Dancing on the Saddles: A Geometric Framework for Stochastic Equilibrium Dynamics[†]

Hanbaek Lee[‡]

University of Cambridge

March 2026

(click here for the latest version)

Abstract

I introduce *conditional saddle paths*—equilibrium trajectories under frozen exogenous states—extending saddle-path analysis to stochastic environments. Across representative-agent, heterogeneous-agent, and search-and-matching economies, this geometric object encodes two fundamental properties of stochastic equilibria. Its *slopes* characterize state dependence: shock responses differ across states if and only if conditional saddles differ in slope. Its *monotonicity* delivers dimension reduction: aggregate capital is a conditional sufficient statistic in Krusell and Smith (1998). In the search-and-matching setting, the saddle geometry serves as a diagnostic for state-contingent policy, revealing non-monotone effects of unemployment benefits on labor market volatility.

Keywords: Conditional saddle path, business cycles, state-dependent dynamics, sufficient statistics, heterogeneous agents, search and matching.

JEL codes: C62, D58, E32, J64.

[†]I am grateful for helpful comments and discussions from Martin M. Andreasen, Timo Boppart, Francisco Buera, Vasco Carvalho, Minsu Chang, Rachel Childers, Elisa Giannone, Matthias Hänsel, Dirk Krueger, Andrea Lanteri, Benjamin Larin, Danial Lashkari, Tim Lee, Benjamin Moll, Juan Carlos Parra-Alvarez, Tommaso Porzio, Itay Saporta-Eksten, and Yucheng Yang, as well as the seminar participants at Birkbeck, University of London, the University of Manchester, and Aarhus University. All errors are my own.

[‡]Email: hl610@cam.ac.uk

1 Introduction

In deterministic models, a single saddle path traces the economy's transition to steady state, and a phase diagram immediately conveys the qualitative dynamics. No comparable geometric framework exists for stochastic environments. Two obstacles stand in the way. First, stochastic switching creates *multiple* equilibrium loci—one for each exogenous state—so no single saddle path exists. Second, in heterogeneous-agent models the endogenous state is an infinite-dimensional wealth distribution, making geometric visualization infeasible and dimension reduction empirical rather than structural. This paper overcomes both obstacles by introducing a single geometric object—the *conditional saddle path*—whose slopes encode state dependence and whose monotonicity delivers dimension reduction.

The paper makes three contributions. First, I define conditional saddle paths as equilibrium trajectories under frozen exogenous states, computed under the decision rules of the stochastic equilibrium. These objects decompose business-cycle fluctuations into movements *along* a conditional saddle (endogenous propagation) and *across* conditional saddles (exogenous transitions), providing a stochastic analogue of deterministic phase diagrams. State-dependent shock responses have a precise geometric signature: they arise if and only if conditional saddle paths differ in slope. Even when impact responses are nearly state-independent, the *shape* of the impulse response—peak timing and persistence—can be highly state-dependent, driven by distance along the saddle rather than slope asymmetry. I develop these results in a canonical RBC model (Sections 2–3).

Second, I provide a dimension-reduction theorem. If an aggregate variable is strictly monotone along each conditional saddle, it is a *conditional sufficient statistic*: it exactly indexes all equilibrium allocations and prices within each frozen-regime spell. I prove that aggregate capital serves this role in Krusell and Smith (1998)—within each spell, the distributional state Φ is uniquely recoverable from K. Regime persistence ensures that this within-spell exactness governs most of the equilibrium dynamics, providing a structural foundation for scalar-state approximations. I also develop a perfect-foresight diagnostic that verifies the key monotonicity condition without solving the stochastic model (Section 4).

Third, I apply the framework to a search-and-matching economy (Section 5), where the saddle geometry serves as a diagnostic for state-contingent policy analysis.

Conditional saddles recover the Beveridge curve as a regime-specific locus, and the surplus compression behind the Shimer puzzle manifests as differential saddle steepness. Comparative statics on unemployment benefits yield a new result: volatility is hump-shaped in the benefit level. The saddle geometry diagnoses why. What governs volatility is the *separation between conditional steady states*, not slope asymmetry alone. At high benefit levels, conditional steady states converge despite widening slope gaps, leaving little room for unemployment to fluctuate between regimes.

Companion work extends the framework to New Keynesian models with an occasionally binding zero lower bound (Lee and Nomura, 2026) and to production network economies with endogenous linkage formation (Lee and Sun, 2026).

Related literature This paper contributes to three strands of literature in macroeconomics. The first is the literature studying global equilibrium dynamics and solution methods under aggregate uncertainty. The challenge of characterizing equilibrium dynamics in models with aggregate shocks has motivated extensive methodological development. Marcet (1988), Den Haan and Marcet (1990), and Krusell and Smith (1998) pioneered the use of bounded rationality approximations. In particular, Krusell and Smith (1998) discovered that a simple linear forecasting rule in aggregate capital achieves remarkable accuracy ($R^2 > 0.9999$) despite the infinite-dimensional state space. Subsequent work has sharpened accuracy and efficiency through diverse approaches (Den Haan, 1996; Reiter, 2010; Auclert et al., 2021; Fernández-Villaverde et al., 2023, among others).

The sequence-space Jacobian (SSJ) method of Auclert et al. (2021) provides a complementary perspective on dimension reduction. SSJ linearizes around a steady state; dimension reduction follows from the additive structure of linearized IRFs. By contrast, conditional saddle paths operate globally and nonlinearly: dimension reduction follows from monotonicity of an aggregate coordinate along the equilibrium manifold. Near a conditional steady state, the saddle's tangent connects to the SSJ representation. Away from it, the conditional saddle captures nonlinear dynamics—including state-dependent response shapes (Section 3)—outside the scope of first-order perturbation.

This paper differs from these computational contributions by providing geometric foundations for *why* dimension reduction works. Aggregate capital's sufficiency follows from geometric properties—nullcline invariance and monotonicity—rather than

numerical happenstance. The result also has computational implications: along each conditional saddle, aggregate capital provides a theoretically justified metric for the repeated transition method of Lee (2025).[1]

Second, this paper extends the geometric tradition of phase-diagram analyses in growth theory (Solow, 1956; Swan, 1956; Ramsey, 1928; Koopmans, 1963; Cass, 1965; Barro and Sala-i Martin, 1992) to stochastic environments. Conditional saddle paths provide a visual tool for understanding how microfounded frictions generate state-dependent dynamics, as documented empirically (Kaplan and Violante, 2014; Vavra, 2014; Berger and Vavra, 2015; Basu and Bundick, 2017; Bloom et al., 2018; Kaplan et al., 2018; Petrosky-Nadeau et al., 2018; Baley and Blanco, 2019; Pizzinelli et al., 2020; Berger et al., 2021; Winberry, 2021; Lee, 2026; Melcangi, 2024).

Finally, in the random dynamical systems literature (Arnold, 1998; Schenk-Hoppé, 1998; Schenk-Hoppé, 2001), the focus is on history-dependent stochastic saddle paths (Yannacopoulos, 2011), which vary with the realized shock sequence.[2] By contrast, conditional saddle paths are history-invariant geometric objects that provide economists with the direct analogue of deterministic phase diagrams (Kamihigashi, 2003, 2005).

2 Conditional saddle path

2.1 Definitions and assumptions

I consider a generic dynamic stochastic model where the corresponding recursive competitive equilibrium (RCE) is characterized by the following aggregate state S and the endogenous and exogenous laws of motion $(\Gamma_{\text{endo}}, \Gamma_{\text{exo}})$:

$$S = [\Phi, A] \tag{1}$$

[1]The idea of solving heterogeneous-agent models by freezing the aggregate state and analyzing transitions appears in several computational approaches, including Bourany (2018), Section 4.3 of a working paper version of Achdou et al. (2021), and Lee (2025). The present paper provides a theoretical foundation: conditional saddle paths formalize the underlying geometric objects, and the sufficiency results explain when dimension reduction is exact.

[2]In Random Dynamical Systems language, the conditional steady state is a deterministic realization of the random fixed point (Schenk-Hoppé and Schmalfuß, 2001), and the conditional saddle path is the invariant manifold of the frozen-regime dynamics.

where Φ is the endogenous aggregate state variable, and A is the exogenous aggregate state variable. The latter may be a multivariate vector that follows a stochastic process. I assume the exogenous aggregate law of motion Γ_{exo} is a Markov chain. For simplicity in the illustration, I assume Γ_{exo} is a two-state Markov chain where $A \in \{B, G\}$, and $\Gamma_{\text{exo}}(A'|A) > 0 \; \forall \, A', A$.

I consider a distributional state space $\mathcal{X}$, whose elements $\Phi \in \mathcal{X}$ summarize the cross-sectional distribution of idiosyncratic household states that are payoff-relevant for equilibrium (e.g., assets and employment/productivity types) together with any endogenous objects needed to evaluate equilibrium decision rules. Formally, one can take $\mathcal{X}$ to be a subset of a metric space of probability measures augmented, if needed, by a finite-dimensional vector of aggregate variables. Throughout, $\Gamma_{\text{endo}}(\cdot, A) : \mathcal{X} \to \mathcal{X}$ denotes the recursive equilibrium law of motion mapping the current distributional state into the next-period state under frozen regime A.

The central object of this paper formalizes a natural thought experiment. Suppose the economy is in state (Φ_0, A) today and the exogenous regime happens to remain at A indefinitely. The resulting equilibrium trajectory—computed under decision rules that fully internalize future regime changes—is the conditional saddle path.

Definition 1 (Conditional saddle path)**.**
Fix a regime $A \in \{B, G\}$ and an initial state $\Phi_0 \in \mathcal{X}$. Let $\{\Phi_t\}_{t \geq 0}$ denote the frozen-regime continuation under A, defined recursively by

$$\Phi_{t+1} = \Gamma_{\text{endo}}(\Phi_t, A), \qquad t \geq 0. \tag{2}$$

The conditional saddle path under A from Φ_0 is defined as the closure of the frozen-regime continuation:

$$\mathcal{M}(\Phi_0, A) := \overline{\{\Phi_t : t \geq 0\}}. \tag{3}$$

Throughout, the distributional state space is a metric space of probability measures $(\mathcal{X}, d)$, where d is compatible with weak convergence of distributions. The closure in Definition 1 is taken with respect to this metric. Based on the conditional saddle path, I define the conditional steady state as follows:

Definition 2 (Conditional steady state)**.**
Given (Φ_0, A), if the frozen-regime continuation $\{\Phi_t\}_{t \geq 0}$ converges, I denote its limit

by

$$\Phi^{cs}(\Phi_0, A) \;:=\; \lim_{t\to\infty} \Phi_t. \tag{4}$$

Definition 1 defines the conditional saddle path directly from the frozen-regime continuation. Fixing A makes the equilibrium law of motion deterministic on the distribution space: starting from Φ_0, the sequence $\{\Phi_t\}_{t\geq 0}$ is generated by repeated application of the endogenous transition operator $\Gamma_{\text{endo}}(\cdot, A)$. In this sense, $\mathcal{M}(\Phi_0, A)$ is the natural analogue of the saddle arm in deterministic saddle-path analysis. It isolates the equilibrium states visited (and their limit points) along the convergence dynamics under regime A.[3]

For the conditional saddle to serve as a useful geometric object—a one-dimensional curve that the economy traverses during each frozen-regime spell—the frozen-regime dynamics must converge to a unique limit and vary smoothly. The following assumption collects these regularity conditions.

Assumption 1 (Regularity of conditional saddle paths)**.**
Fix a regime $A \in \{B, G\}$ and an initial state $\Phi_0 \in \mathcal{X}$. Let $\{\Phi_t\}_{t\geq 0}$ be defined by $\Phi_{t+1} = \Gamma_{\text{endo}}(\Phi_t, A)$.

(i) (Unique existence) The conditional steady state $\Phi^{cs}(\Phi_0, A)$ uniquely exists on $\mathcal{M}(\Phi_0, A)$.

(ii) (Continuity) The map $\Gamma_{\text{endo}}(\cdot, A) : \mathcal{X} \to \mathcal{X}$ is continuous. Moreover, the aggregate variables (e.g., aggregate capital K and consumption C) vary continuously in Φ when restricted to $\mathcal{M}(\Phi_0, A)$.

Assumption 1 collects the regularity properties needed for the paper's geometric arguments. Part (i) requires that, conditional on a regime A and an initial state Φ_0, the frozen-regime continuation converges to a unique conditional steady state. This does not assert global uniqueness across all initial conditions, nor does it rule out multiplicity across regimes.[4] Part (ii) imposes continuity of the frozen-regime law of

[3] Strictly speaking, $\mathcal{M}(\Phi_0, A) = \overline{\{\Phi_t : t \geq 0\}}$ in discrete time is an orbit-closure, not necessarily a smooth manifold. I use saddle-path language to emphasize its one-dimensional, path-like role; in continuous time, the corresponding object is a one-dimensional invariant manifold; see Appendix K for the full continuous-time formulation.

[4] In particular, Proehl (2025) establishes global existence and uniqueness under parametric re-

motion and aggregate observables on the conditional saddle, ensuring that small perturbations of the distributional state produce small changes in aggregate outcomes. Assumption 1(i) thus restricts attention to convergent frozen-regime continuations, ruling out periodic orbits or divergent paths. In the models studied below, convergence follows from standard arguments: diminishing returns in the RBC benchmark and contractivity of the Bewley–Huggett–Aiyagari operator (Bewley, 1986; Huggett, 1993; Aiyagari, 1994) in the heterogeneous-household economy.

With this regularity, the conditional saddle admits a natural notion of position along the curve. Theorem 2 (Section 4) formalizes the key implication: if some aggregate equilibrium variable moves strictly one-way along this curve, it provides a valid coordinate for the entire conditional saddle, yielding an exact one-dimensional representation on $\mathcal{M}(\Phi_0, A)$.

2.2 Representative-agent economy: A canonical RBC

For illustrative purposes, I consider a representative household with temporal log utility. Given initial condition (a_0, A_0), the household maximizes lifetime utility under stochastic aggregate TFP A_t. I present the economy in recursive form and work with a recursive competitive equilibrium. Under standard regularity conditions, this formulation is equivalent to the sequential equilibrium; the sequential formulation is in Appendix A. The recursive form is as follows:

$$v(a; K, A) = \max_{c,a'} \quad \log(c) + \beta \mathbb{E} v(a'; K', A') \tag{5}$$

$$c + a' = a(1 + r(X)) + w(X) \tag{6}$$

$$a' \geq -\overline{a} \tag{7}$$

$$K' = \Gamma_{\text{endo}}(K, A), \quad A' \sim \Gamma_{\text{exo}}(A'|A), \tag{8}$$

where v is the household's value function; a is wealth at the beginning of a period; K is aggregate capital; A is aggregate TFP; Γ_{endo} is the law of motion for K. For illustrative purposes, I assume that TFP A follows a Markov-switching process between

strictions, whereas Assumption 1 (i) allows for global multiplicity as in Walsh and Young (2024). Nevertheless, conditional on a given initial state and a fixed regime, the equilibrium path is assumed to be uniquely determined.

the levels B and G, where $G > B$:

$$\Gamma_{\text{exo}} = \begin{bmatrix} \pi_{BB} & \pi_{BG} \\ \pi_{GB} & \pi_{GG} \end{bmatrix} \tag{9}$$

I consider competitive factor prices given a CRS Cobb–Douglas production function:

$$r(K, A) = A\alpha K^{\alpha-1} - \delta \tag{10}$$

$$w(K, A) = A(1-\alpha)K^{\alpha}, \tag{11}$$

where K is capital stock, which satisfies $K = a$ in equilibrium. The recursive competitive equilibrium (RCE, hereafter) is as follows:

Definition 3 (Recursive competitive equilibrium)**.**
$(c, a', v, r, w, \Gamma_{\text{endo}})$ *is a recursive competitive equilibrium if these functions*

1. *satisfy the individual optimality conditions;*
2. *clear factor markets, resulting in competitive prices;*
3. *satisfy the consistency:*

$$a'(K; K, A) = K' = \Gamma_{\text{endo}}(K, A) \tag{12}$$

Based on the endogenous law of motion Γ_{endo} defined in the recursive competitive equilibrium of Definition 3, I define the conditional K-nullcline.

Definition 4 (Conditional K-nullcline in RBC)**.**
Fix (K_0, A_0) *and a regime* $A \in \{B, G\}$*. Let* $\{K_t\}_{t\geq 0}$ *be the frozen-regime continuation* $K_{t+1} = \Gamma_{\text{endo}}(K_t, A)$*. The conditional* K*-nullcline along the continuation is*

$$\mathcal{N}(A; K_0) := \Big\{(K_t, C_t) : \ t \geq 0, K_{t+1} - K_t = 0\Big\}, \ C_t := C(K_t, A) \tag{13}$$

where C *is the aggregate equilibrium consumption variable. When* $\mathcal{N}(A; K_0)$ *is the graph of a function, define* $C_A^{K\text{cnull}}(\cdot; K_0)$ *by* $\mathcal{N}(A; K_0) = \{(K, C) : \ C = C_A^{K\text{cnull}}(K; K_0)\}$.

In the RBC model, the conditional K-nullcline can be characterized explicitly and is independent of the initial condition K_0.

Proposition 1 (Characterizing the conditional K-nullcline in RBC)**.**
The conditional nullclines of aggregate capital K for $A \in \{B, G\}$ are

$$C_A^{Kcnull}(K) = AK^{\alpha} - \delta K, \tag{14}$$

thus, $C_A^{Kcnull}(K)$ is independent of K_0.

Proof. From the stationary condition $\delta K = I$ and the resource constraint $Y = C + I$, $C = AK^{\alpha} - \delta K$ for each A, which depends on K but not on K_0. ■

In discrete time, capital can jump by a finite amount in one period, so the projected (K, C) trajectory can develop folds or loops. These pathologies are absent in continuous time, where capital moves continuously and cannot skip past the $\dot{K} = 0$ locus. Appendix K develops the continuous-time formulation and shows that monotonicity holds without any segment-crossing restriction. In the discrete-time setting, I rule out these pathologies with the following geometric restriction.

Assumption 2 (No segment crossing of the conditional K-nullcline)**.**
A straight line segment in (K, C) connecting (K_t, C_t) and (K_{t+1}, C_{t+1}) does not intersect the conditional K-nullcline $\forall\, A \in \{B, G\}$:

$$[(K_t, C_t), (K_{t+1}, C_{t+1})] \cap \mathcal{N}(A) = \emptyset, \ \forall t \geq 0. \tag{15}$$

Assumption 2 ensures consecutive points remain on the same side of the K-nullcline, so the nullcline gap $H_t := C_t - C_A^{Kcnull}(K_t)$ satisfies $H_t H_{t+1} > 0$ prior to convergence. Net investment therefore cannot reverse sign. Economically, this rules out single-period overshooting so dramatic that capital jumps past its stationary level and reverses direction—the discrete-time analogue of the continuous-time property that trajectories cannot cross nullclines.[5]

Although Assumption 2 is stated here for the RBC benchmark, it plays a central role in the heterogeneous-household economy (Section 4). The assumption can be verified from model primitives without solving the stochastic equilibrium; Sections 4.1 and 4.2 develop two complementary strategies.

[5]For related geometric restrictions in dynamic optimization, see Reddy et al. (2020) and, on threshold (Skiba) phenomena, Skiba (1978), Dechert and Nishimura (1983), and Wagener (2003).

Proposition 2 (K monotonicity in RBC)**.**
Fix $A \in \{B, G\}$. Under Assumptions 1 and 2, aggregate capital along the frozen-regime continuation converges to K_A^{cs} strictly monotonically.

Proof. Fix A and let $H_t := C_t - C_A^{Kcnull}(K_t)$ denote the nullcline gap. The resource constraint gives $\Delta K_t = -H_t$. By Assumption 2, consecutive points stay on the same side of the K-nullcline, so $H_t H_{t+1} > 0$ prior to convergence. Hence $\{\Delta K_t\}$ has constant sign, giving strict monotonicity. Convergence $K_t \to K_A^{cs}$ follows from Assumption 1. ■

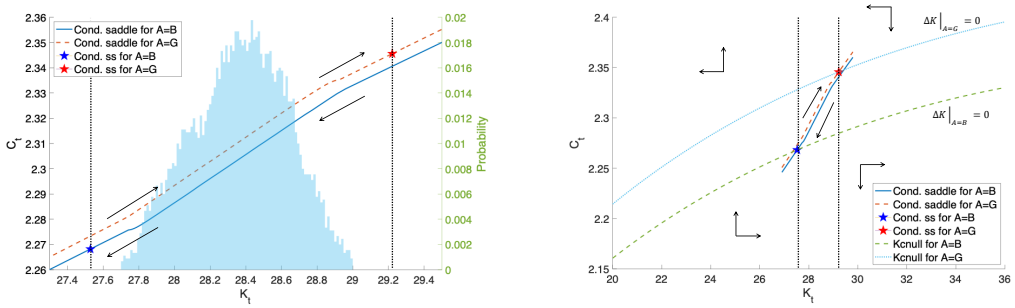
Figure 1 depicts the conditional saddle paths of the calibrated RBC model in the (K, C) plane under $A \in \{B, G\}$, solved globally via the repeated transition method of Lee (2025).[6] Panel (a) zooms into the conditional steady states and overlays a histogram of simulated (K_t, C_t) outcomes; panel (b) shows the global partition with conditional K-nullclines and direction arrows. In the long run, all equilibrium dynamics are confined to the region bounded by the two conditional saddles: within each spell, the economy converges along the active saddle, and regime switches move it across, so the entire ergodic distribution is nested within the saddle boundaries.

MIT shock on the saddle This decomposition yields a sharp characterization of unexpected transitory shocks (MIT shocks). Any unexpected transitory shock—whether TFP, monetary, fiscal, or otherwise—maps into a displacement along a conditional saddle path: there exists a shift along the conditional saddle that replicates the shock's impact, and all subsequent dynamics are along-path convergence. Consequently, the slope and curvature of the conditional saddle at the point of displacement fully determine the economy's propagation of the shock, including its persistence, amplification, and state dependence.

Generalized transition function The conditional saddle paths capture all equilibrium transitions—along saddles (endogenous propagation) and across them (exogenous regime shifts)—enabling generalized transition function (GTF) analysis (Lee, 2025), which encompasses generalized impulse response functions (Koop et al., 1996; Andreasen et al., 2017). The economy's response to a sequence of exogenous shocks is

[6] For each regime A, I hold productivity fixed for 2,500 periods and iterate the equilibrium law of motion from initial capital stocks on both sides of K_A^{cs}, recovering the full conditional saddle in the phase diagram.

Figure 1: Conditional saddle paths and the phase diagram



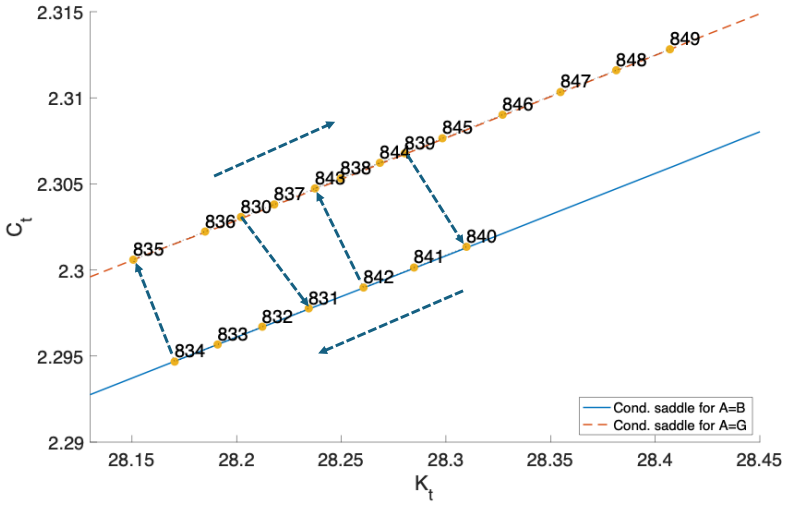
(a) Zoom-in (b) Zoom-out

Notes: Conditional K-nullclines ($\Delta K = 0$) partition the (K, C) plane. Direction arrows show frozen-regime convergence along each conditional saddle. Panel (a) zooms into the neighborhood of the conditional steady states shown in panel (b). The histogram plots the time-series distribution of aggregate capital.

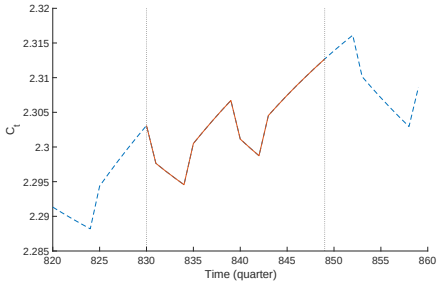
not confined to local dynamics around a steady state: its trajectory is sharply traced in the phase diagram as a sequence of along-saddle and across-saddle movements.

Figure 2 illustrates the equilibrium dynamics in the phase diagram (panel (a)) and in the time domain (panels (b) and (c)) for a 20-quarter subsample. Negative TFP shocks produce downward jumps in consumption across conditional saddle paths, followed by endogenous decline along them; positive shocks reverse the pattern. Compared to time-domain representations, the phase diagram illustrates the full equilibrium dynamics more concisely in a single figure and enables immediate consideration of counterfactual scenarios. For the GIRF illustration, I replace the two-state Markov chain with a finely discretized AR(1) process for log productivity: $\ln A' = \rho \ln A + \epsilon$, $\epsilon \sim N(0, \sigma^2)$. In this case, there is a continuum of frozen-A conditional saddles indexed by A, and the GIRF path moves both along a given conditional saddle (endogenous propagation) and across the foliation as A evolves.

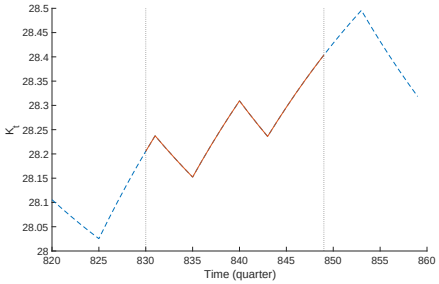
Figures 3 and 4 illustrate post-shock dynamics following a positive TFP shock. Upon impact, consumption jumps upward across conditional saddles. In subsequent periods, forces along and across saddles jointly generate a bow-shaped trajectory in the phase diagram: capital exhibits a hump-shaped response as the shock mean-reverts, while consumption initially rises then falls as TFP crosses the threshold where the direction of along-saddle adjustment reverses.



(a) On the saddles



(b) C - time domain



(c) K - time domain

Figure 2: Equilibrium dynamics: saddle vs. time domain

Notes: Stochastic equilibrium dynamics on the conditional saddle paths (panel (a)) and in the time domain (panels (b) and (c)). Sample: 20 quarters (830–849).

Comparison with the perfect-foresight saddles In the *perfect-foresight (PF)* economy, agents believe the current regime persists forever; the PF steady states solve $1 = \beta(1 - \delta + \alpha A(K_A^{pf})^{\alpha-1})$. Along the *conditional saddle (CS)*, the regime is held fixed ex post, but agents correctly anticipate regime switches under Π, so the conditional steady states K_A^{cs} are pinned by the probability-weighted Euler equations.[7] PF and CS economies share the same K-nullcline (Proposition 3), but differ in their consumption dynamics.

Proposition 3 (K-nullcline invariance over beliefs)**.**

[7]Specifically, $\beta\pi_{AA}(1 + \alpha A(K_A^{cs})^{\alpha-1} - \delta) + \beta\pi_{AA'}\left(c(K_A^{cs}, A')/c_A^{cs}\right)^{-1}(1 + \alpha A'(K_A^{cs})^{\alpha-1} - \delta) = 1$ for each $A \in \{B, G\}$, where $A' \neq A$.

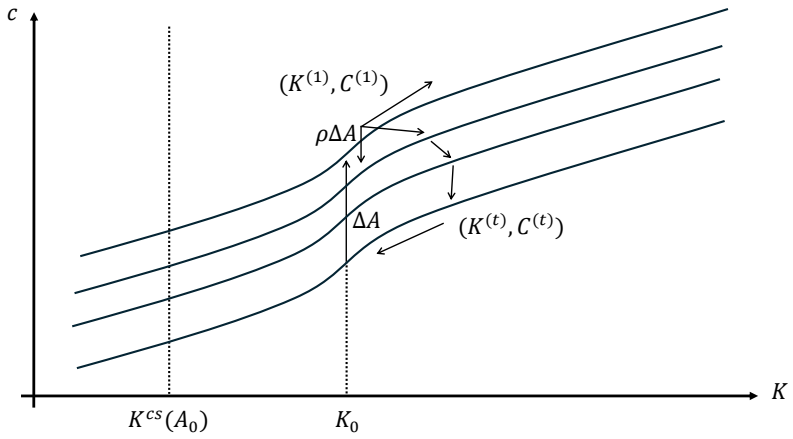


Figure 3: Generalized impulse response function (GIRF) at (K_0, A_0)

Notes: Generalized impulse responses of consumption and capital to a positive TFP shock at initial state (K_0, A_0), where A_0 is the pre-shock productivity level. Upon impact, the economy jumps across conditional saddles to the $A_0 + \Delta$ foliation and subsequently propagates along and across saddles as TFP mean-reverts.

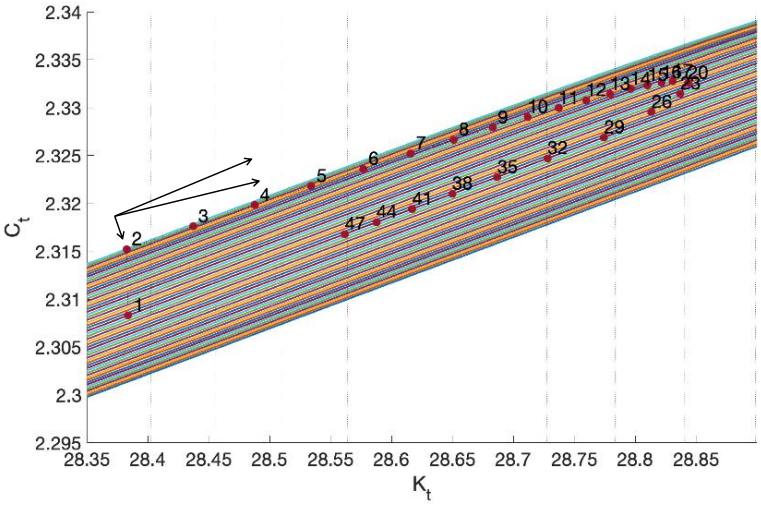
Conditional K-nullclines are identical between the RBC model with aggregate uncertainty and the perfect-foresight counterpart.

Proof.
From the stationary condition $\delta K = I$ and the aggregate resource constraint $Y = C + I$, equation (14) is immediate. Neither condition involves expectations or belief structures, so the nullcline is identical whether or not there is aggregate uncertainty. ■

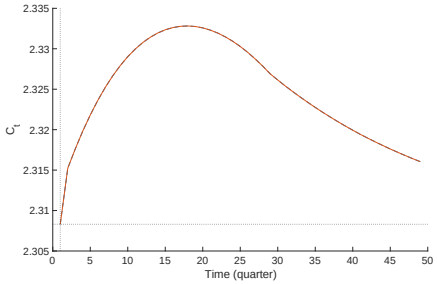
This invariance implies that PF–CS differences are driven entirely by consumption dynamics, not by the feasibility locus.[8] Proposition 4 formalizes this via steady-state orderings. The same invariance extends to the heterogeneous-household environment (Section 4), where it underlies both K-sufficiency and the perfect-foresight diagnostic (Section 4.1).

Proposition 4 (Aggregate uncertainty and the conditional steady states)**.**

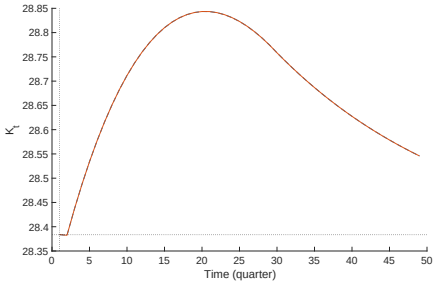
[8]There is no consumption nullcline under aggregate uncertainty; the Euler equations at the conditional steady states play that role.



(a) On the saddles



(b) C - time domain



(c) K - time domain

Figure 4: Impulse responses to a positive TFP shock: saddle vs. time domain

Notes: Impulse responses on conditional saddle paths (panel (a)) and in the time domain (panels (b) and (c)). Sample: 50 quarters.

The following inequalities hold:

$$K_B^{cs} \;<\; K_B^{pf} \;<\; K_G^{pf} \;<\; K_G^{cs}, \qquad c_B^{cs} \;<\; c_B^{pf} \;<\; c_G^{pf} \;<\; c_G^{cs}. \tag{16}$$

Proof.
See Appendix C. ■

PF steady states are nested within the CS steady states: hedging against regime switches pushes the conditional steady states apart (Deaton, 1991; Carroll, 1997). Conversely, the conditional saddle *paths* are nested within the PF paths: at any given K, hedging compresses the consumption gap across regimes. The two nestings reflect

the same mechanism: compressed consumption responses cumulate into amplified long-run capital divergence (Figure 5).

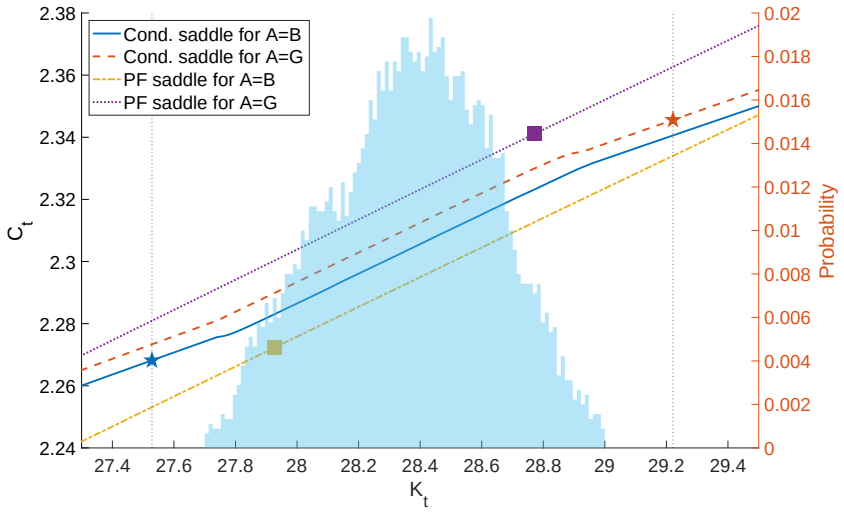


Figure 5: Conditional saddle path comparison: with and without uncertainty

Notes: Conditional saddle paths for $A = B$ (solid) and $A = G$ (dashed) and perfect-foresight saddle paths for $A = B$ (dash-dotted) and $A = G$ (dotted) under standard quarterly calibration.

Conditional boundary condition I work directly with the recursive formulation, which implicitly imposes a stochastic analogue of the deterministic transversality condition.[9] Given conditional saddle paths, the natural counterpart is *conditional transversality*:

$$\lim_{t\to\infty} \beta^t u'(c_t) K_t = 0 \quad \text{under a frozen aggregate regime } A \in \{B, G\} \qquad (17)$$
$$\text{(along the conditional saddle for each } A \in \{B, G\}).$$

Condition (17) rules out explosive continuations along each conditional saddle. Under standard assumptions, it is equivalent to the usual stochastic transversality conditions (Kamihigashi, 2003, 2005) and selects the same bounded RCE. The conditional formulation makes explicit that the boundary condition operates regime-by-regime on $\mathcal{M}(\Phi_0, A)$.

[9]This condition is imposed at the aggregate level and should be distinguished from individual no-Ponzi constraints.

Economies without conditional saddle paths A conditional saddle requires a predetermined endogenous state. Without one, the economy jumps each period to the allocation implied by the current exogenous state, so the conditional "saddle" under frozen A collapses to a singleton. I call such an economy *endogenously stateless*.[10] Models with predetermined states—capital, habits, interest-rate smoothing, or distributional states in HANK—are endogenously stateful, and conditional saddle paths exist under mild regularity conditions.

3 State dependence

In this section, I analyze nonlinear shock responsiveness through the lens of conditional saddle paths. In any stochastic dynamic model that admits conditional saddle paths, the response of an aggregate variable to an exogenous shock is represented by a shift across different conditional saddle paths. If and only if all conditional saddle paths are parallel along the endogenous state, the response is state-independent.

A large empirical literature documents that macroeconomic responses to shocks depend on the phase of the business cycle—recessions amplify negative shocks, expansions dampen them (Neftci, 1984; Auerbach and Gorodnichenko, 2012; Tenreyro and Thwaites, 2016). The following result shows that state dependence has a precise geometric signature: it arises if and only if the conditional saddle paths are not parallel, so the entire phenomenon reduces to comparing slopes across regimes.

Theorem 1 (State-(in)dependence as a geometric condition)**.**
Fix (Φ_0, A_0) and let $\mathcal{M}(\Phi_0, A_0)$ denote the conditional saddle path under the frozen regime A_0. Let $g(\Phi, A)$ be an aggregate equilibrium object (e.g. consumption) defined for $(\Phi, A) \in \mathcal{M}(\Phi_0, A_0) \times \{A_0, A_1\}$. Define the impact gap between regimes A_1 and A_0 at state Φ by

$$\Delta_g(\Phi; A_1, A_0) := g(\Phi, A_1) - g(\Phi, A_0). \tag{18}$$

Then the following are equivalent:

[10] The textbook three-equation New Keynesian model (Gali, 2008) under determinacy is endogenously stateless: $(x_t, \pi_t, i_t) = \Psi s_t$ for exogenous state $s_t := (r_t^n, u_t, \nu_t)$, so freezing s_t produces no nontrivial dynamics.

(i) (State-independent gap) $\Delta_g(\Phi; A_1, A_0)$ *is constant on* $\mathcal{M}(\Phi_0, A_0)$*, i.e. there exists* $c \in \mathbb{R}$ *such that*

$$g(\Phi, A_1) - g(\Phi, A_0) = c \qquad \forall\, \Phi \in \mathcal{M}(\Phi_0, A_0). \tag{19}$$

(ii) (Vertical-translation geometry) Viewed as subsets of $\mathcal{X} \times \mathbb{R}$*,*

$$\mathcal{G}_{A_j} := \{(\Phi, g(\Phi, A_j)) : \ \Phi \in \mathcal{M}(\Phi_0, A_0)\}, \qquad j \in \{0, 1\}, \tag{20}$$

satisfy $\mathcal{G}_{A_1} = \mathcal{G}_{A_0} + (0, c)$*, i.e.* $\mathcal{G}_{A_1}$ *is a constant vertical translation of* $\mathcal{G}_{A_0}$*.*

Proof. $(i) \Rightarrow (ii)$: $g(\Phi, A_1) = g(\Phi, A_0) + c$ for all Φ gives $\mathcal{G}_{A_1} = \mathcal{G}_{A_0} + (0, c)$. $(ii) \Rightarrow (i)$: $\mathcal{G}_{A_1} = \mathcal{G}_{A_0} + (0, c)$ requires $g(\Phi, A_1) = g(\Phi, A_0) + c$ for each Φ. ■

Sharp state independence is a knife-edge. Differential tilt across conditional saddles—whether from real (Winberry, 2021; Lee, 2025), financial (Melcangi, 2024), or labor market frictions (Petrosky-Nadeau et al., 2018; Pizzinelli et al., 2020)—generates state-dependent shock responses. Theorem 1 provides the qualitative characterization; applied work requires a quantitative measure of how far the economy is from the state-independent benchmark.

Remark 1 (Slope ratio as a state-dependence diagnostic)**.**
Near the conditional steady state K_A^{cs}*, the local slope of the conditional saddle in* (K, C) *space is governed by the stable eigenvalue* $\lambda_A^s \in (0, 1)$ *of the linearized frozen-regime dynamics and the nullcline slope* $(C_A^{Kcnull})'(K_A^{cs})$*. Define the slope ratio across regimes,*

$$\mathcal{R}(K) \ := \ \frac{dC/dK|_{\text{saddle},B}}{dC/dK|_{\text{saddle},G}},$$

evaluated at a common K*. Then* $\mathcal{R} = 1$ *corresponds to parallel saddles (state-independent impacts by Theorem 1), while* $|\mathcal{R} - 1|$ *measures the degree of endogenous state dependence.*

This diagnostic connects the geometric framework to the empirical state-dependence literature: $\mathcal{R}$ is a model-implied moment that could in principle be identified from state-dependent local projections (Auerbach and Gorodnichenko, 2012; Tenreyro and Thwaites, 2016).

An extension: asymmetric adjustment cost I modify the budget constraint by adding a quadratic wealth adjustment cost $\mathcal{C}(a', a) = \frac{\widetilde{\mu}}{2}\left(\frac{a'-a}{a}\right)^2 a$, with $\widetilde{\mu} = \mu_+$ for positive investment and $\widetilde{\mu} = \mu_-$ for negative investment ($\mu_+ > \mu_-$).[11] Figure 6 plots the resulting conditional saddle paths. The $A = B$ saddle is substantially steeper than the $A = G$ saddle ($\mathcal{R} \approx 1.3$ at the ergodic mean capital). A one-standard-deviation TFP shock generates a larger consumption response when capital is low (3.42%) than when it is high (2.42%).

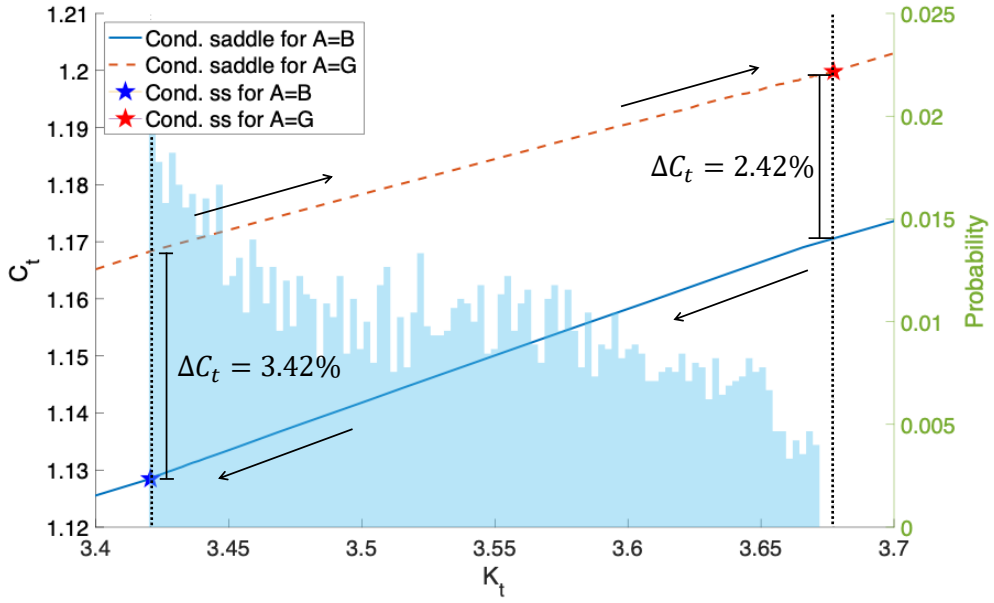


Figure 6: Differently tilted conditional saddle paths — endogenous state dependence

Notes: Conditional saddle paths for $A = B$ (solid) and $A = G$ (dashed) under asymmetric wealth adjustment cost ($\mu_+/\mu_- = 4$). The $A = B$ saddle is steeper, generating state-dependent shock responses.

This geometric representation formalizes the intuition that "climbing up is difficult, falling down is easy" (Petrosky-Nadeau et al., 2018). By Theorem 1, state independence holds if and only if conditional saddles are parallel. Any deviation—whether from financial (Melcangi, 2024), labor market (Petrosky-Nadeau et al., 2018; Pizzinelli et al., 2020), or capital-adjustment frictions—manifests as differential tilt. This provides a structural interpretation of evidence that macroeconomic policy trans-

[11] The adjustment-cost parameters are not calibrated. For illustrative purposes, I set $\mu_+/\mu_- = 4$ ($\mu_- = 1.22$, $\mu_+ = 4.88$). I use an annual calibration ($\beta = 0.96$, $\delta = 0.10$, $\alpha = 0.33$) and a two-state TFP process $A \in \{0.98, 1.02\}$ with persistence 0.90.

mits asymmetrically over the cycle (Auerbach and Gorodnichenko, 2012; Tenreyro and Thwaites, 2016; Jo and Zubairy, 2025). The slope ratio $\mathcal{R}$ of Remark 1 is a scalar diagnostic of state dependence, computable for any globally solved model.

Shape of the response function Even in the standard RBC, where $\mathcal{R} \approx 1$, the *shape* of the response trajectory is highly state dependent. A positive productivity shock displaces the economy across saddles, and the subsequent trajectory depends on K_0's position relative to K_G^{cs}. Far from K_G^{cs}, the marginal product of capital is high, tilting the intertemporal tradeoff toward accumulation and producing a pronounced, late consumption hump (Barro and Sala-i Martin, 1992). Near K_G^{cs}, there is little room to accumulate; the across-saddle reversion dominates quickly and the peak arrives earlier. This channel is orthogonal to the slope asymmetry of Theorem 1. Slope differences govern *impact* magnitudes (cross-saddle dimension); position-dependent convergence speed governs response *shapes* (along-saddle dimension).

Crucially, this state-dependent shape does *not* arise from nonlinearity of the saddle path itself. Even in the log-linearized economy—where the conditional saddle is a straight line—peak timing varies with K_0. The mechanism is the *speed of along-saddle convergence*. Far from K_G^{cs}, the capital gap is large relative to the shock-induced across-saddle displacement, so the along-saddle force (capital accumulation) dominates the across-saddle force (TFP reversion) for many periods before the balance tips, producing a later and more pronounced consumption hump. The shape asymmetry is thus a property of the *race between two convergence rates*—the endogenous along-saddle rate λ and the exogenous across-saddle rate ρ—and how initial distance from steady state weights the former relative to the latter. Nonlinearity (concavity of $f(K)$) amplifies the effect quantitatively but is not its source.

The mechanism can be made precise via the log-linearized dynamics. Near the non-conditional steady state, the Blanchard–Kahn decomposition (Blanchard and Kahn, 1980) on the saddle branch yields

$$\hat{c}_t \;=\; \underbrace{\mathcal{P}(\hat{k}_0)}_{\text{along saddle}}\;\lambda^t \;+\; \underbrace{\mathcal{Q}}_{\text{across saddle}}\;\rho^t, \tag{21}$$

where $\hat{c}_t$ and $\hat{k}_0$ are log-deviations from steady state, $\lambda \in (0,1)$ is the stable eigenvalue of the capital dynamics, $\rho \in (0,1)$ is the TFP persistence, and $\varepsilon > 0$ is the initial shock size. The coefficient $\mathcal{P} = \psi_k[\hat{k}_0 + \theta\varepsilon/(\lambda - \rho)]$ depends on initial capital

through the saddle slope $\psi_k > 0$ and the capital-TFP elasticity $\theta > 0$. The coefficient $\mathcal{Q} = [\psi_a - \psi_k\theta/(\lambda - \rho)]\varepsilon$ does not. The hump shape arises from the tension between two decaying exponentials. The λ^t term captures along-saddle propagation—the endogenous capital accumulation effect—while the ρ^t term captures the direct TFP reversion. The peak occurs where these forces balance. A lower initial capital stock raises $|\mathcal{P}|$ through concavity of $f(K) = AK^\alpha$, tilting the balance later in time.

Proposition 5 (State-dependent peak timing)**.**
Consider the log-linearized RBC dynamics near the non-stochastic steady state. Suppose TFP follows deterministic mean-reversion $\hat{a}_t = \rho^t\varepsilon$ with $\varepsilon > 0$ and $\rho \in (0, 1)$. Let $\lambda \in (0, 1)$ denote the stable eigenvalue of the capital dynamics, with $\lambda > \rho$. Write the consumption path as in (21) and suppose $\mathcal{P}(\hat{k}_0) > 0$.[12] *If the consumption response admits a peak at $t^* > 0$, then*

$$\frac{\partial t^*}{\partial \hat{k}_0} = \frac{-\psi_k}{\mathcal{P}(\hat{k}_0) \cdot \ln(\lambda/\rho)} < 0. \tag{22}$$

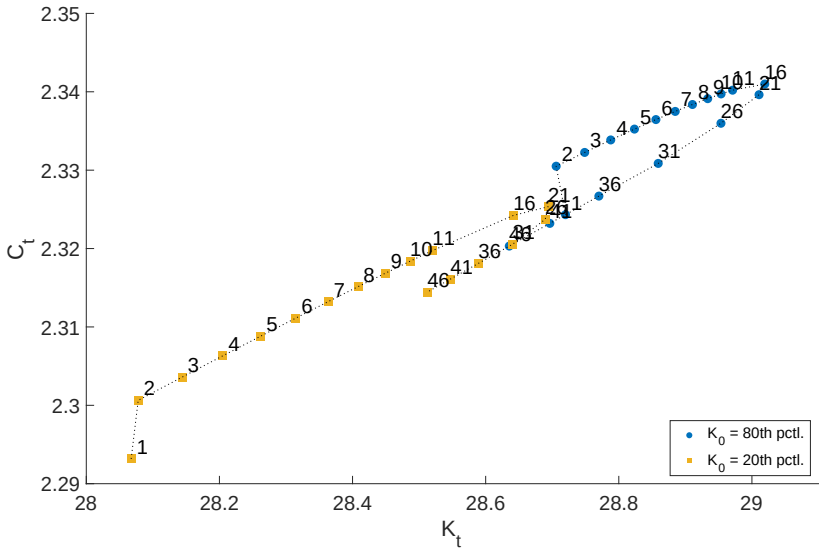
The consumption peak arrives strictly later when initial capital is farther below the steady state.

Proof.
See Appendix D. ■

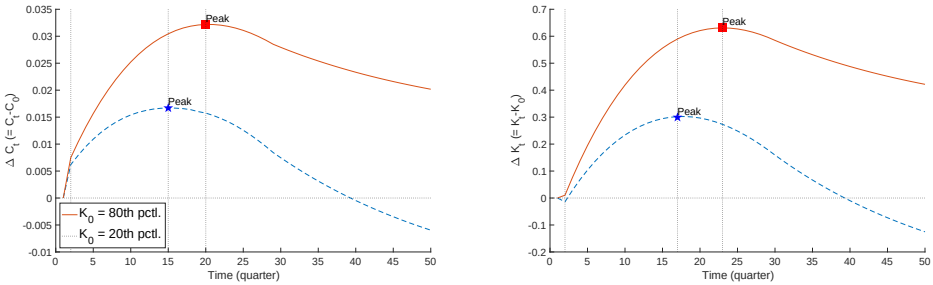
The condition $\lambda > \rho$ holds under standard quarterly RBC calibrations ($\lambda \approx 0.96$, $\rho \leq 0.95$). Proposition 5 is local by construction; Figure 7 confirms that the mechanism operates globally across the ergodic capital distribution.

Figure 7 confirms the logic of Proposition 5 quantitatively from the 20th and 80th percentiles of the ergodic capital distribution. From low K_0, both C and K rise longer before reversion dominates, producing a later and larger consumption peak. From high K_0, the peak arrives earlier—the same saddle geometry, the same shock, the same preferences; only the inherited capital stock differs.

[12] This requires $\hat{k}_0 > -\theta\varepsilon/(\lambda - \rho)$, i.e., initial capital is not farther below steady state than the long-run capital response to the shock. Under standard quarterly calibrations, the ergodic distribution satisfies this with substantial margin.



(a) On the saddles



(b) C — time domain

(c) K — time domain

Figure 7: Shape of the impulse responses: low K_0 vs. high K_0

Notes: Impulse responses to a positive productivity shock from K_0 at the 20th percentile and 80th percentile of the ergodic distribution. Panel (a): post-shock trajectories in (K, C) space on the conditional saddle paths; numbers indicate quarters since impact. Panels (b) and (c): time-domain responses of consumption and capital; vertical markers indicate the peak quarter. Sample: 50 quarters.

4 Heterogeneous-household economy and dimension reduction

In a heterogeneous-household economy, the natural state variable is the entire cross-sectional distribution of wealth—an infinite-dimensional object. Classical aggregation (Gorman, 1961) reduces dimensionality by restricting preferences, but these restrictions fail under uninsurable idiosyncratic risk and borrowing constraints. Yet practi-

tioners routinely approximate the state by a single number, aggregate capital. The conditional saddle framework explains why this works without preference restrictions: if the distribution evolves monotonically along a one-dimensional equilibrium path, then aggregate capital is an exact coordinate for the distribution along that path. This section formalizes the argument.

I consider a continuum of ex-ante homogeneous households of unit measure. The recursive formulation of the households' problem is as follows:

$$v(a, z; \Phi, A) = \max_{c, a'} \quad \log(c) + \beta \mathbb{E} v(a', z'; \Phi', A') \tag{23}$$

$$c + a' = a(1 + r(X)) + w(X)z \tag{24}$$

$$a' \geq 0 \tag{25}$$

$$\Phi' = \Gamma_{\text{endo}}(\Phi, A), \quad A' \sim \Gamma_{\text{exo}}(A'|A). \tag{26}$$

The problem is the same as in the representative-household economy except for (i) uninsurable idiosyncratic labor productivity, which follows a Markov process $z \sim \Gamma_z(z'|z)$; (ii) inclusion of the distribution of individual states Φ in the aggregate endogenous state; and (iii) the corresponding change in the law of motion for the aggregate endogenous state. The RCE is defined as in Krusell and Smith (1998).

The model features two exogenous stochastic processes: idiosyncratic productivity and aggregate TFP. Accordingly, there are two layers of conditional saddle paths: individual and aggregate. The individual saddle path has its own cross-sectional implications, which are analyzed separately in Appendix B. Here I focus on the aggregate conditional saddle path under stochastic TFP, following Krusell and Smith (1998).[13] Consistent with the notation in Section 2.2, I denote $K = K(\Phi)$ and $C = C(\Phi, A)$ as aggregate capital and consumption.

The conditional K-nullcline (Definition 4) extends directly to the heterogeneous-household setting by replacing the scalar state K_0 with the distributional state Φ_0: the nullcline is $\mathcal{N}(\Phi_0, A) := \{(K_t, C_t) : \ t \geq 0, \ K(\Gamma_{\text{endo}}(\Phi_t, A)) - K(\Phi_t) = 0\}$, where $K_t := K(\Phi_t)$ and $C_t := C(\Phi_t, A)$.

Krusell and Smith (1998) posits an endogenous law of motion that tracks aggregate capital K rather than the full distribution Φ: $\log K' = b_0(A) + b_1(A) \log K$, where

[13] In the original model of Krusell and Smith (1998), exogenous individual labor supply co-moves with aggregate TFP. All results are unaffected by this feature, but for expositional brevity I assume labor supply is exogenously fixed.

b_0 and b_1 are state-dependent coefficients. This formulation embeds two distinct assumptions: (i) aggregate capital K suffices to summarize the endogenous aggregate state, and (ii) K follows a log-linear law of motion.

Since Krusell and Smith (1998), it has remained an open question whether the striking accuracy of K-based forecasting ($R^2 > 0.9999$) reflects an approximation that happens to work well or a deeper structural property. Using the conditional-saddle framework, this paper establishes that (i) is exact within each frozen-regime spell. Aggregate capital K uniquely indexes the distributional state on each conditional saddle path, so conditioning on K entails no loss of information for within-spell dynamics. The across-spell bridge to global accuracy—why a single forecasting rule $K' = f(K, A)$ works across all spells—is formally bounded by Proposition 8 below. The across-spell distributional discrepancy vanishes as regime persistence grows, providing a geometric foundation for the near-perfect computational accuracy.

By contrast, (ii) remains a parametric approximation—the true law of motion for K is generally non-Markovian and need not be log-linear—but this is a matter of functional-form choice, not a limitation on the state space. The practical implication is that improving accuracy primarily requires richer functional forms for the law of motion, not additional distributional moments as state variables.

The following theorem provides the first step toward establishing the exact conditional sufficiency of K. The core insight is simple: if the economy converges along a one-dimensional equilibrium path and some aggregate variable moves strictly in one direction throughout, its current value tells you exactly where the economy is along that path. No two distinct distributional states can share the same capital level, because the economy passes through each level exactly once.

Theorem 2 (Monotone aggregate variable as a coordinate)**.**
Fix (Φ_0, A_0) *and* $A \in \{B, G\}$. *Let* $\{\Phi_t\}_{t\geq 0}$ *be the frozen-regime continuation* $\Phi_{t+1} = \Gamma_{\text{endo}}(\Phi_t, A)$ *and let* $\Phi^{cs} := \lim_{t\to\infty} \Phi_t$. *Let* $e : \mathcal{M}(\Phi_0, A) \to \mathbb{R}$ *be an aggregate equilibrium variable. Suppose that the scalar sequence* $\{e(\Phi_t)\}_{t\geq 0}$ *is strictly monotone and converges to* $e(\Phi^{cs})$. *Then* e *uniquely indexes states on the conditional saddle:*

$$e(\Phi) = e(\Phi') \implies \Phi = \Phi' \qquad \forall\, \Phi, \Phi' \in \mathcal{M}(\Phi_0, A). \tag{27}$$

Proof. Define $\psi_A : \mathbb{N} \cup \{\infty\} \to \mathcal{M}(\Phi_0, A)$ by $\psi_A(t) := \Phi_t$ for $t \in \mathbb{N}$ and $\psi_A(\infty) := \Phi^{cs}$.

By definition of $\mathcal{M}(\Phi_0, A)$ and convergence $\Phi_t \to \Phi^{cs}$,

$$\mathcal{M}(\Phi_0, A) = \{\Phi_t : t \geq 0\} \cup \{\Phi^{cs}\}, \tag{28}$$

so ψ_A is surjective. Strict monotonicity of $\{e(\Phi_t)\}$ implies that $e(\Phi_t) \neq e(\Phi_\tau)$ for $t \neq \tau$, hence $\Phi_t \neq \Phi_\tau$ for $t \neq \tau$; therefore ψ_A is injective on $\mathbb{N}$. Moreover, since $\{e(\Phi_t)\}$ is strictly monotone and converges to $e(\Phi^{cs})$, it also holds that $e(\Phi_t) \neq e(\Phi^{cs})$ for all finite t, so ψ_A is injective on $\mathbb{N} \cup \{\infty\}$. Thus ψ_A is a bijection between $\mathbb{N} \cup \{\infty\}$ and $\mathcal{M}(\Phi_0, A)$.

Now let $\Phi, \Phi' \in \mathcal{M}(\Phi_0, A)$. Pick $s, s' \in \mathbb{N} \cup \{\infty\}$ such that $\Phi = \psi_A(s)$ and $\Phi' = \psi_A(s')$. If $e(\Phi) = e(\Phi')$, then $e(\psi_A(s)) = e(\psi_A(s'))$. Since $e \circ \psi_A$ is injective, it follows that $s = s'$, hence $\Phi = \Phi'$. Therefore e is injective on $\mathcal{M}(\Phi_0, A)$, and an inverse map φ_A exists on $e(\mathcal{M}(\Phi_0, A))$. ■

Consequently, there exists a unique inverse map $\varphi_A : e(\mathcal{M}(\Phi_0, A)) \to \mathcal{M}(\Phi_0, A)$ such that $\varphi_A(e(\Phi)) = \Phi$ on $\mathcal{M}(\Phi_0, A)$, and any equilibrium object restricted to $\mathcal{M}(\Phi_0, A)$ can be written as a (single-valued) function of the scalar coordinate e.

The geometric intuition is that the conditional saddle is a one-dimensional curve of equilibrium states, and e acts as a "progress meter" along this curve: it moves strictly in one direction and never reverses. If two distinct states $\Phi \neq \Phi'$ shared the same value $e(\Phi) = e(\Phi')$, both would be at the same "progress" position, yet their subsequent histories could not merge on a one-dimensional path—a contradiction. Hence e uniquely labels states on the conditional saddle.

The following proposition shows that the conditional K-nullclines coincide with those of the representative-agent model.

Proposition 6 (Conditional K-nullclines of Krusell and Smith (1998))**.**
The conditional K-nullclines of the heterogeneous-household model are identical to those of the representative-household model and invariant over the initial distribution Φ_0:

$$C_A^{Kcnull}(K) = AK^\alpha - \delta K, \text{ for } A \in \{B, G\}. \tag{29}$$

Proof.
As in Proposition 3, the stationary condition $\delta K = I$ and the aggregate resource constraint $Y = C + I$ immediately imply the form of the conditional nullclines. ■

The nullcline is pinned by the aggregate resource constraint—an accounting identity indifferent to *who* saves. At any K, output AK^α and depreciation δK are determined by technology alone, so the nullcline $C = AK^\alpha - \delta K$ is distribution-free. This is what enables dimension reduction: the "obstacle" governing the sign of capital accumulation does not move as the distribution evolves along the conditional saddle.

The invariance is specific to the neoclassical production structure with exogenous labor supply. Under separable CRRA preferences with wealth effects on labor, labor supply introduces a distributional channel into the resource constraint. The nullcline then acquires distributional dependence, blocking the analytical route to K-sufficiency (though the perfect-foresight diagnostic of Section 4.1 remains applicable).

Proposition 7 (K monotonicity and injectivity)**.**
Fix $A \in \{B, G\}$. Under Assumptions 1 and 2, aggregate capital converges to K_A^{cs} strictly monotonically along the frozen-regime continuation on $\mathcal{M}(\Phi_0, A)$. Consequently, K is injective on $\mathcal{M}(\Phi_0, A)$.

Proof.
By Proposition 6, the conditional K-nullcline takes the same form as in the RBC benchmark. The proof of strict monotonicity therefore follows that of Proposition 2 verbatim, with the RBC nullcline replaced by $C_A^{Kcnull}(K) = AK^\alpha - \delta K$. Since $\{K_t\}$ is strictly monotone and converges to K_A^{cs}, applying Theorem 2 with $e(\Phi) = K(\Phi)$ yields $K(\Phi) = K(\Phi') \Longrightarrow \Phi = \Phi'$ for all $\Phi, \Phi' \in \mathcal{M}(\Phi_0, A)$. ■

Proposition 6 plays a key role in the proof above because it makes the K-nullcline *sign-determining*: the graph $C = C_A^{Kcnull}(K)$ partitions the (K, C) plane into $\Delta K > 0$ and $\Delta K < 0$ regions by $\Delta K = C_A^{Kcnull}(K) - C$. Assumption 2 rules out pathological discrete-time *crossings* of this partition between successive dates by requiring that the line segment connecting (K_t, C_t) and (K_{t+1}, C_{t+1}) does not intersect the nullcline. Together, these two ingredients imply that net investment cannot reverse sign along the conditional saddle, yielding strictly monotone capital dynamics and hence an exact one-dimensional representation indexed by K.

Proposition 7 implies that, conditional on (Φ_0, A), the mapping $K : \mathcal{M}(\Phi_0, A) \to \mathbb{R}$ is one-to-one. Hence there exists an inverse map $\varphi_{\Phi_0,A}$ such that $\Phi = \varphi_{\Phi_0,A}(K(\Phi))$ for all $\Phi \in \mathcal{M}(\Phi_0, A)$.[14] Equivalently, any equilibrium object restricted to $\mathcal{M}(\Phi_0, A)$

[14]This is a path-specific notion of sufficiency: for a fixed (Φ_0, A), K uniquely indexes states on the entire conditional saddle $\mathcal{M}(\Phi_0, A)$ (i.e. globally *along the transition path*).

admits an exact single-valued representation as a function of (K, A), delivering a one-dimensional state representation along the conditional saddle.

To complement this theoretical result, Figure 8 plots the computed conditional saddle paths in the (K, C) phase diagram. The model is solved globally using the repeated transition method, and the dynamics under each frozen aggregate state are simulated for 2,500 periods. Although the true endogenous state of the model is the full distribution Φ rather than K alone, the figure shows that conditional saddle paths are strictly and monotonically ordered in K, providing clear computational support for the sufficiency result.

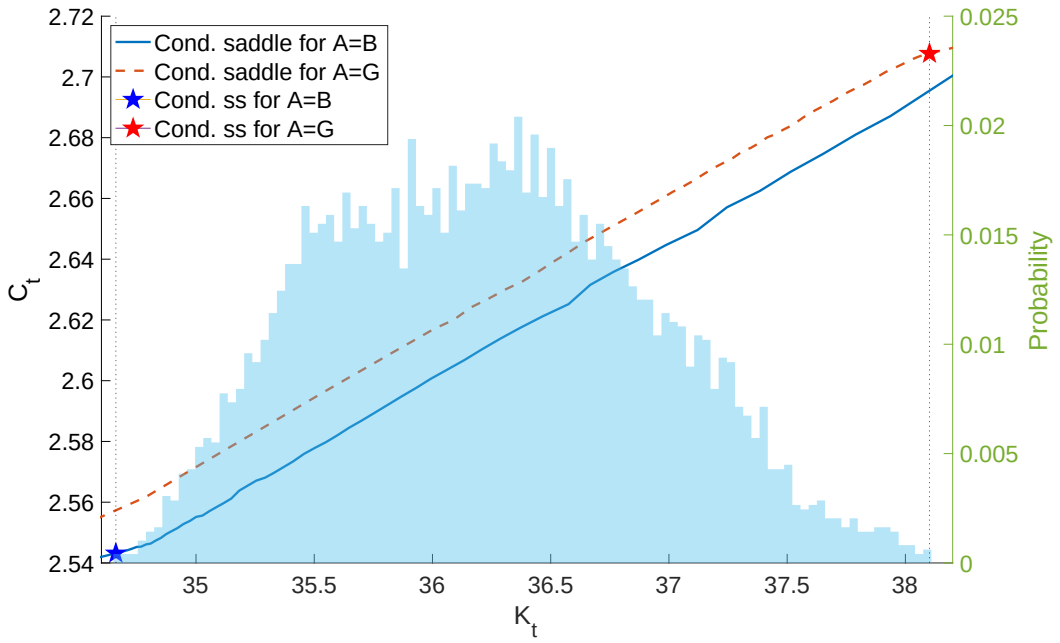


Figure 8: Conditional saddle paths in Krusell and Smith (1998)

Notes: The figure plots conditional saddle paths for $A = B$ (solid) and $A = G$ (dashed) implied by a canonical heterogeneous household business cycle model (Krusell and Smith, 1998). The histogram in the background plots the time-series distribution of the aggregate capital stock.

From spell-wise to global sufficiency Proposition 7 is stated regime-by-regime: K uniquely indexes the conditional saddle $\mathcal{M}(\Phi_0, A)$ under a single frozen regime. The equilibrium path is a concatenation of such one-dimensional segments, each exactly indexed by (K, A), with the handoff between spells pinned down by the scalar pair (K_τ, A_τ). Spell-wise sufficiency is distinct from unconditional global sufficiency, which

would require the same K to pin the same Φ regardless of path history. The following proposition bounds the across-spell discrepancy.

The key intuition is as follows. When the economy is observed at a given K, different prior histories may have produced different underlying distributions Φ—a "cloud" of distributions at the same capital level. Proposition 8 asks: *how tight is this cloud?* Within each regime spell, the economy converges along the conditional saddle, contracting the cloud at rate λ_c per period. A regime spell of length τ therefore shrinks the cloud by λ_c^τ. Since the spell length is random (geometric with persistence ρ_A), the expected contraction factor averages over possible spell lengths. Long spells and fast convergence both make the cloud tighter, so that K alone pins the distribution more precisely.

Proposition 8 (Across-spell approximation bound)**.**
Fix $A \in \{B, G\}$ and suppose Assumptions 1–2 hold. Let $\lambda_c \in (0,1)$ denote the contraction rate of $\Gamma_{\text{endo}}(\cdot, A)$ in the metric d on $\mathcal{X}$ restricted to the conditional saddle,[15] *i.e.,*

$$d\big(\Gamma_{\text{endo}}(\Phi, A),\, \Gamma_{\text{endo}}(\Phi', A)\big) \;\le\; \lambda_c\, d(\Phi, \Phi') \qquad \forall\, \Phi, \Phi' \in \mathcal{M}(\Phi_0, A).$$

Let $\rho_A := \max\{\pi_{BB}, \pi_{GG}\}$ denote the maximal regime persistence. Consider two distributional states Φ, Φ' that share the same aggregate capital $K(\Phi) = K(\Phi')$ but lie on conditional saddles originating from different initial conditions (i.e., reached via different histories of regime switches). Then the expected distributional discrepancy after one spell satisfies

$$\mathbb{E}_\tau\big[\lambda_c^\tau\big] \;=\; \frac{(1-\rho_A)\,\lambda_c}{1-\rho_A\lambda_c},$$

where τ is the (geometric) spell length, so that the expected across-path discrepancy is bounded by $\frac{(1-\rho_A)\,\lambda_c}{1-\rho_A\lambda_c}\, D$, with $D := \sup\{d(\Phi_1, \Phi_2) : K(\Phi_1) = K(\Phi_2),\ \Phi_1 \in \mathcal{M}(\Phi_0, A),\ \Phi_2 \in \mathcal{M}(\Phi_0', A)\}$.[16]

Proof.
See Appendix E. ∎

The supremum D is finite under the natural borrowing limit and compact idiosyncratic support. The bound vanishes as $\rho_A \to 1$ (long spells) or $\lambda_c \to 0$ (fast

[15] Near the conditional steady state, λ_c is governed by the stable eigenvalue $\lambda_s \in (0,1)$; globally, it may differ and must be bounded numerically.

[16] Within a spell of length τ, contractivity reduces discrepancy by λ_c^τ. The spell length is geometric with parameter $1-\rho_A$, so $\mathbb{E}[\lambda_c^\tau] = \sum_{k=1}^{\infty}(1-\rho_A)\rho_A^{k-1}\lambda_c^k = (1-\rho_A)\lambda_c/(1-\rho_A\lambda_c)$.

within-spell convergence).

Under the Krusell and Smith (1998) calibration, quarterly regime persistence is $\rho_A = 0.875$ (average 8-quarter spells) and the distributional contraction rate is $\lambda_c \approx 0.964$, yielding $\mathbb{E}[\lambda_c^\tau] \approx 0.77$. Comparing the globally solved K' at each spell boundary against the prediction from tracking only (K, A) gives an average discrepancy of 2.5×10^{-4} in log terms (1.3% of one standard deviation of $\ln K$), confirming that across-spell error is $O(10^{-4})$. The bound addresses distributional discrepancy from handoffs; functional-form approximation error (e.g., log-linearity) is a separate source.

From near rationality to complete rationality The Krusell and Smith (1998) approach is often labeled near-rational (Akerlof and Yellen, 1985), but this conflates two independent issues. Conditioning on K is exact within each spell—K is injective on the equilibrium manifold, so agents lose no distributional information. What *is* approximate is the log-linear functional form mapping today's K to tomorrow's. The bounded rationality, if any, resides in the parametric law of motion, not in the choice of state variable. Expanding the state space beyond K yields no within-spell improvement; enriching the functional form is the productive margin. The global accuracy ($R^2 > 0.9999$) reflects exact within-spell sufficiency and quantitatively small across-spell errors (Proposition 8).

Connection to block recursive equilibrium The conditional sufficiency of K is closely related to *block recursive equilibrium* (BRE), in which aggregate prices and quantities can be determined from a low-dimensional state without tracking the full distribution. Within each frozen-regime spell, the conditional saddle framework delivers exactly this structure: K pins down all aggregate equilibrium objects, and individual policies are determined given (K, A). Spell-wise K-sufficiency thus establishes that the equilibrium is block recursive *conditional on each regime*—aggregate prices, quantities, and the distributional state Φ itself are all recoverable from (K, A) on the conditional saddle path. Whether the equilibrium is globally block recursive depends on the across-spell discrepancy bounded in Proposition 8. Under the Krusell and Smith (1998) calibration, this discrepancy is quantitatively negligible.

Parametric form of the conditional saddle paths Despite the indexing function K, the specific form of the law of motion remains undetermined. The repeated

transition method (Lee, 2025) exploits the recurrence of equilibrium allocations along the conditional saddle path; Theorem 2 and Proposition 7 provide the theoretical foundation for implementing the RTM using K as the conditional sufficient statistic in heterogeneous-agent models.

Verifying monotonicity in practice The sufficiency results above rely on strict monotonicity of K along frozen-regime continuations (Assumption 2). How should a researcher verify this in a given model? The following two subsections develop complementary strategies. Section 4.1 provides a *perfect-foresight diagnostic*: one solves the deterministic ($\mu = 0$) transition dynamics and checks whether K converges monotonically; if so, a homotopy to $\mu = 1$ guarantees monotonicity under rational expectations. Section 4.2 provides an independent route from static primitives—bounds on household MPCs and price feedback—that does not require solving any transition path.

4.1 Perfect-foresight diagnostic for monotonicity

This section shows that K-monotonicity under rational expectations can be verified from the deterministic perfect-foresight model alone. The key insight is that the K-nullcline—the "obstacle" that the path must not cross—is invariant over beliefs (Proposition 6), so monotonicity can be transported continuously from perfect foresight to rational expectations.

Belief family For $\mu \in [0, 1]$, define the belief transition matrix $\Pi(\mu) := (1-\mu)\, I + \mu\, \Pi$, where I is the identity and Π is the true Markov matrix for $A \in \{B, G\}$. At $\mu = 0$, agents assign probability one to the current regime persisting (perfect foresight); at $\mu = 1$, agents hold rational expectations. For each μ, the μ-*RCE* is the recursive competitive equilibrium under $A' \sim \Pi(\mu)(\cdot|A)$ with the same technology and market structure. Denote the frozen-regime continuation under (μ, A) by $\{(K_t(\mu), C_t(\mu))\}_{t\geq 0}$, with conditional steady state $K_A^{cs}(\mu)$.

Under standard regularity, the frozen-regime continuation and conditional steady state are continuous in μ.[17] The question is whether K-monotonicity, which may hold

[17] Since $\Pi(\mu)$ enters the Bellman equation linearly, the recursive equilibrium operator is continuous in μ. The result follows by the parametric contraction mapping theorem and continuous dependence of the market-clearing fixed point. See Appendix G.

at $\mu = 0$ (perfect foresight), survives the passage to $\mu = 1$ (rational expectations). To formalize this, define the *monotonicity set*

$$\mathcal{S} := \big\{\mu \in [0,1] : \{K_t(\mu)\}_{t\geq 0} \text{ is strictly monotone along the frozen-}A \text{ continuation}\big\}.$$

The next result shows that this verification reduces to a single, easy check: if the deterministic perfect-foresight transition path converges monotonically—a textbook property of the Ramsey model—then monotonicity survives under rational expectations. The intuition is that the K-nullcline—the boundary between the rising-K and falling-K regions of the phase diagram—does not move as beliefs vary from perfect foresight to rational expectations (Proposition 6). If the perfect-foresight path stays on one side of this immovable boundary, then smoothly turning on rational expectations cannot push the path across it, because a crossing would require the nullcline gap to pass through zero—an event that can be ruled out by continuity.

Proposition 9 (Perfect-foresight diagnostic for K-monotonicity)**.**
Fix $A \in \{B, G\}$. If $0 \in \mathcal{S}$, then $\mathcal{S} = [0,1]$. In particular, strict K-monotonicity along perfect-foresight transition paths implies strict K-monotonicity under rational expectations, and Assumption 2 holds.

Proof.
Since $[0,1]$ is connected, it suffices to show $\mathcal{S}$ is nonempty, open, and closed.

Nonempty. $0 \in \mathcal{S}$ by assumption.

Open. Fix $\mu^* \in \mathcal{S}$ and define the nullcline gap $H_t(\mu) := C_t(\mu) - C_A^{Kcnull}(K_t(\mu))$, so that $\Delta K_t(\mu) = -H_t(\mu)$. At μ^*, strict monotonicity of K implies $H_t(\mu^*)$ has constant nonzero sign for all pre-limit t. By Proposition 6, the nullcline C_A^{Kcnull} is independent of μ. Combined with continuous dependence of $(K_t(\mu), C_t(\mu))$ on μ, for any finite T there exists $\varepsilon > 0$ such that $|\mu - \mu^*| < \varepsilon$ implies $H_t(\mu)\, H_t(\mu^*) > 0$ for all $t \leq T$.

For the tail: monotone convergence at μ^* implies the stable eigenvalue of the linearized frozen-regime dynamics at $K_A^{cs}(\mu^*)$ is real and lies in $(0,1)$—otherwise local convergence would be oscillatory. By continuity of the eigenvalue in μ, this property persists for nearby μ, extending monotonicity to the tail. See Appendix H for details.

Closed. Let $\mu_n \to \mu^*$ with $\mu_n \in \mathcal{S}$. By continuity, $H_t(\mu_n) \to H_t(\mu^*)$ for each t. Since each $H_t(\mu_n)$ has constant nonzero sign, the limit inherits constant weak sign—say $H_t(\mu^*) \geq 0$ for all t—so the K-path is weakly monotone.

It remains to show strict inequality. By the eigenvalue argument from the Open step, $H_t(\mu^*) > 0$ for all sufficiently large t, so any zero of H must occur at a finite t^* with $K_{t^*}(\mu^*) \neq K_A^{cs}(\mu^*)$. $H_{t^*} = 0$ means $\Delta K_{t^*} = 0$: the point (K_{t^*}, C_{t^*}) lies on the distribution-free K-nullcline (Proposition 6) at a non-steady-state capital level. This contradicts *unique K-stationarity*: the property that $\Delta K_t = 0$ along the convergent path only at K_A^{cs}. In the representative-agent model, K is Markov under the frozen regime, so $\Delta K_{t^*} = 0$ makes K_{t^*} a fixed point; uniqueness of the conditional steady state (Assumption 1(i)) yields the contradiction directly. In the heterogeneous-agent model, the eigenvalue condition $\lambda_s(\mu^*) \in (0, 1)$ establishes unique K-stationarity locally near $K_A^{cs}(\mu^*)$; the global extension holds under the bounded price-income MPC (Assumption 3) via Lemma 1 (Appendix F). So $H_t(\mu^*) > 0$ for all t and $\mu^* \in \mathcal{S}$. ■

Discussion Proposition 9 reduces verification to a single check: monotonicity of the deterministic PF transition path—a textbook property of the Ramsey saddle (representative agent) or the Bewley–Aiyagari economy under fixed TFP (heterogeneous agents). The proof connects $\mu = 0$ to $\mu = 1$ through a continuous path of equilibria, exploiting nullcline invariance (Proposition 6) as the common "obstacle." For the heterogeneous-agent case, the global step requires unique K-stationarity, verifiable via the MPC bound of Section 4.2; the two routes are complementary.

More generally, the diagnostic extends to any continuous aggregate variable e satisfying unique stationarity ($e(\Phi_{t+1}) = e(\Phi_t)$ only at $\Phi_t = \Phi^{\text{cs}}$), even when the nullcline is not distribution-free. This is relevant for the extensions in Section 4.3: the bond economy and CRRA case involve distributional dependence, but at $\mu = 0$ there is no distributional state, so monotonicity need only be checked along the deterministic transition path.

4.2 Primitive conditions for Assumption 2

An independent verification route works directly under rational expectations ($\mu = 1$) from static allocation primitives—bounds on household marginal propensities to consume and general-equilibrium price feedback—without solving any transition path. Appendix F develops this approach: under a bounded price-income MPC (Assumption 3 therein), Lemma 1 shows that sufficiently weak GE price feedback relative to the nullcline slope ensures Assumption 2, and Corollary 1 provides a local check near the conditional steady state. In the baseline Krusell and Smith (1998) calibration,

the sufficient condition holds with substantial slack (the left-hand side of the bound is below 0.04 under the conservative bound $\bar{m} = 1$, well below the threshold 1).

4.3 Extensions of the dimension-reduction result

Theorem 2 provides verifiable conditions for dimension reduction: identify the K-nullcline and verify monotonicity, using either the perfect-foresight diagnostic (Proposition 9) or the primitive MPC conditions (Appendix F). This subsection applies these tools to several extensions.

Economies with multiple endogenous states Theorem 2 applies to an economy with the multivariate (distributional) endogenous state. As long as there is an aggregate equilibrium variable e that strictly monotonically converges to the conditional steady-state level, then e is a conditional sufficient statistic on each conditional saddle path.

For example, consider the following model that extends the heterogeneous-agent model above by adding endogenous bond holding. The corresponding budget constraint is:

$$c + a' + q(\Phi, A)b' = a(1 + r(\Phi, A)) + b + w(\Phi, A)z \tag{30}$$

where b is bond holding and q is the bond price competitively determined by

$$\int b'(a, z; \Phi, A)d\Phi = 0. \tag{31}$$

The key observation is that Proposition 6 continues to hold: the conditional K-nullcline remains $C_A^{Kcnull}(K) = AK^\alpha - \delta K$, determined solely by the aggregate resource constraint. Since bonds are in zero net supply, their inclusion does not alter the sign-determining property of the nullcline. Thus, the monotonicity argument of Proposition 7 applies, and K remains a conditional sufficient statistic. Lee (2025) confirmed this prediction computationally: equilibrium allocations are strictly monotonically sorted along K in the globally solved model, consistent with Theorem 2.

Remark 2 (Multiple predetermined aggregate states)**.**
When the economy features multiple predetermined aggregate states—e.g., liquid and illiquid capital in a two-asset HANK model—the conditional saddle is generically

higher-dimensional, and Theorem 2 requires a monotone scalar *coordinate along this higher-dimensional object. Existence of such a coordinate is model-specific and cannot be guaranteed by the K-nullcline argument alone. The perfect-foresight diagnostic (Proposition 9) remains applicable provided one can identify a monotone aggregate variable along the deterministic transition path. Establishing general sufficient conditions for dimension reduction on multi-dimensional conditional saddles is an open question.*

Models with endogenous labor supply With GHH preferences $u(c, l_H) = \frac{1}{1-\sigma}(c - \frac{\eta}{1+1/\chi} l_H^{1+1/\chi})^{1-\sigma}$, labor supply depends only on wages and productivity, eliminating wealth effects. The conditional K-nullcline depends on technology parameters and the stationary productivity distribution but not on the wealth distribution (Appendix I), so the monotonicity argument of Theorem 2 applies and K remains a conditional sufficient statistic along each conditional saddle path.

Under separable CRRA utility $u(c, l_H) = \frac{c^{1-\sigma}}{1-\sigma} - \frac{\eta}{1+1/\chi} l_H^{1+1/\chi}$, individual labor supply depends on consumption through the wealth effect, so the K-nullcline acquires distributional dependence. Strict monotonicity cannot be verified analytically in this case, and establishing it formally under CRRA–separable preferences remains an open question. The perfect-foresight diagnostic of Proposition 9 remains applicable since at $\mu = 0$ there is no distributional state, and computational experiments under standard calibrations confirm strictly monotone convergence (Lee, 2025). The derivations for both specifications are in Appendix I.

5 Search and matching

The conditional saddle framework applies to any recursive competitive equilibrium with a predetermined endogenous state, not only neoclassical growth models. This section applies it to a Diamond–Mortensen–Pissarides search-and-matching economy, where the predetermined state is unemployment. Conditional saddle paths recover the Beveridge curve as a regime-specific transition locus, and the slope-asymmetry characterization of Theorem 1 acquires a concrete labor-market interpretation: the surplus compression at the heart of the Shimer puzzle manifests as differential saddle steepness.

The section proceeds in three parts. The baseline model produces conditional sad-

dles whose slope asymmetry generates state dependence in vacancy creation. I then show how stochastic and unemployment-dependent benefits reshape the geometry. Comparative statics reveal a non-monotone relationship between benefit generosity and unemployment volatility, traceable to the compression of the conditional saddle structure.

The DMP environment Consider a discrete-time stochastic Diamond–Mortensen–Pissarides economy.[18] Aggregate productivity follows a Markov chain $z \in \{G, B, \ldots\}$; for exposition I focus on the two extreme states $z \in \{G, B\}$. The aggregate state is $S = [u_{-1},\, z]$, where unemployment $u_{-1} = 1 - n_{-1}$ is the sole predetermined endogenous state. Employment evolves according to

$$n \;=\; (1-\lambda)\, n_{-1} \;+\; q(\theta)\, v, \tag{32}$$

where $\lambda \in (0,1)$ is the exogenous separation rate, $\theta = v/u$ is labor market tightness, and $q(\theta) = \bar{m}\,\theta^{-\xi}$ is the vacancy-filling probability implied by a Cobb–Douglas matching function. Market tightness $\theta(S)$—and hence vacancies $v(S) = \theta(S) \cdot u_{-1}$—is the jump variable. It is pinned down by the free-entry condition for vacancy posting:

$$\frac{\kappa}{q(S)} \;=\; \beta\, \mathbb{E}\left[\left(\frac{c(S)}{c(S')} \right)^{\sigma} \left(z' - w(S') + (1-\lambda)\, \frac{\kappa}{q(S')} \right) \,\middle|\, S \right], \tag{33}$$

where κ is the vacancy posting cost, $w(S)$ is the Nash-bargained wage, and $c(S) = n(S)\, z - \kappa\, v(S) + (1 - n(S))\, b$ is aggregate consumption from the resource constraint.

The conditional saddle path $\mathcal{M}(u_0,\, z)$ is the orbit $\{(u_t, v_t)\}_{t \geq 0}$ generated by (32) and the equilibrium vacancy function $v(\cdot, z)$ under frozen productivity z. Its limit point is $(u_z^{\mathrm{cs}}, v_z^{\mathrm{cs}})$.

Slope asymmetry and the Beveridge curve Figure 9a plots the conditional saddle paths in (u_{-1}, v) space. Both loci slope upward—higher inherited unemployment raises equilibrium vacancies—so state dependence is sign-preserving, as in the RBC benchmark. However, the slopes differ markedly. Under G, matches are prof-

[18] The full model specification and calibration are in Appendix J. I adopt a standard formulation following Shimer (2005) with Nash-bargained wages, exogenous separation, and Cobb–Douglas matching. The computation uses an 11-state Tauchen discretization of log TFP; for expositional simplicity, I present the two extreme states as G and B.

itable and the saddle is steep; under B, the flow surplus is low and the saddle is nearly flat.

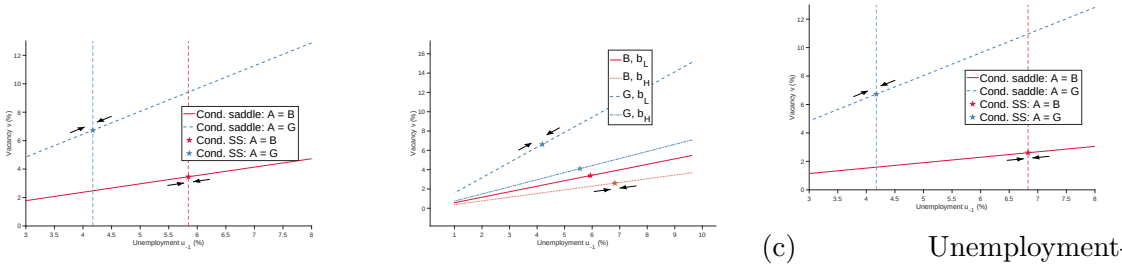
The mechanism is the free-entry condition (33). Each additional unemployed worker represents a potential match, but the value of that match depends on productivity. Under high z, the flow surplus $z - w$ is large, so firms respond aggressively to a marginal increase in the unemployment pool, producing a steep vacancy locus. Under low z, the surplus is compressed—Nash bargaining ensures wages fall less than proportionally to productivity (Shimer, 2005)—so vacancy posting barely responds, producing a nearly flat locus. The slope asymmetry is the visual signature of the surplus compression that underlies the Shimer puzzle.

This asymmetry generates the Beveridge curve "loop" observed in empirical (u, v) data. Starting near the G steady state, a negative shock switches the economy onto the flat B saddle. Vacancies collapse on impact and barely adjust as unemployment accumulates along the flat path. Upon recovery, the economy jumps back onto the steep G saddle at high inherited u_{-1}, driving a sharp vacancy spike and rapid job creation. The counterclockwise rotation through recessions and recoveries is thus the across-saddle jump dynamics of the stochastic equilibrium.

Benefit policy and saddle geometry The baseline with constant benefits (Figure 9a) produces two conditional saddles whose slope difference generates the Beveridge curve loop. I now examine how the benefit structure reshapes this geometry (full specifications in Appendix J).

When benefits follow an independent Markov chain $b \in \{b_H, b_L\}$, the exogenous state expands to $(z, b) \in \{G, B\} \times \{b_H, b_L\}$, producing four conditional saddle paths (Figure 9b). The slope ordering—(G, b_L) steepest, (B, b_H) flattest—separates into two productivity groups, each split by the benefit level.

The most informative case replaces the independent benefit shock with an automatic stabilizer (Figure 9c). Benefits increase linearly in unemployment above a threshold, rising from the baseline b to $b + \bar{b}\, w^{ss}$. Because benefits depend on unemployment rather than forming a separate exogenous state, the model has two conditional saddle paths indexed by TFP alone. Along each saddle, benefits rise as unemployment increases. The effect is asymmetric. In the B regime, the economy settles at higher unemployment, which triggers higher benefits and compresses the surplus. The G saddle retains meaningful slope. This widens the slope asymmetry



(a) Constant b (b) Stochastic b (c) Unemployment-dependent b

Figure 9: Conditional saddle paths in the DMP model under alternative b regimes

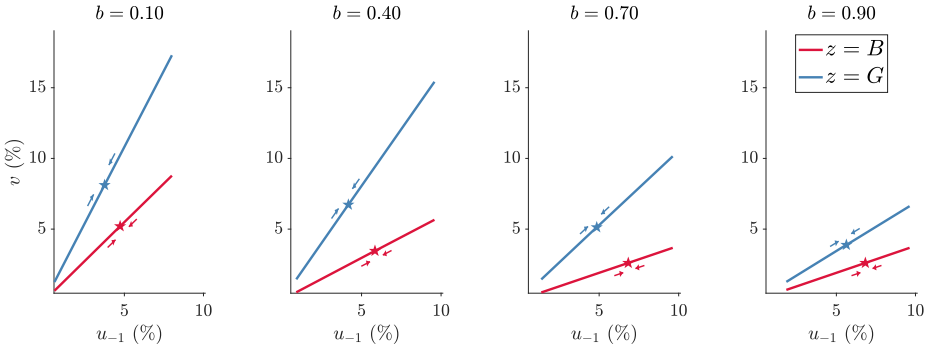
Notes: Conditional saddle paths in (u_{-1}, v) space. Stars mark conditional steady states; arrows indicate the direction of convergence. Panel (a): constant benefits with two conditional saddles under frozen productivity $z \in \{G, B\}$. Panel (b): independently stochastic benefits $b \in \{b_H, b_L\}$ yield four conditional saddles—(G, b_L) steepest, (B, b_H) flattest. Panel (c): benefits depend on unemployment (linearly increasing in u above a threshold), so the exogenous state remains two TFP levels. The slope asymmetry and steady-state unemployment gap both widen relative to panel (a).

and amplifies state dependence relative to the constant-benefit baseline.

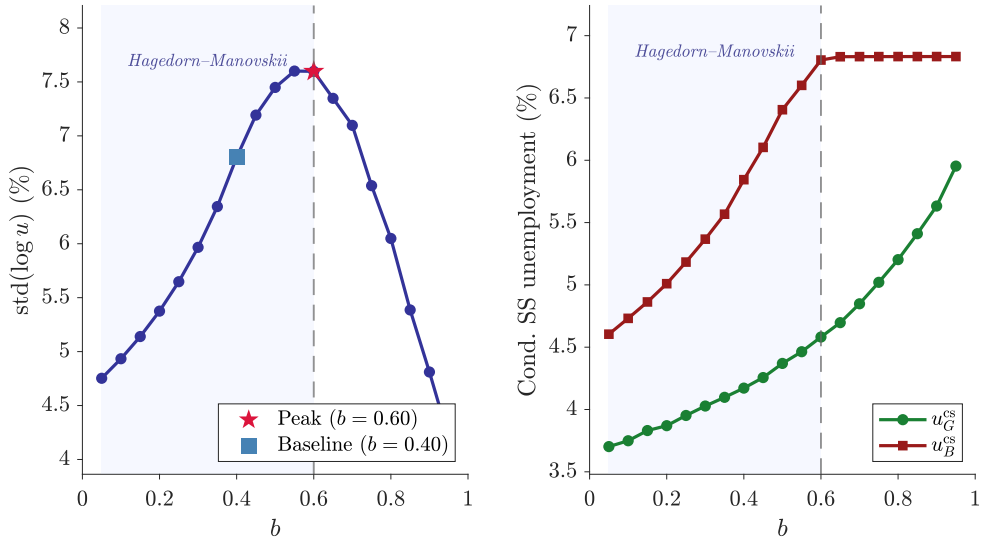
Comparative statics of the benefit level How does the benefit level b reshape the conditional saddle geometry, and what does this imply for unemployment volatility? The Nash wage $w = (1-\eta)b + \eta(z + \kappa\theta)$ is increasing in b, so raising b compresses the match surplus $z - w$. The compression is proportionally more severe in the B regime, where $z - b$ is already small. Figure 10a solves the DMP model for $b \in \{0.10, 0.40, 0.70, 0.90\}$. At low b, both surpluses are large and the saddles are nearly parallel. As b rises, the B saddle flattens while the G saddle retains meaningful slope, widening the slope asymmetry. One might expect monotonically rising volatility, but the effect is *hump-shaped* (Figure 10b): volatility peaks at $b \approx 0.60$ ($\text{std}(\log u) \approx 7.6\%$) and then *declines*. The exact peak location is model- and calibration-specific. The saddle geometry clarifies why high-b calibrations such as Hagedorn and Manovskii (2008) ($b \approx 0.955$) can paradoxically yield lower volatility than moderate ones: at high b, conditional steady states converge despite large slope gaps, compressing the range over which unemployment can fluctuate.[19]

The saddle framework diagnoses this hump by revealing that what governs volatility is the *separation between conditional steady states*, not slope asymmetry alone. The right subplot of Figure 10b plots conditional steady-state unemployment under

[19] In our parameterization, simulated volatility at $b = 0.955$ is about 4.2%, compared with 5.8% at the baseline $b = 0.40$. The saddle diagnostic identifies steady-state convergence as the mechanism.



(a) Conditional saddle paths



(b) Unemployment volatility (left) and conditional steady-state unemployment (right)

Figure 10: Comparative statics of b

Notes: Panel (a): conditional saddle paths in (u_{-1}, v) space for $b \in \{0.10, 0.40, 0.70, 0.90\}$; stars mark conditional steady states. Panel (b), left: std(log u) over 19 values of b; the star marks the peak ($b \approx 0.60$), the square marks the baseline ($b = 0.40$). Panel (b), right: conditional steady-state unemployment under frozen good (u_G^{cs}, green) and bad (u_B^{cs}, red) TFP, extracted from the stochastic policy functions. Both the gap $u_B^{cs} - u_G^{cs}$ and simulated volatility are hump-shaped in b.

frozen good (u_G^{cs}) and bad (u_B^{cs}) TFP for each b. As b rises, u_G^{cs} increases steadily while u_B^{cs} rises faster, widening the gap from under 1 p.p. at $b = 0.05$ to about 2.2 p.p. at $b \approx 0.60$. Beyond that point, u_B^{cs} *saturates*. Agents expect future TFP recovery, so vacancy posting persists even under bad productivity, capping the bad-state conditional steady state near 6.8%. Meanwhile, u_G^{cs} continues to rise as higher b compresses the surplus even under good TFP. The two conditional steady states therefore *con-*

verge at high b: the saddles become highly asymmetric in slope but close together, leaving little room for unemployment to fluctuate between regimes. Both the gap $u_B^{\text{cs}} - u_G^{\text{cs}}$ and simulated volatility are hump-shaped in b, peaking near the same value. The conditional steady-state separation—not the slope difference—is the operative geometric diagnostic for unemployment volatility.

6 Concluding remarks

Conditional saddle paths—equilibrium trajectories under frozen exogenous states—bring the geometric clarity of deterministic phase diagrams to stochastic environments. The slopes of these objects encode state dependence: differential tilt across saddles is the geometric signature of asymmetric shock responses. Their monotonicity delivers dimension reduction: a strictly monotone aggregate variable exactly indexes the equilibrium state within each frozen-regime spell. Persistent spells, exact within-spell sufficiency, and nullcline invariance constraining across-spell variation explain why this conditional result translates into near-global accuracy, providing a structural foundation for scalar-state methods.

The paper illustrates the framework through three applications. The canonical RBC model establishes the basic geometry and reveals that state-dependent response shapes arise from distance along the saddle even when slopes are nearly equal. The heterogeneous-household model of Krusell and Smith (1998) demonstrates that strict monotonicity along conditional saddles delivers exact dimension reduction, providing a structural rationale for scalar-state methods. The search-and-matching application shows the framework's reach beyond neoclassical growth: the Beveridge curve emerges as a regime-specific saddle locus, surplus compression under low productivity manifests as differential saddle steepness, and the saddle geometry provides a direct diagnostic for the non-monotone relationship between benefit generosity and unemployment volatility.

More broadly, conditional saddles offer a complementary lens on stochastic models. They provide a language for interpreting nonlinear dynamics, for assessing when scalar state reductions are exact on the equilibrium path rather than merely accurate, and for characterizing across-path variation as a structural feature of the geometry. A promising direction for future work is to use these geometric objects to sharpen empirical restrictions on state dependence and to discipline the design of state-contingent

stabilization policies.

References

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2021): "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *The Review of Economic Studies*, 89, 45–86.

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

Akerlof, G. A. and J. L. Yellen (1985): "Can Small Deviations from Rationality Make Significant Differences to Economic Equilibria?" *American Economic Review*, 75, 708–720.

Andreasen, M. M., J. Fernández-Villaverde, and J. F. Rubio-Ramírez (2017): "The Pruned State-Space System for Non-Linear DSGE Models: Theory and Empirical Applications," *The Review of Economic Studies*, 85, 1–49.

Arnold, L. (1998): *Random Dynamical Systems*, Springer Monographs in Mathematics, Berlin, Heidelberg: Springer Berlin Heidelberg.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," *Econometrica*, 89, 2375–2408.

Auerbach, A. J. and Y. Gorodnichenko (2012): "Measuring the Output Responses to Fiscal Policy," *American Economic Journal: Economic Policy*, 4, 1–27.

Baley, I. and A. Blanco (2019): "Firm Uncertainty Cycles and the Propagation of Nominal Shocks," *American Economic Journal: Macroeconomics*, 11, 276–337.

Barro, R. J. and X. Sala-i Martin (1992): "Convergence," *Journal of Political Economy*, 100, 223–251.

Basu, S. and B. Bundick (2017): "Uncertainty Shocks in a Model of Effective Demand," *Econometrica*, 85, 937–958.

Berger, D., K. Milbradt, F. Tourre, and J. Vavra (2021): "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy," *American Economic Review*, 111, 2829–78.

Berger, D. and J. Vavra (2015): "Consumption Dynamics During Recessions: Consumption Dynamics During Recessions," *Econometrica*, 83, 101–154.

Bewley, T. (1986): "Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers," *Contributions to Mathematical Economics in Honor of Gérard Debreu*, 79–102.

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

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

Bourany, T. (2018): "Wealth Distribution over the Business Cycle: A Mean-Field Game with Common Noise," Master's thesis, UPMC-Sorbonne Université, supervised by Yves Achdou, Paris-Diderot University (LJLL).

Carroll, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112, 56.

Cass, D. (1965): "Optimum Growth in an Aggregative Model of Capital Accumulation1," *The Review of Economic Studies*, 32, 233–240.

Deaton, A. (1991): "Saving and Liquidity Constraints," *Econometrica*, 59, 1221–1248.

Dechert, W. D. and K. Nishimura (1983): "A Complete Characterization of Optimal Growth Paths in an Aggregated Model with a Non-Concave Production Function," *Journal of Economic Theory*, 31, 332–354.

Den Haan, W. J. (1996): "Heterogeneity, Aggregate Uncertainty, and the Short-Term Interest Rate," *Journal of Business & Economic Statistics*, 14, 399–411.

Den Haan, W. J. and A. Marcet (1990): "Solving the Stochastic Growth Model by Parameterizing Expectations," *Journal of Business & Economic Statistics*, 8, 31–34.

Fernández-Villaverde, J., S. Hurtado, and G. Nuño (2023): "Financial Frictions and the Wealth Distribution," *Econometrica*, 91, 869–901.

Gali, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, Princeton University Press.

Gorman, W. M. (1961): "On a Class of Preference Fields," *Metroeconomica*, 13, 53–56.

Hagedorn, M. and I. Manovskii (2008): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 98, 1692–1706.

Huggett, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

Jo, Y. J. and S. Zubairy (2025): "State-Dependent Government Spending Multipliers: Downward Nominal Wage Rigidity and Sources of Business Cycle Fluctuations," *American Economic Journal: Macroeconomics*, 17, 379–413.

Kamihigashi, T. (2003): "Necessity of transversality conditions for stochastic problems," *Journal of Economic Theory*, 109, 140–149.

——— (2005): "Necessity of the transversality condition for stochastic models with bounded or CRRA utility," *Journal of Economic Dynamics and Control*, 29, 1313–1329.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

Kaplan, G. and G. L. Violante (2014): "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 82, 1199–1239.

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

Koopmans, T. (1963): "On the Concept of Optimal Economic Growth," Cowles Foundation Discussion Papers 163, Cowles Foundation for Research in Economics, Yale University.

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

Lee, H. (2025): "Global Nonlinear Solutions in Sequence Space and the Generalized Transition Function," Working Paper.

——— (2026): "Striking While the Iron Is Cold: Fragility after a Surge of Lumpy Investments," *Working Paper*, 68.

Lee, H. and K. Nomura (2026): "The Spender of Last Resort: Global Equilibrium Dynamics under the Zero Lower Bound," Working Paper.

Lee, H. and Y. Sun (2026): "Endogenous Plucking Through Networks: The Plucking Paradox," Working Paper.

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

Melcangi, D. (2024): "Firms' precautionary savings and employment during a credit crisis," *American Economic Journal: Macroeconomics*, 16, 356–386.

Neftci, S. N. (1984): "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy*, 92, 307–328.

Petrosky-Nadeau, N., L. Zhang, and L.-A. Kuehn (2018): "Endogenous Disasters," *American Economic Review*, 108, 2212–45.

Pizzinelli, C., K. Theodoridis, and F. Zanetti (2020): "STATE DEPENDENCE IN LABOR MARKET FLUCTUATIONS," *International Economic Review*, 61, 1027–1072.

Proehl, E. (2025): "Existence and Uniqueness of Recursive Equilibria with Aggregate and Idiosyncratic Risk," Working Paper.

Ramsey, F. P. (1928): "A Mathematical Theory of Saving," *The Economic Journal*, 38, 543–559.

Reddy, P. B., J. M. Schumacher, and J. C. Engwerda (2020): "A Bendixson Criterion for Optimal Control of a Class of Hybrid Systems," *SIAM Journal on Control and Optimization*, 58, 2219–2252.

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

Schenk-Hoppé, K. R. (2001): "Random Dynamical Systems in Economics," *Stochastics and Dynamics*.

Schenk-Hoppé, K. R. and B. Schmalfuss (2001): "Random fixed points in a stochastic Solow growth model," *Journal of Mathematical Economics*, 36, 19–30.

Schenk-Hoppé, K. R. (1998): "Random attractors–general properties, existence and applications to stochastic bifurcation theory," *Discrete and Continuous Dynamical Systems*, 4, 99–130.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95, 25–49.

Skiba, A. K. (1978): "Optimal Growth with a Convex–Concave Production Function," *Econometrica*, 46, 527–539.

Solow, R. M. (1956): "A Contribution to the Theory of Economic Growth," *The Quarterly Journal of Economics*, 70, 65–94.

Swan, T. W. (1956): "ECONOMIC GROWTH and CAPITAL ACCUMULATION," *Economic Record*, 32, 334–361.

Tenreyro, S. and G. Thwaites (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8, 43–74.

Vavra, J. (2014): "INFLATION DYNAMICS AND TIME-VARYING VOLATILITY: NEW EVIDENCE AND AN SS INTERPRETATION," *The Quarterly Journal of Economics*, 129, 215–258.

Wagener, F. (2003): "Skiba Points and Heteroclinic Bifurcations, with Applications to the Shallow Lake System," *Journal of Economic Dynamics and Control*, 27, 1533–1561.

Walsh, K. J. and E. Young (2024): "Equilibrium Multiplicity in Aiyagari and Krusell-Smith," Working paper.

Winberry, T. (2021): "Lumpy Investment, Business Cycles, and Stimulus Policy," *American Economic Review*, 111, 364–396.

Yannacopoulos, A. N. (2011): "Stochastic Saddle Paths and Economic Theory," in *Dynamics, Games and Science II*, ed. by M. M. Peixoto, A. A. Pinto, and D. A. Rand, Berlin, Heidelberg: Springer Berlin Heidelberg, 735–752.