Identification and Estimation of Dynamic Games with Unknown Information Structure*

Konan Hara† Yuki Ito‡ Paul S. Koh§

October, 2025

Abstract

We develop an empirical framework for analyzing dynamic games when the underlying information structure is unknown to the analyst. We introduce *Markov correlated equilibrium*, a dynamic analog of Bayes correlated equilibrium, and show that its predictions coincide with the Markov perfect equilibrium predictions attainable when players observe richer signals than the analyst assumes. We provide tractable methods for informationally robust estimation, inference, and counterfactual analysis. We illustrate the framework with a dynamic entry game between Starbucks and Dunkin' in the US and study the role of informational assumptions.

Keywords: Dynamic games, Markov, correlated equilibrium, information, partial identification

*We are grateful to Matthew Backus, Gautam Gowrisankaran, Ben Handel, Hidehiko Ichimura, Kei Kawai, Ashley Langer, Sokbae Lee, Qingmin Liu, Juan Pantano, Bernard Salanié, Matthijs Wildenbeest, Tiemen Woutersen, Takuro Yamashita, and seminar participants at UArizona, UC Berkeley, Columbia University, and Midwest econometrics group conference. Lorenzo Magnolfi graciously shared his code in the early stage of this project. We merged our two independent working papers previously circulated as "Estimating Dynamic Games with Unknown Information Structure" by Koh and "Estimation of Dynamic Games With Weak Assumptions on Payoff Type Information" by Hara and Ito. A portion of this paper was completed during Koh's tenure at the U.S. Federal Trade Commission. The views expressed in this article are those of the authors and do not necessarily reflect those of the Federal Trade Commission or any individual Commissioner. All errors are ours.

†Department of Economics, Michigan State University. Email: harakon1@msu.edu

‡Department of Economics, Indiana University, Bloomington. Email: yukiito@iu.edu

§School of Economics, Yonsei University. Email: paulkoh9@gmail.com

1 Introduction

This paper develops a computationally tractable econometric framework for estimating empirical dynamic discrete game models with weak assumptions on players' information. Let

$$f : (\Theta, \mathcal{S}) \rightrightarrows \mathcal{D} \tag{1}$$

describe the data-generating process under Markov perfect equilibrium, where Θ is the set of structural parameters, $\mathcal{S}$ is the set of information structures, and $\mathcal{D}$ is the data; the correspondence in (1) reflects the possibility of multiple equilibria. The standard identification argument relies on inverting (1) under a specific "incomplete information" assumption—namely that players observe their own payoff shocks but not those of their rivals—which simplifies the estimation problem by making it closely parallel to that of single-agent models (Rust, 1994; Aguirregabiria and Mira, 2007; Bajari, Benkard and Levin, 2007; Pakes, Ostrovsky and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008; Bugni and Bunting, 2021; Dearing and Blevins, 2025).[1]

While this assumption has enabled computationally tractable estimation of dynamic games, researchers often have little evidence to justify it. Unlike observable outcomes such as prices, sales, or entry decisions, information sets are inherently unobservable and may reflect complex mixtures of private knowledge, noisy signals, and shared information transmitted through suppliers, trade associations, or industry channels. Firms may also differ in their ability to gather market intelligence: some obtain precise signals about profitability, while others rely on noisy information. In addition, the timing and granularity of what firms know about their rivals can vary widely.

A growing body of research shows that assuming the information structure is *known*, while misspecifying players' information, can generate substantial bias in both static (Grieco, 2014; Pakes et al., 2015; Syrgkanis, Tamer and Ziani, 2021; Magnolfi and Roncoroni, 2023; Koh, 2023; Gualdani and Sinha, 2024) and dynamic settings (Fershtman and Pakes, 2012; Aguirregabiria and Magesan, 2020; Asker, Fershtman, Jeon and Pakes, 2020). The concern is particularly acute in the empirical literature on dynamic games, where most works adopt a specific assumption that limits the

[1] See Aguirregabiria, Collard-Wexler and Ryan (2021) for a recent survey of the literature.

role of players' private information in strategic interactions. With the exception of Ericson and Pakes (1995), which assumes complete information, we are not aware of empirical approaches that allow for alternative information structures. Achieving informationally robust identification by directly recovering the model primitives from equation (1) is, however, infeasible: the set of possible information structures is vast and infinite-dimensional, and the task is further complicated by the inherent challenges of estimating dynamic models of strategic interaction.

In this paper, we develop a tractable empirical framework for analyzing dynamic discrete games while leaving the within-stage information structure *unknown* to the analyst. We assume that the researcher knows only the minimal information available to players but remains agnostic about the possibility that the players may have access to additional signals.[2] The framework typically delivers a set of partially identified parameters that are robust to misspecification of the within-stage information structure. It can be used to construct confidence sets for model parameters that are robust to informational assumptions, thereby assessing the sensitivity of parameter estimates to the researcher's assumptions about players' information.

The main innovation of this paper is the development of *Markov correlated equilibrium*, a dynamic (Markovian) analog of Bayes correlated equilibrium (Bergemann and Morris, 2016). We show that Markov correlated equilibrium can capture the set of Markov perfect equilibrium predictions (the joint distribution of actions and states) that can arise when the players might observe more signals than assumed by the analyst. We establish the informational robustness property of Markov correlated equilibrium by formulating Markov perfect equilibrium as the Bayes Nash equilibrium of the associated reduced one-shot game via the one-shot deviation principle and invoking Theorem 1 of Bergemann and Morris (2016).

We build on the theory of Markov correlated equilibrium to propose computationally tractable approaches for estimation and inference. Roughly speaking, we prove

$$\Theta_I^{MCE} = \bigcup_{\tilde{S} \in \mathcal{S}} \Theta_I^{MPE}\left(\tilde{S}\right), \tag{2}$$

where Θ_I^{MCE} is the identified set of parameters after assuming that a Markov corre-

[2]This assumption restricts the set of possible information structure $\mathcal{S}$ in (1).

lated equilibrium generated the data, and $\Theta_I^{MPE}(\tilde{S})$ is the Markov perfect equilibrium identified set when the information structure is assumed to be $\tilde{S}$. Directly inverting (1) requires constructing the right-hand side of (2), which is computationally infeasible when $\mathcal{S}$ is large. However, we show that constructing the left-hand side of (2) is computationally tractable. In sum, treating the data *as if* they were generated by a Markov correlated equilibrium enables an informationally robust econometric analysis. The identified set (2) is usually non-singleton. We propose a computationally attractive strategy for estimation and inference by formulating the estimation problem as mathematical programs with equilibrium constraints (MPEC) (Su and Judd, 2012) and applying the inference strategy of Horowitz and Lee (2023) by nesting the confidence set for the conditional choice probabilities into the MPEC problem.[3]

We apply our framework to study the dynamic entry and exit decisions of the two largest coffee chains in the US, Starbucks and Dunkin'. As a benchmark, we first estimate the structural parameters under Markov perfect equilibrium, assuming that each player observes their own idiosyncratic payoff shock but not the opponent's, which corresponds to the standard "incomplete information" assumption in the literature. We then employ our Markov correlated equilibrium framework to estimate an identified set that captures all possible parameters that could generate the data under Markov perfect equilibrium when the players may additionally observe information about the opponent's payoff shock. We also estimate an identified set that allows for the possibility that the players may not even observe their payoff shock; in this case, $\mathcal{S}$ spans the entire set of information structures, and the identified set is fully robust to any potential misspecification of the information structure.

Comparison between the estimates from the standard incomplete information Markov perfect equilibrium assumption and our informationally robust Markov correlated equilibrium assumption suggests that the standard incomplete information assumption may be too strong. The estimates from the standard approach suggest that the presence of an opponent has a positive spillover effect on profitability. On the other hand, the Markov correlated equilibrium identified sets do not rule out both positive and negative spillover effects.

[3]The standard estimation strategies for dynamic games, such as the nested fixed-point algorithm of Rust (1987) or the conditional choice probability inversion of Hotz and Miller (1993), are not applicable.

We conduct a counterfactual experiment in which we reduce the firms' entry costs by 10% and simulate the evolution of the number of active firms, averaged over all markets. The possible outcomes range from 0 (no firms active in any market) to 2 (both firms active in all markets). To illustrate the range of counterfactual predictions under our framework, we consider two extreme scenarios. The low (resp. high) scenario considers the expected number of active firms in equilibrium using the lower (resp. upper) bounds of the 95% CIs of the competitive effects parameters, while holding all other parameters at their point estimates.

The standard Markov perfect equilibrium approach predicts that reducing entry costs would raise the average number of active firms from the baseline of 0.37 to between 0.55 (low) and 0.58 (high) by the end of the sample period. Allowing for the possibility that firms observe not only their own but also their rival's payoff shock expands the range to 0.35–0.87 under our Markov correlated equilibrium framework. The range becomes even wider, 0.26–1.32, once we incorporate the uncertainty of the remaining structural parameters, and it grows further if we allow for the possibility that firms may not observe (or may only partially observe) their own payoff shocks.

These results show that our framework can generate policy implications that are both robust to informational assumptions and informative, though the degree of informativeness depends on how flexibly one models the information structure and accounts for parameter uncertainty. At the same time, the robustness provided by Markov correlated equilibrium may admit too many possibilities to be empirically useful, a limitation also noted by Koh (2023), Magnolfi and Roncoroni (2023), and Gualdani and Sinha (2024) in their application of Bayes correlated equilibrium.

Related Literature

Our work advances the literature on the empirical analysis of game-theoretic models in several ways. Our work is most closely related to the econometric frameworks that pursue robustness to misspecification on information structures (Grieco, 2014; Pakes et al., 2015; Syrgkanis, Tamer and Ziani, 2021; Koh, 2023; Magnolfi and Roncoroni, 2023; Canen and Song, 2024; Gualdani and Sinha, 2024; Han, Kaido and Magnolfi, 2024). Building on this literature, we develop a novel framework for estimating a class of dynamic games. More broadly, our work contributes to the literature on

tractable econometric methods for partially identified game-theoretic models under weak assumptions on equilibrium selection (Tamer, 2003; Ciliberto and Tamer, 2009; Beresteanu et al., 2011; Galichon and Henry, 2011).

We contribute to the dynamic games literature by developing a new solution concept with attractive computational properties. Markov correlated equilibrium extends Bayes correlated equilibrium of Bergemann and Morris (2013, 2016) to dynamic Markovian games in the tradition of Maskin and Tirole (2001). While related notions of correlated equilibrium in dynamic games have been studied (Nowak and Raghavan, 1992; Duffie et al., 1994; Solan, 2001; Solan and Vieille, 2002; Mailath and Samuelson, 2006; Von Stengel and Forges, 2008; He and Sun, 2017; Forges, 2020), our formulation is distinct in extending Bayes correlated equilibrium to dynamic stochastic Markov games with incomplete information. Importantly, our solution concept is tailored to the class of dynamic games that serve as the workhorse for empirical studies of dynamic competition (Aguirregabiria et al., 2021).

Our framework maintains tractability by focusing on robustness to within-stage information assumptions. Although one could allow for across-stage incomplete information and dynamic belief updating (e.g., Doval and Ely, 2020; Makris and Renou, 2023), such extensions are typically infeasible in empirical work. Persistent information asymmetries render estimation intractable (Fershtman and Pakes, 2012), and broadening the solution concept would further enlarge the identified set, reducing its empirical informativeness. Our results already demonstrate that Markov correlated equilibrium admits a wide range of outcomes, which can significantly weaken empirical content, as observed by Koh (2023), Magnolfi and Roncoroni (2023), and Gualdani and Sinha (2024) in static settings. For this reason, we restrict attention to the informational assumptions most relevant to existing empirical frameworks, rather than pursuing broader forms of robustness.

Finally, this paper contributes to the empirical literature on dynamic competition and spatial expansion of retail chains (Aguirregabiria and Mira, 2007; Jia, 2008; Holmes, 2011; Ellickson, Houghton and Timmins, 2013; Nishida, 2015; Aguirregabiria and Suzuki, 2016; Arcidiacono, Bayer, Blevins and Ellickson, 2016; Yang, 2020; Beresteanu, Ellickson, Misra and Reeder III, 2024). We examine a novel application: the dynamic expansion of the two largest coffee chains in the US, Starbucks

and Dunkin'. The study most closely related to ours is Fang and Yang (2024), who analyze entry-deterrence motives of the four largest coffee chains in Toronto, Canada, using a dynamic oligopoly framework.

The rest of the paper is organized as follows. In Section 2, we define Markov correlated equilibrium and establish a connection between Markov correlated equilibria and Markov perfect equilibria in a class of dynamic games and discuss its properties. In Section 3, we introduce an econometric model and discuss informationally robust identification. Section 4 proposes strategies for computing the identified set. Section 5 discusses our approach to inference. Section 6 applies our framework to the dynamic entry/exit game by Starbucks and Dunkin' in the US. Section 7 concludes. All proofs are in Appendix A.

2 Theory of Markov Correlated Equilibrium

We consider a class of dynamic Markovian games that has been used as a standard framework for empirical works. We introduce the concept of Markov correlated equilibrium for informationally robust analysis and discuss its properties.

2.1 Setup

Let $t = 1, 2, ..., \infty$ denote discrete time. A *stationary dynamic Markov game of incomplete information* is defined by a pair of basic game G and information structure S. A basic game $G = \langle \mathcal{I}, (\mathcal{A}_i, u_i)_{i \in \mathcal{I}}, \mathcal{X}, \mathcal{E}, \psi, f, \delta \rangle$ specifies the payoff-relevant primitives of the model: $i \in \mathcal{I} = \{1, 2, ..., I\}$ indexes the players; $a_{it} \in \mathcal{A}_i$ denotes player i's action at time t, where $\mathcal{A}_i$ is a finite set of actions available to player i; $a_t = (a_{1t}, ..., a_{It}) \in \mathcal{A} \equiv \times_{i \in \mathcal{I}} \mathcal{A}_i$ denotes an action profile; $x_t \in \mathcal{X}$ denotes a state variable that is publicly observed by the players, and $\mathcal{X}$ is assumed to be finite; $\varepsilon_t \in \mathcal{E}$ denotes a latent state variable that is not directly observed by the players, and $\mathcal{E}$ is assumed to be finite[4]; $\psi_{\varepsilon_t | x_t}$ represents the players' prior on the probability of ε_t conditional on x_t; $f_{x_{t+1} | a_t, x_t, \varepsilon_t}$ specifies the probability of transitioning to state x_{t+1}

[4]The assumption of a finite state space is non-substantive: it simplifies notation and matches the discretization used in estimation. The theory readily extends to a more general state space; see Syrgkanis, Tamer and Ziani (2021) and Magnolfi and Roncoroni (2023) for the case of Bayes correlated equilibrium.

conditional on $(a_t, x_t, \varepsilon_t)$; $u_i : \mathcal{A} \times \mathcal{X} \times \mathcal{E} \to \mathbb{R}$ is the flow payoff function of player i; $\delta \in [0, 1)$ is the players' discount factor.

An information structure $S = ((\mathcal{T}_i)_{i \in \mathcal{I}}, \pi)$ specifies the information-relevant primitives; $\tau_{it} \in \mathcal{T}_i$ denotes player i's private signal at period t, where $\mathcal{T}_i$ represents a finite set of private signals; $\tau_t = (\tau_{1t}, ..., \tau_{It}) \in \mathcal{T} \equiv \times_{i \in \mathcal{I}} \mathcal{T}_i$ is a signal profile; $\pi : \mathcal{X} \times \mathcal{E} \to \Delta(\mathcal{T})$ is a signal distribution that maps the state of the world to a signal profile. Note that players' signals are allowed to be correlated. The interpretation is that each player i does not directly observe the latent state ε_t but receives a private signal τ_{it} whose informativeness about ε_t depends on π.

The game (G, S) is common knowledge to the players. The model proceeds as follows. At the beginning of period t, the common knowledge state $x_t \in \mathcal{X}$ is given and publicly observed. Conditional on x_t, the latent state $\varepsilon_t \in \mathcal{E}$ is drawn from the probability distribution $\psi_{\varepsilon_t | x_t}$. Next, a profile of private signals $\tau_t = (\tau_{1t}, ..., \tau_{It}) \in \mathcal{T}$ is drawn from the signal distribution $\pi_{\tau_t | x_t, \varepsilon_t}$. At this point, each player i observes (x_t, τ_{it}) and uses Bayes' rule to infer ε_t. Then, the players simultaneously choose actions a_{it}, $i = 1, ..., I$, and each player i receives period t payoff $u_i(a_t, x_t, \varepsilon_t) \in \mathbb{R}$. Finally, the observable state x_t transitions to x_{t+1} via the probability kernel $f_{x_{t+1} | a_t, x_t, \varepsilon_t}$. Given x_{t+1}, period $t + 1$ begins.

Note that the primitives imply that the transition probability of state variables factors as $\Pr(x_{t+1}, \varepsilon_{t+1} | a_t, x_t, \varepsilon_t) = \psi_{\varepsilon_{t+1} | x_{t+1}} f_{x_{t+1} | a_t, x_t, \varepsilon_t}$. Assuming that the latent variables are generated independently of the previous-period states is standard in the empirical literature.[5] The assumption also plays a vital role in this paper as it simplifies the characterization of correlated equilibria in dynamic settings.

The players are rational and forward-looking. In period t, each player i chooses action plans to maximize the expected discounted sum of intertemporal payoffs

$$\mathbb{E}\left[\sum_{s=t}^{\infty} \delta^{s-t} u_i(a_s, x_s, \varepsilon_s) | x_t, \tau_{it} \right].$$

The conditioning on (x_t, τ_{it}) reflects the assumption that i observes the public state x_t and her signal τ_{it} before choosing period t action. Since the environment is stationary,

[5] When latent variables are correlated over time, empirical analysis quickly becomes intractable, as computation of perfect Bayesian equilibria requires tracking the full history of players' information and updating beliefs at every contingency (Fershtman and Pakes, 2012).

we suppress time subscripts unless necessary.

Example 1 (Two-player dynamic entry game)**.** As a running example, we consider the two-player dynamic entry/exit model of Pesendorfer and Schmidt-Dengler (2008), which has served as a representative example in the literature (Egesdal, Lai and Su, 2015; Aguirregabiria et al., 2021; Dearing and Blevins, 2025). We also adopt a similar model in our empirical application. There are two players, $i = 1, 2$. In each period, firms simultaneously decide whether to be active ($a_{it} = 1$) or inactive ($a_{it} = 0$). The per-period payoff is

$$u_i(a_{it}, a_{jt}, z_{it}, \varepsilon_{it}) = \begin{cases} (1 - a_{jt})\pi^m + a_{jt}\pi^d + (1 - z_{it})c + \varepsilon_{it} & \text{if } a_{it} = 1 \\ z_{it}\kappa & \text{if } a_{it} = 0 \end{cases}. \tag{3}$$

Here, π^m, π^d, c, and κ represent monopoly profit, duopoly profit, entry cost, and scrap value, respectively. The public state variable is $x_t = (z_{1t}, z_{2t})$, where z_{it} represents whether firm i is an incumbent ($z_{it} = 1$) or not ($z_{it} = 0$). The firms' previous actions determine the incumbency states as $z_{it} = a_{i,t-1}$. In empirical applications, x_t usually contains additional observable market characteristics. The latent state variable is $\varepsilon_t = (\varepsilon_{1t}, \varepsilon_{2t})$, where $\varepsilon_{it} \in \mathcal{E}_i \subseteq \mathbb{R}$ only enters player i's payoff. Each ε_{it} independently follows the standard normal distribution.

Standard empirical dynamic discrete game models assume that each player observes ε_{it} but not ε_{jt} for $j \neq i$. Formally, $\mathcal{T}_i = \mathcal{E}_i$, and the private signal is $\tau_{it} = \varepsilon_{it}$ with probability one. Alternatively, if the information is complete, players publicly observe ε_t, so $\mathcal{T}_i = \mathcal{E}$, and $\tau_{it} = \varepsilon_t$ with probability one.

2.2 Markov Perfect Equilibrium

A Markov strategy of player i is a mapping $\beta_i : \mathcal{X} \times \mathcal{T}_i \to \Delta(\mathcal{A}_i)$ that specifies a probability distribution over actions at each realization of observable state and private signal.[6] A Markov strategy profile $\beta = (\beta_1, ..., \beta_I)$ is a *Markov perfect equilibrium* if

[6]While the standard Markov strategy assumption only allows players to condition on payoff-relevant state (Maskin and Tirole, 2001), our definition allows players' strategy to depend on signals that may be unrelated to the fundamentals (e.g., a public correlation device under complete information), which is a consequence of considering arbitrary information structures. However, our assumptions align with the spirit of Maskin and Tirole (2001) because we rule out arbitrary strate-

it is a subgame perfect equilibrium.[7]

To leverage the results in static games to a dynamic environment, it is useful to characterize the Markov perfect equilibrium conditions using the one-shot deviation principle.[8] The one-shot deviation principle formulation has two advantages. First, we can leverage results established in static environments as it translates a Markov perfect equilibrium as a Bayes-Nash equilibrium of the associated normal-form game. Second, it facilitates equilibrium computation during econometric analysis.

Suppose a strategy profile $\beta = (\beta_1, ..., \beta_I)$ is given. Define the *ex-ante value function* $V_i^\beta(x) \in \mathbb{R}$ as the expected payoff to player i at state x prior to the realization of ε when all players follow the prescriptions in β. For each $i \in \mathcal{I}$, $V_i^\beta \in \mathbb{R}^{|\mathcal{X}|}$ is the unique solution to

$$V_i^\beta(x) = \sum_{\varepsilon \in \mathcal{E}, \tau \in \mathcal{T}, a \in \mathcal{A}} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \beta_{a|x,\tau} \left\{ u_i(a, x, \varepsilon) + \delta \sum_{x' \in \mathcal{X}} V_i^\beta(x') f_{x'|a,x,\varepsilon} \right\}, \quad \forall x \in \mathcal{X} \quad (4)$$

where $\beta_{a|x,\tau} \equiv \prod_{i=1}^I \beta_i(a_i|x, \tau_i)$ denotes the conditional distribution over action profiles induced by the strategy profile β. Next, define the *outcome-specific payoff function* given strategy profile β as

$$v_i^\beta(a, x, \varepsilon) \equiv u_i(a, x, \varepsilon) + \delta \sum_{x' \in \mathcal{X}} V_i^\beta(x') f_{x'|a,x,\varepsilon},$$

which represents the continuation payoff to player i when (a, x, ε) is realized today, and all players follow the prescriptions in β from tomorrow onward.

A strategy profile β in game (G, S) induces a reduced normal-form basic game

gies by requiring actions to depend only on signals of a memoryless latent state. Moreover, we are unaware of a treatment for the construction of the maximally coarse partition of the action space in an incomplete information setup.

[7] This definition is standard in the literature. Beliefs are determined by the signal distribution π, conditional on publicly observed states. With public transitions and common knowledge of probabilities, beliefs need not enter the equilibrium definition. Subgames are well-defined under Markov strategies.

[8] According to the one-shot deviation principle, a strategy profile is subgame perfect precisely when no player can gain from deviating at a single information set. For infinite-horizon games, this principle holds under the "continuity at infinity" condition, which ensures that far-off events matter little (Fudenberg and Tirole, 1991). Dynamic discrete game models satisfy this condition because payoffs are discounted sums of uniformly bounded expected stage payoffs, even if shocks themselves have unbounded support.

$G^\beta = \langle \mathcal{I}, (\mathcal{A}_i, v_i^\beta)_{i\in\mathcal{I}}, \mathcal{X}, \mathcal{E}, \psi \rangle$. Then (G^β, S) describes a "static" game in which player i's "static" payoff function is given by $v_i^\beta : \mathcal{A} \times \mathcal{X} \times \mathcal{E} \to \mathbb{R}$. The following lemma allows us to import the results in static environments into dynamic environments.

Lemma 1 (One-shot deviation principle formulation of Markov perfect equilibrium)**.** *A strategy profile β is a Markov perfect equilibrium of (G, S) if and only if β is a Bayes Nash equilibrium of (G^β, S).*

2.3 Markov Correlated Equilibrium

We introduce a dynamic analog of Bayes correlated equilibrium for dynamic discrete games. A *stationary Markov decision rule* in (G, S) is a mapping

$$\sigma : \mathcal{X} \times \mathcal{E} \times \mathcal{T} \to \Delta(\mathcal{A}),$$

which specifies a probability distribution over action profiles at each realization of state and signals. It is instructive to think of σ as a recommendation strategy of an omniscient mediator who observes (x, ε, τ) and privately recommends action to each player. Suppose the mediator commits to a recommendation strategy σ and announces it to the players. After the state and players' signal (x, ε, τ) are realized, an action profile $a = (a_1, ..., a_I)$ is drawn from the probability distribution $\sigma_{a|x,\varepsilon,\tau}$, and each a_i is privately recommended to each player i. Each player i, having observed (x, τ_i, a_i), decides whether to obey (play a_i) or not (deviate to $a_i' \neq a_i$). If the decision rule σ is such that the players are always obedient, we call σ a Markov correlated equilibrium of (G, S).

We formalize the equilibrium condition as follows. Let $V_i^\sigma \in \mathbb{R}^{|\mathcal{X}|}$ denote the vector of ex-ante value functions induced by σ, obtained as a unique solution to

$$V_i^\sigma(x) = \sum_{\varepsilon\in\mathcal{E},\tau\in\mathcal{T},a\in\mathcal{A}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{a|x,\varepsilon,\tau}\left\{u_i(a, x, \varepsilon) + \delta \sum_{x'\in\mathcal{X}} V_i^\sigma(x') f_{x'|a,x,\varepsilon}\right\}, \quad \forall x \in \mathcal{X}. \tag{5}$$

Define the outcome-specific value function associated with σ as

$$v_i^\sigma(a, x, \varepsilon) \equiv u_i(a, x, \varepsilon) + \delta \sum_{x'\in\mathcal{X}} V_i^\sigma(x') f_{x'|a,x,\varepsilon},$$

which represents the payoff to player i if (a, x, ε) is realized today and σ determines the players' actions in the future. Let $\mathbb{E}^\sigma[v_i^\sigma(a_i', a_{-i}, x, \varepsilon)|x, \tau_i, a_i]$ be the expected payoff to player i from choosing a_i' when i observes (x, τ_i) and receives recommendation a_i. The following definition states that σ is a Markov correlated equilibrium of (G, S) if the players do not deviate from the mediator's recommendations.

Definition 1. A decision rule σ is a *Markov correlated equilibrium* of (G, S) if for each $i \in \mathcal{I}$, $x \in \mathcal{X}$, $\tau_i \in \mathcal{T}_i$, and $a_i \in \mathcal{A}_i$, we have

$$\mathbb{E}^\sigma[v_i^\sigma(a_i, a_{-i}, x, \varepsilon)|x, \tau_i, a_i] \geq \mathbb{E}^\sigma[v_i^\sigma(a_i', a_{-i}, x, \varepsilon)|x, \tau_i, a_i] \tag{6}$$

for each $a_i' \neq a_i$ whenever $\Pr^\sigma(x, \tau_i, a_i) > 0$.

Since

$$\begin{aligned}\mathbb{E}^\sigma[v_i^\sigma(\tilde{a}_i, a_{-i}, x, \varepsilon)|x, \tau_i, a_i] &= \sum_{\varepsilon, a_{-i}} v_i^\sigma(\tilde{a}_i, a_{-i}, x, \varepsilon)\Pr^\sigma(\varepsilon, a_{-i}|x, \tau_i, a_i) \\ &= \sum_{\varepsilon, a_{-i}} v_i^\sigma(\tilde{a}_i, a_{-i}, x, \varepsilon)\left(\frac{\sum_{\tau_{-i}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{(a_i,a_{-i})|x,\varepsilon,\tau}}{\sum_{\bar{\varepsilon},\tau_{-i},\bar{a}_{-i}} \psi_{\bar{\varepsilon}|x}\pi_{\tau|x,\bar{\varepsilon}}\sigma_{(a_i,\bar{a}_{-i})|x,\bar{\varepsilon},\tau}}\right),\end{aligned}$$

canceling out the denominator—which is constant across all possible realizations of $(\varepsilon, \tau_{-i}, a_{-i})$—lets us rewrite the obedience condition (6) as, for each i,

$$\sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{a|x,\varepsilon,\tau}v_i^\sigma(a, x, \varepsilon) \geq \sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{a|x,\varepsilon,\tau}v_i^\sigma(a_i', a_{-i}, x, \varepsilon), \quad \forall x, \tau_i, a_i, a_i'. \tag{7}$$

Equation (7) serves as the operational definition of Markov correlated equilibrium in our econometric analysis.

A Markov correlated equilibrium always exists for any signal distribution π. Nowak and Raghavan (1992) establishes the existence of a stationary correlated equilibrium in a discounted stochastic game with complete information under standard regularity conditions. We will show in the following sections that the set of Markov correlated equilibria under arbitrary information structure always includes a complete information stationary correlated equilibrium.[9] Thus, the set of Markov correlated equilibria is always non-empty.

[9]Existence results for Markov perfect equilibrium under incomplete information do not neces-

2.4 Relationship to Bayes Correlated Equilibrium

As in Section 2.2, a decision rule σ in game (G, S) induces a reduced normal-form basic game $G^\sigma = \langle \mathcal{I}, (\mathcal{A}_i, v_i^\sigma)_{i\in\mathcal{I}}, \mathcal{X}, \mathcal{E}, \psi\rangle$. Then (7) shows that if σ is a Markov correlated equilibrium of (G, S), it is a Bayes correlated equilibrium of (G^σ, S).

Lemma 2 (One-shot deviation principle characterization of Markov correlated equilibrium)**.** *A decision rule σ is a Markov correlated equilibrium of (G, S) if and only if it is a Bayes correlated equilibrium of (G^σ, S).*

When players fully discount the future ($\delta = 0$), a Markov correlated equilibrium collapses to a Bayes correlated equilibrium of the stage game: for each i,

$$\sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{a|x,\varepsilon,\tau}u_i(a,x,\varepsilon) \geq \sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{a|x,\varepsilon,\tau}u_i(a_i',a_{-i},x,\varepsilon), \quad \forall x,\tau_i,a_i,a_i'.$$

Thus, Markov correlated equilibrium is a proper dynamic analog of Bayes correlated equilibrium.

Unfortunately, contrary to its static analog, the Markov correlated equilibrium conditions are not linear with respect to the decision rule, which creates two challenges that are absent in the static case. First, Markov correlated equilibrium is more difficult to compute. While computing a Bayes correlated equilibrium amounts to solving a linear program, computing a Markov correlated equilibrium generally requires solving a non-convex program. Second, the set of Bayes correlated equilibria is convex, but the set of Markov correlated equilibria is generally non-convex. The convexity of the set of Bayes correlated equilibria implies that a mixture of multiple equilibria is observationally equivalent to a single equilibrium; Syrgkanis, Tamer and Ziani (2021) and Magnolfi and Roncoroni (2023) use this property to stay agnostic about the true equilibrium selection rule in the data generating process. However, since Markov correlated equilibrium does not share this property, we assume that a single equilibrium is selected.

sarily imply the existence of a Markov correlated equilibrium for certain information structures. For this reason, we rely on the result of Nowak and Raghavan (1992).

2.5 Informational Robustness

We establish the informational robustness property of Markov correlated equilibrium. Let an equilibrium prediction refer to the joint distribution on actions, states, and signals that are implied by either an equilibrium strategy profile or decision rule. We prove the following claim.

Claim 1 (Markov correlated equilibrium is informationally robust)**.** The set of Markov correlated equilibrium predictions is equal to the set of Markov perfect equilibrium predictions that can arise when players can observe more signals that are unknown to the analyst.

In the empirical application, we consider a scenario where the analyst knows that each firm observes its firm-specific payoff shock but cannot rule out the possibility that they receive some signals about opponents' shocks. Claim 1 says that the analyst can capture the set of all predictions that can arise in such a scenario by resorting to Markov correlated equilibrium. Thus, as in Bayes correlated equilibrium, Markov correlated equilibrium attains a form of informational robustness that allows the analyst to be agnostic about the players observing additional signals after specifying players' minimal information. We formalize Claim 1 by following Bergemann and Morris (2016)'s framework and connecting Markov correlated equilibrium to Bayes correlated equilibrium via the one-shot deviation principle.

First, we introduce a partial order on the set of information structures using the notion of *expansion* as in Bergemann and Morris (2016). Let $\omega \in \Omega \equiv \mathcal{X} \times \mathcal{E}$ for the sake of exposition.

Definition 2 (Expansion)**.** Let $S = (\mathcal{T}, \pi)$ be an information structure. Information structure $S^* = (\mathcal{T}^*, \pi^*)$ is an *expansion* of S, or $S^* \gtrsim_E S$, if there exists $(\tilde{\mathcal{T}}_i)_{i \in \mathcal{I}}$ and $\lambda : \Omega \times \mathcal{T} \to \Delta(\tilde{\mathcal{T}})$ such that $\mathcal{T}_i^* = \mathcal{T}_i \times \tilde{\mathcal{T}}_i$ for all $i \in \mathcal{I}$ and $\pi^*_{\tau, \tilde{\tau}|\omega} = \pi_{\tau|\omega} \lambda_{\tilde{\tau}|\omega,\tau}$.

Intuitively, if $S^* \gtrsim_E S$, the players receive *extra* signals in S^* than they do in S. Specifically, complete information (players observe the latent state ε almost surely) is an expansion of an arbitrary information structure.

Next, we claim that if a strategy profile induces the same joint distribution over states and actions as a decision rule, the two induce the same reduced normal-form

games. Let S^* be an expansion of S. A strategy profile β in (G, S^*) *induces* a decision rule σ for (G, S) if

$$\sigma_{a|x,\varepsilon,\tau} = \sum_{\tilde{\tau}\in\tilde{\mathcal{T}}} \lambda_{\tilde{\tau}|x,\varepsilon,\tau} \prod_{i=1}^{I} \beta_i(a_i|x, \tau_i, \tilde{\tau}_i), \quad \forall a, x, \varepsilon, \tau.$$

The strategy profile β in a game with more signals induces the same joint distribution over actions, states, and signals as the decision rule σ in a game with fewer signals.

Lemma 3 (Identical reduced normal-form basic games)**.** *If β induces σ, then the associated reduced normal-form basic games, G^β and G^σ, are identical.*

Lemma 3 is intuitive: if β induces σ, then β and σ induce identical distributions over action profiles at each state, so the associated ex-ante value functions must be identical to each other, i.e., $V_i^\beta = V_i^\sigma$. Then, the outcome-specific payoff functions are also identical to each other, i.e., $v_i^\beta = v_i^\sigma$, making all primitives of the basic games identical to each other.

Finally, a solution concept generates a *prediction* defined as a probability distribution over actions at each state and signal. Let $\mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S)$ represent the set of predictions that can be induced by a Markov perfect equilibrium in (G, S).[10] Let $\mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S)$ be defined similarly for a Markov correlated equilibrium.

We are ready to state our main theorem. Suppose the analyst knows that the players observe signals specified by S but does not know whether the players have access to extra signals. What is the set of feasible Markov perfect equilibrium predictions? The following theorem says that the set of Markov correlated equilibria of (G, S) captures all Markov perfect equilibrium predictions when the players can observe more signals than in S.

Theorem 1 (Markov correlated equilibrium is informationally robust)**.** *For any basic game G and information structure S,*

$$\mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S) = \bigcup_{S^* \succsim_E S} \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^*).$$

[10] We slightly abuse the notation and use $\mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S)$ to denote a set of conditional distributions over $\mathcal{A}$ defined at each (x, ε, τ).

Theorem 1 is attractive because $\mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G,S)$ is far easier to characterize and compute than $\bigcup_{S^* \succsim_E S} \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}$; the former does not require searching over the set of information structures.[11] Note that when $\delta = 0$, Theorem 1 reduces to

$$\mathcal{P}^{BCE}_{a|x,\varepsilon,\tau}(G,S) = \bigcup_{S^* \succsim_E S} \mathcal{P}^{BNE}_{a|x,\varepsilon,\tau}(G,S^*),$$

which is analogous to Bergemann and Morris (2016) Theorem 1.

In empirical works, it is common to assume that the analyst observes the *conditional choice probabilities* that represent the probability of each action profile at each state that is observable to the econometrician. Thus, it is useful to characterize the implications of Theorem 1 in terms of conditional choice probabilities. Let $\mathcal{P}^{MPE}_{a|x}(G,S)$ denote the set of feasible conditional choice probabilities that can arise under a Markov perfect equilibrium of (G,S). Let $\mathcal{P}^{MCE}_{a|x}(G,S)$ be defined similarly.

Corollary 1. *For any basic game G and information structure S,*

$$\mathcal{P}^{MCE}_{a|x}(G,S) = \bigcup_{S^* \succsim_E S} \mathcal{P}^{MPE}_{a|x}(G,S^*).$$

In Section 3, we discuss how results from Theorem 1 and Corollary 1 can be applied to derive results on the identification of the model parameters.

As in the case of Bayes correlated equilibrium, if players observe more signals, the set of Markov correlated equilibria becomes smaller. In other words, more information implies a smaller set of Markov correlated equilibria.

Theorem 2. *Let G be an arbitrary basic game. If $S^1 \succsim_E S^2$, then $\mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G,S^1) \subseteq \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G,S^2)$.*

2.6 Numerical Example

We continue Example 1 to illustrate the Markov correlated equilibrium identified set. The true parameters in (3) are set to $(\pi^m, \pi^d, c, \kappa) = (1.2, -1.2, -0.2, 0.1)$, and we use

[11] Makris and Renou (2023) obtain a related result in the context of multistage games. However, since their focus is on finite-stage games and they allow for varying information structures *across* stages, our result is not a direct corollary of theirs. Moreover, applying their framework in empirical settings would generally be infeasible, as it requires keeping track of the entire history.

the discount factor $\delta = 0.9$. We generate data by finding a Markov perfect equilibrium under the standard assumption that ε_i is private information to player i. Pesendorfer and Schmidt-Dengler (2008) finds five equilibria, two of which are symmetric to other equilibria. We report the conditional choice probabilities of equilibria (i), (ii), and (iii) of Pesendorfer and Schmidt-Dengler (2008) in Table 1.[12]

Table 1: Probability of entry in three Markov perfect equilibria

	(i)		(ii)		(iii)	
	P1	P2	P1	P2	P1	P2
Out/Out	0.733	0.276	0.615	0.528	0.576	0.576
Out/In	0.613	0.420	0.312	0.840	0.305	0.842
In/Out	0.800	0.223	0.831	0.303	0.842	0.305
In/In	0.752	0.294	0.606	0.578	0.595	0.595

Note: Each value represents the entry probability of each player at each state. P1 and P2 represent player 1 and player 2.

To characterize the bounds on the game outcomes that can arise in a Markov correlated equilibrium, we compute the Markov correlated equilibria that maximize and minimize the expected number of active firms.[13] Using the associated conditional choice probabilities, we simulate the number of active firms over $T = 6$ periods, assuming that the initial state has no incumbent, i.e., $x_1 = (0, 0)$.

Figure 1 plots the simulated average number of active firms for $T = 6$ periods. Note that the outcomes of Pesendorfer and Schmidt-Dengler (2008)'s three Markov perfect equilibria are bounded by those of the Markov correlated equilibria. The Markov correlated equilibrium bounds characterize the sharp bounds on the Markov perfect equilibrium outcomes that can arise when players can receive information about the opponent's payoff shock.

[12] We use Dearing and Blevins (2025)'s MATLAB code to generate the equilibrium conditional choice probabilities.

[13] The procedure is similar to the counterfactual equilibrium calculations in the empirical application detailed in Online Appendix A.2.2.

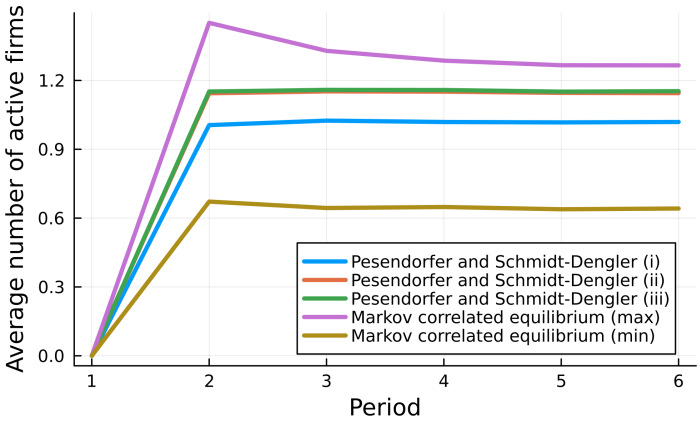


Figure 1: Average number of active firms over time

3 Informationally Robust Identification

In this section, we specialize the model for econometric analysis and discuss informationally robust identification in dynamic discrete games using Markov correlated equilibrium.

3.1 Setup

To facilitate econometric analyses, we impose the following assumptions on the basic game primitives. As commonly assumed in empirical works, we let $\varepsilon \equiv (\varepsilon_1, \varepsilon_2, ..., \varepsilon_I)$ where $\varepsilon_i \in \mathcal{E}_i$ only enters player i's payoff and $f_{x'|a,x,\varepsilon} = f_{x'|a,x}$, i.e., the transition probability of public state x is independent of latent state ε. We assume that the payoff functions and the prior distribution are parameterized by a finite-dimensional vector θ so that $u_i = u_i^\theta$ and $\psi = \psi^\theta$.

The econometrician observes data

$$\{a_{m,t}, x_{m,t} : \ m = 1, 2, ..., M, \ t = 1, 2, ..., T\},$$

which have players' actions $a_{m,t}$ and common knowledge state variables $x_{m,t}$ across M independent markets over T periods. We assume that M is large and that the conditional choice probabilities $\phi : \mathcal{X} \to \Delta(\mathcal{A})$ and the transitional probability function $f : \mathcal{A} \times \mathcal{X} \to \Delta(\mathcal{X})$ can be non-parametrically identified from the data. Finally, we assume that the common discount factor $\delta \in [0, 1)$ is known to the econometrician.

In sum, we treat the basic game $G^\theta = \langle \mathcal{I}, (\mathcal{A}_i, u_i^\theta)_{i\in\mathcal{I}}, \mathcal{X}, \mathcal{E}, \psi^\theta, f, \delta \rangle$ as known to the econometrician up to θ.

We summarize the baseline assumptions for econometric analysis as follows.

Assumption 1 (Baseline assumptions for econometric analysis)**.**

1. *The set of covariates $\mathcal{X}$ and the set of latent states $\mathcal{E}$ are finite.*

2. *The prior distribution ψ^θ and the payoff functions u_i^θ are known up to a finite-dimensional parameter θ.*

3. *The latent state is a vector of player-specific payoff shocks, i.e., $\varepsilon = (\varepsilon_1, ..., \varepsilon_I)$ and $u_i^\theta(a, x, \varepsilon) = u_i^\theta(a, x, \varepsilon_i)$, and does not affect the transition probability of public states, i.e., $f_{x'|a,x,\varepsilon} = f_{x'|a,x}$.*

4. *The conditional choice probabilities $\phi : \mathcal{X} \to \Delta(\mathcal{A})$ and the transition probability function $f : \mathcal{A} \times \mathcal{X} \to \Delta(\mathcal{A})$ are identified from the data.*

The only "unconventional" assumption is that $\mathcal{E}$ is finite. Although unnecessary for the identification arguments, the discrete support assumption is necessary for feasible estimation and makes the connection to the previous section straightforward. Other econometric papers that use Bayes correlated equilibrium also discretize the state space for feasible estimation (Syrgkanis, Tamer and Ziani, 2021; Magnolfi and Roncoroni, 2023; Gualdani and Sinha, 2024). The player-specific payoff shock assumption, which is a conventional assumption, makes the role of the latent state clear, but it is also not required for our identification arguments.

Given a game (G^θ, S), a solution concept determines a set of feasible predictions, which in turn induces a set of feasible conditional choice probabilities. A standard econometric approach is to assume that the information structure S is known and identify the model parameters by inverting the mapping from the model parameters to the observed conditional choice probability vector ϕ.

Definition 3 (Identified set)**.** Given Assumption 1, a solution concept SC, and an information structure S, the identified set of parameters is defined as:

$$\Theta_I^{SC}(S) \equiv \left\{\theta \in \Theta : \ \phi \in \mathcal{P}_{a|x}^{SC}(G^\theta, S)\right\}.$$

In words, a candidate parameter θ enters the identified set if the observed conditional choice probabilities ϕ can be rationalized by some equilibrium (defined by the solution concept) of the model.

3.2 Informationally Robust Identified Set

We consider a scenario where a Markov perfect equilibrium generates the data, but the true information structure is unknown to the analyst; misspecifying the information structure leads to biased parameter estimates. We assume that the analyst only knows that the players minimally observe signals in S but may observe more. For instance, as assumed in Magnolfi and Roncoroni (2023), the analyst might know that each player i observes at least ε_i but does not know whether they can receive extra information about opponents' payoff shocks.

Assumption 2 (Identification under Markov perfect equilibrium)**.** *The data are generated by a Markov perfect equilibrium of* (G^{θ_0}, S^0)*, where* S^0 *is an expansion of* S*.*

It is well-understood that working directly with Assumption 2 is computationally infeasible; when the analyst knows S but does not know the true information structure S^0, the analyst has to consider the set of all information structures that are expansions of S, but the set is too large. We can make the identification problem computationally feasible by replacing Assumption 2 with the following.

Assumption 3 (Identification under Markov correlated equilibrium)**.** *The data are generated by a Markov correlated equilibrium of* $\left(G^{\theta_0}, S\right)$*.*

Theorem 3 (Equivalence of identified sets)**.** *The identified set under Assumptions 1 and 2 is equal to the identified set under Assumptions 1 and 3.*

Theorem 3 says that an econometrician who does not know the true information structure underlying the data-generating process can proceed by treating the data *as if* they were generated by a Markov correlated equilibrium. Syrgkanis, Tamer and Ziani (2021), Magnolfi and Roncoroni (2023), Koh (2023), and Gualdani and Sinha (2024) use similar results in static settings. Theorem 3 extends the results to a class of dynamic games. Note that a Markov correlated equilibrium identified set always includes the parameters that can be obtained under the complete information

Markov perfect equilibrium assumption with mixed strategies (Ericson and Pakes, 1995; Doraszelski and Satterthwaite, 2010).

3.3 Properties of Informationally Robust Identified Set

Identified Set Shrinks with Stronger Informational Assumptions The Markov correlated equilibrium identified set depends on the baseline information structure set by the analyst, which specifies what players minimally observe. Theorem 2 has shown that a stronger assumption on players' information leads to a tighter set of predictions. The following translates this intuition to identified sets.

Theorem 4. *If* $S^1 \gtrsim_E S^2$, *then* $\Theta_I^{MCE}(S^1) \subseteq \Theta_I^{MCE}(S^2)$.

Theorem 4 implies a trade-off between informational robustness and the tightness of the identified set. Using weaker assumptions on information is more robust but increases the size of the identified set. The econometrician obtains the tightest identified set when players' baseline information is complete; the econometrician obtains the loosest identified set when players minimally observe nothing about ε.

Markov Correlated Equilibrium is Not Robust to Multiple Equilibria Issue In the static case, assuming that a single Bayes correlated equilibrium generates the data is innocuous because any mixture of Bayes correlated equilibria is a Bayes correlated equilibrium. However, this property does not carry over to the dynamic case because a mixture of Markov correlated equilibria need not be a Markov correlated equilibrium. This occurs because the Bayes correlated equilibrium obedience condition is linear in the decision rule while the Markov correlated equilibrium obedience condition is nonlinear (non-convex) in the decision rule. Thus, although we remain agnostic about the equilibrium selection rule, we assume that the data are generated by a single equilibrium. Identification with data generated by multiple equilibria is possible (e.g., by using Beresteanu, Molchanov and Molinari (2011)), but far more computationally challenging.[14]

[14] Our framework already incorporates public randomization that is independent across stages. However, because of the relationship between stage-game payoffs and value functions induced by the Markov structure, public randomization does not enable us to convexify the objective set of distributions.

Excluded Variables with Large Support Ensure Point-Identification While the tightness of the identified set is largely an empirical question, the identified set can shrink substantially when excluded variables have large support. In a scenario where players know their own payoff shock, Magnolfi and Roncoroni (2023) prove that players' payoff parameters are point-identified under Bayes correlated equilibrium when there are player-specific excluded variables with large support.[15] Consider, for example, a two-player entry game. Driving a time-persistent variable that enters only player j's payoff to either positive or negative infinity ensures that the player always stays in or stays out, rendering the other player's decision problem a single-agent binary choice problem. The econometrician can use the single-agent data to identify the payoff parameters, including the competitive effects parameters (Tamer, 2003). The same intuition holds in the dynamic case: in markets where one player always stays in or out, the other player solves a single-agent dynamic discrete choice problem while taking the other player's action as fixed. The players' payoff parameters can be identified and estimated using the single-agent dynamic discrete choice data if our interest lies in the case where players know their own payoff shock.[16]

3.4 Numerical Example

Following Pesendorfer and Schmidt-Dengler (2008), we assume that the econometrician knows the values of the scrap value $\kappa = 0.1$ and the discount factor $\delta = 0.9$. However, instead of assuming that the researcher knows the information structure exactly, we assume that the researcher only knows that i minimally observes ε_i and leaves open the possibility that i receives extra signals about ε_{-i}.

Table 2 reports the projection intervals of the Markov correlated equilibrium identified sets obtained under each equilibrium. In all cases, the Markov correlated equilibrium identified set contains the true parameter vector, as expected. However, the size of the identified set can vary depending on which equilibrium generated data. We find that the Markov correlated equilibrium identified set obtained from equilibrium (i)—which is the only nested pseudo-likelihood-stable equilibrium and generates

[15] Parameter governing the correlation of players' payoff types may not be point-identified.

[16] If we are interested in the case where players may not have complete knowledge about their own shock, we will need to extend Theorem 1 of Bergemann and Morris (2016) to a single-agent dynamic discrete choice setup.

highly asymmetric choice probabilities across players—is the tightest.[17] The identified sets under equilibria (ii) and (iii) are indistinguishable because the equilibrium conditional choice probabilities are similar to each other.

Table 2: Markov correlated equilibrium identified sets

	True	(i)	(ii)	(iii)
π^m	1.2	$[0.93, 1.83]$	$[0.88, 2.52]$	$[0.88, 2.52]$
π^d	-1.2	$[-1.87, -1.01]$	$[-3.16, -0.65]$	$[-3.16, -0.65]$
c	-0.2	$[-0.51, 0.37]$	$[-0.67, 1.02]$	$[-0.67, 1.02]$

Note: The table reports the projections of the Markov correlated equilibrium identified set assuming that the data were generated from Pesendorfer and Schmidt-Dengler (2008)'s Markov perfect equilibrium (i), (ii), or (iii).

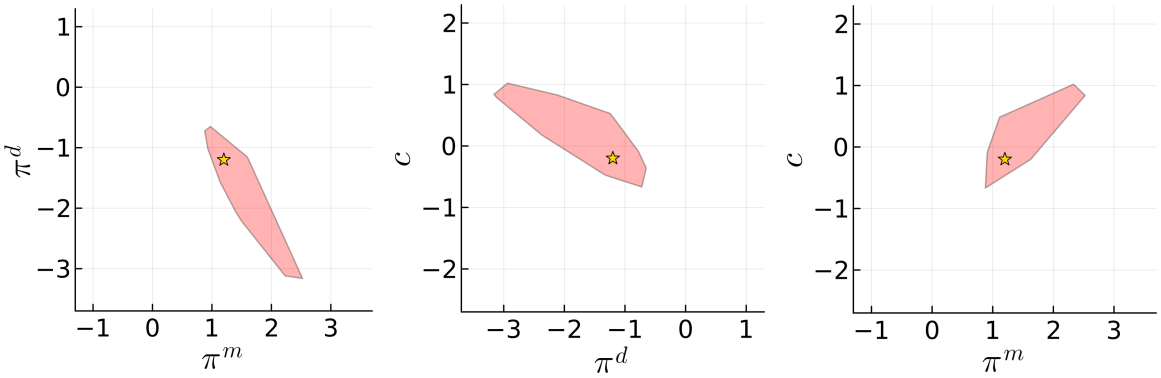


Figure 2: Convex hull of projected identified set

Figure 2 visualizes the results by plotting the convex hull of the sharp identified set corresponding to equilibrium (ii). In each subfigure, the shaded area represents the convex hull of the projected identified sets, and the star represents the location of the true parameters. As expected, the identified set is non-singleton and contains the true parameter.

Our numerical example shows that a strong assumption on information is far from innocuous: removing arbitrary assumptions on who observes what may result

[17]The nested pseudo-likelihood-stability has been the subject of investigation in the dynamic games estimation literature (see Dearing and Blevins (2025) for a review). Equilibrium (i) is nested pseudo-likelihood-stable while equilibria (ii) and (iii) are not. Our Markov correlated equilibrium estimator is not sensitive to the nested pseudo-likelihood-stability of the underlying equilibrium.

in a significant loss of identifying power. Robustness to misspecification on players' information structure expands the identified set, but the degree to which the identified set expands can vary depending on which equilibrium is selected in the data generating process.

4 Estimation

We propose computationally feasible estimation strategies by formulating a mathematical program with equilibrium constraints à la Su and Judd (2012). We assume the econometrician knows the true conditional choice probability vector ϕ (we relax this assumption in the next section). We also assume that the baseline information structure S is fixed and denote the sharp identified set as $\Theta_I \equiv \Theta_I^{MCE}(S)$.

For a candidate parameter θ, we have $\theta \in \Theta_I$ if and only if there exists a decision rule σ that satisfies the Markov correlated equilibrium obedience conditions and induces the observed conditional choice probabilities. We characterize the Markov correlated equilibrium as a set of constraints that the econometrician can plug into optimization software.

Theorem 5 (Sharp Identified Set)**.** *Suppose Assumption 1 holds. Then $\theta \in \Theta_I$ if and only if there exists a pair (σ, V) that satisfies*

$$\sigma_{a|x,\varepsilon,\tau} \geq 0 \text{ for each } a \text{ and } \sum_a \sigma_{a|x,\varepsilon,\tau} = 1, \qquad \forall x, \varepsilon, \tau \tag{8}$$

$$\phi_{a|x} = \sum_{\varepsilon,\tau} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau}, \qquad \forall a, x \tag{9}$$

$$\sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \partial v_i^\theta(a_i', a, x, \varepsilon_i) \leq 0, \qquad \forall i, x, \tau_i, a_i, a_i' \tag{10}$$

$$V_{i,x} = \sum_{\varepsilon,\tau,a} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} u_i^\theta(a, x, \varepsilon_i) + \delta \sum_{a,x'} \phi_{a|x} V_{i,x'} f_{x'|a,x}, \qquad \forall i, x \tag{11}$$

where $v_i^\theta(a, x, \varepsilon_i) \equiv u_i^\theta(a, x, \varepsilon_i) + \delta \sum_{x'} V_{i,x'} f_{x'|a,x}$ is the outcome-specific value function, and $\partial v_i^\theta(a_i', a, x, \varepsilon_i) \equiv v_i^\theta(a_i', a_{-i}, x, \varepsilon_i) - v_i^\theta(a_i, a_{-i}, x, \varepsilon_i)$ represents the deviation payoff from a_i to a_i'.

Constraints (8) mean that σ is a proper conditional probability distribution, as required by the definition of a decision rule; (9) requires that σ induces the observed

conditional choice probabilities ϕ; (10) is the Markov correlated equilibrium obedience condition; (11) is the ex-ante value function equation described in (5).[18] Note that ψ, π, ϕ, δ, and f are treated as known objects in the mathematical program.[19]

Fully Robust Identified Set In the special case where the analyst sets the baseline information structure to players having no minimal signals, S^{null}, the corresponding identified set—"fully robust identified set"—allows for the underlying information structure to be completely arbitrary. It turns out that determining whether a candidate parameter enters the fully robust identified set is computationally easier because the associated program is linear.

Theorem 6 (Fully Robust Identified Set)**.** *Suppose 1 holds.* $\theta \in \Theta_I^{MCE}(S^{null})$ *if and only if there exists* (σ, V) *that satisfies the following linear constraints*

$$\sigma_{a|x,\varepsilon} \geq 0 \text{ for each } a \text{ and } \sum_a \sigma_{a|x,\varepsilon} = 1, \qquad \forall x, \varepsilon$$

$$\phi_{a|x} = \sum_{\varepsilon} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon}, \qquad \forall a, x \tag{12}$$

$$\sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} \partial u_i^\theta(a_i', a, x, \varepsilon_i) + \delta \sum_{a_{-i}, x'} \phi_{a|x} V_{i,x'} \partial f_{x'|a_i', a, x} \leq 0, \qquad \forall i, x, a_i, a_i' \tag{13}$$

$$V_{i,x} = \sum_{\varepsilon, a} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} u_i^\theta(a, x, \varepsilon_i) + \delta \sum_{a, x'} \phi_{a|x} V_{i,x'} f_{x'|a,x}, \qquad \forall i, x \tag{14}$$

where $\partial u_i^\theta(a_i', a, x, \varepsilon_i) \equiv u_i^\theta(a_i', a_{-i}, x, \varepsilon_i) - u_i^\theta(a, x, \varepsilon_i)$ *represents the deviation flow payoff from* a_i *to* a_i'*, and* $\partial f_{x'|a_i', a, x} \equiv f_{x'|a_i', a_{-i}, x} - f_{x'|a,x}$ *is the change of the transition probability from* a_i *to* a_i'*.*

The expressions above reflect the assumption that the players receive no signals about the latent state ε. They only rely on their common prior to update their beliefs after observing private recommendations by the information designer. Since

[18] Note (11) also allows us to express each V_i as a closed-form function of σ. This is similar to expressing the ex-ante value function as a closed-form function of conditional choice probabilities in Aguirregabiria and Mira (2007).

[19] These conditions bear some resemblance to those in Abreu et al. (1990) and Athey et al. (2004). Their approach characterizes the set of continuation payoffs using such conditions, raising the question of whether it could also apply in our setting. However, the Markov structure in our framework imposes a tight restriction between stage-game and continuation payoffs, which renders their approach inapplicable.

$\Theta_I^{MCE}(S) \subseteq \Theta_I^{MCE}(S^{null})$ for any information structure S by Theorem 4, the analyst can exploit computational tractability by first estimating the fully robust identified set and then analyzing cases with minimal information specified, as the search area can be restricted. However, the fully robust identified set may be quite large and empirically uninformative (c.f., Magnolfi and Roncoroni, 2023; Koh, 2023).

5 Inference

The inference problem in this paper is characterized by a moment inequality model with a large number of moment inequalities and a high-dimensional nuisance parameter, making it difficult to apply many inference methods for partially identified models.[20] Inference methods used by Syrgkanis, Tamer and Ziani (2021), Magnolfi and Roncoroni (2023), and Gualdani and Sinha (2024) (e.g., Chernozhukov, Hong and Tamer (2007)'s subsampling approach) are computationally challenging to implement in our setting because the Markov correlated equilibrium constraints are non-convex.

We propose a computationally attractive strategy for inference based on the key insights from Horowitz and Lee (2023).[21] Our approach is intended to minimize the need for repeatedly solving non-convex optimization problems. For simplicity, let us assume that the transition probability function is known to the researcher. However, an extension to the case where f needs to be estimated is straightforward. Our approach is also applicable to Bayes correlated equilibrium.

In our setting, the sampling errors are only associated with the reduced-form parameters, namely the conditional choice probabilities (CCPs) ϕ. If ϕ is known, the identified set can be estimated without statistical uncertainty. Thus, we propose controlling for the sampling uncertainty using a confidence set for the CCPs.[22]

[20] See Canay and Shaikh (2017) and Canay, Illanes and Velez (2023) for recent surveys of the inference methods for partially identified models.

[21] Koh (2023) uses a similar approach in static discrete games settings.

[22] Horowitz and Lee (2023) proposes methods for carrying out (non-asymptotic) inference when the partially identified parameters are characterized by constraints that involve unknown population means that can be estimated from data. Our approach follows their insights in treating the CCPs as unknown parameters that determine the partially identified set. A similar insight has been employed by Kline and Tamer (2016), who propose a Bayesian approach.

5.1 Confidence Set

Suppose Φ_α is a confidence set for the CCPs ϕ with the property that $\phi \in \Phi_\alpha$ with probability at least $1-\alpha$ as the sample size n goes to infinity. Leading examples of Φ_α are box constraints or ellipsoids. Let $\Theta_I(\phi)$ be the identified set when the CCPs are ϕ. Define the confidence set as

$$\widehat{\Theta}_I^\alpha \equiv \bigcup_{\phi \in \Phi_\alpha} \Theta_I(\phi).$$

If Φ_α covers the true ϕ with high probability, then $\widehat{\Theta}_I^\alpha$ covers the true Θ_I with high probability. Furthermore, checking whether $\theta \in \widehat{\Theta}_I^\alpha$ amounts to solving a nonlinear program in Theorem 5, but with ϕ as part of the optimization variables subject to the constraint $\phi \in \Phi_\alpha$.

Theorem 7. *Let Φ_α be a confidence region for ϕ such that*

$$\liminf_{n\to\infty} \Pr\left(\phi \in \Phi_\alpha\right) \geq 1-\alpha.$$

Then

$$\liminf_{n\to\infty} \Pr\left(\Theta_I \subseteq \widehat{\Theta}_I^\alpha\right) \geq 1-\alpha.$$

Moreover, $\theta \in \widehat{\Theta}_I^\alpha$ if and only if there exists (σ, V, ϕ) subject to $\phi \in \Phi_\alpha$, (8), (9), (10), and (11).

5.2 Implementation

We briefly discuss possible approaches for constructing Φ_α. While there are many possible approaches available, it is helpful to construct Φ_α as linear constraints because they are computationally easy to handle. For example, one can construct simultaneous confidence intervals of the form

$$\Phi \equiv \left\{\phi : \hat{L}_{a|x} \leq \phi_{a|x} \leq \hat{U}_{a|x}, \quad \forall a, x\right\}$$

where the vectors $\hat{L}, \hat{U} \in \mathbb{R}^{|\mathcal{A}|\times|\mathcal{X}|}$ are determined by data. Examples include simultaneous confidence intervals for multinomial probabilities proposed by Fitzpatrick and

Scott (1987) or sup-t band by Montiel Olea and Plagborg-Møller (2019); the former is easier to compute since the endpoints of the intervals can be computed in closed-form, but the latter can be tighter.

An alternative is to take a parametric approach.[23] Suppose $\phi = \phi(\gamma)$, where γ is the parameter vector that governs the CCPs. One may estimate γ and construct simultaneous confidence intervals of the form

$$\Phi \equiv \left\{ \gamma : \hat{L}_k \leq \gamma_k \leq \hat{U}_k, \quad \forall k \right\}$$

where $\hat{L}, \hat{U} \in \mathbb{R}^{\dim(\gamma)}$ can be computed using the asymptotic distribution of $\hat{\gamma}$. One may also consider replacing ϕ with its linear approximation $\phi(\gamma) \approx \phi(\hat{\gamma}) + \nabla\phi(\hat{\gamma})^\top(\gamma - \hat{\gamma})$ to further simplify the computation. In the identified set search, we also add the constraints defining ϕ as proper conditional distributions, i.e., $\phi_{a|x} \geq 0$, $\forall a, x$, and $\sum_a \phi_{a|x} = 1$, $\forall x$.

In our empirical application, we construct simultaneous confidence intervals using the sup-t band approach, following Montiel Olea and Plagborg-Møller (2019).

6 Empirical Application

We study a dynamic entry game by the two largest players in the US coffee chain industry, Starbucks and Dunkin', to illustrate the usefulness of our framework. The US coffee chain industry has grown significantly in recent decades, driven by increasing demand for specialty coffee and the rise of a culture emphasizing quality, convenience, and socialization. While companies compete for market share through pricing, product differentiation, marketing strategies, and in-store service, this market is highly concentrated. In particular, Starbucks and Dunkin' have significantly outpaced other competitors by a wide margin. Starbucks and Dunkin' accounted for 65% and 28% (resp. 55% and 35%) of total sales (resp. outlets) from the top 8 coffee chains in the

[23] With limited observations relative to covariate dimension, nonparametric CCP estimation may fail, as some cells contain no data. A flexible multinomial logit specification, though inconsistent under misspecification, offers lower variance and often outperforms nonparametric estimators in small samples. Aguirregabiria and Mira (2007) show this approach performs well in a two-step pseudo–maximum likelihood framework.

US in 2019 (Technomic, 2019).[24]

Little is publicly known about how coffee chains predict opponents' strategic entry decisions. It is reasonable to expect that firms have good information about their stores' (potential) profitability. Still, less is known about how firms learn about their opponents' profitability. Resorting to the standard "incomplete" information assumption—firms observe their payoff shocks but observe nothing about their opponents' shocks—will likely result in a model misspecification. Whether this misspecification is an empirical concern depends on how much it may change the structural estimates and counterfactual outcomes. We apply our Markov correlated equilibrium framework to estimate the parameters and conduct a counterfactual experiment with weaker assumptions on players' information.

6.1 Data and Descriptive Evidence

Our primary dataset is Data Axle's U.S. Historic Business Database, which has annual location information of all business establishments in the U.S. from 1997 to 2023. Figure 3 shows that the number of Starbucks and Dunkin' stores in the US increased dramatically from 1997 to 2023. Figure 4 illustrates how the firms dramatically expanded their geographic footprint over time.

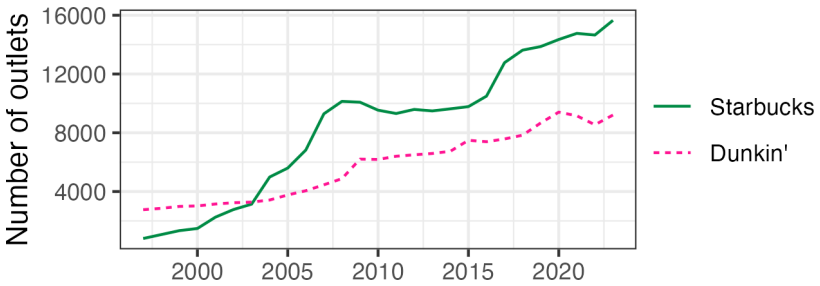


Figure 3: Number of Starbucks and Dunkin' stores in the US over time

We define markets as four-digit 2010 census tracts.[25] Our geographic market definition assumes that coffee chains compete in small geographic areas.[26] A firm is

[24]The third and fourth chains were Tim Hortons and Dutch Bros., which accounted for 2.2% and 1.7% (resp. 2.6% and 1.3%) of the sales (resp. outlets) out of the top 8 chains.

[25]Four-digit census tracts are census-defined geographic units that are contiguous and strike a balance between county (which is too large) and six-digit census tracts (which is too small).

[26]Guler (2018) shows that demand for Starbucks is highly localized. Thomadsen (2005) shows

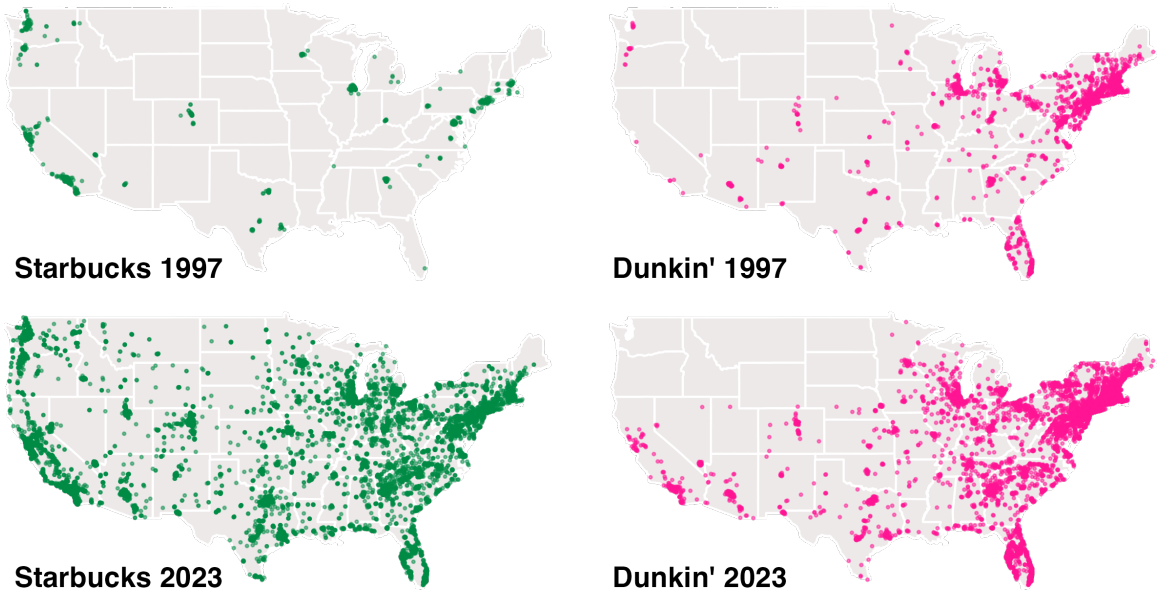


Figure 4: Growth of Starbucks and Dunkin'

active in a market when we observe the firm's outlet in the market. We supplement the firm-location data with the demographic/geographic data from the Longitudinal Track Data Base (Logan, Xu and Stults, 2014) and the neighborhood characteristics data from the National Neighborhood Data Archive. We limit our sample period to 2003–2017 to match the sample period of the neighborhood characteristics data.[27] We exclude the markets without any eating places during the sample period.

Table 3 Panel A shows the active probabilities of each firm from 2018–2023 when the firm either did not have any outlets or had at least one outlet in the corresponding county from 2003–2017. The *ex-post* active probability in a county is much lower when the firm does not have any outlets in that county compared to when it does (1% vs. 20%). Thus, it is reasonable to consider that a firm does not take into account the competitor's situation if the competitor does not have any outlets in the county, and we define a firm as a potential entrant in each market if the firm had at

evidence that only fast-food chains within approximately 0.5 miles—which amounts to 10 minutes walking distance—compete as close substitutes in California; it is likely that coffee consumers travel even less.

[27] We still use the business location data from 2018 to reduce the possibility of incorrectly determining firms not being active due to missing data. We consider a firm to have switched from active to not active in a market (i.e., exited the market) when we do not observe the firm's outlet in the market for the following five years.

least one outlet in the corresponding county during the sample period. We also use own-chain outlets in the county as firm-specific excluded variables that capture the network effects associated with economies of density (Guler, 2018). Table 3 Panel B demonstrates that the number of own-chain outlets in the county has a positive effect on the active probability in markets where both Starbucks and Dunkin' are potential entrants (two-player markets). In contrast, the number of competitor-chain outlets in the county affects the active probability little. We interpret the weak negative relationship as stemming from the firm's strategic behavior in response to the higher competitor's active probability.

Table 3: Active probabilities by the number of outlets in county

	Starbucks	Dunkin'
A. Active probability after 2017		
Own-chain outlets in county $= 0$	0.011	0.011
Own-chain outlets in county > 0	0.211	0.197
B. Active probability in two-player markets		
Own-chain outlets in county $\leq$ median	0.095	0.043
Own-chain outlets in county $>$ median	0.195	0.234
Competitor-chain outlets in county $\leq$ median	0.163	0.146
Competitor-chain outlets in county $>$ median	0.126	0.130

Table 4 reports the summary statistics of a balanced panel of $M = 38,687$ markets observed over $T = 15$ years. Both Starbucks and Dunkin' are potential entrants in 79% of the markets, and 17% and 4% respectively have Starbucks and Dunkin' as a single potential entrant. The firms are active in 14% and 11% of the observations, while entries and exits are rare (1% and 0.2%, respectively), indicating a high correlation in the incumbency status. We assume the firms take the market characteristics as exogenous and time-invariant. For each market, we average characteristics over the observation period and treat them as fixed attributes. The key common characteristic is the number of eating places, which captures overall market attractiveness, as customer traffic tends to rise with the density of food establishments.

Table 4: Summary statistics of 38,687 coffee chain markets from 2003 to 2017

A. Potential entrant	Both	Starbucks	Dunkin'
Market share	0.786	0.172	0.042
B. Player decisions	Starbucks	Dunkin'	
Active probability	0.139	0.113	
Entry probability	0.011	0.007	
Exit probability	0.002	0.002	
C. Market characteristics	Mean	Std. Dev.	Range
Number of eating places	15.6	25.2	$[0.1, 649.8]$
Starbucks outlets in county	592.0	1243.1	$[0.0, 6423.0]$
Dunkin' outlets in county	354.9	750.9	$[0.0, 3948.0]$

6.2 Econometric Model

Our model follows the standard dynamic oligopoly entry/exit model with heterogeneous firms (Aguirregabiria and Mira, 2007; Pakes, Ostrovsky and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008). In each market m and time period t, firm i decides whether to operate in the market ($a_{imt} = 1$) or not ($a_{imt} = 0$). An active firm's flow profit is

$$u_i^\theta(a_{imt} = 1, a_{jmt}, x_{mt}, \varepsilon_{imt}) = \theta_w^\top w_m + \theta_{i,v} v_{im} + \theta_{ec}(1 - z_{imt}) + \theta_{i,ce} a_{jmt} + \varepsilon_{imt}. \quad (15)$$

Here, public state x_{mt} includes three components, common time-invariant market characteristics w_m, firm-specific time-invariant excluded variables (v_{im}, v_{jm}), and incumbency status ($z_{imt} \equiv a_{im,t-1}, z_{jmt} \equiv a_{jm,t-1}$). w_m includes a constant and log number of eating places, and v_{im} is log number of own-chain outlets in the county. Parameters θ_w, $\theta_{i,v}$, θ_{ec}, and $\theta_{i,ce}$ represent coefficients on market characteristics, excluded variable, entry cost, and competitive effects, respectively. We allow the coefficients on the excluded variable and competitive effects to be asymmetric. ε_{imt} is the idiosyncratic per-period payoff shock, which is assumed to be independently drawn from the standard logistic distribution. The state transition is deterministic since the only dynamic state variable is the incumbency status, determined by the firms' actions in the previous period. We set the discount factor to $\delta = 0.9$ and the flow

payoff from being inactive to zero.[28] We discretize market characteristics (w_m, v_{im}) into two bins by grouping the observations above and below the median and replacing the values with within-bin means.

We first estimate the parameters under the standard incomplete information Markov perfect equilibrium assumption using the K-pseudo maximum likelihood approach (Aguirregabiria and Mira, 2007; Bugni and Bunting, 2021). Each player observes their own idiosyncratic shock ε_{imt} but not the opponent's shock ε_{jmt}. We then estimate the Markov correlated equilibrium identified set, maintaining this standard informational assumption. The identified set is informationally robust in the sense that it captures all possible parameters that could generate the data under Markov perfect equilibrium, allowing for the possibility that the players know their own ε_{imt} and may observe information about the opponent's ε_{jmt} before making decisions. We also estimate the fully robust identified set, which additionally encompasses the possibility that the players do not observe or are partially informed about their ε_{imt} (formally defined in Section 4). The identified set is fully robust to any misspecification of the information structure regarding the idiosyncratic shocks $\varepsilon_{mt} = (\varepsilon_{imt}, \varepsilon_{jmt})$. Lastly, we estimate the model without the excluded variable v_{im} with these three information structures above to gauge the importance of the exclusion restriction. We outline our estimation strategy of the Markov correlated equilibrium identified sets and relegate the computational details to Online Appendix A.2.1.

We leverage the excluded variable with large support as discussed in Section 3.3. Our excluded variable v_{im}, the log number of own-chain outlets in the county, has large support because it has no lower bound.[29] Firm i is not a potential entrant (i.e., i's active probability is 0) if the number of own-chain outlets in the county is 0, since $v_{im} = \log 0 = -\infty$. This model specification is motivated by the fact that a firm will be rarely active in markets where the firm did not have any outlets in

[28] Normalizing the flow profit from the outside option to zero is not without loss of generality for counterfactual analysis (Aguirregabiria and Suzuki, 2014; Kalouptsidi, Scott and Souza-Rodrigues, 2017, 2021). We show that our counterfactual experiment of reducing a proportion of entry cost is equivalent to providing an entry subsidy equal to the same proportion of the entry cost net of scrap value without this normalization in Online Appendix A.1.

[29] Given the length of our observation period, we assume that this variable sufficiently captures the time-persistent match effect between firms and counties. While our estimation procedure relies on a specific parametric form of v_{im}, sufficient variation in the excluded variable will tighten the identified set even if we impose another parametric form (see e.g., Magnolfi and Roncoroni (2023)).

the relevant county (Table 3 Panel A). In markets where firm j is not a potential entrant or a_{jmt} is always 0, firm i solves a single-agent dynamic decision problem, as the opponent's idiosyncratic shock ε_{jmt} becomes payoff-irrelevant. We estimate the single-agent discrete choice model, which contains the payoff parameters except the competitive effects parameters in (15),[30] using the two-step pseudo-likelihood method (Aguirregabiria and Mira, 2002). This step allows us to obtain estimates for non-competitive effects parameters, $\theta_w, \theta_{i,v}, \theta_{j,v}$, and θ_{ec}.

We then apply Theorem 7 to estimate the identified set of the competitive effects parameters using markets where both firms are potential entrants. We adopt the non-competitive effects parameters estimated from the single-agent model in the Markov correlated equilibrium identified set estimation to demonstrate the identifying power of our Markov correlated equilibrium framework.[31] We estimate the CCPs using flexible logit specification and construct their simultaneous confidence intervals with a plug-in sup-t implementation (Montiel Olea and Plagborg-Møller, 2019). We estimate a criterion function $Q(\theta)$ that measures the degree of violation in the obedience condition (10) for given competitive effects parameters $\theta = (\theta_{i,ce}, \theta_{j,ce})$. We construct identified set estimates by grid searching θ that satisfies the estimated $Q(\theta) \leq c$ for a small threshold $c > 0$. We use non-zero c to accommodate misspecifications in the structural and CCP logit models. Larger values of c yield more conservative identified sets. We do not incorporate the sampling error associated with the first-stage single-agent model estimation.

We also consider a model without the excluded variables v_{im} and v_{jm}. We grid search for all the payoff parameters $\theta = (\theta_w, \theta_{ec}, \theta_{i,ce}, \theta_{j,ce})$ for this model.

6.3 Estimation Results

Table 5 reports our estimation results with the exclusion restriction. The presence of an opponent has a positive spillover effect on profitability under the standard incom-

[30] We can identify the competitive effects parameters if there is a subset of markets defined by excluded variables where the opponent is always active additionally. It would be challenging to find a sufficient number of such markets to reasonably estimate the competitive effects parameters in our setting (see Table 3 Panel B for the effect of our excluded variable on the active probability).

[31] Aradillas-Lopez and Rosen (2022) also takes a similar two-step approach in static ordered-response games. The authors point identify the non-competitive effects parameters and use the estimates to partially identify the competitive effects parameters.

plete information Markov perfect equilibrium assumption (Column (1)). The 95% confidence intervals (CIs) of the competitive effects parameters become wider as we make the information structure more flexible (compare columns from left to right).[32] The Markov correlated equilibrium estimates suggest the possibility of a large negative opponent effect (columns (2) and (3)). However, the positive effect of Starbucks' presence on Dunkin's profitability is tightly bounded above. The exclusion restriction helps narrow the direction of Dunkin's competitive effects parameter (compare with the results without the excluded variable in Table B1 in Appendix B).

Table 5: Structural parameter estimates

	MPE Standard incomplete (1)	MCE Standard incomplete (2)	MCE Fully robust (3)
Intercept	$[-0.032, -0.002]$	$[-0.040, 0.040]$	$[-0.040, 0.040]$
	-0.016	0.000	0.000
Eating places (log)	$[0.103, 0.110]$	$[0.117, 0.136]$	$[0.117, 0.136]$
	0.106	0.127	0.127
Starbucks outlets in county (log)	$[0.027, 0.031]$	$[0.019, 0.029]$	$[0.019, 0.029]$
	0.029	0.024	0.024
Dunkin' outlets in county (log)	$[0.027, 0.031]$	$[0.029, 0.048]$	$[0.029, 0.048]$
	0.029	0.038	0.038
Entry cost	$[-8.473, -8.365]$	$[-8.323, -8.008]$	$[-8.323, -8.008]$
	-8.408	-8.165	-8.165
Starbucks competitive effect	$[0.027, 0.049]$	$[-0.123, 0.143]$	$[-0.257, 0.449]$
	0.040	-	-
Dunkin' competitive effect	$[0.011, 0.035]$	$[-0.299, 0.038]$	$[-0.300, 0.074]$
	0.024	-	-

Note: Column (1) reports the 95% confidence intervals (CIs) and point estimates for the model with the standard incomplete information Markov perfect equilibrium (MPE) assumption. Columns (2) and (3) use our Markov correlated equilibrium (MCE) framework with the standard incomplete information and without minimal signals, respectively. The MCE column parameters except the competitive effects parameters are obtained using single-agent discrete choice estimation using data where the firms operate as single players. CIs are bootstrapped with resampling at the market level. The competitive effects parameters estimated from the MCE framework do not account for the sampling error associated with the first-stage single-agent dynamic discrete choice estimation.

The market characteristics coefficients and entry cost from the two-player model (column (1)) and the single-agent model (columns (2) and (3) display the same estimates) are similar, which justifies the usage of different subsets of markets to estimate the same parameters.[33] The market characteristics coefficient estimates show positive

[32] We set the criterion function threshold c so that the CIs of the competitive effects parameters become column (1) $\subset$ column (2) $\subset$ column (3).

[33] The CIs do not overlap for some parameters. The discrepancies can be attributed to either

effects of attractive markets and economies of density on profitability. The entry cost estimates are relatively large compared to the other parameter estimates, which is consistent with the low entry probabilities.

6.4 Counterfactual

As a counterfactual experiment, we reduce each firm's entry cost by 10% and simulate the number of active firms from 2003 under three assumptions on information structures used in the estimation. The first is the standard incomplete information Markov perfect equilibrium predictions. We then apply the Markov correlated equilibrium to make informationally robust predictions, assuming either the standard incomplete information or the fully robust specification. For each case, we use the structural parameters estimated under the corresponding solution concept and information structure and apply the same framework to compute the counterfactual predictions.[34]

We consider two extreme scenarios, low and high, to get a sense of a reasonably possible range of the counterfactual predictions. For the low (high) scenario, we use the lower (upper) bounds of the 95% CIs of the competitive effects parameters and find an equilibrium that minimizes (maximizes) the expected number of active firms.[35] As a separate exercise, we also make counterfactual predictions using the lower and upper bounds of the 95% CIs of all structural parameters, rather than only the competitive effects parameters, to incorporate the uncertainty of the non-competitive effects parameters. We use the counterfactual equilibrium CCPs to simulate the evolution of active firms. We detail the counterfactual equilibrium calculations in

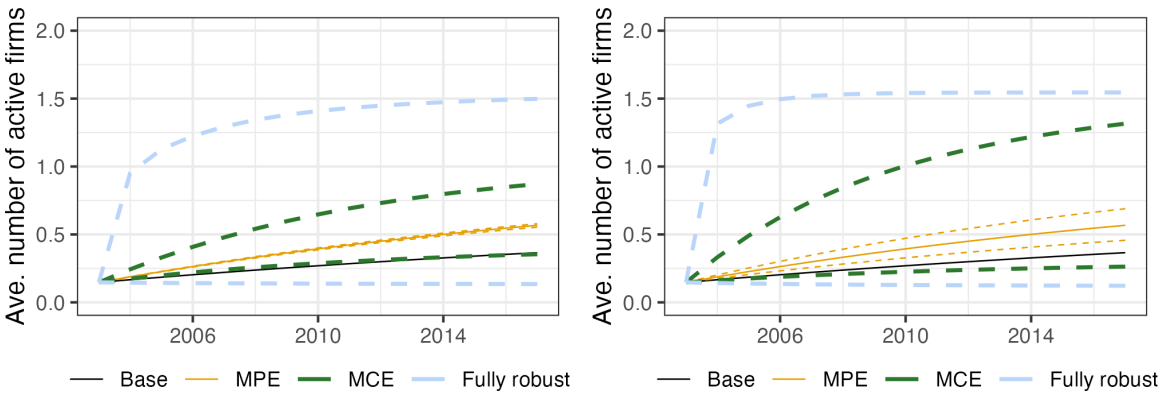
misspecification in the information structure or the differences in the parameters in different markets. We do not provide a formal test of misspecification in the information structure. See Han, Kaido and Magnolfi (2024) for a related test in the Bayes correlated equilibrium context.

[34] In the counterfactual Markov correlated equilibrium predictions, we do not use the *latent information structures* that generate the Markov perfect equilibrium predictions compatible with the data for each parameter value in the Markov correlated equilibrium identified set estimates. Bergemann, Brooks and Morris (2022) shows how to conduct counterfactual analysis with the latent information structure fixed at the pre-counterfactual Markov perfect equilibrium. This approach would not tighten our standard incomplete information Markov correlated equilibrium predictions since the predictions for a given set of structural parameters are tight (Table 6). It may appreciably narrow the fully robust predictions by restricting the range of predictions for a given set of structural parameters.

[35] Note that the pair of values at the lower bounds of the competitive effects parameters for Starbucks and Dunkin' may not necessarily belong to the confidence set.

Online Appendix A.2.2.

Figure 5 plots the counterfactual market evolution from 2003–2017. The y-axis takes the average of all markets, so the minimum and maximum possible values are 0 (no firms active) and 2 (both firms active), respectively. The no policy baseline with the standard incomplete information Markov perfect equilibrium predicts an increase from 0.15 to 0.37 during the sample period (line Base). The Markov perfect equilibrium predicts that the counterfactual policy intervention of reducing the entry cost by 10% increases the number of active firms by 50% in 2017 (line MPE). The range of the two scenarios (dashed curves) is very tight using the CIs of the competitive effects parameters (Figure 5a, average number of active firms in 2017: low, 0.55; high, 0.58) and become somewhat wider using the CIs of all structural parameters (Figure 5b, low, 0.46; high, 0.69).



(a) Use CIs of competitive effects

(b) Use CIs of all parameters

Figure 5: Counterfactual market status (low/high scenarios in dashed lines)

The counterfactual Markov correlated equilibrium predictions, keeping the standard incomplete information, of the average number of active firms in 2017 range from 0.35 (low) to 0.87 (high) using the CIs of the competitive effects parameters (Figure 5a, line MCE). If we are concerned about the informational robustness regarding the competitive effects parameters, this range provides a reasonably plausible set of outcomes, allowing for robust yet still informative policy implications. In contrast, the predictions become much more spread out and less informative, ranging from 0.26 (low) to 1.32 (high), using the CIs of all parameters (Figure 5b, line MCE). The fully robust Markov correlated equilibrium generates very wide predictions even right

after the start of the counterfactual policy intervention (line Fully robust). Moreover, without the exclusion restriction, the counterfactual Markov correlated equilibrium predictions cover almost the entire range of the average number of active firms, $[0, 2]$, even under the standard incomplete information (Figure B1 in Appendix B).

The counterfactual Markov correlated equilibrium predictions involve two sources of multiplicity: partial identification and estimation uncertainty of structural parameters; and flexible counterfactual information structure. The flexible information structure can generate various counterfactual equilibria, resulting in a wide range of predictions. Table 6 tabulates the range of counterfactual predictions for a given set of structural parameters to examine the contribution of these multiplicities. The minimum and maximum number of active firms that can arise from the counterfactual Markov perfect equilibrium predictions coincide for a given set of structural parameters (columns under MPE Standard incomplete), which suggest that we do not have multiple equilibria for a fixed vector of parameters and information structure. Thus, the range of Markov perfect equilibrium predictions is entirely driven by the estimation uncertainty (or 95% CIs) of structural parameters.

Table 6: Contribution of flexible information structure in the counterfactual

		MPE Standard incomplete			MCE Standard incomplete		MCE Fully robust	
Year	Baseline	$\hat{\theta}$ point	$\hat{\theta}$ lower	$\hat{\theta}$ upper	$\hat{\theta}$ lower	$\hat{\theta}$ upper	$\hat{\theta}$ lower	$\hat{\theta}$ upper
A. Use 95% CIs of competitive effects parameters								
2003	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15
2010	0.27	0.39	0.39	0.40	[0.29, 0.32]	[0.64, 0.65]	[0.14, 0.52]	[0.33, 1.41]
2017	0.37	0.57	0.55	0.58	[0.35, 0.38]	[0.87, 0.87]	[0.13, 0.54]	[0.46, 1.50]
B. Use 95% CIs of all parameters								
2003	0.15	0.15	0.15	0.15	0.15	0.15	0.15	0.15
2010	0.27	0.39	0.33	0.47	[0.23, 0.23]	[1.00, 1.01]	[0.13, 0.27]	[0.27, 1.54]
2017	0.37	0.57	0.46	0.69	[0.26, 0.26]	[1.31, 1.32]	[0.12, 0.31]	[0.36, 1.55]

Note: Values are the simulated average number of active firms. The baseline is the standard incomplete information Markov perfect equilibrium (MPE) predictions without entry subsidy. The other columns show the counterfactual predictions with entry subsidy. Column $\hat{\theta}$ point uses the point estimates of the structural parameters. Panel A and B use the lower/upper bound of 95% confidence intervals (CIs) of competitive effects parameters and all parameters, respectively, for $\hat{\theta}$ lower/upper. The range of the MCE counterfactuals shows the minimum and maximum predictions for a given set of structural parameters.

The prediction ranges for a given set of structural parameters remain very tight using the Markov correlated equilibrium, using the standard incomplete information

as the baseline (columns under MCE Standard incomplete). These prediction ranges encompass the counterfactual outcomes arising from any Markov perfect equilibrium when the players minimally observe their payoff shocks and may observe more information about opponents' payoff shocks. The tight counterfactual outcome bounds indicate that the ability to obtain more information about opponents' payoff shocks and the possibility that such information is correlated across firms have little impact on the outcome, given the model parameters fixed. The small competitive effects parameters (in absolute values) can explain the tight counterfactual bounds. If the competitive effects are null, then the firms do not care about the opponents' behavior, and there is no room for an information designer to affect the total number of entrants, on average.[36] Here, the competitive effects parameters range between -0.30 and 0.14, which is small considering the standard deviation of the idiosyncratic payoff shock (the standard deviation of the standard logistic random variable is $\sqrt{\pi^2/3} \approx 1.81$). Consequently, the degree to which the counterfactual outcomes can vary depends largely on the structural parameter values. In contrast, the prediction ranges, when fixing structural parameters, become much wider by allowing the possibility that players may not observe or be partially informed about their own payoff shocks (columns under MCE Fully robust).

7 Conclusion

In this paper, we have studied the identification, estimation, and counterfactual analyses of a common class of dynamic games for empirical analysis, assuming the researcher knows some minimal information available to players. To facilitate informationally robust econometric analysis of dynamic games, we have defined *Markov correlated equilibrium*, a dynamic analog of Bayes correlated equilibrium, and studied its properties. We have discussed new challenges that arise in dynamic environments and proposed practical solutions. We have applied our proposed method to study the dynamic entry and exit decisions of the two largest coffee chains in the US.

We conclude this paper with a discussion of future research directions. First, further refinement of Markov correlated equilibrium (and thus Bayes correlated equi-

[36]The information designer can still induce correlation in the firms' actions. As a result, the CCPs may deviate from those predicted under the standard MPE scenario.

librium) will attract greater interest from empiricists. The type of informational robustness offered by Markov/Bayes correlated equilibrium admits too many possibilities. Markov correlated equilibrium identified sets and counterfactual outcomes can span a wide range, which may not be empirically informative; Magnolfi and Roncoroni (2023) find similar results in their empirical application. Numerical examples from Koh (2023) and Magnolfi and Roncoroni (2023) show that even complete information correlated equilibrium identified sets can be large, indicating that correlation in actions can also weaken the identifying power substantially. In some settings, we may be able to reasonably exclude some information structures from consideration. Additionally, we may want to restrict attention to some equilibrium selection mechanisms to narrow down the possibilities. It will be fruitful to find a tractable approach that can accommodate further assumptions on the underlying information structure and equilibrium selection mechanism.

Second, estimation with Markov correlated equilibrium remains computationally challenging. A similar problem exists for Bayes correlated equilibrium since the analyst needs to run a grid search. Studying how to conduct computationally tractable estimation and inference in this class of partially identified models will be crucial for broadening the scope of applications.

Disclosure Statement During the preparation of this work the authors used Open AI's ChatGPT 5.0 in order to check grammar and improve clarity of the writing. After using this tool/service, the authors reviewed and edited the content as needed and takes full responsibility for the content of the published article. Wharton Research Data Services (WRDS) was used in preparing part of the data set used in the research reported in this manuscript. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third-party suppliers. The authors declare that no financial or personal conflicts of interest influenced the research presented in this paper.

A Proofs

We use BNE, BCE, MPE, and MCE for Bayes Nash equilibrium, Bayes correlated equilibrium, Markov perfect equilibrium, and Markov correlated equilibrium, respectively, for brevity in the following proofs.

A.1 Proof of Lemma 1

Proof. ($\Rightarrow$) Let β be a MPE of (G, S). β induces (G^β, S). The BNE best-response condition is directly implied by the MPE condition of (G, S).

($\Leftarrow$) Let β be a BNE of (G^β, S). The BNE best-response condition implies that there is no profitable one-shot deviation given that each player expects all players to follow the prescription of β in the future. The absence of profitable one-shot deviation implies that β is a subgame perfect equilibrium of (G, S) by the one-shot deviation principle. Therefore, β is a MPE of (G, S). □

A.2 Proof of Lemma 2

Proof. The statement follows by applying the one-shot deviation principle. □

A.3 Proof of Lemma 3

Proof. It is enough to show that $V_i^\beta(x) = V_i^\sigma(x)$ for all $i \in \mathcal{I}$ and $x \in \mathcal{X}$ because then it follows that $v_i^\beta(a, x, \varepsilon) = v_i^\sigma(a, x, \varepsilon)$ for all $i \in \mathcal{I}$, $a \in \mathcal{A}$, $x \in \mathcal{X}$, and $\varepsilon \in \mathcal{E}$.

Recall that V_i^β and V_i^σ satisfy (4) for (G, S^*) and (5) for (G, S), respectively. Since β induces σ, (4) for (G, S^*) is equivalent to

$$\begin{aligned} V_i^\beta(x) &= \sum_{\varepsilon,\tau,\tilde{\tau},a} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\lambda_{\tilde{\tau}|x,\varepsilon,\tau}\beta_{a|x,\tau,\tilde{\tau}} \left(u_i(a, x, \varepsilon) + \delta \sum_{x'} V_i^\beta(x') f_{x'|a,x,\varepsilon} \right) \\ &= \sum_{\varepsilon,\tau,a} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}\sigma_{a|x,\varepsilon,\tau} \left(u_i(a, x, \varepsilon) + \delta \sum_{x'} V_i^\beta(x') f_{x'|a,x,\varepsilon} \right), \end{aligned}$$

where $\beta_{a|x,\tau,\tilde{\tau}} \equiv \prod_{i=1}^I \beta_i(a_i|x, \tau_i, \tilde{\tau}_i)$. Comparison of the above equation for V_i^β and (5) for (G, S) that defines V_i^σ implies that $V_i^\beta = V_i^\sigma$. □

A.4 Proof of Theorem 1

Proof. To prove the statement of the theorem, we make use of Theorem 1 of Bergemann and Morris (2016) (BM), which we restate below.

Lemma 4 (BM Theorem 1)**.** *Let (G, S) be a static game of incomplete information. A decision rule σ is a BCE of (G, S) if and only if, for some expansion S^* of S, there is a BNE of (G, S^*) that induces σ.*

($\subseteq$) Suppose σ is a MCE of (G, S). We want to show that there exists an expansion S^* and a strategy profile β in (G, S^*) such that β is a MPE of (G, S^*) and β induces σ. Since σ is a MCE of (G, S), it is a BCE of (G^σ, S). By BM Theorem 1, there exists an expansion S^* of S and a strategy profile β of (G^σ, S^*) such that β is a BNE of (G^σ, S^*) and β induces σ. But since $G^\sigma = G^\beta$, β is also a BNE of $\left(G^\beta, S^*\right)$, which in turn implies that β is a MPE of (G, S^*).

($\supseteq$) Suppose β is a MPE of (G, S^*). We want to show that if β induces σ in (G, S), then σ is a MCE of (G, S). Since β is a MPE of (G, S^*), it is a BNE of $\left(G^\beta, S^*\right)$. By BM Theorem 1, if σ is induced by β, then σ is a BCE of $\left(G^\beta, S\right)$. But since $G^\beta = G^\sigma$, σ is also a BCE of (G^σ, S), which in turn implies that σ is a MCE of (G, S). □

A.5 Proof of Corollary 1

Proof. ($\subseteq$) Take $\phi \in \mathcal{P}^{MCE}_{a|x}(G, S)$. By definition, there exists $q \in \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S)$ such that q induces ϕ, i.e.,$\phi_{a|x} = \sum_{\varepsilon\in\mathcal{E},\tau\in\mathcal{T}} \psi_{\varepsilon|x}\pi_{\tau|x,\varepsilon}q_{a|x,\varepsilon,\tau}$, $\forall a \in \mathcal{A}, x \in \mathcal{X}$. Then there exists S^* such that $S^* \succsim_E S$ and $q \in \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^*)$, implying that $\phi \in \mathcal{P}^{MPE}_{a|x}(G, S^*)$.

($\supseteq$) Take $\phi \in \bigcup_{S^*\succsim_E S} \mathcal{P}^{MPE}_{a|x}(G, S^*)$. Then there exists S^* such that $S^* \succsim_E S$ and $\phi \in \mathcal{P}^{MPE}_{a|x}(G, S^*)$, which implies that there exists $q \in \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^*)$ such that q induces ϕ. Then since $q \in \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S)$, we have $q \in \mathcal{P}^{MCE}_{a|x}(G, S)$. □

A.6 Proof of Theorem 2

Proof. Let $\phi \in \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S^1)$. We want to show $\phi \in \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S^2)$. From our Theorem 1, we have

$$\mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S^1) = \bigcup_{S^*\succsim_E S^1} \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^*) \text{ and } \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S^2) = \bigcup_{S^{**}\succsim_E S^2} \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^{**}).$$

If $\phi \in \mathcal{P}^{MCE}_{a|x,\varepsilon,\tau}(G, S^1)$, then there exists S^* such that $\phi \in \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^*)$. But since $S^* \succsim_E S^1$ and $S^1 \succsim_E S^2$, we have $S^* \succsim_E S^2$, which implies $\phi \in \bigcup_{S^{**} \succsim_E S^2} \mathcal{P}^{MPE}_{a|x,\varepsilon,\tau}(G, S^{**})$, which is what we wanted to show. □

A.7 Proof of Theorem 3

Proof. We want to show that $\Theta_I^{MCE}(S) = \bigcup_{\tilde{S} \succsim_E S} \Theta_I^{MPE}(\tilde{S})$. Take $\theta \in \Theta_I^{MCE}(S)$. By definition, $\phi \in \mathcal{P}^{MCE}_{a|x}(G^\theta, S)$. By Corollary 1, there exists some $S^* \succsim_E S$ such that $\phi \in \mathcal{P}^{MPE}_{a|x}(G, S^*)$. But this immediately implies $\theta \in \Theta_I^{MPE}(S^*)$, and thus $\theta \in \bigcup_{\tilde{S} \succsim_E S} \Theta_I^{MPE}(\tilde{S})$. The proof of converse follows the above steps in reverse. □

A.8 Proof of Theorem 4

Proof. The statement follows from Theorem 2. □

A.9 Proof of Theorem 5

Proof. The statement directly follows from the fact that the constraints (8), (9), (10), and (11) characterize the MCE identification conditions as implied by the definition of identified sets. Note that the ex-ante value function equation (11) is equivalent to (5) because

$$\begin{aligned}\sum_{\varepsilon,\tau,a} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \sum_{x'} V_{i,x'} f_{x'|a,x} &= \sum_{a,x'} \left(\sum_{\varepsilon,\tau} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \right) V_{i,x'} f_{x'|a,x} \\ &= \sum_{a,x'} \phi_{a|x} V_{i,x'} f_{x'|a,x}.\end{aligned}$$

□

A.10 Proof of Theorem 6

Proof. The ex-ante value function (14) can be derived by expanding (5) and imposing (12) similarly to the proof of Theorem 5 in Appendix A.9. The obedience condition

(13) is derived as follows. The obedience condition in the fully robust case is

$$\sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} \partial v_i^\theta(a_i', a_i, x, \varepsilon_i) \leq 0. \tag{16}$$

But since

$$\begin{aligned}
&\sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} v_i^\theta(a_i', a_{-i}, x, \varepsilon_i) \\
&= \sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} \left(u_i^\theta(a_i', a_{-i}, x, \varepsilon_i) + \delta \sum_{x'} V_{i,x'} f_{x'|a_i', a_{-i}, x} \right) \\
&= \sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} u_i^\theta(a_i', a_{-i}, x, \varepsilon_i) + \sum_{a_{-i}} \left(\sum_{\varepsilon} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} \right) \left(\delta \sum_{x'} V_{i,x'} f_{x'|a_i', a_{-i}, x} \right) \\
&= \sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} u_i^\theta(a_i', a_{-i}, x, \varepsilon_i) + \delta \sum_{a_{-i}, x'} \phi_{a|x} V_{i,x'} f_{x'|a_i', a_{-i}, x},
\end{aligned}$$

where the last equation is obtained by imposing (12), the obedience condition (16) can be replaced by

$$\sum_{\varepsilon, a_{-i}} \psi_{\varepsilon|x} \sigma_{a|x,\varepsilon} \partial u_i^\theta(a_i', a, x, \varepsilon_i) + \delta \sum_{a_{-i}, x'} \phi_{a|x} V_{i,x'} \partial f_{x'|a_i', a, x} \leq 0, \quad \forall i, x, a_i, a_i'$$

where $\partial u_i^\theta(a_i', a, x, \varepsilon_i) \equiv u_i^\theta(a_i', a_{-i}, x, \varepsilon_i) - u_i^\theta(a, x, \varepsilon_i)$ and $\partial f_{x'|a_i', a, x} \equiv f_{x'|a_i', a_{-i}, x} - f_{x'|a, x}$. □

A.11 Proof of Theorem 7

Proof. By construction, $\Pr\left(\Theta_I(\phi) \subseteq \bigcup_{\tilde{\phi} \in \Phi_\alpha} \Theta_I(\tilde{\phi})\right) \geq \Pr(\phi \in \Phi_\alpha)$. By taking the limit inferior on both sides, we obtain the desired result. The remaining statement holds by the definitions of $\Theta_I(\phi)$ and $\widehat{\Theta}_I^\alpha$. □

Table B1: Structural parameter estimates without exclusion restriction

	MPE Standard incomplete (1)	MCE Standard incomplete (2)	MCE Fully robust (3)
Intercept	[0.120, 0.143]	[0.114, 0.356]	[−0.530, 0.390]
	0.132	-	-
Eating places (log)	[0.106, 0.114]	[0.000, 0.164]	[0.000, 0.368]
	0.110	-	-
Entry cost	[−8.871, −8.756]	[−9.051, −5.792]	[−9.087, −0.006]
	−8.813	-	-
Starbucks competitive effect	[0.077, 0.096]	[−0.176, 0.146]	[−0.677, 0.701]
	0.086	-	-
Dunkin' competitive effect	[−0.015, 0.005]	[−0.195, 0.211]	[−0.708, 0.632]
	−0.005	-	-

Note: Column (1) reports the 95% confidence intervals (CIs) and point estimates for the model with the standard incomplete information Markov perfect equilibrium (MPE) assumption. Columns (2) and (3) use our Markov correlated equilibrium (MCE) framework with the standard incomplete information and without minimal signals, respectively. The parameter spaces of the eating places (log) and entry cost are bounded below and above by zero, respectively. CIs are bootstrapped with resampling at the market level.

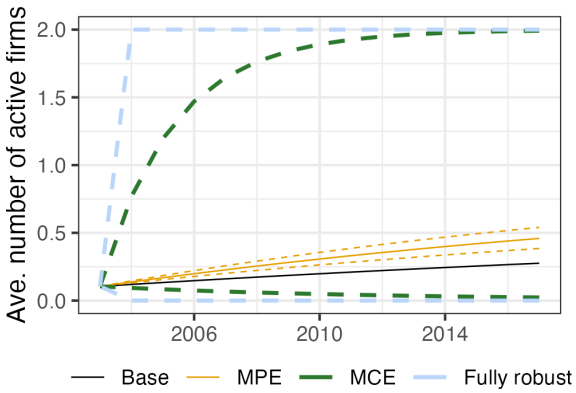


Figure B1: Counterfactual market status without exclusion restriction

B Results Without Exclusion Restriction

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990) "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58 (5), 1041–1063.

Aguirregabiria, Victor, Allan Collard-Wexler, and Stephen P Ryan (2021) "Dynamic

Games in Empirical Industrial Organization," in *Handbook of Industrial Organization*, 4, 225–343.

Aguirregabiria, Victor and Arvind Magesan (2020) "Identification and Estimation of Dynamic Games When Players' Beliefs Are Not in Equilibrium," *Review of Economic Studies*, 87 (2), 582–625, 10.1093/restud/rdz013.

Aguirregabiria, Victor and Pedro Mira (2002) "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models," *Econometrica*, 70 (4), 1519–1543, 10.1111/1468-0262.00340.

——— (2007) "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75 (1), 1–53, 10.1111/j.1468-0262.2007.00731.x.

Aguirregabiria, Victor and Junichi Suzuki (2014) "Identification and Counterfactuals in Dynamic Models of Market Entry and Exit," *Quantitative Marketing and Economics*, 12, 267–304.

——— (2016) "Empirical games of market entry and spatial competition in retail industries," in *Handbook on the Economics of Retailing and Distribution*, 201–232: Edward Elgar Publishing.

Aradillas-Lopez, Andres and Adam M Rosen (2022) "Inference in Ordered Response Games With Complete Information," *Journal of Econometrics*, 226 (2), 451–476.

Arcidiacono, Peter, Patrick Bayer, Jason R Blevins, and Paul B Ellickson (2016) "Estimation of dynamic discrete choice models in continuous time with an application to retail competition," *The Review of Economic Studies*, 83 (3), 889–931.

Asker, John, Chaim Fershtman, Jihye Jeon, and Ariel Pakes (2020) "A Computational Framework for Analyzing Dynamic Auctions: The Market Impact of Information Sharing," *RAND Journal of Economics*, 51 (3), 805–839, 10.1111/1756-2171.12341.

Athey, Susan, Kyle Bagwell, and Chris Sanchirico (2004) "Collusion and price rigidity," *The Review of Economic Studies*, 71 (2), 317–349.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin (2007) "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75 (5), 1331–1370, 10.1111/j.1468-0262.2007.00796.x.

Beresteanu, Arie, Paul B Ellickson, Sanjog Misra, and James C Reeder III (2024) "The dynamics of retail oligopoly," *Available at SSRN 4784452*.

Beresteanu, Arie, Ilya Molchanov, and Francesca Molinari (2011) "Sharp Identification Regions in Models with Convex Moment Predictions," *Econometrica*, 79 (6), 1785–1821.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2022) "Counterfactuals with Latent Information," *American Economic Review*, 112 (1), 343–368, 10.1257/aer.20210496.

Bergemann, Dirk and Stephen Morris (2013) "Robust Predictions in Games With Incomplete Information," *Econometrica*, 81 (4), 1251–1308, 10.3982/ECTA11105.

——— (2016) "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11 (2), 487–522, 10.3982/TE1808.

Bugni, Federico A and Jackson Bunting (2021) "On the Iterated Estimation of Dynamic Discrete Choice Games," *Review of Economic Studies*, 88 (3), 1031–1073.

Canay, Ivan A, Gastón Illanes, and Amilcar Velez (2023) "A user's guide for inference in models defined by moment inequalities," *Journal of Econometrics*, 105558.

Canay, Ivan A and Azeem M Shaikh (2017) "Practical and Theoretical Advances in Inference for Partially Identified Models," *Advances in Economics and Econometrics*, 2, 271–306.

Canen, Nathan and Kyungchul Song (2024) "A Decomposition Approach to Counterfactual Analysis in Game-Theoretic Models," https://arxiv.org/abs/2010.08868.

Chernozhukov, Victor, Han Hong, and Elie Tamer (2007) "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, 75 (5), 1243–1284.

Ciliberto, Federico and Elie Tamer (2009) "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 77 (6), 1791–1828.

Dearing, Adam and Jason R Blevins (2025) "Efficient and convergent sequential pseudo-likelihood estimation of dynamic discrete games," *Review of Economic Studies*, 92 (2), 981–1021.

Doraszelski, Ulrich and Mark Satterthwaite (2010) "Computable Markov-Perfect Industry Dynamics," *RAND Journal of Economics*, 41 (2), 215–243, 10.1111/j.1756-2171.2010.00097.x.

Doval, Laura and Jeffrey C Ely (2020) "Sequential Information Design," *Econometrica*, 88 (6), 2575–2608.

Duffie, Darrell, John Geanakoplos, Andreu Mas-Colell, and Andrew McLennan (1994) "Stationary Markov Equilibria," *Econometrica*, 62 (4), 745–81.

Egesdal, Michael, Zhenyu Lai, and Che-Lin Su (2015) "Estimating Dynamic Discrete-Choice Games of Incomplete Information," *Quantitative Economics*, 6 (3), 567–597.

Ellickson, Paul B, Stephanie Houghton, and Christopher Timmins (2013) "Estimating network economies in retail chains: a revealed preference approach," *The RAND Journal of Economics*, 44 (2), 169–193.

Ericson, Richard and Ariel Pakes (1995) "Markov-perfect industry dynamics: A framework for empirical work," *Review of Economic Studies*, 62 (1), 53–82.

Fang, Limin and Nathan Yang (2024) "Measuring deterrence motives in dynamic oligopoly games," *Management Science*, 70 (6), 3527–3565.

Fershtman, Chaim and Ariel Pakes (2012) "Dynamic Games with Asymmetric Information: A Framework for Empirical Work," *The Quarterly Journal of Economics*, 127 (4), 1611–1661, 10.1093/qje/qjs025.

Fitzpatrick, Simon and Alastair Scott (1987) "Quick Simultaneous Confidence Intervals for Multinomial Proportions," *Journal of the American Statistical Association*, 82 (399), 875–878.

Forges, Françoise (2020) "Correlated Equilibria and Communication in Games," *Complex Social and Behavioral Systems: Game Theory and Agent-Based Models*, 107–118.

Fudenberg, Drew and Jean Tirole (1991) *Game Theory*: MIT press.

Galichon, Alfred and Marc Henry (2011) "Set identification in models with multiple equilibria," *Review of Economic Studies*, 78 (4), 1264–1298.

Grieco, Paul LE (2014) "Discrete Games with Flexible Information Structures: An Application to Local Grocery Markets," *RAND Journal of Economics*, 45 (2), 303–340.

Gualdani, Cristina and Shruti Sinha (2024) "Identification in discrete choice models with imperfect information," *Journal of Econometrics*, 244 (1), 105854.

Guler, Ali Umut (2018) "Inferring the Economics of Store Density from Closures: The Starbucks Case," *Marketing Science*, 37 (4), 611–630.

Han, Sukjin, Hiroaki Kaido, and Lorenzo Magnolfi (2024) "Testing Information Ordering for Strategic Agents," *arXiv preprint arXiv:2402.19425*.

He, Wei and Yeneng Sun (2017) "Stationary Markov Perfect Equilibria in Discounted Stochastic Games," *Journal of Economic Theory*, 169, 35–61.

Holmes, Thomas J (2011) "The diffusion of Wal-Mart and economies of density," *Econometrica*, 79 (1), 253–302.

Horowitz, Joel L and Sokbae Lee (2023) "Inference in a class of optimization problems: Confidence regions and finite sample bounds on errors in coverage probabilities," *Journal of Business & Economic Statistics*, 41 (3), 927–938.

Hotz, V Joseph and Robert A Miller (1993) "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60 (3), 497–529.

Jia, Panle (2008) "What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry," *Econometrica*, 76 (6), 1263–1316.

Kaido, Hiroaki, Francesca Molinari, and Jörg Stoye (2019) "Confidence Intervals for Projections of Partially Identified Parameters," *Econometrica*, 87 (4), 1397–1432.

Kalouptsidi, Myrto, Paul T Scott, and Eduardo Souza-Rodrigues (2017) "On the Non-identification of Counterfactuals in Dynamic Discrete Games," *International Journal of Industrial Organization*, 50, 362–371.

——— (2021) "Identification of Counterfactuals in Dynamic Discrete Choice Models," *Quantitative Economics*, 12 (2), 351–403.

Kennan, John (2006) "A Note on Discrete Approximations of Continuous Distributions," Unpublished Manuscript, University of Wisconsin–Madison.

Kline, Brendan and Elie Tamer (2016) "Bayesian Inference in a Class of Partially Identified Models," *Quantitative Economics*, 7 (2), 329–366.

Koh, Paul S (2023) "Stable Outcomes and Information in Games: An Empirical Framework," *Journal of Econometrics*, 237 (1), 105499.

Logan, John R, Zengwang Xu, and Brian J Stults (2014) "Interpolating Us Decennial Census Tract Data From as Early as 1970 to 2010: A Longitudinal Tract Database," *Professional Geographer*, 66 (3), 412–420.

Magnolfi, Lorenzo and Camilla Roncoroni (2023) "Estimation of Discrete Games with Weak Assumptions on Information," *Review of Economic Studies*, 90 (4), 2006–2041, 10.1093/restud/rdac058.

Mailath, George J and Larry Samuelson (2006) *Repeated Games and Reputations: Long-Run Relationships*: Oxford University Press.

Makris, Miltiadis and Ludovic Renou (2023) "Information Design in Multistage Games," *Theoretical Economics*, 18 (4), 1475–1509.

Maskin, Eric and Jean Tirole (2001) "Markov perfect equilibrium: I. Observable actions," *Journal of Economic Theory*, 100 (2), 191–219.

Montiel Olea, José Luis and Mikkel Plagborg-Møller (2019) "Simultaneous Confidence Bands: Theory, Implementation, and an Application to SVARs," *Journal of Applied Econometrics*, 34 (1), 1–17.

Nishida, Mitsukuni (2015) "Estimating a model of strategic network choice: The convenience-store industry in Okinawa," *Marketing Science*, 34 (1), 20–38.

Nowak, A.S. and T.E.S. Raghavan (1992) "Existence of Stationary Correlated Equilibria with Symmetric Information for Discounted Stochastic Games," *Mathematics of Operations Research*, 17 (3), 519–526.

Pakes, Ariel, Michael Ostrovsky, and Steven Berry (2007) "Simple Estimators for the Parameters of Discrete Dynamic Games (With Entry/Exit Examples)," *RAND Journal of Economics*, 38 (2), 373–399, 10.1111/j.1756-2171.2007.tb00073.x.

Pakes, Ariel, Jack Porter, Kate Ho, and Joy Ishii (2015) "Moment Inequalities and Their Application," *Econometrica*, 83 (1), 315–334.

Pesendorfer, Martin and Philipp Schmidt-Dengler (2008) "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, 75 (3), 901–928.

Rust, John (1987) "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 999–1033.

——— (1994) "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, 4, Chap. 51, 3081–3143, 10.1016/S1573-4412(05)80020-0.

Solan, Eilon (2001) "Characterization of Correlated Equilibria in Stochastic Games," *International Journal of Game Theory*, 30 (2), 259–277.

Solan, Eilon and Nicolas Vieille (2002) "Correlated Equilibrium in Stochastic Games," *Games and Economic Behavior*, 38 (2), 362–399.

Su, Che-Lin and Kenneth L Judd (2012) "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80 (5), 2213–2230.

Syrgkanis, Vasilis, Elie Tamer, and Juba Ziani (2021) "Inference on Auctions with Weak Assumptions on Information," working paper.

Tamer, Elie (2003) "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria," *Review of Economic Studies*, 70 (1), 147–165, 10.1111/1467-937X.00240.

Technomic (2019) "Top 500 Chain Restaurant Report," Accessed via https://www.restaurantbusinessonline.com/financing/largest-coffee-cafe-chains-us.

Thomadsen, Raphael (2005) "The Effect of Ownership Structure on Prices in Geographically Differentiated Industries," *RAND Journal of Economics*, 908–929.

Von Stengel, Bernhard and Françoise Forges (2008) "Extensive-Form Correlated Equilibrium: Definition and Computational Complexity," *Mathematics of Operations Research*, 33 (4), 1002–1022.

Yang, Nathan (2020) "Learning in retail entry," *International Journal of Research in Marketing*, 37 (2), 336–355.

Online Appendix

A Empirical Application Details

A.1 Interpretation of Counterfactual Analysis

As noted in a series of papers (Aguirregabiria and Suzuki, 2014; Kalouptsidi, Scott and Souza-Rodrigues, 2017, 2021), normalizing the flow payoff from the outside option to zero is not without loss of generality for counterfactual analysis. However, we provide a way to interpret our normalization and counterfactual analysis under the assumption that a parametric specification is correct.

Suppose that the flow payoff is given by

$$u_i^\theta(a_{it}, a_{jt}, x_t, \varepsilon_{it}) = \begin{cases} \theta^*_{i,0} + \theta^*_{i,w} w_t + \theta^*_{i,ec}(1 - z_{it}) + \theta^*_{i,ce} a_{jt} + \varepsilon_{it} & \text{if } a_{it} = 1 \\ \theta^*_{i,sv} z_{it} & \text{if } a_{it} = 0 \end{cases}, \quad \text{(A1)}$$

where $\theta^*_{i,ec} < 0$ is the true entry cost and $\theta^*_{i,sv} > 0$ is the true scrap value. A critical assumption in (A1) is that the scrap value is independent of the market characteristics. Payoff (A1) implies that the payoff difference is

$$\begin{aligned} &u_i^\theta(1, a_{jt}, x_t, \varepsilon_{it}) - u_i^\theta(0, a_{jt}, x_t, \varepsilon_{it}) = \\ &\quad \left(\theta^*_{i,0} - \theta^*_{i,sv}\right) + \theta^*_{i,w} w_t + \left(\theta^*_{i,ec} + \theta^*_{i,sv}\right)(1 - z_{it}) + \theta^*_{i,ce} a_{jt} + \varepsilon_{it}. \end{aligned}$$

Thus, the estimated intercept captures $(\theta^*_{i,0} - \theta^*_{i,sv})$ (i.e., fixed cost net of scrap value), and the estimated entry cost captures $(\theta^*_{i,ec} + \theta^*_{i,sv})$ (i.e., entry cost net of scrap value). This is consistent with Aguirregabiria and Suzuki (2014) that the fixed cost, entry cost, and scrap value are not separately identified.

Now consider the counterfactual experiment of reducing the entry cost as in Section 6.4. One can show that this is equivalent to holding $\theta^*_{i,sv}$ fixed and reducing $\theta^*_{i,ec}$ (making it less negative) to $\theta^{**}_{i,ec}$ by

$$\theta^{**}_{i,ec} = \theta^*_{i,ec} - \lambda \underbrace{(\theta^*_{i,ec} + \theta^*_{i,sv})}_{<0},$$

where λ is the proportion of entry cost intended to be cut. Thus, our counterfactual experiment is equivalent to providing an entry subsidy equal to λ of the entry cost net of scrap value.

A.2 Computation

A.2.1 Estimation

We describe our estimation strategy using the excluded variable with large support. We estimate a single-agent discrete choice model with flow profit,

$$u_i^{\theta}(a_{imt}, x_{mt}, \varepsilon_{imt}) = a_{imt}\{\theta_w^{\top} w_m + \theta_{i,v} v_{im} + \theta_{ec}(1 - z_{imt}) + \varepsilon_{imt}\},$$

using markets where firm i is the only potential entrant. Firm j is not a potential entrant in a market if it had no outlets in the corresponding county during the sample period. The competitive effects term, $\theta_{i,ce} a_{jmt}$, is dropped from the two-player flow profit (15) since firm j is not a potential entrant, i.e., a_{jmt} is always 0. We apply the two-step pseudo-likelihood method of Aguirregabiria and Mira (2002) in the estimation. Table A1 in Online Appendix A.3 shows the first-step logit model estimates. Table 5 columns (2) and (3) tabulate the structural estimates $(\hat{\theta}_w, \hat{\theta}_{S,v}, \hat{\theta}_{D,v}, \hat{\theta}_{ec})$, where $i = S$ for Starbucks and $i = D$ for Dunkin', in the rows from "Intercept" to "Entry cost."

We estimate the Markov correlated equilibrium identified set of the competitive effects parameters $\theta = (\theta_{S,ce}, \theta_{D,ce})$ in the flow profit (15) using markets where both firms are potential entrants, considering that the non-competitive effects parameters are given as those estimated from the single-player markets. We adopt the same non-competitive effects parameter estimates $(\hat{\theta}_w, \hat{\theta}_{S,v}, \hat{\theta}_{D,v}, \hat{\theta}_{ec})$ in the two Markov correlated equilibrium identified set estimations to demonstrate the identifying power of our Markov correlated equilibrium framework. We apply Theorem 7 to obtain the confidence region of the identified set of $\theta \in \Theta$, denoted as $\widehat{\Theta}_I$. We take the criterion function approach as commonly done in the partial identification literature (e.g., Chernozhukov, Hong and Tamer (2007)) rather than testing whether $\theta \in \widehat{\Theta}_I$ by solving a feasibility program that directly follows Theorem 7. We define a criterion function $Q : \Theta \mapsto \mathbb{R}_+$ such that $Q(\theta) = 0$ if $\theta \in \widehat{\Theta}_I$ and $Q(\theta) > 0$ otherwise. The crite-

rion function is informative about the degree of violation in the constraints, while the output of a feasibility program is binary ("accepted" or "rejected"). Moreover, $\widehat{\Theta}_I$ can be numerically approximated by considering the sublevel set $\widehat{\Theta}_I(c) \equiv \{\theta : \ Q(\theta) \leq c\}$ with a small threshold $c > 0$.

We define the criterion function so that it measures the degree of violation in the obedience condition (10) for a given θ:

$$Q(\theta) = \min_{q \geq 0, \sigma, V, \phi} \quad q$$

subject to

$$\begin{aligned}
&\sigma_{a|x,\varepsilon,\tau} \geq 0 \text{ for each } a \text{ and } \sum_a \sigma_{a|x,\varepsilon,\tau} = 1, && \forall x, \varepsilon, \tau \\
&\phi_{a|x} = \sum_{\varepsilon,\tau} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \text{ and } \hat{L}_{a|x} \leq \phi_{a|x} \leq \hat{U}_{a|x} && \forall a, x \\
&\sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \partial v_i^\theta(a_i', a, x, \varepsilon_i) \leq q, && \forall i, x, \tau_i, a_i, a_i' \\
&V_{i,x} = \sum_{\varepsilon,\tau,a} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} u_i^\theta(a, x, \varepsilon_i) + \delta \sum_{a,x'} \phi_{a|x} V_{i,x'} f_{x'|a,x}, && \forall i, x
\end{aligned}$$

where $q \geq 0$ measures the degree of violation in the obedience condition, $\hat{L}$ and $\hat{U}$ are the lower and upper simultaneous confidence intervals of the conditional choice probabilities (CCPs) ϕ, and the other notations are the same as they appear in Theorem 5.

In our empirical application, $\psi_{\varepsilon|x} = \psi_\varepsilon$ and $\pi_{\tau|x,\varepsilon} = \pi_{\tau|\varepsilon}$ since idiosyncratic shock ε and signal τ are independent of state x. ψ_ε represents the joint distribution of two independent shocks $(\varepsilon_S, \varepsilon_D)$ each drawn from the standard logistic distribution. We discretize the standard logistic distribution using the best 10-point approximation in Kennan (2006). For the standard incomplete information, $\pi_{\tau|\varepsilon} = 1$ if $\tau = \varepsilon$ and 0 otherwise. For the fully robust specification, $\mathcal{T}$ is a singleton, and τ always takes that value regardless of ε. We construct the simultaneous confidence intervals of the CCPs with the plug-in sup-t implementation of Montiel Olea and Plagborg-Møller (2019). We bootstrap the multinomial logit model shown in Table A2 in Online Appendix A.3 at the market level to estimate the covariance matrix of the CCPs in the plug-in

sup-t implementation.

We grid search 400 ($= 20^2$) Halton draws in the 2-dimensional parameter space of θ and numerically approximate $\widehat{\Theta}_I$ by collecting $\widehat{\Theta}_I(c)$. Theoretically, we expect the standard incomplete information Markov correlated equilibrium identified set to contain the standard incomplete information Markov perfect equilibrium identified point, and the fully robust identified set to contain the standard incomplete information Markov correlated equilibrium identified set. We determine the criterion function threshold c to make our estimation results conform to this theoretical prediction. We set c for the standard incomplete information Markov correlated equilibrium identified set so that its confidence intervals (CIs) contain the CIs of the Markov perfect equilibrium estimates, and then set c for the fully robust identified set so that its CIs contain the standard incomplete information Markov correlated equilibrium identified set CIs.

We also estimate the model without the excluded variable, where we drop the excluded variable term from the flow profit (15) as follows:

$$u_i^\theta(a_{imt}, x_{mt}, \varepsilon_{imt}) = a_{imt}\{\theta_w^\top w_m + \theta_{ec}(1 - z_{imt}) + \theta_{i,ce}a_{jmt} + \varepsilon_{imt}\}.$$

We grid search 1 million ($\approx 16^5$) Halton draws in the 5-dimensional parameter space of $\theta = (\theta_w, \theta_{ec}, \theta_{S,ce}, \theta_{D,ce})$ and collect $\widehat{\Theta}_I(c)$ similarly to the Markov correlated equilibrium identified set estimation described above. Table A3 in Online Appendix A.3 shows the multinomial logit model used in the CCP covariance matrix estimation.[37]

A.2.2 Counterfactual Equilibrium

For the low scenario, we find an equilibrium that minimizes the expected number of active firms for given counterfactual parameters $\tilde{\theta}$ by solving the following problem:

$$\min_{\sigma, V} \quad \frac{1}{|\mathcal{X}|}\sum_x \{0 \cdot \phi_{0,0|x} + 1 \cdot (\phi_{0,1|x} + \phi_{1,0|x}) + 2 \cdot \phi_{1,1|x}\}$$

[37] Alternatively, one may find a projection interval of Θ_I by minimizing $p^\top\theta$ subject to the constraints above, where $p \in \mathbb{R}^{\dim(\theta)}$ is a direction vector (see, e.g., Kaido, Molinari and Stoye (2019)). The projection approach alleviates the need for a grid search over the entire parameter space but can be numerically difficult.

subject to

$$\sigma_{a|x,\varepsilon,\tau} \geq 0 \text{ for each } a \text{ and } \sum_{a} \sigma_{a|x,\varepsilon,\tau} = 1, \qquad \forall x, \varepsilon, \tau$$

$$\phi_{a|x} = \sum_{\varepsilon,\tau} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \qquad \forall a, x$$

$$\sum_{\varepsilon,\tau_{-i},a_{-i}} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} \partial v_i^{\tilde{\theta}}(a_i', a, x, \varepsilon_i) \leq 0, \qquad \forall i, x, \tau_i, a_i, a_i'$$

$$V_{i,x} = \sum_{\varepsilon,\tau,a} \psi_{\varepsilon|x} \pi_{\tau|x,\varepsilon} \sigma_{a|x,\varepsilon,\tau} u_i^{\tilde{\theta}}(a, x, \varepsilon_i) + \delta \sum_{a,x'} \phi_{a|x} V_{i,x'} f_{x'|a,x}, \qquad \forall i, x$$

where $0 \cdot \phi_{0,0|x} + 1 \cdot (\phi_{0,1|x} + \phi_{1,0|x}) + 2 \cdot \phi_{1,1|x}$ is the expected number of active firms at state $x \in \mathcal{X}$ for equilibrium CCPs $\phi_{a_1,a_2|x}$, $|\mathcal{X}|$ is the cardinality of $\mathcal{X}$, and the other notations are the same as they appear in Theorem 5. The objective function takes an unweighted average over x to obtain the (unconditional) expected number of active firms. The constraints, which are the same as Theorem 5, ensure that the resulting decision rule σ satisfies the Markov correlated equilibrium conditions. The specifications of $\psi_{\varepsilon|x}$ and $\pi_{\tau|x,\varepsilon}$ are the same as the estimation, as explained in Online Appendix A.2.1. We find an equilibrium that maximizes this optimization problem for the high scenario.

Let superscripts l and u indicate the lower and upper bounds of the 95% CIs of the structural estimates, respectively, e.g., $\hat{\theta}_{ec}^{l}$ is the lower bound of the 95% CIs of the entry cost parameter estimate. When we only consider the uncertainty of the competitive effects parameters, we set the counterfactual parameters to $\tilde{\theta} = (\hat{\theta}_w, \hat{\theta}_{S,v}, \hat{\theta}_{D,v}, 0.9 \cdot \hat{\theta}_{ec}, \hat{\theta}_{S,ce}^{l}, \hat{\theta}_{D,ce}^{l})$ in the low scenario under the counterfactual experiment of reducing the entry cost by 10%. $0.9 \cdot \hat{\theta}_{ec}$ indicates a 10% reduction of the original entry cost estimate $\hat{\theta}_{ec}$. We use the upper bounds of the 95% CIs of the competitive effects parameters, $(\hat{\theta}_{S,ce}^{u}, \hat{\theta}_{D,ce}^{u})$, instead of the lower bounds $(\hat{\theta}_{S,ce}^{l}, \hat{\theta}_{D,ce}^{l})$ in the high scenario. To additionally incorporate the uncertainty of the non-competitive effects parameters, we set $\tilde{\theta} = (\hat{\theta}_w^{l}, \hat{\theta}_{S,v}^{l}, \hat{\theta}_{D,v}^{l}, 0.9 \cdot \hat{\theta}_{ec}^{l}, \hat{\theta}_{S,ce}^{l}, \hat{\theta}_{D,ce}^{l})$ and $\tilde{\theta} = (\hat{\theta}_w^{u}, \hat{\theta}_{S,v}^{u}, \hat{\theta}_{D,v}^{u}, 0.9 \cdot \hat{\theta}_{ec}^{u}, \hat{\theta}_{S,ce}^{u}, \hat{\theta}_{D,ce}^{u})$ in the low and high scenarios, respectively.

A.3 Additional Estimation Results

Table A1: Binary logit parameter estimates in single-player markets

	Coef.	95% CI
Intercept	-7.071	$[-7.288, -6.826]$
Eating places (log)	0.935	$[0.880, 1.005]$
Starbucks outlets in county (log)	0.170	$[0.135, 0.208]$
Dunkin' outlets in county (log)	0.280	$[0.214, 0.354]$
Incumbent	8.157	$[7.997, 8.317]$
Number of Observations	124,215	
Log-likelihood	-8583.7	

Note: The table reports the coefficient estimates and 95% confidence intervals (CIs) of the binary logit model of a single-player market having the firm active or not. The reference choice is the firm not active. CIs are bootstrapped with resampling at the market level.

Table A2: Multinomial logit parameter estimates in two-player markets

	Starbucks	Dunkin'	Both
Intercept	-7.023	-7.117	-12.045
	$[-7.193, -6.878]$	$[-7.276, -6.934]$	$[-12.502, -11.719]$
Eating places (log)	0.977	0.632	1.422
	$[0.945, 1.015]$	$[0.598, 0.667]$	$[1.334, 1.516]$
Starbucks outlets in county (log)	0.184	-0.223	-0.159
	$[0.164, 0.205]$	$[-0.242, -0.205]$	$[-0.191, -0.128]$
Dunkin' outlets in county (log)	-0.161	0.417	0.203
	$[-0.177, -0.146]$	$[0.400, 0.438]$	$[0.175, 0.228]$
Starbucks is incumbent	8.365	0.408	8.644
	$[8.283, 8.485]$	$[0.114, 0.711]$	$[8.440, 8.868]$
Dunkin' is incumbent	0.795	8.563	8.369
	$[0.536, 1.043]$	$[8.476, 8.644]$	$[8.138, 8.603]$
Number of Observations	456,090		
Log-likelihood	-53066.5		

Note: The table reports the coefficient estimates and 95% confidence intervals (CIs) of the multinomial logit model of a two-player market having Starbucks active, Dunkin' active, both active, or both not active. The reference choice is both not active. CIs are bootstrapped with resampling at the market level.

Table A3: Multinomial logit parameter estimates without exclusion restriction

	Starbucks	Dunkin'	Both
Intercept	-7.095	-6.606	-12.407
	$[-7.186, -7.011]$	$[-6.690, -6.513]$	$[-12.797, -12.062]$
Eating places (log)	1.067	0.640	1.453
	$[1.035, 1.103]$	$[0.599, 0.677]$	$[1.352, 1.574]$
Starbucks is incumbent	8.665	0.491	8.797
	$[8.588, 8.760]$	$[0.215, 0.796]$	$[8.579, 9.018]$
Dunkin' is incumbent	0.731	9.215	8.937
	$[0.479, 0.980]$	$[9.132, 9.297]$	$[8.715, 9.153]$
Number of Observations	697,530		
Log-likelihood	-66013.2		

Note: The table reports the coefficient estimates and 95% confidence intervals (CIs) of the multinomial logit model of a two-player market having Starbucks active, Dunkin' active, both active, or both not active. The reference choice is both not active. CIs are bootstrapped with resampling at the market level.