

Multiplicity and Uniqueness of Equilibria in Continuous-Time Dynamic Discrete Games

Jason R. Blevins and Youngjae Jeong
The Ohio State University

October 24, 2026

Abstract

We consider multiplicity and uniqueness of equilibria in continuous-time dynamic discrete games where players move sequentially according to independent Poisson processes. With multiple equilibria, the structural model is incomplete: the reduced form and counterfactual predictions depend on how equilibria are selected. We provide closed-form sufficient conditions for uniqueness of Markov perfect equilibrium in this class of games. The conditions are semiparametric and can be verified from model primitives. To measure the coverage of the conditions and characterize multiplicity more broadly, we carry out a large-scale numerical search for equilibria of a dynamic duopoly model of entry and exit. Although uniqueness is the norm in our broad search, we find that multiplicity concentrates in economically relevant regions: where the market cannot profitably sustain both firms, where firms are very patient, and at intermediate decision frequencies that balance strategic interaction and commitment. When multiplicity is present, we typically find two stable asymmetric equilibria and one unstable symmetric equilibrium. Monte Carlo and counterfactual simulations show this structure matters for empirical work. A full-solution estimator that assumes uniqueness can converge to the unstable equilibrium, even when the data were generated by a stable one, while a two-step estimator recovers the correct equilibrium.

Keywords: Stochastic games, continuous-time games, dynamic discrete games, Markov perfect equilibrium, equilibrium uniqueness, equilibrium multiplicity, equilibrium selection.

JEL Classification: C57, C62, C63, C73, L13.

1 Introduction

Multiple equilibria complicate both estimation of dynamic discrete games and interpretation of counterfactuals. One motivation for specifying such games in continuous time—where

decision times arrive according to independent Poisson processes, as in the framework of Arcidiacono, Bayer, Blevins, and Ellickson (2016) (henceforth ABBE)—is that players move asynchronously with probability one. This serves the dual purposes of removing simultaneous moves as a source of multiplicity and, because state transitions occur sequentially, reducing the computational burden of calculating players’ expectations over successor states (Doraszelski and Judd, 2012). However, removing simultaneity does not fully eliminate multiplicity, and so continuous-time models, like their discrete-time counterparts, can exhibit multiple Markov perfect equilibria (MPE). This possibility led ABBE and Blevins and Kim (2024) to propose two-step estimation methods that condition on the equilibrium played in the data. Direct theoretical analysis of multiplicity in dynamic games is complicated by the fact that players’ continuation values depend on rivals’ continuation values in a simultaneous system of nonlinear equations. In this paper, we derive what are, to our knowledge, the first closed-form sufficient conditions for uniqueness of MPE in continuous-time dynamic discrete games. Without strong restrictions on the structure of the game, the analytical conditions we can obtain are conservative, so we turn to a large-scale numerical search to determine where multiplicity occurs in the parameter space, understand the economic sources of multiplicity, and quantify its consequences for estimation and counterfactual analysis.

The model is incomplete due to multiple equilibria, which complicates structural estimation (Tamer, 2003). In full-solution estimation, the researcher needs to either compute all equilibria for each candidate parameter vector or impose an equilibrium selection rule (see de Paula, 2013, Aguirregabiria, Collard-Wexler, and Ryan, 2021, and Otsu and Pesendorfer, 2023 for surveys). The literature typically addresses multiplicity by assuming that data are generated from a single equilibrium and then either using two-step or iterative estimation methods (Aguirregabiria and Mira, 2007, Bajari, Benkard, and Levin, 2007, Pesendorfer and Schmidt-Dengler, 2008) that build on conditional choice probability methods for single agent models (Hotz and Miller, 1993), or imposing the equilibrium conditions as constraints during estimation (Su and Judd, 2012, Egedal, Lai, and Su, 2015). The convergence of iterative methods is known to depend on the stability of the mapping used by the estimator (Pesendorfer and Schmidt-Dengler, 2010, Kasahara and Shimotsu, 2012, Su, 2014).

Multiplicity also confounds counterfactual analysis: even if a unique equilibrium is played in the data, it is unclear which equilibrium will be played under counterfactual parameter values. The predicted policy effects, therefore, depend on an equilibrium selection assumption (Aguirregabiria, 2012, Jun and Pinkse, 2020). Sufficient conditions for uniqueness can mitigate these concerns.

Our conditions are semiparametric: they do not require a parametric distribution for the choice-specific errors. Our first result (Theorem 1) applies to the ABBE class of dynamic games with any finite number of players, any finite state space, and binary actions. It

provides a bound on the discount rate ρ in terms of structural primitives that ensures the Bellman system operator is a contraction on a set of value functions that we show contains every equilibrium.

Our second result (Theorem 2) establishes uniqueness for a dynamic duopoly entry/exit model with symmetric firms and offsetting switching costs (equal in magnitude, with opposite sign). The equilibrium system reduces to two equations, allowing us to state a less restrictive contraction condition. When the competitive effect θ_C is moderate relative to the move arrival rate λ (with logit errors, $|\theta_C| < 3.68\lambda$), the equilibrium is unique even for arbitrarily patient firms (Corollary 1). The reduction and discount-free conditions exploit features of the continuous-time model with asynchronous moves: the equilibrium system is linear in rivals' choice probabilities, and firms lose their ability to commit as moves become more frequent. Both conditions are stated in terms of primitives, so applied researchers can verify them to check for uniqueness in estimation or counterfactual analysis. In Section 5.2, we apply Theorem 2 to a counterfactual entry subsidy that eliminates sunk costs.

We investigate the range of parameters for which our conditions hold and further characterize uniqueness and multiplicity by numerically searching for equilibria in the duopoly model over 100 million parameter vectors. We find that uniqueness is common even beyond the regions covered by our theorems, but multiplicity appears in economically relevant cases such as when firms are patient, competition is strong, or entry costs are high relative to scrap values. The multiplicity rate reaches 16% among cases where duopoly is unprofitable, firms discount at standard rates, and moves are not too infrequent. It is more likely when there is a balance between strategic interaction (sufficiently frequent moves) and commitment (sufficiently infrequent moves).

Where multiplicity occurs, we investigate its consequences in terms of economic outcomes, estimation, and counterfactuals. We typically find two stable asymmetric mirror-image equilibria and an unstable symmetric equilibrium. The long-run duopoly share varies by 0.128 across these three equilibria in the median case. In a Monte Carlo experiment, a full-solution maximum likelihood estimator that assumes uniqueness can converge to an unstable equilibrium even when the data were generated by a stable one. In contrast, two-step estimators initialized with first-stage estimates of choice probabilities select the equilibrium that generated the data, and their iterated versions do so when the iteration is contractive. Finally, equilibrium selection can have a large effect on counterfactual predictions involving firm identity, as we demonstrate through an entry subsidy.

Although we focus on dynamic games, equilibrium multiplicity has been studied more thoroughly in static games. Beginning with the foundational work of [Bresnahan and Reiss \(1991\)](#) and [Berry \(1992\)](#), the literature on empirical entry games has focused primarily on equilibrium multiplicity in static models (e.g., [Tamer, 2003](#), [Jia, 2008](#), [Ciliberto and Tamer,](#)

2009, and Sweeting, 2009). More recently, Espín-Sánchez, Parra, and Wang (2023) study static entry games with private information and, based on a measure of competitive strength that orders firms, provide sufficient conditions for a unique equilibrium.

In dynamic games, a number of computational papers have documented the existence of multiple equilibria in discrete-time models. Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) propose a model of learning-by-doing with organizational forgetting and find that multiplicity is concentrated where the inflow of know-how from learning balances the outflow from forgetting. Iskhakov, Rust, and Schjerning (2016) use a recursive lexicographic search algorithm to find all MPE of directional dynamic games—games in which some component of the state evolves irreversibly—and show that discrete-time Bertrand models can support hundreds of millions of equilibria.

Theoretical work has characterized equilibria in related models, including discrete-time finite discounted stochastic games. Haller and Lagunoff (2000) prove that such games generically have a finite number of MPE, and Herings and Peeters (2004) refine this to an odd number. Doraszelski and Escobar (2010) show that such equilibria are generically regular. Doraszelski and Satterthwaite (2010) introduce privately observed, randomly drawn entry costs and scrap values in the Ericson and Pakes (1995) framework to ensure existence of an MPE in cutoff entry/exit strategies.

In certain classes of discrete-time dynamic models, sufficient conditions for uniqueness are known. In a dynamic duopoly model of price competition with learning-by-doing, Cabral and Riordan (1994) prove that the MPE is unique, but show that adding avoidable fixed costs and irreversible exit can result in multiplicity. Besanko et al. (2010) extend this result to a model with organizational forgetting, proving uniqueness in the two limiting cases in which forgetting is either absent or certain. In an oligopoly entry/exit model in which firms make survival decisions sequentially (in order of seniority), Abbring and Campbell (2010) establish existence of an essentially unique symmetric MPE within the class of “last-in first-out” strategies. Abbring, Campbell, Tilly, and Yang (2018) instead assume simultaneous survival decisions and obtain an essentially unique symmetric MPE without last-in first-out restrictions, under the assumption of homogeneous firms. Since these results rely on special structures or equilibrium refinements, they are not directly applicable to the general continuous-time models considered here.

The remainder of this paper is organized as follows. Section 2 introduces the class of continuous-time games we consider and our leading example: a duopoly entry and exit model. Section 3 establishes our main theoretical results, providing sufficient conditions for equilibrium uniqueness. Section 4 develops our computational approach and presents numerical results for the structural parameters and economic quantities that determine

regions of multiplicity and uniqueness.¹ Section 5 discusses the implications of equilibrium multiplicity for structural estimation and counterfactual analysis. Section 6 concludes.

2 Continuous-Time Dynamic Discrete Games

In this section we briefly review the class of continuous-time dynamic discrete games introduced by ABBE and present the two-player entry/exit model that will serve as our leading example throughout the paper.

2.1 General Framework

There are N players in the model, indexed by $i = 1, \dots, N$. Each player makes discrete choices over an infinite time horizon at stochastically determined decision epochs that arrive according to independent Poisson processes with rates λ_i . Therefore, the move times of players are almost surely sequential rather than simultaneous. There is a commonly observed state vector x , which is a member of the finite state space $\mathcal{X} \equiv \{x_1, \dots, x_K\}$. State transitions occur endogenously as a result of the actions of players.

At a decision epoch for player i in state x_k , the player observes a vector of private choice-specific errors $\varepsilon_{ik}(t)$ and chooses a discrete action j from the choice set $\mathcal{J}$. In this paper, we limit consideration to binary-choice models with $\mathcal{J} = \{0, 1\}$. Action $j = 0$ represents continuation, with no corresponding state change, while $j = 1$ causes the state to jump immediately and deterministically from x_k to $x_{\ell(i,k)}$, where $\ell(i, \cdot): \{1, \dots, K\} \rightarrow \{1, \dots, K\}$ denotes player i 's switching state map. We encode this map as a deterministic switching matrix S_i with $[S_i]_{k, \ell(i,k)} = 1$ for each row k and zeros elsewhere.

While the model remains in some state x_k between state transitions, players continuously accrue flow payoffs u_{ik} . At a decision epoch t where player i chooses action j in state x_k , that player receives a lump-sum payoff

$$\psi_{ijk} + \varepsilon_{ijk}(t),$$

where ψ_{ijk} is a deterministic, instantaneous choice-specific payoff and $\varepsilon_{ijk}(t)$ is the corresponding choice-specific component of the error vector. We normalize the payoff for the continuation action $j = 0$ by setting $\psi_{i0k} = 0$ for all i and k (Assumption 5 below). Therefore, with binary actions we simply drop the j subscript and write $\psi_{ik} \equiv \psi_{i1k}$ for the switching cost and $\psi_i = (\psi_{i1}, \dots, \psi_{iK})^\top$ for the vector of switching costs across states.

The errors are i.i.d. across epochs (Assumption 6 below) and do not directly enter other players' continuation values, so it is irrelevant whether rivals observe past error vectors. We

¹Complete replication code is available at <https://github.com/jrblevin/ctunique-replication>.

assume each player observes only its own errors (i.e., the errors are private information). Since the model is stationary, we omit the time argument and write ε_{ijk} and ε_{ik} to denote a choice-specific error component and vector, respectively, at a generic decision epoch.

We now formally state the main assumptions of the model, closely following ABBE.

Assumption 1 (Discrete States). The state space is finite: $K \equiv |\mathcal{X}| < \infty$.

Assumption 2 (Discount Rates). Players use a common discount rate $\rho \in (0, \infty)$.

Assumption 3 (Move Arrival Rates). Move arrival times follow independent Poisson processes with rate parameters λ_i satisfying $0 < \lambda_i < \infty$ for $i = 1, \dots, N$.

Assumption 4 (Bounded Payoffs). Flow payoffs and instantaneous payoffs are finite: $|u_{ik}| < \infty$ and $|\psi_{ijk}| < \infty$ for $i = 1, \dots, N$, $j = 0, 1$, and $x_k \in \mathcal{X}$.

Assumption 5 (Costless Continuation). For $i = 1, \dots, N$ and for all $x_k \in \mathcal{X}$, $\psi_{i0k} = 0$.

Assumption 6 (Distribution of Private Information). The error vectors $\varepsilon_{ik}(t) \equiv (\varepsilon_{i0k}(t), \varepsilon_{i1k}(t))$ are i.i.d. across players i , states x_k , and decision epochs t . The distribution is absolutely continuous with respect to Lebesgue measure, with finite first moments and support equal to $\mathbb{R}^{|\mathcal{J}|}$.

Under Assumption 6, the error difference $\varepsilon_{i0k} - \varepsilon_{i1k}$ is absolutely continuously distributed with support $\mathbb{R}$, and we let F and f denote its CDF and density.

Assumption 7 (Bounded Error-Difference Density). The density f is bounded: $c_F \equiv \sup_z f(z) < \infty$.

Assumption 8 (Exchangeable Errors with Lipschitz Density). The pair $(\varepsilon_{i0k}, \varepsilon_{i1k})$ is exchangeable and the density f is Lipschitz continuous with constant c'_F .

Assumption 9 (Type I Extreme Value Errors). For each i and k , the choice-specific errors ε_{i0k} and ε_{i1k} are independent and each follows the type I extreme value distribution with mean 0 and scale parameter 1.

Remark. Assumption 9 implies Assumptions 7 and 8 with $c_F = 1/4$ and $c'_F = \sqrt{3}/18$ (Lemma 5 in Appendix A). Although this assumption is standard in applied work and yields closed-form choice probabilities and surplus functions, our uniqueness results in Section 3 do not require it. We state it for use in examples and in the numerical work of Section 4.

This framework is a special case of the ABBE model, with no exogenous state changes and binary actions. The first restriction is for simplicity only. Our main uniqueness condition holds if exogenous state changes are allowed (see the remarks following Theorem 1). Generalizing to more than two choices is a more substantive extension that we leave for future work.

2.2 Value Functions and Equilibrium

A *stationary Markov strategy* for player i is a function $\delta_i: \mathcal{X} \times \mathbb{R}^{|\mathcal{J}|} \rightarrow \mathcal{J}$ that assigns an action to each combination of the state x_k and realized error vector ε_{ik} . A strategy δ_i induces conditional choice probabilities (CCPs)

$$(1) \quad \sigma_{ijk} \equiv \Pr[\delta_i(x_k, \varepsilon_{ik}) = j \mid x_k],$$

where the probability is taken over the distribution of ε_{ik} in Assumption 6. Since actions are binary, we track only the switching probabilities $\sigma_{ik} \equiv \sigma_{i1k}$, with $1 - \sigma_{ik} = \sigma_{i0k}$ being the probability of continuation. We denote profiles of CCPs and strategies by $\sigma = (\sigma_1, \dots, \sigma_N)$ and $\delta = (\delta_1, \dots, \delta_N)$, respectively, and write σ_{-i} for the collection of CCPs of player i 's rivals.

Let V_{ik} denote the value function for player i in state x_k given the rivals' switching probabilities σ_{-i} (we suppress the dependence of V_{ik} on σ_{-i} , which is implicit):

$$(2) \quad V_{ik} = \frac{1}{\rho + \eta} \left[u_{ik} + \sum_{m \neq i} \lambda_m \{ \sigma_{mk} V_{i, \ell(m, k)} + (1 - \sigma_{mk}) V_{ik} \} \right. \\ \left. + \lambda_i \mathbb{E}_\varepsilon \max\{ \psi_{ik} + \varepsilon_{i1k} + V_{i, \ell(i, k)}, \varepsilon_{i0k} + V_{ik} \} \right],$$

where $\eta \equiv \sum_{m=1}^N \lambda_m$ is the total move arrival rate, so that values are discounted at the effective rate $\rho + \eta$. The first term inside the square brackets, u_{ik} , is the flow payoff in state x_k . The middle terms reflect the possibility that rival m receives a decision opportunity (at rate λ_m) and either switches (with probability σ_{mk}), moving the state to $x_{\ell(m, k)}$, or continues (with probability $1 - \sigma_{mk}$), leaving the state unchanged. These middle terms are linear in rival players' choice probabilities since at most one player moves per decision epoch. In a simultaneous-move discrete-time model, however, the continuation value of player i is an expectation over the joint actions of all rivals, so rival choice probabilities enter inside the expected maximum over player i 's own actions. The final term is player i 's own expected option value when it receives a decision opportunity, where ψ_{ik} is the switching cost and the expectation $\mathbb{E}_\varepsilon$ is taken over the choice-specific errors.

Let $E_{ik}(V_i)$ denote McFadden's social surplus function (McFadden, 1978, 1981), which gives the expected maximum utility for player i in state x_k :

$$\begin{aligned} E_{ik}(V_i) &\equiv \mathbb{E}_\varepsilon \max\{ \psi_{ik} + \varepsilon_{i1k} + V_{i, \ell(i, k)}, \varepsilon_{i0k} + V_{ik} \} \\ &= L(\psi_{ik} + V_{i, \ell(i, k)}, V_{ik}) \\ &= V_{ik} + e(\psi_{ik} + V_{i, \ell(i, k)} - V_{ik}), \end{aligned}$$

where

$$L(z_1, z_2) \equiv \mathbb{E}_\varepsilon \max\{ z_1 + \varepsilon_{i1k}, z_2 + \varepsilon_{i0k} \} \quad \text{and} \quad e(z) \equiv L(z, 0).$$

Here $L(z_1, z_2)$ is the surplus from a binary choice with deterministic payoff components z_1 and z_2 , and $e(z)$ is shorthand we use throughout the analysis that follows. The final equality for E_{ik} holds because adding a constant to both payoffs shifts the surplus by that amount. Both functions L and e are finite under Assumption 6 and do not depend on (i, k) since the errors are identically distributed across players and states. Collecting these values into a vector $E_i(V_i)$, with $V_i = (V_{i1}, \dots, V_{iK})^\top$, we can write $E_i(V_i) = V_i + e(\psi_i + S_i V_i - V_i)$, where e applies componentwise.

Let $b_i: \mathcal{X} \times \mathbb{R}^{|\mathcal{J}|} \times [0, 1]^{(N-1)K} \rightarrow \mathcal{J}$ denote the *best-response function* for player i in state x_k upon observing the error vector ε_{ik} given the rivals' CCPs σ_{-i} . Player i chooses to switch ($j = 1$) if and only if the choice-specific value of switching is at least that of continuing:

$$(3) \quad b_i(x_k, \varepsilon_{ik}, \sigma_{-i}) = 1\{\psi_{ik} + \varepsilon_{i1k} + V_{i,\ell(i,k)} \geq \varepsilon_{i0k} + V_{ik}\},$$

where V_{ik} and $V_{i,\ell(i,k)}$ are the values induced by σ_{-i} through (2).

Definition. A *Markov perfect equilibrium* (MPE) is a profile of stationary Markov strategies $\delta^* = (\delta_1^*, \dots, \delta_N^*)$ such that, for each player i and each (x_k, ε_{ik}) ,

$$\delta_i^*(x_k, \varepsilon_{ik}) = b_i(x_k, \varepsilon_{ik}, \sigma_{-i}^*),$$

where $\sigma^* = (\sigma_1^*, \dots, \sigma_N^*)$ is the CCP profile derived from δ^* via (1).

An MPE requires that each player's strategy is a best response to the rivals' CCPs and that the CCPs are themselves consistent with the strategies. Because player i 's payoffs depend on the rivals only through the distribution of the rivals' actions in each state, the strategies δ^* enter the best-response problem only through their induced CCPs σ^* . The equilibrium can therefore be equivalently characterized in distributional-strategy form (Milgrom and Weber, 1985). Following Aguirregabiria and Mira (2007) and ABBE, we adopt this characterization and treat MPE as a fixed point of the equilibrium condition in CCPs σ . Applying (3), the equilibrium switching probability is

$$(4) \quad \sigma_{ik} = \Pr[\psi_{ik} + \varepsilon_{i1k} + V_{i,\ell(i,k)} \geq \varepsilon_{i0k} + V_{ik}] = F(\psi_{ik} + V_{i,\ell(i,k)} - V_{ik}),$$

where F is the CDF of the error difference $\varepsilon_{i0k} - \varepsilon_{i1k}$.

Example 1 (Logit). Under Assumption 9, the error difference $\varepsilon_{i0k} - \varepsilon_{i1k}$ follows the standard logistic distribution, so $F = \Lambda$, where $\Lambda(z) \equiv \exp(z)/(1 + \exp(z)) = 1/(1 + \exp(-z))$ is the logistic CDF, and the switching probabilities (4) take the familiar logit form. The surplus function has the log-sum-exp closed form $E_i(V_i) = \log[\exp(\psi_i + S_i V_i) + \exp(V_i)]$, with the exponential and logarithm applied componentwise.

The values V_{ik} and $V_{i,\ell(i,k)}$ in (4) are induced by the rivals' CCPs σ_{-i} . We stack the switching probabilities (4) for each player i across states x_k . A fixed point $\sigma^* = (\sigma_1^*, \dots, \sigma_N^*)$ of this system characterizes an MPE. Under Assumptions 1–6, at least one such fixed point exists, by Theorem 1 of Blevins (2026).

The equilibrium conditions can be expressed in terms of either CCPs or value functions. Given an equilibrium CCP profile σ , each player's value function is uniquely determined as the solution of a linear system with a nonsingular coefficient matrix, by Theorem 2 of Blevins (2026), which generalizes Proposition 6 of ABBE. Conversely, given the value functions, the CCPs are given directly by (4). The value functions are therefore sufficient to determine the CCPs, and the two fixed-point problems have equivalent solutions: σ^* is an equilibrium CCP profile if and only if the associated value functions solve the Bellman system obtained by substituting (4) into (2). In the remainder of the paper we therefore characterize equilibria as solutions of this value-function system.

Substituting the explicit forms of the choice probabilities and social surplus functions, we can express the equilibrium system of Bellman equations entirely in terms of value functions. Player m 's CCPs induce the $K \times K$ transition probability matrix

$$(5) \quad \Sigma_m(V_m) \equiv I + D_m(V_m)(S_m - I),$$

where $D_m(V_m) \equiv \text{diag}(\sigma_{m1}, \dots, \sigma_{mK})$ holds the switching probabilities (4). Row k of $\Sigma_m(V_m)$ places probability σ_{mk} on the switching destination $\ell(m, k)$ and probability $1 - \sigma_{mk}$ on k , so each $\Sigma_m(V_m)$ is a stochastic matrix. The Bellman equation (2) for player i can then be stated simply in vectorized form as

$$(6) \quad V_i = \frac{1}{\rho + \eta} \left[u_i + \sum_{m \neq i} \lambda_m \Sigma_m(V_m) V_i + \lambda_i E_i(V_i) \right].$$

Defining $V \equiv (V_1^\top, \dots, V_N^\top)^\top \in \mathbb{R}^{NK}$, the equilibrium system can be written compactly as

$$(7) \quad V = T(V),$$

where $T: \mathbb{R}^{NK} \rightarrow \mathbb{R}^{NK}$ is the *Bellman system operator* defined by stacking (6) over players $i = 1, \dots, N$.

2.3 A Dynamic Duopoly Model of Entry and Exit

Our leading example, to which our second uniqueness condition and our numerical work apply, is a two-player instance of this framework in an entry/exit setting. When active, the firms operate in the same product market, providing the same good or service.

The state of the market is represented by a state vector $x \in \mathcal{X} \equiv \{(0, 0), (0, 1), (1, 0), (1, 1)\}$, where the first component represents the entry status of firm 1—with 0 denoting inactive

Initial State		Decision		New State	
k	x_k	Firm i	Action j	$\ell(i, k)$	$x_{\ell(i, k)}$
1	(0, 0)	1	1	3	(1, 0)
2	(0, 1)	1	1	4	(1, 1)
3	(1, 0)	1	1	1	(0, 0)
4	(1, 1)	1	1	2	(0, 1)
1	(0, 0)	2	1	2	(0, 1)
2	(0, 1)	2	1	1	(0, 0)
3	(1, 0)	2	1	4	(1, 1)
4	(1, 1)	2	1	3	(1, 0)

Table 1: State transitions for switching actions ($j = 1$) for each player $i = 1, 2$. Continuation actions ($j = 0$) result in no state change.

and 1 denoting active—and the second component represents the entry status of firm 2. It will be convenient to index the states by integers as x_k for $k = 1, 2, 3, 4$ where $x_1 = (0, 0)$, $x_2 = (0, 1)$, $x_3 = (1, 0)$, and $x_4 = (1, 1)$.

For each firm, the switching action ($j = 1$) toggles its own activity status: an inactive firm enters the market and an active firm exits the market. The continuation action ($j = 0$) leaves the firm's status—and hence the state—unchanged. The resulting transitions are shown in Table 1.

In terms of the switching matrices defined in Section 2.1, the transitions in Table 1 correspond to

$$S_1 = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad S_2 = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

The flow payoffs for the two firms can be represented by vectors as follows:

$$(8) \quad u_1 = \begin{pmatrix} 0 \\ 0 \\ \theta_M \\ \theta_M + \theta_C \end{pmatrix} \quad \text{and} \quad u_2 = \begin{pmatrix} 0 \\ \theta_M \\ 0 \\ \theta_M + \theta_C \end{pmatrix} \quad \text{for} \quad \begin{matrix} x_1 = (0, 0) \\ x_2 = (0, 1) \\ x_3 = (1, 0) \\ x_4 = (1, 1) \end{matrix}.$$

We assume $\theta_M > 0$ and $\theta_C < 0$ in the duopoly model: θ_M is the flow profit from operating in the market as a monopolist, and θ_C is the change in flow profit from operating in the market as a duopolist after the entry of the rival firm.

The switching-cost vectors are

$$(9) \quad \psi_1 = \begin{pmatrix} \theta_{EC} \\ \theta_{EC} \\ \theta_{SV} \\ \theta_{SV} \end{pmatrix} \text{ and } \psi_2 = \begin{pmatrix} \theta_{EC} \\ \theta_{SV} \\ \theta_{EC} \\ \theta_{SV} \end{pmatrix} \text{ for } \begin{matrix} x_1 = (0, 0) \\ x_2 = (0, 1) \\ x_3 = (1, 0) \\ x_4 = (1, 1) \end{matrix}.$$

Here θ_{EC} denotes the cost of entry (switching from 0 to 1) and θ_{SV} denotes the scrap value received when the firm exits (switching from 1 to 0).

For the duopoly entry/exit model, $N = 2$ and $K = 4$, so (7) is the full eight-equation equilibrium system. For firm 1, equation (6) reads

$$(10) \quad V_1 = \frac{1}{\rho + \eta} [u_1 + \lambda_2 \Sigma_2(V_2)V_1 + \lambda_1 E_1(V_1)],$$

where the 4×4 transition probability matrix determined by firm 2's switching probabilities is

$$(11) \quad \Sigma_2(V_2) = \begin{bmatrix} 1-F(\theta_{EC}+V_{22}-V_{21}) & F(\theta_{EC}+V_{22}-V_{21}) & 0 & 0 \\ F(\theta_{SV}+V_{21}-V_{22}) & 1-F(\theta_{SV}+V_{21}-V_{22}) & 0 & 0 \\ 0 & 0 & 1-F(\theta_{EC}+V_{24}-V_{23}) & F(\theta_{EC}+V_{24}-V_{23}) \\ 0 & 0 & F(\theta_{SV}+V_{23}-V_{24}) & 1-F(\theta_{SV}+V_{23}-V_{24}) \end{bmatrix}$$

and firm 1's social surplus vector $E_1(V_1) = V_1 + e(\psi_1 + S_1 V_1 - V_1)$ is, explicitly,

$$E_1(V_1) = \begin{bmatrix} V_{11} + e(\theta_{EC} + V_{13} - V_{11}) \\ V_{12} + e(\theta_{EC} + V_{14} - V_{12}) \\ V_{13} + e(\theta_{SV} + V_{11} - V_{13}) \\ V_{14} + e(\theta_{SV} + V_{12} - V_{14}) \end{bmatrix}.$$

Similarly for firm 2,

$$(12) \quad V_2 = \frac{1}{\rho + \eta} [u_2 + \lambda_1 \Sigma_1(V_1)V_2 + \lambda_2 E_2(V_2)],$$

where

$$(13) \quad \Sigma_1(V_1) = \begin{bmatrix} 1-F(\theta_{EC}+V_{13}-V_{11}) & 0 & F(\theta_{EC}+V_{13}-V_{11}) & 0 \\ 0 & 1-F(\theta_{EC}+V_{14}-V_{12}) & 0 & F(\theta_{EC}+V_{14}-V_{12}) \\ F(\theta_{SV}+V_{11}-V_{13}) & 0 & 1-F(\theta_{SV}+V_{11}-V_{13}) & 0 \\ 0 & F(\theta_{SV}+V_{12}-V_{14}) & 0 & 1-F(\theta_{SV}+V_{12}-V_{14}) \end{bmatrix}$$

and $E_2(V_2) = V_2 + e(\psi_2 + S_2 V_2 - V_2)$.

3 Conditions for Equilibrium Uniqueness

Having characterized the equilibrium system, and with existence of an equilibrium guaranteed, we now establish conditions under which the equilibrium is unique. Equilibria are solutions

of the fixed-point problem $V = T(V)$, where T is the Bellman system operator defined in (7). Our approach proceeds in two steps. First, we derive an a priori bound on equilibrium values: the span—the difference between the largest and smallest components—of *any* equilibrium value function is bounded by a constant M determined by the model primitives, so every equilibrium lies in the sublevel set $\mathcal{V}_M$ of value functions whose spans are at most M . This set is *invariant* under T in the sense that T maps it into itself. Second, we provide sufficient conditions on the model parameters under which T is a contraction on this set—that is, there exists $\gamma < 1$ such that $\|T(V) - T(V')\|_\infty \leq \gamma \|V - V'\|_\infty$ for all V, V' in $\mathcal{V}_M$. By the Banach fixed-point theorem, T then has a unique fixed point in the set, and since every equilibrium lies in $\mathcal{V}_M$, the equilibrium is globally unique. Establishing contractivity requires bounding the Lipschitz constant of T , which in turn requires controlling the derivatives of the choice probabilities and the social surplus functions. The $\Sigma_m(V_m)$ matrices, which depend on the rivals' values, couple the players' Bellman equations and make the contraction property more difficult to establish than in single-agent dynamic programming problems.

Notably, our results are semiparametric with respect to the error distribution and do not rely on the logit structure of Assumption 9. Theorem 1 requires Assumption 7 (bounded density), and Theorem 2 adds Assumption 8 (exchangeability and Lipschitz density).

Let $\bar{\lambda} \equiv \max_i \lambda_i$ denote the maximal move arrival rate, let $\lambda_{-i} \equiv \sum_{m \neq i} \lambda_m$ denote the total arrival rate of player i 's rivals, and let $\tilde{\lambda} \equiv \max_i \lambda_{-i}$. In the duopoly model, λ_{-i} is the rate of the single rival, and $\tilde{\lambda} = \max\{\lambda_1, \lambda_2\} = \bar{\lambda}$. For any vector $z \in \mathbb{R}^n$, define the *span seminorm* (see Puterman, 2005, pp. 195–196) as²

$$\text{sp}(z) \equiv \max_k z_k - \min_k z_k.$$

Using this notation, define

$$(14) \quad U \equiv \max_i \text{sp}(u_i), \quad \Psi \equiv \max_i \text{sp}(\psi_i), \quad \text{and} \quad R(V) \equiv \max_i \text{sp}(V_i),$$

the maximum range of the flow payoffs, switching costs, and value functions over players, respectively. For the duopoly entry/exit model, U and Ψ reduce to

$$U = \max\{|\theta_M|, |\theta_M + \theta_C|, |\theta_C|\} \quad \text{and} \quad \Psi = |\theta_{EC} - \theta_{SV}|.$$

The following lemma establishes the preliminary results that form the basis of our uniqueness conditions. Part (a) bounds the span of $T(V)$ in terms of the span of V for arbitrary $V \in \mathbb{R}^{NK}$ (not necessarily an equilibrium). The bound is an affine contraction in the span with modulus $\eta/(\rho + \eta)$, and the constant M defined in the lemma is its steady state. Part (b) establishes that the sublevel set $\mathcal{V}_M = \{V : R(V) \leq M\}$ is invariant under

²This is a seminorm: $\text{sp}(z) \geq 0$, $\text{sp}(z + z') \leq \text{sp}(z) + \text{sp}(z')$, and $\text{sp}(\alpha z) = |\alpha| \text{sp}(z)$. Its kernel consists of the constant vectors: $\text{sp}(z) = 0$ if and only if z is constant.

T , which will allow us to apply the Banach fixed-point theorem on $\mathcal{V}_M$. Part (c) shows that every equilibrium must lie in $\mathcal{V}_M$.

Lemma 1 (Operator Bound and Invariant Set). *Suppose Assumptions 1–6 hold. Let $M \equiv (U + \bar{\lambda}\Psi)/\rho$ and $\mathcal{V}_M \equiv \{V \in \mathbb{R}^{NK} : R(V) \leq M\}$.*

(a) *For any $V \in \mathbb{R}^{NK}$,*

$$(15) \quad R(T(V)) \leq \frac{\eta}{\rho + \eta} R(V) + \frac{U + \bar{\lambda}\Psi}{\rho + \eta}.$$

(b) *$T(\mathcal{V}_M) \subseteq \mathcal{V}_M$.*

(c) *If V^* is an equilibrium, so that $T(V^*) = V^*$, then $V^* \in \mathcal{V}_M$.*

The proof of Lemma 1, along with all other proofs, can be found in Appendix A.

Three features of the bound M are worth noting. First, it depends on move arrival rates only through $\bar{\lambda} = \max_i \lambda_i$. Therefore, two parameter vectors with the same total rate $\sum_i \lambda_i$ but differing in how the rates are split among players can yield different bounds. Second, as players become more patient ($\rho \rightarrow 0$), the modulus $\eta/(\rho + \eta)$ in (15) approaches one, so the contraction in the span becomes weaker and its steady state M tends to infinity. For very patient players, small differences in flow payoffs across states can result in large differences in the value function, making it more difficult to guarantee uniqueness. Third, even when the span of flow payoffs U is small, variation in switching costs across states inflates the bound through the Ψ term, so small differences in flow payoffs may not yield a tight bound on the span of the value function.

3.1 Sufficient Condition for Uniqueness

Since every equilibrium lies in $\mathcal{V}_M$ (Lemma 1(c)), it suffices to establish uniqueness on $\mathcal{V}_M$. Theorem 1 gives a sufficient condition for T to be a contraction on $\mathcal{V}_M$. The condition requires the discount rate to be sufficiently large relative to a quantity depending on the span of flow payoffs U , switching costs Ψ , and the move arrival rates.

Theorem 1 (Contraction-Based Uniqueness). *Suppose Assumptions 1–7 hold. If*

$$(16) \quad \rho > \sqrt{2c_F \bar{\lambda}(U + \bar{\lambda}\Psi)},$$

then there is a unique equilibrium.

Remark. Similar contraction arguments date back to Shapley (1953), who showed, in two-player zero-sum stochastic games, that a positive probability of the game stopping at each stage (equivalent to discounting when the probability is constant) makes the value

operator a contraction. Theorem 1 gives an analogous but stronger requirement—discounting large enough relative to the strategic interaction—under which the equilibrium is globally unique in our nonzero-sum, continuous-time setting. Curtat (1996), Horst (2005) introduce similar diagonal-dominance conditions in the context of discrete-time stochastic games with continuous states and actions, to ensure existence rather than uniqueness of equilibrium.

Remark. The factor of $2c_F$ appears in (16) because each rival’s choice probability $\sigma_{mk} = F(\psi_{mk} + V_{m,\ell(m,k)} - V_{mk})$ depends on the difference of two components of V_m . The factor $\tilde{\lambda}$ under the square root is related to the frequency of rival moves. The factor $U + \bar{\lambda}\Psi$ is related to the equilibrium value span M . Flow payoffs are incurred directly, while switching costs are scaled by the rate $\bar{\lambda}$ since they are incurred at decision epochs.

Remark. The number of players enters the condition in (16) only through $\tilde{\lambda}$. This is another implication of asynchronous moves: each rival’s transition term appears additively in the Bellman equation for player i , so the rivals’ move rates are summed in the Lipschitz bound. When players have a common arrival rate $\lambda_i = \lambda$, we have $\tilde{\lambda} = (N - 1)\lambda$ and the condition becomes $\rho > \sqrt{2c_F(N - 1)\lambda(U + \bar{\lambda}\Psi)}$. On the other hand, in single-agent models we have $\tilde{\lambda} = 0$ and (16) holds for all $\rho > 0$.

Remark. We abstract away from exogenous state changes in Section 2.1. However, Lemma 1 and Theorem 1 still hold if they are incorporated. Suppose, following ABBE, that the state transitions exogenously from x_k to x_l at finite rates q_{kl} . For any $\bar{q} \geq \max_k \sum_{l \neq k} q_{kl}$, the total rate of events increases to $\eta = \bar{q} + \sum_m \lambda_m$. Let P denote the corresponding transition matrix implied by uniformization at rate $\bar{q}$. The main change is then the addition of a term $\bar{q}PV_i$ inside the square brackets in (6). Since P does not depend on V and stochastic matrices are nonexpansive under both sp and $\|\cdot\|_\infty$, this term contributes $\bar{q}R(V)$ to the span bound of Lemma 1 and $\bar{q}\|V_i - V'_i\|_\infty$ to the Lipschitz bound of Theorem 1, so both proofs proceed as before with the new value of η .

3.2 Uniqueness under Offsetting Switching Costs

The sufficient condition in the previous section bounds the Lipschitz constant of the Bellman system operator T on the NK -dimensional set $\mathcal{V}_M$. We now return to the duopoly entry/exit model of Section 2.3 and show that when switching costs satisfy an offsetting condition, the eight-equation equilibrium system can be reduced to two equations, and the resulting two-dimensional map admits a sharper contraction condition. The key restriction in this section is that entry costs and scrap values are equal in magnitude.

Assumption 10 (Offsetting Switching Costs). $\theta_{EC} = -\kappa$ and $\theta_{SV} = \kappa$ for some $\kappa \geq 0$.

This condition is satisfied, for example, when entry requires an investment of κ that is fully recovered upon exit, so that entry involves no sunk costs. It yields the following

simplified structure for the choice probabilities, social surplus, and transition probability matrices in the game.

Lemma 2 (Identities with Offsetting Switching Costs). *In the duopoly entry/exit model, under Assumptions 1–8 and 10:*

(a) $F(\theta_{EC} + d) + F(\theta_{SV} - d) = 1$ for all $d \in \mathbb{R}$.

(b) $e(\theta_{SV} - d) - e(\theta_{EC} + d) = \kappa - d$ for all $d \in \mathbb{R}$.

(c) For all $V \in \mathbb{R}^8$, the 2×2 blocks of $\Sigma_2(V_2)$ for state pairs (1, 2) and (3, 4) each have identical rows and therefore

$$(17) \quad [\Sigma_2(V_2) V_1]_1 = [\Sigma_2(V_2) V_1]_2 \quad \text{and} \quad [\Sigma_2(V_2) V_1]_3 = [\Sigma_2(V_2) V_1]_4.$$

Similarly, rows 1 and 3 of $\Sigma_1(V_1)$ are identical, as are rows 2 and 4, and therefore

$$(18) \quad [\Sigma_1(V_1) V_2]_1 = [\Sigma_1(V_1) V_2]_3 \quad \text{and} \quad [\Sigma_1(V_1) V_2]_2 = [\Sigma_1(V_1) V_2]_4.$$

Using the identities in Lemma 2, the eight-dimensional equilibrium system can be reduced to a two-dimensional problem in value differences, which makes it easier to establish equilibrium uniqueness analytically.

Lemma 3 (Constant Competitive Effect). *In the duopoly entry/exit model, suppose Assumptions 1–8 and 10 hold. In any equilibrium,*

$$(19) \quad (V_{14} - V_{12}) - (V_{13} - V_{11}) = (V_{24} - V_{23}) - (V_{22} - V_{21}) = \theta_C / (\rho + \eta).$$

For firm 1, $V_{13} - V_{11}$ represents the value gained by entering when firm 2 is out of the market, and $V_{14} - V_{12}$ is the value gained by entering when firm 2 is active. Lemma 3 establishes that in equilibrium the difference between these two differences—the effect of competition on the value of entry—is the same for both firms and equals the competitive effect θ_C discounted at the effective rate $\rho + \eta = \rho + \lambda_1 + \lambda_2$. The equality in the lemma is exact due to linearity of each Bellman equation in the rival’s choice probabilities, together with Lemma 2(a); these terms need not cancel under simultaneous moves. The following theorem uses this simplification to reduce the equilibrium mapping to two dimensions and establish a sufficient condition for contractivity.

Theorem 2 (Uniqueness with Offsetting Switching Costs). *In the duopoly entry/exit model, suppose Assumptions 1–8 and 10 hold and that $\lambda_1 = \lambda_2 = \lambda$. If*

$$(20) \quad c_F |\tilde{\Delta}| + (c_F^2 + c_F') \tilde{\Delta}^2 < 1, \quad \text{where} \quad \tilde{\Delta} \equiv \lambda \theta_C / [(\rho + \lambda)(\rho + 2\lambda)],$$

then there is a unique equilibrium. The unique equilibrium is symmetric: $V_{11} = V_{21}$, $V_{12} = V_{23}$, $V_{13} = V_{22}$, and $V_{14} = V_{24}$.

Corollary 1. *In the duopoly entry/exit model, suppose Assumptions 1–8 and 10 hold with $\lambda_1 = \lambda_2 = \lambda$. If*

$$c_F \frac{|\theta_C|}{2\lambda} + (c_F^2 + c_F') \left(\frac{\theta_C}{2\lambda} \right)^2 < 1,$$

then there is a unique equilibrium for every discount rate $\rho > 0$.

In contrast to Theorem 1, this condition does not involve the discount rate. It requires that the strength of competition is moderate relative to the move arrival rate: with logit errors, it holds whenever $|\theta_C| < 3.68\lambda$. In (20), $\tilde{\Delta} = \lambda\Delta/(\rho + \lambda)$, where $\Delta = \theta_C/(\rho + \eta)$ is the competitive effect of Lemma 3. The first term can be related to the static case: for a two-player binary game of incomplete information with interaction effect $\tilde{\Delta}$, $c_F|\tilde{\Delta}| < 1$ implies a unique equilibrium (see Aradillas-López, 2020, eq. 15). The second term captures the dynamic effect of each firm’s own future decisions and the curvature of the error CDF F . With logit errors, it tightens the condition from $|\tilde{\Delta}| < 4$ to $|\tilde{\Delta}| < 1.84$.

Remark. The left-hand side of condition (20) is nonmonotone in λ . It is strictly increasing in $|\tilde{\Delta}|$, so both $|\tilde{\Delta}|$ and the left-hand side are maximized at $\lambda^* = \rho/\sqrt{2}$. The condition is easier to satisfy at extreme values of λ , since $\tilde{\Delta} \rightarrow 0$ both as $\lambda \rightarrow 0$, where strategic interaction is removed, and as $\lambda \rightarrow \infty$, where commitment becomes more difficult. The condition is harder to satisfy at intermediate rates $\lambda \approx \rho/\sqrt{2}$, where strategic interaction and commitment are both present.

4 Numerical Experiments

In this section, we investigate numerically the prevalence of uniqueness and multiplicity in the duopoly model of entry and exit. We evaluate the theoretical conditions, measure how far uniqueness extends beyond them, and characterize the parameter regions where multiplicity arises.

4.1 Computational Approach

We solve for equilibria of the dynamic duopoly model of Section 2.3 by applying Newton’s method to the system $H(V) = V - T(V) = 0$ from multiple starting values. We employ the analytical Jacobian, which is derived in Appendix C. Mapping the boundary between uniqueness and multiplicity across the parameter space requires solving the eight-equation system hundreds of billions of times, making the quadratic convergence speed of Newton’s method important. Computation at this scale is possible because the continuous-time model is computationally light: with sequential moves, the cost of each Bellman evaluation grows

linearly, rather than exponentially, in the number of players (Doraszelski and Judd, 2012; ABBE).

Although the problem of computing all Nash equilibria of static games has a rich history (see McKelvey and McLennan, 1996, for a survey), this is not the case for dynamic games. Homotopy methods have been applied to trace out equilibrium correspondences in dynamic games (Borkovsky, Doraszelski, and Kryukov, 2010, 2012), but all-solutions homotopies are available only for systems with special structure, such as polynomial systems. Our Bellman system is not polynomial, so applying homotopy methods would essentially be a different multi-start search heuristic.

We first generate 100,000,000 parameter vectors from a Halton low-discrepancy sequence over $\theta_M \in [0, 5]$, $\theta_C \in [-5, 0]$, $\theta_{EC} \in [-5, 0]$, $\theta_{SV} \in [0, 5]$, $\rho \in [0.01, 0.5]$, and $\lambda \in [0.001, 5]$. We focus on the case of symmetric rates $\lambda_1 = \lambda_2 = \lambda$ for simplicity and assume that the errors follow a type I extreme value distribution, as stated in Assumption 9. With the logit scale normalized to one, payoffs are interpreted in units of the errors. The time unit is left unspecified, so rates and flow payoffs matter only through ratios such as λ/ρ and θ_M/ρ ; references to standard discount rates below assume annual units. For each parameter vector, we perform a multi-start search from 5,000 initial value functions.³ Six starting values are deterministic: the zero vector, the vector of perpetuity values u_{ik}/ρ , and refinements of these based on 20 and 100 iterations of the Bellman system operator T . Since $\lambda_1 = \lambda_2$, the game is expected to admit a symmetric equilibrium. To target such equilibria, we also draw 500 starting values from the subspace of symmetric value functions. The remaining 4,494 starting values are quasi-random Halton draws from a rectangle that contains every equilibrium.⁴ For each asymmetric solution found, we exchange the firm values and refine by Newton’s method to recover the mirror-image counterpart. Distinct equilibria are classified by clustering the conditional choice probabilities in the sup norm with a tolerance of 10^{-3} . We also classify each equilibrium V^* as stable or unstable under iteration of the Bellman system operator T : V^* is stable if the spectral radius ρ_{DT} of the Jacobian $DT(V^*)$ is less than one, so that iterating T from nearby starting values converges to V^* .

³In total, the search performs over 500 billion Newton solves. The implementation uses Fortran 2008 with OpenMP parallelization. Our results were obtained using GNU Fortran 15.1 and LAPACK 3.8.0. The Newton solver uses a relative convergence tolerance of 10^{-13} , with at most 500 iterations per solve. The search is fully deterministic: parameter vectors and starting values are generated by Halton sequences indexed by integer seeds, so the results are exactly reproducible. The full search required 37.9 hours on a single 48-core cluster node (Intel Xeon Platinum 8260).

⁴Following the proof of Lemma 1, we can obtain level bounds: each transition term of the Bellman system operator lies between $\min_k V_{ik}$ and $\max_k V_{ik}$, and the logit surplus satisfies $\max\{z, 0\} \leq \log(1 + e^z) \leq \max\{z, 0\} + \log 2$, so any fixed point satisfies $\rho \max_k V_{ik} \leq \max_k u_{ik} + \lambda_i(\max\{0, \theta_{EC}, \theta_{SV}\} + \log 2)$ and $\rho \min_k V_{ik} \geq \min_k u_{ik}$. Every equilibrium therefore satisfies $\min\{0, \theta_M, \theta_M + \theta_C\}/\rho \leq V_{ik} \leq [\max\{0, \theta_M, \theta_M + \theta_C\} + \lambda_i(\max\{0, \theta_{EC}, \theta_{SV}\} + \log 2)]/\rho$ for all i, k .

Table 2: Distribution of the number of equilibria

Number of Equilibria	Count	Percent
1	96,694,508	96.69%
3	3,303,595	3.30%
5	1,884	0.0019%
7	12	< 0.0001%
Total	99,999,999	100.00%

Notes: Based on 100,000,000 parameter vectors from a Halton sequence. The search was incomplete in 1 case, finding no equilibrium. This case is excluded from the total.

For robustness, after completing searches from each of the 5,000 starting values, we check three additional conditions to determine whether to continue with additional starting values or stop. If an even number of equilibria (including zero) is found, no symmetric equilibrium is found, or a unique unstable equilibrium is found, then the search continues up to two more times. Each of these conditions indicates a missed equilibrium: existence and generic oddness (by analogy with the discrete-time result of [Herings and Peeters, 2004](#)) imply an odd count, a symmetric game is expected to have a symmetric equilibrium, and in the standard configuration we find below, a lone unstable equilibrium suggests the stable mirror-image pair may have been missed. For each continuation, we add 2,000 Halton draws (refined by 20 iterations of T on the first retry and 100 on the second) and 500 symmetric starting values. Even a thorough multi-start search is not guaranteed to find all equilibria (see [Besanko et al., 2010](#)), so in the following “unique” means that the search described returned a single equilibrium.

4.2 Prevalence of Uniqueness

Table 2 reports the distribution of the number of equilibria across the sampled parameter vectors. Overall, the equilibrium is unique in 96.69% of cases, but this rate reflects the wide sampling region. In 3.30% of cases, the solver found three equilibria. This is by far the most common outcome among non-unique cases. The remaining 0.0019% of cases have five or seven equilibria. In all but one case, the solver successfully returned one or more equilibria, and in all successful cases the number of equilibria was odd.⁵ When we restrict attention to economically relevant parameter values where monopoly is viable but duopoly is

⁵Equilibrium existence is guaranteed by Theorem 1 of [Blevins \(2026\)](#) under Assumptions 1–6, so the single case with no equilibrium found represents an incomplete search due to numerical issues encountered with extreme parameter values.

unprofitable ($\theta_M + \theta_C < 0$), firms discount at standard rates ($\rho \leq 0.05$), and moves are not too infrequent relative to discounting ($\lambda \geq 0.1$, so $\lambda \geq 2\rho$), the multiplicity rate increases to 16%. Instead of focusing on the overall uniqueness rate, below we will characterize where in the parameter space multiplicity is likely to occur.

4.3 Evaluation of the Theoretical Conditions

Using the numerical results, we evaluate the sufficient conditions from Theorem 1 and Theorem 2. Since the numerical model uses the logit specification of Assumption 9, we check the conditions with $c_F = 1/4$ and $c'_F = \sqrt{3}/18$ (Lemma 5). To allow for the possibility that uniqueness extends beyond the regions characterized by the theorems, we evaluate each condition at several thresholds τ . In the case of Theorem 1, we can restate the condition as $\tilde{\lambda}(U + \tilde{\lambda}\Psi)/(2\rho^2) < \tau$, where $\tau = 1$ is equivalent to the condition of the theorem and $\tau > 1$ represents a relaxation. For Theorem 2, we apply the contraction inequality (20) exactly and instead relax the switching-cost requirement of Assumption 10. The assumption requires $\theta_{EC} + \theta_{SV} = 0$, so we consider relaxations $|\theta_{EC} + \theta_{SV}| < \tau$ with $\tau > 0$ to allow switching costs that do not offset exactly. Note that the threshold τ is not comparable across the two cases: for Theorem 1 it scales the sufficient inequality, while for Theorem 2 it bounds the deviation from offsetting switching costs in payoff units.

The results are shown in Table 3, which classifies predictions as “positive” or “negative.” A positive prediction here means that the specified uniqueness condition is satisfied. A negative prediction, on the other hand, simply means that the condition is false but does not imply that uniqueness fails. Precision is defined as $TP/(TP + FP)$, where TP and FP are the numbers of true positives and false positives, respectively. Precision is thus the fraction of positive predictions that in fact have a unique equilibrium. Recall is defined as $TP/(TP + FN)$, the fraction of unique cases captured by the condition, where FN denotes the number of false negatives. Since the conditions are sufficient for uniqueness, false positives should not occur, and we expect precision to be 100%. Recall instead measures how much of the overall uniqueness region the condition includes.

There are no false positives at our smallest thresholds, so the numerical results are consistent with both theorems. For Theorem 1, we encounter the first false positive at $\tau = 2.2$, so the lower bound on ρ^2 is conservative by roughly a factor of two in our search. The theorem imposes no special structure on payoffs and applies to games with any finite number of players, any finite state space, and any bounded error-difference density, but the cost is that the coverage is limited to 1.0% of all unique-equilibrium cases. Theorem 2 trades off generality for a condition that does not require strong discounting. For $\tau = 0.2$, we retain perfect precision and include 7.7% of all unique cases. The first false positive appears at $\tau = 0.43$. With $\tau = 3$, the condition identifies 81.4% of all unique-equilibrium cases, but at

Table 3: Coverage of the sufficient conditions for uniqueness

Threshold τ	Predicts Unique	Precision	False Positives	Recall
<i>Panel A. Theorem 1: $\tilde{\lambda}(U + \bar{\lambda}\Psi)/(2\rho^2) < \tau$</i>				
1	999,421	100.0%	0	1.0%
2	1,813,168	100.0%	0	1.9%
5	3,782,160	99.8%	6,067	3.9%
10	6,341,101	99.0%	60,365	6.5%
20	10,293,998	97.9%	216,434	10.4%
<i>Panel B. Theorem 2: (20) and $\theta_{\text{EC}} + \theta_{\text{SV}} < \tau$</i>				
0.2	7,414,923	100.0%	0	7.7%
0.5	17,969,963	99.9%	36	18.6%
1.0	34,048,658	99.9%	13,519	35.2%
2.0	60,531,716	99.6%	213,733	62.4%
3.0	79,447,837	99.1%	727,282	81.4%
∞	94,580,777	98.1%	1,827,215	95.9%

Notes: Precision is the fraction of cases predicted unique that are in fact unique. Recall is the fraction of all unique cases that the condition identifies. Precision is shown as 100.0% only when there are no false positives. In Panel A, $\tau = 1$ is the exact sufficient condition of Theorem 1, and $\tau > 1$ relaxes it by the factor τ . In Panel B, we impose the requirement (20) of Theorem 2 exactly while τ relaxes the switching-cost requirement $\theta_{\text{EC}} + \theta_{\text{SV}} = 0$.

the expense of more false positives.

The false positives that appear as Assumption 10 is relaxed, and that remain when it is dropped entirely ($\tau = \infty$), show that the inequality from Theorem 2 is not by itself sufficient for uniqueness. However, near the exactly offsetting case it appears to be close to sharp. The inequality (20) holds at 94.6% of the 7.8 million vectors with $|\theta_{\text{EC}} + \theta_{\text{SV}}| < 0.2$, and at every vector in this slice with multiple equilibria the left-hand side is at least 1.19, so the inequality could not be relaxed far without introducing false positives. The low recall at small τ in Panel B of Table 3 therefore reflects the small number of sampled vectors with nearly offsetting switching costs, rather than looseness in the inequality. This suggests that a version of Theorem 2 with slack in the restriction $\theta_{\text{EC}} + \theta_{\text{SV}} = 0$ may hold.




Figure 1: Empirical probability of multiple equilibria as a function of model parameters, averaged over all other parameters.

4.4 Determinants of Multiplicity

The numerical results exhibit several patterns in the parameters and underlying economic motives that determine multiplicity. Figure 1 provides a marginal summary of how each model parameter affects multiplicity. Overall, multiplicity decreases with the discount rate ρ (a) and scrap value θ_{SV} (e). Multiplicity increases with the competitive effect $|\theta_C|$ (d) and entry cost $|\theta_{EC}|$ (f). The effect of the move arrival rate λ (b) is nonmonotone. Multiplicity is more likely to occur at intermediate values of λ than at extreme values due to a mechanism we return to below. The monopoly flow profit θ_M (c) gradually increases and then decreases multiplicity. When θ_M is small, firms have little incentive to enter, but when θ_M is large, duopoly becomes viable ($\theta_M + \theta_C > 0$) and both firms find it profitable to enter.

In Figure 2, we examine the joint relationships between various structural parameters. We focus on three determinants of multiplicity: patience and strategic interaction, competition intensity, and switching costs.

Patience and strategic interaction. Figure 2(a) shows the probability of multiple equilibria as a function of the discount rate ρ and the move arrival rate λ . Multiplicity is concentrated in the lower-left corner, where firms are patient and moves are infrequent. As expected from the theoretical results, and previously shown in the marginal plot (Figure 1(b)), the effect of λ is nonmonotone. Increasing λ has two effects: it increases strategic interaction, so firms interact more frequently, but it also decreases firms' ability to commit to actions, since each firm expects the other to re-evaluate and respond more frequently. In the limit, as $\lambda \rightarrow \infty$, switching costs dominate flow payoffs and firms expect rival states to persist only briefly, so best responses become nearly independent of rival states. Under offsetting switching costs, Corollary 1 shows that the equilibrium is then unique. Therefore, multiplicity is highest at intermediate values of λ while uniqueness is more likely for very low or very high values of λ . As a result, the multiplicity contours form arcs whose peaks shift to the right as ρ increases.

Competition intensity. Figure 2(b) shows the multiplicity probability as a function of $|\theta_C|$ and θ_M . The ratio of these parameters is a measure of competition intensity, and indeed there is a sharp threshold just below the diagonal, which represents the boundary between profitable ($|\theta_C|/\theta_M < 1$) and unprofitable ($|\theta_C|/\theta_M > 1$) duopoly. When duopoly is profitable, it is more likely for the model to admit a unique equilibrium where both firms are active. Otherwise, the market cannot support both firms, creating a coordination problem that gives rise to multiple asymmetric equilibria. These equilibria have a self-fulfilling character, where each firm stays active only when it expects its rival to exit. We see the same pattern in Figure 2(d), which plots the competition intensity ratio against λ . Multiplicity is rare when $|\theta_C|/\theta_M \ll 1$ regardless of the move arrival rate.

Switching costs. Figure 2(c) shows how multiplicity relates to switching costs. Multiplicity is most likely in the upper left (high entry cost, low scrap value) and least likely in the lower right (low entry cost, high scrap value), with clear contour lines indicating that the probability of multiplicity depends heavily on the net switching cost $\theta_{EC} + \theta_{SV}$. High scrap values induce exit, which reduces the coordination problem discussed above, while high entry costs exacerbate it. This also highlights the role of switching costs in commitment. When firms have to pay high entry costs and recover little value upon exit (upper-left corner), state transitions are costly in both directions. Market structure adjusts slowly, so each firm's current status serves as a commitment, and as a result, it is easier to support multiple equilibria. On the other hand, when firms can enter easily and recover large values when they exit (lower-right corner), switching is relatively costless and commitment is weak, making it more difficult to support multiple equilibria.




Figure 2: Empirical probability of multiple equilibria as a function of important model parameter pairs. Shading indicates the fraction of parameter vectors in each bin with multiple equilibria, from white (all unique) to black (all multiple). Based on 100,000,000 parameter vectors from a Halton sequence with $\lambda_1 = \lambda_2$.

Distributions of payoff-related quantities. Figure 3 displays the densities of four payoff-related quantities conditional on unique-equilibrium and multiple-equilibria cases. The dark gray densities are based on the 3.31% of parameter vectors that exhibit multiple equilibria. First, the monopoly profit θ_M (a) is nearly uniformly distributed among unique-equilibrium cases, but is unimodal under multiplicity. The duopoly profit $\theta_M + \theta_C$ (b) shows more separation. The distribution for multiple-equilibria cases is concentrated at negative values (unprofitable duopoly) while the unique-equilibrium distribution resembles the unconditional triangular distribution (since duopoly profit is the sum of two uniformly distributed variables whose supports have equal widths). The distribution of the competition intensity ratio $|\theta_C|/\theta_M$ (c) for multiple-equilibria cases is concentrated just above the threshold near one, while the density for unique-equilibrium cases is highest below one. Finally, the distribution




Figure 3: Conditional densities of payoff-related quantities by equilibrium multiplicity. Each panel shows the distribution of a quantity among parameter vectors with a unique equilibrium (light) and among those with multiple equilibria (dark).

of the net switching cost $\theta_{EC} + \theta_{SV}$ (d) is shifted toward negative values for multiple-equilibria cases, where entry costs exceed scrap values (increasing commitment). The distribution conditional on uniqueness is again nearly triangular.

4.5 Equilibrium Selection and Long-Run Outcomes

In this section we quantify how the equilibria differ economically in three ways: in long-run market structure, in producer surplus, and in which firm occupies the market. Each distinct equilibrium induces a continuous-time Markov chain over market states x_k . For each equilibrium at each of the 3,305,491 parameter vectors where we found multiple equilibria, we compute the stationary distribution along with two measures of producer surplus: the stationary flow of operating profits and the present value $V_{11} + V_{21}$ in state $x_1 = (0, 0)$. The latter includes switching costs and the expected surplus from the choice-specific errors. Table 4 reports the distribution of spreads: the maximum minus the minimum over equilibria at each parameter vector.

The long-run probability of duopoly differs across all equilibria, including the unstable

Table 4: Spreads of long-run outcomes across equilibria

Long-Run Outcome	All Equilibria			Stable Pair		
	Median	90%	Max	Median	90%	Max
<i>Aggregate outcomes</i>						
Duopoly share	0.128	0.337	0.892	0.000	0.000	0.000
Monopoly share	0.125	0.336	0.892	0.000	0.000	0.000
Empty-market share	0.001	0.011	0.099	0.000	0.000	0.000
Producer surplus, flow / U	0.164	0.406	0.897	0.000	0.000	0.000
Producer surplus, PDV / (U/ρ)	0.063	0.250	0.746	0.000	0.000	0.000
<i>Firm identity</i>						
Firm 1 active share	0.811	0.988	1.000	0.810	0.987	1.000
Firm 1 value, PDV / (U/ρ)	0.238	0.468	0.859	0.238	0.469	0.859

Notes: Spread (maximum minus minimum) of each outcome across equilibria at each of the 3,305,491 parameter vectors with multiple equilibria. Stable Pair columns condition on the 3,288,866 vectors with exactly two stable equilibria. Surplus and value spreads are expressed in units of the span of flow payoffs U or the present value U/ρ , so that entries are comparable across parameter vectors. PDV denotes present discounted value.

one, by 0.128 at the median across parameter vectors with multiple equilibria, and by 0.337 at the 90th percentile. Since the shares sum to one, the duopoly and monopoly spreads can differ by at most the empty-market spread. The latter is small for most parameter vectors, resulting in the nearly identical marginal distributions in Table 4. What equilibrium selection changes, therefore, is how the stationary distribution splits between duopoly and monopoly, not whether the market is served. Equilibrium selection also affects total producer surplus: the flow measure differs by 0.164 at the median (normalized by the span of flow payoffs U).

Across all sampled parameter vectors, we found that 99.999% of unique equilibria are stable. Among parameter vectors that produce multiple equilibria, 99.49% exhibit a standard configuration with three equilibria, two stable and one unstable. Table 5 displays the full distribution of equilibrium counts and stability among multiple-equilibria cases.⁶ Stability is important for both computation, since only stable equilibria can be reached by iterating T

⁶The 0.45% of multiple-equilibria cases with three unstable equilibria have the same structure as the standard configuration, one symmetric equilibrium and a mirror-image pair, but the paired equilibria are no longer stable. The smallest spectral radius among the three equilibria is 1.21 at the median. These cases are characterized by low discount and move arrival rates (median $\rho = 0.02$ and $\lambda = 0.15$) and strong competition effects (median $|\theta_C|/\theta_M = 3.1$).

Table 5: Stability of multiple equilibria

All Equilibria		Symmetric Equilibria		Occurrence	
Total	Stable	Stable	Unstable	Count	Percent
3	0	0	1	14,878	0.45%
3	2	0	1	3,288,717	99.49%
5	2	0	1	144	0.0044%
5	3	1	0	1,740	0.05%
7	2	0	1	5	0.0002%
7	4	0	1	7	0.0002%

from all nearby starting values, and estimation, where an analogous property of the iterative continuous-time nested pseudo-likelihood (CTNPL) estimator determines its convergence behavior.

Within the standard three-equilibrium configuration, selection between the two stable equilibria is primarily distributional in our symmetric model: it determines which firm occupies the market, but not the distributions of aggregate market structure or total producer surplus. The two stable equilibria are mirror images with swapped firm roles, so the aggregate-outcome entries in the stable-pair columns of Table 4 are identically zero. Firm 1’s stationary active share, on the other hand, differs across equilibria by 0.810 at the median and 0.987 at the 90th percentile. Being the firm with the higher activity probability is worth 0.238 at the median in present discounted value terms (normalized by U/ρ).

5 Implications for Estimation and Counterfactuals

There are well-known consequences of multiplicity for applied researchers: estimation is complicated by the lack of a well-defined reduced form, and counterfactual predictions depend on equilibrium selection. In this section, we quantify both of these using the equilibria computed in Section 4 in combination with a targeted Monte Carlo experiment and a simple counterfactual exercise.

5.1 Estimation under Multiplicity

We consider maximum likelihood estimation (MLE) of $\theta = (\theta_M, \theta_C, \theta_{EC})$ using data on the states and actions observed in 200 independent markets over a time interval $[0, 25]$. Following the literature, we take (ρ, λ) as known, and we fix the scrap value θ_{SV} at its true value to normalize payoffs for identification (Magnac and Thesmar, 2002, Aguirregabiria and Suzuki,




Figure 4: Expected log-likelihood of each equilibrium branch (stable branches solid gray, unstable branch dashed) along a θ_C slice through a bifurcation point, for 200 markets observed over $[0, 25]$ under the stable DGP. The envelope over branches (solid black, oracle MLE objective) is maximized at the true parameter value. Circles mark the branch selected by the solver's default starting value (naïve MLE objective).

2014).⁷

Four data-generating processes (DGPs) are specified along a θ_C slice that passes through a bifurcation in the equilibrium mapping.⁸ Two DGPs are on the unique-equilibrium side of the bifurcation, and the remaining two share a single multiple-equilibria parameter vector, but one generates data from a stable equilibrium and the other from the unstable equilibrium (the stable and unstable DGPs, respectively).

The four DGPs can be understood visually in Figure 4, which plots the expected log-likelihood of each equilibrium branch along the θ_C slice, with the data fixed at the stable DGP (true $\theta_C = -4$). Data are simulated from the equilibrium CCPs corresponding to the top branch (solid gray line). The remaining parameters are held fixed at true values $(\theta_M, \theta_{EC}, \theta_{SV}, \rho, \lambda) = (2, -2, 1, 0.05, 1)$. The unstable DGP has the same true $\theta_C = -4$, but data are simulated instead from the CCPs corresponding to the unstable branch (dashed gray line). The remaining two DGPs have true parameters on the unique-equilibrium side of

⁷For any b , replacing $(\theta_M, \theta_{EC}, \theta_{SV})$ by $(\theta_M + \rho b, \theta_{EC} - b, \theta_{SV} + b)$ shifts each firm's value function by b in the states where it is active and leaves all equilibrium choice probabilities unchanged. Fixing θ_{SV} removes this indeterminacy.

⁸We assume all markets play the same equilibrium, allowing data to be pooled. Otsu, Pesendorfer, and Takahashi (2016) consider testing whether different markets play different equilibria.

the bifurcation at $\theta_C = -3.1$, with one intentionally near the branching point ($\theta_C = -3$, the near-boundary DGP) and the other far inside the uniqueness region ($\theta_C = -2$, the interior DGP). Neither theorem covers the DGP parameters, so uniqueness at the near-boundary and interior DGPs was verified numerically.⁹

We carry out a Monte Carlo experiment that compares three estimators across the four DGPs. First, we consider a naïve nested fixed point MLE (Rust, 1987) that assumes uniqueness. For each trial parameter vector, this estimator finds a single equilibrium by applying Newton’s method to the equilibrium system, initialized at $V_0 = [0, 0, \theta_M/\rho, (\theta_M + \theta_C)/\rho, 0, \theta_M/\rho, 0, (\theta_M + \theta_C)/\rho]^\top$, the vector of perpetuity values u_{ik}/ρ , with components (stacked across players) corresponding to the present discounted values of monopoly and duopoly flows in the respective states. Another default a researcher might choose is the zero vector. Both are symmetric across firms, a point we return to below.

The second estimator is the oracle MLE, which serves as an infeasible benchmark: for each parameter vector we “warm-start” the Newton solver at the true DGP equilibrium’s value function. The third estimator is the CTNPL estimator of Blevins and Kim (2024), an iterative estimator that solves a linear valuation system given the current choice probabilities at each step, maximizes the implied pseudo-likelihood over θ , and then updates the CCPs using the best-response mapping. We initialize this estimator at nonparametric frequency estimates of the CCPs, which consistently estimate those of the equilibrium that generated the data without requiring the researcher to know it in advance. We compare CTNPL after a single iteration (the “two-step” estimator) and after iterating to convergence.¹⁰

In Table 6, we compare the estimates over 500 Monte Carlo replications and report the bias and root mean squared error (RMSE) of each estimator across the four DGPs. In all cases, we initialize the optimizer at the corresponding true parameter vector to separate equilibrium selection from optimization error; however, the initializing equilibria (value functions or CCPs) differ across estimators as previously described. The final column of Table 6 reports the percentage of replications for which the estimator selects the correct DGP equilibrium.

As expected, the oracle MLE tracks the correct equilibrium branch as the parameters vary near the truth. This estimator maximizes the sample counterpart of the upper envelope in Figure 4, which is maximized at the true parameter vector. Conditional on correct equilibrium selection, the identified parameters are estimated precisely. The gap between

⁹We selected the DGP to be representative: a standard discount rate, $\rho = 0.05$, flow payoffs under which monopoly is profitable but duopoly is not, and a spread in long-run duopoly share across equilibria near the median among multiple-equilibria cases.

¹⁰We terminate the iteration when either the maximum change in the CCPs falls below 10^{-6} or after 2000 stages. All replications converged except in the near-boundary DGP, where 99.80% converged (median 119 iterations); replications that did not converge are included in Table 6 at their estimates after 2000 stages.

Table 6: Estimation under multiplicity: Monte Carlo

Estimator	Bias (RMSE)			Selects
	θ_M	θ_C	θ_{EC}	DGP Eq.
<i>Unique (interior): $\theta_C = -2.00$, $\rho_{DT} = 0.976$, $\rho_{NPL} = 0.478$</i>				
Naïve MLE	+0.005 (0.099)	-0.004 (0.134)	+0.000 (0.067)	100.00%
Oracle MLE	+0.005 (0.099)	-0.004 (0.134)	+0.000 (0.067)	100.00%
CTNPL (one iter.)	-0.002 (0.103)	+0.004 (0.138)	-0.001 (0.067)	100.00%
CTNPL (converged)	+0.005 (0.099)	-0.004 (0.134)	+0.000 (0.067)	100.00%
<i>Unique (near boundary): $\theta_C = -3.00$, $\rho_{DT} = 0.990$, $\rho_{NPL} = 0.968$</i>				
Naïve MLE	+0.002 (0.082)	-0.006 (0.148)	-0.002 (0.064)	100.00%
Oracle MLE	+0.002 (0.082)	-0.006 (0.148)	-0.002 (0.064)	100.00%
CTNPL (one iter.)	-0.007 (0.089)	+0.010 (0.159)	-0.003 (0.064)	100.00%
CTNPL (converged)	-0.002 (0.079)	+0.011 (0.130)	-0.001 (0.064)	100.00%
<i>Multiple (stable DGP): $\theta_C = -4.00$, $\rho_{DT} = 0.976$, $\rho_{NPL} = 0.065$</i>				
Naïve MLE	-0.904 (0.905)	+0.334 (0.395)	-0.882 (0.887)	0.00%
Oracle MLE	-0.000 (0.101)	+0.008 (0.220)	-0.009 (0.119)	100.00%
CTNPL (one iter.)	-0.017 (0.110)	+0.042 (0.240)	-0.018 (0.135)	99.60%
CTNPL (converged)	-0.001 (0.110)	+0.009 (0.232)	-0.009 (0.125)	100.00%
<i>Multiple (unstable DGP): $\theta_C = -4.00$, $\rho_{DT} = 1.126$, $\rho_{NPL} = 1.463$</i>				
Naïve MLE	-0.000 (0.079)	-0.008 (0.197)	-0.002 (0.069)	100.00%
Oracle MLE	-0.000 (0.079)	-0.008 (0.197)	-0.002 (0.069)	100.00%
CTNPL (one iter.)	-0.013 (0.087)	+0.021 (0.210)	-0.003 (0.069)	100.00%
CTNPL (converged)	-0.161 (0.175)	+0.789 (0.797)	+0.016 (0.073)	0.00%

Notes: Bias (RMSE in parentheses) over 500 replications; 200 markets observed continuously over $[0, 25]$; (ρ, λ) known and θ_{SV} fixed at its true value; optimizers start at the true parameters. The DGPs vary θ_C along a slice through a multiple-equilibria parameter vector (multiplicity boundary at $\theta_C \approx -3.1$). Section 5.1 defines the estimators. At the stable DGP the naïve solver’s default start selects an equilibrium 0.31 from the DGP equilibrium in maximum CCP distance. The final column reports the share of replications in which the estimated CCPs lie within 0.1 in maximum absolute CCP distance of the DGP equilibrium. ρ_{DT} and ρ_{NPL} are the spectral radii of the Jacobian $DT(V^*)$ and the CTNPL operator at the DGP equilibrium, respectively.

the oracle MLE and other estimators therefore represents the equilibrium selection error.

The naïve MLE, however, assumes there is always a unique equilibrium and maximizes the likelihood of whichever equilibrium its inner solver returns. As expected, this estimator performs well under both unique-equilibrium DGPs. The results for the multiple-equilibria DGPs are somewhat counterintuitive: the naïve MLE performs well under the unstable DGP but has large biases under the stable DGP.

The reason is that Newton’s method solves the equilibrium system directly, so the equilibrium it returns depends on the starting value rather than on stability under iteration of T . If $\lambda_1 = \lambda_2$, the equilibrium system commutes with the firm-relabeling map Ξ from the proof of Theorem 2 (which does not depend on Assumption 10), and so do its Jacobian and each Newton step. Thus the iterates generated by Newton’s method when initialized at a symmetric value function must be symmetric and can only converge to a symmetric equilibrium. Since this slice has only one symmetric equilibrium branch—the unstable branch—the naïve MLE tracks it, as shown by the circles in Figure 4. The naïve objective coincides with the oracle’s for the unstable DGP. As the final column of Table 6 shows, the selection is correct in 100% of cases for the unstable DGP and is consistently wrong for the stable DGP. The estimator’s consistency in the former case therefore hinges on a coincidence between the equilibrium that generated the data and the one equilibrium a symmetric initialization can reach.

This dependence on a coincidence is not specific to our DGP. In the equilibrium sets computed in Section 4, the parameter vectors with multiple equilibria have exactly one symmetric equilibrium, which is unstable except in the five-equilibria cases with three stable equilibria (Table 5). Outside these cases, the coincidence breaks down when data are generated by a stable equilibrium.

CTNPL selects the equilibrium from preliminary frequency estimates of the CCPs. With one iteration, it achieves bias and RMSE nearly as low as the oracle MLE under the unique-equilibrium DGPs. Under the multiple-equilibria DGPs, the results differ substantially after iteration to convergence. The relevant diagnostic is not the spectral radius ρ_{DT} , which governs stability under iteration of T , but rather the spectral radius ρ_{NPL} of the CTNPL operator, which maps the current choice probabilities to their update after one iteration, including the re-estimation of θ . This is the continuous-time analogue of the discrete-time NPL convergence condition discussed by Kasahara and Shimotsu (2012). We report both measures in each DGP panel heading in Table 6 for comparison.¹¹ At all four DGPs,

¹¹See Aguirregabiria and Mira (2002, 2007), Kasahara and Shimotsu (2008, 2012), Pesendorfer and Schmidt-Dengler (2010), Egesdal et al. (2015), Aguirregabiria and Marcoux (2021), Bugni and Bunting (2021), Blevins and Kim (2024), and Dearing and Blevins (2025), among others, for discussions of convergence of iterative estimators.

the two radii fall on the same side of unity.¹² Since the CTNPL iteration is contractive ($\rho_{\text{NPL}} = 0.065$) in the stable DGP, iterating to convergence selects the correct equilibrium. For the unstable DGP, the CTNPL operator is also unstable ($\rho_{\text{NPL}} = 1.463$), so further iteration selects the wrong equilibrium and increases bias in the estimated parameters. When the equilibrium is unique, both for the interior DGP and even near the multiplicity boundary (where $\rho_{\text{NPL}} = 0.968$), the converged CTNPL estimates remain on par with the MLE. In practice, although the researcher does not know which equilibrium was realized in the DGP, ρ_{NPL} can be evaluated at the consistent one-step CTNPL estimates before iterating further.

5.2 Counterfactual Selection Ambiguity

Setting aside estimation, equilibrium selection remains an issue for counterfactual policy analysis. Even if the structural parameters can be consistently estimated, the equilibrium played under counterfactual parameters is indeterminate under multiplicity (Aguirregabiria, 2012, Jun and Pinkse, 2020). To measure the effects of this indeterminacy, we implement an entry subsidy in the dynamic duopoly model: we set $\theta_{\text{EC}} \rightarrow \theta_{\text{EC}} + s$ at the first 10,000 parameter vectors in the Halton sequence with multiple equilibria, at least one of which is stable. We take the baseline equilibrium to be the stable equilibrium favoring firm 1 (a labeling normalization) and compute the full set of counterfactual equilibria following the same multi-start procedure used in Section 4.

Even with the baseline equilibrium known, the predicted effect of the subsidy depends on which counterfactual equilibrium arises in the market. The resulting ambiguity can be seen in Table 7, which reports quantiles of the distribution (across parameter vectors) of spreads of the predicted effects across counterfactual equilibria. Each spread is the largest difference between the predictions of any two selection rules.

We highlight three main results. First, the subsidy itself changes the coordination problem in the market and reduces multiplicity: the share of parameter vectors with multiple counterfactual equilibria falls from 86% at $s = 0.25$ to 53% at $s = 1.0$. Indeed, the subsidy raises the net switching cost $\theta_{\text{EC}} + \theta_{\text{SV}}$ —a key determinant of multiplicity, as shown in Figure 2(c). Second, among stable counterfactual equilibria, aggregate predictions such as the duopoly share do not depend on selection: because the stable equilibria typically occur in mirror-image pairs, the distribution is zero up to the 90th percentile. The nonzero maximum arises due to the five-equilibria cases with three stable equilibria. Third, counterfactual effects involving firm identity—which firm serves the market—are very sensitive to equilibrium selection even when attention is restricted to stable equilibria, as essentially the entire market

¹²These are distinct operators and their radii need not agree. ρ_{DT} is bounded below by the effective discount factor $(\lambda_1 + \lambda_2)/(\rho + \lambda_1 + \lambda_2) = 0.976$ here, since adding a constant c to every component of V shifts $T(V)$ by $c\eta/(\rho + \eta)$, so the constant vector is an eigenvector of DT with eigenvalue $\eta/(\rho + \eta)$.

Table 7: Selection ambiguity of counterfactual predictions under multiplicity

Subsidy	Min	10%	Median	90%	Max	n
<i>Change in duopoly share, all counterfactual eq.</i>						
+0.25	0.000	0.022	0.129	0.338	0.775	8,641
+0.5	0.000	0.023	0.132	0.339	0.775	7,340
+1.0	0.000	0.026	0.137	0.346	0.771	5,268
Offsetting	0.001	0.037	0.169	0.365	0.713	1,977
<i>Change in duopoly share, stable counterfactual eq.</i>						
+0.25	0.000	0.000	0.000	0.000	0.342	8,640
+0.5	0.000	0.000	0.000	0.000	0.342	7,339
+1.0	0.000	0.000	0.000	0.000	0.344	5,266
Offsetting	0.000	0.000	0.000	0.000	0.000	1,977
<i>Change in firm 1 active share, stable counterfactual eq.</i>						
+0.25	0.003	0.376	0.804	0.988	1.000	8,640
+0.5	0.010	0.367	0.802	0.988	1.000	7,339
+1.0	0.007	0.365	0.804	0.990	1.000	5,266
Offsetting	0.039	0.378	0.818	0.995	1.000	1,977

Notes: Entry subsidy $\theta_{EC} \rightarrow \theta_{EC} + s$ at 10,000 parameter vectors with multiple equilibria and at least one stable equilibrium. For each, we select a stable equilibrium as the baseline. The Offsetting row sets $s = -(\theta_{EC} + \theta_{SV})$ where this is positive, so that $\theta_{EC} + s + \theta_{SV} = 0$. Each panel reports the spread (maximum minus minimum) across the counterfactual equilibria of the stated effect. All panels condition on the presence of multiple counterfactual equilibria, which results in fewer cases n as the fixed subsidy increases. The stable panels further require at least two stable counterfactual equilibria.

is reallocated between monopoly equilibria.

These findings suggest that researchers should carefully consider possible multiplicity of counterfactual equilibria, especially for policy questions involving firm identity. Label-invariant aggregates coincide across stable counterfactual equilibria in our symmetric model, although this need not hold with asymmetric firms.

For counterfactual parameters satisfying the conditions of Theorem 1 or Theorem 2, the analysis can proceed without ambiguity. To illustrate, consider the offsetting subsidy $s = -(\theta_{EC} + \theta_{SV})$, which makes entry costs fully recoverable upon exit, as in Assumption 10. The uniqueness condition of Theorem 2 then depends only on θ_C , λ , and ρ , which are not affected by the subsidy. At the 9,090 parameter vectors where $s > 0$, the condition holds in 61% of cases, and the search finds a unique equilibrium in 78%. Table 7 reports the spreads for the rest.

6 Conclusion

We derive closed-form sufficient conditions for equilibrium uniqueness in continuous-time dynamic discrete games. The first condition applies to games with any finite number of players, any finite state space, and binary actions. It guarantees uniqueness when the discount rate is sufficiently large relative to move arrival rates and variation in payoffs. The second condition ensures equilibrium uniqueness in a duopoly model without sunk costs. When competition is moderate relative to the frequency of moves, uniqueness holds even for arbitrarily patient firms.

We then conduct a large-scale computational study of the determinants of uniqueness and multiplicity in a duopoly model of entry and exit. Multiplicity occurs primarily in economically relevant regions involving patient firms, markets where duopoly is unprofitable, and intermediate move arrival rates. Elsewhere, we find uniqueness across most of the sampled parameter space, well beyond the regions covered by our two conditions. Extending the conditions to cover this larger region is a natural direction for future work.

Where multiplicity does occur, we investigate its economic and econometric consequences. Even at the same parameter value, multiple equilibria may differ meaningfully in terms of long-run market structure or producer surplus. The naïve full-solution estimator, which assumes uniqueness, can track an unstable equilibrium branch even if the data-generating equilibrium is stable. The CTNPL estimator, which estimates the equilibrium from the data, remains reliable when multiplicity is present, but we recommend using the iterative version only when the CTNPL operator is contractive, which can be checked at the one-step estimates. For counterfactuals, predictions of aggregate quantities coincide across stable counterfactual equilibria under firm symmetry, while predictions of quantities involving firm

identity largely depend on equilibrium selection. Because our theoretical conditions can be checked directly from the model primitives, they aid researchers in confirming uniqueness rather than imposing it.

A Proofs

This section contains the proofs of all formal results stated in the main text. Following these, we present two supporting lemmas.

Proof of Lemma 1. Write $R = R(V)$ for brevity. For any $V \in \mathbb{R}^{NK}$, multiplying (6) by $\rho + \eta$ gives, for player i in state k , $(\rho + \eta)[T(V)]_{ik} = u_{ik} + \sum_{m \neq i} \lambda_m [\Sigma_m(V_m)V_i]_k + \lambda_i E_{ik}(V_i)$. For any two states k and l , subtracting yields:

$$\begin{aligned} (\rho + \eta) \{[T(V)]_{ik} - [T(V)]_{il}\} &= (u_{ik} - u_{il}) \\ &+ \sum_{m \neq i} \lambda_m \{[\Sigma_m(V_m)V_i]_k - [\Sigma_m(V_m)V_i]_l\} + \lambda_i [E_{ik}(V_i) - E_{il}(V_i)]. \end{aligned}$$

Stochastic matrices are nonexpansive under sp .¹³ Therefore, each transition term satisfies $\text{sp}(\Sigma_m(V_m)V_i) \leq \text{sp}(V_i) \leq R$. For the surplus term, note that $E_{ik}(V_i) = L(\psi_{ik} + V_{i,\ell(i,k)}, V_{ik})$, where $\ell(i, k)$ denotes the destination state when player i switches from state k . By Lemma 4(c), L is Lipschitz continuous with constant 1 in the ℓ^∞ norm. Comparing states k and l , the second arguments of L differ by at most $\text{sp}(V_i) \leq R$, while the first arguments differ by at most $\text{sp}(\psi_i) + \text{sp}(V_i) \leq \Psi + R$. Therefore, $|E_{ik}(V_i) - E_{il}(V_i)| \leq \max\{R, \Psi + R\} = \Psi + R$. Combining results gives

$$(\rho + \eta) |[T(V)]_{ik} - [T(V)]_{il}| \leq U + \sum_{m \neq i} \lambda_m R + \lambda_i (\Psi + R) = U + \eta R + \lambda_i \Psi.$$

Taking the maximum over players i and state pairs (k, l) gives

$$(\rho + \eta) R(T(V)) \leq U + \eta R + \bar{\lambda} \Psi.$$

Dividing by $\rho + \eta$ gives part (a).

For part (b), if $R(V) \leq M$ then $(\rho + \eta) R(T(V)) \leq U + \eta M + \bar{\lambda} \Psi = (\rho + \eta)M - \rho M + U + \bar{\lambda} \Psi = (\rho + \eta)M$, where the last equality uses $\rho M = U + \bar{\lambda} \Psi$ (from the definition of M). Hence $R(T(V)) \leq M$.

For part (c), at an equilibrium $V^* = T(V^*)$, so $R(V^*) = R(T(V^*))$. Substituting into (15) and rearranging: $\rho R(V^*)/(\rho + \eta) \leq (U + \bar{\lambda} \Psi)/(\rho + \eta)$, giving $R(V^*) \leq M$. ■

¹³If P is a stochastic matrix and z is a conformable vector, then each entry of Pz is a convex combination of the components of z . It follows that $\max(Pz) \leq \max z$ and $\min(Pz) \geq \min z$. Subtracting yields $\text{sp}(Pz) \leq \text{sp}(z)$. Since each entry of Pz lies in $[\min z, \max z]$, we also have $\|Pz\|_\infty \leq \|z\|_\infty$.

Proof of Theorem 1. We show that T is a contraction on the invariant set $\mathcal{V}_M$ defined in Lemma 1. Let $V, V' \in \mathcal{V}_M$. From (6), the operator difference for player i is

$$T_i(V) - T_i(V') = \frac{1}{\rho + \eta} \left\{ \sum_{m \neq i} \lambda_m [\Sigma_m(V_m)V_i - \Sigma_m(V'_m)V'_i] + \lambda_i [E_i(V_i) - E_i(V'_i)] \right\}.$$

We bound the two types of terms on the right separately.

Rivals' transition terms. For each rival $m \neq i$, we add and subtract $\Sigma_m(V_m)V'_i$:

$$\Sigma_m(V_m)V_i - \Sigma_m(V'_m)V'_i = \Sigma_m(V_m)(V_i - V'_i) + (\Sigma_m(V_m) - \Sigma_m(V'_m))V'_i.$$

Since $\Sigma_m(V_m)$ is stochastic (row sums equal 1), the first part satisfies

$$\|\Sigma_m(V_m)(V_i - V'_i)\|_\infty \leq \|V_i - V'_i\|_\infty.$$

For the second part, let $\Delta\Sigma_m \equiv \Sigma_m(V_m) - \Sigma_m(V'_m)$. From (5), row k of $\Delta\Sigma_m$ is $[\sigma_{mk}(V_m) - \sigma_{mk}(V'_m)](\iota_{\ell(m,k)} - \iota_k)^\top$, where ι_k denotes the k th standard basis vector in $\mathbb{R}^K$, so

$$(\Delta\Sigma_m V'_i)_k = [\sigma_{mk}(V_m) - \sigma_{mk}(V'_m)](V'_{i,\ell(m,k)} - V'_{ik}).$$

Each switching probability has the form $\sigma_{mk} = F(\psi_{mk} + V_{m,\ell(m,k)} - V_{mk})$, where the argument depends on V_m through a difference of two components. The change in this argument is bounded by $|(V_{m,\ell(m,k)} - V'_{m,\ell(m,k)}) - (V_{mk} - V'_{mk})| \leq 2\|V_m - V'_m\|_\infty$, and composing with the Lipschitz constant c_F of F (Lemma 4(d)) gives

$$|\sigma_{mk}(V_m) - \sigma_{mk}(V'_m)| \leq 2c_F\|V_m - V'_m\|_\infty.$$

Since $V' \in \mathcal{V}_M$, we also have $|V'_{i,\ell(m,k)} - V'_{ik}| \leq \text{sp}(V'_i) \leq R(V') \leq M$, giving

$$\|\Delta\Sigma_m V'_i\|_\infty \leq 2c_F\|V_m - V'_m\|_\infty \cdot M.$$

Own surplus term. Write $E_{ik}(V_i) = L(\psi_{ik} + V_{i,\ell(i,k)}, V_{ik})$ as in the proof of Lemma 1. Since V_i and V'_i are evaluated at the same state k , the switching cost ψ_{ik} cancels and the arguments of L differ by $(V_{i,\ell(i,k)} - V'_{i,\ell(i,k)}, V_{ik} - V'_{ik})$. By Lemma 4(c),

$$|E_{ik}(V_i) - E_{ik}(V'_i)| \leq \max\{|V_{i,\ell(i,k)} - V'_{i,\ell(i,k)}|, |V_{ik} - V'_{ik}|\} \leq \|V_i - V'_i\|_\infty.$$

Combining. Collecting the terms above:

$$\|T_i(V) - T_i(V')\|_\infty \leq \frac{\lambda_{-i} + \lambda_i}{\rho + \eta} \|V_i - V'_i\|_\infty + \frac{2c_F M}{\rho + \eta} \sum_{m \neq i} \lambda_m \|V_m - V'_m\|_\infty,$$

recalling that $\lambda_{-i} = \sum_{m \neq i} \lambda_m$. Since $\|V_i - V'_i\|_\infty \leq \|V - V'\|_\infty$ and $\|V_m - V'_m\|_\infty \leq \|V - V'\|_\infty$ for all m , while $\lambda_{-i} + \lambda_i = \eta$ and $\lambda_{-i} \leq \tilde{\lambda}$, taking the maximum over players i gives, for the full NK -dimensional operator,

$$\|T(V) - T(V')\|_\infty \leq \underbrace{\frac{\eta + 2c_F \tilde{\lambda} M}{\rho + \eta}}_{\gamma} \|V - V'\|_\infty.$$

Deriving the condition. The contraction condition $\gamma < 1$ requires $\eta + 2c_F \tilde{\lambda} M < \rho + \eta$, which simplifies to $2c_F \tilde{\lambda} M < \rho$. Substituting $M = (U + \bar{\lambda} \Psi)/\rho$ from Lemma 1:

$$\frac{2c_F \tilde{\lambda} (U + \bar{\lambda} \Psi)}{\rho} < \rho,$$

which rearranges to (16).

Under this condition, T is a contraction on $\mathcal{V}_M$ with respect to $\|\cdot\|_\infty$. The set $\mathcal{V}_M$ is nonempty (it contains $V = 0$, since $M \geq 0$) and closed in $\mathbb{R}^{NK}$ (since R is continuous), hence complete. By Lemma 1(b), T maps $\mathcal{V}_M$ to itself. The Banach fixed-point theorem therefore gives a unique fixed point in $\mathcal{V}_M$. By Lemma 1(c), every equilibrium lies in $\mathcal{V}_M$, so uniqueness is global. ■

Proof of Lemma 2. Part (a): $F(-\kappa + d) + F(\kappa - d) = F(-\kappa + d) + F(-(-\kappa + d)) = 1$, by Lemma 4(e).

Part (b): $e(\theta_{SV} - d) - e(\theta_{EC} + d) = e(\kappa - d) - e(-(\kappa - d)) = \kappa - d$ by Lemma 4(f).

Part (c): Consider $\Sigma_2(V_2)$ from (11), which is block diagonal with 2×2 blocks for state pairs (1, 2) and (3, 4). Within each block, the entry probability is $p = F(\theta_{EC} + d)$ for the relevant value difference d , and the exit probability is $F(\theta_{SV} - d)$. By part (a), the exit probability equals $1 - p$, so each block takes the form $\begin{bmatrix} 1-p & p \\ 1-p & p \end{bmatrix}$ (rank 1, identical rows). The product $\Sigma_2 V_1$ therefore yields the same value for states 1 and 2, and for states 3 and 4. Upon inspecting (13), the argument for $\Sigma_1(V_1)$ is analogous: rows 1 and 3 are identical, as are rows 2 and 4. ■

Proof of Lemma 3. Define the following value differences for firm 1:

$$\begin{aligned} v_{11} &= V_{13} - V_{11} && \text{(entry value, rival out),} \\ v_{12} &= V_{14} - V_{12} && \text{(entry value, rival in),} \\ v_{13} &= V_{12} - V_{11} && \text{(rival's entry effect, firm 1 out),} \\ v_{14} &= V_{14} - V_{13} && \text{(rival's entry effect, firm 1 in),} \end{aligned}$$

and analogously for firm 2: $v_{21} = V_{22} - V_{21}$, $v_{22} = V_{24} - V_{23}$, $v_{23} = V_{23} - V_{21}$, $v_{24} = V_{24} - V_{22}$. Note the identity $v_{i4} = (v_{i2} - v_{i1}) + v_{i3}$. In this notation, (19) states that $v_{12} - v_{11}$ and

$v_{22} - v_{21}$ both equal $\theta_C/(\rho + \eta)$. We show that $v_{i2} - v_{i1}$ is the same constant for each firm i in any equilibrium, without assuming equal arrival rates or a symmetric equilibrium. Focusing on $x_3 = (1, 0)$, where firm 1's switching action is exit ($x_3 \rightarrow x_1$) and firm 2's is entry ($x_3 \rightarrow x_4$), the instantaneous Bellman equation for firm 1 (given in (30) in Appendix B) can be written:

$$(21) \quad \rho V_{13} = \theta_M + \lambda_2 F(\underbrace{\theta_{EC}}_{-\kappa} + \underbrace{V_{24} - V_{23}}_{v_{22}}) \underbrace{(V_{14} - V_{13})}_{v_{14}} + \lambda_1 e(\underbrace{\theta_{SV}}_{\kappa} + \underbrace{V_{11} - V_{13}}_{-v_{11}}) \\ = \theta_M + \lambda_2 p_1 v_{14} + \lambda_1 e(\kappa - v_{11}),$$

where we define $p_1 = F(\theta_{EC} + v_{22})$ as firm 2's entry probability in state x_3 . Similarly, for $x_1 = (0, 0)$, we have

$$\rho V_{11} = \lambda_2 p_0 v_{13} + \lambda_1 e(-\kappa + v_{11}),$$

where $p_0 = F(\theta_{EC} + v_{21})$ is firm 2's entry probability in state x_1 . Subtracting and applying the surplus identity of Lemma 2(b), $e(\kappa - v_{11}) - e(-\kappa + v_{11}) = \kappa - v_{11}$, gives

$$(22) \quad (\rho + \lambda_1) v_{11} = \theta_M + \lambda_1 \kappa + \lambda_2 p_1 v_{14} - \lambda_2 p_0 v_{13}.$$

Similarly, using parts (a) and (b) of Lemma 2 again, subtracting the Bellman equations for states $x_4 = (1, 1)$ and $x_2 = (0, 1)$ gives:

$$(23) \quad (\rho + \lambda_1) v_{12} = \theta_M + \theta_C + \lambda_1 \kappa - \lambda_2(1 - p_1) v_{14} + \lambda_2(1 - p_0) v_{13}.$$

Subtracting (22) from (23), the p_0 and p_1 terms collapse: $-\lambda_2(1 - p_1)v_{14} - \lambda_2 p_1 v_{14} = -\lambda_2 v_{14}$ and $\lambda_2(1 - p_0)v_{13} + \lambda_2 p_0 v_{13} = \lambda_2 v_{13}$. The result is:

$$(\rho + \lambda_1)(v_{12} - v_{11}) = \theta_C - \lambda_2 v_{14} + \lambda_2 v_{13}.$$

Substituting the identity $v_{14} = (v_{12} - v_{11}) + v_{13}$, the v_{13} terms cancel:

$$(\rho + \eta)(v_{12} - v_{11}) = \theta_C.$$

The same argument applies to firm 2, swapping x_2 with x_3 and λ_1 with λ_2 , and yields the same denominator $\rho + \eta$. Defining $\Delta = \theta_C/(\rho + \eta)$, we have $v_{i2} = v_{i1} + \Delta$ for $i = 1, 2$, which is (19). ■

Proof of Theorem 2. The proof proceeds in two steps, using the value differences v_{ij} and the constant Δ defined in the proof of Lemma 3. We first use Lemma 3 to reduce the eight-dimensional system to a two-dimensional fixed-point problem. We then show that the resulting map is a contraction.

Step 1: Reduction to two dimensions. Lemma 3 reduces the four unknowns per firm to two: v_{i1} and v_{i3} (since $v_{i2} = v_{i1} + \Delta$ and $v_{i4} = \Delta + v_{i3}$). We now show that v_{i3} depends only on v_{i1} , further reducing to a single unknown per firm. The “firm 1 out” states are $x_1 = (0, 0)$ and $x_2 = (0, 1)$. From (10), their Bellman equations are

$$(\rho + \eta)V_{11} = \lambda_2[\Sigma_2(V_2) V_1]_1 + \lambda_1 E_{11} \quad \text{and} \quad (\rho + \eta)V_{12} = \lambda_2[\Sigma_2(V_2) V_1]_2 + \lambda_1 E_{12},$$

where the flow payoffs are $u_{11} = u_{12} = 0$ since firm 1 is inactive in both states. By (17), $[\Sigma_2(V_2) V_1]_1 = [\Sigma_2(V_2) V_1]_2$, so subtracting gives

$$(24) \quad (\rho + \eta)v_{13} = \lambda_1(E_{12} - E_{11}).$$

For the social surplus terms, the representation $E_{ik}(V_i) = V_{ik} + e(\psi_{ik} + V_{i,\ell(i,k)} - V_{ik})$ gives

$$E_{11} = V_{11} + e(-\kappa + v_{11}) \quad \text{and} \quad E_{12} = V_{12} + e(-\kappa + v_{12}),$$

so $E_{12} - E_{11} = v_{13} + e(-\kappa + v_{12}) - e(-\kappa + v_{11}) = v_{13} + e(-\kappa + v_{11} + \Delta) - e(-\kappa + v_{11})$, where the second step uses $v_{12} = v_{11} + \Delta$ from Lemma 3. Substituting this into (24), recalling that $\eta = \lambda_1 + \lambda_2$, and collecting v_{13} terms yields:

$$v_{13} = \phi_1(v_{11}) \equiv \frac{\lambda_1[e(-\kappa + v_{11} + \Delta) - e(-\kappa + v_{11})]}{\rho + \lambda_2}.$$

Substituting $v_{14} = \Delta + v_{13}$ and $v_{13} = \phi_1(v_{11})$ into (22):

$$\begin{aligned} (\rho + \lambda_1)v_{11} &= \theta_M + \lambda_1\kappa + \lambda_2 p_1(\Delta + \phi_1(v_{11})) - \lambda_2 p_0 \phi_1(v_{11}) \\ &= \theta_M + \lambda_1\kappa + \lambda_2 p_1 \Delta + \lambda_2(p_1 - p_0) \phi_1(v_{11}). \end{aligned}$$

For brevity, we now write $\tilde{v}_i = v_{i1}$. The entry probabilities can be expressed in terms of $\tilde{v}_2$ using $v_{22} = \tilde{v}_2 + \Delta$: $p_0(\tilde{v}_2) = F(-\kappa + \tilde{v}_2)$ and $p_1(\tilde{v}_2) = F(-\kappa + \tilde{v}_2 + \Delta)$. Substituting these in yields a fixed-point equation $\tilde{v}_1 = G_1(\tilde{v}_1, \tilde{v}_2)$ with

$$(25) \quad G_1(\tilde{v}_1, \tilde{v}_2) = \frac{1}{\rho + \lambda_1} [\theta_M + \lambda_1\kappa + \lambda_2 p_1(\tilde{v}_2)\Delta + \lambda_2(p_1(\tilde{v}_2) - p_0(\tilde{v}_2)) \phi_1(\tilde{v}_1)].$$

The analogous reduction for firm 2—swapping x_2 with x_3 and λ_1 with λ_2 , and using (18) in place of (17)—yields $\tilde{v}_2 = G_2(\tilde{v}_1, \tilde{v}_2)$. Under the equal-rate assumption $\lambda_1 = \lambda_2 = \lambda$, $\phi_1 = \phi_2 \equiv \phi$ and $G_2(\tilde{v}_1, \tilde{v}_2) = G_1(\tilde{v}_2, \tilde{v}_1)$.

Step 2: Contraction. Since f is continuous, G is C^1 on the convex domain $\mathbb{R}^2$. Therefore, if $\|J_G\|_\infty$, the maximum absolute row sum, is uniformly bounded above by a constant $\gamma < 1$, then the mean-value inequality (Ortega and Rheinboldt, 1970, 3.2.3) implies that G is a contraction with modulus γ in the ℓ^∞ norm. The Jacobian of $G = (G_1, G_2)$ has the form $J_G = \begin{bmatrix} a & b \\ b' & a' \end{bmatrix}$ where $a = \partial G_1 / \partial \tilde{v}_1$, $b = \partial G_1 / \partial \tilde{v}_2$, and the symmetry $G_1(\tilde{v}_1, \tilde{v}_2) = G_2(\tilde{v}_2, \tilde{v}_1)$

gives $a' = a(\tilde{v}_2, \tilde{v}_1)$ and $b' = b(\tilde{v}_2, \tilde{v}_1)$. Therefore both row sums satisfy the same bound, and it suffices to bound $|a| + |b|$.

From (25), G_1 depends on $\tilde{v}_1$ only through $\phi_1(\tilde{v}_1)$ and on $\tilde{v}_2$ only through $p_0(\tilde{v}_2)$ and $p_1(\tilde{v}_2)$. We will use the following three bounds:

- (i) $|p_1(\tilde{v}_2) - p_0(\tilde{v}_2)| = |F(-\kappa + \tilde{v}_2 + \Delta) - F(-\kappa + \tilde{v}_2)| \leq c_F |\Delta|$ by Lemma 4(d),
- (ii) $|\phi_1(\tilde{v}_1)| \leq \lambda_1 |\Delta| / (\rho + \lambda_2)$, since $e'(z) = F(z) \in (0, 1)$ by Lemma 4(a),
- (iii) $|\phi_1'(\tilde{v}_1)| \leq c_F \lambda_1 |\Delta| / (\rho + \lambda_2)$,

where (iii) follows by differentiating ϕ_1 and applying Lemma 4(a) to obtain $\phi_1'(\tilde{v}_1) = \frac{\lambda_1 [F(-\kappa + \tilde{v}_1 + \Delta) - F(-\kappa + \tilde{v}_1)]}{\rho + \lambda_2}$, and then applying Lemma 4(d) again.

Next, to bound a we differentiate (25) with respect to $\tilde{v}_1$ and apply (i) and (iii) to find:

$$(26) \quad a = \frac{\lambda_2(p_1 - p_0)}{\rho + \lambda_1} \phi_1'(\tilde{v}_1) \quad \text{and} \quad |a| \leq \frac{c_F^2 \lambda_1 \lambda_2 \Delta^2}{(\rho + \lambda_1)(\rho + \lambda_2)}.$$

Similarly, to bound b we differentiate (25) with respect to $\tilde{v}_2$, writing $p_0' = dp_0/d\tilde{v}_2$ and $p_1' = dp_1/d\tilde{v}_2$. By Assumptions 7 and 8, $p_0' = f(-\kappa + \tilde{v}_2) \leq c_F$, $p_1' = f(-\kappa + \tilde{v}_2 + \Delta) \leq c_F$, and f is Lipschitz with constant c_F' , so $|p_1' - p_0'| = |f(-\kappa + \tilde{v}_2 + \Delta) - f(-\kappa + \tilde{v}_2)| \leq c_F' |\Delta|$. Applying these together with (ii):

$$(27) \quad b = \frac{\lambda_2}{\rho + \lambda_1} [p_1' \Delta + (p_1' - p_0') \phi_1] \quad \text{and} \quad |b| \leq \frac{c_F \lambda_2 |\Delta|}{\rho + \lambda_1} + \frac{c_F' \lambda_1 \lambda_2 \Delta^2}{(\rho + \lambda_1)(\rho + \lambda_2)}.$$

Combining (26) and (27) with $\lambda_1 = \lambda_2 = \lambda$:

$$\begin{aligned} \|J_G\|_\infty &\leq |a| + |b| \leq \frac{c_F^2 \lambda^2 \Delta^2}{(\rho + \lambda)^2} + \frac{c_F \lambda |\Delta|}{\rho + \lambda} + \frac{c_F' \lambda^2 \Delta^2}{(\rho + \lambda)^2} \\ &= \frac{c_F \lambda |\Delta|}{\rho + \lambda} + \frac{(c_F^2 + c_F') \lambda^2 \Delta^2}{(\rho + \lambda)^2}. \end{aligned}$$

Since $\Delta = \theta_C / (\rho + 2\lambda)$, we have $\lambda \Delta / (\rho + \lambda) = \tilde{\Delta}$, and therefore

$$\|J_G\|_\infty \leq c_F |\tilde{\Delta}| + (c_F^2 + c_F') \tilde{\Delta}^2.$$

Under condition (20) of the theorem, $\|J_G\|_\infty < 1$ and so G is a contraction on $(\mathbb{R}^2, \|\cdot\|_\infty)$. By the Banach fixed-point theorem (Ortega and Rheinboldt, 1970, 5.1.3), G has a unique fixed point $(\tilde{v}_1^*, \tilde{v}_2^*)$.

We now show that this uniqueness extends to the full eight-dimensional system. Because every equilibrium induces a fixed point of G , all equilibria share the value differences $(\tilde{v}_1^*, \tilde{v}_2^*)$ and hence, by Lemma 3 and Step 1, all of the differences v_{ij} . The levels are determined as well. In the instantaneous Bellman equation (21) (see also Appendix B), the right-hand side

depends only on value differences, so with $\rho > 0$ each value V_{ik} is uniquely determined. Hence there is at most one equilibrium. Since equilibrium existence is guaranteed by Theorem 1 of [Blevins \(2026\)](#), it follows that there is exactly one equilibrium.

It remains to show that the equilibrium is symmetric. Let $\Xi: \mathbb{R}^8 \rightarrow \mathbb{R}^8$ denote the map that swaps the firm labels together with the states x_2 and x_3 :

$$\Xi V \equiv (V_{21}, V_{23}, V_{22}, V_{24}, V_{11}, V_{13}, V_{12}, V_{14}).$$

With $\lambda_1 = \lambda_2$, the Bellman system commutes with this relabeling, $T(\Xi V) = \Xi T(V)$: for each state, the Bellman equation for firm 1 becomes the equation for firm 2 after swapping x_2 with x_3 and λ_1 with λ_2 (cf. (30) and (33) in Appendix B). Therefore ΞV^* is an equilibrium whenever V^* is, and uniqueness implies $\Xi V^* = V^*$: $V_{11} = V_{21}$, $V_{12} = V_{23}$, $V_{13} = V_{22}$, and $V_{14} = V_{24}$. ■

Proof of Corollary 1. Let $g(\rho)$ denote the left-hand side of (20), holding θ_C and λ fixed. Note that g is non-increasing in ρ , so

$$\sup_{\rho > 0} g(\rho) \leq \lim_{\rho \rightarrow 0^+} g(\rho) = \frac{c_F |\theta_C|}{2\lambda} + \frac{(c_F^2 + c'_F) \theta_C^2}{4\lambda^2} < 1,$$

where the last inequality is the condition of the corollary. Therefore (20) holds for every $\rho > 0$ and the conclusion follows from Theorem 2. ■

Supporting Lemmas

Parts (a) and (b) of the following lemma are related to the two-alternative case of the Williams–Daly–Zachary theorem: the gradient of the surplus function with respect to the deterministic utility components is the vector of choice probabilities ([McFadden, 1981](#), Section 5.8, Theorem 5.1); see also [Rust \(1994, Theorem 3.1\)](#).

Lemma 4 (Error Distribution Properties). *Under Assumptions 6, 7, and 8:*

- (a) $e'(z) = F(z) \in (0, 1)$ for all $z \in \mathbb{R}$.
- (b) For all $z_1, z_2, c \in \mathbb{R}$, $L(z_1 + c, z_2 + c) = L(z_1, z_2) + c$, so that $L(z_1, z_2) = z_2 + e(z_1 - z_2)$, with $\frac{\partial L}{\partial z_1} = F(z_1 - z_2)$ and $\frac{\partial L}{\partial z_2} = 1 - F(z_1 - z_2)$.
- (c) L is Lipschitz continuous with constant 1 in the ℓ^∞ norm: $|L(z_1, z_2) - L(z'_1, z'_2)| \leq \max\{|z_1 - z'_1|, |z_2 - z'_2|\}$ for all $z_1, z_2, z'_1, z'_2 \in \mathbb{R}$.
- (d) $|F(z) - F(z')| \leq c_F |z - z'|$ for all $z, z' \in \mathbb{R}$.
- (e) $F(z) + F(-z) = 1$ for all $z \in \mathbb{R}$.

(f) $e(z) - e(-z) = z$ for all $z \in \mathbb{R}$.

Specifically, (a), (b), and (c) require only Assumption 6; (d) requires Assumptions 6 and 7; (e) and (f) require Assumptions 6 and 8.

Proof. The properties are independent of the player and state, so for brevity write $\varepsilon_0 = \varepsilon_{i0k}$ and $\varepsilon_1 = \varepsilon_{i1k}$, so that $e(z) = E_\varepsilon \max\{z + \varepsilon_1, \varepsilon_0\}$.

Part (a): For each realization of $(\varepsilon_0, \varepsilon_1)$, the integrand $\max\{z + \varepsilon_1, \varepsilon_0\}$ in the definition of e is 1-Lipschitz in z . The derivative is $1\{z + \varepsilon_1 > \varepsilon_0\}$ wherever $z + \varepsilon_1 \neq \varepsilon_0$, an event of probability one. The difference quotients are therefore bounded by one and converge almost surely, so dominated convergence (Durrett, 2019, Theorem 1.6.7) gives $e'(z) = \Pr[\varepsilon_0 - \varepsilon_1 < z] = \Pr[\varepsilon_0 - \varepsilon_1 \leq z] = F(z)$, where the second equality holds because the difference is continuously distributed. The value $F(z)$ lies in $(0, 1)$ because the difference has full support.

Part (b): Since $\max\{z_1 + c + \varepsilon_1, z_2 + c + \varepsilon_0\} = \max\{z_1 + \varepsilon_1, z_2 + \varepsilon_0\} + c$, taking expectations gives $L(z_1 + c, z_2 + c) = L(z_1, z_2) + c$. Setting $c = -z_2$ gives $L(z_1, z_2) = z_2 + L(z_1 - z_2, 0) = z_2 + e(z_1 - z_2)$. Differentiating and applying part (a) yields $\partial L / \partial z_1 = e'(z_1 - z_2) = F(z_1 - z_2)$ and $\partial L / \partial z_2 = 1 - e'(z_1 - z_2) = 1 - F(z_1 - z_2)$.

Part (c): For each realization of $(\varepsilon_0, \varepsilon_1)$, $|\max\{z_1 + \varepsilon_1, z_2 + \varepsilon_0\} - \max\{z'_1 + \varepsilon_1, z'_2 + \varepsilon_0\}| \leq \max\{|z_1 - z'_1|, |z_2 - z'_2|\}$. Taking expectations preserves the inequality.

Part (d): $|F(z) - F(z')| = |\int_{z'}^z f(s) ds| \leq c_F |z - z'|$.

Part (e): By exchangeability, $\varepsilon_0 - \varepsilon_1$ and $\varepsilon_1 - \varepsilon_0$ have the same distribution, so the difference is symmetric about zero. Since the difference is continuously distributed, $F(-z) = \Pr[\varepsilon_0 - \varepsilon_1 \leq -z] = \Pr[\varepsilon_0 - \varepsilon_1 \geq z] = 1 - F(z)$.

Part (f): Using exchangeability and $\max\{a, b\} - z = \max\{a - z, b - z\}$, we have $e(-z) = E_\varepsilon \max\{\varepsilon_1 - z, \varepsilon_0\} = E_\varepsilon \max\{\varepsilon_0 - z, \varepsilon_1\} = E_\varepsilon \max\{\varepsilon_0, \varepsilon_1 + z\} - z = e(z) - z$. ■

Lemma 5 (Logit Constants). *If Assumption 9 holds, then $F = \Lambda$, $L(z_1, z_2) = \log[\exp(z_1) + \exp(z_2)]$, and $e(z) = L(z, 0) = \log[1 + \exp(z)]$. Furthermore, Assumptions 7 and 8 are satisfied with $c_F = \sup_z \Lambda'(z) = 1/4$ and $c'_F = \sup_z |\Lambda''(z)| = \sqrt{3}/18$.*

Proof. The errors are independent across choices with a common marginal distribution, so the pair $(\varepsilon_{i0k}, \varepsilon_{i1k})$ is exchangeable. The difference of two independent type I extreme value random variables follows the standard logistic distribution, so $F = \Lambda$, and for independent mean-zero type I extreme value errors, the surplus function is $L(z_1, z_2) = \log[\exp(z_1) + \exp(z_2)]$. For the constants, write $t = \Lambda(z) \in (0, 1)$, so that $\Lambda'(z) = t(1 - t)$, which is maximized at $t = 1/2$, giving $c_F = 1/4$. Differentiating again gives $\Lambda''(z) = \Lambda'(z)(1 - 2\Lambda(z)) = t(1 - t)(1 - 2t)$. The derivative with respect to t is $6t^2 - 6t + 1$, which has roots $t = (3 \pm \sqrt{3})/6$ at which $|\Lambda''(z)| = \sqrt{3}/18$. This is the Lipschitz constant since Λ'' is continuous and bounded. ■

B Instantaneous Bellman Equations

The Bellman equations (10) and (12) can be restated in instantaneous (ρ -discounted) form by multiplying by $\rho + \eta$ and subtracting $\eta V_{ik} = (\lambda_1 + \lambda_2)V_{ik}$ from each equation. Substituting the choice probabilities $\sigma_{ik} = F(\psi_{ik} + V_{i,\ell(i,k)} - V_{ik})$ from (4) and the surplus $E_{ik}(V_i) = V_{ik} + e(\psi_{ik} + V_{i,\ell(i,k)} - V_{ik})$ yields the following system:

$$(28) \quad \rho V_{11} = \lambda_2 F(\theta_{EC} + V_{22} - V_{21})(V_{12} - V_{11}) + \lambda_1 e(\theta_{EC} + V_{13} - V_{11})$$

$$(29) \quad \rho V_{12} = \lambda_2 F(\theta_{SV} + V_{21} - V_{22})(V_{11} - V_{12}) + \lambda_1 e(\theta_{EC} + V_{14} - V_{12})$$

$$(30) \quad \rho V_{13} = \theta_M + \lambda_2 F(\theta_{EC} + V_{24} - V_{23})(V_{14} - V_{13}) + \lambda_1 e(\theta_{SV} + V_{11} - V_{13})$$

$$(31) \quad \rho V_{14} = \theta_M + \theta_C + \lambda_2 F(\theta_{SV} + V_{23} - V_{24})(V_{13} - V_{14}) + \lambda_1 e(\theta_{SV} + V_{12} - V_{14})$$

$$(32) \quad \rho V_{21} = \lambda_1 F(\theta_{EC} + V_{13} - V_{11})(V_{23} - V_{21}) + \lambda_2 e(\theta_{EC} + V_{22} - V_{21})$$

$$(33) \quad \rho V_{22} = \theta_M + \lambda_1 F(\theta_{EC} + V_{14} - V_{12})(V_{24} - V_{22}) + \lambda_2 e(\theta_{SV} + V_{21} - V_{22})$$

$$(34) \quad \rho V_{23} = \lambda_1 F(\theta_{SV} + V_{11} - V_{13})(V_{21} - V_{23}) + \lambda_2 e(\theta_{EC} + V_{24} - V_{23})$$

$$(35) \quad \rho V_{24} = \theta_M + \theta_C + \lambda_1 F(\theta_{SV} + V_{12} - V_{14})(V_{22} - V_{24}) + \lambda_2 e(\theta_{SV} + V_{23} - V_{24})$$

Each equation has up to three terms: a flow payoff (zero when the firm is inactive), a rival-transition term (weighted by λ_{-i}), and a surplus term (weighted by λ_i). The structural symmetry between the two firms is apparent: (30) for firm 1 (monopolist in $x_3 = (1, 0)$) has the same structure as (33) for firm 2 (monopolist in $x_2 = (0, 1)$), with λ_1 and λ_2 interchanged.

C Analytical Jacobian

For computational efficiency, we derive the analytical Jacobian of the equilibrium system. Throughout this appendix we impose the logit specification of Assumption 9, as maintained in the numerical work, with $F = \Lambda$ and $e(z) = \log[1 + \exp(z)]$. The equilibrium conditions $H(V) = V - T(V) = 0$ have Jacobian $J = I - DT$, where DT is the Jacobian of T . The Jacobian has block structure

$$DT = \begin{bmatrix} DT_{11} & DT_{12} \\ DT_{21} & DT_{22} \end{bmatrix}$$

where $DT_{im} = \frac{\partial T_i}{\partial V_m}$ (each block is 4×4).

We consider each block in turn. The block $DT_{11} = \frac{\partial T_1}{\partial V_1}$ contains the derivatives of firm 1's values with respect to firm 1's own values. Recalling the definition of T_1 in (10) and differentiating gives:

$$DT_{11} = \frac{1}{\rho + \eta} \left[\lambda_2 \Sigma_2(V_2) + \lambda_1 \frac{\partial E_1}{\partial V_1} \right].$$

To express this matrix concisely, we write out the switching probabilities σ_{ik} of (4) for firm i in state k in full:

$$\begin{aligned}\sigma_{11} &= \Lambda(\theta_{\text{EC}} + V_{13} - V_{11}), & \sigma_{12} &= \Lambda(\theta_{\text{EC}} + V_{14} - V_{12}), \\ \sigma_{13} &= \Lambda(\theta_{\text{SV}} + V_{11} - V_{13}), & \sigma_{14} &= \Lambda(\theta_{\text{SV}} + V_{12} - V_{14}), \\ \sigma_{21} &= \Lambda(\theta_{\text{EC}} + V_{22} - V_{21}), & \sigma_{22} &= \Lambda(\theta_{\text{SV}} + V_{21} - V_{22}), \\ \sigma_{23} &= \Lambda(\theta_{\text{EC}} + V_{24} - V_{23}), & \sigma_{24} &= \Lambda(\theta_{\text{SV}} + V_{23} - V_{24}).\end{aligned}$$

To keep the notation compact, we write $\tilde{\sigma}_{ik} \equiv 1 - \sigma_{ik}$ for the continuation probability of firm i in state x_k , so that $\sigma_{ik} + \tilde{\sigma}_{ik} = 1$.

The first term, $\lambda_2 \Sigma_2(V_2)$, comes from differentiating the matrix-vector product $\Sigma_2(V_2)V_1$ with respect to V_1 . Since Σ_2 depends only on V_2 , we have $\frac{\partial}{\partial V_1}[\Sigma_2(V_2)V_1] = \Sigma_2(V_2)$. The second term involves derivatives of the social surplus. For $E_{1k} = \log[\exp(\psi_{1k} + V_{1,\ell(1,k)}) + \exp(V_{1k})]$, where $\ell(1, k)$ is the destination state when firm 1 switches from state k , we apply Lemma 4(b) with $F = \Lambda$, $z_1 = \psi_{1k} + V_{1,\ell(1,k)}$, and $z_2 = V_{1k}$:

$$\begin{aligned}\frac{\partial E_{1k}}{\partial V_{1k}} &= 1 - \Lambda(\psi_{1k} + V_{1,\ell(1,k)} - V_{1k}) = \tilde{\sigma}_{1k}, \\ \frac{\partial E_{1k}}{\partial V_{1,\ell(1,k)}} &= \Lambda(\psi_{1k} + V_{1,\ell(1,k)} - V_{1k}) = \sigma_{1k}, \\ \frac{\partial E_{1k}}{\partial V_{1k'}} &= 0 \quad \text{for } k' \neq k, \ell(1, k).\end{aligned}$$

Thus each row of $\frac{\partial E_1}{\partial V_1}$ has exactly two nonzero entries: the continuation probability $\tilde{\sigma}_{1k}$ on the diagonal and the switching probability σ_{1k} at column $\ell(1, k)$. The complete 4×4 matrix structure is:

$$\frac{\partial E_1}{\partial V_1} = \begin{bmatrix} \tilde{\sigma}_{11} & 0 & \sigma_{11} & 0 \\ 0 & \tilde{\sigma}_{12} & 0 & \sigma_{12} \\ \sigma_{13} & 0 & \tilde{\sigma}_{13} & 0 \\ 0 & \sigma_{14} & 0 & \tilde{\sigma}_{14} \end{bmatrix}.$$

Note that by Lemma 4(b), each row sums to 1. Indeed, comparing with (13), the surplus derivative matrix is the transition matrix itself, $\partial E_1 / \partial V_1 = \Sigma_1(V_1)$, with the $F(\cdot)$ entries of (13) written here as the switching probabilities σ_{1k} .

Combining the two terms, the complete expression for DT_{11} is:

$$DT_{11} = \frac{1}{\rho + \eta} \begin{bmatrix} \lambda_2 \tilde{\sigma}_{21} + \lambda_1 \tilde{\sigma}_{11} & \lambda_2 \sigma_{21} & \lambda_1 \sigma_{11} & 0 \\ \lambda_2 \sigma_{22} & \lambda_2 \tilde{\sigma}_{22} + \lambda_1 \tilde{\sigma}_{12} & 0 & \lambda_1 \sigma_{12} \\ \lambda_1 \sigma_{13} & 0 & \lambda_2 \tilde{\sigma}_{23} + \lambda_1 \tilde{\sigma}_{13} & \lambda_2 \sigma_{23} \\ 0 & \lambda_1 \sigma_{14} & \lambda_2 \sigma_{24} & \lambda_2 \tilde{\sigma}_{24} + \lambda_1 \tilde{\sigma}_{14} \end{bmatrix}.$$

Next, we turn to the block $DT_{12} = \frac{\partial T_1}{\partial V_2}$, the derivative of firm 1's values with respect to firm 2's values. From the definition of T_1 , we can see that this involves differentiating

$\Sigma_2(V_2)V_1$ with respect to V_2 . The (k, l) element of DT_{12} is given by:

$$(DT_{12})_{kl} = \frac{\lambda_2}{\rho + \eta} \frac{\partial}{\partial V_{2l}} [(\Sigma_2 V_1)_k] = \frac{\lambda_2}{\rho + \eta} \sum_{k'=1}^4 \frac{\partial(\Sigma_2)_{kk'}}{\partial V_{2l}} V_{1k'}.$$

For each element of Σ_2 that depends on V_2 , we use $\Lambda'(z) = \Lambda(z)(1 - \Lambda(z))$. The derivatives of the choice probabilities are:

$$\frac{\partial \sigma_{2k}}{\partial V_{2l}} = \begin{cases} \sigma_{2k} \tilde{\sigma}_{2k} & \text{if } V_{2l} \text{ appears with a positive sign in } \sigma_{2k} \\ -\sigma_{2k} \tilde{\sigma}_{2k} & \text{if } V_{2l} \text{ appears with a negative sign in } \sigma_{2k} \\ 0 & \text{otherwise.} \end{cases}$$

Since rows 1 and 2 of Σ_2 depend only on (V_{21}, V_{22}) and rows 3 and 4 depend only on (V_{23}, V_{24}) , the complete 4×4 matrix is:

$$DT_{12} = \frac{\lambda_2}{\rho + \eta} \begin{bmatrix} \sigma_{21} \tilde{\sigma}_{21} (V_{11} - V_{12}) & \sigma_{21} \tilde{\sigma}_{21} (V_{12} - V_{11}) & 0 & 0 \\ \sigma_{22} \tilde{\sigma}_{22} (V_{11} - V_{12}) & \sigma_{22} \tilde{\sigma}_{22} (V_{12} - V_{11}) & 0 & 0 \\ 0 & 0 & \sigma_{23} \tilde{\sigma}_{23} (V_{13} - V_{14}) & \sigma_{23} \tilde{\sigma}_{23} (V_{14} - V_{13}) \\ 0 & 0 & \sigma_{24} \tilde{\sigma}_{24} (V_{13} - V_{14}) & \sigma_{24} \tilde{\sigma}_{24} (V_{14} - V_{13}) \end{bmatrix}.$$

Each entry is weighted by firm 2's switching probability sensitivity $\sigma_{2k} \tilde{\sigma}_{2k}$ times the difference in firm 1's values that appears in row k of Σ_2 .

By symmetry with DT_{12} , the block $DT_{21} = \frac{\partial T_2}{\partial V_1}$ has (k, l) element:

$$(DT_{21})_{kl} = \frac{\lambda_1}{\rho + \eta} \sum_{k'=1}^4 \frac{\partial(\Sigma_1)_{kk'}}{\partial V_{1l}} V_{2k'}.$$

Since rows 1 and 3 of Σ_1 depend only on (V_{11}, V_{13}) and rows 2 and 4 depend only on (V_{12}, V_{14}) , the complete 4×4 matrix is:

$$DT_{21} = \frac{\lambda_1}{\rho + \eta} \begin{bmatrix} \sigma_{11} \tilde{\sigma}_{11} (V_{21} - V_{23}) & 0 & \sigma_{11} \tilde{\sigma}_{11} (V_{23} - V_{21}) & 0 \\ 0 & \sigma_{12} \tilde{\sigma}_{12} (V_{22} - V_{24}) & 0 & \sigma_{12} \tilde{\sigma}_{12} (V_{24} - V_{22}) \\ \sigma_{13} \tilde{\sigma}_{13} (V_{21} - V_{23}) & 0 & \sigma_{13} \tilde{\sigma}_{13} (V_{23} - V_{21}) & 0 \\ 0 & \sigma_{14} \tilde{\sigma}_{14} (V_{22} - V_{24}) & 0 & \sigma_{14} \tilde{\sigma}_{14} (V_{24} - V_{22}) \end{bmatrix}.$$

Each entry is weighted by firm 1's switching probability sensitivity $\sigma_{1k} \tilde{\sigma}_{1k}$ times the difference in firm 2's values that appears in row k of Σ_1 .

For the block $DT_{22} = \frac{\partial T_2}{\partial V_2}$, from $T_2 = \frac{1}{\rho + \eta} [u_2 + \lambda_1 \Sigma_1(V_1)V_2 + \lambda_2 E_2(V_2)]$, by symmetry with DT_{11} :

$$DT_{22} = \frac{1}{\rho + \eta} \left[\lambda_1 \Sigma_1(V_1) + \lambda_2 \frac{\partial E_2}{\partial V_2} \right]$$

where the social surplus derivative matrix for firm 2 is:

$$\frac{\partial E_2}{\partial V_2} = \begin{bmatrix} \tilde{\sigma}_{21} & \sigma_{21} & 0 & 0 \\ \sigma_{22} & \tilde{\sigma}_{22} & 0 & 0 \\ 0 & 0 & \tilde{\sigma}_{23} & \sigma_{23} \\ 0 & 0 & \sigma_{24} & \tilde{\sigma}_{24} \end{bmatrix}.$$

Combining the two terms, the complete expression for DT_{22} is:

$$DT_{22} = \frac{1}{\rho + \eta} \begin{bmatrix} \lambda_1 \tilde{\sigma}_{11} + \lambda_2 \tilde{\sigma}_{21} & \lambda_2 \sigma_{21} & \lambda_1 \sigma_{11} & 0 \\ \lambda_2 \sigma_{22} & \lambda_1 \tilde{\sigma}_{12} + \lambda_2 \tilde{\sigma}_{22} & 0 & \lambda_1 \sigma_{12} \\ \lambda_1 \sigma_{13} & 0 & \lambda_1 \tilde{\sigma}_{13} + \lambda_2 \tilde{\sigma}_{23} & \lambda_2 \sigma_{23} \\ 0 & \lambda_1 \sigma_{14} & \lambda_2 \sigma_{24} & \lambda_1 \tilde{\sigma}_{14} + \lambda_2 \tilde{\sigma}_{24} \end{bmatrix}.$$

Since $\partial E_2 / \partial V_2 = \Sigma_2(V_2)$ as well, comparing with (11), the two diagonal blocks coincide: $DT_{22} = \frac{1}{\rho + \eta} [\lambda_1 \Sigma_1(V_1) + \lambda_2 \Sigma_2(V_2)] = DT_{11}$.

References

- Abbring, J. H. and J. R. Campbell (2010). Last-in first-out oligopoly dynamics. *Econometrica* 78, 1491–1527.
- Abbring, J. H., J. R. Campbell, J. Tilly, and N. Yang (2018). Very simple Markov-perfect industry dynamics: Theory. *Econometrica* 86, 721–735.
- Aguirregabiria, V. (2012). A method for implementing counterfactual experiments in models with multiple equilibria. *Economics Letters* 114, 190–194.
- Aguirregabiria, V., A. Collard-Wexler, and S. P. Ryan (2021). Dynamic games in empirical industrial organization. In K. Ho, A. Hortaçsu, and A. Lizzeri (Eds.), *Handbook of Industrial Organization*, Volume 4, Chapter 4, pp. 225–343. Elsevier.
- Aguirregabiria, V. and M. Marcoux (2021). Imposing equilibrium restrictions in the estimation of dynamic discrete games. *Quantitative Economics* 12(4), 1223–1271.
- Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica* 70, 1519–1543.
- Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53.
- Aguirregabiria, V. and J. Suzuki (2014). Identification and counterfactuals in dynamic models of market entry and exit. *Quantitative Marketing and Economics* 12, 267–304.
- Aradillas-López, A. (2020). The econometrics of static games. *Annual Review of Economics* 12(1), 135–165.
- Arcidiacono, P., P. Bayer, J. R. Blevins, and P. B. Ellickson (2016). Estimation of dynamic discrete choice models in continuous time with an application to retail competition. *The Review of Economic Studies* 83, 889–931.
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370.

- Berry, S. T. (1992). Estimation of a model of entry in the airline industry. *Econometrica* 60, 889–917.
- Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010). Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica* 78, 453–508.
- Blevins, J. R. (2026). Identification and estimation of continuous-time dynamic discrete choice games. *Quantitative Economics* 17, 254–296.
- Blevins, J. R. and M. Kim (2024). Nested pseudo likelihood estimation of continuous-time dynamic discrete games. *Journal of Econometrics* 238, 105576.
- Borkovsky, R. N., U. Doraszelski, and Y. Kryukov (2010). A user’s guide to solving dynamic stochastic games using the homotopy method. *Operations Research* 58, 1116–1132.
- Borkovsky, R. N., U. Doraszelski, and Y. Kryukov (2012). A dynamic quality ladder model with entry and exit: Exploring the equilibrium correspondence using the homotopy method. *Quantitative Marketing and Economics* 10, 197–229.
- Bresnahan, T. F. and P. C. Reiss (1991). Empirical models of discrete games. *Journal of Econometrics* 48, 57–81.
- Bugni, F. A. and J. Bunting (2021). On the iterated estimation of dynamic discrete choice games. *The Review of Economic Studies* 88, 1031–1073.
- Cabral, L. M. B. and M. H. Riordan (1994). The learning curve, market dominance, and predatory pricing. *Econometrica* 62, 1115–1140.
- Ciliberto, F. and E. Tamer (2009). Market structure and multiple equilibria in the airline industry. *Econometrica* 77, 1791–1828.
- Curtat, L. O. (1996). Markov equilibria of stochastic games with complementarities. *Games and Economic Behavior* 17, 177–199.
- de Paula, A. (2013). Econometric analysis of games with multiple equilibria. *Annual Review of Economics* 5, 107–131.
- Dearing, A. and J. R. Blevins (2025). Efficient and convergent sequential pseudo-likelihood estimation of dynamic discrete games. *The Review of Economic Studies* 92, 981–1021.
- Doraszelski, U. and J. F. Escobar (2010). A theory of regular Markov perfect equilibria in dynamic stochastic games: Genericity, stability, and purification. *Theoretical Economics* 5, 369–402.

- Doraszelski, U. and K. L. Judd (2012). Avoiding the curse of dimensionality in dynamic stochastic games. *Quantitative Economics* 3, 53–93.
- Doraszelski, U. and M. Satterthwaite (2010). Computable Markov-perfect industry dynamics. *The RAND Journal of Economics* 41, 215–243.
- Durrett, R. (2019). *Probability: Theory and Examples* (Fifth ed.). Cambridge University Press.
- Egedal, M., Z. Lai, and C.-L. Su (2015). Estimating dynamic discrete-choice games of incomplete information. *Quantitative Economics* 6, 567–597.
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *The Review of Economic Studies* 62, 53–82.
- Espín-Sánchez, J.-A., A. Parra, and Y. Wang (2023). Equilibrium uniqueness in entry games with private information. *The RAND Journal of Economics* 54, 512–540.
- Haller, H. and R. Lagunoff (2000). Genericity and Markovian behavior in stochastic games. *Econometrica* 68, 1231–1248.
- Herings, P. J.-J. and R. J. A. P. Peeters (2004). Stationary equilibria in stochastic games: Structure, selection, and computation. *Journal of Economic Theory* 118, 32–60.
- Horst, U. (2005). Stationary equilibria in discounted stochastic games with weakly interacting players. *Games and Economic Behavior* 51(1), 83–108.
- Hotz, V. J. and R. A. Miller (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies* 60, 497–529.
- Iskhakov, F., J. Rust, and B. Schjerning (2016). Recursive lexicographical search: Finding all Markov perfect equilibria of finite state directional dynamic games. *The Review of Economic Studies* 83, 658–703.
- Jia, P. (2008). What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica* 76, 1263–1316.
- Jun, S. J. and J. Pinkse (2020). Counterfactual prediction in complete information games: Point prediction under partial identification. *Journal of Econometrics* 216, 394–429.
- Kasahara, H. and K. Shimotsu (2008). Pseudo-likelihood estimation and bootstrap inference for structural discrete Markov decision models. *Journal of Econometrics* 146(1), 92–106.

- Kasahara, H. and K. Shimotsu (2012). Sequential estimation of structural models with a fixed point constraint. *Econometrica* 80, 2303–2319.
- Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica* 70, 801–816.
- McFadden, D. L. (1978). Modelling the choice of residential location. In A. Karlqvist, L. Lundqvist, F. Snickars, and J. W. Weibull (Eds.), *Spatial Interaction Theory and Planning Models*, Volume 3 of *Studies in Regional Science and Urban Economics*, pp. 75–96. Amsterdam: North-Holland.
- McFadden, D. L. (1981). Econometric models of probabilistic choice. In C. F. Manski and D. L. McFadden (Eds.), *Structural Analysis of Discrete Data with Econometric Applications*, pp. 198–272. Cambridge, MA: MIT Press.
- McKelvey, R. D. and A. McLennan (1996). Computation of equilibria in finite games. In H. M. Amman, D. A. Kendrick, and J. Rust (Eds.), *Handbook of Computational Economics*, Volume 1, pp. 87–142. Amsterdam: Elsevier.
- Milgrom, P. R. and R. J. Weber (1985). Distributional strategies for games with incomplete information. *Mathematics of Operations Research* 10, 619–632.
- Ortega, J. and W. Rheinboldt (1970). *Iterative Solution of Nonlinear Equations in Several Variables*. Academic Press.
- Otsu, T. and M. Pesendorfer (2023). Equilibrium multiplicity in dynamic games: Testing and estimation. *The Econometrics Journal* 26, C26–C42.
- Otsu, T., M. Pesendorfer, and Y. Takahashi (2016). Pooling data across markets in dynamic Markov games. *Quantitative Economics* 7, 523–559.
- Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies* 75, 901–928.
- Pesendorfer, M. and P. Schmidt-Dengler (2010). Sequential estimation of dynamic discrete games: A comment. *Econometrica* 78, 833–842.
- Puterman, M. L. (2005). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley-Interscience.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55, 999–1013.

- Rust, J. (1994). Structural estimation of Markov decision processes. In R. F. Engle and D. L. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, pp. 3081–3143. Amsterdam: North-Holland.
- Shapley, L. S. (1953). Stochastic games. *Proceedings of the National Academy of Sciences* 39, 1095–1100.
- Su, C.-L. (2014). Estimating discrete-choice games of incomplete information: Simple static examples. *Quantitative Marketing and Economics* 12(2), 167–207.
- Su, C.-L. and K. L. Judd (2012). Constrained optimization approaches to estimation of structural models. *Econometrica* 80, 2213–2230.
- Sweeting, A. (2009). The strategic timing incentives of commercial radio stations: An empirical analysis using multiple equilibria. *The RAND Journal of Economics* 40, 710–742.
- Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies* 70, 147–165.