Stochastic Evolution and Mutation

Game Theory · Module 7 — Evolutionary Game Theory

Node 7.5

Position in Knowledge Graph

System: Game Theory **Structural Domain:** Evolutionary Game Theory **Node:** 7.5

The replicator dynamics and ESS of 7.1-7.2 are deterministic: given an initial population state, the trajectory is fully determined. This is a useful idealization for infinite populations but misses two features that real evolutionary systems always have: **randomness** and **mutation**. Stochastic evolution adds both. The population state is now a random process – a Markov chain on the simplex (or on a finite set of states when the population is finite) – and mutation ensures that no strategy is ever fully eliminated, so the system can explore the full strategy space even from a "locked-in" state. The central concept is **stochastic stability** (Foster-Young 1990, Kandori-Mailath-Rob 1993, Young 1993): an equilibrium is stochastically stable if it receives positive probability in the stationary distribution as the mutation rate vanishes. Stochastic stability selects among multiple equilibria – it is the evolutionary equilibrium-selection mechanism that the deterministic replicator dynamics (which may have multiple attractors) cannot provide. The result is striking: in many games, stochastic stability selects the **risk-dominant** equilibrium, not the Pareto-dominant one. This has deep implications for cooperation theory (Module 5) and learning theory (Module 9).

Formal Definition

Finite-population Markov chain. Consider a population of N agents, each choosing from a finite strategy set $S = \{1, \ldots, K\}$. The population state is a vector $x = (x_1, \ldots, x_K)$ with $x_k \in \{0, 1, \ldots, N\}$ and $\sum_k x_k = N$, where x_k is the number of agents playing strategy k. The state space is the set of such vectors, denoted Ω.

Best-response dynamics with mutation. In each period: 1. One randomly selected agent revises their strategy. 2. With probability $1 - \epsilon$, the agent switches to a **best response** to the current population profile. 3. With probability $\epsilon > 0$, the agent **mutates**: chooses a strategy uniformly at random (or according to some fixed mutation distribution).

This defines an ergodic Markov chain on Ω for any $\epsilon > 0$ – ergodic because mutation ensures every state is reachable from every other state.

Stationary distribution. For $\epsilon > 0$, the chain has a unique stationary distribution π^ϵ on Ω.

Stochastic stability. A state $\omega \in \Omega$ is **stochastically stable** if

$$\lim_{\epsilon \to 0} \pi^\epsilon(\omega) > 0.$$

That is, ω retains positive probability in the limit of vanishing mutation. The set of stochastically stable states is the "long-run prediction" of the evolutionary process.

Resistance and the minimum-cost tree characterization. The **resistance** of a transition from state ω to ω' is the minimum number of simultaneous mutations required for the chain to move from ω to ω' in one step. Young (1993) showed that the stochastically stable states can be characterized by a **minimum-cost spanning tree** construction on the graph of absorbing states of the unperturbed ($\epsilon = 0$) dynamics:

- Identify the **absorbing states** of the unperturbed best-response chain (the states from which no best-response revision leads to a different state – these are the pure-strategy Nash equilibria when the population fills a single strategy).
- For each pair of absorbing states, compute the resistance: the minimum number of mutations needed to move from one basin of attraction to the other.
- Construct a directed graph with absorbing states as nodes and resistances as edge weights. The stochastically stable states are those that are the "roots" of minimum-cost spanning trees in this graph.

Theorem (Young 1993; Kandori-Mailath-Rob 1993). In 2x2 symmetric coordination games played by a finite population with best-response revision and uniform mutation: - Both pure Nash equilibria are absorbing states of the unperturbed dynamics. - The stochastically stable equilibrium is the **risk-dominant** Nash equilibrium – the one with the larger basin of attraction under best-response dynamics.

Intuition

The intuition for stochastic stability is vivid. Imagine a finite population sitting at a Nash equilibrium. With zero mutation, they stay there forever. With small mutation, occasionally an agent randomly switches to a different strategy. If the mutation count is below a threshold, best-response dynamics push the population back to the equilibrium. But if enough simultaneous mutations occur (a "cluster" of deviators), the population can tip into the basin of attraction of a different equilibrium.

Stochastic stability asks: which equilibrium is hardest to dislodge? The answer is the one that requires the most simultaneous mutations to escape. In a coordination game, the risk-dominant equilibrium has a larger basin of attraction – more agents need to simultaneously deviate to push the population out. So it is stochastically stable.

The risk-dominant equilibrium is defined by Harsanyi and Selten (1988) as the equilibrium with the larger "basin" in the best-response sense. For a 2x2 coordination game with pure equilibria (A, A) and (B, B):

$$(A, A) \text{ is risk-dominant iff } u(A, A) - u(B, A) > u(B, B) - u(A, B),$$

i.e., the "incentive to play A when your opponent plays A" exceeds the "incentive to play B when your opponent plays B." Risk dominance measures the robustness of an equilibrium to strategic uncertainty about the opponent's choice.

The Pareto-dominant equilibrium (the one with the highest total payoff) may or may not be risk-dominant. In the Stag Hunt game – $(4, 4)$ for mutual Stag, $(3, 3)$ for mutual Hare, with miscoordination giving $(0, 3)$ – the Pareto-dominant equilibrium (Stag, Stag) is risk-dominated because the risk of playing Stag against a Hare-player is high. Stochastic stability selects (Hare, Hare) – the safer, Pareto-inferior equilibrium.

This is a profound result with implications for cooperation theory. It says that **evolutionary selection favors risk over reward**. In the long run, populations converge to the equilibrium that is safest against deviations, not the one that is best for everyone. Cooperation is Pareto-dominant but often risk-dominated; stochastic stability explains why defection and safe play are so persistent in evolutionary and social systems.

Structural Role

7.5 provides the equilibrium-selection theory that the deterministic framework of 7.1-7.4 lacks. The replicator dynamics can have multiple stable fixed points; stochastic stability selects among them. The selection criterion – risk dominance, not Pareto dominance – is the evolutionary answer to the equilibrium-selection problem that the folk theorem (5.4) raised.

The connections are:

- **7.1-7.2 (Replicator dynamics, ESS)** give the deterministic attractors; 7.5 selects among them.

- **7.4 (Population games)** provides the large-population framework where the finite-population Markov chain converges to the replicator; 7.5 studies the finite-population corrections.
- **5.4 (Folk theorem)** raises the multiplicity problem; 7.5 is one evolutionary answer.
- **9.5 (Convergence to equilibrium)** studies the learning-dynamics analogue: do agents learn to play risk-dominant or Pareto-dominant equilibria?
- **7.6 (Adaptive dynamics)** studies a related but distinct process: strategy mutation in continuous strategy spaces.

Mathematical Development

Moran process and finite-population dynamics

The **Moran process** is a specific finite-population evolutionary model:

1. Population of N agents, each playing one of K strategies.
2. Each period: one agent is selected to reproduce (proportional to fitness) and one is selected to die (uniformly), maintaining constant N.
3. The offspring inherits the parent's strategy, with mutation probability ϵ: with probability ϵ, the offspring plays a uniformly random strategy; with probability $1-\epsilon$, they copy the parent.

The Moran process defines an ergodic Markov chain on the state space Ω for $\epsilon > 0$. For $\epsilon = 0$, the absorbing states are the **monomorphic states** where all N agents play the same strategy.

Fixation probability. Starting from a single mutant playing strategy B in a population of A-players, the probability that B eventually fixes (takes over the whole population) is:

$$\rho_{A\to B} = \frac{1}{1 + \sum_{j=1}^{N-1} \prod_{i=1}^{j} (f_A(i)/f_B(i))}$$

where $f_A(i)$ and $f_B(i)$ are the fitness of A and B players when there are i copies of B in the population. Under weak selection (fitness close to neutral), this simplifies to tractable expressions.

KMR model (Kandori-Mailath-Rob 1993)

The **KMR model** is the foundational framework for stochastic stability. A population of N agents plays a 2x2 coordination game. In each period, one agent revises. Without mutation: the agent plays best response to the current population frequencies. With mutation rate ϵ: the agent plays a uniformly random strategy.

Absorbing states of the unperturbed chain. The states where all N agents play A or all play B (both pure NE).

Resistance calculation. From the all-A state to the all-B state: best-response dynamics switch from A to B when the proportion of B-players exceeds the threshold q^*. So the minimum number of mutations to tip from A to B is $\lceil Nq^* \rceil$. From B to A: the minimum mutations are $\lceil N(1-q^*) \rceil$.

Risk-dominant selection. (A, A) is stochastically stable iff $q^* < 1/2$ – equivalently, the basin of attraction of A is larger than that of B – equivalently, (A, A) is risk-dominant. The result follows from the minimum-resistance tree: the cost of dislodging A exceeds the cost of dislodging B iff $q^* < 1/2$.

Young's minimum-cost tree characterization

Theorem (Young 1993). In a general finite-state Markov chain with multiple absorbing classes and small perturbations, the stochastically stable states are determined by the minimum-cost spanning trees rooted at each absorbing class. Specifically, state ω is stochastically stable iff the minimum total resistance of a tree rooted at ω (where all edges point toward ω) is minimized over all absorbing classes.

Beyond 2x2: stochastic stability in general games

For general $K \times K$ games, the stochastic stability analysis is more complex. The minimum-cost tree may select equilibria that are not obviously "risk-dominant" in the Harsanyi-Selten sense (which was defined only for 2x2 games). The general principle is: **stochastically stable equilibria are the ones with the largest basins of attraction under the unperturbed dynamics**, measured by the minimum number of mutations needed to escape.

For potential games (games with a Lyapunov function), the stochastically stable states coincide with the maximizers of the potential function – a clean characterization. For non-potential games, the analysis requires the full tree construction.

Worked Example

Stag Hunt with finite population

Stage game:

	Stag	Hare
Stag	$(4, 4)$	$(0, 3)$
Hare	$(3, 0)$	$(3, 3)$

Both (Stag, Stag) and (Hare, Hare) are Nash equilibria. (Stag, Stag) is Pareto-dominant; (Hare, Hare) is risk-dominant.

Risk dominance check. $u(\text{Stag}, \text{Stag}) - u(\text{Hare}, \text{Stag}) = 4 - 3 = 1$. $u(\text{Hare}, \text{Hare}) - u(\text{Stag}, \text{Hare}) = 3 - 0 = 3$. Since $1 < 3$, Hare is risk-dominant.

Best-response threshold. An agent plays Stag iff the proportion of Stag-players in the population exceeds q^*. Best response: Stag iff $4q > 3q + 3(1 - q) = 3$, so $4q > 3$, $q > 3/4$. The threshold is $q^* = 3/4$. Basin of Stag: $q > 3/4$, a small region. Basin of Hare: $q < 3/4$, a large region.

KMR analysis with $N = 12$. All-Stag absorbs. To tip to Hare: need $\lceil 12 \cdot 3/4 \rceil = 9$ mutations (from Stag to Hare). All-Hare absorbs. To tip to Stag: need $\lceil 12 \cdot 1/4 \rceil = 3$ mutations.

Resistance of dislodging Stag: 9. Resistance of dislodging Hare: 3. Since $3 < 9$, Hare is harder to dislodge. Wait – the resistance of dislodging state X is the number of mutations needed to leave it. Resistance of leaving all-Hare is 3 (need 3 Stag mutants to tip). Resistance of leaving all-Stag is 9 (need 9 Hare mutants to tip). So all-Stag is *easier* to leave (lower resistance to escape). The stochastically stable state is the one that is *harder* to leave: **(Hare, Hare)** with escape resistance 3... no, 9 is higher than 3. Let me be more careful.

Resistance from all-Stag to all-Hare = 9 (cost of escaping Stag). Resistance from all-Hare to all-Stag = 3 (cost of escaping Hare). The stochastically stable state is the one where the cost of escaping is higher. That is all-Stag? No – in the minimum-cost tree, the stochastically stable state ω minimizes the total resistance of the tree *rooted at* ω. With two states, the tree rooted at all-Hare has one edge: from all-Stag to all-Hare, resistance 9. The tree rooted at all-Stag has one edge: from all-Hare to all-Stag, resistance 3.

Minimum cost: tree rooted at all-Stag has cost 3. Tree rooted at all-Hare has cost 9. The minimum-cost root is all-Stag (cost $3 < 9$). So all-Stag is stochastically stable?

Wait, I need to recheck the standard result. KMR says risk-dominant is stochastically stable, and Hare is risk-dominant. Let me re-examine.

The convention: in Young's tree construction, a tree rooted at ω has directed edges pointing *toward* ω from all other states. The cost of the tree is the sum of edge resistances. The stochastically stable state minimizes this cost.

Tree rooted at Hare: edge from Stag to Hare. How many mutations are needed to move from all-Stag to the basin of Hare? We need enough Hare-mutants in the all-Stag population so that best-response dynamics carry the population to all-Hare. This requires $\lceil N(1-q^*)\rceil = \lceil 12 \cdot 1/4\rceil = 3$ mutations. So resistance = 3.

Tree rooted at Stag: edge from Hare to Stag. We need enough Stag-mutants in the all-Hare population so that best-response carries to all-Stag. This requires $\lceil N \cdot q^*\rceil = \lceil 12 \cdot 3/4\rceil = 9$ mutations. So resistance = 9.

Minimum cost: rooted at Hare, cost 3. Rooted at Stag, cost 9. So **(Hare, Hare)** is stochastically stable (minimum cost tree root). This confirms risk dominance selects Hare.

Interpretation. In the long run, with occasional random mutations, the population spends most of its time at Hare. The Pareto-superior Stag equilibrium is visited occasionally (when a lucky cluster of Stag-mutations occurs) but is quickly destabilized (a few Hare mutations tip the population back). This is the evolutionary tragedy: the safe outcome beats the good outcome.

Poker Anchor

Concept mapping

Formal object	Poker instantiation
Population state	The distribution of strategy types at a given stake: nits, LAGs, TAGs, maniacs
Mutation	A player randomly trying a new strategy (experimenting, tilting, studying a new approach)
Best-response revision	A player switching to the most profitable strategy given the current population
Stochastically stable state	The long-run composition of the player pool
Risk-dominant equilibrium	The "safe" strategy that is hardest to dislodge from the player pool
Pareto-dominant equilibrium	The "optimal" strategy that makes everyone better off if everyone used it

Concrete situation: why TAG is the stochastically stable style

At mid-stakes online poker, the player pool has converged over the past decade toward a tight-aggressive (TAG) meta. This is not the theoretically optimal strategy (GTO mixed strategies are), but it is the risk-dominant one: TAG play is safe against a wide range of opponents, and a population of TAG players is hard to dislodge (deviating to a more exploitative or more passive style is costly against TAG opponents). The GTO equilibrium might be Pareto-superior (if everyone played GTO, everyone's edge would be maximized against the recreational pool), but it is risk-dominated: a single player deviating from GTO toward exploitation can do very well against TAGs, but the TAGs respond, and the deviation is punished.

Occasional "mutations" happen: a player studies a new style, or tilts into a different strategy, or a new training tool shifts the meta. But most mutations are quickly punished by the TAG population (a LAG mutant loses to the nit-heavy TAG field; a nit mutant wins less because they cannot exploit). Only when a critical mass of mutants appears (e.g., a new solver-informed cohort enters the pool) does the population tip to a new equilibrium.

What this understanding buys you

You can predict meta stability. A meta is stochastically stable when the dominant strategy has a large basin of attraction – hard to dislodge by individual deviations. If you can estimate the basin size (how many players need to deviate to tip the population), you can predict how stable the current meta is and how likely a meta shift is.

Mutations are the engine of meta evolution. Every player who experiments, studies a new approach, or enters the pool with a fresh strategy is a mutant in the evolutionary model. The more experimentation there is, the faster the population explores the strategy space and the more likely a meta shift becomes.

Cross-Links

- **7.1 Replicator dynamics** – the deterministic limit of the stochastic process.
- **7.2 ESS** – stochastic stability refines ESS (an ESS is stochastically stable in many models).
- **7.4 Population games** – the large-population framework.
- **5.4 Folk theorem** – stochastic stability selects among repeated-game equilibria.
- **9.5 Convergence to equilibrium** – learning-dynamics analogue of stochastic stability.
- **Statistics – Markov chains** – the mathematical tools (ergodicity, stationary distributions, large deviations).

Structural Compression

Invariant: In finite-population evolutionary models with small mutation rates, the long-run behavior is concentrated on **stochastically stable** states – the absorbing states of the unperturbed dynamics that require the most mutations to dislodge; in 2x2 coordination games, stochastic stability selects the risk-dominant equilibrium, not the Pareto-dominant one.

Mechanism: Ergodic Markov chain with absorbing states under zero mutation and full exploration under positive mutation; minimum-cost spanning trees on the graph of absorbing states characterize the stationary distribution in the vanishing-mutation limit; the cost of dislodging an equilibrium is the minimum number of simultaneous mutations needed to tip the population out of its basin of attraction.

Cross-System Echo: Stochastic stability is the game-theoretic analogue of **metastability** in statistical physics. A physical system at low temperature settles into local energy minima; occasional thermal fluctuations (analogous to mutations) cause transitions between minima; the globally stable state is the one with the deepest energy well (largest basin), which is the minimum of the free energy. The KMR model is a birth-death chain with structure identical to the Kramers escape problem in physics: the escape rate from a metastable state is exponential in the barrier height (number of mutations). The risk-dominant equilibrium is the "ground state" of the evolutionary system – the analogy is exact for potential games.

Exercises

1. For the Stag Hunt game with payoffs $(4,4),(0,3),(3,0),(3,3)$, compute the best-response threshold q^* and verify that (Hare, Hare) is risk-dominant and stochastically stable.
2. In a coordination game with payoffs (a,a) for (A, A) and (b,b) for (B, B) and $(0,0)$ for miscoordination, derive the condition on a and b for (A, A) to be risk-dominant. Show that risk dominance and Pareto dominance coincide iff the game is trivial (one equilibrium dominates the other in both senses).
3. For the KMR model with $N = 20$ and the Prisoner's Dilemma (unique NE at (D, D)), what is the stochastically stable state? Verify that stochastic stability does not help cooperation here.
4. Compute the fixation probability of a single Hawk mutant in a population of $N = 10$ Dove players in the Hawk-Dove game with payoffs $V = 2, C = 4$ under the Moran process with neutral drift.
5. In a 3x3 coordination game with three pure NE, use Young's minimum-cost tree construction to determine the stochastically stable state. (Choose payoffs that make the construction non-trivial.)

6. Show that in a potential game, the stochastically stable states under log-linear dynamics coincide with the maximizers of the potential function.
7. **Argue, structurally,** why stochastic stability selects for robustness (risk dominance) rather than efficiency (Pareto dominance). What does this tell us about the long-run prospects for cooperation in evolutionary systems, and how does this connect to Module 5's theory of repeated-game cooperation?

Game Theory · Module 7 — Evolutionary Game Theory · Node 7.5
Concepts: markov-chains, nash-equilibrium, stability, stochastic-optimization
Prerequisites: 7.1, 7.2, 7.4
Enables: 7.6, 9.5

hbar.university — own.your.cognition