuation value is already available from the next period on the path—not only saving computation but also providing an exact continuation value with no interpolation error. The matching step is therefore needed only for the $N_S - 1$ *unrealized* shock states. With two aggregate states, one of the two conditional value functions is exact and only one must be obtained by matching; with N_S states, $N_S - 1$ must be obtained by matching.

[16] When idiosyncratic and aggregate shocks co-move, as in Krusell and Smith (1998), the joint transition probability $\Gamma(s', S'|s, S)$ depends on both S and S', so the appropriate state representation is (Φ, S_{-1}, S) rather than (Φ, S). In this case, the matching step should condition on the *current* exogenous state S when

Step 1b: Expectations. Construct the expected future value (or marginal value) function by combining the value functions from the matched periods:

$$\mathbb{E}_t[V'(a')] = \sum_{S'} \pi(S'|S_t) \cdot V_{\tilde{t}(S')}(a') \tag{14}$$

where $\pi(S'|S_t)$ is the transition probability and $V_{\tilde{t}(S')}$ is the value function from the matched period in the previous iteration.

Step 1c: Individual optimization. Given the constructed expectations, solve the household's problem at period t to obtain updated policy functions $\{c_t, a'_t\}$.

Step 2: Forward simulation (from $t = 1$ to $t = T$). Using the updated policy functions, simulate the economy forward from the initial distribution to compute the implied capital path $\{K_t^*\}_{t=1}^T$.

Step 3: Convergence check and update Compute the discrepancy between the predicted and realized paths (here, the mean squared error in aggregate capital):

$$\text{Error} = \frac{1}{T} \sum_{t=1}^{T} \left(K_t^{(n)} - K_t^* \right)^2 \tag{15}$$

If the error falls below a pre-specified tolerance, the algorithm terminates. Otherwise, update the predicted path using a damped convex combination:

$$K_t^{(n+1)} = \xi K_t^{(n)} + (1 - \xi) K_t^* \tag{16}$$

where $\xi \in (0, 1)$ is a damping parameter (typically 0.5–0.9). Return to Step 1.

The converged solution satisfies $K_t^{(n)} = K_t^*$ for all t: the aggregate path that agents expect equals the aggregate path that their optimal decisions imply.

Extension to heterogeneous-agent models The algorithm above extends directly to the generic model framework in Section 2.1. The only modification is that the endogenous aggregate state Φ becomes the distribution of individual states, so the forward simulation step involves simulating the distribution rather than a single agent's decisions. For models with non-trivial market clearing conditions, the forward pass also computes implied prices from the market clearing identities, as described in Section 2.6.

searching for counterfactual candidates under S'—that is, search over $\{\tau : S_\tau = S' \text{ and } S_{\tau-1} = S_t\}$. This refinement requires longer simulations to identify suitable counterfactual candidates but does not alter the structure of the algorithm.

Computational cost Each RTM iteration has cost $O(T \times N_S \times n_x)$, where T is the simulation length, N_S is the number of exogenous states, and n_x is the cost of solving the individual problem at one period. The matching step has cost $O(N_S \log T)$ per period via binary search on pre-sorted lists.[17] The total cost is comparable to other simulation-based methods such as Krusell and Smith (1998), but with two important differences: (i) the RTM avoids the cost of specifying and fitting a law of motion, and (ii) for models with non-trivial market clearing, it avoids the additional $O(T \times n_{\text{iter}})$ cost of nested price-finding loops that state-space methods incur, where n_{iter} is the number of iterations needed to clear markets in each period. For the RBC model above, Online Appendix D compares the RTM's accuracy and computation time against linearization, the OccBin method of Guerrieri and Iacoviello (2015), and the GDSGE global solution of Cao et al. (2023). Section 4 provides further accuracy validation for the leading application.

2.5 Sufficient statistic

The matching step (13) requires comparing endogenous aggregate states across periods. In heterogeneous-agent models, the endogenous state is the distribution Φ, which is infinite-dimensional. Directly comparing distributions across periods is feasible—the matching-robustness exercises in Section 4.4 and Online Appendix C demonstrate this by comparing full-distribution matching in the portfolio model and enriched moment-based matching in the HANK model—but computationally expensive, requiring longer simulations to adequately cover the higher-dimensional matching space.

A more efficient approach uses a *sufficient statistic*: a scalar variable $e_t = e(\Phi_t)$ that captures the information in Φ_t relevant for the continuation objects used in the matching step. The following proposition, proved in Online Appendix A, establishes conditions under which matching on e_t recovers the same continuation objects as direct distributional matching along the ergodic path. Proposition 2 then quantifies the interpolation error when only approximate matches are available.

Proposition 1 (Sufficient statistic for matching)**.**
Let Assumptions 1–4 hold. Let $e_t = e(\Phi_t)$ be a measurable function of the endogenous aggregate state, and suppose that for each exogenous state S and every individual state (a, s) in the support of Φ_t,

$$e_t < e_{t'} \iff V_t(a, s) < V_{t'}(a, s) \quad \textit{for all } t, t' \in \mathcal{T}_S \equiv \{t \mid S_t = S\},$$

[17] Binary search repeatedly halves the candidate list, locating the two bracketing periods in at most $\lceil \log_2 T \rceil$ comparisons. The initial sorting is done once at cost $O(T \log T)$.

or the reverse inequality holds uniformly (strict monotonicity in e_t). Then within each $\mathcal{T}_S$, periods sharing the same (e_t, S_t) yield the same continuation objects: $V_t(a,s) = V_{t'}(a,s)$ 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 in Proposition 2.

Proof.
See Online Appendix A. ■

Proposition 1 is a statement about the equivalence of matching for continuation objects along the ergodic set. It does not require that Φ be uniquely determined by e; the requirement is only that periods with identical (e_t, S_t) along the ergodic path yield identical continuation values. In finite samples, the RTM uses bracketing and interpolation in e; Proposition 2 quantifies the resulting approximation error through the maximum gap h_T between adjacent matched sufficient statistics.

In practice, the natural candidate for the sufficient statistic is the first moment of the individual endogenous state. For models with a single asset (e.g., Krusell and Smith 1998), $e_t = K_t = \int a \, d\Phi_t$ is aggregate capital. For models with multiple assets (e.g., the two-asset HANK model in Section 3), I show empirically that aggregate capital $K_t = \int k \, d\Phi_t$ remains a valid sufficient statistic despite the model's joint distribution over (k, b, z).

The monotonicity condition is a strong pointwise requirement—it must hold at every individual state in the support simultaneously—but it is not an assumption imposed on the model. Rather, it is a verifiable property of the converged solution, checked *ex post* on the equilibrium path using the Spearman rank correlation diagnostic (Remark 1). The condition holds naturally in models where aggregate wealth is the dominant driver of continuation values, which encompasses most workhorse heterogeneous-agent models.

Pointwise monotonicity implies a testable necessary condition on the distribution-weighted marginal value, stated in the following remark. I report these tests for all applications, consistently finding Spearman correlations exceeding 0.999 in absolute value (Section 4.4), providing strong empirical support for the pointwise condition.[18]

Remark 1 (Empirical test for monotonicity)**.**
For each exogenous state partition $\mathcal{T}_S = \{t \mid S_t = S\}$, a necessary condition for the pointwise monotonicity in Proposition 1 is

$$\left|\text{Spearman}\left(\{e_t\}_{t\in\mathcal{T}_S},\ \{\textstyle\int V_t'(x)\, d\Phi_t\}_{t\in\mathcal{T}_S}\right)\right| = 1.$$

[18] A Spearman rank correlation near unity verifies the rank-preservation (order-equivalence) component of the monotonicity condition. The smoothness condition in Proposition 2 below is a separate, stronger requirement; it is assessed directly through the formal error bound estimates in Section 4.4.

Failure of this condition (correlation significantly below unity) rules out e_t as a sufficient statistic. Near-unity values are consistent with the pointwise condition, because violation of pointwise monotonicity while maintaining integral monotonicity would require offsetting non-monotonicities across individual states—a configuration not encountered in the applications considered; the Spearman diagnostic in Section 4.4 provides a model-specific test of this condition.

In finite simulations, exact matches are rare and the RTM instead interpolates between bracketing candidates. The following proposition bounds the resulting interpolation error.

Proposition 2 (Interpolation error bound)**.**
Let Assumptions 1–4 hold. Suppose:

(i) For each exogenous state $S \in \mathcal{S}$, define the conditional expectation

$$\mathcal{E}(\mathbf{x}; e, S) \equiv \mathbb{E}\big[V(\mathbf{x}'; \Phi', S') \,\big|\, e, S\big].$$

This function is Lipschitz continuous in e with constant L_S:

$$\big|\mathcal{E}(\mathbf{x}; e, S) - \mathcal{E}(\mathbf{x}; e', S)\big| \leq L_S \,|e - e'| \quad \text{for all } \mathbf{x},\ e, e' \in [\underline{e}, \bar{e}].$$

(ii) The converged simulation path $\{(e_t, S_t)\}_{t=1}^{T}$ visits each state S at times $\mathcal{T}_S \subset \{1, \ldots, T\}$, and the maximum gap between adjacent sorted sufficient statistics is

$$h_S \equiv \max_{i=1,\ldots,|\mathcal{T}_S|-1} \big(e_{(i+1)} - e_{(i)}\big),$$

where $e_{(1)} \leq e_{(2)} \leq \cdots \leq e_{(|\mathcal{T}_S|)}$ are the sorted values of $\{e_t : t \in \mathcal{T}_S\}$.

Then, for the RTM with linear interpolation between the two bracketing candidates, the pointwise interpolation error in the conditional expectation satisfies

$$\max_t \big|\mathcal{E}^{\mathrm{RTM}}(\mathbf{x}; e_t, S_t) - \mathcal{E}(\mathbf{x}; e_t, S_t)\big| \leq \frac{L}{2} \cdot h_T \quad \text{for all } \mathbf{x}, \tag{17}$$

where $L = \max_S L_S$ and $h_T = \max_S h_S$ is the overall maximum gap, which shrinks with the simulation length T.

Proof.
See Online Appendix A. ■

The Lipschitz condition is mild: it requires that conditional expectations vary smoothly with the sufficient statistic, which holds under standard regularity (bounded returns, continuous

technology). The practical implication is that the RTM's approximation error is controllable: it shrinks with the simulation length T at the one-dimensional rate, regardless of how rich the underlying heterogeneity is. Both components of the bound—the Lipschitz constant L and the maximum gap h_T—are computable from the converged solution, yielding an ex post error certificate rather than just an asymptotic guarantee. Online Appendix C reports empirical estimates of these error-bound components.

Convergence rate and the role of the sufficient statistic If the ergodic distribution of e conditional on S has a density bounded below by a positive number, then the average gap $\bar{h}_S$ between consecutive sorted values of e conditional on S satisfies $\bar{h}_S = O(N_S/T)$, where $N_S \equiv |\mathcal{S}|$ denotes the number of exogenous states. The maximum gap h_S used in Proposition 2 carries an additional logarithmic factor, $h_S = O(N_S \log T/T)$, reflecting the extreme-value behavior of the largest spacing.[19] This makes precise three features: (a) the interpolation error decreases at rate $1/T$ (up to a logarithmic factor), the standard rate for one-dimensional non-parametric interpolation; (b) without the sufficient statistic, interpolation on the d-dimensional ergodic set would yield errors of order $O(T^{-1/d})$—the sufficient statistic reduces the effective dimension to one; and (c) the bound scales linearly with N_S, so models with richer exogenous processes require proportionally longer simulations. The total approximation error thus decomposes into a deterministic interpolation component of order $L \cdot h_T$ (Proposition 2) and a stochastic coverage component that vanishes as $T \to \infty$.[20]

Proposition 3 (Consistency)**.**
Suppose the conditions of Propositions 1 and 2 hold, with bounded returns and compact state spaces. Suppose additionally that (i) the Bellman operator $\mathcal{T}$ is a contraction 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

[19] For n draws from a density bounded below on a compact interval, the maximum spacing is $O(\log n/n)$ almost surely (Devroye, 1981). Since $n \approx T/N_S$ within each exogenous state partition, $h_S = O(N_S \log(T/N_S)/T)$.

[20] When the economy traverses the endogenous state space slowly (slow mixing), a finite simulation covers only a narrow region of the ergodic set. From the perspective of backward discounting, the cost of limited coverage is mitigated: distant, uncovered states contribute little to current decisions. However, from the perspective of forming conditional expectations, limited coverage is costly: the matching step requires counterfactual candidates—periods with similar endogenous states under different exogenous shocks—and slow traversal may leave the current endogenous region unvisited under some shock realizations. This is why slow traversal ultimately requires longer simulations for accurate conditional expectations, even though the discounted value of distant states is small.

with a uniform constant L over $f \in \mathcal{F}$;[21] *and (iii) $h_T \to 0$ almost surely as $T \to \infty$.*[22] *If the RTM iterations converge to a fixed point V_T^{RTM} of the approximate operator $\widetilde{\mathcal{T}}_T$ (which uses interpolated continuation values from the simulation of length T), then*

$$\|V_T^{\mathrm{RTM}} - V^*\|_\infty \leq \frac{L \cdot h_T}{2(1-\beta)} \xrightarrow{a.s.} 0 \quad \text{as } T \to \infty, \tag{18}$$

where V^* is the true recursive competitive equilibrium value function.

Proof.
See Online Appendix A. ■

Proposition 3 establishes that the RTM is consistent: as the simulation length grows, the approximation error vanishes. The rate $h_T = O(N_S \log T/T)$ from the convergence rate discussion above yields $\|V_T^{\mathrm{RTM}} - V^*\|_\infty = O(N_S \log T/T)$. The contraction assumption is standard in dynamic programming but restrictive—it requires bounded returns and a compact state space. The proposition establishes accuracy conditional on convergence: it bounds the distance between the RTM fixed point and the true value function, but does not prove that the RTM iterations converge to that fixed point. Section 4 provides extensive empirical evidence that the method converges accurately in the applications considered.

Sufficiency in the RTM versus state-space context Sufficiency here is weaker than in a state-space approach: e_t only needs to index the value functions along the simulated path for the matching step, not summarize Φ for a recursive law of motion. Closeness of e_t does not imply closeness of Φ in any distributional metric. What ultimately matters for the RTM is not distributional similarity per se, but similarity of continuation values: the Lipschitz condition in Proposition 2 requires only that continuation values vary smoothly with e_t along the ergodic set. The matching-robustness checks in Section 4.4 verify this: full-distribution matching in the portfolio model and enriched moment-based matching in the HANK model produce nearly identical aggregate paths.

The monotonicity condition can fail when the endogenous aggregate state has multiple independent components that affect continuation values in distinct directions—for instance,

[21] Sufficient primitive conditions for (ii) include a Feller transition kernel together with a density of e conditional on S bounded below; $\mathcal{F}$ can then be taken as bounded continuous functions sharing a fixed uniform modulus, which both operators preserve. These are heuristic sufficient conditions rather than a proof from primitives.

[22] This condition follows under the compact-support and density-bounded-away-from-zero conditions discussed above, together with Harris recurrence of the simulated process (Assumption 3). Harris recurrence supplies repeated visits to positive-measure regions; the mesh-rate statement additionally uses the one-dimensional spacing result for the sufficient statistic.

if both mean wealth and wealth inequality independently shift the interest rate in opposing ways. In such cases, the matching step would require multiple sufficient statistics or direct matching on Φ_t, at the cost of longer simulation paths (the convergence rate discussion above). For the applications in this paper, no such divergence between e_t-based and Φ_t-based solutions is detected. Section 2.7 documents a concrete failure case where the monotonicity condition breaks down.

Matching with interpolation In practice, the matching step uses linear interpolation between the two closest candidates rather than exact matching. For each future shock S', I identify all periods $\{s : S_s = S'\}$, sort them by e_s, find the two periods $\widetilde{t}^{dn}$ and $\widetilde{t}^{up}$ bracketing e_{t+1} from below and above, and interpolate individual policies between them using the linear weight $\omega = (e_{t+1} - e_{\widetilde{t}^{dn}})/(e_{\widetilde{t}^{up}} - e_{\widetilde{t}^{dn}})$. This ensures smooth updating even when the ergodic set is discretely sampled. The matching is performed once per period based on the sufficient statistic e_t, and the resulting interpolation weights are applied identically to all aggregate variables (value functions, policies, prices). This is consistent with the sufficient statistic theory: if continuation values are monotone in e_t, then equilibrium prices that clear markets at those continuation values inherit the same ordering.

When the sufficient statistic satisfies the monotonicity condition in Proposition 1, the two bracketing candidates provide a monotone bracket for the continuation object: if $e_\ell < e_{t+1} < e_u$ with $S_\ell = S_u$, then $V_{e_\ell}(a, s)$ and $V_{e_u}(a, s)$ bound $V_{t+1}(a, s)$ from above and below (or vice versa) for each (a, s). Linear interpolation between these bounds is therefore stable—any convex combination lies within the monotone bracket, preventing overshoot—and the interpolation error is controlled by Proposition 2.

The matching step can be interpreted as non-parametric interpolation on the simulated endogenous state space. The converged equilibrium path provides a grid of sample points $\{(e_s, V_s)\}$, and the RTM constructs continuation values at arbitrary query points e_{t+1} via piecewise linear interpolation between the two nearest neighbors with the same exogenous state.[23] When a sufficient statistic reduces the matching to one dimension, this reduces to standard one-dimensional interpolation, avoiding the curse of dimensionality associated with high-dimensional nearest-neighbor methods.[24] Because the sufficient statistic reduces the effective matching dimension to one, a single simulation path provides adequate coverage of the endogenous state space at moderate T. The multi-seed robustness test in Section 4

[23] I am grateful to an anonymous referee for suggesting this characterization of the RTM as non-parametric interpolation on a simulation-generated grid, which clarifies the connection to standard non-parametric methods and motivates the error bound analysis.

[24] The connection to non-parametric interpolation on simulated grids relates to the literature on adaptive grids (Judd et al., 2011; Maliar et al., 2011), where computation is concentrated on the realized ergodic set rather than a pre-specified rectangular grid.

confirms that the converged solution does not depend on the particular shock draw.

2.6 Non-trivial market clearing conditions

Many economically interesting models feature market clearing conditions that cannot be solved in closed form. Examples include the bond market in portfolio choice models (Krusell and Smith, 1997), and the goods and labor markets under endogenous supply and demand. In state-space methods, these conditions require an inner loop at each period to find the market-clearing price, substantially increasing computational cost.

The RTM handles non-trivial market clearing by tracking the implied price path along the sequence and updating it across iterations. Rather than solving for the exact market-clearing price at each period, the RTM:

1. Computes the implied price from the current iteration's allocations and the market clearing identity.
2. Updates the guessed price path using a damped convex combination.
3. Relies on the stability and uniqueness of the equilibrium so that the implied price update approaches the market-clearing price over iterations.

To motivate why the implied price approaches the market-clearing price under iteration, consider the market clearing condition $Q^D(p_t, X_t, X_{t+1}) - Q^S(p_t, X_t, X_{t+1}) = 0$, where Q^D and Q^S are demand and supply and p_t is the market-clearing price as in (5). Suppose Q^D lacks a closed-form characterization. At iteration n, the RTM fixes demand at the guessed price $p_t^{(n)}$ and computes the implied price p_t^* that clears the supply side. As the predicted price path converges, the implied price converges to the true market-clearing price:

$$\begin{aligned}
&\lim_{n\to\infty} Q^D(p_t^{(n)}, X_t, X_{t+1}) - Q^S(p_t^*, X_t, X_{t+1}) = 0 \\
&\implies Q^D(\lim_{n\to\infty} p_t^{(n)}, X_t, X_{t+1}) - Q^S(p_t^*, X_t, X_{t+1}) = 0 \\
&\implies Q^D(p_t, X_t, X_{t+1}) - Q^S(p_t^*, X_t, X_{t+1}) = 0 \\
&\implies p_t^* = p_t \quad \text{(by uniqueness of the equilibrium).}
\end{aligned} \tag{19}$$

The passage from the first to the second line exchanges the limit and the aggregate demand function, which involves two layers not explicit in the generic form (5). First, at the aggregate level, aggregate demand is an integral of individual demands over the cross-sectional distribution. Because the compact budget set uniformly bounds individual demands across all iterations, the limit passes inside the integral. Second, at the individual level, the continuity of policy functions in prices ensures that each individual's demand converges

pointwise as the guessed price converges.[25]

The standard alternative—solving for the exact market-clearing price at each period via a demand-and-supply fixed point—requires a nested inner loop within each iteration, which can be costly when the market clearing condition lacks a closed-form solution. The implied price approach avoids this nested loop while, under stability and uniqueness of the RCE, approaching the market-clearing price under iteration (see the convergence argument above). The computational savings are particularly large for models with multiple market clearing conditions, where cross-market general equilibrium effects arise: a change in one price (e.g., the bond price) alters household portfolio choices, shifting the capital distribution, which changes wages and the interest rate, which feeds back into bond demand. These cross-market interactions are highly nonlinear when occasionally binding constraints are active. Because the implied-price approach updates all prices simultaneously and solves the full nonlinear model at each iteration, these interactions are captured directly in the forward simulation rather than approximated locally. In the two-asset HANK model of Section 3, two non-trivial market clearing conditions (labor and bonds) are updated simultaneously without nested loops, whereas a state-space method would require solving a multi-dimensional price-finding system at each period.

A specific challenge arises when the market clearing condition involves a zero-net-supply asset, such as bonds. The identity $q^b \cdot 0 = 0$ provides no information about the bond price. I address this through a fictitious-supply tatonnement, introducing a fictitious bond supply $\bar{B} > 0$ that anchors the price update formula; in equilibrium, household net bond demand converges to zero, so the fictitious supply does not distort allocations. The details are provided in Online Appendix G.

2.7 Extensions

The generic framework above is deliberately kept simple. The RTM has been applied to a range of richer environments, with sample codes available in the supplementary material:

- *Heterogeneous firms:* models with firm-level investment and idiosyncratic productivity (Khan and Thomas, 2008);
- *Multiple aggregate shocks:* joint TFP and fiscal spending shocks (up to 21 combined states), as well as aggregate uncertainty (stochastic volatility) shocks;
- *Frictional labor markets:* Diamond-Mortensen-Pissarides (DMP) models with bilaterally

[25]The first step applies the Lebesgue dominated convergence theorem, with the uniform bound on individual demands from the compact budget set serving as the dominating function. The second step follows from the Maximum Theorem: the smooth objective function and compact budget correspondence guarantee continuity of policy functions in prices.

determined wages;

- *Sticky prices:* New Keynesian models with Rotemberg or Calvo pricing (the HANK application in Section 3 uses Rotemberg);
- *Multi-dimensional endogenous aggregate states:* models where the distribution interacts with additional slow-moving aggregate variables such as capital and bonds.

A case of RTM failure The method can fail when multiple slow-moving endogenous aggregate states interact. This arises in a New Keynesian model with Calvo pricing when the nominal interest rate adjusts slowly. Under Calvo pricing, price dispersion (Δ) becomes a slow-moving endogenous aggregate state. Using Δ_{t-1} as the sufficient statistic, the RTM converges when the Taylor rule is contemporaneous. However, when interest rate smoothing is introduced ($(1+i_t) = (1+i_{t-1})^{\rho_r}[\text{Taylor rule}]^{1-\rho_r}$ with $\rho_r = 0.85$), the lagged interest rate i_{t-1} becomes a second slow-moving endogenous state. Matching on Δ_{t-1} alone no longer suffices: two matched periods may have different i_{t-1} values, introducing errors that amplify through the persistent Taylor rule. Extending the matching to the joint space (Δ_{t-1}, i_{t-1}) using Delaunay triangulation, kernel regression, and alternative damping schedules all produce the same sustained divergence after brief initial progress (Online Appendix C). The breadth of matching-geometry variants tried makes a limitation-of-matching explanation harder to sustain than a model-side explanation (non-existence of an ergodic RCE under this configuration), though neither can be ruled out from numerical evidence alone. By contrast, the Rotemberg specification used in the HANK application (Section 3) avoids this issue because inflation is a jump variable determined by the New Keynesian Phillips curve, and no price dispersion state arises. Online Appendix C provides quantitative evidence: the non-persistent Taylor rule specification converges (with non-monotonic oscillations), while the persistent specification diverges exponentially after brief initial progress. That said, the two-asset HANK application (Section 3) shows that multi-dimensional slow-moving components—aggregate capital and bonds—can be successfully handled by the RTM when a scalar sufficient statistic remains valid.

3 Leading application: Two-asset HANK model

This section applies the RTM to a two-asset heterogeneous-agent New Keynesian (HANK) model. The model features two-dimensional aggregate uncertainty (TFP and monetary policy), endogenous portfolio choice between risky capital and safe bonds, multiple occasionally binding constraints, non-trivial market clearing in both the labor and bond markets, convex capital adjustment costs, endogenous labor supply, Rotemberg price adjustment, and a Taylor-type

monetary policy rule.

3.1 Model

Households A unit continuum of households indexed by idiosyncratic labor productivity z chooses consumption c, labor supply n, next-period capital k', and next-period bond holdings b' to solve (with individual state (k, b, z) and aggregate state X):

$$V(k,b,z;X) = \max_{c,n,k',b'} \log(c) - \frac{\eta}{1+\frac{1}{\chi}} n^{1+\frac{1}{\chi}} + \beta \mathbb{E} V(k',b',z';X') \tag{20}$$

$$\text{s.t.} \quad c + k' + q^b(X) b' + \Psi(k',k) = (1+r(X))k + b + w(X) z n + D(X) \tag{21}$$

$$k' \geq \underline{k}, \quad b' \geq \underline{b} \tag{22}$$

where χ is the Frisch elasticity of labor supply, η controls labor disutility, $\Psi(k',k) = \frac{\mu}{2}\left(\frac{k'-k}{k}\right)^2 k$ is a convex capital adjustment cost with parameter $\mu > 0$.[26] $D(X)$ is a lump-sum dividend equal to firm profits net of the Rotemberg adjustment cost:

$$D(X) = (1 - mc(X))Y(X) - \tfrac{\psi}{2}(\pi(X) - \pi^*)^2 Y(X). \tag{23}$$

The aggregate state $X = (\Phi, A, m^{MP})$ includes the joint distribution Φ of individual states (k, b, z), aggregate TFP A, and the monetary policy shock m^{MP}. For notational brevity, the dependence on X is suppressed in what follows.

The first-order conditions yield the Frisch labor supply $n = \left(\frac{wz}{\eta c}\right)^{\chi}$ and two inter-temporal optimality conditions:

$$\frac{1 + \Psi_1(k',k)}{c} = \beta \mathbb{E}\left[\frac{1 + r' + \Psi_2(k'',k')}{c'}\right] + \lambda, \quad \lambda \geq 0,\ (k' - \underline{k})\lambda = 0 \tag{24}$$

$$\frac{q^b}{c} = \beta \mathbb{E}\left[\frac{1}{c'}\right] + \phi, \quad \phi \geq 0,\ (b' - \underline{b})\phi = 0 \tag{25}$$

where λ and ϕ are the Lagrange multipliers on the capital and bond borrowing constraints, respectively. Here $\Psi_1(k',k) = \mu(k'-k)/k$ is the marginal adjustment cost and $\Psi_2(k'',k') = -\partial\Psi(k'',k')/\partial k' = (\mu/2)\left[(k''/k')^2 - 1\right]$ captures the marginal benefit of higher k' through reduced future adjustment costs. Because the adjustment cost $\Psi(k'',k')$ has k' in the denominator, choosing $k' = 0$ would generate an infinitely large adjustment cost in the following period. I impose a small positive lower bound $\underline{k} > 0$ on capital holdings to ensure

[26] Due to the capital adjustment cost, the individual state space cannot be reduced to a single wealth variable $s = b + k$, which adds an additional computational dimension to the portfolio choice problem presented in Krusell and Smith (1997).

the adjustment cost ratio k'/k is well-defined.

Firms A continuum of monopolistically competitive intermediate goods firms produces output using aggregate capital K and labor N with technology $Y = AK^{\alpha}N^{1-\alpha}$. Firms face Rotemberg quadratic price adjustment costs $\frac{\psi}{2}(\pi - \pi^*)^2 Y$, where π is inflation and π^* is the inflation target. Optimal pricing yields the New Keynesian Phillips curve (NKPC):

$$mc = \frac{(\varepsilon - 1) + \psi(1+\pi)(\pi - \pi^*) - \mathbb{E}\left[\beta \frac{C}{C'}\psi(1+\pi')(\pi' - \pi^*)\frac{Y'}{Y}\right]}{\varepsilon} \tag{26}$$

where mc is real marginal cost and ε is the elasticity of substitution across varieties. Cost minimization by intermediate goods firms yields the standard factor pricing conditions:

$$r = mc \cdot \alpha A(K/N)^{\alpha-1} - \delta, \qquad w = mc \cdot (1-\alpha)A(K/N)^{\alpha} \tag{27}$$

where factors are paid their marginal products scaled by real marginal cost, reflecting the monopolistic markup ($mc = 1$ recovers the competitive case).

Monetary policy The central bank sets the nominal interest rate according to a Taylor rule:

$$1 + i = \frac{1+\pi}{q^b_{ss}} \cdot \left(\frac{1+\pi}{1+\pi^*}\right)^{\phi_\pi} \cdot \left(\frac{Y}{Y_f}\right)^{\phi_Y} \cdot e^{m^{MP}} \tag{28}$$

where q^b_{ss} is the steady-state bond price, $Y_f = AK^{\alpha}_{ss}N^{1-\alpha}_{ss}$ is a reference output level that scales with current TFP A but uses steady-state factor inputs (so the output gap Y/Y_f captures demand-driven deviations from potential), and m^{MP} is an i.i.d. monetary policy shock.[27] The bond is real (one consumption unit next period; Euler equation (25)). With the nominal gross rate $1 + i_t$ set by the Taylor rule, realized inflation yields the ex post Fisher relation $1 + i_t = (1 + \pi_{t+1})/q^b_t$.

Market clearing Four markets clear simultaneously:

$$K' = \int k'\, d\Phi, \quad 0 = \int b'\, d\Phi, \quad N = \int zn\, d\Phi \tag{29}$$

$$Y = C + K' - (1-\delta)K + \int \Psi(k', k)\, d\Phi + \tfrac{\psi}{2}(\pi - \pi^*)^2 Y \tag{30}$$

[27] The steady-state bond price q^b_{ss} is determined endogenously by the bond market clearing condition, reflecting precautionary saving motives in the heterogeneous-agent economy.

The resource constraint (30) is consistent with the household budget constraint through Walras' law. The bond market clearing condition $\int b' d\Phi = 0$ (zero net supply) is non-trivial, as q^b cannot be directly backed out from an invertible identity; this is addressed through a fictitious-supply tatonnement (Online Appendix G).

Aggregate shocks TFP follows a 2-state symmetric Markov chain $A \in \{0.99, 1.01\}$ with quarterly persistence 0.875. The monetary policy shock is a 3-state i.i.d. process $m^{MP} \in \{-m_0, 0, m_0\}$ with $m_0 = 0.0025$. The combined exogenous state has $2 \times 3 = 6$ states, with transition matrix given by the Kronecker product $\Pi^A \otimes \Pi^{MP}$.

3.2 Calibration

Table 2 reports parameter values at a quarterly frequency. The model shares the two-asset structure of Bayer et al. (2019) and Luetticke (2021), but whereas those papers solve with linearized methods, the global nonlinear solution here reveals state-dependent dynamics that linearization cannot capture. Standard parameters (α, δ, χ, ε, ϕ_π, ϕ_Y, ρ_z, σ_z, TFP process) follow the HANK literature. The capital adjustment cost parameter $\mu = 0.01$ is standard in the macro literature and generates a fraction of capital-constrained households of approximately 9%, with an additional 12% bond-constrained with a positive multiplier. Two parameters are calibrated to match equilibrium targets: $\beta = 0.991$ delivers a risk-free rate of approximately 1.5% annualized, and $\underline{b} = -2.00$ targets a fraction of bond-constrained households of approximately 41%. The interplay of binding constraints and portfolio choice generates a mean equity risk premium of approximately 4.1 percentage points annualized. The Calvo-equivalent stickiness $\theta = 0.66$ (average price duration of 3 quarters) maps to the Rotemberg cost parameter $\psi = (\varepsilon - 1)\theta/[(1-\theta)(1-\beta\theta)] \approx 50.5$, where the mapping equates the log-linearized slopes of the Calvo and Rotemberg Phillips curves at zero trend inflation.

3.3 RTM implementation

The RTM solves this model by iterating on five aggregate paths: capital $\{K_t\}$, labor supply $\{N_t\}$, bond price $\{q_t^b\}$, inflation $\{\pi_t\}$, and marginal cost $\{mc_t\}$. The backward–forward–update structure follows Section 2.4, with the sufficient statistic K used for matching (validated by the monotonicity test in Section 4.4).

The key model-specific features are the NKPC expectations and the bond market clearing. In the backward pass, the matched future periods provide not only household continuation values (24)–(25) but also the aggregate variables (π, C, Y) needed to evaluate the NKPC (26). In the forward pass, the bond price is updated via the fictitious-supply tatonnement

Table 2: Parameter values: Two-asset HANK model

Parameter	Symbol	Value	Parameter	Symbol	Value
Household preferences			*New Keynesian*		
Discount factor*	β	0.991	Elast. of substitution	ε	10.000
Frisch elasticity	χ	1.000	Price stickiness	ψ	50.500
Labor disutility	η	8.000	Taylor: inflation	ϕ_π	1.500
			Taylor: output gap	ϕ_Y	0.125
Production			Inflation target	π^*	0.005
Capital share	α	0.360			
Depreciation	δ	0.025	*Asset markets*		
Adj. cost	μ	0.010	Capital lower bound	$\underline{k}$	0.100
			Bond borrowing limit*	$\underline{b}$	-2.000
Idiosyncratic shocks			Tatonnement step	$\bar{B}$	20.000
Persistence	ρ_z	0.900			
Std. deviation	σ_z	0.050	*Aggregate shocks*		
			TFP grid	A	$\{0.990, 1.010\}$
			TFP persistence	ρ_A	0.875
			MP shock half-width	m_0	0.0025

Notes: *Calibrated. The adjustment cost is $\Psi(k', k) = (\mu/2)((k' - k)/k)^2 k$. β targets a risk-free rate of ≈1.5% annualized. $\underline{b}$ targets a bond-constrained fraction of ≈41%. η normalizes hours. ψ is the Rotemberg parameter equivalent to Calvo $\theta = 0.66$ (3-quarter average price duration). TFP process follows Krusell and Smith (1998).

(Online Appendix G):

$$q_t^{b,\text{new}} = \frac{\int \left[\omega + wzn + D - c - k' - \Psi(k', k)\right] d\Phi + q_t^b \bar{B}}{\bar{B}} \tag{31}$$

where $\omega = (1 + r)k + b$ is beginning-of-period wealth, D is the lump-sum dividend (23), and $\bar{B} = 20$ is the fictitious bond supply. Inflation is recovered by inverting the Taylor rule (28) using the ex post Fisher relation.

The algorithm iterates with damped updates until the mean squared difference between predicted and realized paths falls below the convergence tolerance. The baseline solution ($T =$ 2,000, seed 100) converges in 677 iterations (approximately 9,000 seconds on a MacBook Pro M3 Pro). Regressing $\log K_{t+1}$ on $(\log K_t, \log A_t, \log K_t \cdot \log A_t)$ yields $R^2 > 0.999$, indicating that aggregate capital dynamics are nearly linear despite the model's high-dimensional distribution over (k, b, z). The RTM does not use this regression as a forecasting rule; it serves as a diagnostic for the aggregate law of motion. The validity of K as a sufficient statistic for the matching step is separately verified by the formal monotonicity test in Section 4.4.

3.4 Equilibrium dynamics

Nonlinear dynamics In a state-space approach, aggregate variables involved in non-trivial market clearing—such as the bond price and labor—would each require a separate parametric law of motion as a function of K. Table 3 assesses the adequacy of such approximations by fitting linear and quadratic regressions of each aggregate variable on K, conditioning on TFP state.

Table 3: Nonlinearity diagnostic: R^2 improvement from quadratic term

Variable (regressed on K)	R^2_{linear}	$R^2_{\text{quadratic}}$	Improvement
K_{t+1}	0.999+	0.999+	<0.001
q^b	0.929	0.933	0.003
w	0.934	0.934	<0.001
r	0.936	0.936	<0.001
Risk premium	0.662	0.679	0.017
π	0.612	0.616	0.004
N	0.307	0.308	<0.001
C	0.996	0.996	<0.001
I	0.663	0.664	0.001
Frac. k-constrained	0.336	0.526	0.190
Frac. b-constrained	0.998	0.998	<0.001

Notes: Each regression is estimated separately within each TFP state to control for the exogenous productivity level; the table reports the unweighted average R^2 across the two TFP states (results are similar within each state). Lower R^2 values for π, risk premium, N, and I reflect the direct influence of the monetary policy shock state, which varies within each TFP partition and is not captured by K alone. The fraction of b-constrained households achieves $R^2 = 0.998$ linearly, reflecting the tight link between aggregate capital and the extensive margin of constraint binding.

Several variables exhibit meaningful nonlinearities even at the linear level: the bond price q^b and wage w achieve linear R^2 values of 0.929 and 0.934, respectively, indicating that a linear law of motion would leave non-trivial residual variation. The risk premium achieves a linear R^2 of only 0.662, with a noticeable quadratic improvement (0.017), reflecting genuinely nonlinear dynamics arising from constraint–portfolio interactions. The fraction of capital-constrained households exhibits the most pronounced nonlinearity: a quadratic R^2 improvement of 0.190, reflecting the nonlinear extensive margin of constraint binding. Labor N and investment I exhibit even lower linear R^2 values (0.307 and 0.663), reflecting the strong influence of the monetary policy shock state beyond what aggregate K alone captures. The non-improvability of these fits by low-order polynomials underscores the value of a global solution method that does not rely on parametric laws of motion.

Asymmetry across business cycle phases The model generates pronounced asymmetry between recessions and expansions. The annualized risk premium—the total capital return

(rental rate plus dividends per unit of capital) minus the bond yield—averages approximately 4.1 percentage points and is state-dependent. During expansions ($A = 1.01$), aggregate capital is high and constraints are loose, allowing wider portfolio dispersion and a larger capital-bond spread; the risk premium reaches approximately 4.3 percentage points. During recessions ($A = 0.99$), bond scarcity intensifies: constrained households substitute toward bonds, driving up bond demand and compressing the capital-bond spread to approximately 3.8 percentage points. This bond-scarcity channel—absent in single-asset models, models without occasionally binding constraints, and linearized solutions—generates a sizable and state-dependent risk premium. Section 3.6 examines the state-dependent dynamics in detail.

Distribution dynamics To capture distributional dynamics from both the extensive margin (fraction of constrained households) and the intensive margin (how tightly constrained they are), I define aggregate constraint multipliers: $\bar{\lambda} \equiv \int \lambda \, d\Phi$ for the capital constraint and $\bar{\phi} \equiv \int \phi \, d\Phi$ for the bond borrowing constraint. These continuous measures integrate over the entire wealth distribution, including near-binding households with high marginal propensities to consume. The resulting characterization of constraint dynamics is consistent with the hand-to-mouth behavior emphasized by Kaplan and Violante (2014) and Kaplan et al. (2018), but computed here within a fully nonlinear equilibrium. In particular, the state-dependent analysis in Section 3.6 focuses on $\bar{\lambda}$, which varies meaningfully with the business cycle and governs the key nonlinear transmission mechanisms documented below.

3.5 Generalized transition function

To analyze how the distributional dynamics documented above shape the economy's response to shocks, I use the *generalized transition function* (GTF), which traces the stochastic path of any aggregate variable from a specific initial state (Φ_0, S_0). For any variable of interest v, the GTF at horizon j is

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

where $\Phi_j = \Phi_j(S_1, \ldots, S_j; \Phi_0)$ is the distribution at horizon j, determined by the equilibrium law of motion and the realized shock sequence from Φ_0, and S_j follows a j-step Markov chain Γ^j from the initial exogenous state S_0. Since S_j is a random variable, g_j is itself also a random variable for each j. Fixing the initial shock S_1 and taking the difference $\mathbb{E}[g_j|\Phi_0, S_0, S_1] - \mathbb{E}[g_j|\Phi_0, S_0]$ yields the *generalized impulse response function* (GIRF) of Koop et al. (1996) and Andreasen et al. (2017). Conditioning on different initial distributions Φ_0—for example, distributions with high versus low aggregate constraint intensity $\bar{\lambda} \equiv \int \lambda \, d\Phi$—

reveals how the economy's response to identical shocks varies with the distributional state. Because the RTM computes the full global solution path, the GTF can be constructed directly from the converged paths. Online Appendix E provides the formal definition along with an extension to initial conditions off the recursive competitive equilibrium. Using the GTF, I highlight the key economic insights from the model below, with the first two explored further in Online Appendices D and E.

3.6 Global nonlinear equilibrium analysis

State-dependent shock responses Using the GTF, I condition on $\bar{\lambda} \equiv \int \lambda \, d\Phi$—the aggregate capital constraint multiplier defined in Section 3.4—to classify simulation periods into tight-constraint and loose-constraint regimes. Because $\bar{\lambda}$ covaries strongly with TFP (tight constraints coincide with recessions), I condition *within each TFP state*: for each TFP realization, periods are ranked by $\bar{\lambda}$ and split into tight and loose groups. This within-TFP conditioning isolates the constraint channel from the TFP level effect, ensuring that the comparison reflects genuine state dependence rather than differences in the exogenous shock. When $\bar{\lambda}$ is high within a given TFP state, more households face binding or near-binding capital constraints, raising cross-sectional average MPCs and amplifying the aggregate demand response to shocks. This amplification mechanism is analogous to Petrosky-Nadeau et al. (2018), though driven here by capital constraint intensity rather than labor market disasters.

Constraint–portfolio interactions generate a sizable and state-dependent risk premium—the annualized total capital return (rental rate plus dividends per unit of capital) minus the bond yield—that is entirely absent in linearized solutions. To ensure robust identification of the state-dependent GIRFs, I use a percentile-averaging procedure: paired-simulation GIRFs are computed at each period within each group, and the resulting impulse responses are averaged across all conditioning periods in the group. This eliminates the sensitivity to any single conditioning period that arises in nonlinear stochastic environments.[28]

Table 4 and Figure 2 document the state-dependent GIRFs of a contractionary monetary policy shock (25 basis points) using two complementary conditioning approaches. Table 4 conditions *within each TFP state* separately (top/bottom quartile of $\bar{\lambda}$ within each TFP realization), providing the cleanest identification of the constraint channel by holding TFP fixed. Figure 2 conditions on $\bar{\lambda}$ *across all periods* regardless of TFP (top/bottom 10%), providing a single tight-versus-loose comparison alongside a linearized benchmark—the natural

[28]The qualitative patterns—risk premium amplification, output asymmetry, and directional differences in consumption and bond prices—are robust across alternative threshold choices (top/bottom 10%, 25%, 30%) and across simulation lengths ($T = 2{,}000$, $T = 3{,}000$, and $T = 5{,}000$). Panel B of Table 4 reports the corresponding state-dependent responses to TFP shocks.

counterpart, since the linearized IRF is state-independent by construction. Both use paired stochastic simulations in which both paths receive the same random draws; only the MP shock at $j = 1$ differs.

Table 4 (Panel A) reveals that the risk premium rises by +0.095 pp under tight constraints in recessions but is essentially unchanged under loose constraints (+0.002 pp), and output contracts 1.7 times as much (−0.134% versus −0.083%). In expansions, the pattern persists: the risk premium rises by +0.058 pp under tight constraints but barely moves under loose constraints (+0.003 pp). Bond price and consumption responses also differ in direction across constraint regimes in both TFP states, though these differentials are small in absolute magnitude. Figure 2 confirms these patterns in the unconditional comparison. These regime-dependent differences—large magnitude differences in the risk premium and output, and directional differences in bond prices and consumption—illustrate that $\bar{\lambda}$ governs not only the strength but partly the direction of monetary transmission, features that first-order perturbation methods cannot capture by construction. The consumption and bond-price differentials, though statistically significant (Online Appendix D.1), are small in absolute terms; the risk premium and output differentials carry the economic weight of the state-dependence finding.

Panel B reports the corresponding analysis for TFP shocks. The state dependence is more moderate: the risk premium response to a positive TFP shock is +0.461 pp under tight constraints versus +0.432 pp under loose constraints—a gap of 0.03 pp that is persistent but small relative to the overall response. Output and consumption responses are nearly identical across constraint regimes. The contrast with Panel A is informative: monetary policy transmission is strongly state-dependent because the MP shock operates through the bond market, where constrained households' portfolio rebalancing frictions create qualitatively different responses. TFP shocks, by contrast, affect all households through the production side (wages, returns to capital), a channel that is less sensitive to the distribution of binding constraints.

Comparison with linearized benchmark To isolate the role of nonlinearity in aggregate dynamics, I construct a linear aggregate law of motion by projecting the converged RTM paths onto (K_t, A_t, m_t^{MP}) via OLS. This asks: what would the impulse responses look like if an econometrician summarized the model's aggregate dynamics with linear regressions? This is not an independent perturbation solution—which is unavailable for models with occasionally binding constraints—but rather a linear approximation to the RTM's own aggregate law of motion, analogous to the linear perceived law of motion in Krusell and Smith (1998). Figure 2 overlays the linearized IRF (gray dotted) on the RTM state-dependent GIRFs (red solid:

Table 4: State-dependent impact responses

	Low TFP ($A = 0.99$)		High TFP ($A = 1.01$)	
	Tight	Loose	Tight	Loose
Panel A: MP shock (contractionary, 25 bp)				
Output (% dev.)	−0.134	−0.083	−0.110	−0.075
Consumption (% dev.)	+0.022	−0.025	+0.013	−0.015
Inflation (pp ann.)	−0.216	−0.178	−0.201	−0.166
Risk premium (pp ann.)	+0.095	+0.002	+0.058	+0.003
Panel B: TFP shock ($A : 0.99 \leftrightarrow 1.01$)				
Output (% dev.)	+2.293	+2.312	−2.337	−2.299
Consumption (% dev.)	+0.501	+0.486	−0.484	−0.500
Inflation (pp ann.)	−0.068	−0.054	+0.050	+0.062
Risk premium (pp ann.)	+0.461	+0.432	−0.428	−0.452

Notes: Percentile-averaged impact responses at horizon $j = 1$, computed via paired stochastic simulations (500 per conditioning point). The conditioning is *within each TFP state*: for each TFP realization separately, periods are ranked by the aggregate capital constraint multiplier $\bar{\lambda} \equiv \int \lambda \, d\Phi$ and split into "Tight" (top quartile) and "Loose" (bottom quartile). This isolates the constraint channel from TFP level effects, since $\bar{\lambda}$ covaries with TFP. Each paired simulation draws the same random sequences for both the shocked and unshocked paths; only the forced shock at $j = 1$ differs. In Panel A, TFP evolves identically in both paths while the MP state is forced to be contractionary (shocked) versus neutral (unshocked) at $j = 1$. In Panel B, the MP shock is drawn identically while TFP is forced to transition at $j = 1$. Under "Low TFP," Panel B reports the response to a positive TFP shock ($A : 0.99 \to 1.01$); under "High TFP," it reports the response to a negative TFP shock ($A : 1.01 \to 0.99$). Risk premium is defined as the total capital return (rental rate plus dividends per unit of capital) minus the bond yield. Figure 2 plots the full impulse response paths for Panel A using an alternative conditioning: top/bottom 10% of $\bar{\lambda}$ across all periods regardless of TFP state, providing a single tight-versus-loose comparison alongside a linearized benchmark.

tight constraints; blue dashed: loose constraints) for a contractionary monetary policy shock.

Linearization eliminates state dependence entirely, producing a single impulse response that cannot capture differential transmission across constraint regimes. It misses the risk premium amplification, the output asymmetry, and the directional differences in bond-price and consumption responses across constraint regimes. This comparison illustrates the value of a global nonlinear solution: the state-dependent dynamics documented above are invisible to linear methods, even when applied to the same converged equilibrium paths.

Heterogeneous portfolio rebalancing The global solution reveals heterogeneous portfolio dynamics across wealth groups, illustrated in Figure 3. Defining the leverage ratio as b'/k':

Level heterogeneity. Wealth-poor households (bottom tercile) are heavily leveraged, with $b'/k' \approx -17$: they borrow near the bond constraint and hold minimal capital. Middle-class households maintain moderate leverage ($b'/k' \approx -3.2$), while wealth-rich households (top

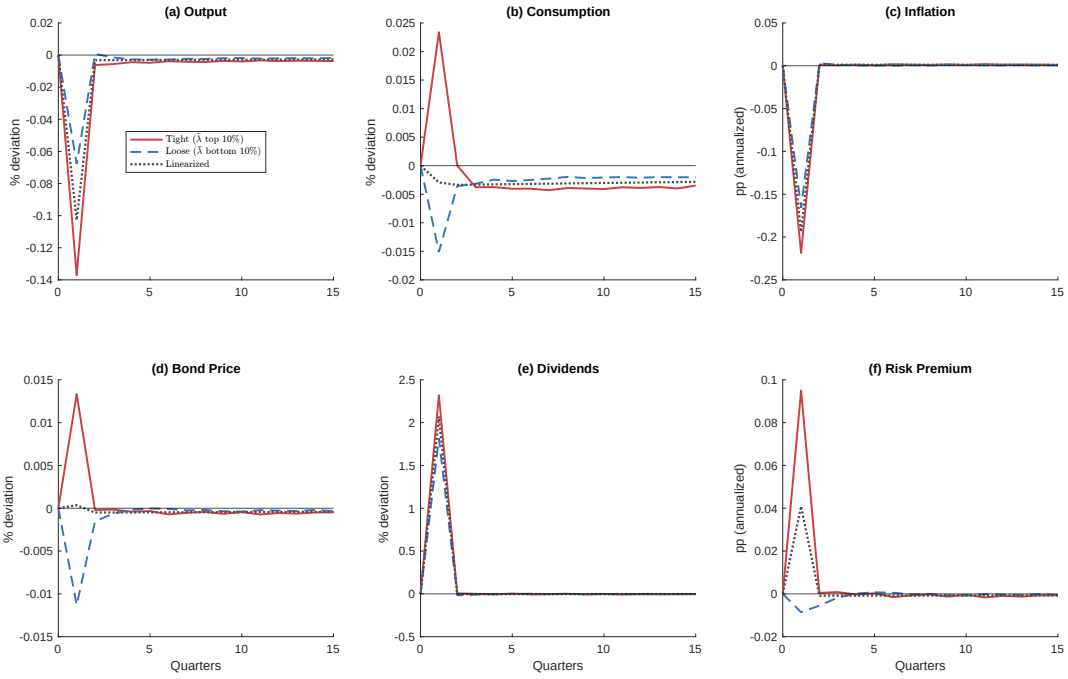

Figure 2: State-dependent GIRFs of a contractionary monetary policy shock (25 bp)

Notes: Red solid lines: tight constraint (high $\bar{\lambda}$, top 10%); blue dashed lines: loose constraint (low $\bar{\lambda}$, bottom 10%); gray dotted lines: linearized approximation. Unlike Table 4, which conditions within each TFP state separately, this figure conditions on $\bar{\lambda}$ across all periods regardless of TFP state, yielding a single tight-versus-loose comparison. This unconditional approach allows a direct comparison with the linearized benchmark (gray dotted), which is state-independent by construction. 2,000 paired stochastic simulations per conditioning point. The linearized model projects the RTM's aggregate dynamics onto linear functions of (K, A, m^{MP}) via OLS. Online Appendix D.1 provides 95% bootstrap confidence intervals confirming statistical significance of the between-regime differences.

tercile) hold modest positive bond positions ($b'/k' \approx 0.07$) (panel (a)). This cross-sectional dispersion in leverage—absent in representative-agent models—implies that aggregate shocks affect household balance sheets asymmetrically, with poor households bearing disproportionate exposure to interest rate changes.

Heterogeneous rebalancing. Leverage rebalancing is concentrated among wealth-poor and middle-class households, who transition in and out of binding constraints over the business cycle. Panel (b) shows deviations from each tercile's ergodic mean leverage: middle-class households exhibit the largest fluctuations (std = 0.95), followed by wealth-poor households (std = 0.58), while wealth-rich households barely move (std = 0.002). The middle class is most volatile because households entering and exiting binding constraints concentrate in this tercile; wealth-rich households' large capital buffers insulate them from constraint dynamics. This concentration of leverage adjustment among constrained and near-constrained households underlies the state-dependent risk premium and monetary transmission documented above. Online Appendix F documents the same patterns in a simpler portfolio choice model without

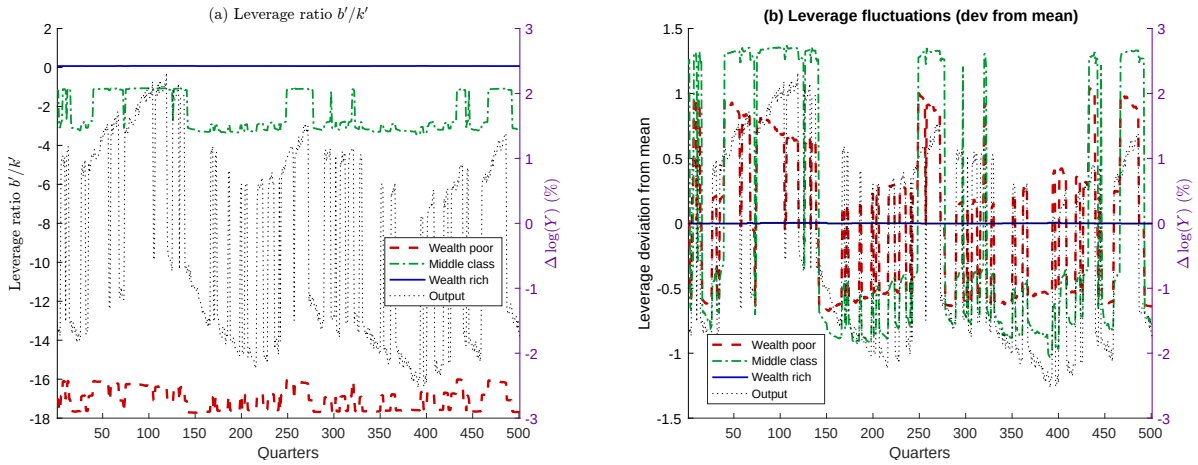

Figure 3: Heterogeneous portfolio rebalancing over the business cycle

Notes: Panel (a): leverage ratio b'/k' by wealth tercile along the ergodic equilibrium path (left axis), plotted against output deviations (right axis). Wealth-poor households (bottom tercile, red dashed) are heavily leveraged ($b'/k' \approx -17$) with large cyclical fluctuations. Middle-class households (green dash-dot) maintain moderate leverage (≈ -3.2). Wealth-rich households (blue solid) hold slightly positive bond positions (≈ 0.07). Panel (b): deviations from each tercile's ergodic mean leverage. Middle-class households exhibit the largest fluctuations as households transition in and out of binding constraints, while wealth-rich households barely move. Online Appendix F documents the same patterns in a simpler portfolio choice model without the NK block.

the NK block, confirming that these dynamics arise from the interaction of binding constraints with endogenous portfolio choice.

Bootstrap confidence intervals for the state-dependent monetary-policy GIRFs are in Online Appendix D, and Online Appendix F provides a simpler portfolio-choice counterpart without the NK block. The nonlinear features documented in this section—state-dependent responses, endogenous risk premia, and constraint-driven amplification—depend critically on the solution's accuracy, which Section 4 validates through comprehensive tests.

4 Accuracy and robustness

This section evaluates accuracy and robustness along four dimensions: Euler equation errors (Section 4.1); external benchmarks, including an off-equilibrium forward simulation test (Section 4.2); robustness to random seeds, simulation lengths, and interpolation schemes (Section 4.3); and the necessary-condition check for the sufficient statistic (Section 4.4).

4.1 Euler equation errors

I compute Euler equation errors following the methodology of Judd (1992) for the two-asset HANK model of Section 3. For each household with state (k, b, z) at each period t in the

converged solution, I evaluate the residual from the capital Euler equation (24):

$$EE_t(k, b, z) = \log_{10} \left| \frac{\beta \mathbb{E}\left[(1 + r' + \Psi_2)u'(c')\right]}{(1 + \Psi_1)u'(c)} - 1 \right| \tag{32}$$

for unconstrained households (those with $\lambda = 0$). Specifically, I evaluate the error at each grid point (k, b, z) for each period t in the post-burn-in sample, then weight by the stationary-equilibrium distribution to compute the period-level error; the reported statistics are over the time series of these distribution-weighted errors. The distribution-weighted Euler equation errors across the simulation path yield a mean error of -4.82 in $\log_{10}$ units. Following the classification in Judd (1992), errors below -3 are considered "good" and errors below -5 are "excellent"; the mean error is close to the "excellent" threshold. The reported Euler-equation error concerns the capital Euler equation only. The bond Euler/KKT condition (25) is satisfied exactly by construction in the implementation, rather than separately validated by this exercise: the solution procedure enforces the bond block directly through the budget-residual step and the multiplier ϕ with complementary slackness ($\phi \geq 0$ binding for constrained households, $\phi = 0$ for unconstrained households). For this reason, no separate bond-Euler error statistic is reported here.[29]

A brief remark on the formal consistency result (Proposition 3): the proposition establishes convergence under contraction of the Bellman operator with modulus β, bounded returns, and a compact state space. For the HANK model with log utility, occasionally binding constraints, and endogenous labor, verifying these assumptions analytically is difficult—the state space is truncated by grid bounds rather than naturally compact, and boundedness of returns relies on the grid structure. The empirical accuracy evidence in this section—Euler equation errors near the "excellent" threshold, multi-seed robustness, off-equilibrium generalization, and the monotonicity diagnostic—provides complementary assurance that the converged solution is accurate, even where the formal assumptions are not directly verifiable. The proposition's role is to establish that the RTM's approximation principle is sound in a well-defined class of models; the accuracy tests confirm that this principle delivers reliable results in the leading application.

For the portfolio choice model with endogenous labor supply (Krusell and Smith 1997 extension), I additionally compute Euler equation errors using the Den Haan (2010) protocol, which evaluates errors at off-equilibrium points by perturbing the aggregate capital stock.

[29] The capital and bond Euler equations are solved asymmetrically in the implementation: consumption c is determined directly from the capital Euler equation, and the capital Euler error (-4.82) reflects interpolation error across the simulation. The bond block is satisfied by construction because b' is determined as the budget residual given c and k', and ϕ is defined as the bond Euler residual with complementary slackness at the borrowing constraint.

The distribution-weighted errors remain below -4 across the capital grid, confirming accuracy away from the equilibrium path.

4.2 External benchmarks

Den Haan (2010) protocol for Krusell and Smith (1998) (hereafter KS1998) I apply the RTM to the canonical Krusell and Smith (1998) model and benchmark against the published accuracy results in Den Haan (2010). The RTM solution achieves a law-of-motion R^2 of 0.99997, and the maximum absolute percentage error in the aggregate capital path relative to the direct simulation is 0.08%, which is competitive with the best methods reported in that comparison study.

Off-equilibrium forward simulation I extend the Den Haan (2010) accuracy test to evaluate how well the converged RTM policy functions generalize to off-equilibrium aggregate states. Starting from the steady-state distribution, I advance the cross-sectional distribution using the converged policy functions. At each period, the aggregate state is matched to the equilibrium path via iso-shock interpolation at the *simulated* aggregate capital K_t^{sim} rather than the equilibrium capital K_t^{eq}. Any drift between K_t^{sim} and K_t^{eq} reflects accumulated error from evaluating policy functions at off-equilibrium aggregate states.

Table 5: Off-equilibrium forward simulation accuracy

Variable	$\max \lvert e_t \rvert / \bar{x}$	RMSE / $\bar{x}$
Panel A: Krusell–Smith (1998), $T = 5{,}000$, $N_S = 2$		
Capital K	0.282%	0.039%
Panel B: Two-Asset HANK, $T = 5{,}000$, $N_S = 6$		
Capital K	0.389%	0.134%
Labor N	0.144%	0.043%
Bond price q	0.046%	0.013%
Inflation π (ann.)	0.108 pp	0.030 pp
Risk premium (ann.)	0.180 pp	0.059 pp

Notes: For each variable x_t, the forward simulation advances using the converged RTM policy functions evaluated at the *simulated* aggregate state K_t^{sim} (via iso-shock matching), rather than the equilibrium path K_t^{eq}. The error $e_t = x_t^{\text{sim}} - x_t^{\text{eq}}$ measures how well the converged policy functions generalize to off-equilibrium aggregate states. Both models use the same shock sequence as the equilibrium path. For inflation and risk premium, errors are reported in percentage points (annualized). Panel B uses policies from the $T = 5{,}000$ solution, which provides the finest state-space coverage for the matching step. The initial cross-sectional distribution is the equilibrium distribution, obtained by evolving the stationary-equilibrium distribution forward for 500 periods using the on-equilibrium policy functions before the test window begins.

Table 5 reports the results, following Den Haan (2010) in focusing on error magnitudes rather than R^2. For KS1998, the maximum capital error is 0.28% of steady-state K with

RMSE of 0.04%. For the HANK model (using policies from the $T = 5{,}000$ solution, which provides the finest state-space coverage for the matching step), errors are larger but remain moderate: the maximum capital error is 0.39% with RMSE of 0.13%. The risk premium—the most nonlinear variable—achieves a maximum error of 0.18 pp and RMSE of 0.06 pp (annualized). The larger errors for the HANK model relative to KS1998 reflect a fundamental difference in model structure, not in solution quality: small off-equilibrium deviations in K shift the fraction of constrained households and trigger portfolio rebalancing, so that errors compound nonlinearly over the simulation—an effect largely absent in the near-linear KS1998 model. The per-period RMSE of 0.13% for capital confirms that the policy functions are highly accurate period by period.

RBC with irreversible investment For the representative-agent RBC model with an occasionally binding irreversibility constraint (McGrattan, 1996; Christiano and Fisher, 2000), I compare the RTM against three alternatives: the linearized solution, OccBin (Guerrieri and Iacoviello, 2015), and GDSGE (Cao et al., 2023). As reported in Online Appendix D, among the four methods compared at similar computation times (≈ 90 seconds), the RTM exhibits the lowest prediction error (discrepancy between predicted and realized equilibrium paths) and Euler equation error. The RTM's accuracy advantage is most pronounced in higher-order moments (skewness and kurtosis), which are sensitive to the correct treatment of occasionally binding constraints.

RBC perturbation benchmark For the canonical RBC model with log utility and inelastic labor supply, where the first-order perturbation solution (Blanchard and Kahn, 1980) is essentially exact, Figure 4 confirms that the RTM recovers the known solution: the two consumption policy functions agree closely across the ergodic capital range, and forward-simulating both under the identical shock sequence yields a capital path R^2 exceeding 0.999.

Heterogeneous-firm model For the Khan and Thomas (2008) model with non-trivial market clearing, the RTM converges in approximately one-tenth the time of the Krusell and Smith (1997) algorithm by bypassing the period-by-period price-finding loop.

4.3 Robustness to shock draws and simulation length

A key concern about simulation-based methods is whether the solution depends on the particular random draw of aggregate shocks or on the simulation length. I address this through systematic robustness checks for the two-asset HANK model.

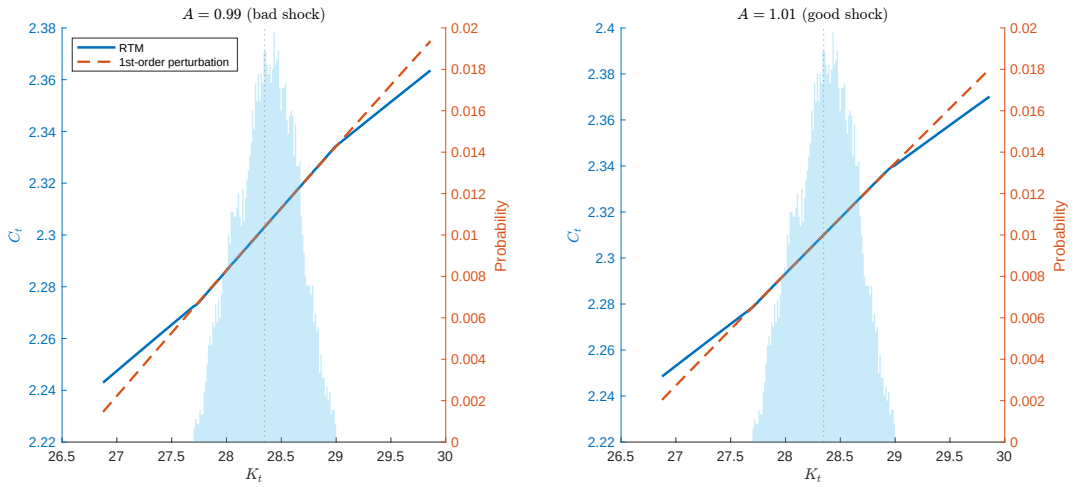

Figure 4: RTM vs. first-order perturbation: consumption policy function $C(K)$

Notes: Solid lines: RTM (extracted via iso-shock interpolation). Dashed lines: analytical first-order perturbation. Shaded histogram: ergodic distribution of aggregate capital.

Multi-seed comparison I solve the model using eight different random seeds ($T = 2{,}000$ each) and three simulation lengths ($T \in \{2{,}000, 3{,}000, 5{,}000\}$, seed 100). Table 6 reports the results.

Table 6: Solution robustness: Two-asset HANK model

	Mean K	Mean C	Mean q
Panel A: Multi-seed ($T = 2{,}000$)			
Seed 50	10.965	0.858	0.9913
Seed 100	10.930	0.857	0.9913
Seed 150	10.900	0.856	0.9906
Seed 200	10.905	0.856	0.9907
Seed 250	10.842	0.855	0.9905
Seed 500	10.870	0.855	0.9906
Seed 800	10.909	0.857	0.9907
Seed 1000	10.942	0.857	0.9907
CV	0.34%	0.12%	0.03%
Panel B: Simulation length (seed 100)			
$T = 2{,}000$	10.930	0.857	0.991
$T = 3{,}000$	10.933	0.857	0.991
$T = 5{,}000$	10.954	0.858	0.991
CV	0.10%	0.05%	0.01%

Notes: Panel A (eight seeds): each seed generates a different random path for the combined TFP $\times$ monetary policy shock; all use $T = 2{,}000$ with 500-period burn-in. Panel B: all use seed 100 with 500-period burn-in. CV = coefficient of variation. All solutions achieve LoM $R^2 > 0.9999$.

The coefficient of variation across seeds is 0.34% for capital, 0.12% for consumption, and 0.03% for the bond price, confirming robustness to the particular shock realization. The

cross-seed variation reflects different ergodic samples from different shock sequences under finite simulation length; all seeds converge to the same tolerance ($\|\text{error}\|_{\text{norm}} \leq 10^{-6}$). Across simulation lengths, the CV is at most 0.10% for all equilibrium aggregates. All solutions achieve $R^2 > 0.9999$, and $T = 2{,}000$ provides sufficient coverage.

Interpolation scheme and mesh convergence Online Appendix C reports two additional robustness exercises. First, replacing the baseline two-point linear interpolation with inverse-distance weighted (k-NN), cubic spline, or local linear interpolation produces very similar results—law-of-motion R^2 and aggregate paths are nearly indistinguishable across schemes, while Euler equation errors remain in the same accuracy range—confirming that the RTM's accuracy derives from the quality of the simulated grid rather than the specific interpolation rule. For models with higher-dimensional ergodic sets or sparser state-space coverage, richer interpolation schemes may become beneficial; the modular structure of the RTM accommodates such extensions without altering the algorithm's core logic. Second, solving both models at multiple simulation lengths confirms that the mean mesh width $\bar{h}$ decreases at rate $O(1/T)$, as predicted by the convergence rate discussion in Section 2.5: doubling T halves $\bar{h}$ at every step. Euler equation errors remain bounded and stable across all T values, and the baseline choice of $T = 2{,}000$ provides a favorable cost–accuracy trade-off for both models. The formal error bound $L \cdot h/2$ from Proposition 2 is conservative: for the HANK model, the actual Euler equation error ($10^{-4.82}$) is orders of magnitude below the upper bound.

4.4 Monotonicity diagnostic

The sufficient statistic approach requires that the distribution-weighted marginal value functions be strictly monotone in K, conditional on the exogenous state. This is a weaker condition than requiring K to be a good predictor of the future law of motion (as in Krusell and Smith 1998): the RTM requires only that two periods with similar K and identical exogenous state have similar *continuation values*, not that a parametric regression of K' on K fits well. The Spearman rank correlation provides a distribution-free test of this monotonicity condition. For the HANK model, the Spearman rank correlation between the distribution-weighted capital Euler residual $\sum_i [(1+r)/c_i]\Phi_i$ and aggregate capital averages -0.9997 across the six combined exogenous states, indicating near-perfect monotonicity. The bond Euler residual shows an equally strong correlation (mean -0.9997). Despite the model's high-dimensional endogenous aggregate state (the joint distribution over two asset dimensions and idiosyncratic productivity), a scalar sufficient statistic delivers an accurate solution.

As a direct numerical validation, Online Appendix C reports two matching-robustness

exercises. First, for the Krusell and Smith (1997) model, replacing K-based matching with matching on the full cross-sectional distribution Φ using the sup-norm distance produces nearly identical convergence trajectories and aggregate paths, with a maximum capital deviation of only 0.24% of steady state. Second, for the HANK model, enriching the matching criterion with aggregate bonds and cross-sectional second moments produces maximum capital deviation of only 0.01% of steady state, with near-zero deviations in all other variables. Both exercises confirm that K-only matching introduces no material bias.

5 Concluding remarks

This paper develops the repeated transition method (RTM), a global nonlinear solution framework for dynamic stochastic general equilibrium models. By matching periods with similar endogenous aggregate states along a simulated equilibrium path, the RTM constructs state-contingent expectations without parameterizing laws of motion or assuming perfect foresight. A formal sufficient statistic result shows that a scalar summary of the endogenous distribution recovers, at exact matches, the same continuation values as matching on the full infinite-dimensional distribution, and consistency is established under standard contraction conditions.

The leading application to a two-asset HANK model demonstrates that global nonlinear solutions yield qualitatively different conclusions from linearized methods: a sizable risk premium averaging 4.1 percentage points annualized, and state-dependent monetary policy transmission where the risk premium responds sharply to a contractionary MP shock under tight constraints but is essentially zero under loose constraints, and output contracts more strongly under tight than under loose constraints, even after controlling for TFP. Bond price and consumption responses also differ in direction across regimes. These findings demonstrate that global methods reveal equilibrium features—state-dependent risk premia, output asymmetry, and amplified transmission under tight constraints—that linear approximations systematically miss in heterogeneous-agent economies with binding constraints and endogenous portfolio choice.

More broadly, the RTM extends global nonlinear analysis to heterogeneous-agent recursive competitive equilibria featuring occasionally binding constraints, endogenous portfolio choice, non-trivial market clearing, and rich cross-sectional heterogeneity. Sample codes covering heterogeneous firms, frictional labor markets, sticky prices, uncertainty shocks, and multi-dimensional endogenous states are available in the supplementary material.

Three directions for future work stand out. First, characterizing when multiple sufficient statistics are genuinely required within a recursive competitive equilibrium, and identifying economic environments in which they arise. Second, improving coverage of rare aggregate

states through targeted shock insertion or multi-path pooling, which would be essential for disaster-risk applications. Third, speeding up the matching step through deep learning techniques (Azinovic et al., 2022; Fernández-Villaverde et al., 2023; Han et al., 2026) or lower-level programming languages. Together, these extensions would push global nonlinear analysis toward classes of models and applications currently beyond its practical reach.

References

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2018). When Inequality Matters for Macro and Macro Matters for Inequality. *NBER Macroeconomics Annual 32*, 1–75. _eprint: https://doi.org/10.1086/696046.

Algan, Y., O. Allais, and W. J. Den Haan (2008). Solving heterogeneous-agent models with parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control 32*(3), 875–908.

Algan, Y., O. Allais, and W. J. Den Haan (2010). Solving the incomplete markets model with aggregate uncertainty using parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control 34*(1), 59–68. Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.

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.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica 89*(5), 2375–2408.

Auclert, A., M. Rognlie, and L. Straub (2024). Fiscal and Monetary Policy with Heterogeneous Agents. *Working Paper*.

Azinovic, M., L. Gaegauf, and S. Scheidegger (2022). Deep equilibrium nets. *International Economic Review 63*(4), 1471–1525.

Bayer, C. and R. Luetticke (2020). Solving discrete time heterogeneous agent models with aggregate risk and many idiosyncratic states by perturbation. *Quantitative Economics 11*(4), 1253–1288.

Bayer, C., R. Luetticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Blanchard, O. J. and C. M. Kahn (1980). The solution of linear difference models under rational expectations. *Econometrica 48*(5), 1305–1311.

Boppart, T., P. Krusell, and K. Mitman (2018, April). Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control 89*, 68–92.

Brumm, J., C. Krause, A. Schaab, and S. Scheidegger (2022). Sparse grids for dynamic economic models. Working paper, SSRN 3979412.

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.

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.

Den Haan, W. J. (1996). Heterogeneity, aggregate uncertainty, and the short-term interest rate. *Journal of Business & Economic Statistics 14*(4), 399–411.

Den Haan, W. J. (1997). Solving Dynamic Models with Aggregate Shocks and Heterogeneous Agents. *Macroeconomic Dynamics 1*(2), 355–386. Publisher: Cambridge University Press.

Den Haan, W. J. (2010, January). Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents. *Journal of Economic Dynamics and Control 34*(1), 79–99.

Den Haan, W. J. and A. Marcet (1990). Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics 8*(1), 31–34.

Devroye, L. (1981). Laws of the iterated logarithm for order statistics of uniform spacings. *Annals of Probability 9*(5), 860–867.

Duarte, V. (2018). Machine learning, the real business cycle, and asset pricing. *Working Paper*.

Elenev, V., T. Landvoigt, and S. Van Nieuwerburgh (2021). A Macroeconomic Model With Financially Constrained Producers and Intermediaries. *Econometrica 89*(3), 1361–1418.

Fair, R. and J. Taylor (1983). Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models. *Econometrica 51*(4), 1169–85.

Fernández-Villaverde, J. (2025). Deep learning for solving economic models. *Journal of Economic Literature*. Forthcoming.

Fernández-Villaverde, J., S. Hurtado, and G. Nuño (2023). Financial frictions and the wealth distribution. *Econometrica 91*(3), 869–901.

Ferreira, M. H., T. Haber, and H. Lee (2025). Dry Firms, Deep Recessions: Corporate Payouts and Aggregate Dynamics. *Working Paper*.

Gomes, F. and A. Michaelides (2007, 11). Asset pricing with limited risk sharing and heterogeneous agents. *The Review of Financial Studies 21*(1), 415–448.

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.

Han, J., Y. Yang, and W. E (2026). DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks. *Quantitative Economics*.

Heaton, J. and D. Lucas (2000). Portfolio choice and asset prices: The importance of entrepreneurial risk. *The Journal of Finance 55*(3), 1163–1198.

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

Judd, K. L., L. Maliar, and S. Maliar (2011, July). Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models: Approaches for solving dynamic models. *Quantitative Economics 2*(2), 173–210.

Juillard, M. (1996). Dynare : a program for the resolution and simulation of dynamic models with forward variables through the use of a relaxation algorithm. CEPREMAP Working Papers (Couverture Orange) 9602, CEPREMAP.

Kahou, M. E., J. Fernandez-Villaverde, J. Perla, and A. Sood (2021). Exploiting Symmetry in High-Dimensional Dynamic Programming. *Working Paper*, 51.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

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.

Lee, H. and K. Nomura (2024). The Spender of Last Resort: Global Equilibrium Dynamics under the Zero Lower Bound. *Working Paper*, 31.

Lee, H., P. Schnattinger, and F. Zanetti (2024). Endogenous Separations and Non-monotone Beveridge Curve Shifts. *Working Paper*, 31.

Levintal, O. (2018). Taylor projection: A new solution method for dynamic general equilibrium models. *International Economic Review 59*(3), 1345–1373.

Luetticke, R. (2021, April). Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics 13*(2), 1–25.

Maliar, L., S. Maliar, and P. Winant (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics 122*, 76–101.

Maliar, S., L. Maliar, and K. Judd (2011, February). Solving the multi-country real business cycle model using ergodic set methods. *Journal of Economic Dynamics and Control 35*(2), 207–228.

Marcet, A. (1988). Solving nonlinear stochastic models by parameterizing expectations. *Manuscript. Pittsburgh: Carnegie Mellon Univ*.

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

Meyn, S. P. and R. L. Tweedie (1993). *Markov Chains and Stochastic Stability*. Springer London.

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

Reiter, M. (2001, April). Recursive Solution Of Heterogeneous Agent Models. Computing in Economics and Finance 2001 167, Society for Computational Economics.

Reiter, M. (2009, March). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control 33*(3), 649–665.

Reiter, M. (2010). Solving the incomplete markets model with aggregate uncertainty by backward induction. *Journal of Economic Dynamics and Control 34*(1), 28–35. Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.

Schaab, A. and A. T. Zhang (2022). Dynamic programming in continuous time with adaptive sparse grids. Working paper, Columbia Business School.

Stokey, N. L., R. E. Lucas, and E. C. Prescott (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Winberry, T. (2018). A method for solving and estimating heterogeneous agent macro models. *Quantitative Economics 9*(3), 1123–1151.