

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

Hanbaek Lee[‡]

January 9, 2026

([click here for the latest version](#))

Abstract

This paper extends deterministic saddle-path analysis to stochastic environments by introducing *conditional saddle paths*: equilibrium manifolds under frozen exogenous states. This concept yields a global geometric representation of stochastic equilibrium dynamics, in which equilibrium fluctuations decompose into movements *along* (endogenous propagation) and *across* (exogenous state transitions) conditional saddle paths. The framework delivers two theoretical results. First, state-dependent impulse responses arise from differences in the slopes of conditional saddle paths. Second, if an aggregate equilibrium variable varies strictly monotonically along conditional saddle paths, it uniquely indexes equilibrium states and thus provides an exact one-dimensional sufficient statistic. Applying this result, I prove that aggregate capital is a sufficient statistic in a canonical heterogeneous-household model ([Krusell and Smith, 1998](#)).

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

JEL codes: C62, D31, E32.

[†]I am grateful for insightful comments from Timo Boppart, Francisco Buera, Vasco Carvalho, and Yucheng Yang. All errors are my own.

[‡]University of Cambridge. Email: hl610@cam.ac.uk

1 Introduction

Understanding equilibrium dynamics in models with aggregate uncertainty remains a central challenge in macroeconomics. Unlike deterministic models, where saddle-path diagrams provide immediate geometric intuition, stochastic equilibrium models lack comparable geometric frameworks. This makes it difficult to develop intuition about how economies respond to shocks, how different states interact, and why certain computational methods work. The challenge becomes particularly acute in heterogeneous-agent models, where the natural state variable is an infinite-dimensional wealth distribution — yet computational work routinely achieves dimension reduction to low-dimensional aggregates. Can we extend geometric saddle-path analysis to stochastic environments? What does such a framework reveal about equilibrium dynamics, state-dependent responses, and the success of computational approximations?

This paper makes two contributions. First, it extends the saddle-path analysis of equilibrium dynamics from deterministic models to stochastic environments by introducing the notion of *conditional saddle path*: equilibrium trajectories of the regime-frozen economy (holding the aggregate state fixed at a level). This geometric object decomposes business-cycle fluctuations into movements *along* a conditional saddle (endogenous propagation) and *across* conditional saddles (transitions when the exogenous state changes), providing a unified framework for analyzing state-dependent equilibrium dynamics. The framework enables visualization of generalized impulse responses and nonlinear transition dynamics in phase diagrams, analogous to how saddle-path diagrams illuminate deterministic models. Based on this geometric framework, I establish that state-dependent shock responses arise essentially from differences in the local slopes along the conditional saddle paths.

Second, it provides a general dimension-reduction theorem on these manifolds. If an aggregate equilibrium variable is strictly monotone and convergent along the conditional saddle, then it is injective on the invariant equilibrium set and hence uniquely indexes all equilibrium allocations and prices. As a consequence, the equilibrium can be represented *exactly* — not merely approximately — with a one-dimensional sufficient statistic. Applying this theoretical result, I establish that aggregate capital is a sufficient statistic for canonical heterogeneous-agent models (Krusell and Smith, 1998). This provides a geometric foundation for the empirical success of scalar approximations and offers verifiable conditions based on primitive assumptions rather

than numerical iteration.

These results extend beyond the canonical one-asset setting. In heterogeneous-agent models with multiple endogenous state variables, such as risky and riskless assets, liquid and illiquid assets, domestic and foreign bonds, the same logic applies (Krusell and Smith, 1997; Mendoza, 2010; Khan and Thomas, 2013; Kaplan and Violante, 2014; Berger and Vavra, 2015; Kaplan et al., 2018): whenever a monotone-convergent aggregate coordinate exists along conditional saddle paths, equilibrium dynamics remain exactly traceable by a scalar index.

Conditional saddle paths are equilibrium manifolds defined under frozen exogenous aggregate states: they describe counterfactual equilibrium continuations — how the economy would evolve if the aggregate state were held fixed at a given value — computed under the same decision rules that govern the stochastic equilibrium with regime switching. These counterfactual paths are economically meaningful because equilibrium decisions internalize the possibility of future regime changes, and the frozen-regime objects formalize the corresponding equilibrium “thought experiments.”

Conditional saddle paths are closely connected to the random dynamical systems (RDS) notion of invariant manifolds for dynamics under exogenous forcing (Arnold, 1998). In the present setting, the forcing is the aggregate Markov state, and a recursive competitive equilibrium (RCE) implies a time-homogeneous endogenous law of motion; conditional saddle paths are the invariant manifolds of the resulting frozen-regime equilibrium maps. Both perspectives organize stochastic dynamics using invariant geometric objects rather than local linearizations. The distinction in this paper is therefore not the mathematics of invariance, but the equilibrium discipline: the invariant-manifold structure is pinned down by optimality, market clearing, and equilibrium consistency, and it is precisely what makes global transition functions and impulse responses well defined and comparable across aggregate states within the recursive equilibrium framework.

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. Marcer (1988), Den Haan and Marcer (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. Thereafter, the literature has dramatically developed to sharpen the accuracy and improve the computational efficiency by incorporating moment-based approximations, exact aggregation, functional approximations, sequence-space approaches, and machine/deep learning (Den Haan, 1996, 1997; Reiter, 2001; Algan et al., 2008, 2010; Den Haan and Rendahl, 2010; Reiter, 2010; Ahn et al., 2018; Boppart et al., 2018; Elenev et al., 2021; Auclert et al., 2021; Cao et al., 2023; Azinovic et al., 2022; Fernández-Villaverde et al., 2023; Han et al., 2025; Payne et al., 2025).

This paper differs from these computational contributions by providing geometric foundations for *why* dimension reduction works. Rather than developing new algorithms, I introduce conditional saddle paths as an geometric framework for understanding equilibrium dynamics — analogous to how saddle-path diagrams provide intuition in deterministic models. The framework reveals that aggregate capital’s sufficiency in a canonical heterogeneous-household model follows from geometric properties (nullcline invariance and monotonicity) rather than numerical happenstance.

The exact sufficiency result also has computational implications. Lee (2025) develops a repeated transition method that constructs conditional expectations for individual-level problems by identifying periods in simulations, where aggregate states are similar, enabling reuse of computed transitions across these similar states. The method requires a metric for determining when aggregate states are “close enough” to pool. My sufficiency result establishes that aggregate capital distance provides a theoretically justified metric — agents’ problems are identical whenever capital stocks coincide, regardless of distributional differences. This validates distance-based pooling strategies and enables efficient implementation without requiring explicit distributional tracking.

This paper builds on the traditional use of geometric methods to analyze economic dynamics in deterministic environments. Phase-diagram analyses of the Solow–Swan (Solow, 1956; Swan, 1956) and Ramsey–Cass–Koopmans (Ramsey, 1928; Cass, 1965; Koopmans, 1963) growth models provide foundational intuition about convergence and stability. I extend this geometric approach to stochastic environments by introducing conditional saddle paths. This framework offers a geometric interpretation of stochastic equilibrium dynamics, including nonlinear and state-dependent impulse

responses. In particular, it provides a useful visual tool for understanding how microfounded frictions generate state-dependent dynamics, as documented in the recent literature (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; Melcangi, 2024; Winberry, 2021; Lee, 2026).

Finally, in mathematics, the literature on random dynamical systems provides foundational tools for geometric analyses of stochastic dynamic processes. However, its focus has not been on history-invariant saddle paths of the type that arise in recursive competitive equilibrium (Arnold, 1998). For example, Yannacopoulos (2011) introduces the notion of stochastic saddle paths, which are conceptually distinct from conditional saddle paths in that they vary with the realized history of shocks.¹

By contrast, conditional saddle paths furnish economists with a geometric representation of stochastic equilibrium dynamics that is directly analogous to the role of phase diagrams in deterministic models. While this framework relies on standard regularity conditions to ensure well-defined and bounded equilibrium paths (Kamihigashi, 2003, 2005), the emphasis here is not on establishing existence results, but rather on providing geometric tools for understanding and analyzing stochastic equilibrium dynamics.

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 law of motions (Γ_{endo} , Γ_{exo}):

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

where Φ is the endogenous aggregate state variable, and A is the exogenous aggregate state variable. The latter admits a multivariate vector that follows a stochastic process. I assume the exogenous aggregate law of motion Γ_{exo} is a Markov chain.

¹Stochastic saddle paths depend on the specific sequence of past realizations, whereas conditional saddle paths are invariant to history.

For simplicity in the illustration, I assume Γ_{exo} is a two-state Markov chain where $A \in \{B, G\}$, and $\Gamma_{exo}(A'|A) > 0$ for $\forall A', A$.

As a pre-step to define the conditional saddle path, I first introduce the reachable equilibrium set from an initial aggregate state.

Definition 1 (Reachable equilibrium set).

Let (Φ_t, A_t) denote the endogenous and exogenous aggregate state under a stationary RCE, where $\Phi_{t+1} = \Gamma_{endo}(\Phi_t, A_t)$ and A_{t+1} follows the Markov matrix Π . Fix an initial condition (Φ_0, A_0) . The reachable equilibrium set $\mathcal{R}(\Phi_0, A_0) \subseteq \mathcal{X} \times \{B, G\}$ is the smallest set containing (Φ_0, A_0) such that for every $(\Phi, A) \in \mathcal{R}(\Phi_0, A_0)$ and every $A' \in \{B, G\}$ with $\Pi_{AA'} > 0$,

$$(\Gamma_{endo}(\Phi, A), A') \in \mathcal{R}(\Phi_0, A_0).$$

The reachable equilibrium set $\mathcal{R}(\Phi_0, A_0)$ is best viewed as an *extended equilibrium history*. Starting from the initial state (Φ_0, A_0) , it contains not only the states visited along the realized sequence of aggregate regimes, but also all *counterfactual continuations* induced by alternative regime realizations permitted by the Markov matrix Π . Formally, $\mathcal{R}(\Phi_0, A_0)$ collects every equilibrium state (Φ, A) that can arise from (Φ_0, A_0) under some feasible history of aggregate regimes, evaluated under the *same* stationary RCE decision rules. This construction is useful because it lets us define geometric equilibrium objects (conditional saddle paths, impulse responses, transition paths) without tying them to one particular realized history, even though a realized path typically visits each distribution Φ at most once.

Definition 2 (Conditional saddle path).

Given (Φ_0, A_0) , for each $A \in \{B, G\}$ define the A -slice of \mathcal{R} as

$$\mathcal{R}_A(\Phi_0, A_0) := \{\Phi \in \mathcal{X} : (\Phi, A) \in \mathcal{R}(\Phi_0, A_0)\}.$$

The conditional saddle path under regime A is the largest subset $\mathcal{M}_A(\Phi_0, A_0) \subseteq \mathcal{R}_A(\Phi_0, A_0)$ that is forward invariant under the frozen-regime mapping:

$$\Phi \in \mathcal{M}_A(\Phi_0, A_0) \implies \Gamma_{endo}(\Phi, A) \in \mathcal{M}_A(\Phi_0, A_0).$$

Equivalently,

$$\mathcal{M}_A(\Phi_0, A_0) := \bigcup \left\{ \mathcal{S} \subseteq \mathcal{R}_A(\Phi_0, A_0) : \Gamma_{\text{endo}}(\mathcal{S}, A) \subseteq \mathcal{S} \right\}.$$

Definition 2 extracts from $\mathcal{R}(\Phi_0, A_0)$ the forward-invariant subset relevant for frozen-regime dynamics. Intuitively, fixing A turns the stochastic economy into a deterministic dynamical system on the distribution space: given Φ , the RCE law of motion produces the next endogenous state $\Gamma_{\text{endo}}(\Phi, A)$. The conditional saddle path $\mathcal{M}_A(\Phi_0, A_0)$ is the maximal set of equilibrium states that is closed under this frozen-regime evolution. As in deterministic saddle-path analysis, restricting attention to \mathcal{M}_A isolates the economically meaningful equilibrium evolution toward the long-run behavior under regime A .

Based on the conditional saddle path, conditional steady state can be defined as follows:

Definition 3 (Conditional steady state).

Given (Φ_0, A_0) , a conditional steady state under regime A is any $\Phi_A^{cs} \in \mathcal{M}_A(\Phi_0, A_0)$ satisfying

$$\Gamma_{\text{endo}}(\Phi_A^{cs}, A) = \Phi_A^{cs}.$$

Assumption 1 (Geometric regularity of conditional saddle paths).

Fix (Φ_0, A_0) . For each $A \in \{B, G\}$, the conditional saddle $\mathcal{M}_A := \mathcal{M}_A(\Phi_0, A_0)$ satisfies:

- (i) (Uniqueness) There exists a unique $\Phi_A^{cs} \in \mathcal{M}_A$ with $\Gamma_{\text{endo}}(\Phi_A^{cs}, A) = \Phi_A^{cs}$.
- (ii) (Stability) $\Gamma_{\text{endo}}^{(t)}(\Phi, A) \rightarrow \Phi_A^{cs}$ for all $\Phi \in \mathcal{M}_A$.
- (iii) (Single-valued) There exists a homeomorphism $\psi_A : (0, 1] \rightarrow \mathcal{M}_A$ with $\psi_A(1) = \Phi_A^{cs}$.

In Assumption 1 (iii), the index set $(0, 1]$ is a labeling device: $\psi_A(1) = \Phi_A^{cs}$ fixes a normalization, but Φ_A^{cs} need not be a boundary point of \mathcal{M}_A in any economic projection (e.g. the curve may approach $K(\Phi_A^{cs})$ from both sides).

Proehl (2025) establishes existence and uniqueness of recursive equilibrium in heterogeneous-agent models under moderate parametric restrictions. Their result concerns global uniqueness and does not depend on the initial condition. By contrast, Assumption 1 (i) in this paper is conditional on the initial state and therefore

does not rule out the possibility of global multiplicity, as in [Walsh and Young \(2024\)](#).² A theoretical characterization of the conditions under which Assumption 1 (i) holds is beyond the scope of this paper; accordingly, the analysis proceeds under this assumption. In addition, Assumption 1 (ii) restricts attention to non-explosive equilibrium dynamics.

Assumption 1 (iii) is essentially *necessary* for the paper’s notion of global impulse responses and transition functions to be well defined as single-valued equilibrium objects. This is also the implicit convention in most of the macroeconomic literature: impulse responses and transition dynamics are typically reported as *functions* of an initial condition and a shock experiment, not as set-valued correspondences. If the relevant frozen-regime invariant equilibrium set were *thick* in the distribution space, (e.g. higher-dimensional, folded, or branching) then the same reported conditioning information could be consistent with multiple distinct distributions that generate different future allocations under the same experiment. In that case, a *GIRF* or *transition function* would be inherently set-valued unless one conditions on the full distribution and/or imposes an explicit equilibrium selection rule. Assumption 1 (iii) rules out this pathology on the conditional saddle, thereby aligning the analysis with the standard practice of treating global dynamics as uniquely defined paths.

With this geometric regularity in place, the conditional saddle admits a natural notion of “position along the curve” toward the conditional steady state. Later, Theorem 2 formalizes the key implication: if some aggregate equilibrium variable moves strictly one-way along this curve, then it provides a valid coordinate for the entire conditional saddle. In that case, knowing the scalar value is equivalent to knowing the full distributional state within the relevant equilibrium set, so the stochastic equilibrium admits an exact one-dimensional representation on \mathcal{M}_A .

2.2 Representative-agent economy: A canonical RBC

A representative household with temporal log utility is considered. Given initial condition (a_0, A_0) , the household maximizes life-time 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 (including an appropriate transversality condition), this recursive formulation is equivalent to the sequential

²Nevertheless, conditional on any given initial state, the equilibrium is assumed to be unique.

equilibrium; for completeness the sequential formulation is provided in Appendix A. The recursive form of the household's problem is as follows:

$$v(a; K, A) = \max_{c, a'} \log(c) + \beta \mathbb{E}v(a'; K', A') \quad (2)$$

$$c + a' = a(1 + r(X)) + w(X) \quad (3)$$

$$a' \geq -\bar{a} \quad (4)$$

$$K' = \Gamma_{endo}(K, A), \quad A' \sim \Gamma_{exo}(A'|A), \quad (5)$$

where v is the household's value function; a is the wealth in 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 the levels B and G , where $G > B$:

$$\Gamma_{exo} = \begin{bmatrix} \pi_{BB} & \pi_{BG} \\ \pi_{GB} & \pi_{GG} \end{bmatrix} \quad (6)$$

I consider the following competitive factor prices given CRS Cobb-Douglas production function:

$$r(K, A) = A\alpha K^{\alpha-1} - \delta \quad (7)$$

$$w(K, A) = A(1 - \alpha)K^\alpha, \quad (8)$$

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

Definition 4 (Recursive competitive equilibrium).

$(c, a', v, r, w, \Gamma_{endo})$ is a recursive competitive equilibrium if these functions

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

$$a'(K; K, A) = K' = \Gamma_{endo}(K, A) \quad (9)$$

In the model, the household ex-ante takes into account the aggregate uncertainty

in its decision. The conditional saddle path is the sequence of outcomes implied by such ex-ante decisions when the ex-post exogenous aggregate states are *frozen* at A . Figure 1 illustrates two conditional saddle paths corresponding to the aggregate productivity states B and G . The figure is based on a standard quarterly calibration of the RBC model described above. The model is solved globally using the repeated transition method of Lee (2025). Conditional saddle paths are then recovered by fixing the aggregate productivity state for 2,000 periods and tracing the resulting equilibrium trajectories of capital and consumption.³

The global dynamics of the recursive competitive equilibrium are fully characterized by movements along and across these conditional saddle paths. When aggregate productivity remains constant, the economy evolves along the corresponding saddle path. If the economy is initialized between the two conditional saddle paths and $A = G$, capital converges strictly monotonically upward toward the conditional steady state K_G^{cs} ; if $A = B$, capital converges strictly monotonically downward toward K_B^{cs} . If the initial state lies outside this region, the economy gradually transitions into it while moving along the relevant saddle paths.

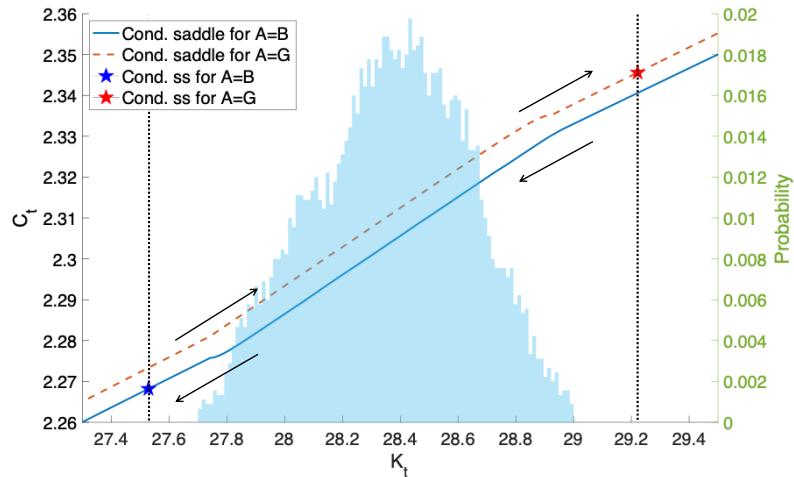


Figure 1: Conditional saddle paths and steady states

Notes: The figure plots conditional saddle paths for $A = B$ (solid) and $A = G$ (dashed) implied in the canonical RBC model under standard quarterly calibration. The histogram in the background plots the time-series distribution of the aggregate capital stock.

³For each productivity state, the dynamics are computed twice: once starting from a sufficiently low initial capital stock to obtain the upward-converging path, and once from a sufficiently high initial capital stock to obtain the downward-converging path.

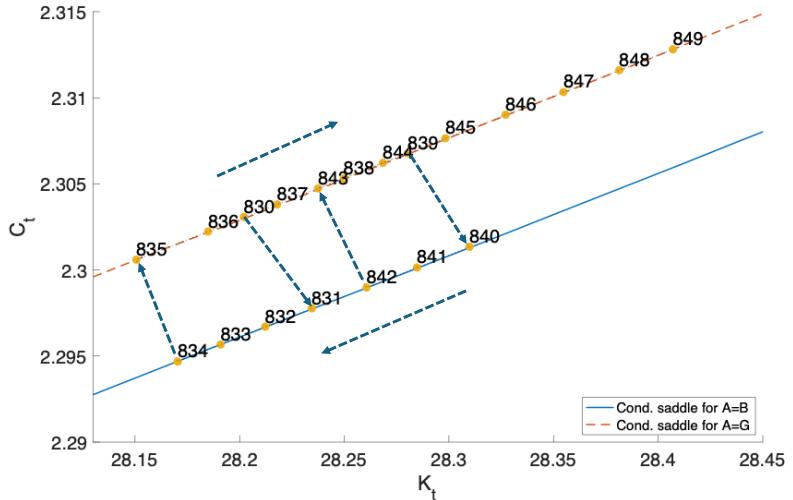
Along the conditional saddle, the economy evolves according to the law of motion Γ_{endo} for the endogenous state. In particular, the aggregate capital K strictly monotonically converges to the conditional steady state — if K is lower than a conditional steady state of the conditional saddle the economy is in, it increases, and vice versa for the case of K higher than it. This strictly monotonically converging pre-determined variable is the key to link the heterogeneous-agent model and the representative-agent model as in the current setup — the variable becomes the sufficient statistic for the heterogeneous-agent model.

When the level of TFP exogenously changes, the economy jumps from one saddle to the other vertically — capital responds from the following period as it is a pre-determined variable one period ahead. In this dynamics, the key force to move the economy is the law of motion for the exogenous state Γ_{exo} . Combined together, both law of motions generate bounded stochastic fluctuations in the economy.

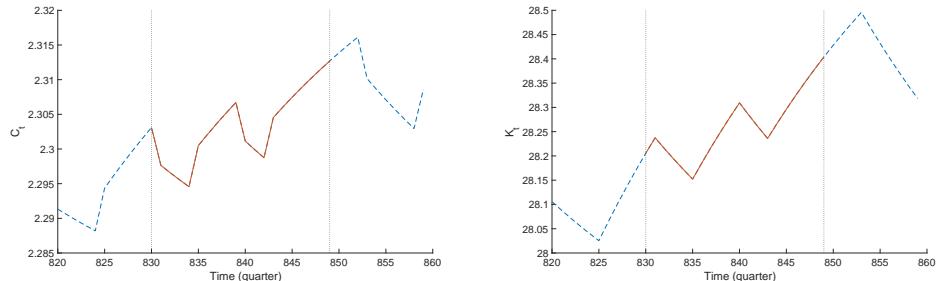
Generalized transition function The conditional saddle path captures the entire transitions across different endogenous (along the saddle) and exogenous states (across the saddle), allowing generalized transition function (GTF) analysis (Lee, 2025), which encompasses generalized impulse response functions (Koop et al., 1996; Andreasen et al., 2017). The economy’s response to a sequence of exogenous shocks is not confined to local dynamics around the steady state, and its trajectory is sharply traced in the phase diagram.

Figure 2 illustrates the equilibrium dynamics in consumption and capital stock in phase diagram (panel (a)) and in time domain (panels (b) and (c)) for an arbitrary subsample period (830 – 849, 20 quarters) in the equilibrium path. In period 830, an exogenous negative TFP shock hits leading to a downward jump in consumption across the conditional saddle paths. Then, for four consecutive periods, consumption and capital stock endogenously decline along the saddle. In period 835, an exogenous positive TFP shock hits, shifting the economy to the other conditional saddle path, followed by upward endogenous adjustments in consumption and capital. Compared to the time-domain figures, phase diagram more concisely illustrates the equilibrium dynamics in the economy within a figure. Moreover, it enables to immediately consider the counter-factual scenarios from eyeballing off the realized exogenous paths.

As the information on all possible global equilibrium dynamics are contained in the conditional saddle paths, generalized impulse response function (GIRF) can be also



(a) On the saddles



(b) C - time domain (c) K - time domain

Figure 2: Equilibrium dynamics: saddle vs. time domain

Notes: The figure illustrates the stochastic equilibrium dynamics of consumption and capital stock on the conditional saddle paths (panel (a)) and in the time domain (panels (b) and (c)) implied by a canonical RBC model under standard quarterly calibration. The sample period covers 20 quarters from 830 to 849.

conveniently represented. For the GIRF illustration I replace the two-state Markov chain with a continuous AR(1) process for A : $A' = (1 - \rho)\mu + \rho A + \epsilon$, $\epsilon \sim N(0, \sigma)$. 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.

Figure 3 illustrates the post-shock dynamics in the phase diagram following a positive TFP shock, as aggregate productivity gradually reverts to its steady-state level. Upon impact, consumption jumps upward, inducing a vertical shift of the economy across conditional saddle paths. In subsequent periods, forces operating

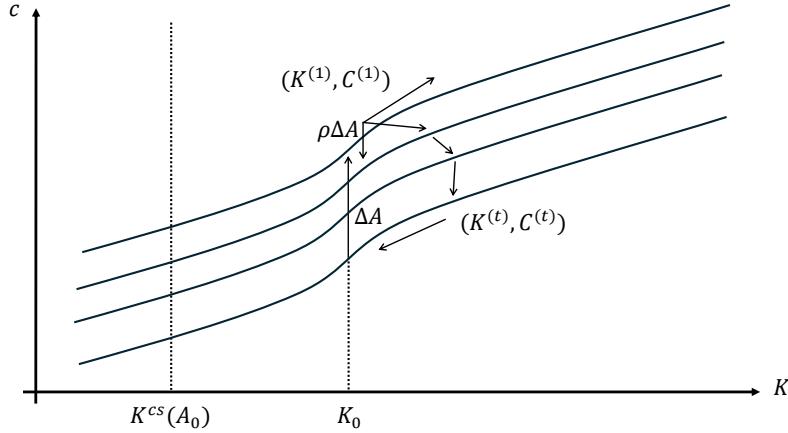


Figure 3: Generalized impulse response function (GIRF) at (K_0, A_0)

Notes: The figure illustrates the generalized impulse responses of consumption and capital stock on the conditional saddle paths to a positive TFP shock.

both along and across these saddles jointly shape the post-shock dynamics, generating a right-bending, bow-shaped trajectory. Consequently, the capital stock K exhibits an inverted-U-shaped impulse response over time, while consumption c displays a downward-sloping response in the time domain.

Figure 4 presents the simulated impulse responses of consumption and capital following a positive TFP shock, shown in the phase diagram (panel (a)) and in the time domain (panels (b) and (c)) over periods 1–50. Consistent with the phase-diagram intuition, the bow-shaped capital trajectory translates into a hump-shaped impulse response in the time domain, as shown in panel (c). Consumption also exhibits a hump-shaped response: initially, high-TFP conditional saddle paths push consumption upward, but as TFP gradually decays and crosses a threshold level, the direction of adjustment reverses, pulling consumption back down.

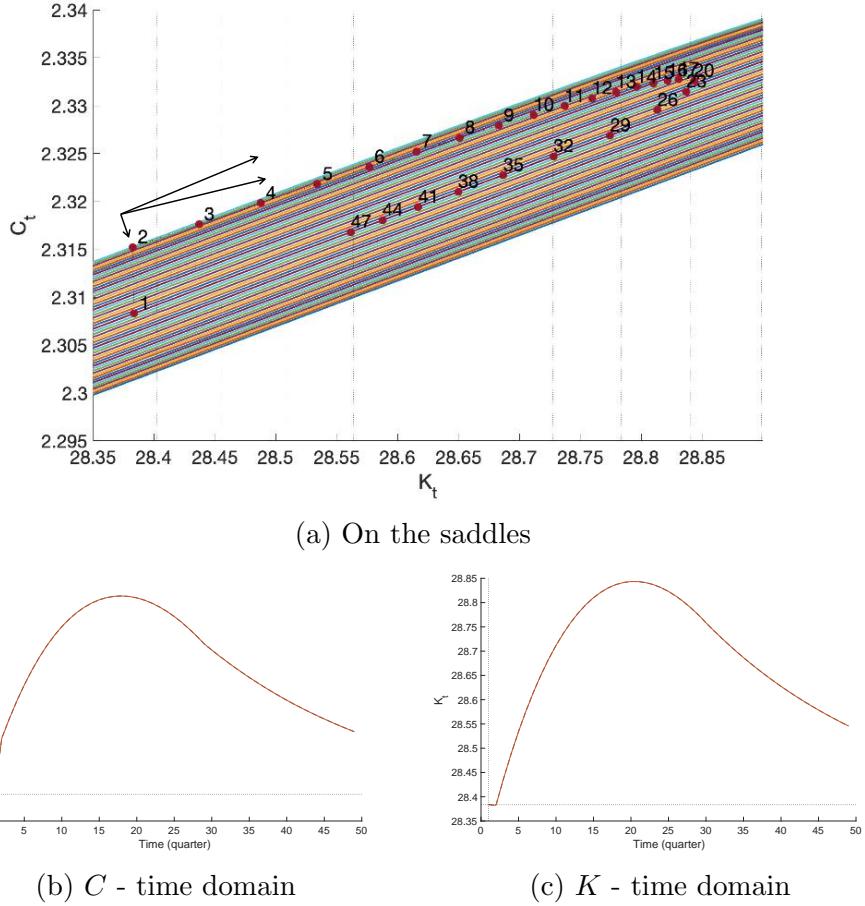


Figure 4: Impulse responses to a positive TFP shock: saddle vs. time domain

Notes: The figure illustrates the impulse responses of consumption and capital stock on the conditional saddle paths (panel (a)) and in the time domain (panels (b) and (c)). The sample period covers 50 quarters from 1 to 50.

Comparison with the perfect-foresight saddles The household's inter-temporal optimality conditions evaluated at the conditional steady states are as follows:

$$\begin{aligned}
 & \underbrace{\beta\pi_{BB}(1 + \alpha B(K_B^{cs})^{\alpha-1} - \delta) + \beta\pi_{BG} \left(\frac{c(K_B^{cs}, G)}{c_B^{cs}} \right)^{-1} (1 + \alpha G(K_B^{cs})^{\alpha-1} - \delta)}_{\text{Discounted return for } B \rightarrow B} = 1 \quad (10) \\
 & \underbrace{\beta\pi_{GG}(1 + \alpha G(K_G^{cs})^{\alpha-1} - \delta) + \beta\pi_{GB} \left(\frac{c(K_G^{cs}, B)}{c_G^{cs}} \right)^{-1} (1 + \alpha B(K_G^{cs})^{\alpha-1} - \delta)}_{\text{Discounted return for } G \rightarrow G} = 1 \quad (11)
 \end{aligned}$$

These equations pin down the conditional steady-state conditions for the aggregate capital stocks K_B^{cs} and K_G^{cs} :

$$K_B^{cs} = \left(\frac{\alpha(\pi_{BB}B + \pi_{BG}G)}{1/\beta - (1-\delta)(\pi_{BB} + \pi_{BG}(g_c^{cs}(B))^{-1})} \right)^{\frac{1}{1-\alpha}} \quad (12)$$

$$K_G^{cs} = \left(\frac{\alpha(\pi_{GG}G + \pi_{GB}B)}{1/\beta - (1-\delta)(\pi_{GG} + \pi_{GB}(g_c^{cs}(G))^{-1})} \right)^{\frac{1}{1-\alpha}} \quad (13)$$

where $g_c^{cs}(A)$, $A \in \{B, G\}$ are gross consumption growth rates when a regime change happens in each conditional steady state.

In the *perfect-foresight (PF)* economy, agents believe regime switching is impossible and, consistent with this belief, the regime does not switch. In contrast, along the *conditional-saddle (CS)* dynamics, the regime is held fixed ex post for the purpose of tracing the conditional manifold, but agents still correctly anticipate that regime switches are possible under Π . Thus PF and CS share the same feasibility accounting and the same K -nullcline under a given productivity level, yet generally differ in their Euler-implied consumption dynamics and hence in their steady states and saddle paths.

From the stationary condition for the capital stock ($\delta K = I$) and the national accounting identity ($Y = C + I$), the conditional nullclines of aggregate capital K for $A \in \{B, G\}$ are as follows:

$$c_B^{Kcnnull} = B(K_B^{Kcnnull})^\alpha - \delta K_B^{Kcnnull} \quad (14)$$

$$c_G^{Kcnnull} = G(K_G^{Kcnnull})^\alpha - \delta K_G^{Kcnnull} \quad (15)$$

Then, I consider a perfect-foresight dynamics where there is no uncertainty in the economy and TFP A is fixed at either B or G . I denote allocations for this perfect-foresight economy with superscript pf . Following a canonical neoclassical growth model analysis, I derive the steady states for the perfect-foresight counterpart with different TFP levels:

$$(K_B^{pf}, c_B^{pf}) = \left(\left(\frac{\alpha B}{1/\beta + \delta - 1} \right)^{\frac{1}{1-\alpha}}, B \left(\frac{\alpha B}{1/\beta + \delta - 1} \right)^{\frac{\alpha}{1-\alpha}} - \delta \left(\frac{\alpha B}{1/\beta + \delta - 1} \right)^{\frac{1}{1-\alpha}} \right) \quad (16)$$

$$(K_G^{pf}, c_G^{pf}) = \left(\left(\frac{\alpha G}{1/\beta + \delta - 1} \right)^{\frac{1}{1-\alpha}}, G \left(\frac{\alpha G}{1/\beta + \delta - 1} \right)^{\frac{\alpha}{1-\alpha}} - \delta \left(\frac{\alpha G}{1/\beta + \delta - 1} \right)^{\frac{1}{1-\alpha}} \right) \quad (17)$$

Notably, two economies share the same capital (conditional) nullclines for each productivity levels, as formalized in Proposition 1.

Proposition 1 (K -nullcline invariance over beliefs).

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

Proof.

From the stationary condition $\delta K = I$, equations (14) and (15) are immediate for both models. Therefore, the conditional nullclines are the same. \blacksquare

This property provides two important insights. First, the differences in the limit behaviors of two economies (steady states) and the corresponding saddle paths are crucially determined by the consumption dynamics – consumption nullclines.⁴ In Proposition 2, I compare the rankings between the baseline model and the perfect-foresight model.

Second, both with and without aggregate uncertainty, the saddle-path dynamics require the capital stock K to converge strictly monotonically to the (conditional) steady state. Even when the baseline model is extended to a heterogeneous-household environment, as in Section 4, the K -nullcline remains invariant, preserving the monotonicity of K along conditional saddle paths. This property is crucial for establishing the sufficiency of K in summarizing aggregate fluctuations in the heterogeneous-household setting.

Proposition 2 (Aggregate uncertainty and the conditional steady states).

The following inequalities hold:

$$K_B^{cs} < K_B^{pf} < K_G^{pf} < K_G^{cs}, \quad c_B^{cs} < c_B^{pf} < c_G^{pf} < c_G^{cs}.$$

Proof.

See Appendix C. \blacksquare

Proposition 2 shows that the (conditional) steady state associated with low TFP features lower capital stock and consumption under the aggregate uncertainty than

⁴Rigorously speaking, there is no consumption nullcline in the model with aggregate uncertainty. The Euler equations at the conditional steady states play the same role as the consumption nullclines without uncertainty.

under perfect foresight; the one with high TFP features the opposite. Therefore, conditional steady states of capital and consumption with aggregate uncertainty are nested by the counterparts in the steady state with perfect foresight. These are fully driven by rational expectations for future regime shifting under aggregate uncertainty.

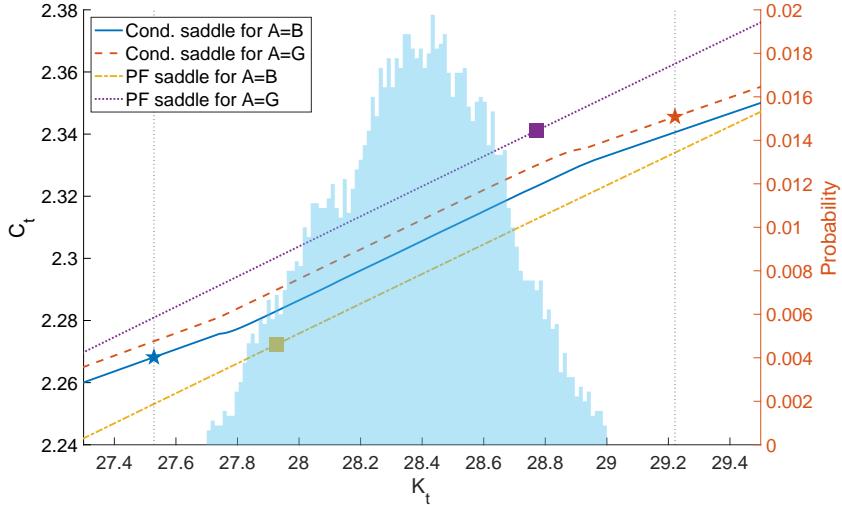


Figure 5: Conditional saddle path comparison: with and without uncertainty

Notes: The figure plots 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) implied by a canonical RBC model under standard quarterly calibration. The histogram in the background plots the time-series distribution of the aggregate capital stock.

Figure 5 plots the conditional saddle paths for the model with aggregate uncertainty and the perfect-foresight deterministic saddle paths. As shown in Proposition 2, conditional steady states for both capital stocks and consumption with perfect foresight are nested by the ones with aggregate uncertainty.

Conditional boundary condition In the canonical deterministic neoclassical growth model, the sequential formulation of competitive equilibrium requires a transversality condition to pin down a unique equilibrium path.⁵ The transversality condition guarantees bounded equilibrium paths and ensures the equivalence between the sequential and recursive formulations of equilibrium. In this paper, I define the stochastic equilibrium recursive form without explicitly constructing the sequential counterpart,

⁵This condition is imposed at the aggregate level and should be distinguished from individual no-Ponzi constraints.

implicitly assuming a stochastic analogue of the transversality condition.

Given the notion of conditional saddle paths, a natural stochastic counterpart to the deterministic boundary condition is conditional transversality:

$$\lim_{t \rightarrow \infty} \beta^t u'(c_t) K_t = 0 \quad \text{for a TFP frozen at } A \in \{B, G\} \quad (18)$$

(along the conditional saddle for each $A \in \{B, G\}$).

That is, conditional transversality requires lifetime utility to remain bounded along each conditional steady-state path. This requirement is consistent with the internal consistency of recursive competitive equilibrium, in which agents correctly anticipate the long-run evolution of the economy conditional on the current aggregate state. If a boundary condition were not conditional-saddle-specific — so that some realizations of exogenous states generated explosive paths — the value function would fail to be well defined along those paths.

The conditional transversality condition (18) is observationally equivalent to the standard transversality conditions (Kamihigashi, 2003, 2005), which requires boundedness for all possible optimal paths — both select identical bounded recursive competitive equilibria. The conditional formulation, however, makes explicit the role of the boundary condition in ensuring bounded dynamics along each conditional saddle, which is essential for the dimension-reduction results that follow.

Economies without conditional saddle paths A conditional saddle requires a predetermined endogenous state: without one, the economy jumps each period to the unique bounded allocation implied by the current exogenous state, leaving no nontrivial “along-the-saddle” dynamics.

Remark 1 (Endogenous memorylessness implies degenerate conditional saddle paths). *Consider an RCE with state (Φ_t, A_t) . Call the economy endogenously stateless if the endogenous state is memoryless: $\Phi_t = \phi(A_t)$ for some function ϕ , so that all equilibrium allocations depend on A_t only. Then the economy is saddleless: the “conditional steady state” under frozen A is simply $\phi(A)$, and the invariant set is a singleton. Conversely, a non-degenerate conditional saddle requires at least one predetermined endogenous state.*

Example: Gali (2008) three-equation NK model. The textbook New Keynesian model

$$x_t = \mathbb{E}_t x_{t+1} - \frac{1}{\sigma} (i_t - \mathbb{E}_t \pi_{t+1} - r_t^n), \quad (19)$$

$$i_t = \varphi_\pi \pi_t + \varphi_x x_t + \nu_t, \quad (20)$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \kappa x_t + u_t, \quad (21)$$

under determinacy yields $(x_t, \pi_t, i_t) = \Psi s_t$ for $s_t := (r_t^n, u_t, \nu_t)$ and some matrix Ψ . The equilibrium is endogenously stateless: freezing the exogenous state collapses dynamics to an immediate jump to the conditional steady point. Fluctuations occur entirely *across* conditional points, not along a saddle.

Models with predetermined states—capital, habits, interest-rate smoothing, or distributional states in HANK—are endogenously stateful, and conditional saddle paths generically exist.

3 State dependence in a shock response

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, a response of an aggregate variable to an exogenous shock is represented by shifts across different conditional saddle paths. Then, if and only if all the conditional saddle paths are parallel along the endogenous state, the response becomes state-independent.

Theorem 1 (State-(in)dependence as a geometric restriction).

Fix (Φ_0, A_0) and let $\mathcal{M}_{A_0}(\Phi_0, A_0)$ 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 $(\Phi, A) \in \mathcal{M}_{A_0}(\Phi_0, A_0) \times \{A_0, A_1\}$. Define the impact gap between regimes A_1 and A_0 at state Φ by

$$\Delta_g(\Phi; A_1, A_0) := g(\Phi, A_1) - g(\Phi, A_0).$$

Then the following are equivalent:

(i) (State-independent gap) $\Delta_g(\Phi; A_1, A_0)$ is constant on $\mathcal{M}_{A_0}(\Phi_0, A_0)$, i.e. there

exists $c \in \mathbb{R}$ such that

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

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

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

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): if $g(\Phi, A_1) = g(\Phi, A_0) + c$ for all Φ , then $(\Phi, g(\Phi, A_1)) = (\Phi, g(\Phi, A_0) + c)$ for all Φ , hence $\mathcal{G}_{A_1} = \mathcal{G}_{A_0} + (0, c)$. (ii) \Rightarrow (i): if $\mathcal{G}_{A_1} = \mathcal{G}_{A_0} + (0, c)$, then for each Φ we must have $g(\Phi, A_1) = g(\Phi, A_0) + c$, so the gap is constant. \blacksquare

Sharp state independence may be a knife-edge property of an RCE. However, Theorem 1 provides an insight regarding conditions under which state dependence becomes amplified. In particular, when a conditional saddle path is more steeply tilted with respect to the endogenous state, as illustrated in Figure 6, a shock responsiveness becomes state dependent. Such differential slopes may arise from various real (Winberry, 2021; Lee, 2025), financial (Melcangi, 2024), labor market frictions (Petrosky-Nadeau et al., 2018; Pizzinelli et al., 2020), and the scope of the relevant shocks include TFP shocks and fiscal/monetary policy shocks (Tenreyro and Thwaites, 2016; Lee, 2025). The sufficiency and the necessity of tilt in the state dependence implies that models with such nature necessarily implies different slopes of the conditional saddle paths.

An example: asymmetric adjustment cost As an example where state dependence arises due to the tilt in the conditional saddle, I consider an extended RBC model with asymmetric adjustment cost. Specifically, the representative household's

budget constraint is modified in the following way:

$$c + a' + \mathcal{C}(a', a) = a(1 + r(X)) + w(X) \quad (22)$$

$$\mathcal{C}(a', a) = \frac{\tilde{\mu}}{2} \left(\frac{a' - a}{a} \right)^2 a \quad (23)$$

$$\tilde{\mu} = \begin{cases} \mu_+ & \text{if } a' > a \\ \mu_- & \text{if } a' < a \end{cases} \quad (24)$$

where \mathcal{C} is the wealth adjustment cost which indirectly reflects the frictional capital market. The adjustment cost is asymmetric between positive and negative investment, as specified in equation (24). In particular, I consider the case $\mu_+ > \mu_-$.⁶

Figure 6 plots the conditional saddle paths under asymmetric adjustment costs. Relative to the saddle path associated with $A=G$, the conditional saddle path for $A=B$ is substantially steeper. As a result, a one-standard-deviation TFP shock generates a larger consumption response when the capital stock is low (3.42%) than when it is high (2.42%).

This geometric representation formalizes the intuition that “climbing up is difficult, falling down is easy”—a recurring theme in the literature on asymmetric business cycles and endogenous disasters (Petrosky-Nadeau et al., 2018). In this framework, asymmetric frictions generate differently tilted conditional saddles: the steeper slope under low productivity implies an aggressive contraction across manifolds upon negative shocks, while recovery proceeds more slowly along a flatter path. The resulting gap in shock responsiveness (3.42% versus 2.42% in the example above) arises precisely because the geometric distance between conditional saddles varies with the economy’s position along them.

This geometric framework suggests that whenever a model in the literature features endogenous state-dependent dynamics, such behavior must fundamentally originate from the differential tilt of the conditional saddle paths. Theorem 1 establishes that state independence is equivalent to a vertical-translation geometry, in which equilibrium manifolds remain parallel along the endogenous state. It follows that any deviation into state-dependent responsiveness—regardless of whether the underlying friction is financial, labor-related, or based on adjustment costs—necessarily implies

⁶The adjustment-cost parameters are not calibrated. For illustrative purposes, I set $\mu_+ = 4$ and $\mu_- = 1$. All other parameters follow a standard quarterly RBC calibration.

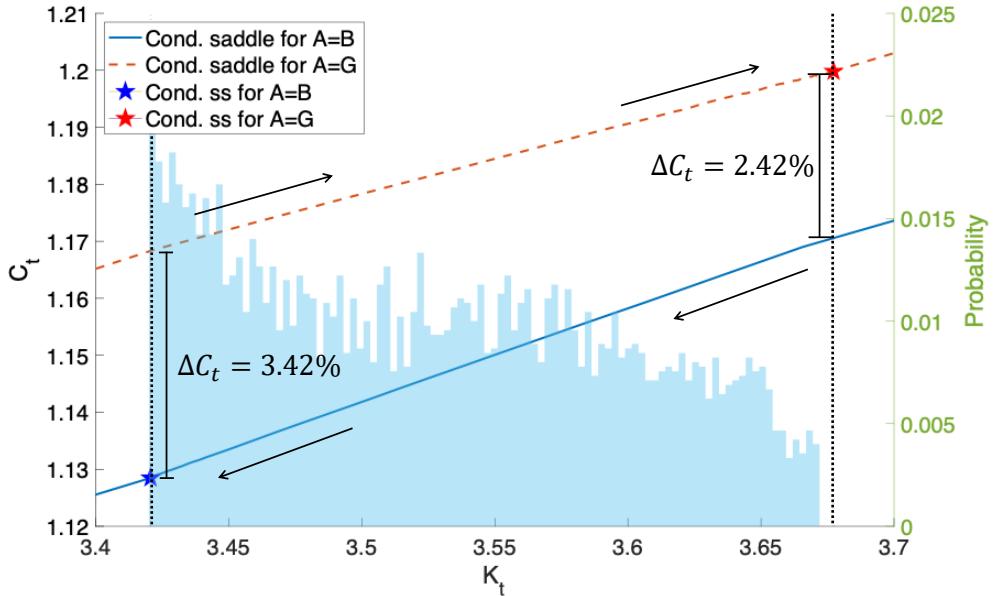


Figure 6: Differently tilted conditional saddle paths – endogenous state dependence

Notes: The figure plots conditional saddle paths for $A = B$ (solid) and $A = G$ (dashed) implied by an RBC model with asymmetric wealth adjustment cost. The conditional saddle path for $A = B$ is steeper than the one for $A = G$. The histogram in the background plots the time-series distribution of the aggregate capital stock.

a geometric structure where the local slopes of the conditional saddles diverge across aggregate regimes. In this sense, the “tilt” (the local rate of endogenous accumulation or decumulation) serves as the universal geometric signature of endogenous state dependence.

Furthermore, this framework clarifies the distinction between endogenous and exogenous state dependence. While endogenous state dependence—the sensitivity of shocks to the current endogenous state—necessarily originates from the differential tilt of the conditional saddle paths, exogenous state dependence arises from the interaction and magnitude of the shifts across these manifolds. Even in economies where saddles are geometrically parallel—resulting in state-independent impact gaps across capital stocks—the equilibrium path may still exhibit nonlinear responses if the exogenous states are “outstretched.” In such cases, the vertical distance between saddles reflects a larger exogenous regime shift, creating greater fluctuations that are decomposed into horizontal propagation along the saddles and vertical jumps across them. Consequently, the conditional saddle representation provides a unified visual

tool for assessing both the slope-driven propagation of endogenous frictions and the level-driven impact of exogenous aggregate shocks.

4 Heterogeneous-household economy and dimension reduction

This section introduces conditional saddle path in a canonical heterogeneous-household economy. I consider a continuum of unit measure of ex-ante homogeneous households. The recursive formulation of the households' problem is as follows:

$$v(a, z; \Phi, A) = \max_{c, a'} \log(c) + \beta \mathbb{E}v(a', z'; \Phi', A') \quad (25)$$

$$c + a' = a(1 + r(X)) + w(X)z \quad (26)$$

$$a' \geq 0 \quad (27)$$

$$\Phi' = \Gamma_{endo}(\Phi, A), \quad A' \sim \Gamma_{exo}(A'|A). \quad (28)$$

The problem is the same as in the representative-household economy except for 1) uninsurable idiosyncratic labor productivity, which follows a Markov process $z \sim \Gamma_z(z'|z)$; 2) inclusion of distribution of individual states Φ in the aggregate endogenous state; and 3) the corresponding change in the law of motions for the endogenous aggregate state. The RCE is defined as in [Krusell and Smith \(1998\)](#).

The model includes two different stochastic exogenous processes: idiosyncratic productivity and aggregate TFP. Therefore, there are two layers of conditional saddle paths: one is individual conditional saddle path, and the other is aggregate conditional saddle path. The individual saddle path has its own cross-sectional implication which deserves a separate analysis, but it is out of this paper's focus. So, the corresponding analysis is included in Appendix. I elaborate on the aggregate conditional saddle with stochastic TFP process, where the model closely follows [Krusell and Smith \(1998\)](#).⁷

[Krusell and Smith \(1998\)](#) posits an endogenous law of motion that tracks aggregate capital K rather than the full distribution Φ : $\log K' = \alpha(A) + \beta(A) \log K$, for $A \in \{B, G\}$, where α and β are state-dependent coefficients. This formulation embeds

⁷In the original model of [Krusell and Smith \(1998\)](#), the exogenous individual labor supply co-moves with the exogenous aggregate TFP. All the results stay unaffected after including this feature in the model, but for the expositional brevity, I assume the labor supply is exogenously fixed.

two key assumptions: (i) that aggregate capital K is (approximately) sufficient to summarize endogenous aggregate dynamics, and (ii) that the law of motion is log-linear. Using the conditional-saddle framework, this paper establishes that K is an exact sufficient statistic. By contrast, log-linearity follows from the specific model structure in [Krusell and Smith \(1998\)](#) and typically requires numerical verification. The following theorem provides the first step toward establishing the exact sufficiency of K .

Theorem 2 (Monotone aggregate variable as a coordinate).

Given (Φ_0, A_0) , suppose an aggregate equilibrium variable $e : \mathcal{M}_A(\Phi_0, A_0) \rightarrow \mathbb{R}$ varies strictly monotonically, i.e., there exists a homeomorphism $\psi_A : (0, 1] \rightarrow \mathcal{M}_A(\Phi_0, A_0)$ such that the function $s \mapsto e(\psi_A(s))$ is strictly monotone on $(0, 1]$. Then e uniquely indexes endogenous states on the conditional saddle:

$$e(\Phi) = e(\Phi') \implies \Phi = \Phi'. \quad (29)$$

Hence there exists a (unique) inverse map ϕ_A such that $\phi_A(e(\Phi)) = \Phi$ on $\mathcal{M}_A(\Phi_0, A_0)$, and any equilibrium object restricted to the conditional saddle can be written as a function of the single scalar e .

Proof. Because \mathcal{M}_A is parameterized by $\psi_A : (0, 1] \rightarrow \mathcal{M}_A$, strict monotonicity of $s \mapsto e(\psi_A(s))$ implies injectivity of e on \mathcal{M}_A . Since $e : \mathcal{M}_A \rightarrow e(\mathcal{M}_A)$ is surjective by definition of the image, injectivity implies that e is a bijection between \mathcal{M}_A and $e(\mathcal{M}_A)$. Hence an inverse $\varphi_A : e(\mathcal{M}_A) \rightarrow \mathcal{M}_A$ exists. \blacksquare

The intuition behind the theoretical result is as follows: Fix a regime $A \in \{B, G\}$ and restrict attention to the set of equilibrium states that lie on the corresponding conditional saddle (the invariant equilibrium manifold). Starting from any initial equilibrium state Φ_0 on this set, the equilibrium law of motion determines a unique subsequent sequence $\{\Phi_t\}_{t \geq 0}$ that remains on the invariant conditional saddle path. Then, it converges to the unique conditional steady state. The key geometric observation is that the conditional saddle is effectively a string: it can be viewed as a single curve of equilibrium states. As a consequence, points on the conditional saddle are naturally ordered by their position along this curve.

Now consider the aggregate equilibrium variable e . By assumption, along any equilibrium history on the conditional saddle, the scalar $e_t \equiv e(\Phi_t)$ moves strictly

monotonically over time and converges to its conditional steady-state value. Intuitively, e acts as a “progress meter” along the conditional saddle: it always moves in one direction toward its limiting value and never reverses.

Suppose, toward a contradiction, that e were not injective on the conditional saddle. Then there exist two distinct equilibrium states $\Phi \neq \Phi'$ on the conditional saddle such that $e(\Phi) = e(\Phi')$. Consider the two equilibrium histories generated by these initial states. Since both histories begin with the same value of the progress meter and e_t must move strictly one-way toward the same limit along the conditional saddle, the two histories would have to remain synchronized in their progress toward the conditional steady state. But this is impossible because distinct starting points on a one-dimensional conditional saddle generate equilibrium histories that cannot merge. This contradiction implies that no two distinct equilibrium states on the conditional saddle can share the same value of e .

Therefore e uniquely labels equilibrium states on the conditional saddle: if $e(\Phi) = e(\Phi')$, then $\Phi = \Phi'$. This establishes the injectivity claim in Theorem 2 and explains why a strictly monotone, convergent aggregate variable can serve as a one-dimensional indexing variable for equilibrium dynamics on the conditional saddle.

Remark 2 (Injectivity and sufficiency).

Because e is injective on \mathcal{M}_A , any equilibrium object restricted to \mathcal{M}_A can be written as a function of (e, A)

Specifically, the following variables can be defined:

$$v^e(\cdot, \cdot; \tilde{e}, A) \in V(\tilde{e}, A) := \{v(\cdot, \cdot; \Phi, A) \mid \forall \Phi \in \mathcal{M}_A(\Phi_0, A_0) \text{ s.t. } e(\Phi) = \tilde{e}\} \quad (30)$$

$$r^e(\tilde{e}, A) \in R(\tilde{e}, A) := \{r(\Phi, A) \mid \forall \Phi \in \mathcal{M}_A(\Phi_0, A_0) \text{ s.t. } e(\Phi) = \tilde{e}\} \quad (31)$$

$$w^e(\tilde{e}, A) \in W(\tilde{e}, A) := \{w(\Phi, A) \mid \forall \Phi \in \mathcal{M}_A(\Phi_0, A_0) \text{ s.t. } e(\Phi) = \tilde{e}\} \quad (32)$$

$$\Gamma_{endo}^e(\tilde{e}, A) \in G(\tilde{e}, A) := \{e(\Gamma(\Phi, A)) \mid \forall \Phi \in \mathcal{M}_A(\Phi_0, A_0) \text{ s.t. } e(\Phi) = \tilde{e}\}. \quad (33)$$

V, R, W and G are nonempty by Assumption 1 and singletons by Theorem 2. Therefore, the recursive problem below is equivalent to the original recursive formulation

in equilibrium, as they yield the same equilibrium allocations.

$$v^e(a, z; e, A) = \max_{c, a'} \log(c) + \beta \mathbb{E} v^e(a', z'; e', A') \quad (34)$$

$$c + a' = a(1 + r^e(e, A)) + w^e(e, A)z \quad (35)$$

$$a' \geq 0 \quad (36)$$

$$e' = \Gamma_{endo}^e(e, A), \quad A' \sim \Gamma_{exo}(A'|A), \quad (37)$$

Now, I show that the conditional K nullclines are the same as in the representative agent model in the following proposition.

Proposition 3 (Conditional K -nullclines of Krusell and Smith (1998)).

The heterogeneous household model's conditional K nullclines are identical to the counterparts of the model with the representative household and invariant over the initial distribution Φ_0 :

$$C_B^{K_{null}} = B(K_B^{K_{null}})^\alpha - \delta K_B^{K_{null}} \quad (38)$$

$$C_G^{K_{null}} = G(K_G^{K_{null}})^\alpha - \delta K_G^{K_{null}} \quad (39)$$

where $C_A^{K_{null}}$ is aggregate consumption.

Proof.

As in Proposition 1, the stationary condition $\delta K = I$ immediately implies the form of the conditional nullclines. ■

Proposition 3 shows that the K -nullcline in the Krusell-Smith model is distribution-invariant. While this provides strong intuition for monotonicity, it does not by itself rule out discrete-time pathologies where the saddle path touches the nullcline at an intermediate point $\Phi \neq \Phi_A^{cs}$ (a “turning point”) or crosses it repeatedly (a “spiral”). Such behaviors correspond to folds or loops in the (K, C) phase diagram, which would defeat a single-valued global representation indexed by K .

To rule out these pathologies rigorously, I impose the *single-crossing condition* (SC):

$$C(\Phi, A) = C_A^{K_{null}}(K(\Phi)) \implies \Phi = \Phi_A^{cs} \quad \forall \Phi \in \mathcal{M}_A. \quad (\text{SC})$$

In continuous-time formulations, this condition is often redundant because direction reversals require \dot{K} to pass through zero continuously, which by uniqueness of the

steady state occurs only at Φ_A^{cs} . Condition (SC) is the necessary discrete-time analogue to ensure well-behaved saddle geometry.⁸

Remark 3 (Verifying the single-crossing condition).

Condition (SC) has a direct geometric interpretation: the saddle path in the (K, C) phase plane must not cross the K -nullcline $C = AK^\alpha - \delta K$ except at the steady state. This rules out:

1. **Spirals:** Trajectories orbiting the steady state, crossing the nullcline repeatedly.
2. **Wobbles:** Trajectories reversing direction (turning points) away from the steady state.

Figure 7 confirms that in the canonical Krusell-Smith economy, the conditional saddle paths lie entirely within the “northeast” and “southwest” cones defined by the nullcline, satisfying (SC).

Proposition 4 (K monotonicity and injectivity).

Fix $A \in \{B, G\}$. Under Assumption 1 and Condition (SC), aggregate capital converges to K_A^{cs} strictly monotonically along the frozen-regime continuation. Consequently, K is injective on \mathcal{M}_A .

Proof.

Define $\Delta K(\Phi) := C_A^{K\text{-null}}(K(\Phi)) - C(\Phi, A)$, so that $K_{t+1} - K_t = \Delta K(\Phi_t)$. By (SC), $\Delta K(\Phi) = 0$ if and only if $\Phi = \Phi_A^{cs}$.

Define $\delta := \Delta K \circ \psi_A : (0, 1] \rightarrow \mathbb{R}$. Since ψ_A is a homeomorphism and both K and $C(\cdot, A)$ are continuous, δ is continuous on $(0, 1]$. By (SC), $\delta(s) = 0$ only at $s = 1$.

By the Intermediate Value Theorem, δ cannot change sign on $(0, 1)$ without passing through zero. Hence δ has constant sign on $(0, 1)$: either $\delta(s) > 0$ for all $s < 1$ (so K is strictly increasing) or $\delta(s) < 0$ for all $s < 1$ (so K is strictly decreasing).

Since K moves strictly monotonically along the topological string \mathcal{M}_A , it maps distinct states to distinct capital levels. Thus K is injective on \mathcal{M}_A . ■

Proposition 4 establishes that aggregate capital K is an exact sufficient statistic for endogenous state dynamics in a canonical heterogeneous-household business-cycle

⁸Discrete time allows “jumps” over the nullcline or “turning points” where ΔK reverses sign without passing through a stationary state. Condition (SC) explicitly rules out such pathologies, aligning with the computational practice in heterogeneous-agent macroeconomics and the original Krusell and Smith (1998) framework.

model. To complement this theoretical result, Figure 7 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 2,000 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.

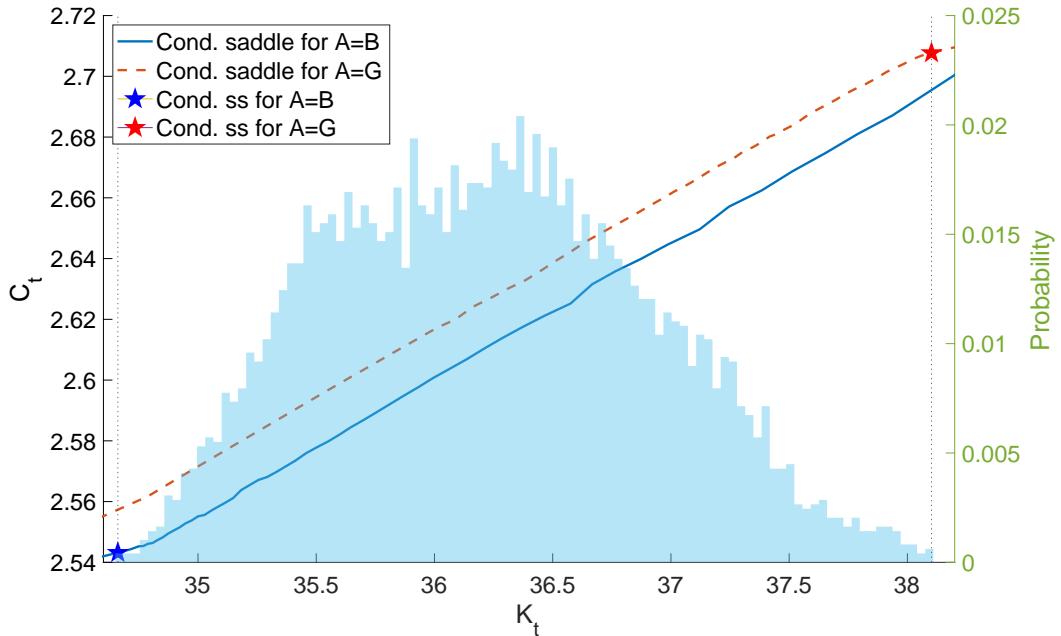


Figure 7: Conditional saddle paths in Krusell and Smith (1998)

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 near rationality to complete rationality Krusell and Smith (1998) assumes the specific parametric law of motion to compute the heterogeneous-household model. Then, they confirm the accuracy of the law of motion through the accurate consistency between the realized and assumed dynamics. After the celebrated contribution, the approach is often labeled as near or bounded rational approach. The conditional-saddle framework shows that, restricted to the conditional saddle, conditioning on K can be exact (a sufficient statistic) for equilibrium objects. The

additional log-linear functional form used by Krusell and Smith (1998) remains a quantitative approximation.

Parametric form of the conditional saddle paths Despite the indexing function K , specific conditional saddle paths and the form of the law of motion remains undetermined. This problem exists even for representative agent business cycle models. Lee (2025) develops repeated transition method that utilizes the recurrence of the equilibrium allocations along the conditional saddle path. For heterogeneous-agent models, the existence of the indexing variable (sufficient statistic) starkly eases the implementation of the RTM. Theorem 2 and Proposition 4 theoretically supports the implementation feasibility of the RTM using the sufficient statistic.

4.1 Extensions

Theorem 2 provides verifiable conditions for dimension reduction: identify the K -nullcline and verify monotonicity. This section applies the diagnostic to several extensions, illustrating both the framework's scope and its limits.

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 sufficient statistic.

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

$$c + a' + q(\Phi, A)b' = a(1 + r(\Phi, A)) + b + w(\Phi, A)z \quad (40)$$

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

$$\int b'(a, z; \Phi, A)d\Phi = 0. \quad (41)$$

As long as the production function and the labor supply is assumed as in the canonical setup as above, the K nullcline stays invariant from the representative counterpart, enabling K sufficiency. Lee (2025) showed that the equilibrium allocations are strictly

monotonically sorted along K through the globally solved computational outcome. This is because the conditional saddle's strict monotone property along K implied by its conditional nullcline is unaffected by the inclusion of bond dimension. It confirms the prediction of Theorem 2.

Models with endogenous labor supply A heterogeneous-household model with endogenous labor supply and CRRA–GHH preferences is a natural extension of the exogenous-labor-supply framework studied in this paper. GHH utility is particularly useful in this context because it eliminates wealth effects in labor-supply decisions, which could otherwise disrupt the monotonicity properties through distributional channels. Consider preferences given by

$$u(c, l_H) = \frac{1}{1-\sigma} \left(c - \frac{\eta}{1 + \frac{1}{\chi}} l_H^{1+\frac{1}{\chi}} \right)^{1-\sigma}, \quad (42)$$

and the budget constraint

$$c + a' = a(1 + r(\Phi, A)) + w(\Phi, A)z l_H. \quad (43)$$

where z denotes idiosyncratic labor productivity. The individual labor-supply optimality condition then implies the following aggregate labor supply $L(\Phi, A)$:

$$l_H(a, z; \Phi, A) = \left(\frac{z}{\eta} \right)^\chi w(\Phi, A)^\chi \quad (44)$$

$$L(\Phi, A) = \int z l_H(a, z; \Phi, A) d\Phi = w(\Phi, A)^\chi \int \frac{z^{\chi+1}}{\eta^\chi} d\Phi_z \quad (45)$$

where $M = \int \frac{z^{\chi+1}}{\eta^\chi} d\Phi_z$ and Φ_z is the cumulative distribution function of the stationary productivity distribution. Then, from the optimality condition in the production sector with respect to the labor demand,

$$(1 - \alpha) A K^\alpha L(\Phi, A)^{-\alpha} = w(\Phi, A) = \left(\frac{L(\Phi, A)}{M} \right)^{\frac{1}{\chi}}. \quad (46)$$

Combining these two conditions, the conditional aggregate capital nullcline is as follows:

$$c_A^{Kcnnull} = A(M^{\frac{1}{1+\alpha\chi}}((1-\alpha)A)^{\frac{\chi}{1+\alpha\chi}})^{1-\alpha}(K_A^{Kcnnull})^{\alpha+(1-\alpha)\frac{\alpha\chi}{1+\alpha\chi}} - \delta K_A^{Kcnnull}, \quad (47)$$

which satisfies the same monotone convergence to the conditional steady state.

Distributional disturbance Now consider the following CRRA-utility setup:

$$u(c, l_H) = \frac{1}{1-\sigma}c^{1-\sigma} - \frac{\eta}{1+\frac{1}{\chi}}l_H^{1+\frac{1}{\chi}}. \quad (48)$$

The corresponding individual labor supply is as follows:

$$l_H(a, z; \Phi, A) = \left(\frac{z}{\eta c(a, z; \Phi, A)^\sigma} \right)^\chi w(\Phi, A)^\chi \quad (49)$$

Following the same step as in the GHH case above, I obtain the following conditional K nullcline condition:

$$c_A^{Kcnnull} = A(M(\Phi^{Kcnnull})^{\frac{1}{1+\alpha\chi}}((1-\alpha)A)^{\frac{\chi}{1+\alpha\chi}})^{1-\alpha}(K_A^{Kcnnull})^{\alpha+(1-\alpha)\frac{\alpha\chi}{1+\alpha\chi}} - \delta K_A^{Kcnnull}, \quad (50)$$

where $M(\Phi^{Kcnnull}) = \int \frac{z^{\chi+1}}{\eta^\chi c^{\sigma\chi}} d\Phi^{Kcnnull}$. Because of heterogeneous wealth effects, the conditional capital nullcline depends on the distribution of individual states $\Phi^{Kcnnull}$ associated with the level K^{cnnull} . As a result, strict monotonicity of the capital dynamics cannot be verified analytically. Nevertheless, [Lee \(2025\)](#) provides computational evidence that the recursive competitive equilibrium of the model with this utility specification exhibits strictly monotone convergence of aggregate capital K under the standard calibration.

5 Concluding remarks

This paper develops a geometric framework for stochastic equilibrium dynamics by introducing *conditional saddle paths*: invariant equilibrium manifolds defined under frozen exogenous states. This object extends the familiar saddle-path intuition from deterministic models to environments under aggregate uncertainty. In the resulting

phase-diagram representation, business-cycle fluctuations decompose into movements *along* a conditional saddle (endogenous propagation within a regime) and *across* conditional saddles (transitions in exogenous states). The framework clarifies why impulse responses can be state-dependent: such dependence is a geometric property of the equilibrium manifold and arises precisely when conditional saddles differ in slope rather than by mere vertical translation.

When an aggregate equilibrium variable varies strictly monotonically along a conditional saddle, it provides a global coordinate: it uniquely indexes equilibrium states and therefore summarizes all equilibrium allocations and prices on the relevant invariant set. Applying this logic, I provide a theoretical proof of the sufficiency of aggregate capital in a canonical heterogeneous-household model. Beyond the Krusell–Smith benchmark, the same reasoning applies in multi-asset and richer heterogeneous-agent environments whenever a monotone-convergent aggregate coordinate exists.

More broadly, conditional saddles offer a complementary lens on stochastic models: they provide a language for interpreting nonlinear dynamics and for assessing when scalar state approximations are exact rather than merely accurate. They also provide a natural geometry for *state-contingent* policy analysis: by making state dependence explicit in the phase diagram, the framework clarifies when the same intervention should be expected to have different quantitative effects across regimes and over the cycle. 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 in heterogeneous-agent economies under aggregate uncertainty.

References

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

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

Algan, Y., O. Allais, and W. J. Den Haan (2010). Solving the incomplete markets

model with aggregate uncertainty using parameterized cross-sectional distributions. *Journal of Economic Dynamics and Control* 34(1), 59–68. Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.

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

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(5), 2375–2408.

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

Baley, I. and A. Blanco (2019, January). Firm uncertainty cycles and the propagation of nominal shocks. *American Economic Journal: Macroeconomics* 11(1), 276–337.

Basu, S. and B. Bundick (2017). Uncertainty shocks in a model of effective demand. *Econometrica* 85(3), 937–958.

Berger, D., K. Milbradt, F. Tourre, and J. Vavra (2021, September). Mortgage prepayment and path-dependent effects of monetary policy. *American Economic Review* 111(9), 2829–78.

Berger, D. and J. Vavra (2015, January). Consumption Dynamics During Recessions: Consumption Dynamics During Recessions. *Econometrica* 83(1), 101–154.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018). Really Uncertain Business Cycles. *Econometrica* 86(3), 1031–1065.

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

Cao, D., W. Luo, and G. Nie (2023, January). Global DSGE Models. *Review of Economic Dynamics*, S1094202523000017.

Cass, D. (1965, 07). Optimum growth in an aggregative model of capital accumulation1. *The Review of Economic Studies* 32(3), 233–240.

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

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

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

Den Haan, W. J. and P. Rendahl (2010, January). Solving the incomplete markets model with aggregate uncertainty using explicit aggregation. *Journal of Economic Dynamics and Control* 34(1), 69–78.

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

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

Han, J., Y. Yang, and W. E (2025). DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks. Working Paper.

Kamihigashi, T. (2003). Necessity of transversality conditions for stochastic problems. *Journal of Economic Theory* 109(1), 140–149.

Kamihigashi, T. (2005). Necessity of the transversality condition for stochastic models with bounded or crra utility. *Journal of Economic Dynamics and Control* 29(8), 1313–1329.

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

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

Khan, A. and J. K. Thomas (2013). Credit shocks and aggregate fluctuations in an economy with production heterogeneity. *Journal of Political Economy* 121(6), 1055–1107.

Koop, G., M. Pesaran, and S. M. Potter (1996). Impulse response analysis in nonlinear multivariate models. *Journal of Econometrics* 74(1), 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. (1997, June). Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns. *Macroeconomic Dynamics* 1(02).

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

Lee, H. (2025). Global nonlinear solutions in sequence space and the generalized transition function. Working Paper.

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

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(1), 356–386.

Mendoza, E. G. (2010, December). Sudden stops, financial crises, and leverage. *American Economic Review* 100(5), 1941–66.

Payne, J., A. Rebei, and Y. Yang (2025). Deep Learning for Search and Matching Models. Working Paper.

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

Pizzinelli, C., K. Theodoridis, and F. Zanetti (2020). State dependence in labor market fluctuations. *International Economic Review* 61(3), 1027–1072.

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

Ramsey, F. P. (1928, 12). A mathematical theory of saving. *The Economic Journal* 38(152), 543–559.

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

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

Solow, R. M. (1956). A contribution to the theory of economic growth. *The Quarterly Journal of Economics* 70(1), 65–94.

Swan, T. W. (1956). Economic growth and capital accumulation. *Economic Record* 32(2), 334–361.

Tenreyro, S. and G. Thwaites (2016, October). Pushing on a String: US Monetary Policy Is Less Powerful in Recessions. *American Economic Journal: Macroeconomics* 8(4), 43–74.

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics* 129(1), 215–258.

Walsh, K. J. and E. Young (2024). Equilibrium multiplicity in aiyagari and krusell-smith. Working paper.

Winberry, T. (2021, January). Lumpy Investment, Business Cycles, and Stimulus Policy. *American Economic Review* 111(1), 364–396.

Yannacopoulos, A. N. (2011). Stochastic saddle paths and economic theory. In M. M. Peixoto, A. A. Pinto, and D. A. Rand (Eds.), *Dynamics, Games and Science II*, Berlin, Heidelberg, pp. 735–752. Springer Berlin Heidelberg.