NBER WORKING PAPER SERIES

DYNAMIC PRICE COMPETITION:
THEORY AND EVIDENCE FROM AIRLINE MARKETS

Jose M. Betancourt
Ali Hortaçsu
Aniko Oery
Kevin R. Williams

Working Paper 30347
http://www.nber.org/papers/w30347

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2022, Revised April 2023

The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research. We thank the anonymous airline for giving us access to the data used in this study. Under the agreement with the authors, the airline had “the right to delete any trade secret, proprietary, or Confidential Information” supplied by the airline. All authors have no material financial relationships with entities related to this research. We gratefully acknowledge financial and computational support from the Yale School of Management. We thank Alessandro Bonatti, Chris Conlon, Jim Dana, Volker Nocke, Ariel Pakes, Ben Polak, Michael Riordan, Satoru Takahashi, Mike Whinston, and several conference and seminar participants.

Dynamic Price Competition: Theory and Evidence from Airline Markets
Jose M. Betancourt, Ali Hortaçsu, Aniko Oery, and Kevin R. Williams
NBER Working Paper No. 30347
August 2022, Revised April 2023
JEL No. C70,C73,D21,D22,D43,D60,L13,L93

ABSTRACT

We introduce a model of dynamic pricing in perishable goods markets with competition and provide conditions for equilibrium uniqueness. Pricing dynamics are rich because both own and competitor scarcity affect future profits. We identify new competitive forces that can lead to misallocation due to selling units too quickly: the Bertrand scarcity trap. We empirically estimate our model using daily prices and bookings for competing U.S. airlines. We compare competitive equilibrium outcomes to those where firms use pricing heuristics based on observed internal pricing rules at a large airline. We find that pricing heuristics increase revenues (4-5%) and consumer surplus (3%).

Jose M. Betancourt
Yale University
jose.betancourtvalencia@yale.edu

Ali Hortaçsu
Kenneth C. Griffin Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
hortacsu@uchicago.edu

Aniko Oery
Yale School of Management
165 Whitney Ave
New Haven, CT 06520
aniko.oery@yale.edu

Kevin R. Williams
Yale School of Management
165 Whitney Avenue
New Haven, CT 06520
and NBER
kevin.williams@yale.edu

1 Introduction

Firms dynamically adjust prices in many markets with perishable inventory. Examples range from seats on airplanes and trains, tickets for entertainment events, reservations for cruises, to inventory in retailing. In these markets, prices may adjust for a variety of reasons. First, prices reflect changing opportunity costs due to demand uncertainty—in the presence of scarcity, the cost of selling a unit of inventory today depends on a firm's ability to sell it in the future. Second, demand may change over time in predictable ways. If consumers with high willingness to pay tend to arrive late, firms have an incentive to save inventory. Finally, in all aforementioned examples, prices may adjust in response to competitive interactions. Yet, much of the theoretical and empirical literature on dynamic pricing in perishable goods markets has abstracted from competition entirely. In this paper we introduce a framework to study dynamic price competition. We explore how dynamic price competition can adversely affect market efficiency and how alternative pricing mechanisms can improve overall welfare.

In order to understand how competition can generate market inefficiencies, consider two competing airlines selling seats for a given departure date. Demand is uncertain. A social planner would initially set high prices because of the option value of allocating seats to future consumers with higher valuations. However, when competing firms engage in dynamic pricing, intense competition can lead to inefficiently low prices early on. As a result, there is an over-provision of bookings far from departure and an under-provision of bookings close to departure. We call this effect the *Bertrand scarcity trap*. Although this inefficiency can arise even when underlying demand is constant, it can worsen when late-arriving buyers have higher valuations.

We formalize this intuition and empirically measure the welfare consequences of firms engaging in dynamic pricing by building a rich model of dynamic price competition. We provide sufficient conditions for equilibrium existence and uniqueness and for convergence to a system of ordinary differential equations (ODEs). This characterization allows us to derive new theoretical insights on dynamic price competition and to simulate market outcomes without explicitly solving stage game equilibria. We demonstrate the usefulness of this result in our empirical application that involves billions of stage games. We apply our model to one of the best known examples of dynamic pricing under competition—the U.S. airline industry. We use novel data that provide daily bookings and prices for multiple, competing U.S. airlines. With demand estimated, we compare market outcomes under the competitive equilibrium outcome to a scenario where airlines do not react to competitor scarcity and use a discrete set of fares. This scenario matches the constraints that we observe in one airline's internal pricing systems. Our main empirical finding is that the adoption of such heuristics softens competition and allows airlines to alleviate the Bertrand scarcity trap. We estimate that this is welfare improving, increasing revenues (by 4-5%) and consumer surplus (by 3%), relative to the competitive outcome.

Our framework extends single-agent dynamic pricing models (e.g., Gallego and Van Ryzin, 1994; Zhao and Zheng, 2000; Talluri and Van Ryzin, 2004) to oligopoly. We consider uncertain, time-varying demands for an arbitrary number of differentiated products and firms. Firms are exogenously endowed with limited initial inventory. Demand satisfies general regularity conditions. Our framework accommodates demand models beyond independence of irrelevant alternatives, including the form of nested logit demand used in our empirical application. Consumers arrive randomly according to a time-varying Poisson process with preferences that depend on their arrival time. Upon arrival, each consumer decides whether to purchase an available product or exit the market by selecting an outside option. Within a period, firms simultaneously choose prices after observing all remaining inventories. Then, demand is realized and remaining inventories are updated. This process repeats until the deadline or until all products are sold out. Unsold inventory is scrapped after the deadline. We characterize Markov-perfect equilibria, where the payoff-relevant state is the vector of remaining inventories and time.

The Markovian structure allows us to summarize the impact of today's prices on the continuation game in a "scarcity matrix" that depends on the current state. We call the marginal impact on a firm's continuation profit of selling a unit of an own product *own-scarcity effect*. Similarly, we define the impact on a firm's continuation profit of a competitor selling a unit *competitor-scarcity effect.* These scarcity effects define the scarcity matrix, where the number of scarcity effects is equal to the number of firms times the products in the game.

Stage games are complex because each firm's payoff is affected not only by its own residual demands but also by competitors' demands through the competitor-scarcity effects. Prices can be of strategic complements or strategic substitutes.[1] Stage game payoffs are generally not (log) supermodular (Milgrom and Roberts, 1990), nor are they of the form considered in Caplin and Nalebuff (1991) and Nocke and Schutz (2018). To make progress, we use a fixed-point theorem in Kellogg (1976) to derive sufficient conditions for both existence and uniqueness of stage game equilibria. We provide some guidance on when multiplicities may arise.

Our main theorem provides conditions under which the dynamic equilibrium is unique. We prove that the continuous-time limit of the unique discrete-time equilibrium is characterized by a system of ODEs. We first use the characterization to show that well-known results from the single-firm dynamic pricing setting do not carry over to the oligopoly case. For example, value functions are not concave in time, nor are they monotonic or concave in capacity. All scarcity effects can be positive or negative, and non-monotonic in time.

Second, we prove near the deadline that a sale by the firm with the lowest inventory remaining softens competition the most. This result substantiates how intuition from single-firm dynamic pricing models may not hold in an oligopoly, e.g., that scarcity implies high prices. Instead, limited own inventory can result in rivals charging relatively high prices (and the firm relatively low prices) in an attempt to get the firm to sell out. Firms generally benefit from asymmetries. Competition is fiercest when firms have the same number of units remaining. This generalizes Dudey (1992) and Martínez-de Albéniz and Talluri (2011), who focus on undifferentiated products and homogeneous consumers with deterministic constant demand. Moreover, some of the economic forces in our model occur when a single firm faces long-lived buyers and therefore, essentially competes against its future self (Board and Skrzypacz, 2016; Gershkov et al., 2018; Dilme and Li, 2019).[2] As in Dilme and Li (2019), a firm may have fire sales (offer low prices) in order to create future scarcity. Our model has additional dimensions of scarcity that create complementary incentives to increase prices in an attempt to shift demand to competitors.[3]

[1]Downward sloping best-response curves may arise even though we do not endogenize the initial capacity choice as in Dana and Williams (2022).

[2]Board and Skrzypacz (2016) and Gershkov et al. (2018) consider forward-looking buyers when the firm can fully commit to a selling mechanism and hence, resist the temptation to fire-sale.

[3]There exists a large literature on dynamic price competition in other settings, e.g., Maskin and Tirole (1988); Bergemann and Välimäki (2006) who do not consider limited capacity. In related work, Dana (1999a) and Dana (1999b) allow firms to choose prices and quantities before demand uncertainty is resolved. We abstract from

Third, we identify a new welfare effect under price competition—the Bertrand scarcity trap. We provide examples in which dynamic price competition leads to inefficient rationing because low prices result in over-provision early on and under-provision close to the perishability date. It is possible that welfare under dynamic price competition can be lower than if all products were managed by a single firm. Whether or not restricting competitive interactions yields higher welfare depends on demand, which we measure in an economically important industry.

We apply our framework to study airline pricing. While this industry has been noted for significant price dispersion within and across routes (Borenstein and Rose, 1994; Stavins, 2001; Gerardi and Shapiro, 2009; Berry and Jia, 2010; Sengupta and Wiggins, 2014), limited access to necessary data has made it difficult to study dynamic pricing in oligopoly. Most dynamic pricing studies involving perishable inventory consider single-firm settings (e.g., Lazarev, 2013; Williams, 2022; Aryal et al., 2022; Cho et al., 2018; D'Haultfœuille et al., 2022). Exceptions include Puller et al. (2012), who study the role of ticket characteristics in explaining fares, Siegert and Ulbricht (2020), who document airline pricing patterns based on market structure, and Chen and Jeziorski (2023), who analyze dynamic pricing in a duopoly airline route with collected data.[4] We use novel data provided to us by a large U.S. airline that contain not only flight-level prices and bookings for its own flights, but also the same granular data for its competitors.[5] The data are similar to the Nielsen data used to study retail markets, except that our data are not anonymized. Our sample covers 50 routes and nine months of departures in 2019.

We use the data to document new facts on airline pricing and to estimate a Poisson demand model, where aggregate demand uncertainty is captured through Poisson arrivals, and preferences are modeled through discrete choice nested logit demand using a rich set of observable characteristics. Instead of fixing the market size, as commonly done in empirical work, we use search data for one airline to inform arrival process parameters. We scale up the estimated arrival process to account for unobserved searches, e.g., searches conducted on online travel agencies or a rival's website. We show that our results are robust to the choice of scaling param-

forward-looking buyers due to recent empirical evidence that shows limited reshopping using clickstream data in airline markets (Hortaçsu et al., 2021b).

[4]Dynamic pricing without perishable capacity has also been studied. Sweeting et al. (2020) studies limit pricing in airline markets. Kehoe et al. (2018) consider dynamic pricing and the role of information applied to e-commerce.

[5]The airline has elected to remain anonymous.

eter as well as the inclusion of unobserved preferences that are potentially correlated with price (beyond hundreds of fixed effects). We find significant variation in willingness to pay across routes and days before departure. In general, demand becomes more inelastic as the departure date approaches. Average own-price elasticities are -1.4.

With demand estimated, we simulate a number of counterfactuals. We use our ODE equilibrium characterization to solve dynamic games (route-departure dates) with large state spaces—some feature over 131 million states. In total, our analysis studies thousands of dynamic games, involving over 59 billion stage games. We first compare dynamic competitive equilibrium outcomes to another useful theoretical benchmark—uniform pricing. With uniform pricing, each firm commits to a single price for each flight. We find that uniform price competition results in higher total welfare than dynamic price competition, which contrasts with work in single-firm settings, including Hendel and Nevo (2013) in retailing, Castillo (2022) in ride-share, and Williams (2022) for single-carrier airline markets. The reason is that uniform pricing shifts the distribution of sales to later periods where consumers are less price sensitive. This effect outweighs the welfare loss stemming from the inability to react to scarcity.

We then investigate the use of airline pricing heuristics. We base our analysis on the observed pricing technology of one airline—we observe both the pricing heuristic's code and associated documentation. The pricing heuristic is not a reinforcement learning algorithm (Calvano et al., 2020; Asker et al., 2021; Hansen et al., 2021) as it focuses on perishable inventory subject to time-varying opportunity costs of remaining capacity. Our empirical findings suggest a strategic reason for why airlines use pricing heuristics. Heuristics soften price competition early on and alleviate the Bertrand scarcity trap. Firms benefit from shifting the distribution of bookings to periods with less price sensitive demands (by up to 5% revenue increase). Two opposing forces affect late-arriving, price insensitive consumers. Heuristics result in higher prices than under uniform pricing, but lower prices than the competitive equilibrium close to departure. Output under heuristics is higher close to departure and therefore, in aggregate, we find that heuristics also benefit consumers (3% higher consumer surplus).

Although our analysis identifies that pricing heuristics may improve welfare relative to the perfect information, competitive benchmark, we find that the overall benefits of competition are significant. The competitive equilibrium yields substantially higher welfare (16%) than joint-profit maximization by a single firm. The dynamic competitive equilibrium outcome obtains 88% of the first best under a social planner. Heuristics obtain 93% of the first best.

2 A Model of Dynamic Price Competition

In this section, we build and analyze a model of dynamic price competition with finite inventory and a sales deadline. In Section 2.1, we set up the dynamic pricing game and detail our demand assumptions. In Section 2.2, we consider the special case where all products are owned by a single firm. We then characterize the equilibrium of the oligopoly case in Section 2.3. We discuss stage game properties in Section 2.4 and properties of equilibrium dynamics in Section 2.5. Finally, we formalize the Bertrand scarcity trap in Section 2.6.

2.1 Model

Firms, products, and timing. We consider a set $\mathscr{F} := \{1,\ldots,F\}$ of firms and a set $\mathscr{J} := \{1,\ldots,J\}$ of products. Products in $\mathscr{J}_f$ are owned by firm f, where $\left(\mathscr{J}_f\right)_{f\in\mathscr{F}}$ is a partition of $\mathscr{J}$. That is, $\mathscr{J} = \bigcup_{f\in\mathscr{F}} \mathscr{J}_f$ and $\mathscr{J}_f \cap \mathscr{J}_{f'} = \emptyset$ for $f \neq f'$ so that each product is sold by exactly one firm. Each firm f is endowed with initial inventories $K_{j,0} \in \mathbb{N}$, $j \in \mathscr{J}_f$. Any remaining inventory at the deadline $T > 0$ is scrapped with zero value. We study a discrete-time environment with periods $t \in \{0, \Delta, \ldots, T-\Delta\}$, $\Delta > 0$, and later consider the continuous-time limit as $\Delta \to 0$. In every period t, firms simultaneously set prices of their products $\mathbf{p}_{f,t} := (p_{j,t})_{j\in\mathscr{J}_f}$. Then, a single consumer arrives with probability $\Delta\lambda_t$, where λ_t is smooth in t. Therefore, each consumer can be indexed by her arrival time t. Consumer t either buys a unit of an available product upon arrival or leaves the market. In the following, we impose assumptions on purchase probabilities.

Demand. If all products are available, then consumer t, facing a price vector $\mathbf{p} := \left(p_j\right)_{j\in\mathscr{J}}$, buys product j with probability $s_j(\mathbf{p}; \boldsymbol{\theta}_t, \mathscr{J})$, where $\boldsymbol{\theta}_t \in \mathscr{T} \subset \mathbb{R}^n$ is a vector of $n \geq 1$ parameters that are smooth and deterministic in t. We impose the following regularity conditions on s_j.

Assumption 1. *For all $\boldsymbol{\theta} \in \mathscr{T}$ and $\mathbf{p} \in \mathbb{R}^{\mathscr{J}}$, the following hold:*

i) For any j, $\lim_{p_j \to \infty} s_j(\mathbf{p}; \boldsymbol{\theta}, \mathscr{J}) p_j = 0$. For any subset $\mathscr{A} \subset \mathscr{J}$ and $j \in \mathscr{A}$, the limit[6]

$$s_j(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A}) := \lim_{\substack{p_{j'} \to \infty \\ j' \notin \mathscr{A}}} s_j(\mathbf{p}; \boldsymbol{\theta}, \mathscr{J}) \in [0,1]$$

exists and is smooth in $\boldsymbol{\theta}$ and $\mathbf{p}^{\mathscr{A}} \in \mathbb{R}^{\mathscr{A}}$, where $p_{j'}^{\mathscr{A}} = p_{j'}$ for all $j' \in \mathscr{A}$;

ii) For all j, $s_j(\mathbf{p}; \boldsymbol{\theta}, \mathscr{J})$ is strictly decreasing in p_j and strictly increasing in $p_{j'}$ for $j' \neq j$;

iii) For all $\mathscr{A} \subset \mathscr{J}$, $j \in \mathscr{A}$ and $\underline{\mathbf{p}} \in \mathbb{R}^{\mathscr{A}}$, there exists a $C > 0$ such that for all $\mathbf{p}^{\mathscr{A}} \geq \underline{\mathbf{p}}$,

$$C s_j(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A}) < \frac{\partial s_0}{\partial p_j}(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A}) \qquad \text{for all } j,$$

where $s_0(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A}) := 1 - \sum_{j' \in \mathscr{A}} s_{j'}(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A})$.

Assumption 1-i) ensures that demand is well-defined when products sell out, i.e., when these products' prices are set to equal infinity. Assumption 1-ii) states that all products are imperfect substitutes. Assumption 1-iii) can be viewed as a generalized concavity assumption as it assures that profit-maximizing prices of a static multi-product problem are interior and uniformly bounded from above. This relatively weak assumption essentially asserts that $\frac{\partial s_0}{\partial p_j}(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A}) / s_j(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A})$ remains bounded from zero when $p_j^{\mathscr{A}}$ is large. The condition also implies that $\frac{\partial s_0}{\partial p_j}(\mathbf{p}^{\mathscr{A}}; \boldsymbol{\theta}, \mathscr{A}) > 0$, i.e., the outside option is an imperfect substitute for all products.

We show in Appendix A that Assumption 1 implies that a single firm's profit-maximizing prices solve a system of first-order conditions (FOCs). To write these FOCs in matrix form, we denote the vector of choice probabilities of available products by $\mathbf{s}^{\mathscr{A}}(\cdot) := \left(s_j(\cdot)\right)_{j \in \mathscr{A}}$ and define for any $\boldsymbol{\theta}$, $\mathscr{A}$, and $\mathbf{p} \in \mathbb{R}^{\mathscr{A}}$, the vector of inverse quasi own-price elasticities of demand as[7]

$$\hat{\boldsymbol{\epsilon}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A}) := \left(\left(D_{\mathbf{p}} \boldsymbol{s}^{\mathscr{A}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A})\right)^{\intercal}\right)^{-1} \boldsymbol{s}^{\mathscr{A}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A}).$$

Therefore, the FOCs of a single firm's profit maximization problem $\max_{\mathbf{p} \in \mathbb{R}^{\mathscr{A}}} \mathbf{s}^{\mathscr{A}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A})^{\intercal}(\mathbf{p} -$

[6]The limit takes all prices of products $j' \notin \mathscr{A}$ to infinity where the order does not matter.

[7]The Jacobian of the demand vector is invertible by Assumption 1-iii) as we explain in the Appendix A.2. We restrict attention to available products to assure that the Jacobian matrix of the demand vector is invertible.

c) for an arbitrary marginal cost vector $\mathbf{c} \in \mathbb{R}^{\mathscr{A}}$ can be written as a classic markup formula:

$$\mathbf{p} - \mathbf{c} = -\hat{\boldsymbol{\epsilon}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A}).$$

This system of equations has a unique solution if we impose the following assumption:

Assumption 2. *The vector of inverse quasi own-price elasticities* $\hat{\boldsymbol{\epsilon}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A})$ *satisfies*

$$\det\Big(-D_{\mathbf{p}}\hat{\boldsymbol{\epsilon}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A}) - I\Big) \neq 0$$

for all $\mathbf{p} \in \mathbb{R}^{\mathscr{A}}$, $\boldsymbol{\theta} \in \mathscr{T}$, *and* $\mathscr{A} \subset \mathscr{J}$, *where* $I \in \mathbb{R}^{\mathscr{A} \times \mathscr{A}}$ *is the identity matrix.*

Assumption 2 is exactly the assumption made in Kellogg (1976) that implies that $\mathbf{c} - \hat{\boldsymbol{\epsilon}}(\mathbf{p}; \boldsymbol{\theta}, \mathscr{A})$ has a unique fixed point. This assumption replaces the commonly made assumption of quasi-concavity or log-concavity which is, for example, not satisfied for multinomial logit (see, e.g., Hanson and Martin, 1996). Our demand assumptions hold for commonly used demand systems, e.g., logit demand. We show in Appendix C that our assumptions also hold for the form of nested logit demand that we consider in our empirical application.

We omit the conditioning arguments $\boldsymbol{\theta}$ and/or $\mathscr{A}$ in all expressions whenever the meaning is unambiguous. When the time index is relevant, we write $s_{j,t}(\mathbf{p}) := s_j(\mathbf{p}; \boldsymbol{\theta}_t, \mathscr{A}_t)$.

Markov perfect equilibrium. The payoff-relevant state in the pricing game is given by the vector of inventories $\mathbf{K} := (K_j)_{j \in \mathscr{J}}$ and time t. We study Markov perfect equilibria in which each firm's strategy is measurable with respect to $(\mathbf{K}, t)$. We denote a Markov pricing strategy of firm f by $\mathbf{p}_{f,t}(\mathbf{K}) = \big(p_{j,t}(\mathbf{K})\big)_{j \in \mathscr{J}_f}$.

2.2 The Single-Firm Case

We start with a special case of our model when all products are owned by a single firm. This allows us to introduce supply-side notation that we carry over to the oligopoly case. We demonstrate that the single-firm case is well behaved and exhibits "nice" properties. Unfortunately, all of these properties can fail in the oligopoly case.

Consider a single firm M that offers all J products for sale. The firm's continuation payoff at time $t \leq T-\Delta$, given capacity vector $\mathbf{K}$, satisfies the dynamic program

$$\Pi_{M,t}(\mathbf{K};\Delta) = \max_{\mathbf{p}} \Delta\lambda_t \sum_{j\in\mathscr{J}} \underbrace{s_{j,t}(\mathbf{p})\Big(p_j + \Pi_{M,t+\Delta}(\mathbf{K}-\mathbf{e}_j;\Delta)\Big)}_{\text{payoff from selling product } j} + \underbrace{\Big(1-\Delta\lambda_t \sum_{j\in\mathscr{J}} s_{j,t}(\mathbf{p})\Big)}_{\text{probability of no purchase}} \Pi_{M,t+\Delta}(\mathbf{K};\Delta),$$

where $\mathbf{e}_j \in \mathbb{N}^{\mathscr{J}}$ is the unit vector of zeros with a one in the jth position. The firm receives a revenue of p_j and a continuation value in period $t+\Delta$ with one fewer unit of j if it sells. If the firm does not sell, the capacity vector remains unchanged, and time moves forward by Δ. The firm faces two boundary conditions: (i) $\Pi_{M,T}(\mathbf{K};\Delta) = 0$ for all $\mathbf{K}$ and (ii) $\Pi_{M,t}(\mathbf{K};\Delta) = -\infty$ if $K_j < 0$ for a $j \in \mathscr{J}$. The boundary conditions ensure that remaining inventory is scrapped with zero value after the deadline T and that the firm cannot oversell.

Hence, the optimal price at each state $(\mathbf{K}, t)$ solves a static maximization problem parameterized by the demand parameters $\boldsymbol{\theta}$ and $\boldsymbol{\omega} = (\omega_j)_{j\in\mathscr{J}}$, where $\omega_j = \Pi_{M,t}(\mathbf{K};\Delta) - \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j;\Delta)$ is commonly referred to as the *opportunity cost of selling product* j:[8]

$$\max_{\mathbf{p}} \sum_{j\in\mathscr{J}} s_j(\mathbf{p};\boldsymbol{\theta},\mathscr{J})\big(p_j - \omega_j\big).$$

We denote the profit-maximizing price for parameters $\boldsymbol{\theta}$ and $\boldsymbol{\omega}$ by $\mathbf{p}^M(\boldsymbol{\omega},\boldsymbol{\theta}) := (p_j^M(\boldsymbol{\omega},\boldsymbol{\theta}))_{j\in\mathscr{J}}$. By Kellogg (1976), Assumption 2 implies that there is a unique optimal price vector which is continuous in $\boldsymbol{\omega}$ and $\boldsymbol{\theta}$. The solution of the optimal control problem can be characterized by an ordinary differential equation (ODE) as $\Delta \to 0$ by the following Lemma.

Lemma 1. $\Pi_{M,t}(\mathbf{K}) := \lim_{\Delta\to 0} \Pi_{M,t}(\mathbf{K};\Delta)$ *solves the ordinary differential equation*

$$\dot{\Pi}_{M,t}(\mathbf{K}) = -\lambda_t \sum_{j\in\mathscr{J}} s_{j,t}\left(\mathbf{p}^M(\boldsymbol{\omega}_t(\mathbf{K}),\boldsymbol{\theta})\right)\Big(p_j^M(\boldsymbol{\omega}_t(\mathbf{K}),\boldsymbol{\theta}) - \omega_{j,t}(\mathbf{K})\Big), \tag{1}$$

[8]Note that strictly speaking, the opportunity cost of selling product j is given by $\omega_j - \sum_{j'\neq j} \frac{s_{j'}'(\mathbf{p})}{1-s_j(\mathbf{p})}\omega_{j'}$ as by selling product j, the firm forgoes the opportunity to sell any other product to the customer.

where $\boldsymbol{\omega}_t(\mathbf{K}) := \left(\Pi_{M,t}(\mathbf{K}) - \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j)\right)_{j\in\mathscr{J}}$ *with boundary conditions (i)* $\Pi_{M,T}(\mathbf{K}) = 0$ *for all* $\mathbf{K}$*, and (ii)* $\Pi_{M,t}(\mathbf{K}) = -\infty$ *if* $K_j < 0$ *for a* $j \in \mathscr{J}$*.*

Each state $(\mathbf{K}, t)$ defines a set of available products $\mathscr{A}(\mathbf{K}) = \{j : K_j \neq 0\}$ and a vector of opportunity costs of available products $\boldsymbol{\omega}_{\mathscr{A},t}(\mathbf{K}) := \left(\Pi_{M,t}(\mathbf{K}) - \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j)\right)_{j\in\mathscr{A}}$. Profit-maximizing prices $\mathbf{p}^M_t(\mathbf{K}) \in \mathbb{R}^{\mathscr{A}(\mathbf{K})}$ of available products solve

$$\mathbf{p} = \underbrace{\boldsymbol{\omega}_{\mathscr{A},t}(\mathbf{K})}_{\text{opportunity costs}} - \underbrace{\hat{\epsilon}(\mathbf{p};\boldsymbol{\theta},\mathscr{A})}_{\text{inverse quasi own-price elasticities}} . \tag{2}$$

Hence, the optimal pricing policy $\mathbf{p}^M_t(\mathbf{K})$ is continuous in time and its dynamics are governed by the evolution of quasi own-price elasticities and the opportunity costs. The opportunity costs in turn depend on the stochastic process of remaining inventory $\mathbf{K}_t = \left(K_{j,t}\right)_{j\in\mathscr{J}}$.[9] Proposition 1 summarizes properties of the single-firm dynamic pricing model.

Proposition 1. *The solution to the continuous-time single-firm revenue maximization problem in Lemma 1 satisfies the following:*

i) $\Pi_{M,t}(\mathbf{K})$ *is decreasing in* t *for* $\mathbf{K} \neq \mathbf{0}$ *and increasing in* K_j*, for all* $j \in \mathscr{J}$ *and* $t < T$*;*

ii) $\omega_{j,t}(\mathbf{K})$ *is decreasing in* t *for* $\mathbf{K} \neq \mathbf{0}$ *and decreasing in* K_j*, for all* j *and* $t < T$*;*

iii) The stochastic process $\omega_{j,t\wedge\tau}(\mathbf{K}_t)$*,* $\tau := \inf\{t \geq 0 | K_{j,t} \leq 1\}$*, is a submartingale.*

Statements i) and ii) of Proposition 1 imply that more inventory and more time remaining increase continuation profits, continuation profits are concave in capacity, and that opportunity costs are decreasing towards the deadline if $\mathbf{K}$ is held fixed. These properties generalize Gallego and Van Ryzin (1994), who consider a single product. Statement (iii) asserts that, on average, opportunity costs are increasing. This formal result implies that given constant $\boldsymbol{\theta}_t \equiv \boldsymbol{\theta}$, price paths are on average increasing in time by Equation 2. That is, demand uncertainty alone causes prices to increase on average over time. The price of a product may decrease when only one unit of it remains because the firm cannot benefit from scarcity once the product sells out.[10]

[9] Note that we abuse notation slightly by denoting the optimal price policy $\mathbf{p}^M_t(\mathbf{K})$, while also denoting the static optimal price parameterized by $(\boldsymbol{\omega},\boldsymbol{\theta})$ by $\mathbf{p}^M(\boldsymbol{\omega},\boldsymbol{\theta})$.

[10] This result explains the inverted U-shape found when plotting average dynamic prices over time (McAfee and Te Velde, 2006, e.g.,). One a product sells out, its price is excluded from the average price.

2.3 Equilibrium Characterization of the Oligopoly Pricing Game

We derive an analogous system of ODEs to Equation 1 for the limit of continuation profits of all firms f, denoted $\Pi_{f,t}(\mathbf{K};\Delta)$. One of our key insights is to show that equilibrium prices in state $(\mathbf{K},t)$ correspond to an equilibrium of a stage game that is parameterized by *scarcity effects.* We define the scarcity effect of product j on firm f in state $(\mathbf{K},t)$ to be

$$\omega^f_{j,t}(\mathbf{K};\Delta) := \Pi_{f,t+\Delta}(\mathbf{K};\Delta) - \Pi_{f,t+\Delta}(\mathbf{K}-\mathbf{e}_j;\Delta).$$

This scarcity effect captures the impact that one fewer unit of product j has on the continuation value of firm f. We call $\omega^f_{j,t}(\mathbf{K};\Delta)$ for $j \in \mathscr{J}_f$ *own-scarcity effects* and $\omega^f_{j,t}(\mathbf{K};\Delta)$ for $j \notin \mathscr{J}_f$ *competitor-scarcity effects*. The competitor-scarcity effects capture the impact of a sale of a competitor's product on own continuation profits. Today's prices affect future payoffs through these scarcity effects. Formally, given a pricing strategy $\mathbf{p}_t(\mathbf{K}) := \left(p_{j,t}(\mathbf{K})\right)_{j\in\mathscr{J}}$, firm f's value function can be recursively written as (omitting Δ for readability)

$$\begin{aligned}\Pi_{f,t}(\mathbf{K};\Delta) = \quad & \Delta\lambda_t \Bigg(\underbrace{\sum_{j\in\mathscr{J}_f} s_{j,t}(\mathbf{p}_t(\mathbf{K}))\,(p_{j,t}(\mathbf{K}) + \Pi_{f,t+\Delta}(\mathbf{K}-\mathbf{e}_j;\Delta)) +}_{\text{payoff from own sale}} \\ & \underbrace{\sum_{j'\neq\mathscr{J}\setminus\mathscr{J}_f} s_{j',t}(\mathbf{p}_t(\mathbf{K}))\,\Pi_{f,t+\Delta}(\mathbf{K}-\mathbf{e}_{j'};\Delta)}_{\text{payoff if competitor sells } j'} \Bigg) + \underbrace{\left(1-\Delta\lambda_t \sum_{j'\in\mathscr{J}} s_{j'}(\mathbf{p}_t(\mathbf{K}))\right)}_{\text{probability of no purchase}} \Pi_{f,t+\Delta}(\mathbf{K};\Delta).\end{aligned}$$

Subtracting $\Pi_{f,t+\Delta}(\mathbf{K};\Delta)$ does not change pricing incentives in state $(\mathbf{K},t)$, so firm f chooses its prices $\mathbf{p}_t(\mathbf{K})$ to maximize $\Pi_{f,t}(\mathbf{K};\Delta) - \Pi_{f,t+\Delta}(\mathbf{K};\Delta)$ which is equal to

$$\Delta\lambda_t \Bigg(\sum_{j\in\mathscr{J}_f} s_{j,t}\,(\mathbf{p}_t(\mathbf{K}))\Big(p_{j,t}(\mathbf{K}) - \omega^f_{j,t}(\mathbf{K};\Delta)\Big) - \textstyle\sum_{j'\notin\mathscr{J}_f} s_{j',t}\,(\mathbf{p}_t(\mathbf{K}))\,\omega^f_{j',t}(\mathbf{K};\Delta)\Bigg). \tag{3}$$

The first part of Equation 3 is analogous to the single-firm setting—expected demand times a markup. The second part of the equation is new and measures how a firm is affected by competitor scarcity, weighted by competitor demand.

For any equilibrium Markov pricing strategy $(\mathbf{p}^*_t(\cdot))_{t=0,\ldots T-\Delta}$ of the dynamic pricing game,

$\mathbf{p}_t^*(\mathbf{K})$ is an equilibrium of a stage game in which each firm f maximizes Equation 3. For each state $(\mathbf{K}, t)$, the stage game can be parameterized by a scarcity matrix

$$\Omega_t(\mathbf{K};\Delta) = \left(\omega_{j,t}^f(\mathbf{K};\Delta)\right)_{f,j} \in \mathbb{R}^{\mathscr{F}\times\mathscr{J}}$$

and demand parameters $\boldsymbol{\theta}$. We denote equilibrium prices of a stage game with parameters $\Omega, \boldsymbol{\theta}$ by $\mathbf{p}^*(\Omega, \boldsymbol{\theta})$. Unfortunately, stage game equilibria are generally not unique as we show in Section 2.4. However, if the time horizon is not too long and for sufficiently small Δ, $\Omega_t(\mathbf{K};\Delta)$ stays in a neighborhood that guarantees that all stage games admit a unique equilibrium. Therefore, equilibrium continuation profits can be characterized by a system of ODEs.

We additionally derive an ODE for equilibrium price paths using the FOCs of the stage game. Specifically, consider a capacity vector $\mathbf{K}$. The equilibrium stage game payoff of firm f given $\mathbf{K}$ can be written in matrix form. Given the demand vector $\mathbf{s}_f^{\mathscr{A}(\mathbf{K})}(\cdot) = (s_j(\cdot))_{j\in\mathscr{J}_f\cap\mathscr{A}(\mathbf{K})}$ of available products owned by firm f and the limiting equilibrium price vector $\mathbf{p}_t^{\mathscr{A}(\mathbf{K}),*}$ of available products, the equilibrium payoff is given by

$$\mathbf{s}_f^{\mathscr{A}(\mathbf{K})}(\mathbf{p}_t^{\mathscr{A}(\mathbf{K}),*};\boldsymbol{\theta}_t,\mathscr{A}(\mathbf{K}))^\intercal \mathbf{p}_{f,t}^{\mathscr{A}(\mathbf{K}),*} - \mathbf{s}(\mathbf{p}_t^{\mathscr{A}(\mathbf{K}),*};\boldsymbol{\theta}_t,\mathscr{A}(\mathbf{K}))^\intercal \boldsymbol{\omega}_t^f(\mathbf{K}).$$

We omit the sub- and superscripts and arguments $\mathscr{A}$ and $\mathbf{K}$ since we are holding the capacity vector fixed for the following argument. It follows that $\mathbf{p}_t^*$ satisfies the FOC

$$\mathbf{g}_f(\mathbf{p},\Omega_t,\boldsymbol{\theta}_t) := \underbrace{\left(\left(D_{\mathbf{p}_f}\mathbf{s}_f(\mathbf{p};\boldsymbol{\theta}_t)\right)^\intercal\right)^{-1} D_{\mathbf{p}_f}\left(\mathbf{s}(\mathbf{p};\boldsymbol{\theta}_t)^\intercal \boldsymbol{\omega}_t^f\right)^\intercal}_{\substack{\text{net opportunity costs}\\ \text{of selling}}} - \underbrace{\left(\left(D_{\mathbf{p}_f}\mathbf{s}_f(\mathbf{p};\boldsymbol{\theta}_t)\right)^\intercal\right)^{-1}\mathbf{s}_f(\mathbf{p};\boldsymbol{\theta}_t)}_{\substack{\text{inverse quasi}\\ \text{own-price elasticities}}} \equiv \mathbf{p}.$$

This implies that for $\mathbf{g} := (\mathbf{g}_f)_{f\in\mathscr{F}}$, $\frac{\partial}{\partial t}\left(\mathbf{g}(\mathbf{p}_t,\Omega_t,\boldsymbol{\theta}_t) - \mathbf{p}_t\right) = 0$ defines an ODE for equilibrium prices. We summarize our main theoretical result in Theorem 1.

Theorem 1 (Continuous-time Limit)**.** *For every* $\mathbf{K}$*, there exists a* $T_0(\mathbf{K}) > 0$*, non-increasing in* $\mathbf{K}$*, so that for any* $T \leq T_0(\mathbf{K})$*, there exists a unique equilibrium of the dynamic pricing game for sufficiently small* Δ *with the following properties:*

i) Each equilibrium value function $\Pi_{f,t}^*(\mathbf{K};\Delta)$ *converge to a limit* $\Pi_{f,t}^*(\mathbf{K})$ *as* $\Delta \to 0$ *that* *solves the ordinary differential equation*

$$\dot{\Pi}_{f,t}(\mathbf{K}) = -\lambda_t \left(\sum_{j\in\mathscr{J}_f} s_j(\mathbf{p}^*(\Omega_t(\mathbf{K}),\boldsymbol{\theta}_t))\left(p_j^*(\Omega_t(\mathbf{K}),\boldsymbol{\theta}_t) - (\Pi_{f,t}(\mathbf{K}) - \Pi_{f,t}(\mathbf{K}-\mathbf{e}_j))\right) - \sum_{j'\notin\mathscr{J}_f} s_{j'}(\mathbf{p}^*(\Omega_t(\mathbf{K});\boldsymbol{\theta}_t))\left(\Pi_{f,t}(\mathbf{K}) - \Pi_{f,t}(\mathbf{K}-\mathbf{e}_{j'})\right)\right),$$

where $\Omega_t(\mathbf{K})) = \lim_{\Delta\to 0}\Omega_t(\mathbf{K};\Delta)$*, with boundary conditions (i)* $\Pi_{f,T}(\mathbf{K}) = 0$ *for all* $\mathbf{K}$*, (ii)* $\Pi_{f,t}(\mathbf{K}) = -\infty$ *if* $K_j < 0$ *for a* $j \in \mathscr{J}_f$*, and (iii)* $\Pi_{f,t}(\mathbf{K}-\mathbf{e}_{j'}) = \Pi_{f,t}(\mathbf{K})$ *if* $K_{j'} = 0$ *for a* $j' \notin \mathscr{J}_f$*,* $K_j \geq 0$ *for all* $j \in \mathscr{J}_f$*;*

ii) For each capacity vector $\mathbf{K}$*, the vector of equilibrium prices* $\mathbf{p}_\mathbf{t}^*(\mathbf{K}) \in \mathbb{R}^{\mathscr{A}(\mathbf{K})}$ *of available products satisfies the ordinary differential equation*

$$\dot{\mathbf{p}}_t = -\left(D_\mathbf{p}\mathbf{g}(\mathbf{p}_t,\Omega_t,\boldsymbol{\theta}_t) - I\right)^{-1}\left(\sum_{j,f} D_{\omega_j^f}\mathbf{g}(\mathbf{p}_t,\Omega_t,\boldsymbol{\theta}_t)\dot{\omega}_{j,t}^f + D_{\boldsymbol{\theta}}\mathbf{g}(\mathbf{p}_t,\Omega_t,\boldsymbol{\theta}_t)\right),$$

where we omit all arguments $\mathbf{K}$ *and take* $\dot{\omega}_{j,t}^f(\mathbf{K}) = \dot{\Pi}_{f,t}(\mathbf{K}) - \dot{\Pi}_{f,t}(\mathbf{K}-\mathbf{e}_j)$ *as given, with boundary condition* $\mathbf{p}_T = \mathbf{g}(\mathbf{p}_T,\mathbf{O},\boldsymbol{\theta}_T)$*.*

Using this system of ODEs, we can calculate $\Pi_{f,t}(\mathbf{K})$, $\mathbf{p}_t(\mathbf{K})$, and $\Omega_t(\mathbf{K})$, for all t without explicitly calculating stage game equilibria except for at time T. This single stage game is easy to analyze because $\Omega_T = \mathbf{0}$. Theorem 1 offers a powerful tool to simulate equilibria in games with large stage spaces because stage games are not solved explicitly. It also allows us to derive equilibrium properties that we discuss next.

2.4 Properties of the Stage Game

In this section, we examine how own/competitor scarcity affects the nature of price competition and uniqueness of stage game equilibria. We illustrate stage game properties using a simple duopoly example with two products and logit demand of the form $s_f(\mathbf{p}) = \frac{\exp(1-p_f)}{1+\sum_{f'\in\{1,2\}}\exp(1-p_{f'})}$ for

$f \in \{1,2\}$. Stage games are parameterized by a 2×2 dimension matrix

$$\Omega = \begin{pmatrix} \omega_1^1 & \omega_2^1 \\ \omega_1^2 & \omega_2^2 \end{pmatrix}.$$

Each firm maximizes $s_f(\mathbf{p})\left(p_f - \omega_f^f\right) - s_{f'}(\mathbf{p})\,\omega_{f'}^f$ with corresponding FOC:

$$\frac{\partial s_{f'}}{\partial p_f}(\mathbf{p})\left(\frac{\partial s_f}{\partial p_f}(\mathbf{p})\right)^{-1}\omega_{f'}^f + \omega_f^f - s_f(\mathbf{p})\left(\frac{\partial s_f(\mathbf{p})}{\partial p_f}\right)^{-1} = p_f.$$

Note that the left-hand side of the above equation defines $g_f(p_1, p_2)$. An increase in the competitor price increases firm f's best response price if $\frac{\partial g_f}{\partial p_{f'}} > 0$. That is, the competitor's price is a strategic complement. However, if $\frac{\partial g_f}{\partial p_{f'}} < 0$, then an increase in the competitor price decreases firm f's best response price. That is, the competitor's price is a strategic substitute.

To show when strategic complements or strategic substitutes arise for this example, we calculate (see Figure 14 in Appendix D for graphs)

$$\frac{\partial}{\partial p_{f'}} g_f(\mathbf{p}) = \frac{\partial}{\partial p_{f'}}\left(\frac{\partial s_{f'}}{\partial p_f}(\mathbf{p})\left(\frac{\partial s_f}{\partial p_f}(\mathbf{p})\right)^{-1}\right)\omega_{f'}^f - \frac{\partial}{\partial p_{f'}}\left(\underbrace{s_f(\mathbf{p})\left(\frac{\partial s_f(\mathbf{p})}{\partial p_f}\right)^{-1}}_{\substack{\text{inverse quasi} \\ \text{own-price elasticity}}}\right). \tag{4}$$

First, note that the inverse-quasi own-price elasticity $s_f(\mathbf{p})\left(\frac{\partial s_f(\mathbf{p})}{\partial p_f}\right)^{-1} = -(1 - s_f(\mathbf{p}))^{-1}$ is decreasing in the competitor's price $p_{f'}$. Hence, if the competitor-scarcity effect $\omega_{f'}^f$ is zero, then the competitor price is a strategic complement. This game coincides with a static price competition game with imperfect substitutes.[11] Caplin and Nalebuff (1991) and Nocke and Schutz (2018) study such oligopoly games if demand satisfies the property of independence of irrelevant alternatives. Second, note that $\frac{\partial s_{f'}}{\partial p_f}(\mathbf{p})\left(\frac{\partial s_f}{\partial p_f}(\mathbf{p})\right)^{-1} = -\frac{\exp(1-p_{f'})}{1+\exp(1-p_{f'})}$ is increasing in $p_{f'}$. Hence, if the competitor-scarcity effect $\omega_{f'}^f$ is positive, the competitor price is also a strategic complement.

In the dynamic game, $\omega_{f'}^f$ measures the impact on a firm's profit if the rival sells. Because scarcity typically (but not always as we show in Section 2.5) increases prices, $\omega_{f'}^f$ is typically

[11] As noted in Vives (2018) and Nocke and Schutz (2018), static oligopoly games in multi-product environments are generally not games of strategic complements.

negative. If $\omega^f_{f'}$ is large and negative, then the competitor price can be a strategic substitute. Intuitively, the competitor's market share increases if the competitor decreases its price and hence, the firm's marginal benefit of selling decreases if the competitor-scarcity effect is large. As a result, the firm increases its own price.

Figure 1: Multiplicities in stage-game equilibria

(a) Parametrization of $(\omega^1_2(x), \omega^2_1(x))$

(b) Multiplicity in equilibrium prices

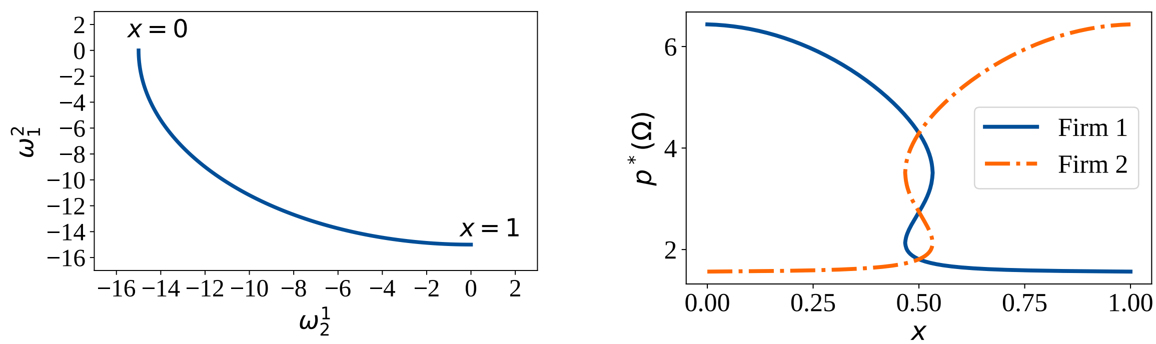

Note: In these graphics we parameterize (ω^2_1, ω^1_2) with a curve $(\omega^1_2(x), \omega^2_1(x)) = \left(-15\cos\left(\frac{\pi}{2}x\right), -15\sin\left(\frac{\pi}{2}x\right)\right), x \in [0,1]$, where we set $(\omega^1_1, \omega^2_2) = (0,0)$, and assume logit demand with $\delta = (1,1)$, $\alpha_t = 1$. Panel (a) depicts the parameterized curve and panel (b) equilibrium prices of both firms given (ω^2_1, ω^1_2) at varying values of x.

These strategic effects are the reason why Proposition 1 does not extend to oligopoly and why multiplicity of equilibria may arise. Note that while own-scarcity effects simply shift best-response functions, competitor-scarcity effects also change their slopes.[12] Our sufficient condition for uniqueness of stage games (Assumption 3 in Appendix A) requires

$$\left(s_1(\mathbf{p}) + \alpha\omega^1_2 s_0(\mathbf{p})\right)\left(s_2(\mathbf{p}) + \alpha\omega^2_1 s_0(\mathbf{p})\right) \neq 1 + \frac{1 - s_1(\mathbf{p}) - s_2(\mathbf{p})}{s_1(\mathbf{p})s_2(\mathbf{p})} \text{ for all } \mathbf{p}.$$

Indeed, this condition does not depend on firms' own-scarcity effects (ω^1_1 and ω^2_2), but it can be violated if competitor-scarcity effects are large. We demonstrate that multiplicity may arise using logit demand. We set own-scarcity effects equal to zero and parameterize competitor-scarcity effects using a continuous function. We plot a parameterization of a path of (ω^2_1, ω^1_2) in Figure 1-(a) along with the corresponding equilibrium prices for both firms in Figure 1-(b). Note that moving along x results in price jumps to a different equilibrium at around $x = 0.55$.

[12] We illustrate this in Figure 15 in Appendix D.

While this example suggests that the dynamic pricing game might not converge to a system of ODEs, we show in Appendix A that any scarcity matrix in a neighborhood of $\Omega = \mathbf{O}$ results in a unique equilibrium. Thus, Lemma 1 can be generalized to an oligopoly as long as the time horizon is not too long. Otherwise, scarcity effects can become sufficiently large so that multiplicities arise.

2.5 Properties of Equilibrium Dynamics

Using our equilibrium characterization, we can show that the general insights from the single-firm setting (Proposition 1) do not extend to oligopoly. We continue with the duopoly example of the previous section. We fix firm 2's initial capacity to be $K_2 = 3$ and vary firm 1's initial capacity K_1 (either 2 or 4). In Figure 2-(a), we plot firm 1's profits over time, and in panel (b), we plot firm 1's own-scarcity effects over time. The graphs show that firm 1 expects higher profits with $K_1 = 4$ than with $K_1 = 2$ far from the deadline, however, this reverses close to the the deadline. As a result, (i) value functions are non-monotonic in own capacity; (ii) own-scarcity effects are non-monotonic in own capacity; and (iii) all scarcity effects can be positive or negative (shown in Figure 16 in Appendix D).

Figure 2: Simulated profits and own-scarcity effects when $K_2 = 3$ and K_1 varies

(a) Firm 1 equilibrium profit, $\Pi_t^1(K_1, 3)$ (b) Firm 1 own-scarcity effect, $\omega_{1,t}^1(K_1, 3)$

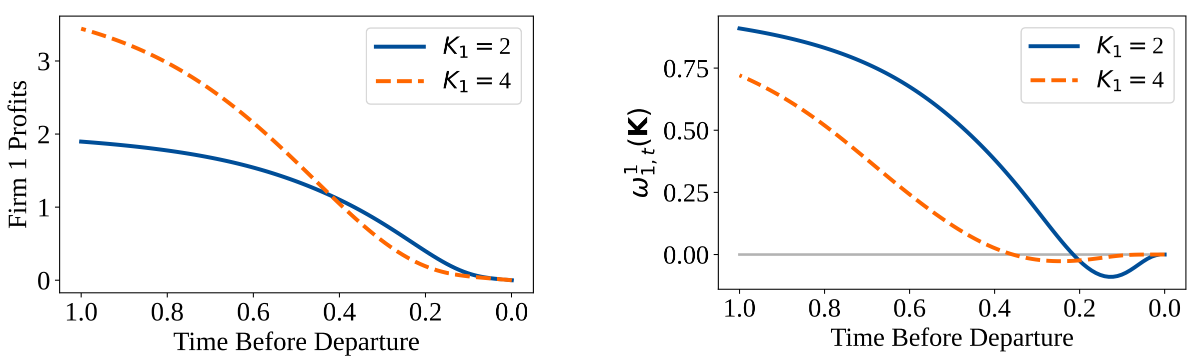

Notes: The simulations assume $\delta = (1,1)$, $\alpha_t \equiv 1$ and logit demand. Panel (a) shows firm 1's profits over time, $t \in [0,1]$, for $\mathbf{K} = (2,3)$ and $\mathbf{K} = (4,3)$. Panel (b) shows firm 2's profits over time, $t \in [0,1]$, for the same states.

Although this simple example shows that the properties of Proposition 1 do not hold for the dynamic pricing game, we make progress by showing that equilibrium dynamics are affected by scarcity in an intuitive way. We show that with a constant distribution of demand over time, prices increase the most, i.e., competition softens the most, after a sale of the product with the smallest inventory remaining. Formally, close to the deadline, we compare the evolution over time of equilibrium price paths $\mathbf{p}_t^*(\mathbf{K})$ across different capacity vectors $\mathbf{K}$. The price level at the deadline T is independent of the capacity vector as it is simply the equilibrium price in a Bertrand game with zero marginal cost for all available products. Hence, it suffices to compare the order of change of $\mathbf{p}_t^*$ close to the deadline for different capacity vectors $\mathbf{K}$. We show that the order of change for all available products is determined by the inventory remaining of the product with the smallest inventory remaining. Hence, price changes are relatively large if the inventory of this product changes. We illustrate this in Figure 17 in Appendix D and state the formal result here.

Proposition 2. *Let $\lambda_t \equiv \lambda$, $\boldsymbol{\theta}_t \equiv \boldsymbol{\theta}$. Then, for $\mathbf{K}$ with $\underline{K} := \min_j K_j$, the following holds:*

$$p_{j,t}(\mathbf{K}) = p_{j,T}^* + \mathcal{O}(|T-t|^{\underline{K}}),\ t \to T \text{ for all } j,$$

i.e., price changes close to the deadline are at most of order $\underline{K}$. *If $\lim_{t\to T}(\Pi_{f,t})^{(\underline{K})}(\mathbf{K}-\mathbf{e}_{j'}) \neq 0$ for all f and j' with $K_{j'} = \underline{K}$, then*

$$p_{j,t}(\mathbf{K}) = p_{j,T}^* + \Theta(|T-t|^{\underline{K}}),\ t \to T \text{ for all } j,$$

i.e., price changes are exactly of order $\underline{K}$.

Proposition 2 is a generalization of the equilibrium property of Dudey (1992) and Martínez-de Albéniz and Talluri (2011) who consider undifferentiated products and deterministic demand. Without uncertainty and fixed demand, the firm with the least capacity deterministically sells out its products first before the next product is offered at an acceptable price. Note that while Proposition 2 only covers the case where the distribution of demands does not change over time, we empirically verify that this property also holds if demand becomes more inelastic over time.

2.6 The Bertrand Scarcity Trap

In this section, we discuss a novel welfare effect that arises in the presence of scarcity. Price competition can cause capacity to be mis-allocated to consumers with relatively low willingness to pay even if there are opportunities to sell to higher-valuation consumers in the future. In order to formalize this idea, assume that a consumer who arrives in period t has consumption utility for product j given by $u_{j,t} = \delta_j - \alpha_t p_t + \epsilon_{j,t}$ where $(\epsilon_{j,t})_{j\in\mathscr{J}}$ is drawn from some continuous distribution, e.g., type-1 extreme value. Then, $s_{j,t}(\mathbf{p};\mathscr{A}_t)$ denotes the probability that $u_{j,t} \geq \max_{i\in\mathscr{A}_t} u_{j,t}$. We further define a state-$(\mathbf{K},t)$ outcome to be a tuple $(j, u_{j,t})$ where j is an available product and $u_{j,t}$ is the realized utility level. An allocation rule a maps each state $(\mathbf{K},t)$ to a probability measure on the set of feasible outcomes in that state. An allocation rule can be induced by Bertrand price competition, a single firm's pricing decision, a social planner, or alternative pricing mechanisms. We formalize how an allocation rule can "add inefficiencies" over time. We denote the continuation welfare in state $(\mathbf{K},t)$, given an allocation rule a, by $W_t^a(\mathbf{K},t)$ and introduce the following definition.

Definition 1. We define a *state-*$(\mathbf{K},t)$ *constrained-efficient price given allocation rule* a to be

$$\bar{\mathbf{p}}_t^a(\mathbf{K}) = \left(\bar{p}_{j,t}^a(\mathbf{K})\right)_{j\in\mathscr{A}(\mathbf{K})} \in$$

$$\arg\max_{\mathbf{p}} \underbrace{\frac{1}{\alpha_t}\mathbb{E}\left[\max_{i\in\mathscr{A}(\mathbf{K})}\left(\delta_i - \alpha_t p_i + \epsilon_{i,t}\right)\right]}_{\text{consumer surplus}} + \sum_{j\in\mathscr{A}(\mathbf{K})} s_j(\mathbf{p};\mathscr{A}(\mathbf{K}))\left(p_j - \underbrace{\left(W_{t+\Delta}^a(\mathbf{K}) - W_{t+\Delta}^a(\mathbf{K}-\mathbf{e}_j)\right)}_{\substack{\text{forgone future} \\ \text{welfare of selling } j}}\right).$$

Thus, a price vector is state-$(\mathbf{K},t)$ constrained efficient given a if it maximizes the sum of consumer surplus and producer surplus minus the forgone continuation welfare if a unit is sold (taking future allocations as given) in state $(\mathbf{K},t)$. This allows us to distinguish between two types of inefficiencies that can occur in each stage game.

Definition 2. We say product j is *over-provided* if $p_j < \bar{p}_{j,t}^a(\mathbf{K})$ and it is *under-provided* if $p_{j,t} > \bar{p}_t^a(\mathbf{K})$.

Under static Bertrand price competition with differentiated products, we expect under-provision of products given that firms maintain some market power, and any restriction on competition typically exacerbates this inefficiency. However, in dynamic settings subject to scarcity, inefficient rationing can naturally occur due to over-provision of a product early on. Bertrand competition tends to exacerbate over-provision, so restricting competition can be welfare improving. We call this observation the *Bertrand scarcity trap*.

Theorem 2 (The Bertrand Scarcity Trap)**.** *Dynamic price competition with scarcity and a deadline can entail over-provision of products in some states* $(\mathbf{K}, t)$*, regardless if products are differentiated or not.*

We prove the theorem by constructing two simple examples where the Bertrand scarcity trap occurs—one with undifferentiated products, and one with differentiated products.

2.6.1 Illustrative Example

The following illustrative example shows that the Bertrand scarcity trap can be so severe that a single firm may price more efficiently than two competing firms. Consider two undifferentiated

Figure 3: Illustrative example of the Bertrand scarcity trap

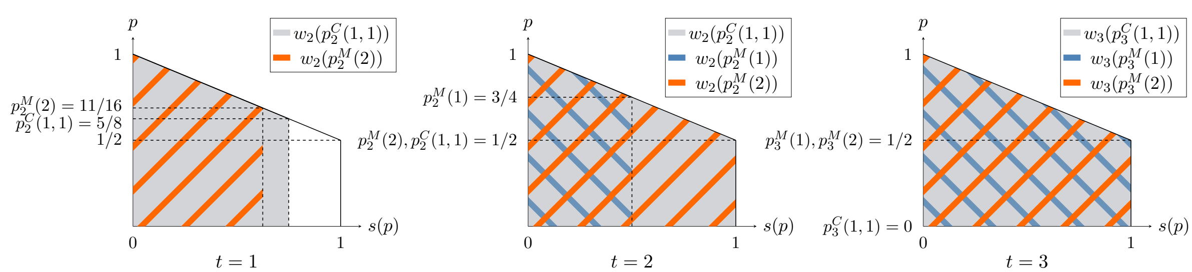

Notes: The graph depicts demand curves, single-firm optimal prices, and competitive prices in the three periods. The orange and blue regions represent per-period welfare given a single firm with two and one unit left, respectively. The grey region represents per-period welfare if two competing firms are active.

products that are available for sale over three sequential markets, $t = 1, 2, 3$. In every period t, a single short-lived consumer with i.i.d. unit demand arrives. If p is the lowest available price, a consumer in period t buys with probability $s_t(p) = 2(1-p)\mathbb{1}(p > 0.5) + \mathbb{1}(p \leq 0.5)$ as illustrated in Figure 3. We compare two market structures. In one market structure, each product is sold by competing firms, in the other, a single firm sells both products. We denote per-period welfare given a price p by $w_t(p)$, continuation welfare for a single merged firm with remaining capacity K by $W_t^M(K)$, and continuation welfare for two firms with capacity vector $(1,1)$, by $W_t^c(1,1)$. Prices are denoted by $p_t^M(K)$ and $p_t^c(1,1)$, analogously. In Figure 3, the per-period welfare is illustrated by the filled regions under the demand curves.

In the last period ($t = 3$), the monopoly price is $p_3^M(1) = \frac{1}{2}$, and monopoly profits are $\frac{1}{2}$. The equilibrium price and profits with Bertrand competition are 0. Total welfare is 0.75 in both settings. Thus, welfare in the last period is unaffected by market structure and maximized. Note that if two firms compete, a firm can gain $\frac{1}{2}$ in profits if the other firm sells in period 1 or 2.

Period-2 demand is identical to period 3. However, if a single firm has only one unit remaining, the firm sets a higher price equal to $p_2^M(1) = 0.75$ because it knows that there is another chance to sell this unit in period 3 yielding expected profits of $\frac{1}{2}$. The constrained-optimal price is also 0.75, so the single firm is pricing efficiently.[13] With $K=2$, the price is $p_2^M(2)=\frac{1}{2}$, which is also constrained optimal as only one unit can be sold in period 3. With competition, prices are the same, i.e., $p_2^c(1,1)=\frac{1}{2}$.[14] Finally, note that the continuation welfare given the allocation rule with a single firm and $K=1$ is $W_2^M(1)=0.8125$, with a single firm and $K=2$ is $W_2^M(2)=1.5$, and with competition is $W_2^c(1,1)=1.5$, respectively.

Moving to the first period, one can show that $p_1^M(2) = 0.6875 > p_1^c(1,1) = 0.625$. The constrained-efficient price for both monopoly and the competitive setting is 0.6875. Hence, the single firm is exactly solving a social planner's problem as all prices are constrained-efficient, while a competitive market is over-providing the product at $t=1$. Therefore, the equilibrium is subject to the Bertrand scarcity trap.

2.6.2 Bertrand Scarcity Trap with Logit Demand

Next, we show that the Bertrand scarcity trap can indeed occur within our framework. To do so, we simulate equilibria using the ODE characterization in Theorem 1. We assume a logit

[13] This is because $p=0.75$ maximizes $2p(1-p)+(1-p)^2+(1-2(1-p))0.75$.

[14] Note that firms do not have an incentive to deviate to higher or lower prices. In general, competition with undifferentiated products can lead to multiplicities and non-existence of symmetric pure-strategy equilibria as also shown in Talluri and Van Ryzin (2004) and Dudey (1992).

demand system with time-dependent price sensitivity α_t:

$$s_f(\mathbf{p}) = \frac{\exp\left(1-\alpha_t p_f\right)}{1+\sum_{f'\in\{1,2\}} \exp\left(1-\alpha_t p_{f'}\right).}$$

We simulate equilibrium and constrained-efficient prices for two different sets of parameters: (a) logit demand that is constant over time, and (b) logit demand with increasing willingness to pay over time. Figure 18 in Appendix D contains additional examples, including one that assumes nested logit demand. We focus on the simplest case with initial capacities of $\mathbf{K}_0 = (1,1)$. Note that this implies that prices are decreasing over time. The solid lines in Figure 4 represent equilibrium prices $p_t^*(1,1)$ and $p_t^*(1,0) = p^*(0,1)$. The dotted lines represent constrained-efficient prices given Bertrand competition. In panel (a), we only observe under-provision in all states—competitive prices are too high. In panel (b), we show that with increasing demand over time, there is over-provision early on and under-provision closer to the deadline for all capacities. Hence, scenario (b) is subject to the Bertrand scarcity trap.

Figure 4: Equilibrium prices and constrained-efficient prices under Bertrand competition

(a) Constant logit demand | (b) Logit demand with increasing α_t

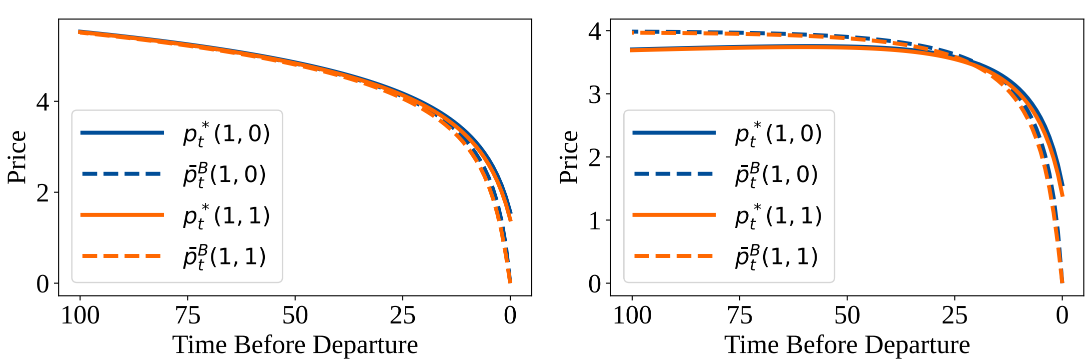

Notes: The simulations assume $\mathbf{K} = (1,1)$, $\lambda_t \equiv 1$, $\alpha_t = 1 + a(T-t)$. The parameters are as follows. In panel (a): $a = 0$; in panel (b): $a = \frac{1}{100}$.

Our simulations indicate that the Bertrand scarcity trap can be more severe when demand becomes more inelastic over time and when products are closer substitutes. Intuitively, this occurs because early sellouts are more costly from a welfare perspective.[15] However, it is

[15]Simulations in Figure 18 in Appendix D indicate that stronger competition between the inside goods lead to more severe over-provision early on.

not true that the Bertrand scarcity trap always occurs if demand (in terms of arrival rates or preferences) is “sufficiently increasing” or if products are close substitutes.

Finally, we note that because the Bertrand scarcity trap can result in over-provision early on and under-provision close to the deadline, efficiency may be improved by “compressing” prices over time. For this reason, we consider both uniform pricing and the use of pricing heuristics in our empirical analysis.

3 Data and Descriptive Evidence

We apply our framework to the US airline industry, a significant contributor to US economic activity, with over 811 million passengers flown, representing $196 billion in revenues for domestic travel in 2019 alone.[16] Although numerous insights on airline markets have been derived from publicly available data published by the Bureau of Transportation Statistics (BTS) (e.g., Berry, 1992; Borenstein and Rose, 1994, etc.), a key barrier in studying airline pricing dynamics has been a lack of available data. Publicly available data (such as the BTS DB1B and BTS T100 tables) are aggregated and are not suited for studying dynamic pricing at high frequencies.

Our study exploits new data provided to us through a research partnership with a large U.S. airline that enables us to investigate dynamic pricing in markets with competition. We merge internal data for a single airline with data that our research partner acquires from third parties. The third-party data has strong parallels with other contributed data sets, such as the the Nielsen scanner data that is commonly used in retail studies. Firms supply information on core business activities (such as pricing and demand information) to third parties that then assemble the data to be used for market intelligence purposes. Our data has all the same features of the Nielsen data, except that ours are not anonymized.

We first discuss our route selection criteria. We then provide an overview on the specific data variables that we use in our study before presenting preliminary evidence. Our data cover the first nine months of departures in 2019.

[16]Source: https://www.bts.gov/newsroom/2019-annual-and-4th-quarter-us-airline-financial-data and https://www.bts.gov/newsroom/final-full-year-2019-traffic-data-us-airlines-and-foreign-airlines.

3.1 Route Selection

We use the publicly available DB1B data to select routes to study. These data contain 10% of bookings in the US but lack information on the booking date, departure date, and flights involved. Our analysis concentrates on nonstop flight competition. We use the following selection criteria applied to the 2019 DB1B data:

i) We limit ourselves to routes where nonstop service is provided by exactly two airlines and remove routes where nonstop service is not provided by our airline partner. For expositional purposes we always refer to the second airline as "the competitor";

ii) We eliminate routes where the total number of nonstop quarterly traffic is less than 2,000 and greater than 50,000;

iii) We calculate the fraction of passengers who are not making connections for a given route, i.e., consumers traveling from A to B who do not make an intermediate connection at C. We remove routes in which the fraction of nonstop travel is less than 50%.

After imposing these selection criteria, we eliminate routes with incomplete third-party data. From this final set of over 500 routes, we randomly select 50 routes.[17]

Figure 5: Summary Analysis from the DB1B Data

(a) PDF of Nonstop Traffic (b) CDF of Passenger-Weighted Fares (c) Monthly Departure Count

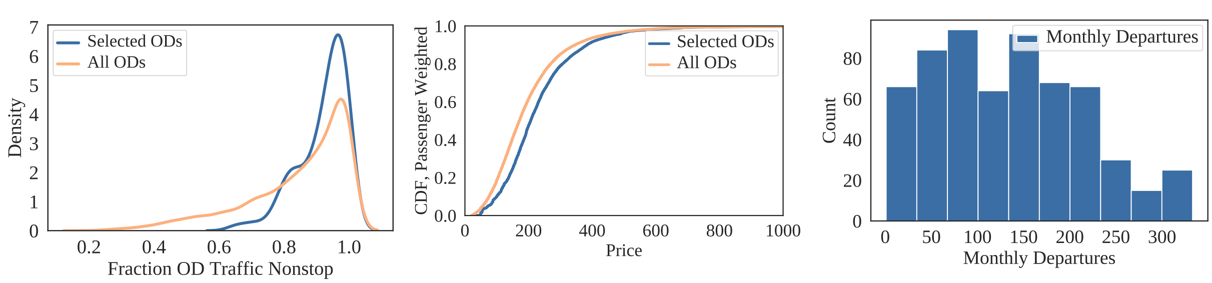

Note: Panel (a) records the PDF of nonstop traffic among local traffic in the DB1B data (orange) and for selected routes (blue). Panel (b) plots the CDF of prices for selected routes (blue) and all dual-carrier markets (orange). Panel (c) reports the number of aggregate monthly departures for the routes in our sample.

In Figure 5 we provide summary analysis of the 50 routes in our sample with a comparison to the DB1B data. In panel (a), we show the distribution of passengers flying nonstop (see

[17]We obtained data for only these 50 routes.

Section 3.2.1 for a formal definition). Our sample is skewed to the right on purpose as this limits the role of connecting options which we do not model. In panel (b) we show that the distribution of fares for our selected routes is similar to the universe of duopoly routes. Finally, in panel (c) we use the BTS T100 segment data to plot the total number of monthly departures for the routes in our sample. Over half of our sample contains routes in which there are fewer than five daily frequencies across both airlines between the origin and destination. Several routes feature twice-daily service (one flight per airline). At the other extreme, one route in our data contains nearly 10 flights per day.

3.2 Data Description

We next describe the novel data that we use in our empirical analysis.

3.2.1 Bookings Data

The bookings data detail flight-level sales counts at a daily frequency. We focus exclusively on economy-class bookings. We observe separate booking counts for different types of itineraries booked, but we do not observe the itineraries themselves. For example, consider an origin-destination pair (OD), or a route, denoted by $A \rightarrow B$. We observe the number of passengers who book nonstop travel for flights on route $A \rightarrow B$. We refer to these passengers as *local* passengers. Due to data availability, we do not observe booking counts of the form $A \rightarrow C \rightarrow B$. These are also local passengers who are traveling from A to B, but they make an intermediate connection at C. As previously stated, our route selection criteria are meant to minimize the importance of this type of traffic for consumer demand.

Our data also contains *flow* passenger counts. These are passengers who are traveling on itineraries such as $A \rightarrow B \rightarrow C$ or $C \rightarrow A \rightarrow B$. That is, these passengers are either using B as a connection point between $A \rightarrow C$, or, they are using A as a connection point between $C \rightarrow B$. This traffic is important to consider because it affects the overall availability of seats on a flight from $A \rightarrow B$. We again do not observe the exact itineraries booked, so we model flow traffic as exogenous reductions in remaining capacity. As a robustness check, we subtract all flow passenger bookings to define an alternative initial capacity condition.

Table 1: Synthesized Example of Booking Data

Route	Airline	Flight Num.	Dep. Date.	Days from Dep.	Local Pax.	Flow Pax.
$A \to B$	ID 1	FL 1	9/1/2019	90	1	0
$A \to B$	ID 2	FL 2	9/1/2019	90	0	1
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$A \to B$	ID 1	FL 1	9/1/2019	17	1	2
$A \to B$	ID 2	FL 2	9/1/2019	17	1	0
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$A \to B$	ID 1	FL 1	9/1/2019	0	2	3
$A \to B$	ID 2	FL 2	9/1/2019	0	3	1

Note: Synthesized data example of the booking data, which is at the route-airline-flight number-departure date-day before departure level. We observe the number of consumers who book nonstop travel (local pax.) as well as the number of consumers who book connecting itineraries for the routes we study (flow pax.).

We present a synthesized example of the booking data in Table 1. Our data are at the route-airline-flight number-departure date-day before departure level. Three other features of the bookings data are worth noting. First, the booking counts in our sample are complete in that they contain both tickets purchased directly with the airline and any alternative booking channel, e.g., an online travel agency.[18] Second, because we observe all bookings, we can construct each flight's load factor over time (total seats booked / capacity). This is a key state variable in our study. Third, as previously mentioned, we observe booking counts but not the associated itineraries themselves. For example, if a consumer purchases a round-trip ticket, we observe the associated booking counts for the outbound and return legs, but we do not know that those two "counts" are associated with a single itinerary. This may be a significant concern if there are significant round-trip discounts. We confirm that this is not the case for the routes we study because we observe all fares offered in a given market.

3.2.2 Pricing Data

We observe all potential fares consumers may face for the routes in our sample. It is typically the case that more than one fare may be available for purchase at a given point in time. This is a consequence of how airlines price, and need not conform to tickets of different qualities, such as

[18] In the industry, these bookings are often referred to as *direct* and *indirect* bookings, respectively.

first class versus economy tickets. We briefly describe airline pricing practices in Section 3.4. Each itinerary has a discrete set of potential prices consumers may face over time. We refer to the set of potential fares as a "fare menu." Our focus is on nonstop, economy class fare menus, which we observe for all airlines and routes in our sample. Note that these menus are never flight-specific, but they may be departure-date-specific. That is, all flights with the same departure date share the same fare menu. In practice, it is common that many departure dates share the same menu. We observe fare menus at the daily level and therefore, we observe if any fare value changes (see below for more details).

Figure 6 contains an example fare menu. The vertical axis denotes a fare level, and the horizontal axis denotes days from departure (DFD). For this particular airline-route, there are at most eight unique fares.[19] If the fare-DFD is filled in, this means it is possible that consumers may be offered this fare. Whether a fare is actually offered depends on how the airline's pricing heuristic allocates remaining inventory to these fares. The white area in the bottom right denotes advance purchase (AP) restrictions—it is not possible to book the lowest fares close to departure and therefore, they will never be available for purchase.

Figure 6: Fare Menu Example

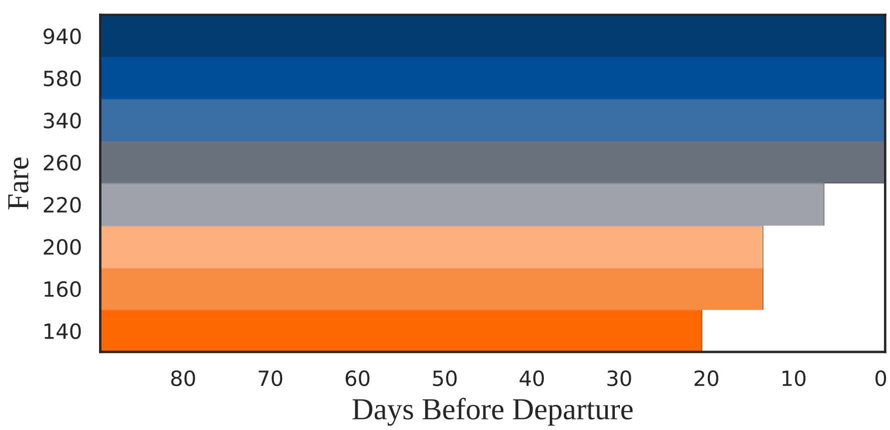

Note: Example pricing menu over time. Prices rounded to nearest $20.

In addition to the fare menus themselves, we observe the lowest available fare (LAF) that can be purchased at a given point in time. These data are flight-time specific because they depend on

[19]Note that within a row, it is possible for the fare level to change by day before departure. This is because each row corresponds to a fare class and not necessarily a unique fare. We observe whenever this occurs.

remaining capacity. The LAF for each flight is what we use as "price" in our demand analysis. We assume that all consumers purchase at the LAF because of recent empirical evidence that shows that over 91% of consumers purchase the LAF (Hortaçsu et al., 2021b).

3.2.3 Arrivals Data

We do not assume market sizes as is common in empirical studies. Instead, we leverage click-stream data provided to us by the air carrier. We use these data to construct a measure of arrivals, e.g., the number of consumers who searched for tickets at the route, departure date, day before departure level.

Allowing for market size variation is important in our setting because flights are subject to demand shocks. However, using search data for a single carrier understates true arrivals because consumers may search and purchase directly with competitor airlines or with online travel agencies. We address this complication by scaling up observed searches using a hyperparameter and conducting robustness exercises which we describe below.

3.3 Summary Analysis

Table 2: Summary statistics

Data Series	Variable	Mean	Std. Dev.	Median	5th pctile	95th pctile
Fares						
	One-Way Fare ($)	233.7	111.4	218.6	92.1	390.7
	Num. Fare Changes	6.4	2.4	6.0	3.0	11.0
Bookings						
	Booking Rate-local	0.2	0.6	0.0	0.0	1.0
	Booking Rate-all	0.5	1.2	0.0	0.0	3.0
	Ending LF (%)	72.1	19.8	76.0	32.9	98.0

Note: One-Way fare is for the lowest economy class ticket available for purchase. Number of fare changes records the number of price adjustments observed for each flight. Booking rate-local excludes flow traffic. Booking rate-all includes both local and flow traffic. Ending load factor (LF) reports the percentage of seats booked at departure time.

We provide a summary of the main data in Table 2. We limit our analysis to the last 90 days before departure due to overwhelming sparsity in bookings beyond 90 days. Average fares across airlines in our sample are $234. On average, each flight experiences over six price adjustments in 90 days. The average daily booking rate is less than one. Roughly 40% of observed bookings is local traffic, which is the demand we model in Section 4. The remaining are flow bookings, which we model using exogenous reductions in remaining capacity over time. At the departure time, average load factors are 72%. Roughly 3.5% of flights in our sample eventually sell out. We do not model overbooking but note that this practice may be welfare improving as it could allow firms to reallocate capacity to consumers with the highest valuations.

Figure 7: Prices and Bookings by Day Before Departure

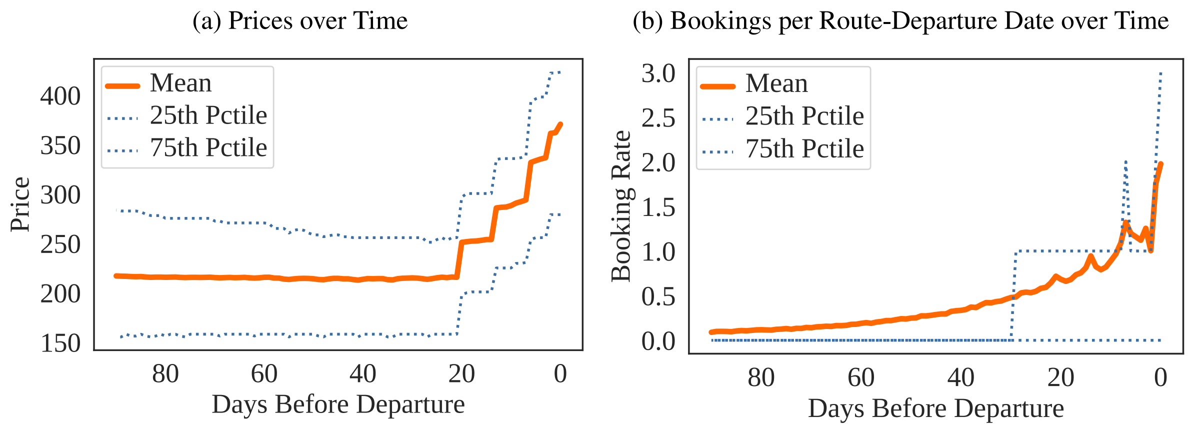

Note: Panel (a) shows the average and interquartile range of flight prices over time. Panel (b) shows the average and interquartile range of flight booking rates per route-departure date over time. Greater than 30 days before departure, the 25th and 75th percentiles coincide.

In Figure 7 we plot average fares and flight-level booking rates by day before departure. The left panel (a) shows that average fares are fairly flat between 90 and 21 days before departure. The top end of the distribution is decreasing in this time window. There are noticeable "steps" in the last 21 days before departure which highlights the use of advance purchase (AP) discounts as described in Section 3.2.2. In our sample, we typically observe AP requirements at 21, 14, 7, and 3 days before departure. Note that fares increase by over 70% in three months. In the right panel (b) we highlight that bookings increase as the departure date approaches. This coincides with increasing prices suggesting that demand becomes more inelastic over time. The booking rate is greater than one per flight over the last month before departure.

In Figure 8 we compare outcomes across competitors. The left panel (a) provides a scatter plot of load factors at departure across airlines. The orange squares present route-level averages.

Note there exists a large mass of points both above and below the 45-degree line. We find that no airline consistently sells a larger fraction of capacity than the other carrier for all routes. The blue dots correspond to route-departure date averages by airline. We do observe a few departure dates in which airlines sell all their inventory. In our empirical analysis, we do not model flights after sell out. In the right panel (b), we plot the average fare difference across airlines over time when exactly two flights are offered. Note that fares tend to be similar—the average difference is less than \$10. One airline has relatively higher prices well in advance of departure and relatively lower prices close to departure. More importantly, we find that prices across airlines are nearly equal 50% of the time. This is a consequence of airlines having identical prices in their fare menus.

Figure 8: Load Factor and Price Differences across Carriers

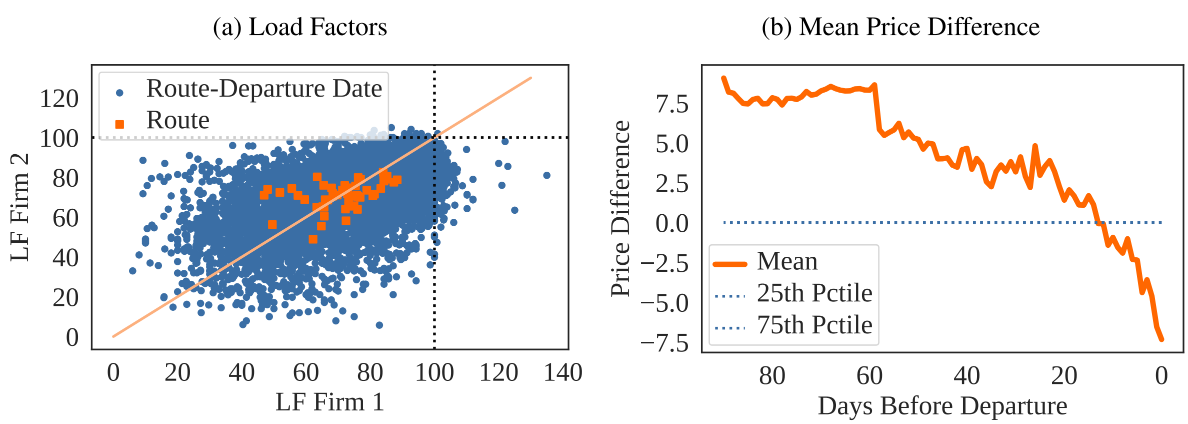

Note: Panel (a) shows the average load factor (across all flights) at the route-departure date level for both competitors in blue. The orange squares report average route-level load factors. The diagonal line is the 45-degree line. Panel (b) shows the average and the 25th and 75 percentiles of the difference in prices for markets in which each firm offers exactly one flight.

3.4 Pricing Heuristics

We observe how one airline has designed and operationalized its pricing system. We use this information to guide our supply-side models. Reviewing the airline's pricing documentation and optimization code itself, we find that airline pricing practices differ substantially from our competitive equilibrium benchmark model in a number of important ways. First, no competitor pricing information enters the pricing algorithm. The algorithm used by our airline is not a reinforcement learning algorithm as studied in recent applied theory work. It originated in operations research to approximate the single-firm optimal control problem solution and focuses on time-varying opportunity costs due to aggregate demand uncertainty. In addition to not incorporating competitor prices, the heuristic does not internalize the existence of any competition, regardless of market structure. Every flight is optimized in isolation.

How is this consistent with observing identical prices across airlines (Figure 8-b) when the observed pricing heuristic does not incorporate any competitor information? Price matching is possible because managers that decide fare menus have matched competitor fares. That is, the example fare menu shown in Figure 6 for a single airline may partially or completely match the fare menu chosen by a competitor airline. Hortaçsu et al. (2021b) shows that all major airlines have teams that monitor competitor fare menus and either initialize or respond to competitor actions. While fare menus naturally limit which prices consumers may face and enable price matching, we confirm that this is not strategically chosen by the pricing heuristic.

We do not endogenize the fare menus in our analysis because nearly 40% of prices within a fare menu do not change within our sample. That is not to say flight prices (LAF) do not change over time, but rather, the discrete set of fares from which LAF is determined by do not change across calendar time. We find that 20% of fare menu prices change once. The maximum number of unique prices across fare menus we observe in the sample is 11; however, this represents only 0.01% of observed fare menus. Studying the strategic implications of designing fare menus is beyond the scope of this paper.

Given these observations, we design two heuristics for our counterfactual analysis that reflect these observed constraints. We refer to the heuristics as the "lagged model" and the "deterministic model." Both heuristics use discrete fare menus and do not internalize competitors as a strategic player. The heuristics differ in how the single-agent dynamic control problem is constructed.[20] In the lagged model, each firm, having observed its competitor's last period price (LAF), assumes that this price will also be charged in the current and all future periods. If that LAF disappears due to an AP opportunity expiring, the next lowest fare is assumed to be offered for the remaining time. Each firm then calculates its residual demand curves in all remaining

[20]We do not consider the pricing heuristic used in Hortaçsu et al. (2021b) due to data availability, e.g., the demand forecasts used by all airlines.

periods and solves a single-firm dynamic programming problem. In the deterministic model, each firm simply assumes its competitor will price at the lowest possible fare on the menu in all remaining periods. That is, the lagged model relies on LAFs whereas the deterministic model leverages the fare menu information only. We assume that firms believe their competitor will be at the minimum fare on the menu because more than 50% of observed prices are at the minimum value on their respective pricing menus. We assume all firms use the same pricing heuristic.[21]

4 Demand Model and Estimates

4.1 Empirical Specification

We model nonstop air travel demand using a flexible nested logit demand model. Our approach differs from some recent empirical work on airlines that assume a mixed-logit specification of "business" and "leisure" travelers (Lazarev, 2013; Williams, 2022; Aryal et al., 2022; Chen and Jeziorski, 2023; Hortaçsu et al., 2021b). In these studies, two price sensitivity parameters are estimated as well as the fractions of arriving consumer types over time. We do not consider mass-point random coefficients, rather, we allow for time-specific price coefficients. We pursue this approach because it maps to our theoretical model and results in unique equilibrium price paths. We found that mixed-logit models commonly result in four stage-game equilibria, requiring both full characterization of stage game equilibria and an equilibrium selection mechanism. Our approach allows for a rich demand system and avoids solving stage game equilibria explicitly due to Theorem 1.

We define a market as an origin-destination (r), departure date (d), and day before departure (t) combination. Each flight j, leaving on date d, is modeled across time $t \in \{0,...,T\}$. The first period of sale is $t = 0$, and the flight departs at T. Demand is modeled at the daily level over a 90-day horizon. Arriving consumers choose a flight that maximizes their individual utilities from the choice set $\mathscr{J}_{t,d,r}$, or select the outside option, $j = 0$.

[21] See Brown and MacKay (2021) on work on pricing algorithm choice.

We specify consumer arrivals to be

$$\lambda_{t,d,r} = \exp\left(\tau_r^{\text{OD}} + \tau_d^{\text{DD}} + \tau_{t,d}^{\text{SD}} + f(t)\right),$$

where the τs denote fixed effects for the route (OD), departure date (DD), and search date (SD). We model $f(\cdot)$ using a polynomial series of degree three. We scale up these estimated arrival rates using hyperparameters to account for unobserved searches. Smoothness of $f(\cdot)$ allows us to use our ODE equilibrium characterization.

Conditional on arrival, we specify consumer utilities as

$$u_{i,j,t,d,r} = \mathbf{x}_{j,d,r}\boldsymbol{\beta} - \alpha_t p_{j,t,d,r} + \zeta_{i,J} + (1-\sigma)\varepsilon_{i,j,t,d,r},$$

where $\zeta_{i,J} + (1-\sigma)\varepsilon_{i,j,t,d,r}$ follows a type-1 extreme value distribution, and $\zeta_{i,J}$ is an idiosyncratic preference for the inside goods. Products are partitioned into two nests. The outside good belongs to its own nest, and all inside goods to the second nest. The parameter $\sigma \in [0,1]$ denotes correlation in preferences within the nests. We allow the consumer price sensitivity (α_t) to vary over time using three-day intervals of time; hence, we estimate 30 price sensitivity parameters. We include a number of covariates in $\mathbf{x}$ that are assumed to not vary across t: departure week of the year, departure day of the week, route, carrier, and departure time fixed effects. In Section 4.4, we discuss an extension of this baseline model that includes an additional unobservable (ξ) that is potentially correlated with price. Arriving consumers solve their utility maximization problems: consumer i chooses flight j if and only if

$$u_{i,j,t,d,r} \geq u_{i,j',d,t,r}, \forall j' \in \mathscr{J}_{t,d,r} \cup \{0\}.$$

We define (dropping the t,d,r subscripts)

$$D_{\mathscr{J}} := \sum_{j \in \mathscr{J}} \exp\left(\frac{\mathbf{x}_j\boldsymbol{\beta} - \alpha p_j}{1-\sigma}\right),$$

so that the probability that a consumer purchases j within the set of inside goods is equal to

$$s_{j|\mathscr{J}} := \frac{\exp\left(\frac{\mathbf{x}_j\boldsymbol{\beta}-\alpha p_j}{1-\sigma}\right)}{D_{\mathscr{J}}}.$$

It follows that the probability that a consumer purchases any inside good product is equal to

$$s_{\mathscr{J}} := \frac{D_{\mathscr{J}}^{1-\sigma}}{1+D_{\mathscr{J}}^{1-\sigma}}.$$

Overall product shares are equal to $s_j = s_{j|\mathscr{J}} \cdot s_{\mathscr{J}}$, which are at the market level (t, d, r). Our assumptions imply that demand is distributed Poisson with a flight-level booking rate of $\min\left\{\lambda_{t,d,r} \cdot s_{j,t,d,r}, K_{j,t,d,r}\right\}$, where $K_{j,t,d,r}$ denotes remaining inventory.

4.2 Estimation Procedure

We estimate the model using a two-step estimation approach. In the first step, we estimate the arrival process parameters using Poisson regressions. We then estimate preferences of the Poisson demand model using maximum likelihood. We estimate standard errors using bootstrap.

We follow Hortaçsu et al. (2021b) in constructing arrivals using clickstream data for one airline. We count the number of searches corresponding to each market (r, d, t), and then scale up estimated arrival rates to account for unobserved searches. We apply the property of Poisson distributions that the sum of Poisson variables is Poisson with added intensities. We assume that consumers who search/purchase through alternative platforms (travel agents, other airlines' websites) have the same underlying preferences. We use the fraction of direct bookings by day before departure as weights when we scale up the estimated arrival rates. This adjusts arrivals for a single carrier. In our preferred specification, we then double these arrival rates to account for competitor indirect and direct searches, both of which are unobserved to us. Our demand estimates do not vary substantially under alternative scaling parameters (see Section 4.4).

4.3 Identification

In empirical work, it is customary to treat the market size as given. We use arrivals data to discipline our demand estimates and recover changes in willingness to pay over time. Without access to arrivals data, it is difficult to estimate preferences in models with demand uncertainty because we cannot distinguish between data generated by many arrivals and price sensitive consumers, or few arrivals and price insensitive consumers (Vulcano et al., 2012). Our arrivals data estimate that market participation increases over time in all routes studied, which informs the attractiveness of travel relative to the normalized outside option. For example, if we assumed market sizes were constant over time when they are actually increasing over time, we would estimate early demand as being too elastic and late demand as being too inelastic.

Stochastic demand allows us to measure demand response to price changes. In the model, every booking changes the opportunity cost for the next unit and results in a discontinuous price jump. This also reflects how the observed pricing system operates: a booking can close the availability of a given fare on the fare menu at a random time. Our identification of demand of the baseline model is inspired by the regression discontinuity literature. Within intervals of time (three day intervals), where α_t is fixed and the arrival process has been estimated, bookings cause jumps in opportunity costs and hence, prices (as well as choice set variation). Over 60% of observed price changes occur before advance purchase discounts expires. The timing of advance purchase discounts for a given route does not change over time, and it is uncommon for fares on the fare menus to change over time (see Section 3.4). Our identification approach relies on measuring the observed demand response around price jumps. As a result, we can measure interval-specific price elasticities. A price change for one firm informs substitution patterns to other products versus the outside good (σ). The fixed effects are identified by booking rate differences across weeks of the year, route, times of the day, and airlines. In our demand model extension that allows for an additional unobservable to be correlated with price, we instrument for demand using cost (scarcity) shifters (see next subsection).

4.4 Demand Estimates

We summarize our demand estimates in Table 3. The nesting parameter is estimated to be 0.5 implying substantial substitution within inside goods. We estimate a significant change in consumer price sensitivity over time, which is shown in Figure 9-(a). Almost all of our controls are significant, with day of the week and week of the year having the strongest influence on market shares. The competitor FEs are less important in driving variation in shares. We estimate the average own-price elasticity to be -1.44 (s.d. = 0.81), indicating slightly more elastic demand than in Hortaçsu et al. (2021b), which uses similar data for single-carrier routes.

Table 3: Demand Estimates Summary Table

Variable	Symbol	Estimate	Std. Error.	Range	% Sig.
Nesting Parameter	σ	0.498	0.010	–	–
Price Sensitivity	α	–	–	[-0.511,-0.074]	100.0
Competitor FE	–	–	–	[0.000,0.071]	100.0
Day of Week FE	–	–	–	[-1.637,-0.961]	100.0
Departure Time FE	–	–	–	[-0.462,-0.050]	100.0
Route FE	–	–	–	[-0.177,0.226]	94.4
Week FE	–	–	–	[-0.953,0.699]	86.0
Sample Size	N		2,814,686		
Average Elasticity	e^D		-1.438		

Note: Demand estimates for the 50 routes in our sample. Demand model contains 304 parameters. Standard errors for the two-step estimation procedure computed using 100 bootstrap samples. Log-likelihood value equal to -1,319,714.4.

In Figure 9-(b), we plot average adjusted arrival rates as well as different percentiles (5%, 25%, 75%, 95%) across markets. For each route, we estimate just a few arrivals 90 days before departure that rise to over 10 passengers per day close to departure. Note that while the 75th percentile closely follows the mean, the top part of the distribution is substantially higher, which corresponds to the routes with a larger number of departures and overall bookings. Empirical shares are also increasing, with the aggregate inside share peaking at 25% within the last few days of departure.

Before turning to counterfactuals, we briefly discuss additional demand results. Our demand estimates are robust to the choice in scaling factor. We find that average demand elasticities with scaling parameters between 1.0 and 3.5 are between -1.40 (s.d. = 0.78) and -1.46 (s.d. = 0.82). We have also estimated a demand model that incorporates an additional unobservable (ξ) that is potentially correlated with p. We use a 2-step estimation procedure where we first estimate the arrival process parameters and then use a control function to estimate the demand parameters using quasi-maximum likelihood estimation. Included in our set of instruments is a polynomial expansion of remaining inventory, indicators for AP fares, and the number of flights available for purchase. With this approach, we estimate average demand elasticities to be -1.59, with a standard deviation of 0.87. These estimates are also robust to the choice in scaling parameter.[22]

Figure 9: Price Sensitivity and Arrival Rate Parameters

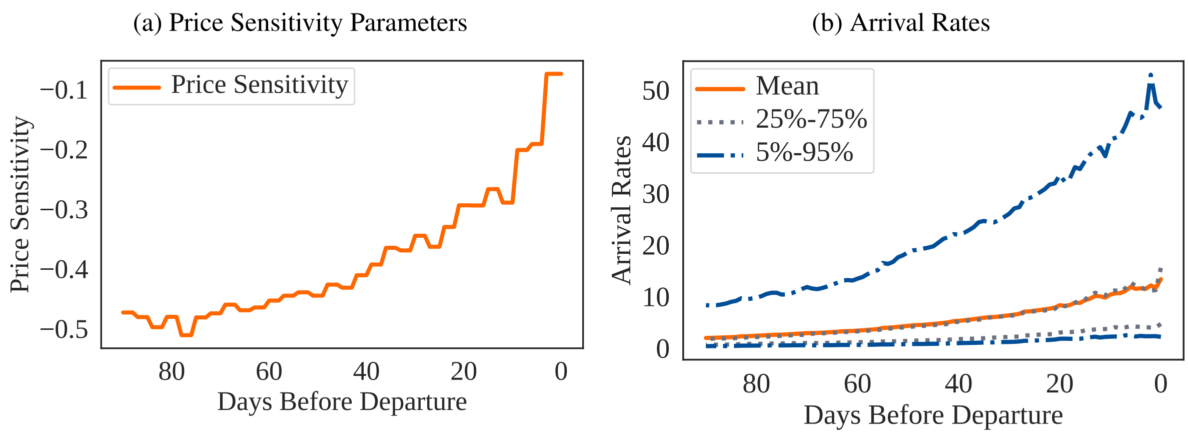

Note: Panel (a) shows our estimates of the price sensitivity parameters in 3-day groupings. Panel (b) shows fitted values of arrival rates over time adjusted for unobserved searches. The mean is the average arrival rate across all markets. The percentiles are also over markets.

5 Counterfactual Analysis

With demand estimated, we conduct a series of counterfactual exercises. We first discuss our counterfactual setup and implementation strategy before presenting our findings.

[22]We do not use this specification as our baseline model as the pricing equation does not account for the presence of multiple unobservables in shifting price (Petrin and Train, 2010). Nonetheless, Hortaçsu et al. (2021a) provide some Monte Carlo evidence that the magnitude of the bias may not be severe in similar data generating environments. Implementing all counterfactuals using these demand estimates, we find quantitatively similar effects.

5.1 Counterfactual Design

Competitive Dynamic Pricing. We approximate the continuous-time model to solve for equilibrium prices for every route-departure date in our sample. Both firms start with initial capacities K_f and $K_{f'}$, which we take from the data. We apply an ODE solver algorithm to the system of ODEs derived in Theorem 1. We tune the solver based on estimated arrival rates. For example, when the daily arrival rate is low, we can approximate the model using a coarse time interval Δ. When the arrival rate is high, we adjust the interval so that pricing decisions are at a finer level. Our approximation ensures that the expected number of arrivals within a time interval is less than one. Within an interval, the ODE system is calculated using a fourth-order Runge-Kutta (RK4) solver.[23] By leveraging our equilibrium characterization, we avoid explicitly solving the following approximate number of stage games in our non-stationary, finite-horizon setting:

$$\underbrace{R}_{\text{Routes}} \times \underbrace{D}_{\text{Dep. Dates}} \times \underbrace{T}_{\text{Days From Dep.}} \times \underbrace{\frac{1}{\Delta}}_{\text{Time Interval}} \times \underbrace{J \times F}_{\dim(\Omega)} \times \underbrace{\Pi_{f,j} K_0}_{\text{initial capacities}} \approx 59 \text{ billion.}$$

We store Ω_t and $\mathbf{p}_t$ every 24 hours due to memory constraints (the policy functions alone would be over 3TB in size). Our implementation of the dynamic competitive equilibrium outcome sets prices to be constant within a day. This maps well to our empirical setting as prices are adjusted daily—our airline reopitmizes remaining inventory once a day. To simulate demand given the policy functions, we use multinomial distributions after drawing arrivals from their corresponding Poisson distributions. When demand exceeds remaining inventory, we assume random rationing in all counterfactuals.

Because our equilibrium characterization involves differentiable pricing policies, we verify that our simulated pricing policies exhibit no jumps. We do not detect multiplicities in solving the dynamic game for all route-departure date combinations across all counterfactuals.

Uniform Pricing. We begin by comparing the dynamic competitive equilibrium outcome to

[23] In order to test the accuracy of the RK4 solver for our case, we analytically solve a simple single-firm problem with constant demand and compared the solver's performance to using Euler's method. The relative error of the Euler method solver was about 10^{-3}, while for the RK4 solver it was 10^{-9}, resulting in a solution 6 orders of magnitude more accurate.

another natural benchmark model—uniform pricing. With uniform pricing, firms set a single price for each route-departure date. To solve for the uniform pricing equilibrium, we use iterative best response. Given each competitor price from the last iteration, we simulate 10,000 flights to compute expected revenues. We solve for the optimal price and iterate across best-response functions until convergence.

Pricing Heuristics. Recall that the two pricing heuristics we describe in Section 3.4 capture key characteristics of the observed pricing technology for one airline: (i) it does not condition on competitor scarcity, (ii) it does not consider competition as a strategic player, and (iii) prices are determined from a discrete set. For the lagged model, we initialize the fare entering residual demand to be the lowest fare on the competitor's pricing menu. Expectations of residual demand adjust after AP fares expire. We use our fare data to construct route-specific fare menus that firms take as given.

Social Planner and Single-Firm Solutions. We also solve the continuous-time social planner and single firm problems. For the single-firm scenario, we assume a single firm manages both flights and jointly maximizes flight revenues. The social planner also manages both flights while maximizing total welfare. We use the ODE characterization in Lemma 1 for the the simulations of the single-firm counterfactual. The social planner's problem can also be characterized by ODEs (see Appendix A and C).

Additional Implementation Details. To implement all counterfactuals, we conduct 10,000 Monte Carlo experiments for every route, departure date combination. We smooth α_t using a polynomial regression in order to avoid discontinuities in the time derivatives ($R^2 = 0.974$). We simulate all counterfactuals twice, once where flow traffic is subtracted from initial observed capacity in advance, and one where flow traffic is modeled through Poisson processes that make inventory units disappear independent of the price. We report the latter specification here. Appendix D contains the former approach. Both the direction and magnitude of all effects are similar across specifications.

In the main text, we focus on 18 duopoly markets where each airline offers exactly one flight. In Appendix D, we report results for all routes. The reason we separate the counterfactuals is that with more than two flights, solving for equilibria of the dynamic pricing game becomes computationally challenging even with our characterization. The dimensionality of the state space and policy functions grows exponentially. For example, consider a route where airlines operate 100 seat aircraft. Let $a{:}b$ denote the number of flights operated by airline a and airline b, respectively. Our main analysis focuses on 1:1 routes. A 2:1 has 150 times more states (13 bil.). A 2:2 route has 20,000 times more states (1.7 tril.). Examining more complex market structures likely requires state aggregation, which we leave for future research.

In Appendix D, we consider all routes, but we reduce the number of flights studied. To do this, we adjust the choice set, utilities, and capacities for routes where an airline offers multiple flights a day. Appendix D contains details of the procedure as well as the counterfactual results. Both the direction and magnitude of the overall welfare effects for the entire sample are consistent with the 18 routes reported here.

5.2 Competitive Forces in Dynamic Pricing Games

Before reporting our welfare estimates, we describe the competitive forces present in our estimated model. We find both positive and negative scarcity effects. In fact, they vary within a particular route-departure date (see Figure 19-(a) in Appendix D). However, in general, we find that scarcity effects tend to not change signs. Checking the cross derivatives of the best response, only 0.5% of stage games (5,900,000 games) are not a game of strategic complements. This is likely also the reason why we do not observe multiplicity of equilibria. Own-scarcity effects tend to remain positive because own scarcity drives up prices. Similarly, competitor-scarcity effects tend to remain negative because the sale of a competitor typically increases future prices. Continuation profits are indeed non-monotonic in own capacity.

We plot average scarcity effects in Figure 10. Own-scarcity effects are largest around three weeks before departure. This is because selling a unit decreases a firm's continuation payoff the most when inventory is sufficiently scarce. However, note that the curves are U-shaped and start tending toward zero within a few weeks of departure. This occurs because once a firm sells out, it is not reported in the average, i.e., peak flights with higher scarcity effects drop out of the sample. Scarcity effects also drop toward zero close to the deadline as there are fewer opportunities to sell.

Figure 10: Benchmark Model Scarcity Effects

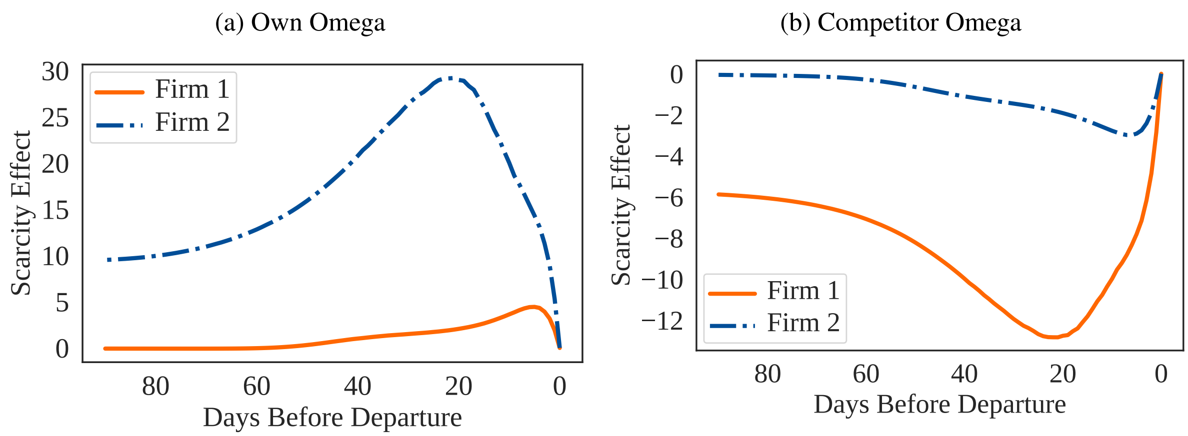

Note: Panel (a) reports the own-firm scarcity effect over time for both firms. Panel (b) reports the cross-firm competitor scarcity effect over time for both firms.

We also find that scarcity effects are asymmetric across airlines due to preferences and initial capacity differences. Competitor-scarcity effects tend to be larger for our airline research partner than its competitors. This asymmetry implies that the sale of one airline softens competition more than a sale of the other airline. At the same time, own-scarcity effects tends to be larger for competitor airlines. This is because competitor airlines typically offer planes with lower capacities than our research partner.

Finally, we investigate the insights from Proposition 2 in our data in Figure 19 in Appendix D. We find that the sale from the firm with the minimum inventory remaining results in the largest price effect. Price increases when the firm with the minimum inventory remaining sells are over five times ($30) greater than the price increase ($5) if the firm with more seats remaining sells closer to departure.

5.3 Welfare Comparison to Uniform Pricing

We report market outcomes comparing the dynamic competitive outcome to uniform pricing in Table 4. Average prices in our benchmark simulations ($226) are close to observed prices ($234). Prices are 10% higher under uniform pricing ($250). Although uniform pricing features higher average prices, revenues are substantially lower for both firms. The revenue effects are

Table 4: Counterfactual Results for Single Product, Duopoly Routes

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	225.9	5369.1	5557.1	16104.1	27030.2	19.3	70.7	9.2
Uniform	250.1	4464.3	4756.1	18357.7	27578.1	18.6	69.9	8.0
% Diff.	10.7	-16.9	-14.4	14.0	2.0	-3.8	-0.8	-1.2

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.

driven by relatively higher fares for early-arriving, price sensitive customers and relatively lower fares for late-arriving, price insensitive customers under uniform pricing. Dynamic pricing expands output due to lower prices early on. This can be seen in Figure 11-(a), which shows purchase probabilities over time. Figure 20 in Appendix D plots sellouts and load factors.

Figure 11: Competitive Dynamic Pricing and Uniform Pricing

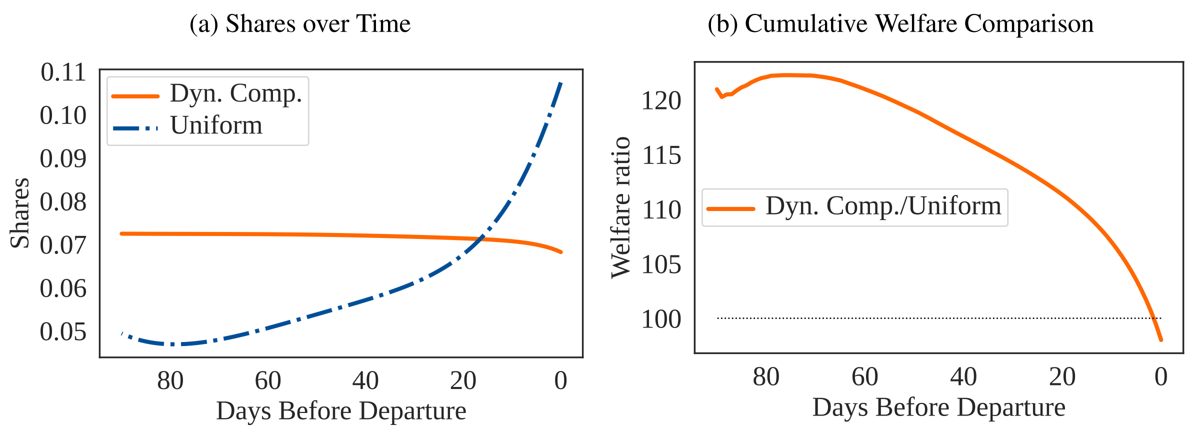

Note: Panel (a) shows the average shares over time for the benchmark and uniform models. Panel (b) shows the ratio of average cumulative welfare for the benchmark model with respect to the uniform one.

While we find that total output is higher, total welfare is lower under the dynamic competitive equilibrium. This contrasts with recent empirical studies in the single-firm setting, where dynamic pricing has been found to increase welfare (Hendel and Nevo, 2013; Castillo, 2022; Williams, 2022).[24] Consumer surplus is 14% higher with uniform pricing, which is larger in magnitude than the associated revenue losses (between 14-17%) of not adjusting prices based

[24]Note that in both settings the welfare comparison is theoretically ambiguous.

on demand and scarcity. The welfare loss is driven by high fares close to the departure date, when demand intensity and willingness to pay are high. This can be seen in Figure 11-(b), which plots the ratio of cumulative welfare under dynamic pricing over uniform pricing.

These welfare effects are robust along a number of dimensions. First, at the route level, we find no routes for which welfare under dynamic pricing is higher than under uniform pricing. Second, our results are robust to how we handle flow traffic (see Figure 21 and Table 6 in Appendix D). Finally, our results also hold when we investigate the entire data sample (see Table 10 in Appendix D).

5.4 Welfare Comparison to Pricing Heuristics

Counterfactual results comparing the dynamic competitive equilibrium outcome to the use of heuristics appear in Table 5. Figure 12 plots market outcomes over time. We normalize market outcomes under the dynamic competitive equilibrium model to 100 and report percentage differences for the heuristics. At a high level, we find that heuristics: (i) raise prices, (ii) raise total revenues, (iii) raise consumer surplus, and (iv) result in higher total welfare than under the dynamic competitive equilibrium. Our results provide a strategic reason for the use of pricing heuristics—they result in higher total revenues than under dynamic price competition.[25] Interestingly, both models result in significant price matching, at the same frequency as observed in the data (between 58-60%).

Table 5: Heuristic Counterfactuals for Single Product, Duopoly Routes

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Lagged	107.1	104.4	105.1	102.9	103.6	99.4	99.9	98.6
Deterministic	100.6	99.8	100.6	107.8	104.7	103.3	101.1	106.1

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.

[25] We do not study pricing heuristic adoption because of the difficulty in analyzing outcomes if only one firm uses a heuristic while the other firm solves the full dynamic problem.

We find that the deterministic model primarily benefits consumers (+8%) because it places maximal downward pressure on prices—firms believe their competitors will always be capacity unconstrained and charge the lowest fares on the menu. This expands output on average, increases load factors, and increases the frequency of sell outs. Average prices are slightly higher than under the competitive equilibrium outcome due to the use of discrete prices. We find that average prices are within 1% of the competitive equilibrium outcome, and while overall revenues improve, firm 1's revenues are essentially unchanged (-.2%). Firm 2's revenues increase by 0.6%.

Figure 12: Competitive Dynamic Pricing and Pricing Heuristics

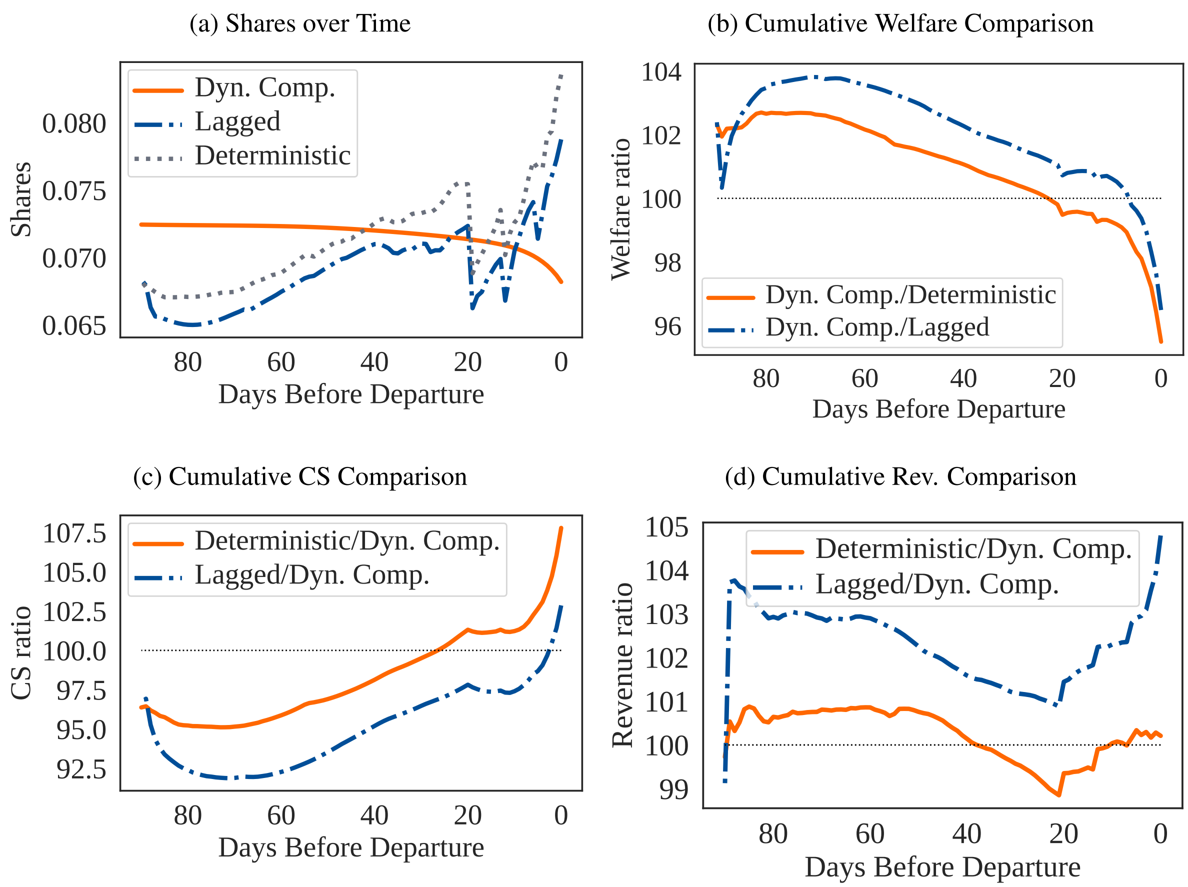

Note: (a) Product shares over time. (b) Cumulative welfare of competitive equilibrium relative to heuristics over time. (c) Cumulative consumer surplus of competitive equilibrium relative to heuristics over time. (d) Cumulative revenue of competitive equilibrium relative to heuristics over time.

The lagged pricing model adjusts residual demand outside of predefined advance purchase requirements. We find that this heuristic benefits firms and consumers alike, however, it supplies more benefits to firms (4-5% revenue increase) than to consumers (3% surplus increase). The reason is that demand shocks tend to raise own-scarcity effects until sufficiently close to the deadline. This causes a firm to raise its price. Its competitor, having observed the increased fare, is incentivized to also increase its price. On the flip side, this effect also causes prices to decrease towards the departure date as scarcity effects tend to decline. Hence, the gap in prices between the deterministic and lagged models closes.

These findings are also robust across routes and how we handle flow traffic. We find only one route in which welfare under heuristics is lower than under the dynamic competitive outcome. In Figure 7 in Appendix D, we show that the revenue and consumer surplus effects under heuristics with restricted capacities yields similar findings.

5.5 Discussion of Findings

Our framework allows us to uncover why pricing heuristics and uniform pricing increase welfare relative to the competitive equilibrium outcome. We show that some airline markets are subject to the Bertrand scarcity trap. We calculate the constrained efficient prices (see Definition 1) and compare it to the competitive equilibrium outcome. Figure 13 shows an example route that is subject to the Bertrand scarcity trap: in a significant fraction of states we observe over-provision for at least one firm well before the departure date. More precisely, we calculate the fraction of states at time t with $\min_j \mathbf{K} = K$ where at least one of the competitive prices is lower than its corresponding constrained-efficient price. This fraction is represented as a shade of grey, with white being 100% of states and black being 0% of states. The orange (solid) line tracks the median flight, and the blue (dotted) lines represent the 10th and 90th percentiles across simulations. For this route, competitive prices are inefficiently low for the first 30 days prior to departure. Close to departure, we find that both airlines tend to under-provide.

Across all routes and simulations, we estimate that 3% of posted prices are subject to over-provision due to the Bertrand scarcity trap.[26] Hence, our welfare effects are driven by both over-provision of products early on and under-provision close to the departure date. In fact, we find

[26] Figure 24 in Appendix D plots the percentage of flights subject to the Bertrand scarcity trap over time.

Figure 13: Bertrand Scarcity Trap in Example Route

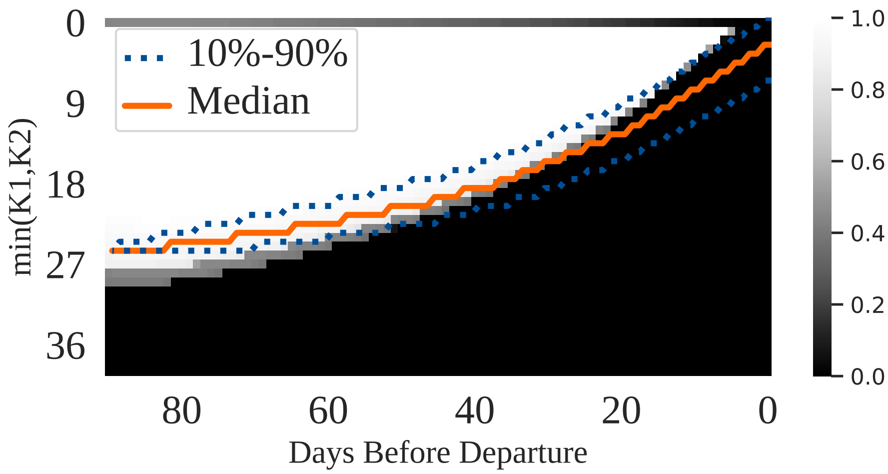

Note: The heatmap at (t, K) represents the fraction of states at time t such that $\min_j \mathbf{K} = K$ that are subject to the Bertrand Scarcity Trap. The orange (solid) and blue (dotted) lines represent the median and 10th-90th percentiles of minimum capacity observed in simulations, respectively.

that there is under-provision close to the deadline in all counterfactuals performed (competitive outcome, uniform pricing, heuristics) compared to the social planner's solution. Under the dynamic competitive outcome, welfare is 88% of the first-best outcome, and heuristics obtain 92.5% of the first-best (see Tables 8-9 in Appendix D).

Our results also showcase the benefits of price competition in airline markets more generally. Simulating revenues where both flights are managed by a single firm, we find that welfare goes down 16% relative to the dynamic competitive equilibrium outcome. That is, while the competitive outcome involves some inefficiencies due to the Bertrand scarcity trap and from under-provision close to the departure date, competitive forces have a net positive effect on welfare and come closer to the efficient outcome.

6 Conclusion

In this paper we introduce a framework to study dynamic pricing of perishable goods in an oligopoly. We show that dynamic price competition may involve inefficiencies because of a new competitive force that we call the *Bertrand scarcity trap*: competition can lead to over-provision of bookings early on leading to under-provision of bookings close to departure . We empirically estimate our model using daily prices and bookings for multiple, competing airlines. We compare market outcomes under the competitive equilibrium outcome to a scenario where airlines do not respond to competitor scarcity and use discrete prices. Our examination of one airline's internal pricing systems demonstrates that such constraints exist in its system. Our main empirical finding is that the adoption of such constraints in the pricing systems soften competition and allow airlines to avoid under-provision of bookings close to departure when demand is more price insensitive. In the data, this under-provision is caused in part by the Bertrand scarcity trap, but also by market power due to product differentiation. Competing heuristic pricing mechanisms increase both firm revenues and consumer welfare because capacity is used more effectively.

We see several potential directions for future work. First, it may be possible to examine dynamic price competition in settings with larger state spaces (more firms and/or products) by combining aspects of our equilibrium characterization with state aggregation methods. Second, we believe relevant extensions of our framework include endogenizing capacity choice, allowing for dynamic versioning (across cabins or ticket qualities), and examining network effects across routes. Finally, we consider short-lived buyers. A fruitful area for future research is to consider dynamic pricing games where buyers may strategically delay their purchasing decisions. In such a model, overselling and cancellations, can play an important role.

References

Aryal, Gaurab, Charles Murry, and Jonathan W Williams, "Price discrimination in international airline markets," *arXiv preprint arXiv:2102.05751*, 2022.

Asker, John, Chaim Fershtman, and Ariel Pakes, "Artificial Intelligence and Pricing: The Impact of Algorithm Design," Technical Report, National Bureau of Economic Research 2021.

Bergemann, Dirk and Juuso Välimäki, "Dynamic price competition," *Journal of Economic Theory*, 2006, *127* (1), 232–263.

Berry, Steven and Panle Jia, "Tracing the woes: An empirical analysis of the airline industry," *American Economic Journal: Microeconomics*, 2010, *2* (3), 1–43.

Berry, Steven T, "Estimation of a Model of Entry in the Airline Industry," *Econometrica*, 1992, pp. 889–917.

Board, Simon and Andrzej Skrzypacz, "Revenue management with forward-looking buyers," *Journal of Political Economy*, 2016, *124* (4), 1046–1087.

Borenstein, Severin and Nancy L Rose, "Competition and price dispersion in the US airline industry," *Journal of Political Economy*, 1994, *102* (4), 653–683.

Brown, Zach Y and Alexander MacKay, "Competition in pricing algorithms," Technical Report, National Bureau of Economic Research 2021.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolo, and Sergio Pastorello, "Artificial intelligence, algorithmic pricing, and collusion," *American Economic Review*, 2020, *110* (10), 3267–97.

Caplin, Andrew and Barry Nalebuff, "Aggregation and imperfect competition: On the existence of equilibrium," *Econometrica: Journal of the Econometric Society*, 1991, pp. 25–59.

Castillo, Juan Camilo, "Who Benefits from Surge Pricing?," *Available at SSRN 3245533*, 2022.

Chen, Nan and Przemyslaw Jeziorski, "Consequences of dynamic pricing in competitive airline markets," 2023.

Cho, Sungjin, Gong Lee, John Rust, and Mengkai Yu, "Optimal dynamic hotel pricing," in "2018 Meeting Papers," Vol. 179 2018.

Dana, James D, "Equilibrium price dispersion under demand uncertainty: the roles of costly capacity and market structure," *The RAND Journal of Economics*, 1999, pp. 632–660.

_ , "Using yield management to shift demand when the peak time is unknown," *The Rand Journal of Economics*, 1999, pp. 456–474.

_ **and Kevin R Williams**, "Intertemporal price discrimination in sequential quantity-price games," *Marketing Science*, 2022.

de Albéniz, Victor Martínez and Kalyan Talluri, "Dynamic price competition with fixed capacities," *Management Science*, 2011, *57* (6), 1078–1093.

D'Haultfœuille, Xavier, Ao Wang, Philippe Février, and Lionel Wilner, "Estimating the Gains (and Losses) of Revenue Management," *arXiv preprint arXiv:2206.04424*, 2022.

Dilme, Francesc and Fei Li, "Revenue management without commitment: Dynamic pricing and periodic flash sales," *The Review of Economic Studies*, 2019, *86* (5), 1999–2034.

Dontchev, Asen L and R Tyrrell Rockafellar, *Implicit functions and solution mappings*, Vol. 543, Springer, 2009.

Dudey, Marc, "Dynamic edgeworth-bertrand competition," *The Quarterly Journal of Economics*, 1992, *107* (4), 1461–1477.

Fleming, Wendell H and Halil Mete Soner, *Controlled Markov processes and viscosity solutions*, Vol. 25, Springer Science & Business Media, 2006.

Gallego, Guillermo and Garrett Van Ryzin, "Optimal dynamic pricing of inventories with stochastic demand over finite horizons," *Management science*, 1994, *40* (8), 999–1020.

Gerardi, Kristopher S and Adam Hale Shapiro, "Does competition reduce price dispersion? New evidence from the airline industry," *Journal of Political Economy*, 2009, *117* (1), 1–37.

Gershkov, Alex, Benny Moldovanu, and Philipp Strack, "Revenue-maximizing mechanisms with strategic customers and unknown, markovian demand," *Management Science*, 2018, *64* (5), 2031–2046.

Hansen, Karsten T, Kanishka Misra, and Mallesh M Pai, "Frontiers: Algorithmic collusion: Supra-competitive prices via independent algorithms," *Marketing Science*, 2021, *40* (1), 1–12.

Hanson, Ward and Kipp Martin, "Optimizing multinomial logit profit functions," *Management Science*, 1996, *42* (7), 992–1003.

Hendel, Igal and Aviv Nevo, "Intertemporal price discrimination in storable goods markets," *American Economic Review*, 2013, *103* (7), 2722–51.

Horn, Roger A and Charles R Johnson, *Matrix analysis*, Cambridge university press, 2012.

Hortaçsu, Ali, Olivia R Natan, Hayden Parsley, Timothy Schwieg, and Kevin R Williams, "Incorporating Search and Sales Information in Demand Estimation," Technical Report, National Bureau of Economic Research 2021.

_ , _ , _ , _ , and _ , "Organizational structure and pricing: Evidence from a large us airline," Technical Report, National Bureau of Economic Research 2021.

Kehoe, Patrick J, Bradley J Larsen, and Elena Pastorino, "Dynamic competition in the era of big data," in "Technical report, working paper," Stanford University Stanford, CA., 2018.

Kellogg, RB, "Uniqueness in the Schauder fixed point theorem," *Proceedings of the American Mathematical Society*, 1976, *60* (1), 207–210.

Konovalov, Alexander and Zsolt Sándor, "On price equilibrium with multi-product firms," *Economic Theory*, 2010, *44* (2), 271–292.

Lazarev, John, "The welfare effects of intertemporal price discrimination: an empirical analysis of airline pricing in US monopoly markets," 2013.

Maskin, Eric and Jean Tirole, "A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles," *Econometrica: Journal of the Econometric Society*, 1988, pp. 571–599.

McAfee, R Preston and Vera Te Velde, "Dynamic pricing in the airline industry," *Handbook on economics and information systems*, 2006, *1*, 527–67.

Milgrom, Paul and John Roberts, "Rationalizability, learning, and equilibrium in games with strategic complementarities," *Econometrica: Journal of the Econometric Society*, 1990, pp. 1255–1277.

Nocke, Volker and Nicolas Schutz, "Multiproduct-firm oligopoly: An aggregative games approach," *Econometrica*, 2018, *86* (2), 523–557.

Petrin, Amil and Kenneth Train, "A control function approach to endogeneity in consumer choice models," *Journal of marketing research*, 2010, *47* (1), 3–13.

Puller, Steven L, Anirban Sengupta, and Steven N Wiggins, "Does scarcity drive intra-route price dispersion in airlines?," 2012.

Sengupta, Anirban and Steven N Wiggins, "Airline pricing, price dispersion, and ticket characteristics on and off the internet," *American Economic Journal: Economic Policy*, 2014, *6* (1), 272–307.

Siegert, Caspar and Robert Ulbricht, "Dynamic oligopoly pricing: Evidence from the airline industry," *International Journal of Industrial Organization*, 2020, *71*, 102639.

Stavins, Joanna, "Price discrimination in the airline market: The effect of market concentration," *Review of Economics and Statistics*, 2001, *83* (1), 200–202.

Sweeting, Andrew, James W Roberts, and Chris Gedge, "A model of dynamic limit pricing with an application to the airline industry," *Journal of Political Economy*, 2020, *128* (3), 1148–1193.

Talluri, Kalyan and Garrett Van Ryzin, "Revenue management under a general discrete choice model of consumer behavior," *Management Science*, 2004, *50* (1), 15–33.

Vives, Xavier, "Strategic complementarities in oligopoly," in "Handbook of Game Theory and Industrial Organization, Volume I," Edward Elgar Publishing, 2018.

Vulcano, Gustavo, Garrett Van Ryzin, and Richard Ratliff, "Estimating primary demand for substitutable products from sales transaction data," *Operations Research*, 2012, *60* (2), 313–334.

Williams, Kevin R., "The Welfare Effects of Dynamic Pricing: Evidence From Airline Markets," *Econometrica*, 2022, *90* (2), 831–858.

Zhao, Wen and Yu-Sheng Zheng, "Optimal dynamic pricing for perishable assets with nonhomogeneous demand," *Management science*, 2000, *46* (3), 375–388.

A Proofs

A.1 Technical results

A.1.1 Continuous time limit

We use the following result for the proofs of Lemma 1 and Theorem 1.

Lemma 2. *Consider a continuous price function* $(\Omega, \boldsymbol{\theta}) \mapsto \mathbf{p}^*(\Omega, \boldsymbol{\theta}) = (p_j^*(\Omega, \boldsymbol{\theta}))_j$ *on a compact set* $\mathcal{O}$*, and a bounded and continuous function* $\mathbf{A} : \mathbb{R}^{\mathcal{J}} \times \mathbb{R}^{\mathcal{F} \times \mathcal{J}} \times \mathcal{T} \to \mathbb{R}^{\mathcal{F}}$*. Let* $\Pi_{f,t}(\mathbf{K}; \Delta)$*,* $f \in \mathcal{F}$*, be a solution to the difference equations*

$$\left(\frac{\Pi_{f,t+\Delta}(\mathbf{K};\Delta) - \Pi_{f,t}(\mathbf{K};\Delta)}{\Delta} \right)_f = -\lambda_t \mathbf{A}\Big(\mathbf{p}^*\big(\Omega(\mathbf{K};\Delta)\big), \boldsymbol{\theta}_t\big), \Omega(\mathbf{K};\Delta), \boldsymbol{\theta}_t\Big)$$

where $\Omega(\mathbf{K};\Delta) = (\omega^f_{j,t}(\mathbf{K};\Delta))_{f,j}$*,* $\omega^f_{j,t}(\mathbf{K};\Delta) := \Pi_{f,t+\Delta}(\mathbf{K};\Delta) - \Pi_{f,t+\Delta}(\mathbf{K} - \mathbf{e}_j;\Delta)$*, with boundary conditions (i)* $\Pi_{f,T}(\mathbf{K};\Delta) = 0$*, (ii)* $\Pi_{f,t}(\mathbf{K};\Delta) = -\infty$ *if* $K_j < 0$ *for a* $j \in \mathcal{J}_f$*, and (iii)* $\Pi_{f,t}(\mathbf{K} - \mathbf{e}_j;\Delta) = \Pi_{f,t}(\mathbf{K};\Delta)$ *if* $K_j = 0$ *for a* $j \notin \mathcal{J}_f$*,* $K_{j'} \geq 0$ *for all* $j' \in \mathcal{J}_f$*. Then,* $(\Pi_{f,t}(\mathbf{K};\Delta))_f$ *converges and any limit* $(\Pi_{f,t}(\mathbf{K}))_f$ *satisfies*

$$\big(\dot{\Pi}_{f,t}(\mathbf{K})\big)_f = -\lambda_t \mathbf{A}\Big(\mathbf{p}^*\big(\Omega(\mathbf{K}), \boldsymbol{\theta}_t\big), \Omega(\mathbf{K}), \boldsymbol{\theta}_t\Big),$$

where $\Omega(\mathbf{K}) = (\omega^f_{j,t}(\mathbf{K}))_{f,j}$*,* $\omega^f_{j,t}(\mathbf{K}) := \Pi_{f,t}(\mathbf{K}) - \Pi_{f,t+\Delta}(\mathbf{K} - \mathbf{e}_j)$*, with boundary conditions (i)* $\Pi_{f,T}(\mathbf{K}) = 0$*, (ii)* $\Pi_{f,t}(\mathbf{K};) = -\infty$ *if* $K_j < 0$ *for a* $j \in \mathcal{J}_f$*, and (iii)* $\Pi_{f,t}(\mathbf{K} - \mathbf{e}_{j'}) = \Pi_{f,t}(\mathbf{K})$ *if* $K_{j'} = 0$ *for a* $j' \notin \mathcal{J}_f$*,* $K_j \geq 0$ *for all* $j \in \mathcal{J}_f$*.*

Proof. Since $\mathbf{A}$ is bounded, the difference equations show that $(\Pi_f(\mathbf{K};\Delta))_{f \in \mathcal{F}, \mathbf{K} \leq \mathbf{K}_0}$ is equicontinuous and equibounded in t as $\Delta \to 0$. Hence, by the Arzela-Ascoli Theorem, there exist limit points $(\Pi_f(\mathbf{K}))_{f \in \mathcal{F}, \mathbf{K} \leq \mathbf{K}_0}$. We claim that

$$\big(\Pi_{f,t}(\mathbf{K})\big)_f = \int_t^T \lambda_u \mathbf{A}\Big(\mathbf{p}^*\big(\Omega_u(\mathbf{K}), \boldsymbol{\theta}_u\big), \Omega_u(\mathbf{K}), \boldsymbol{\theta}_u\Big) du. \tag{5}$$

To this end, we note that if we let $\lceil u \rceil_\Delta$ to be the smallest number that is divisible by Δ and larger or equal than u

$$\left(\Pi_{f,t}(\mathbf{K};\Delta)\right)_f = \int_t^T \lambda_{\lceil u \rceil_\Delta} \mathbf{A}\Big(\mathbf{p}^*\big(\Omega_{\lceil u \rceil_\Delta}(\mathbf{K};\Delta), \boldsymbol{\theta}_{\lceil u \rceil_\Delta}\big), \Omega_{\lceil u \rceil_\Delta}(\mathbf{K};\Delta), \boldsymbol{\theta}_{\lceil u \rceil_\Delta}\Big) du. \tag{6}$$

We take the limit $\Delta \to 0$ on both sides. The left-hand side of (6) converges to the left-hand side of (5). On the right-hand side, $\Omega_{\lceil u \rceil_\Delta}(\mathbf{K};\Delta)$ converges to $\Omega_u(\mathbf{K})$. Hence, by continuity of $\mathbf{p}^*$ and $\mathbf{A}$ the integrand in (6) converges to the integrand in (5). By the dominated convergence theorem the right-hand side of (6) converges to the right-hand side of (5). Thus, any limiting value function exists and must satisfy (5). ■

A.1.2 Continuity of stage game prices

Lemma 3. *Let $\mathscr{P} \subset \mathbb{R}^{\mathscr{J}}$ be compact and $\mathscr{O}$ a compact set of $(\Omega, \boldsymbol{\theta})$. Further, let $g : \mathscr{P} \times \mathscr{O} \to \mathscr{P}, (\mathbf{q};\Omega,\boldsymbol{\theta}) \mapsto \mathbf{p}$ be (i) continuous in $\mathbf{q}$, (ii) continuous in Ω and $\boldsymbol{\theta}$, (iii) such that it implicitly defines a unique $\mathbf{p}^*(\Omega,\boldsymbol{\theta})$ satisfying $g(\mathbf{p}^*(\Omega,\boldsymbol{\theta});\Omega,\boldsymbol{\theta}) = \mathbf{p}^*(\Omega,\boldsymbol{\theta})$ for all $(\Omega,\boldsymbol{\theta}) \in \mathscr{O}$. Then, $\mathbf{p}^*(\Omega,\boldsymbol{\theta})$ depends continuously on Ω and $\boldsymbol{\theta}$.*

Proof. Consider the graph of $p^*(\Omega,\boldsymbol{\theta})$: $G = \{(\mathbf{p},\Omega,\boldsymbol{\theta}) : g(\mathbf{p};\Omega,\boldsymbol{\theta}) = p\}$. By the continuity of g, G is closed in $\mathscr{P} \times \mathscr{O}$. Since $p^*(\Omega,\boldsymbol{\theta})$ stays in the compact set $\mathscr{P}$ and is single-valued, G is upper hemicontinuous as a correspondence. Hence, p^* is continuous as a function.[27] ■

A.2 Proofs of Single Firm Model

A.2.1 Proof of Lemma 1

In the following we omit the conditioning argumet $\mathscr{A}$.

Step 1: All profit-maximizing prices $\mathbf{p}^M$ are interior. First, we show that given $\boldsymbol{\omega}$ and $\boldsymbol{\theta}$,

$$\mathbf{p}^M \in \arg\max_{\mathbf{q}} \sum_{j \in \mathscr{J}} s_j(\mathbf{q};\boldsymbol{\theta})(q_j - \omega_j)$$

[27] We thank Satoru Takahashi for helping us to simplify this proof.

is bounded from below by a vector $\underline{\mathbf{p}} = (\underline{p} + \omega_1, \ldots, \underline{p} + \omega_J)$, $\underline{p} \in \mathbb{R}$. We proceed with a proof by contradiction. Suppose such a $\underline{\mathbf{p}}$ did not exist. Then, for any $\underline{p} \in \mathbb{R}$ there exists an optimal price vector $\mathbf{p}^M$ and a j such that $p_j^M - \omega_j = \min_{j'}(p_{j'}^M - \omega_{j'}) < \underline{p}$. At this optimal price $\mathbf{p}^M$ (which could include (minus) infinite prices), the derivative of the stage game profit with respect to any price dimension has to be smaller than or equal to zero by optimality. The derivative with respect to p_j at $\mathbf{p}^M$ (or as we converge to $\mathbf{p}^M$ if it includes (minus) infinite prices) is

$$
\begin{aligned}
&\lim_{\mathbf{p}\to\mathbf{p}^M} \sum_{k\neq j} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_k - \omega_k) + s_j(\mathbf{p};\boldsymbol{\theta}) + \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_j - \omega_j) && \geq \\
&\lim_{\mathbf{p}\to\mathbf{p}^M} -(p_j - \omega_j)\underbrace{\left(\left|\frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\right| - \sum_{k\neq j}\frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\right)}_{=\frac{\partial s_0}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})>0 \text{ by Assumption 1-iii)}} + s_j(\mathbf{p};\boldsymbol{\theta}) && \geq \\
&\lim_{\mathbf{p}\to\mathbf{p}^M} -\underline{p}\frac{\partial s_0}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) + s_j(\mathbf{p};\boldsymbol{\theta}) \xrightarrow[\underline{p}\to-\infty]{} \infty \qquad \text{by Assumption 1-iii).}
\end{aligned}
$$

Thus, for sufficiently small $\underline{p}$, this yields a contradiction, i.e. any optimal price vector $\mathbf{p}^M$ is bounded by a vector $\underline{\mathbf{p}}$ from below.

Next, we show that given $\boldsymbol{\omega}$ and $\boldsymbol{\theta}$, any profit maximizing price vector $\mathbf{p}^M$ is bounded by a vector $\bar{\mathbf{p}} = (\bar{p} + \omega_1, \ldots, \bar{p} + \omega_J)$, $\bar{p} \in \mathbb{R}$. We again proceed with a proof by contradiction. Suppose such a $\bar{\mathbf{p}}$ did not exist. Then, for any $\bar{p} \in \mathbb{R}$, there exists an optimal price vector $\mathbf{p}^M$ and a j such that $p_j^M - \omega_j = \max_{j'}\left(p_{j'}^M - \omega_{j'}\right) > \bar{p}$. At the optimal price $\mathbf{p}^M$ (which could include (minus) infinite prices), the derivative of the stage game profit with respect to any price dimension has to be greater than or equal to zero by optimality. There exists a constant $C > 0$ satisfying Assumption 1-iii) as we have established a lower bound $\underline{p}$ for $\mathbf{p}^M$. The derivative with respect to p_j at $\mathbf{p}^M$ (or as we converge to $\mathbf{p}^M$ if it includes (minus) infinite prices) is

$$\begin{aligned}
&\lim_{\mathbf{p}\to\mathbf{p}^M} \underbrace{\sum_{k\neq j} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_k-\omega_k)}_{\geq 0} + s_j(\mathbf{p};\boldsymbol{\theta}) + \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_j-\omega_j) && \leq \\
&\lim_{\mathbf{p}\to\mathbf{p}^M} \sum_{k\neq j} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_j-\omega_j) + C^{-1}\frac{\partial s_0}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) + \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_j-\omega_j) && = \\
&\lim_{\mathbf{p}\to\mathbf{p}^M} \underbrace{\frac{\partial s_0}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})}_{>0}(C^{-1}-(p_j-\omega_j)) \leq \lim_{\mathbf{p}\to\mathbf{p}^M} \frac{\partial s_0}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(C^{-1}-\overline{p}) \xrightarrow[\overline{p}\to\infty]{} -\infty && .
\end{aligned}$$

by Assumption 1-iii). Thus, for sufficiently large $\bar{p}$, this yields a contradiction. Hence, any optimal price vector $\mathbf{p}^M$ is bounded by a vector $\bar{\mathbf{p}}$ from above.

Step 2: Uniqueness of profit-maximizing price $\mathbf{p}^M$. It follows from Step 1 that any profit-maximizing price $\mathbf{p}^M$ of the stage game must satisfy the FOCs of the firm. Assumption 1-iii) implies that the Jacobian matrix of demand $D_{\mathbf{p}^{\mathscr{A}}}\mathbf{s}(\mathbf{p}^{\mathscr{A}};\boldsymbol{\theta},\mathscr{A})$ is diagonally dominant since $\frac{\partial s_0}{\partial p_j}(\mathbf{p}^{\mathscr{A}};\boldsymbol{\theta},\mathscr{A}) = \left|\frac{\partial s_j}{\partial p_j}(\mathbf{p}^{\mathscr{A}};\boldsymbol{\theta},\mathscr{A})\right| - \sum_{j'\in\mathscr{A}\setminus\{j\}} \frac{\partial s_{j'}}{\partial p_j}(\mathbf{p}^{\mathscr{A}};\boldsymbol{\theta},\mathscr{A}) > 0$.[28] Then, $D_{\mathbf{p}}\mathbf{s}(\mathbf{p}^{\mathscr{A}};\boldsymbol{\theta},\mathscr{A})$ is non-singular by the Levy-Desplanques Theorem (see, e.g., Theorem 6.1.10. in Horn and Johnson (2012)). Hence, the FOCs can be written as Equation 2. Because of Assumption 2 there is a unique solution to this system of equations by Lemma 2 (Kellogg (1976)) in Konovalov and Sándor (2010).

Step 3: Convergence. We can apply the Implicit Function Theorem to Equation 2 by Assumption 2 and it follows that the unique optimal price $\mathbf{p}^M(\Omega,\boldsymbol{\theta})$ is continuous in Ω and $\boldsymbol{\theta}$. Convergence to Equation 1 follows by Lemma 2.

A.2.2 Proof of Proposition 1

Proof. i) To see that $\Pi_{M,t}(\mathbf{K})$ is decreasing in t, note that in Equation 1, setting $p_j > (\Pi_{M,t}(\mathbf{K}-\Pi_{M,t}(\mathbf{K}-\mathbf{e}_j)))$ results in a positive stage-game payoff, so $\dot{\Pi}_{M,t}(\mathbf{K}) < 0$.

[28]Consistent with the common convention, the Jacobi matrix of a vector-valued function $f(\mathbf{x}) \in \mathbb{R}^n$, $\mathbf{x} \in \mathbb{R}^n$ is $D_{\mathbf{x}}f(\mathbf{x}) := \left(\frac{\partial f_i}{\partial x_j}\right)_{i,j}$, i denoting rows and j columns, and bold vectors $\mathbf{x}$ are column vectors.

Next, we show that $\Pi_{M,t}(\mathbf{K}) > \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j)$ for all j by induction in $\sum_j K_j$.

Induction start: It is immediate that $\Pi_{M,t}(\mathbf{e}_j) \geq \Pi_{M,t}(\mathbf{0}) = 0$ for all j and $t \leq T$.

Induction hypothesis: Assume that $\Pi_{M,t}(\mathbf{K}) > \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j)$ for all $\mathbf{K}$ with $\sum_j K_j = \bar{K}$ and $j \in \mathscr{J}$.

Induction step: Now, consider a capacity vector $\mathbf{K}$ with $\sum_j K_j = \bar{K}+1$. The solution of the differential equation for the profits is

$$\Pi_{M,t}(\mathbf{K}) = \int_t^T \lambda_z \sum_j s_j(\mathbf{p}_z^M(\mathbf{K}))\,(p_{j,z}^M(\mathbf{K}) + \Pi_{M,z}(\mathbf{K}-\mathbf{e}_j)) \cdot e^{-\int_t^z \lambda_u \sum_{j'} s_{j'}(\mathbf{p}_u^M(\mathbf{K}))du} dz.$$

By sub-optimality of the prices $\mathbf{p}_t^M(\mathbf{K}-\mathbf{e}_k)$ given capacity vector $\mathbf{K}$, we have for all k

$$\begin{aligned}\Pi_{M,t}\big(\mathbf{K}\big) \geq \\ \int_t^T \lambda_z \sum_j s_j\big(\mathbf{p}_z^M(\mathbf{K}-\mathbf{e}_k)\big)\,\big(p_{j,z}^M(\mathbf{K}-\mathbf{e}_k) + \underbrace{\Pi_{M,z}(\mathbf{K}-\mathbf{e}_j)}_{\substack{> \Pi_{M,z}\left(\mathbf{K}-\mathbf{e}_k-\mathbf{e}_j\right) \\ \text{by induction hypothesis}}}\big) \cdot e^{-\int_t^z \lambda_u \sum_{j'} s_{j'}(\mathbf{p}_u^M(\mathbf{K}-\mathbf{e}_k))du} dz \\ > \Pi_{M,t}\big(\mathbf{K}-\mathbf{e}_k\big).\end{aligned}$$

ii) Next, we show that $\Pi_{M,t}(\mathbf{K}) - \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j) \leq \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j) - \Pi_{M,t}(\mathbf{K}-2\mathbf{e}_j)$ for all j. To this end, let

$$H(\mathbf{x};\boldsymbol{\theta}) = -\max_{\mathbf{p}} \sum_j s_j(\mathbf{p};\boldsymbol{\theta})(p_j - x_j).$$

Note that H is concave as a minimum of affine functions, strictly increasing in $\mathbf{x}$. Since H is concave and continuous, by the Fenchel-Moreau Theorem, it admits the representation

$$H(\mathbf{x};\boldsymbol{\theta}) = \inf_{\mathbf{s}} (\mathbf{s}\cdot\mathbf{x} - H^*(\mathbf{s};\boldsymbol{\theta}))$$

where $H^*(\mathbf{s};\boldsymbol{\theta}) = \inf_{\mathbf{x}} (\mathbf{x}\cdot\mathbf{s} - H(\mathbf{x};\boldsymbol{\theta}))$ is the concave conjugate of H. Moreover,

$$\dot{\Pi}_{M,t}(\mathbf{K}) = \lambda_t H(\nabla\Pi_t(\mathbf{K});\boldsymbol{\theta}_t)$$

where $\nabla\Pi_{M,t}(\mathbf{K}) = \big(\Pi_{M,t}(\mathbf{K}) - \Pi_{M,t}(\mathbf{K}-\mathbf{e}_j)\big)_j$. Thus, $\Pi_{M,t}(\mathbf{K})$ is the value function for the optimal control problem

$$\Pi_{M,t}(\mathbf{K}) = \sup_{\mathbf{s}\in\mathscr{A}} \mathbb{E}\left[\int_t^T \lambda_u H^*(\mathbf{s}_u;\boldsymbol{\theta}_u)du \,\middle|\, \mathbf{X}_t^{\mathbf{s}} = \mathbf{K}\right] =: \sup_{\mathbf{s}} J_t(\mathbf{K},\mathbf{s})$$

where $\mathbf{X}_t^{\mathbf{a}}$ is the process which jumps by $-\mathbf{e}_j$ at rate $\lambda_t s_{j,t}$ and $\mathbf{s}\in\mathscr{A}$ are processes adapted with respect to the filtration on the probability space supporting $\mathbf{X}^{\mathbf{s}}$, with the property $s_{j,t}=0$ if $X_{j,t}^{\mathbf{s}}=0$ (Theorem 8.1 in Fleming and Soner (2006)). Let $\mathbf{s}_{\mathbf{K}}^*$ be the optimal control in the previous equation and $\mathbf{s}_{K-2}^*$ be the optimal control when $\mathbf{K}$ is replaced by $\mathbf{K}-2\mathbf{e}_j$. Then, note that since $\mathbf{s}_{\mathbf{K}}^*, \mathbf{s}_{\mathbf{K}-2\mathbf{e}_j}^* \in \mathscr{A}$, $\frac{\mathbf{s}_{\mathbf{K}}^*+\mathbf{s}_{\mathbf{K}-2\mathbf{e}_j}^*}{2}\in\mathscr{A}$ because the process $\left(\mathbf{X}_s^{\frac{\mathbf{s}_{\mathbf{K}}^*+\mathbf{s}_{\mathbf{K}-2\mathbf{e}_f}^*}{2}}\right)_s$ can be chosen as $\left(\frac{\mathbf{X}_s^{\mathbf{s}_{\mathbf{K}}^*}+\mathbf{X}_s^{\mathbf{s}_{\mathbf{K}-2\mathbf{e}_f}^*}}{2}\right)_s$ ("coupling argument"). Hence,

$$\begin{aligned}
&\Pi_{M,t}(\mathbf{K})+\Pi_{M,t}(\mathbf{K}-2\mathbf{e}_j)-2\Pi_{M,t}\big(\mathbf{K}-\mathbf{e}_j\big) && \leq \\
&J_t(\mathbf{K},s_{\mathbf{K}}^*)+J_t(\mathbf{K}-2\mathbf{e}_j,s_{\mathbf{K}-2\mathbf{e}_j}^*)-2\,J_t\left(\mathbf{K}-\mathbf{e}_j,\frac{\mathbf{s}_{\mathbf{K}}^*+\mathbf{s}_{\mathbf{K}-2\mathbf{e}_j}^*}{2}\right) && \leq \\
&\mathbb{E}\left[\int_t^T \lambda_u\left(H^*(\mathbf{s}_{\mathbf{K},u}^*)+H^*(\mathbf{s}_{\mathbf{K}-2\mathbf{e}_j,u}^*)-2H^*\left(\frac{\mathbf{s}_{\mathbf{K},u}^*+\mathbf{s}_{\mathbf{K}-2\mathbf{e}_f,u}^*}{2}\right)\right)du \,\middle|\, \mathbf{X}_t^{\mathbf{s}_{\mathbf{K}}^*}=\mathbf{K},\,\mathbf{X}_t^{\mathbf{s}_{\mathbf{K}-2\mathbf{e}_j}^*}=\mathbf{K}-2\mathbf{e}_j,\right] && \leq 0.
\end{aligned}$$

iii) To show that $\omega_{j,t\wedge\tau}^M(\mathbf{K}_t)$ is a submartingale, we show that for any capacity vector $\bar{\mathbf{K}}$ with $\bar{K}_j\geq 2$:

$$\lim_{\Delta\to 0}\frac{\mathbb{E}_0\big[\omega_{j,t+\Delta}^M(\mathbf{K}_{t+\Delta})-\omega_{j,t}^M(\mathbf{K}_t)\big|\mathbf{K}_t=\bar{\mathbf{K}}\big]}{\Delta}\geq 0.$$

To this end, first, note that $\mathbf{K}_t$ is right-continuous in t. Consider $\bar{\mathbf{K}}$ with $\bar{K}_j\geq 2$. Then, we have that

$$\begin{aligned}
&\lim_{\Delta\to 0}\frac{\mathbb{E}_0\big[\omega_{j,t+\Delta}^M(\mathbf{K}_{t+\Delta})-\omega_{j,t}^M(\mathbf{K}_t)|\mathbf{K}_t=\bar{\mathbf{K}}\big]}{\Delta} = \\
&\lim_{\Delta\to 0}\frac{\mathbb{E}_0\big[\omega_{j,t+\Delta}^M(\mathbf{K}_{t+\Delta})-\omega_{j,t}^M(\mathbf{K}_{t+\Delta})|\mathbf{K}_t=\bar{\mathbf{K}}\big]}{\Delta}+\lim_{\Delta\to 0}\frac{\mathbb{E}_0\big[\omega_{j,t}^M(\mathbf{K}_{t+\Delta})-\omega_{j,t}^M(\mathbf{K}_t)|\mathbf{K}_t=\bar{\mathbf{K}}\big]}{\Delta} = \\
&\dot{\omega}_{j,t}^M(\bar{\mathbf{K}})+\lambda_t\sum_{j'}s_{j',t}(\mathbf{p}_t^M(\bar{\mathbf{K}}))\big(\omega_{j,t}^M(\bar{\mathbf{K}}-\mathbf{e}_{j'})-\omega_{j,t}^M(\bar{\mathbf{K}})\big)
\end{aligned}$$

by right-continuity of the process $\mathbf{K}_t$. By (1), we can write

$$\dot{\omega}^M_{j,t}(\bar{\mathbf{K}}) = -\lambda_t\Big[\sum_{j'} s_{j',t}(\mathbf{p}^M_t(\bar{\mathbf{K}}))\big(p^M_{j',t}(\bar{\mathbf{K}})-\omega^M_{j',t}(\bar{\mathbf{K}})\big)-s_{j',t}(p^M_t(\bar{\mathbf{K}}-\mathbf{e}_j))\big(p^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)-\omega^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)\big)\Big].$$

and we know that

$$\begin{aligned}-\omega^M_{j',t}(\bar{\mathbf{K}})+\omega^M_{j,t}(\bar{\mathbf{K}})-\omega^M_{j,t}(\bar{\mathbf{K}}-\mathbf{e}_{j'}) &= \Pi^M(\bar{\mathbf{K}}-\mathbf{e}_{j'})-\Pi^M(\bar{\mathbf{K}}-\mathbf{e}_j)-\Pi^M(\bar{\mathbf{K}}-\mathbf{e}_{j'})+\Pi^M(\bar{\mathbf{K}}-\mathbf{e}_{j'}-\mathbf{e}_j)\\ &= -\omega^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)\end{aligned}$$

Hence, $\lim\limits_{\Delta\to 0}\frac{\mathbb{E}_0\left[\omega^M_{j,t+\Delta}(\mathbf{K}_{t+\Delta})-\omega^M_{j,t}(\mathbf{K}_t)\mid\mathbf{K}_t=\bar{\mathbf{K}}\right]}{\Delta}$ is equal to

$$-\lambda_t\Big[\sum_{j'} s_{j',t}\big(\mathbf{p}^M_t(\bar{\mathbf{K}})\big)\big(p^M_{j',t}(\bar{\mathbf{K}})-\omega^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)\big)-s_{j',t}\big(\mathbf{p}^M_t(\bar{\mathbf{K}}-\mathbf{e}_j)\big)\big(p^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)-\omega^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)\big)\Big]$$

Then, note that by optimality of $\mathbf{p}^M_t(\bar{\mathbf{K}}-\mathbf{e}_j)$,

$$\sum_{j'} s_{j',t}\big(\mathbf{p}^M_t(\bar{\mathbf{K}})\big)\big(p^M_{j',t}(\bar{\mathbf{K}})-\omega^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j)\big)\le\sum_{j'} s_{j',t}\big(\mathbf{p}^M_t(\bar{\mathbf{K}}-\mathbf{e}_j)\big)\big(p^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_j))-\omega^M_{j',t}(\bar{\mathbf{K}}-\mathbf{e}_{j'})\big).$$

Hence, $\lim\limits_{\Delta\to 0}\frac{\mathbb{E}_0\left[\omega^M_{j,t+\Delta}(\mathbf{K}_{t+\Delta})-\omega^M_{j,t}(\mathbf{K}_t)\mid\mathbf{K}_t=\bar{\mathbf{K}}\right]}{\Delta}\ge 0$. ■

A.3 Proofs of Oligopoly Model

A.3.1 Existence, Uniqueness, and Continuity of the Stage Game Equilibrium

In this section we analyze the stage game parameterized by $\boldsymbol{\theta}$ and Ω with several lemmata. Recall we defined in Section 2.3 the function

$$\mathbf{g}_f(\mathbf{p},\boldsymbol{\theta},\Omega):=\underbrace{\Big(\big(D_{\mathbf{p}_f}\mathbf{s}_f(\mathbf{p};\boldsymbol{\theta})\big)^{\intercal}\Big)^{-1}D_{\mathbf{p}_f}\big(\mathbf{s}(\mathbf{p};\boldsymbol{\theta})^{\intercal}\boldsymbol{\omega}^f\big)^{\intercal}}_{\substack{\text{net opportunity costs}\\ \text{of selling}}}-\underbrace{\Big(\big(D_{\mathbf{p}_f}\mathbf{s}_f(\mathbf{p};\boldsymbol{\theta})\big)^{\intercal}\Big)^{-1}\mathbf{s}_f(\mathbf{p};\boldsymbol{\theta})}_{\substack{\text{inverse quasi}\\ \text{own-price elasticities}}}.$$

Then, the FOC for firm f's stage game payoff maximization problem is $\mathbf{g}_f(\mathbf{p},\boldsymbol{\theta},\Omega)=\mathbf{p}_f$. We show next that the following conditions on the stage game guarantee uniqueness and existence of a stage game equilibrium if equilibrium prices are interior.

Assumption 3. *The following two conditions hold,*

i) $\det\left(D_{\mathbf{p}_f}\mathbf{g}_f(\mathbf{p},\boldsymbol{\theta},\Omega)-I_{\mathcal{J}_f}\right)\neq 0$ *for all* $\mathbf{p}$ *and* f*;*

ii) $\det\left(D_{\mathbf{p}}\big(\mathbf{g}(\mathbf{p},\boldsymbol{\theta},\Omega)\big)-I_{\mathcal{J}}\right)\neq 0$ *for all* $\mathbf{p}$*, where* $\mathbf{g}(\mathbf{p},\boldsymbol{\theta},\Omega):=\left(\mathbf{g}_f(\mathbf{p},\boldsymbol{\theta},\Omega):f\in\mathcal{F}\right)\in\mathbb{R}^{\mathcal{J}}$.

Lemma 4. *Given Assumptions 1, 2 and 3, the stage game admits a unique equilibrium. The equilibrium price vector is finite for all available products.*

Proof. **Step 1: All equilibrium prices $\mathbf{p}^*$ are interior.**

First, we show that for fixed Ω and $\boldsymbol{\theta}$, any equilibrium price vector $\mathbf{p}^*$ is bounded from below by a vector $\bar{\mathbf{p}}=((\bar{p}+\omega_j^f)_{j\in\mathcal{J}_f}:f\in\mathcal{F})$, $\bar{p}\in\mathbb{R}$. We proceed with a proof by contradiction. Suppose such a $\underline{\mathbf{p}}$ did not exist. Then, for any $\underline{p}$ there exists an equilibrium price vector $\mathbf{p}^*$ and a j such that $p_j^*-\omega_j^f=\min_{f'}\min_{k\in\mathcal{J}_{f'}}p_k^*-\omega_k^{f'}<\underline{p}$. Additionally, let $k^*=\operatorname{argmax}_{k\notin\mathcal{J}_f}\omega_k^f$. At this equilibrium price vector $\mathbf{p}^*$ (which could include (minus) infinite prices), the derivative of firm f's stage game profit with respect to all firm f's prices has to be smaller or equal to zero by optimality. The derivative with respect to p_j at $\mathbf{p}^*$ (or as we converge to $\mathbf{p}^*$ if it includes (minus) infinite prices) is

$$\lim_{\mathbf{p}\to\mathbf{p}^*}\tfrac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_j-\omega_j^f)+\sum_{k\in\mathcal{J}_f\setminus\{j\}}\tfrac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_k-\omega_k^f)-\sum_{k\notin\mathcal{J}_f}\tfrac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\omega_k^f+s_j(\mathbf{p};\boldsymbol{\theta}) \geq$$

$$\lim_{\mathbf{p}\to\mathbf{p}^*}-\left(\underbrace{\left|\frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\right|-\sum_{k\in\mathcal{J}_f\setminus\{j\}}\frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})}_{\geq 0\text{ by Assumption 1-iii)}}\right)\left(p_j-\omega_j^f+\underbrace{\frac{\sum_{k\notin\mathcal{J}_f}\frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})}{\left|\frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\right|-\sum_{k\in\mathcal{J}_f\setminus\{j\}}\frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})}}_{\in(0,1)\text{ by Assumption 1-iii)}}|\omega_{k^*}^f|\right)+s_j(\mathbf{p};\boldsymbol{\theta}) \geq$$

$$\lim_{\mathbf{p}\to\mathbf{p}^*}-\left(\left|\tfrac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\right|-\sum_{k\in\mathcal{J}_f\setminus\{j\}}\tfrac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})\right)\left(\underline{p}+|\omega_{k^*}^f|\right)+s_j(\mathbf{p};\boldsymbol{\theta})\xrightarrow[\underline{p}\to-\infty]{\infty}$$

by Assumption 1-iii). Thus, for sufficiently small $\underline{p}$, this yields a contradiction, i.e. any equilibrium price vector $\mathbf{p}^*$ is bounded by a vector $\underline{\mathbf{p}}$ from below.

Next, we show that for fixed Ω and $\boldsymbol{\theta}$, any equilibrium price vector $\mathbf{p}^*$ is bounded from above by a vector $\bar{\mathbf{p}} = ((\bar{p} + \omega_j^f)_{j \in \mathscr{J}_f} : f \in \mathscr{F})$, $\bar{p} \in \mathbb{R}$, by contradiction. Suppose such a $\bar{\mathbf{p}}$ did not exist. Then, for any $\bar{p}$, there exists an equilibrium price vector $\mathbf{p}^*$ and a j such that $p_j^* - \omega_j^f = \max_{f'} \max_{k \in \mathscr{J}_{f'}} p_k^* - \omega_k^{f'} > \bar{p}$, $j \in \mathscr{J}_f$. At the equilibrium price $\mathbf{p}^*$ (which could include (minus) infinite prices), the derivative of firm f's stage game profit with respect to all firm f's prices has to be greater or equal to zero by optimality. There exists a constant $C > 0$ satisfying Assumption 1-iii) as we have established a lower bound $\underline{p}$ for $\mathbf{p}^*$. Additionally, let $k^* = \operatorname{argmax}_{k \notin \mathscr{J}_f} |C^{-1} + \omega_k^f|$. The derivative of firm f's payoff with respect to p_j at $\mathbf{p}^*$ (or as we converge to $\mathbf{p}^*$ if it includes (minus) infinite prices) is

$$
\begin{aligned}
&\lim_{\mathbf{p} \to \mathbf{p}^*} \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_j - \omega_j^f) + \sum_{k \in \mathscr{J}_f \setminus \{j\}} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})(p_k - \omega_k^f) - \sum_{k \notin \mathscr{J}_f} \frac{\partial s_k}{\partial p_j}(\mathbf{p})\omega_k^f + s_j(\mathbf{p}) \leq \\
&\lim_{\mathbf{p} \to \mathbf{p}^*} \left(\frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) + \sum_{k \in \mathscr{J}_f \setminus \{j\}} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) \right)(p_j - \omega_j^f) + C^{-1} \left(\left| \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) \right| - \sum_{k \in \mathscr{J}_f \setminus \{j\}} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) \right) \\
&+ \sum_{k \notin \mathscr{J}_f} \frac{\partial s_k}{\partial p_j}(\mathbf{p}) \left| - C^{-1} - \omega_k^f \right| \leq \\
&\lim_{\mathbf{p} \to \mathbf{p}} \left(\left| \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) \right| - \sum_{k \in \mathscr{J}_f \setminus \{j\}} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) \right) \\
&\left(C^{-1} - \bar{p} + \underbrace{\frac{\sum_{k \notin \mathscr{J}_f} \frac{\partial s_k}{\partial p_j}(\mathbf{p})}{\left| \frac{\partial s_j}{\partial p_j}(\mathbf{p};\boldsymbol{\theta}) \right| - \sum_{k \in \mathscr{J}_f \setminus \{j\}} \frac{\partial s_k}{\partial p_j}(\mathbf{p};\boldsymbol{\theta})}}_{\in (0,1)} \left| C^{-1} + \omega_{k^*}^f \right| \right) \xrightarrow[\bar{p} \to \infty]{} -\infty.
\end{aligned}
$$

Thus, for sufficiently large $\bar{p}$, this yields a contradiction. Hence, any equilibrium price vector $\mathbf{p}^*$ is bounded by a vector $\bar{\mathbf{p}} = ((\bar{p} + \omega_j^f)_{j \in \mathscr{J}_f} : f \in \mathscr{F})$ from above.

All in all, it follows that the best response of each firm must be within a box with extreme points $\bar{\mathbf{p}}$ and $\underline{\mathbf{p}}$.

Step 2: Uniqueness of equilibrium price $\mathbf{p}^*$.

It follows from Step 1 that any equilibrium price $\mathbf{p}^*$ of the stage game is a solution to the system of FOCs. Assumption 2 ensures that the Jacobi matrix $D_{\mathbf{p}^f} s(\mathbf{p}^f; \boldsymbol{\theta})$ non-singular by the Levy-Desplanques Theorem (see e.g. Theorem 6.1.10. in Horn and Johnson (2012)). Hence, the FOCs can be written as $\mathbf{g}(\mathbf{p}) = \mathbf{p}$ where $\mathbf{g}$ is as defined in Assumption 3. By Assumption 3-ii), there is a unique solution to this system of equations by Lemma 2 (Kellogg (1976)) in Konovalov and Sándor (2010). Further, by Assumption 3-i) and Kellogg (1976), there is a unique solution of the first order condition of each firm's optimization problem, given by $\mathbf{g_f}(\mathbf{p}) = \mathbf{p}^f$. Thus, for any competitor prices, there exists a unique best response of each firm f, which solves $\mathbf{g}_f(\mathbf{p}) = \mathbf{p}^f$ and the unique solution to $\mathbf{g}(\mathbf{p}) = \mathbf{p}$ must be an equilibrium. ■

Lemma 5. *Let Assumptions 1, 2 and 3 hold for a compact, path-connected set $\mathcal{O}$ of $(\Omega, \boldsymbol{\theta})$. Then the unique equilibrium price vector $\mathbf{p}^*(\Omega, \boldsymbol{\theta})$ is continuous in $(\Omega, \boldsymbol{\theta})$ on $\mathcal{O}$.*

Proof. Let Assumptions 1, 2, and Assumption 3 hold for a compact, path-connected set $\mathcal{O}$ of $(\Omega, \boldsymbol{\theta})$. Then, by Lemma 4, all stage games with parameters $(\Omega, \boldsymbol{\theta}) \in \mathcal{O}$ admit a unique and finite equilibrium that are uniformly bounded on $\mathcal{O}$. Hence, we can apply Lemma 3. ■

A.3.2 Proof of Theorem 1

First, note that by Assumption 3 and smoothness of demand functions, there exists a neighborhood around $(\mathbf{0}, \boldsymbol{\theta}_T)$ so that the stage game has a unique equilibrium. Let us denote that neighborhood $\mathcal{O}$.

The discrete-time game can be written as in Lemma 2 where A is negative the expression in Equation 3 divided by Δ. The corresponding function A is bounded in the neighborhood $\mathcal{O}$ because demand is bounded by 0 and 1. Thus, the discrete time derivatives $\frac{\Pi_{f,t+\Delta}(\mathbf{K};\boldsymbol{\Delta}) - \Pi_{f,t}(\mathbf{K};\Delta)}{\Delta}$ are uniformly bounded and value functions are equicontinuous and equibounded as $\Delta \to 0$. As a result, $\omega^f_{j,t}(\mathbf{K};\Delta)$ is equicontinuous and equibounded. Hence, there is a $T_0(\mathbf{K})$ so that all $(\Omega_t(\mathbf{K};\Delta), \boldsymbol{\theta}_t)$ are in $\mathcal{O}$. If $\mathbf{K}' > \mathbf{K}$, the number of states is simply increasing and hence, $T_0(\mathbf{K}) > T_0(\mathbf{K}')$.

For this fixed $T_0(\mathbf{K})$, the dynamic game for fixed Δ can be analysed recursively. The boundary condition at the deadline T is given by: $\Pi_{f,T}(\mathbf{K};\Delta) = 0$ for all $\mathbf{K}$. The resulting scarcity effects are $\omega^f_{j,T-\Delta}(\mathbf{K};\Delta) = 0$ for all $j \in \mathcal{J} \cap \mathcal{A}(\mathbf{K})$. In period t, given $\Omega_t(\mathbf{K};\Delta)$ for all possible capacity vectors $\mathbf{K}$, we can recursively define equilibrium prices $p_t^*(\mathbf{K};\Delta)$ as equilibrium prices of the stage game where each firm f maximizes

$$\sum_{j\in\mathcal{J}_f\cap\mathcal{A}(\mathbf{K})} s_{j,t}(\mathbf{p})(p_j-\omega_{j,t}^f(\mathbf{K};\Delta))-\sum_{j'\notin\mathcal{J}_f\cap\mathcal{A}(\mathbf{K})} s_{j,t}(\mathbf{p})\omega_{j',t}^f(\mathbf{K};\Delta)).$$

This procedure yields a unique equilibrium as long as all stage games have unique equilibria, which is guaranteed for sufficiently small Δ if $T<T_0(\mathbf{K})$.

By Lemma 4 and Lemma 5, the stage games for $(\Omega,\boldsymbol{\theta})\in\mathcal{O}$ have a unique solution $p^*(\Omega,\boldsymbol{\theta})$ that is also continuous in $(\Omega,\boldsymbol{\theta})$. Then, convergence follows by Lemma 2.

The differential equation for the equilibrium prices follows then immediately from the taking the time derivative of the system of FOCs $\mathbf{g}(\mathbf{p}_t^*(\mathbf{K}),\boldsymbol{\theta}_t,\Omega_t(\mathbf{K}))-\mathbf{p}_t^*(\mathbf{K})\equiv\mathbf{0}$.

A.3.3 Proof of Proposition 2

Let $\lambda_t=\lambda$, $\boldsymbol{\theta}_t=\boldsymbol{\theta}$. So, we will drop the parameter $\boldsymbol{\theta}$ in the notation in this proof. For t close to T, we have established in Theorem 1 that the equilibirum of the stage game is unique and the price vectors $\mathbf{p}_t^*(\mathbf{K})=\mathbf{p}^*(\Omega_t(\mathbf{K}))$ are implicitly defined by a system of equations given by

$$\left(D_{\mathbf{p}_f}\mathbf{s}_f\left(\mathbf{p}_t^*(\mathbf{K})\right)\right)^{\mathsf{T}}\mathbf{p}_t^*(\mathbf{K})-\left(D_{\mathbf{p}_f}\mathbf{s}(\mathbf{p}_t^*(\mathbf{K}))\right)^{\mathsf{T}}\boldsymbol{\omega}_t^f(\mathbf{K})+\mathbf{s}_f\left(\mathbf{p}_t^*(\mathbf{K})\right)=0 \qquad \forall f.$$

We omit $\mathbf{K}$ for readability whenever possible. The only time-dependent variables are then $\Omega_t=(\boldsymbol{\omega}_t^f)_{f\in\mathcal{F}}$. Hence, $\mathbf{p}_t^*$ and Ω_t are continuous in t. Due to the ODE, Ω_t is continuously differentiable, so $\mathbf{p}_t^*$ is continuously differentiable. Inductively it follows that as we take derivatives of the ordinal differential equation, if Ω_t is n times continuously differentiable, then $\mathbf{p}_t^*$ is n times continuously differentiable. The n-th time derivative $(p_t^*)^{(n)}$ depends on the time derivatives $\Omega_t,\ldots,\Omega_t^{(n)}$ and is well defined because the implicit function is smooth in $\mathbf{p}$ and Ω. We are interested in the limit as $t\to T$. We show by induction in n that if $K_j>n$ for all j, then as $t\to T$, $(\omega_{j,t}^f)^{(n)}(\mathbf{K})=0$ for all f,j which implies the claim by Taylor's theorem.

Induction start: First, $\lim_{t \to T} \Omega_t = 0$. Furthermore, we can write for all f and j:

$$\begin{aligned}\dot{\omega}^f_{j,t}(\mathbf{K}) =& \dot{\Pi}_{f,t}(\mathbf{K}) - \dot{\Pi}_{f,t}(\mathbf{K}-\mathbf{e}_j)\\ =& -\lambda\Big[\underbrace{\mathbf{s}_f\left(\mathbf{p}^*(\Omega_t(\mathbf{K}))\right)^{\intercal}\mathbf{p}^*_f(\Omega_t(\mathbf{K})) - \mathbf{s}(\mathbf{p}^*(\Omega_t(\mathbf{K})))^{\intercal}\boldsymbol{\omega}^f_t(\mathbf{K})}_{=:G^1_f(\Omega_t(\mathbf{K}))}\\ &-\underbrace{\left(\mathbf{s}_f\left(\mathbf{p}^*\left(\Omega_t(\mathbf{K}-\mathbf{e}_j)\right)\right)^{\intercal}\mathbf{p}^*_f\left(\Omega_t(\mathbf{K}-\mathbf{e}_j)\right) - \mathbf{s}\left(\mathbf{p}^*\left(\Omega_t(\mathbf{K}-\mathbf{e}_j)\right)\right)^{\intercal}\boldsymbol{\omega}^f_t(\mathbf{K}-\mathbf{e}_j)\right)}_{=:G^1_f(\Omega_t(\mathbf{K}-\mathbf{e}_j))}\Big]\end{aligned}$$

Thus, as $t \to T$, $\dot{\omega}^f_{j,t}(\mathbf{K}) = 0$ if $K_j > 1$. If $j \in \mathscr{J}_f$ and $K_j = 1$, then $\dot{\omega}^f_{j,t}(\mathbf{K}) < 0$. If $j \notin \mathscr{J}^f$ and $K_j = 1$, then by the competition effect $\dot{\omega}^f_{j,t}(\mathbf{K}) > 0$. This implies that $\dot{p}^*_{j,T}(\mathbf{K}) < 0$ if $K_j = 1$ and $\dot{p}^*_{j,T}(\mathbf{K}) = 0$ otherwise.

Induction assumption: Letting for $\Omega^{(m)}_t(\mathbf{K})$ be that matrix of m-th derivatives of $\omega^f_j(\mathbf{K})$, we can write for all f and j

$$(\omega^f_{j,t})^{(n-1)}(\mathbf{K}) = -\lambda\left[G^{n-1}_f\left(\left(\Omega^{(m)}_t(\mathbf{K})\right)^{n-2}_{m=0}\right) - G^{n-1}_f\left(\left(\Omega^{(m)}_t(\mathbf{K}-\mathbf{e}_j)\right)^{n-2}_{m=0}\right)\right]$$

where $G^{n-1}_f((\Omega^{(m)}_t(\mathbf{K}-\mathbf{e}_j))^{n-2}_{m=0}) = \frac{\partial^{n-2}}{(\partial t)^{n-2}} G^1_f(\Omega_t(\mathbf{K}))$. If $K_j > n-1$ for all j, then as $t \to T$, $(\omega^f_{j,t})^{(n-1)}(\mathbf{K}) = 0$ for all f, j.

Induction step: Given the induction assumption, we can also calculate the next order derivative recursively

$$(\omega^f_{j,t})^{(n)}(\mathbf{K}) = -\lambda\left[G^n((\Omega^{(m)}_t(\mathbf{K}))^{n-1}_{m=0}) - G^n(\Omega^{(m)}_t(\mathbf{K}-\mathbf{e}_j))^{n-1}_{m=0}))\right].$$

Then, note if $\min_i K_i > n$, then $(\omega^f_{j,t})^{(n)}(\mathbf{K}) = 0$ by the Induction Assumption. If $\min_i K_i = n$,

$$(\omega^f_{j,t})^{(n)}(\mathbf{K}) = -\lambda\left[-G^n(\Omega^{(m)}_t(\mathbf{K}-\mathbf{e}_j))^{n-1}_{m=0}))\right] = -\lambda\frac{\partial^{n-1}}{(\partial t)^{n-1}} G^1_f((\mathbf{K}-\mathbf{e}_j)).$$

A.3.4 Welfare Dynamics

Consider a discrete choice model as specified in Section 2.6. Let us denote the per-period welfare given prices $\mathbf{p}$ by $w_t(\mathbf{p})$. Further, let us assume that there is a unique solution to the following maximization problem for all $\boldsymbol{\omega} = (\omega_j)_{j\in\mathcal{J}} \in \mathcal{R}^{\mathcal{J}}$:

$$\arg\max_{\mathbf{p}} w_t(\mathbf{p}) - \sum_{j\in\mathcal{J}} s_j(\mathbf{p};\boldsymbol{\theta})\omega_j \tag{7}$$

and let us denote the solution for parameters $\boldsymbol{\theta}$ and $\boldsymbol{\omega}$ by $\mathbf{p}^w(\boldsymbol{\omega},\boldsymbol{\theta}) := (p_j^M(\boldsymbol{\omega},\boldsymbol{\theta}))_{j\in\mathcal{J}}$. Then, there exists a unique welfare-maximizing price path and corresponding continuation welfare $W_t(\mathbf{K};\Delta)$ and an analogous result to Lemma 1 holds.

Lemma 6. *Let us assume that there is a unique welfare-maximizing price path. Then,* $W_t(\mathbf{K}) := \lim_{\Delta\to 0} W_t(\mathbf{K};\Delta)$ *solves the ordinary differential equation*

$$\dot{W}_t(\mathbf{K}) = -\lambda_t\Big(w_t\left(\mathbf{p}^w(\boldsymbol{\omega}_t^w(\mathbf{K}),\boldsymbol{\theta})\right) - \sum_{j\in\mathcal{J}} s_j(\mathbf{p}^w(\boldsymbol{\omega}_t^w(\mathbf{K}),\boldsymbol{\theta});\mathcal{A}(\mathbf{K}))\omega_{j,t}^w(\mathbf{K})\Big),$$

with boundary conditions (i) $W_T(\mathbf{K}) = 0\ \forall\mathbf{K}$ *and (ii)* $W_t(\mathbf{K}) = -\infty$ *if* $K_j < 0$ *for a* $j\in\mathcal{J}$*, where* $\boldsymbol{\omega}_t^w(\mathbf{K}) = (\omega_{j,t}^w(\mathbf{K}))_{j\in\mathcal{J}}$ *with* $\omega_{j,t}^w(\mathbf{K}) := W_t(\mathbf{K}) - W_t(\mathbf{K}-\mathbf{e}_j)$.

The proof follows annalogously to the proof of Lemma 1 from Lemma 2. However, in order to be able to apply this result to our examples and empirical application, we need to show that the welfare-maximizing price path is unique. To this end, we show that Equation 7 holds for the specific form of nested logit demand functions that we use in our empricial application in Section C.

B A Mark-up Formula for IIA Demand

Given the commonly made assumption of "independence of irrelevant alternatives (IIA)," we can derive a clean explicit mark-up formula for dynamic price competition. The IIA assumption in our setting can be stated as follows:

Assumption 4 (Independence of Irrelevant Alternatives)**.** $\frac{\partial}{\partial p_j}\frac{s_{j_1}(\mathbf{p})}{s_{j_2}(\mathbf{p})} = 0$ *for* $j \neq j_1, j_2 \in \mathscr{J} \cup \{0\}$.

Given Assumptions 1, 2 and 4, we can show that the game with multi-product firms can be transformed into a game of single-product firms.

Proposition 3 (Mark-up formula under IIA)**.** *Let Assumptions 1, 2 and Assumption 4 hold and* $-\frac{\partial}{\partial p_j}\frac{s_j(\mathbf{p})}{\frac{\partial s_j}{\partial p_j}} \neq 1$ *for all* $\mathbf{p}$*. Then, there exists an equilibrium of the stage game for any scarcity matrix* Ω*. All equilibrium prices* $\mathbf{p}^*(\Omega, \boldsymbol{\theta})$ *coincide with the equilibrium prices of a game with a set* $\mathscr{J}$ *of players who each simultaneously choose a price* p_j *maximizing*

$$s_j(\mathbf{p})\big(p_j - c_j\left(\mathbf{p}_{-j};\Omega,\boldsymbol{\theta}\right)\big)$$

with a cost function

$$c_j\left(\mathbf{p}_{-j};\Omega,\boldsymbol{\theta}\right) := \omega^f_j - \sum_{j' \in \mathscr{J}_f \setminus \{j\}} \tilde{s}_{j,j'}(\mathbf{p}_{-j})(p_{j'} - \omega^f_j) + \sum_{j' \notin \mathscr{J}_f} \tilde{s}_{j,j'}(\mathbf{p}_{-j})\omega^f_{j'} \tag{8}$$

and $\tilde{s}_{j,j'}(\mathbf{p}_{-j}) := \frac{s_{j'}(\mathbf{p})}{1-s_j(\mathbf{p})}$.

Proof. Let Assumptions 1, 2 and 4 hold. First, note that Assumption 4 implies that for $j_1, j_2 \neq k$

$$\frac{s_{j_1}(\mathbf{p})}{s_{j_2}(\mathbf{p})} = \frac{\frac{\partial s_{j_1}}{\partial p_k}(\mathbf{p})}{\frac{\partial s_{j_2}}{\partial p_k}(\mathbf{p})}.$$

By Step 1 in the proof of Lemma 4 and by Assumption 2, any equilibrium price vector of the stage game $\mathbf{p}^*(\Omega;\boldsymbol{\theta})$ must satisfy for all $j \in \mathscr{J}_f$ the FOCs of firm f's payoff given by:

$$p_j - \omega^f_j + \sum_{j' \in \mathscr{J}_f \setminus \{j\}} \frac{\frac{\partial s_{j'}(\mathbf{p})}{\partial p_j}}{\frac{\partial s_j(\mathbf{p})}{\partial p_j}}(p_{j'} - \omega^f_{j'}) - \sum_{j' \notin \mathscr{J}_f} \frac{\frac{s_{j'}(\mathbf{p})}{\partial p_j}}{\frac{\partial s_j(\mathbf{p})}{\partial p_j}}\omega^f_{j'} = -\frac{s_j(\mathbf{p})}{\frac{\partial s_j(\mathbf{p})}{\partial p_j}}.$$

Since $\frac{\partial s_j}{\partial p_j}(\mathbf{p}) = -\sum_{k\in\mathscr{J}\setminus\{j\}} \frac{\partial s_k}{\partial p_j}(\mathbf{p}) - \frac{\partial s_0}{\partial p_j}$, this can be rewritten as

$$p_j - \omega_j^f - \sum_{j'\in\mathscr{J}_f\setminus\{j\}} \frac{1}{\sum_{k\in\mathscr{J}\setminus\{j\}} \frac{s_k(\mathbf{p})}{s_{j'}(\mathbf{p})} + \frac{s_0(\mathbf{p})}{s_{j'}(\mathbf{p})}}(p_{j'} - \omega_{j'}^f) + \sum_{j'\notin\mathscr{J}_f} \frac{1}{\sum_{k\in\mathscr{J}\setminus\{j\}} \frac{s_k(\mathbf{p})}{s_{j'}(\mathbf{p})} + \frac{s_0(\mathbf{p})}{s_{j'}(\mathbf{p})}} \omega_{j'}^f = -\frac{s_j(\mathbf{p})}{\frac{\partial s_j(\mathbf{p})}{\partial p_j}}$$

$$\Leftrightarrow p_j - \omega_j^f - \sum_{j'\in\mathscr{J}_f\setminus\{j\}} \frac{s_{j'}(\mathbf{p})}{1-s_j(\mathbf{p})}(p_{j'} - \omega_{j'}^f) + \sum_{j'\notin\mathscr{J}_f} \frac{s_{j'}(\mathbf{p})}{1-s_j(\mathbf{p})} \omega_{j'}^f = -\frac{s_j(\mathbf{p})}{\frac{\partial s_j(\mathbf{p})}{\partial p_j}}.$$

By Assumption 4, for $j' \neq j$, $\frac{\partial}{\partial p_j} \frac{s_{j'}(\mathbf{p})}{1-s_j(\mathbf{p})} = 0$, we can define $\tilde{s}_{j,j'}((p_{j'})_{j'\neq j}) := \frac{s_{j'}(\mathbf{p})}{1-s_j(\mathbf{p})}$ and

$$c((p_{j'})_{j'\neq j}; \Omega) := \omega_j^f + \sum_{j'\in\mathscr{J}_f\setminus\{j\}} \tilde{s}_{j,j'}((p_{j'})_{j'\neq j})(p_{j'} - \omega_j^f) - \sum_{j'\notin\mathscr{J}_f} \tilde{s}_{j,j'}((p_{j'})_{j'\neq j})\omega_{j'}^f.$$

Thus, the FOCs of the stage game are equivalent to the first order conditions of a game with $\mathscr{J}$ players where each player j's payoff is given by

$$s_j(\mathbf{p})\big(p_j - c((p_{j'})_{j'\neq j}; \Omega)\big).$$

We call this game the "auxiliary game with J players." Note that the derivative of player j's payoff is greater or equal than zero if and only if

$$\frac{\partial s_j(\mathbf{p})}{\partial p_j}\big(p_j - c\left((p_{j'})_{j'\neq j}; \Omega\right)\big) + s_j(\mathbf{p}) \geq 0.$$

Hence any equilibrium of the stage game is an equilibrium of a game with J players with the above payoffs and vice versa.

In order to show existence of equilibria of the stage game, it is sufficient to show existence of equilibria of the auxiliary game with J players and the above payoffs. First, recall that by Step 1 in the proof of Lemma 4, all best response prices are interior and hence, if an equilibrium exists, it must satisfy the FOCs. Further, since we assume $-\frac{\partial}{\partial p_j} \frac{s_j(\mathbf{p})}{\frac{\partial s_f}{\partial p_f}} \neq 1$ for all $\mathbf{p}$, the first-order condition has a unique solution which must be a maximizer of player j's payoff function. All in all, the best response function of player j, $\mathscr{R}_j$, maps a compact set of prices $\mathbf{q}$ into a compact set of prices $\mathbf{p}$. For $\epsilon > 0$, consider the mapping

$$\Phi : (\mathbf{p}, \mathbf{q}) \mapsto \left(p_j - \epsilon \Big(p_j - c_j(\mathbf{q}_{-j}; \Omega, \boldsymbol{\theta}) + \frac{s_j(\mathbf{q}_{-j}, p_j)}{\frac{\partial s_j(\mathbf{q}_{-j}, p_j)}{\partial p_j}} \Big) \right)_{j \in \mathcal{J}}$$

Then $D_{\mathbf{p}}\Phi$ is a diagonal matrix with diagonal entries

$$\phi_j := 1 - \epsilon \Bigg(1 + \underbrace{\frac{\partial}{\partial p_j} \frac{s_j(\mathbf{q}_{-j}, p_j)}{\frac{\partial s_j(\mathbf{q}_{-j}, p_j)}{\partial p_j}}}_{\geq 0} \Bigg)$$

Let $\epsilon > 0$ be so that $\phi_j > 0$ for all j. Then all diagonal entries are in $(0, 1-\epsilon)$ and Φ is Lipschitz continuous with Lipschitz constant $\max_j \phi_j$. Further $D_{\mathbf{q}}\Phi$ is bounded because it is continuous. Then, the implicit function theorem in the form of Theorem 1.A.4 in Dontchev and Rockafellar (2009) implies continuity of $\mathcal{R} = ((\mathcal{R}_j)_j)$. Hence, by Brouwer's fixed-point theorem $\mathcal{R} = ((\mathcal{R}_j)_j)$ has a fixed point. ■

Proposition 3 implies that even with multiple firms and products, the first-order conditions (FOCs) that implicitly define the best response functions of the firms, can be written in a markup formulation for each product, with $\epsilon_j(\mathbf{p}) = \frac{\partial s_j(\mathbf{p})}{\partial p_j} \frac{p_j}{s_j(\mathbf{p})}$ being the elasticity of demand, as

$$\frac{p_j^*(\Omega, \boldsymbol{\theta}) - c_j(\mathbf{p}_{-j}; \Omega, \boldsymbol{\theta})}{p_j^*(\Omega, \boldsymbol{\theta})} = -\frac{1}{\epsilon_j(\mathbf{p}^*(\Omega, \boldsymbol{\theta}))}.$$

C Nested Logit Calculations

Since our empirical application uses a nested logit specification, we verify in the following that all assumptions made in the model are satisfied for a nested logit demand model given by

$$s_j(\mathbf{p}) = \underbrace{\frac{e^{\frac{\delta_j - \alpha p_j}{1-\sigma}}}{\sum_{j \in \mathcal{J}} e^{\frac{\delta_j - \alpha p_j}{1-\sigma}}}}_{=: s_{j|\mathcal{J}}(\mathbf{p})} \frac{\left(\sum_{i \in \mathcal{J}} e^{\frac{\delta_i - \alpha p_i}{1-\sigma}} \right)^{1-\sigma}}{1 + \left(\sum_{i \in \mathcal{J}} e^{\frac{\delta_i - \alpha p_i}{1-\sigma}} \right)^{1-\sigma}} \qquad s_0(\mathbf{p}) = \frac{1}{1 + \left(\sum_{i \in \mathcal{J}} e^{\frac{\delta_i - \alpha p_i}{1-\sigma}} \right)^{1-\sigma}}.$$

Note that the same properties follow for regular logit by setting $\sigma = 0$ and replacing α with $\frac{\alpha}{\rho}$. To simplify notation, let $D_{\mathscr{J}} := \sum_{i \in \mathscr{J}} e^{\frac{\delta_i - \alpha p_i}{1-\sigma}}$ and $G := \sigma \frac{1+D_{\mathscr{J}}^{1-\sigma}}{D_{\mathscr{J}}^{1-\sigma}} + 1 - \sigma$. Then,

$$\begin{aligned} \frac{\partial s_j}{\partial p_j} = -\frac{\alpha}{1-\sigma} s_j \left(1 - \left(\sigma s_{j|\mathscr{J}} + (1-\sigma) s_j\right)\right) &= \frac{\alpha}{1-\sigma}(G s_j^2 - s_j) \\ \frac{\partial s_j}{\partial p_{j'}} = \frac{\alpha}{1-\sigma} s_{j'} \left(\sigma s_{j|\mathscr{J}} + (1-\sigma) s_j\right) &= \frac{\alpha}{1-\sigma} G s_{j'} s_j. \end{aligned}$$

It is easy to check that Assumptions 1-i) and ii) are satisfied. We show that Assumption 1-iii) is satisfied. Letting $\underline{s}_0 \equiv s_0\left(\underline{\mathbf{p}}\right)$, we can set $C = \alpha \underline{s}_0 > 0$ since then

$$\frac{\partial s_0}{\partial p_j} = \alpha s_j s_0 > C s_j.$$

Then,

$$(D_{\mathbf{p}} \mathbf{s}(\mathbf{p}; \boldsymbol{\theta}))^{-1} = -\frac{1}{\alpha s_0} \cdot \begin{pmatrix} 1 + \sigma D_{\mathscr{J}}^{\sigma-1} + (1-\sigma)\frac{s_0}{s_1} & 1 + \sigma D_{\mathscr{J}}^{\sigma-1} & \dots & 1 + \sigma D_{\mathscr{J}}^{\sigma-1} \\ 1 + \sigma D_{\mathscr{J}}^{\sigma-1} & \ddots & & 1 + \sigma D_{\mathscr{J}}^{\sigma-1} \\ \vdots & & \ddots & \vdots \\ 1 + \sigma D_{\mathscr{J}}^{\sigma-1} & \dots & \dots & 1 + \sigma D_{\mathscr{J}}^{\sigma-1} + (1-\sigma)\frac{s_0}{s_J} \end{pmatrix}$$

$$(D_{\mathbf{p}} \mathbf{s}(\mathbf{p}; \boldsymbol{\theta}))^{-1} = -\frac{1-\sigma}{\alpha} \cdot \begin{pmatrix} \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} + \frac{1}{s_1} & \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} & \dots & \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} \\ \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} & \ddots & & \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} \\ \vdots & & \ddots & \vdots \\ \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} & \dots & \dots & \frac{G+\sigma+D_{\mathscr{J}}^{1-\sigma}}{1-\sigma} + \frac{1}{s_J} \end{pmatrix}$$

Hence, $\hat{\boldsymbol{\epsilon}} = ((D_p \mathbf{s}(\mathbf{p}; \boldsymbol{\theta}))^{\intercal})^{-1} \mathbf{s}(\mathbf{p}; \boldsymbol{\theta}) = -\frac{1}{\alpha s_0} \mathbf{1}$ and noting that $\frac{\partial}{\partial p_j}\left(\frac{1}{s_0}\right) = -\alpha \frac{s_j}{s_0}$,

$$D_{\mathbf{p}} \hat{\boldsymbol{\epsilon}} = \begin{pmatrix} \frac{s_1}{s_0} & \dots & \frac{s_J}{s_0} \\ & \ddots & \\ \frac{s_1}{s_0} & \dots & \frac{s_J}{s_0} \end{pmatrix}.$$

It follows that Assumption 2 is satisfied:

$$\det\left(-D_{\mathbf{p}}\hat{\boldsymbol{\epsilon}} - I\right) = (-1)^J \frac{1}{s_0} \neq 0.$$

Finally, we show that Equation 7 holds. To this end, note that we can write the static consumer surplus for our demand specification as

$$CS = \frac{1}{\alpha_t}\log\left(1 + D_{\mathscr{J}}^{1-\sigma}\right) = \frac{1}{\alpha_t}\log\left(1 + \left(\sum_{j\in\mathscr{J}}\exp\left(\frac{\delta_j - \alpha_t p_j}{1-\sigma}\right)\right)^{1-\sigma}\right)$$

and hence $D_{\mathbf{p}}CS = -\mathbf{s}(\mathbf{p})$. We can write the per-period welfare as

$$w_t(\mathbf{p}) = CS + \sum_{j\in\mathscr{J}} s_j(\mathbf{p};\boldsymbol{\theta})p_j.$$

Then, the objective function in Equation 7 boils down to

$$CS + \sum_{j\in\mathscr{J}} s_j(\mathbf{p};\boldsymbol{\theta})(p_j - \omega_j).$$

The FOC of the optimization problem is then given by $\mathbf{p} = \boldsymbol{\omega}$. This implies the uniqueness of the solution for all $\boldsymbol{\omega}$.

It follows immediately that all properties are satisfied for all subsets $\mathscr{A} \subset \mathscr{J}$.

D Additional Tables and Figures

D.1 Simulations

Figure 14: Strategic complements and substitutes in the stage game

(a) Strategic complements

(b) Strategic substitutes

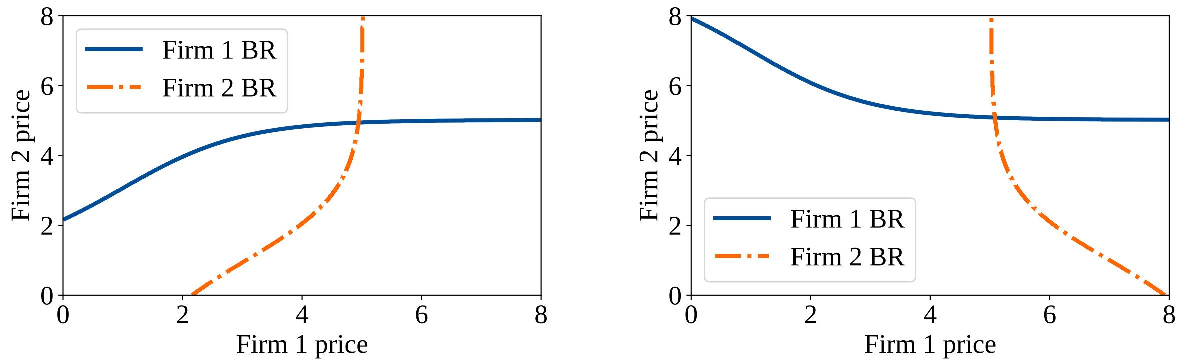

Notes: The simulations assume $\delta = (1,1)$, $\alpha_t \equiv 1$ and logit demand with scaling factor $\rho = 1$, as well as $\omega_1^1 = \omega_2^2 = 4$. Panel (a) shows both firms' best response functions for $\omega_2^1 = \omega_1^2 = 4$. Panel (b) shows both firms' best response functions for $\omega_2^1 = \omega_1^2 = -4$.

Figure 15: Effects of own and competitor scarcity on prices

(a) Exogenous Ω

(b) Dynamic game Ω

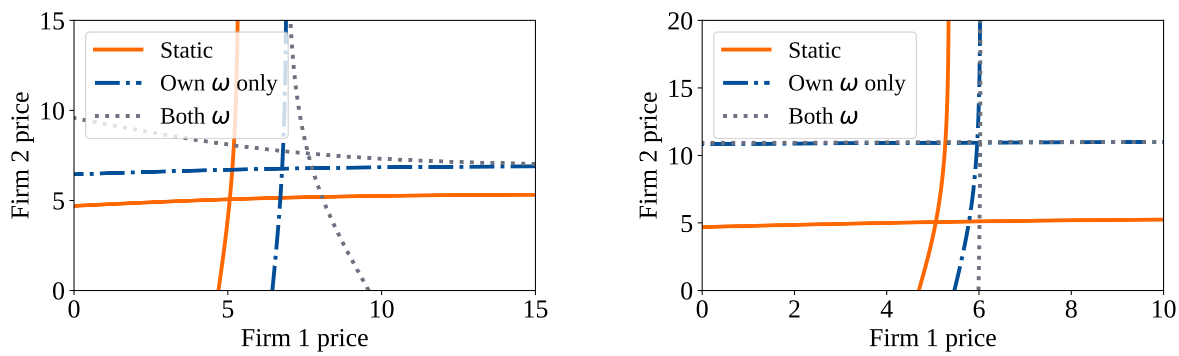

Notes: The simulations assume $\delta = (1,1)$, $\alpha_t \equiv 1$ and logit demand with scaling factor $\rho = 4$. Panel (a) shows both firms' best response functions for $\omega_1^1 = \omega_1^2 = 2$ and $\omega_2^1 = \omega_2^2 = -6$ when no ωs are considered in the profits (orange), when only the own ωs are considered (blue), and when both ωs are considered (grey). Panel (b) shows an analogous figure for the Ω matrix obtained at $t = 0$ in the dynamic duopoly game with $T = 2$ and $\lambda_t \equiv 10$ at the state $\mathbf{K} = (20,1)$.

Figure 16: Simulated prices and scarcity effects

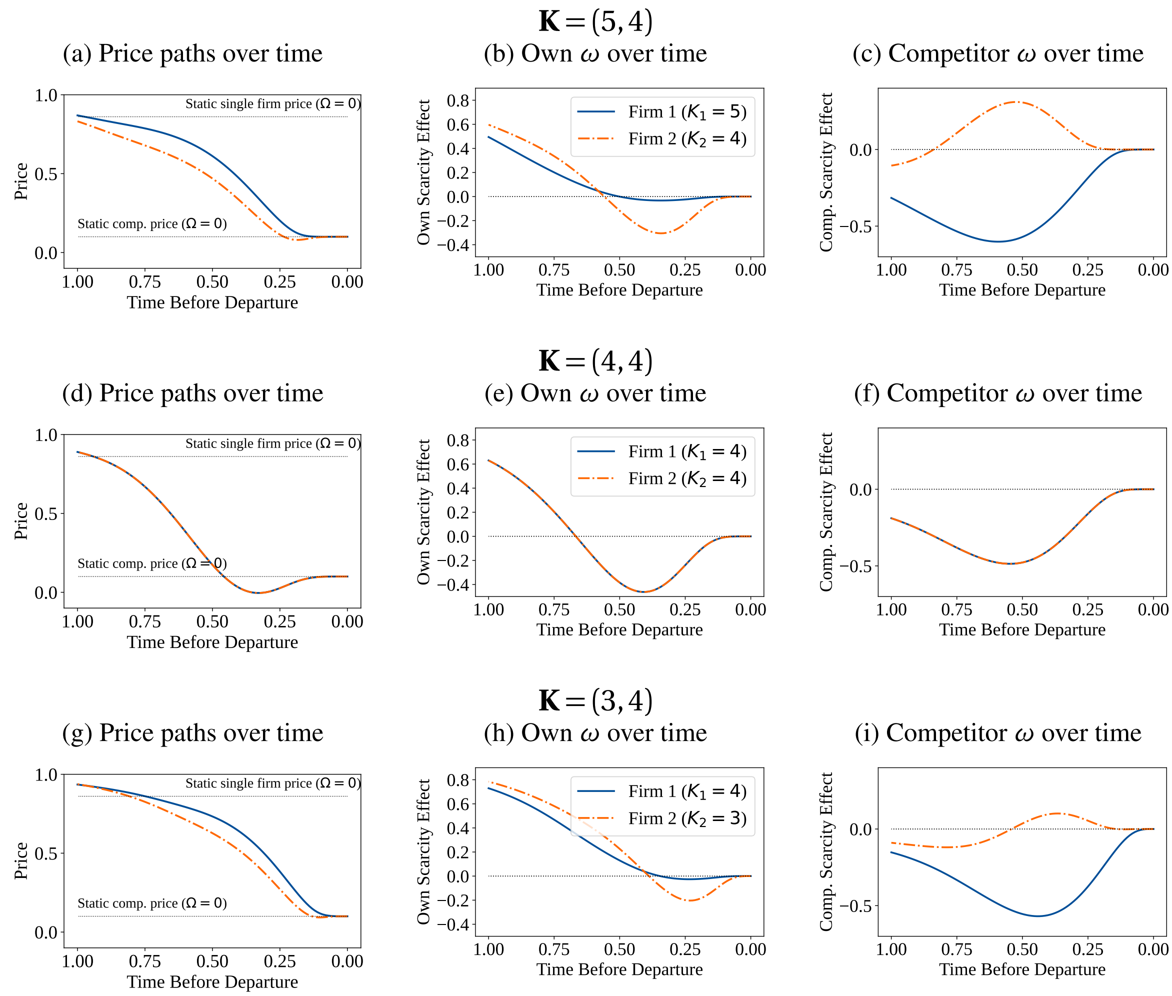

Notes: The simulations assume $\delta = (1,1)$, $\alpha_t \equiv 1$ and logit demand with scaling factor $\rho = 0.05$.

Figure 17: Price paths for varying levels of capacity

(a) Sale of a product with minimum inventory (b) Sale of a product without minimum inventory

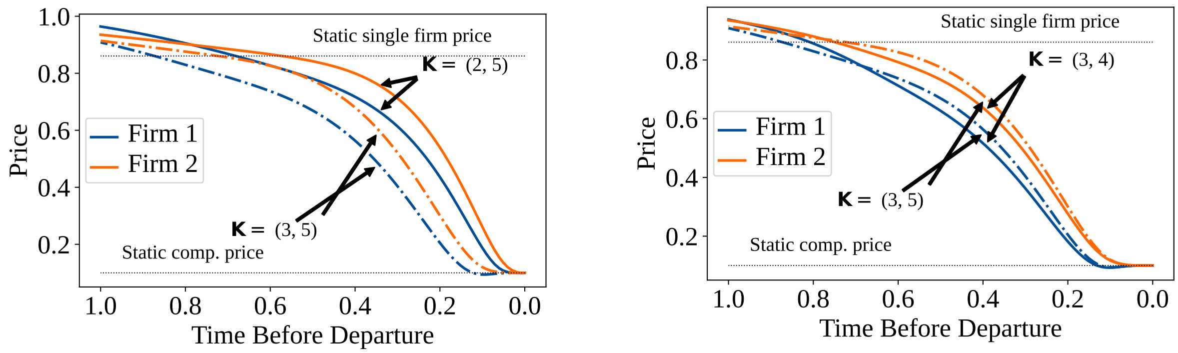

Notes: These simulations correspond to logit demand with parameter values $\delta_j = 1$, $\alpha = 1$, $\lambda = 10$ and scale factor $\rho = 0.05$. Panel (a) shows both firm's price paths for $\mathbf{K} = (3,5)$ and $\mathbf{K} = (2,5)$. Panel (b) shows both firm's price paths for $\mathbf{K} = (3,5)$ and $\mathbf{K} = (3,4)$.

Figure 18: Bertrand scarcity trap for with increasing arrivals and constant nested logit

(a) Constant logit demand with increasing λ_t (b) Constant nested logit demand

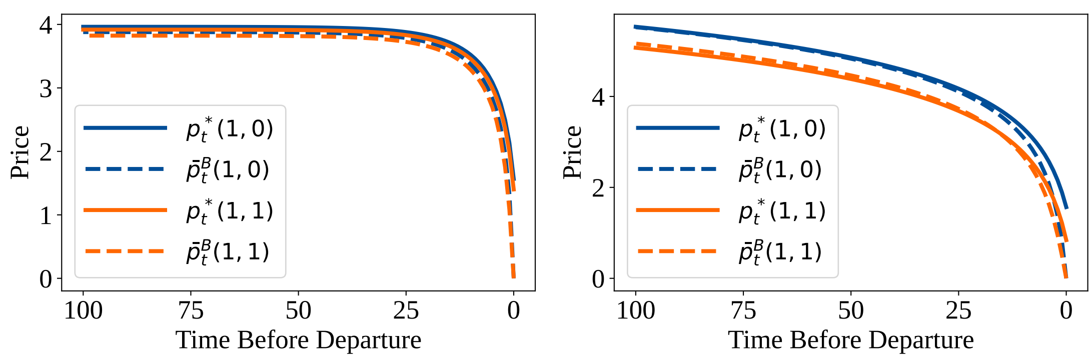

Note: The simulations assume $\mathbf{K} = (1,1)$. Panel (a) uses the same demand system as in Figure 4 with $\lambda_t = 2\exp(-(T-t)/100)$; panel (b) uses constant nested logit demand given by $s_f(\mathbf{p}) = \frac{\exp\left(\frac{1-p_j}{1-\sigma}\right)}{\sum_{f'\in\{1,2\}}\exp\left(\frac{1-p_{f'}}{1-\sigma}\right)} \cdot \frac{\left(\sum_{f'\in\{1,2\}}\exp\left(\frac{1-p_{f'}}{1-\sigma}\right)\right)^{1-\sigma}}{1+\left(\sum_{f'\in\{1,2\}}\exp\left(\frac{1-p_{f'}}{1-\sigma}\right)\right)^{1-\sigma}}$.

D.2 Empirical Evidence of Dynamic Pricing Forces

Figure 19: Example of a negative own Opportunity Costs

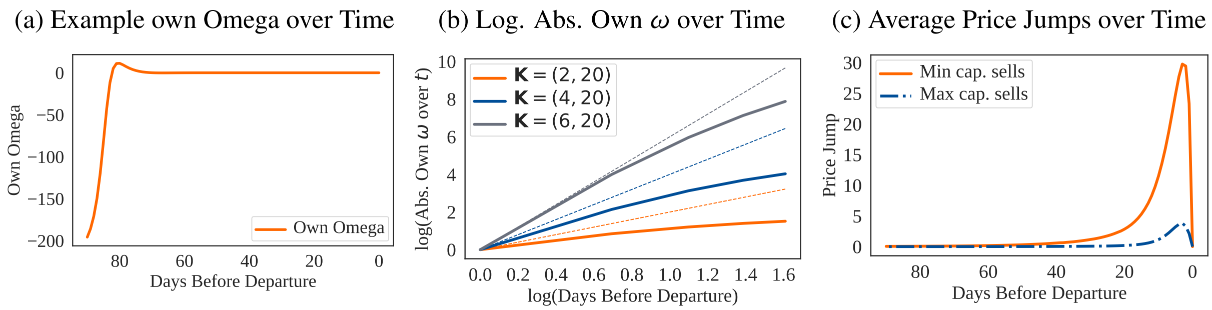

Note: Panel (a) shows the own ω over time for a given state in one of our benchmark solutions. Panel (b) shows the log of the absolute value of the own ω over time for three states in one of our Benchmark solutions. The dotted lines represent the behavior these curves would follow if the omegas were proportional to $|T-t|^{\min(\mathbf{K})}$. Panel (c) shows the price change if the firm with the minimum and maximum capacities sell a unit.

D.3 Welfare Calculations with Restricted Capacities

Figure 20: Competitive Dynamic Pricing and Uniform Pricing

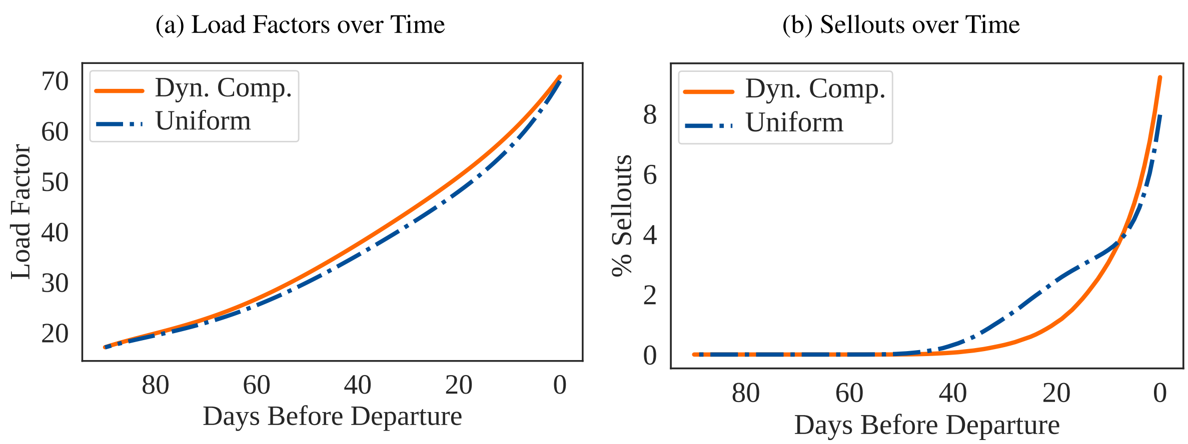

Note: Panel (a) shows average load factors over time for uniform pricing and dynamic price competition. Panel (b) shows the average sellouts over time for the same two models.

Figure 21: Counterfactual Summary Plots, Restricted Capacities

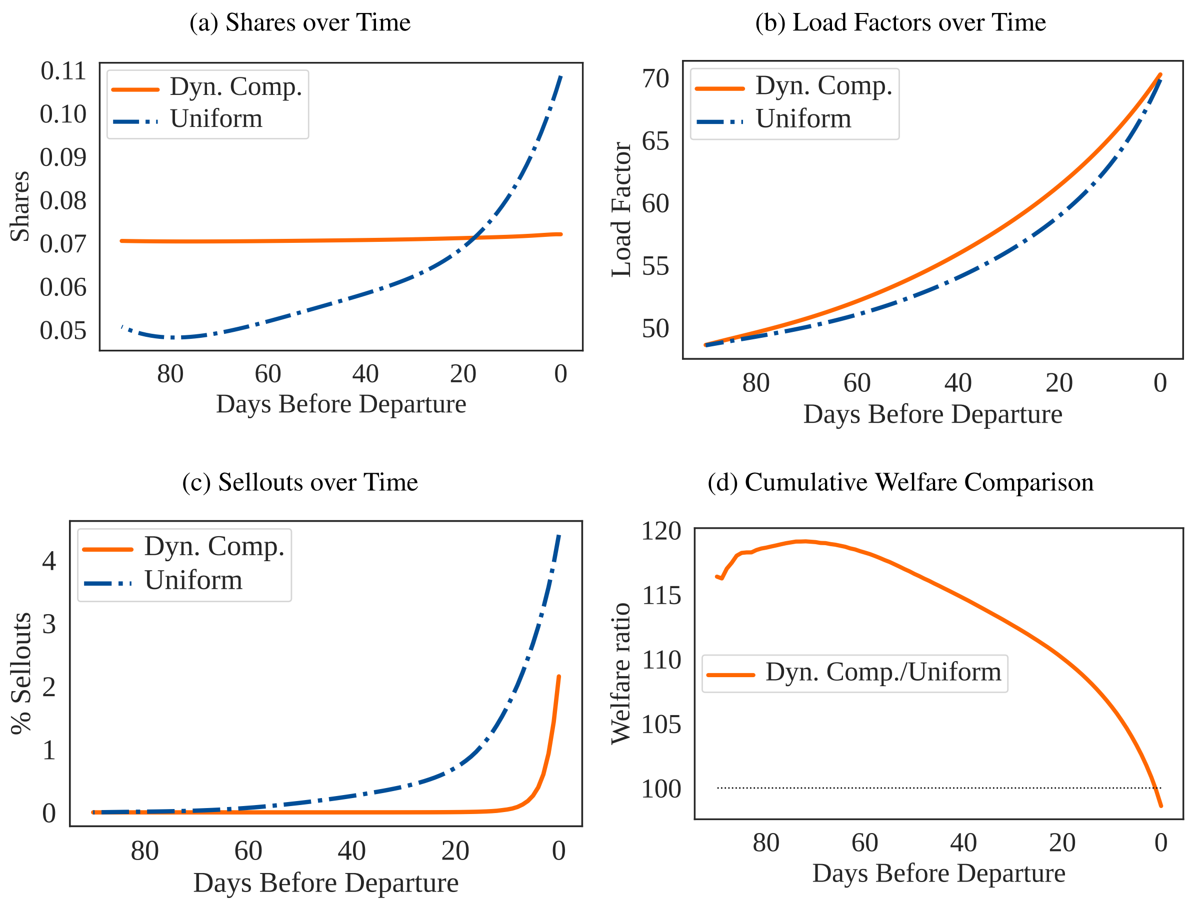

Note: Panel (a) shows the average shares over time for the benchmark and uniform models. Panel (b) shows the average load factors over time for the same two models. Panel (c) shows the average sellouts over time for the same two models. Panel (d) shows the ratio of average cumulative welfare for the benchmark model with respect to the uniform one.

Table 6: Counterfactual Results for Single Product, Duopoly Routes, Restricted Capacities

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	228.4	5216.1	5535.3	16561.3	27312.8	18.9	70.3	2.2
Uniform	243.7	4362.3	4533.4	18804.3	27700.0	18.5	69.8	4.4
% Diff.	6.7	-16.4	-18.1	13.5	1.4	-2.0	-0.5	2.2

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.

Figure 22: Competitive Dynamic Pricing and Pricing Heuristics, Restricted Capacities

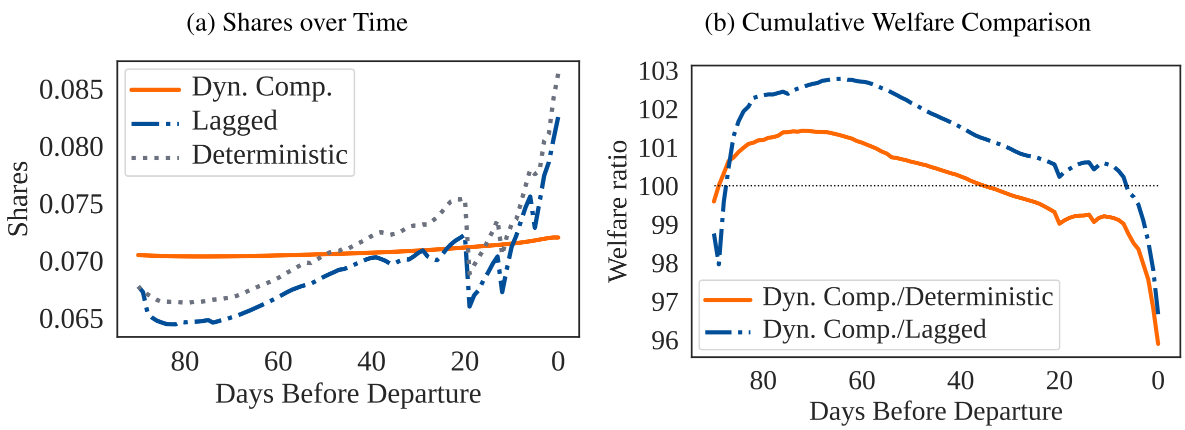

Note: (a) shows purchase probabilities for for three models over time. (b) shows the cumulative welfare under dynamic price competition relative to heuristics over time.

Table 7: Heuristic Counterfactuals for Single Product, Duopoly Routes, Restricted Capacities

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Lagged	105.3	103.0	104.6	103.2	103.5	99.1	99.8	97.8
Deterministic	99.6	99.7	100.2	107.1	104.3	102.4	101.0	123.2

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.

Figure 23: Heuristic Counterfactuals Results over Time, Restricted Capacities

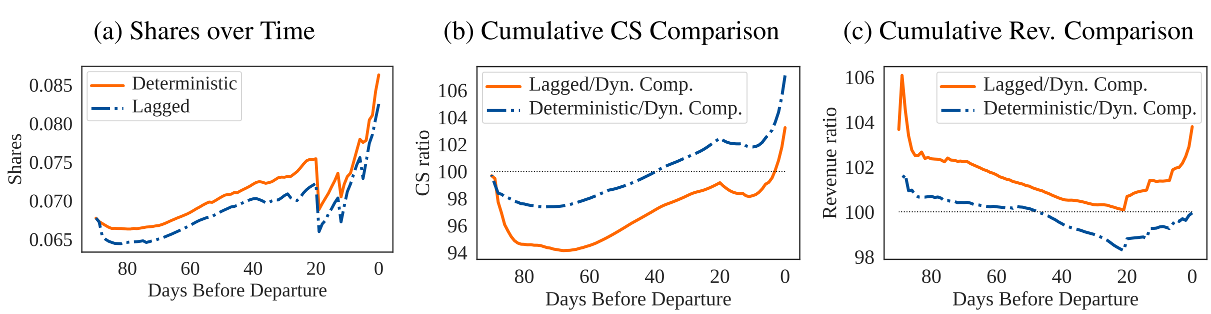

Note: Panel (a) shows the average shares over time for the two heuristic models. Panel (b) shows the ratios of cumulative consumer surplus for the two models with respect to the benchmark. Panel (c) shows the ratios of cumulative revenue for the two models with respect to the benchmark.

Figure 24: Bertrand Scarcity Trap over Time

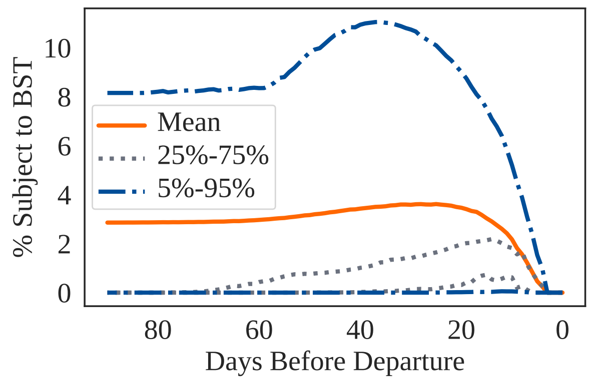

Note: Plotted is the percentage of simulated observations subject to the Bertrand Scarcity Trap. The percentage is first calculated for a given route (r) and days from departure (DFD). The mean and percentiles shown are then taken across routes.

Table 8: Social Planner and Single Firm Counterfactuals for Single Product, Duopoly Routes

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Social Planner	12.0	9.5	20.0	179.9	113.2	171.2	120.7	322.5
Single Firm	151.7	107.5	112.1	70.1	86.2	71.7	91.2	67.1

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.

Table 9: Social Planner and Single Firm Counterfactuals for Single Product, Duopoly Routes, Restricted Capacities

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	100.0	100.0	100.0	100.0	100.0	100.0	100.0	100.0
Social Planner	19.4	13.9	27.9	178.9	116.8	167.9	120.0	413.6
Single Firm	150.9	108.0	110.7	69.4	85.1	71.2	91.1	36.1

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.

D.4 Welfare Calculations for Entire Sample

i) In our counterfactuals we consider only two products. In order to include routes that have more than one flight per carrier per day, we adjust the choice set, utilities, and capacities for all routes.

ii) We take the mean utilities (δ) across observed flights for each route-carrier-departure date.

iii) We use the maximum observed capacity for each route-carrier-departure date. Although it may be natural to sum the capacities when restricting the choice set, we have found that large capacities presents a significant computational burden.

iv) We use the observed arrival process for each route-departure date. We do not adjust the estimated arrival processes as the inside good shares tend to be small. That is, because most consumers choose the outside good, we do not scale down arrival rates to account for smaller choice sets.

Table 10: Counterfactual Results for Entire Sample

	Price	Firm 1 Rev.	Firm 2 Rev.	CS	Welfare	Q	LF	Sellouts
Dyn. Comp.	220.4	5489.5	5925.6	16504.4	27919.5	20.1	80.1	20.4
Uniform	261.8	4729.7	5192.9	18711.8	28634.4	18.8	78.6	15.0
% Diff.	18.8	-13.8	-12.4	13.4	2.6	-6.4	-1.5	-5.4

Note: Price is the average across routes (r) after computing the average across firms (f), departure dates (DD), days before departure (DFD) and simulation number (n) within a route. Firm revenues are similarly defined, except aggregated over DFD. CS is the expected consumer surplus, computed the same way as revenues. Welfare is the sum of revenues and CS. Q is the total number of seats sold. LF is the average fraction of seats sold (including flow traffic) at the departure time. Sellouts is the fraction of flights sold out.