Dynamic Equilibrium Strategies in Two-Sided Markets
Research Paper

Janik Bürgermeister[1], Martin Bichler[1], and Maximilian Schiffer[1]

Technical University of Munich, Domputer Science, Munich, Germany
{janik.buergermeister, martin.bichler}@tum.de

Abstract. We analyze when predatory pricing is rational in platform competition. We develop a multi-stage Bayesian model with continuous types and actions, expanding on earlier research that assumed complete information. By leveraging deep reinforcement learning, we identify equilibria in these multi-stage games. This framework also allows us to factor in risk aversion and cost asymmetries. For the complete-information scenario with two platforms, we demonstrate monotonicity in both single- and multi-stage games, which ensures a unique equilibrium and that learning dynamics reliably converge to the Nash equilibrium. In the Bayesian versions of the game, we numerically approximate a Bayes-Nash equilibrium. Our findings reveal that when there is uncertainty about competitors' costs, monopolies occur in roughly 60% of cases, even when the cost distributions are symmetric. In contrast, under complete-information monopolies only appear under asymmetries. Moreover, within the Bayesian framework, asymmetries tend to increase monopolization, while risk aversion has the opposite effect.

Keywords: Two-sided markets, Predatory Pricing, Bayesian multi-stage games, Learning in games

20th International Conference on Wirtschaftsinformatik
September 2025, Münster, Germany

1 Introduction

Two-sided markets revolve around platforms that connect two groups, where each group's value depends on the other. This cross-side network effect, encourages consumers to join if there are many sellers, and vice versa. To attract users, platforms often subsidize one group (e.g., social media users) to attract the revenue-generating side (e.g., advertisers). The analysis of platforms has drawn significant interest in Information Systems (Li & Wang 2024, Jung et al. 2019, Constantinides et al. 2018, Cennamo et al. 2018).

Predatory pricing has long been a key topic in oligopoly competition literature and more recently in platform competition and regulatory discussions (Ganesh 2024). It is defined as "reducing prices in the short run to eliminate competition and benefit in the long run" Atad & Yehezkel (2024). For example, Apple introduced its music streaming and video services with free or discounted offers to drive user adoption and expand its software ecosystem. Similarly, in 2015, Google launched Google Photos with unlimited free storage and no ads. Though unprofitable initially, these strategies aimed to attract users, establish market dominance, and enable future monetization.

Predation can drive competitors out of the market, potentially harming consumers. Antitrust laws, such as the Sherman Act in the U.S. and Article 102 of the TFEU, prohibit such abusive practices. A 2020 report by the U.S. House Judiciary Committee's Antitrust Subcommittee highlights predation as a significant risk in digital platform markets.[1]

Price-cost margins are commonly used to assess market power and detect predatory pricing. However, in two-sided markets, below-cost pricing may reflect cross-subsidization rather than predation. Measuring market power in such markets has led to varying approaches (Weyl 2010, Behringer & Filistrucchi 2015). Predatory pricing is generally recognized when total platform revenue falls below total costs, but under what conditions are such strategies rational?

A game-theoretical model can help to identify such conditions. While most platform competition studies use static games (Jullien et al. 2021), analyzing predatory strategies requires dynamic models. However, literature on dynamic platform competition remains limited (Fumagalli & Motta 2013, Amelio et al. 2020, Halaburda et al. 2020). These models are typically finite-horizon extensive-form games or infinite-horizon models seeking Markov Perfect equilibria (Halaburda et al. 2020, Atad & Yehezkel 2024). Even with complete information, finding equilibria analytically is challenging, often requiring numerical methods, without explicit equilibrium predictions.

Complete-information models simplify strategic interactions, as competing platforms rarely know each other's exact costs. This uncertainty complicates pricing decisions. Bayesian game theory addresses such uncertainty, modeling auctions and contests as single-stage Bayesian games with continuous types (values) and actions (bids). Introducing uncertainty can significantly alter equilibrium predictions compared to complete-information models (Sano 2012, Guler et al. 2016). Recently, *multi-stage Bayesian games* have been introduced as a generalization of Bayesian and stochastic games, incorporating continuous types and actions (Myerson & Reny 2020). While these models better capture strategic challenges, equilibrium characterization remains largely unexplored.

[1] https://www.govinfo.gov/content/pkg/CPRT-117HPRT47832/pdf/CPRT-117HPRT47832.pdf

Contributions

We analyze when predatory or exclusionary pricing emerges as an equilibrium in platform competition. To this end, we introduce a multi-stage platform competition model with continuous types and actions, departing from prior studies that rely solely on complete-information models. Our findings highlight the importance of uncertainty and show that monopolies frequently arise when platform cost distributions are symmetric.

We first analyze the complete-information platform competition model to examine Nash equilibrium uniqueness. Not all games can be learned (Milionis et al. 2022), and strict monotonicity is the only known condition ensuring a unique Nash equilibrium and convergence of various learning algorithms (Nagurney & Zhang 2012). We prove that the complete-information model by Armstrong & Wright (2007) (AW) is strictly monotone, even for asymmetric platforms. This is crucial, as it guarantees equilibrium learnability. Monotonicity is a strong condition that rarely holds (Bichler et al. 2023).

For extensions to more platforms or multi-stage Bayesian models with bankruptcy, proving strict monotonicity is challenging. While equilibrium existence in multi-stage games is established (Myerson & Reny 2020), equilibrium uniqueness and computation remains largely unexplored. To address this, we employ deep reinforcement learning (DRL)(Pieroth et al. 2024), enabling ex-post verification of approximate equilibria. Notably, we consistently find the same equilibrium across different hyperparameters, suggesting possible uniqueness.

In our numerical experiments we observe fast converges to the analytical equilibrium in the complete-information AW model. In the symmetric case, where both platforms have identical costs, predatory pricing does not arise. However, for asymmetric platforms, predatory pricing emerges as an equilibrium, even with a slight advantage for the predator. In some cases, predators achieve positive profits in the first stage before securing monopoly status, qualifying as exclusionary pricing.

In multi-stage Bayesian models where platforms face cost uncertainty, monopolies emerge even with symmetric priors. For risk-neutral platforms, monopolies arise in about 60% of cases. Asymmetries further increase monopoly prevalence. For strategic decisions, platform managers are likely risk-averse. We observe that risk aversion raises prices and reduces monopoly formation in equilibrium.

This paper is among the first to apply DRL for equilibrium analysis in multi-stage games. A key objective is to assess whether these techniques effectively identify equilibria in general platform competition models. While strict monotonicity holds in the symmetric complete-information model, proving similar properties for models with bankruptcy or multi-stage Bayesian settings is far more challenging.

Remarkably, DRL consistently converges to equilibrium across all models, highlighting its potential for equilibrium analysis. These advanced learning techniques facilitate equilibrium analysis in models previously beyond analytical reach. This approach also enables the study of uncertainty and risk aversion in strategic decision-making.

2 Related Literature

Next, we provide a literature overview on platform competition and multi-stage games.

Research on two-sided market competition has grown significantly. For a broad overview, we refer to Jullien et al. (2021). Seminal models include Rochet & Tirole (2003), which focuses on pricing strategies and welfare, and Armstrong (2006), which highlights subsidizing one side to attract the other. Armstrong & Wright (2007) (AW) extends this to indirect network effects and has since influenced a wide literature (Cennamo et al. 2018, Constantinides et al. 2018, Jung et al. 2019, Li & Wang 2024).

Several empirical studies have examined predatory pricing in platform markets. Bhattacharjea (2018) uses case studies of Uber and Ola to illustrate challenges to traditional competition law. Connor (2022) discusses Portugal's ban on below-cost pricing, emphasizing the tension between competition and market distortion. Behringer & Filistrucchi (2015) demonstrate that a monopolist may price below marginal cost on one side, indicating that one-sided pricing tests may misidentify predatory behavior.

Amelio et al. (2020) investigates how incumbent platforms may engage in exclusionary pricing to deter the entry of competitors in a complete-information extensive-form game. An incumbent platform sets its prices before a potential entrant decides whether to enter, observing the set prices and incurring a fixed entry cost. Agents on both sides of the market observe all prices and decide which platform to join. All players are assumed to have full knowledge of the game structure, the actions taken, and the payoffs.

Atad & Yehezkel (2024) models infinitely repeated platform competition, where the platform that dominated in the previous period becomes the incumbent with a focal advantage. It examines the impact of antitrust policies that prohibit predatory pricing, either symmetrically (for both platforms) or asymmetrically (only for the incumbent). The main insight is that symmetric regulation fails to mitigate inefficient incumbency and even reduces consumer surplus, while asymmetric regulation increases consumer surplus and can improve welfare in a static market size. The authors employ a Markov Perfect Equilibrium, and incorporate stochastic platform quality and market size variations.

Myerson & Reny (2020) introduced multi-stage games with continuous type and action spaces. Players act over several stages, receiving private signals. Nature may choose states probabilistically at each stage, influenced by the game's history. Strategies map signals to probability distributions over actions. Such multi-stage games are a generalization that includes Bayesian games, games with perfect information, and finite-horizon stochastic games which include all finitely-repeated games. Multi-stage games emphasize the sequence of moves and the evolution of the game state, in contrast to the repeated games with incomplete information described in (Aumann et al. 1995).

Multi-stage games are an excellent model for the strategic situation in platform markets, considering the uncertainties that platforms face. Unfortunately, equilibrium problems in such multi-stage games are systems of non-linear differential equations for which exact analytical solutions often remain elusive. As a result, we don't have explicit equilibrium predictions. Recently, Pieroth et al. (2024) introduced a method to compute equilibrium based on DRL. The method includes an equilibrium verifier that certifies whether a strategy profile found is an approximate Bayes-Nash equilibrium.

3 The Model

We extend the model of AW to more than two platforms, and introduce asymmetries in platforms costs. Furthermore, we show that the two-player case's baseline model is strictly monotonous. This guarantees the existence and uniqueness of a Nash equilibrium and makes the equilibrium learnable for gradient-type algorithms. Finally, we introduce extensions to incorporate multiple stages and incomplete information.

3.1 Extension of the AW Model to Multiple Platforms

Two groups, buyers (B) and sellers (S), interact via N_P different platforms. The active players of the game are the platforms. Each platform might have different costs for buyers $f_{j,B}$ and for sellers $f_{j,B}$. They seek to find prices $p_{j,B}, p_{j,S}$ to maximize their profit $\pi_j(p)$, given their demand from each market-side ($n_{j,B}, n_{j,S}$).

$$\pi_j : \mathbb{R}^{N_P \times 2} \to \mathbb{R} \quad \pi_j(p) = \sum_{k \in \{B,S\}} (p_{j,k} - f_{j,k}) n_{j,k}(p). \tag{1}$$

We write $p_j = (p_{j,B}, p_{j,S})$ and use p_{-j} for the vector where we dropped the prices of platform j. The analytical extension of the demand functions $n_{j,k}$ is more involved.

Armstrong & Wright (2007) use the Hotelling model to derive the demand. In this model, platforms are located at either end of a unit interval, where individual buyers and sellers are uniformly distributed on. The natural higher-dimensional extensions are simplexes. In particular, we model each buyer i with a location $r_{B_i} \in \mathbb{R}^{N_P}$ relative to the platforms. These locations are assumed to be uniformly i.i.d. on the unit simplex $\Delta(N_P) = \left\{ x \in \mathbb{R}^{N_P}, \sum_{j=1}^{N_P} x_j = 1, x_j \geq 0 \right\}$.

Each standard basis vector $e_j \in \mathbb{R}^{N_P}$ corresponds to a platform. The locations give rise to distances $d_{j,B_i} := ||r_{B_i} - e_j||_1$ between buyers and a platform. The terms for the seller are defined analogously (r_{S_k}, d_{j,S_k}). When a member i of group k subscribes to platform j with distance $x = d_{j,k_i}$, they get a group-specific utility of

$$U_{j,k}(x) = u_{0,k} - p_{j,k} - t_k x + b_k n_{j,l}. \tag{2}$$

The network effects are modeled by $b_k n_{j,l}$, where $n_{j,l}$ is proportional to the number of agents from the other group on platform j. The transport costs or preferences of a platform for a group k member are modeled by $t_k x$.

In what follows, we assume that buyers and sellers single-home, i.e., a group member only joins the platform where they receive the highest utility. It follows that

$$n_{j,k} = \mathbb{P}(U_{j,k_i} > \max_{l \neq j} U_{l,k_i}) \tag{3}$$

where the randomness stems from the location distribution.

To derive the theoretical results as in Armstrong & Wright (2007), the following assumptions are necessary: The first assumption is that transport costs outweigh the benefits from the other group, and the second assumption ensures that the platform's profit functions are concave w.r.t. their own prices.

Assumption 1 *1.* $t_S > b_S$, $t_B > b_B$, *2.* $4t_S t_B > (b_S + b_B)^2$.

For two platforms the derivation of closed form solution for the demand is illustrated in Armstrong & Wright (2007). However, for more platforms a analytical solution is intractable. Therefore, we resort to numerical methods for these cases. The market is in equilibrium if no platform can increase its profit by unilaterally deviating from its prices.

Definition 1. *The price vector* p^* *is a Nash equilibrium (NE) if and only if*

$$\pi_j(p_j, p^*_{-j}) \leq \pi_j(p^*) \; \forall j \in [N_P], \; p_j \in \mathbb{R}^2 \, .$$

The following proposition provides the equilibrium prices $p_{i,k}$.

Proposition 1 (Armstrong & Wright (2007)). *Assume assumptions 1 holds. Furthermore, assume that platforms have symmetric costs, i.e.,* $f_{j,k} = f_{i,k} = f_k$*. Then, the two platforms offer the same pair of prices,* p_S *and* p_B*, and half the agents from each group join each platform. If* $f_S + t_S \geq b_B$ *and* $f_B + t_B \geq b_S$*, equilibrium prices are* $p_S = f_S + t_S - b_B, \; p_B = f_B + t_B - b_S$*.*

This theoretical result can be extended to asymmetric costs. However, the representation of the equilibrium prices is lengthy, and we omit it for brevity.

3.2 Convergence of Gradient-Based Algorithms

We now examine the convergence properties of gradient-based methods in the previous model. We show that these methods converge to unique equilibrium prices when applied by each player.

We begin by considering a game with N players, where each player's continuously differentiable utility is given by $\pi_i(x_i, x_{-i})$. Gradient-based algorithms assume that players follow the gradient of their utility function with respect to their own actions. The corresponding dynamical system is given by:

$$\dot{x} = F(x) = -(\nabla_{x_j} \pi_j(x))_{j \in [n]}, \quad x(0) = x_0 \in K, \tag{4}$$

where $F(x) : K \mapsto \mathbb{R}^n$. We need to establish a connection between the dynamical system and the original game. Note that, generally, an equilibrium of eq. (4) ($\dot{x} = 0$) must not coincide with a Nash equilibrium of the corresponding game.
Games with a continuous action space can be modeled as a *variational inequality* (VI).

Definition 2. *Let* $F : K \to \mathbb{R}^n$*, where* $K \subseteq \mathbb{R}^n$ *is a non-empty closed set. A (finite) variational inequality* $VI(K, F)$ *is the problem of finding a vector* $x^* \in \mathbb{R}^n$ *such that*

$$F(x^*)^T (x - x^*) \geq 0 \quad \forall x \in K. \tag{5}$$

We call x^* *the solution and* K *the feasible set of the VI.*

For a game, the corresponding VI is defined via $F(x) = -(\nabla_{x_j} \pi_j(x))_{j \in [n]}$. Importantly, if $F(x)$ is strictly monotone, the solution to the variational inequality is unique and corresponds to the unique Nash equilibrium of the game (Rosen 1965).

Definition 3. *A function $F : \mathbb{R}^n \to \mathbb{R}^n$ is called monotone if*

$$\langle F(x) - F(y), x - y \rangle \geq 0 \, \forall x, y \in \mathbb{R}^n$$

and strictly monotone if the inequality holds strictly.
A game is called (strictly) monotone if $F(x) = (-\nabla_{x_i} \pi_i)_{i \in N}$ is (strictly) monotone.

After establishing uniqueness and a correspondence between the dynamical system and the original game, we address convergence. For this, we need to introduce the concept of a *monotone attractor*, which is a point in a dynamical system to which nearby trajectories will converge in a monotonic fashion. The strict monotonicity of the game is a sufficient condition for the convergence of gradient-based methods to the equilibrium:

Proposition 2. *(Nagurney & Zhang 2012)*
If $F(x)$ is strictly monotone, then x^* is a global strictly monotone attractor for (4).

We are now able to formulate our first result.

Theorem 1. *Let the platform's profit be given as in Equation* (1)*. If Assumption 1 holds, then the corresponding game is strictly monotone for two platforms.*

This ensures that gradient-based algorithms converge to the unique equilibrium price (Mertikopoulos & Zhou 2019). This result holds even for a more general model where costs may vary across groups and platforms, i.e., $f_{i,k} \neq f_{j,k}$. Consequently, the costs affect only the equilibrium price, not its stability.

Proof (Proof Sketch). We use a Lemma from Nagurney & Zhang (2012) and show that the game hessian is positive definite.

3.3 Multi-Stage Platform Competition

We now study a game where players engage in the single-stage game repeatedly over a finite number of rounds. Their total profit is given by the sum of the profits in each stage.

If the stage game has a unique Nash equilibrium, then in the finitely repeated version, the unique subgame perfect NE (SPNE) involves players playing the stage game's NE in every period (Benoit et al. 1984). We show that without access to previous prices, the multi-stage platform game is also strictly monotone implying stability of the NE.

Corollary 1. *The platform competition game is strictly monotone under Assumptions 1 and 2. Consequently, playing the single-stage equilibrium at each stage of the multi-stage game constitutes the unique Nash equilibrium, and gradient-type algorithms are guaranteed to converge to it.*

Proof (Proof Sketch). The total profit for each platform in the multi-stage game is the sum of profits across all stages. The game Hessian is a block-diagonal matrix, with each block corresponding to a single-stage game. Since each block is positive definite by the proof of Theorem 1, it follows that the entire game Hessian is positive definite, implying strict monotonicity of the game.

Furthermore, if platforms have access to previous prices the unique Nash equilibrium remains playing the single-stage Nash equilibrium at each stage, which can be proven via backward induction. However, this result is only true, if platforms do not drop out due to losses they make. Now, we extend the model to account for the possibility of bankruptcy, where platforms exit the market if their profit falls below a predefined credit limit at any stage. A credit limit of x_i means that platform i drops out if its profit at any stage is less than x_i. Players observe which competitors remain active at the start of each round. Let $\mathcal{P}^{(t)} \subseteq [N_P]$ denote the set of active players at stage t

$$\mathcal{P}^{(t+1)} = \mathcal{P}^{(t)} \cap \{i : \pi_{it} \geq x_i\}.$$

Players can exit the competition but cannot reenter. If a platform becomes a monopolist, it receives a reward of r_M, reflecting monopoly prices that are sufficiently high for a platform to prefer the monopoly outcome.

Let $M_i^{(t)} \in \{0, 1\}$ indicate whether player i is a monopolist at stage t.

$$M_i^{(t+1)} = \mathbb{I}_{\{\pi_j^{(t)} < x_j, \forall j \neq i, j \in \mathcal{P}^{(t)}\}} \cdot \mathbb{I}_{\{\pi_i^{(t)} \geq x_i\}},\ M_i^{(0)} = 0.$$

Monopolists retain their status until the game's end. The single-stage profit function adjusts accordingly:

$$\pi_{it}^M(p) = \begin{cases} r_M, & \text{if } M_i^{(t)} = 1 \text{ and } i \in \mathcal{P}^{(t)} \\ \sum_{k \in \{B,S\}} (p_{i,k}^{(t)} - f_{i,k}) n_{i,k}(p^{(t)}), & \text{if } M_i^{(t)} = 0 \text{ and } i \in \mathcal{P}^{(t)} \\ 0, & \text{if } i \notin \mathcal{P}^{(t)} \end{cases}$$

The transition from oligopoly to a monopoly in later stages introduces discontinuities into the objective function, complicating the analysis of monotonicity of the finite-horizon repeated game, which we refer to as repeated platform competition with bancruptcy. Not much is known about convergence of learning algorithms to equilibrium in such repeated complete-information games.

3.4 Multi-Stage Bayesian Platform Competition

So far, we have considered only complete information games. For our numerical analysis, we define a multi-stage game within the framework of Myerson & Reny (2020) that allows for bankruptcy and incomplete information. Players observe their own marginal costs $f_i \sim F_i$, have access to all prices from previous rounds, and know which players are currently active. However, the generality of the model comes at a cost. While the existence of equilibria in multi-stage games has been established, the problem of finding such equilibria remains largely unexplored in this general game class.

Recent advances in DRL, particularly algorithms like REINFORCE and Proximal Policy Optimization (PPO), have successfully identified equilibria in multi-stage games (Pieroth et al. 2024). Building on this approach, we numerically compute equilibria.

4 Numerical Analysis

Our theoretical analysis focused on gradient-based algorithms, which converge to equilibrium in monotone games. In bankruptcy settings, actions must be state-dependent, as players adjust pricing based on the number of active competitors. We examine DRL as a suitable approach, as it can capture dependencies on past prices and can be used for Bayesian games where platform characteristics, such as marginal costs, are uncertain.

4.1 Evaluation

First, we discuss the metrics we use to analyze convergence of DRL algorithms to equilibrium. The ex-ante *utility loss* measures the loss of an agent by not playing a best-response to the opponents' prices p_{-i} (Srinivasan et al. 2018, Brown et al. 2018)

$$\ell_i(p_i, p_{-i}) = \sup_{p_i'} \pi_i(p_i', p_{-i}) - \pi_i(p_i, p_{-i}). \tag{6}$$

Note that a price vector p is an ε-NE if and only if $\ell_i(p_i, p_{-i}) \leq \varepsilon$ for all platforms i. In more complex environments where an analytical solution is unavailable, we employ an adaptive grid search to compute the best response to the opponents' prices. This allows us to always compute wether learned strategies are close to an equilibrium.

4.2 Single-Stage, Complete-Information Game

In our first set of results, we analyze the single-stage complete information model with two to three possibly asymmetric platforms, using PPO and REINFORCE. For two platforms we investigate convergence in terms of the mean and standard deviation of L_2^{avg} and ℓ_i across ten runs. We observe convergence to an ϵ-NE ($\epsilon \leq 0.004$) within 400 iterations. Furthermore, learned prices closely align with the analytical predictions. Extensions to three platforms are difficult to analyze analytically. We verify ex-post convergence to ϵ-Nash equilibria by computing the best response to opponents' prices via an adaptive grid search. We observe convergence to an ϵ-NE with a higher ϵ ($\epsilon \leq 0.04$) within the same number of iterations.

4.3 Multi-Stage, Complete-Information Game

This section presents numerical simulations of the multi-stage game extension to analyze how game variations influence platform behavior and equilibrium. We focus on three key factors: price transparency, bankruptcy, and asymmetric credit limits. Equilibrium was computed across environments with up to ten stages, where T denotes the number of stages. As the number of stages increases, the complexity of the environment rises, requiring more iterations for convergence. Each analysis runs for $600T$ iterations. Results are consistent, independent of T. If a monopoly arose, then this happened already in the second stage. Since two-stage games already capture the essential dynamics, we focus on them subsequently. Next, we describe the main insights from the simulations qualitatively. The detailed numerical results were omitted due to place constraints.

Symmetric Multi-Stage Game With symmetric costs and credit limits, we do not observe predatory pricing behavior in the complete-information model. Players learn the equilibrium strategies corresponding to a single-stage game in each stage as predicted by Corollary 1. Transparency accelerates convergence towards the single-stage Nash equilibrium, significantly reducing the required iterations. Again, no tacit collusion or signaling arises, indicating that transparency does not lead to coordinated pricing strategies. Instead, platforms integrate competitor pricing into decision-making, sustaining competitive behavior without collusive outcomes.

Multi-Stage Game with Asymmetric Credit Limits Now, we investigate how advantages in credit limit and cost influence the learned prices. We start with asymmetric credit limits, where one platform is assigned a higher credit limit than the other.

Predation only occurs if the monopoly reward is sufficiently high, making the monopoly an attractive alternative to the oligopoly. When the disparity in credit limits is significant, consistent predatory pricing occurs. We observe that the stronger player prices aggressively (below marginal costs) on both sides of the market, capturing all of the demand in the first stage. This drives the other platform out of the market and leads to a negative profit for the predator in the first stage. Afterward, the predator can offset his losses from the first stage in the second stage by enjoying monopoly benefits.

Multi-Stage Game with Asymmetric Marginal Costs The costs of platform 1 are fixed and we vary the costs of platform 0. Both platforms have a credit limit of zero. A substantial marginal cost advantage enables persistent predation. In this case, both players deviate from single-stage equilibrium pricing. The weaker platform prices near its marginal costs on both sides, whereas the stronger platform sets prices below the weaker player's marginal costs. We observe that seller prices are lower and buyer prices are higher than in the single-stage equilibrium. Overall, predation in the first stage reduces the total price level. Consequently, the first-stage profits of the predator are reduced compared to the single-stage equilibrium. If the cost advantage is large, the price level remains above the total marginal cost level of the predator. Thus, the predator does not incur losses. A regulation enforcing total prices above total costs cannot prevent exclusionary pricing in this case.

4.4 Multi-Stage Bayesian Game

Next, we examine the impact of uncertainty regarding competitors' costs. In our multi-stage Bayesian game analysis, platform marginal costs are drawn from a uniform distribution at the game's start and remain constant thereafter.

Single Stage As a baseline, we consider a symmetric single-stage game where costs are drawn from uniform distributions. Both players share identical cost distributions, but distinct distributions for each side of the market. Note that instead of learning prices $p_i \in \mathbb{R}^2$, players now learn equilibrium price functions $p_i(f_i) : \mathbb{R}^2 \to \mathbb{R}^2$. The learned strategies align with the complete information game equilibrium. The expected equilibrium prices closely match the corresponding complete-information equilibrium. The strategies are symmetric across players, and prices increase linearly as sampled costs rise. On the sellers' side, prices are proportional to $f_{i,B} + f_{i,S}$, whereas on the buyers' side, prices are proportional to $\alpha f_{i,B} + \beta f_{i,S}$, where $\alpha \neq \beta$ and $\alpha, \beta > 0$.

The Symmetric Multi-Stage Model We extend the previous model to multiple stages, where the same symmetric marginal cost distributions are used as before. Each platform has a credit limit of 0, meaning players are eliminated if they incur a loss.

We observe that players are equally likely to become a monopolist, with an approximate 30-30-40 split between monopoly player 0, monopoly player 1, and oligopoly. With a small probability of 2%, both players exit the market.

While first-stage strategies are asymmetric, overall profits remain symmetric. Pricing strategies depend on each platform's marginal costs, with prices increasing as costs rise.

Platform 1 buyer-side prices exceed marginal costs, while seller-side prices fall below them. Conversely, platform 2 buyer-side prices fall below marginal costs, while seller-side prices exceed them. In comparison to the single-stage equilibrium, prices are reduced. For example, while prices for sellers were ranging from 1.2 to 1.3 in the single-stage equilibrium, they range from 0.7 to 1.3 in the first stage of the multi-stage equilibrium with dropout. In the second-stage oligopoly, learned prices approximate those from the corresponding single-stage game but exhibit higher variance.

The Asymmetric Multi-Stage Model Next, we investigate the Bayesian game with asymmetric costs, where one platform has an advantage on both market sides, with expected costs aligning with Section 4.3. Platform 0 has lower costs in approximately 96% of games, while platform 1 benefits in only 4%. These values serve as upper bounds on monopoly likelihood, as a platform must have lower costs to dominate.

Platform 0 becomes a monopolist in 75% of games, platform 1 in 2%, and 23% of games remain oligopolistic in the second stage. Platform 0's average prices are lower than platform 1's. Due to asymmetric costs, platforms adopt asymmetric strategies, with first-stage strategies dependent on marginal costs.

Risk-Aversion Risk neutrality may be too strong an assumption for strategic decisions, making it valuable to examine how risk aversion influences predation.

We incorporate risk aversion into platform utility functions using a standard exponential utility function: $\hat{\pi}_j = (1 - \exp(-\alpha \pi_j))/\alpha$ where α represents the coefficient of risk aversion. This corresponds to a Constant Absolute Risk Aversion (CARA) utility function. Higher α values indicate greater risk aversion.

Experimental results ($\alpha \in [0, 2]$) show that increasing α reduces the fraction of monopolies per player. For example, the monopoly rate drops from 30% at $\alpha = 0$ to approximately 15% for $\alpha \approx 1$. Additionally, the probability of both players exiting falls below 1%. Moreover, as risk aversion increases, asymmetries between players diminish, and prices align more closely with marginal costs.

5 Discussion

We now discuss the limitations of our model and outline avenues for future research.

First, for tractability, we assume credit limits are exogenously given. While this simplifies the analysis, it abstracts from the strategic and financial mechanisms that may influence credit availability. Future work could endogenize credit constraints by linking them to financing conditions or firm-specific characteristics. This extension would also offer a more grounded explanation for why start-ups or smaller platforms are particularly vulnerable to predation by larger, well-established firms.

Second, we assume that market entry does not occur, reflecting the presence of substantial entry barriers, notably network effects. For instance, launching a new communication platform to compete with an incumbent like WhatsApp is extremely challenging due to the difficulty of displacing existing user bases. Future research could examine how limited or costly entry affects predatory dynamics. However, if entry barriers are high, predation strategies would largely persist unchanged.

Third, we do not consider models where demand in future periods could be influenced by past demand. While introducing such dynamics makes the model more realistic, the core strategic outcome remains similar. For example, a dominant firm with a cost advantage may set prices that a weaker rival cannot match, causing the latter's demand to erode over time and eventually leading to market exit. In our model, this outcome occurs in a single period, offering a more immediate and tractable representation. Thus, while an extended model with demand persistence may better capture gradual decline, it ultimately converges to similar qualitative conclusions, albeit over a longer horizon.

6 Conclusion

We extend a static platform competition model into a multi-stage game with continuous types and actions, departing from prior work that relies on complete-information, single-stage models. We identify conditions under which predatory strategies are rational.

For asymmetric single- and multi-stage competition models without bankruptcy, we prove strict monotonicity, ensuring equilibrium uniqueness and convergence of gradient-based algorithms. Numerically, we demonstrate that DRL converges to approximate equilibrium in extensions with additional platforms and multi-stage Bayesian models.

In complete-information multi-stage models, exclusionary pricing emerges as an equilibrium only in asymmetric settings. When costs are asymmetric, the stronger platform prices above total costs, implying that even regulatory price floors may not prevent monopolization in later stages.

With incomplete information, monopolies arise even in symmetric models. Depending on cost draws, one or both platforms may incur losses in the first stage and exit. In equilibrium, each platform faces a 30% chance of loss, with a 2% probability that both go bankrupt. While risk aversion reduces monopoly prevalence, the effect persists, and cost asymmetries further increase monopoly likelihood.

Future research could examine how switching costs and multi-homing affect the success of predatory platform pricing. Additionally, exploring alternative regulatory interventions remains an important direction.

References

Amelio, A., Giardino-Karlinger, L. & Valletti, T. (2020), 'Exclusionary pricing in two-sided markets', *International Journal of Industrial Organization* **73**, 102592.

Armstrong, M. (2006), 'Competition in two-sided markets', *The RAND Journal of Economics* **37**(3), 668–691.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/j.1756-2171.2006.tb00037.x*

Armstrong, M. & Wright, J. (2007), 'Two-sided Markets, Competitive Bottlenecks and Exclusive Contracts', *Economic Theory* **32**(2), 353–380.
URL: *https://link.springer.com/10.1007/s00199-006-0114-6*

Atad, O. & Yehezkel, Y. (2024), 'Regulating Platform Competition in Markets with Network Externalities: Will Predatory Pricing Restrictions Increase Social Welfare?', *The Journal of Industrial Economics* **72**(3).
URL: *https://onlinelibrary.wiley.com/doi/abs/10.1111/joie.12385*

Aumann, R. J., Maschler, M. & Stearns, R. E. (1995), *Repeated games with incomplete information*, MIT press.

Behringer, S. & Filistrucchi, L. (2015), 'Areeda–turner in two-sided markets', *Review of Industrial Organization* **47**(2), 131–148.

Benoit, J.-P., Krishna, V. et al. (1984), 'Finitely repeated games'.

Bhattacharjea, A. (2018), 'Predatory pricing in platform competition: The case of uber and ola in india', *Competition and Regulation in Network Industries* **19**(1-2), 114–131.

Bichler, M., Lunowa, S. B., Oberlechner, M., Pieroth, F. R. & Wohlmuth, B. (2023), 'On the Convergence of Learning Algorithms in Bayesian Auction Games'.

Brown, N., Lerer, A., Gross, S. & Sandholm, T. (2018), 'Deep counterfactual regret minimization', *arXiv preprint arXiv:1811.00164* **97**, 793–802.

Cennamo, C., Ozalp, H. & Kretschmer, T. (2018), 'Platform architecture and quality trade-offs of multihoming complements', *Information Systems Research* **29**(2), 461–478.

Connor, K. (2022), 'Portugal bans below-cost pricing by online travel platforms', *MLex Market Insight* . Accessed: 2025-01-29.

Constantinides, P., Henfridsson, O. & Parker, G. G. (2018), 'Introduction—platforms and infrastructures in the digital age'.

Fumagalli, C. & Motta, M. (2013), 'A simple theory of predation', *The Journal of Law and Economics* **56**(3), 595–631.

Ganesh, A. (2024), 'Predatory pricing in platform markets: a modified test for firms within the scope of Article 3 of the DMA and super-dominant platform firms under Article 102 TFEU', *European Competition Journal* **0**(0).
URL: *https://doi.org/10.1080/17441056.2024.2428032*

Guler, K., Bichler, M. & Petrakis, I. (2016), 'Ascending combinatorial auctions with risk averse bidders', *Group Decision and Negotiation* **25**, 609–639.

Halaburda, H., Jullien, B. & Yehezkel, Y. (2020), 'Dynamic competition with network externalities: how history matters', *The RAND Journal of Economics* **51**(1), 3–31.

Jullien, B., Pavan, A. & Rysman, M. (2021), Two-sided markets, pricing, and network effects, *in* 'Handbook of Industrial Organization', Vol. 4, Elsevier, pp. 485–592.
URL: *https://linkinghub.elsevier.com/retrieve/pii/S1573448X21000078*

Jung, D., Kim, B. C., Park, M. & Straub, D. W. (2019), 'Innovation and policy support for two-sided market platforms: Can government policy makers and executives optimize both societal value and profits?', *Information Systems Research* **30**(3), 1037–1050.

Li, Z. & Wang, G. (2024), 'Regulating powerful platforms: Evidence from commission fee caps', *Information Systems Research* .

Mertikopoulos, P. & Zhou, Z. (2019), 'Learning in games with continuous action sets and unknown payoff functions', *Mathematical Programming* **173**(1), 465–507.

Milionis, J., Papadimitriou, C., Piliouras, G. & Spendlove, K. (2022), 'Nash, conley, and computation: Impossibility and incompleteness in game dynamics', *arXiv preprint arXiv:2203.14129* **120**(41), e2305349120.

Myerson, R. B. & Reny, P. J. (2020), 'Perfect Conditional ϵ-Equilibria of Multi-Stage Games With Infinite Sets of Signals and Actions', *Econometrica* **88**(2), 495–531.

Nagurney, A. & Zhang, D. (2012), *Projected dynamical systems and variational inequalities with applications*, Vol. 2, Springer Science & Business Media.

Pieroth, F. R., Kohring, N. & Bichler, M. (2024), 'Equilibrium computation in multi-stage auctions and contests', *arXiv preprint arXiv:2312.11751* .

Rochet, J.-C. & Tirole, J. (2003), 'Platform Competition in Two-Sided Markets', *Journal of the European Economic Association* **1**(4), 990–1029.
URL: *https://academic.oup.com/jeea/article/2280902/Platform*

Rosen, J. B. (1965), 'Existence and uniqueness of equilibrium points for concave n-person games', *Econometrica: Journal of the Econometric Society* pp. 520–534.

Sano, R. (2012), 'Non-bidding equilibrium in an ascending core-selecting auction', *Games and Economic Behavior* **74**(2), 637–650.

Srinivasan, S., Lanctot, M., Zambaldi, V., Pérolat, J., Tuyls, K., Munos, R. & Bowling, M. (2018), 'Actor-critic policy optimization in partially observable multiagent environments', *Advances in neural information processing systems* **31**.

Weyl, E. G. (2010), 'A price theory of multi-sided platforms', *American Economic Review* **100**(4), 1642–1672.