

Discussion Paper Series – CRC TR 224

Discussion Paper No. 020
Project B 03

All-Pay Oligopolies: Price Competition With
Unobservable Inventory Choices

Joao Montez¹
Nicolas Schutz²

May 2018

¹ Faculty of Business and Economics, University of Lausanne and CEPR. Email: joao.montez@unil.ch
² University of Mannheim. Email: schutz@uni-mannheim.de

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

All-Pay Oligopolies: Price Competition with Unobservable Inventory Choices*

Joao Montez[†] Nicolas Schutz[‡]

May 27, 2018

Abstract

We study a class of games where stores source unobservable inventories in advance, and then simultaneously set prices. Our framework allows for firm asymmetries, heterogeneous consumer tastes, endogenous consumer information through advertising, and salvage values for unsold units. The payoff structure relates to a complete-information all-pay contest with outside options, non-monotonic winning and losing functions, and conditional investments. In the generically unique equilibrium, stores randomize their price choice and, conditional on that choice, serve all their targeted demand—thus, some inventories may remain unsold. As inventory costs become fully recoverable, the equilibrium price distribution converges to an equilibrium of the associated Bertrand game (where firms first choose prices and then produce to order). This suggests that with production in advance, the choice between a Cournot analysis and a Bertrand-type analysis, as properly generalized in this paper, should depend on whether or not stores observe rivals’ inventories before setting prices.

Keywords: Oligopoly, inventories, production in advance, all-pay contests, Bertrand convergence.

1 Introduction

In retail markets, each independent store typically chooses not only a sale price but also an inventory level—a quantity sourced from suppliers to be made readily available to consumers.

*We thank Simon Anderson, Pierre Boyer, Bruno Jullien, Sébastien Mittraille, Hans-Theo Normann, Martin Peitz, Régis Renault, Patrick Rey, Andre Speit, Yossi Spiegel, Thomas Tröger, Thomas von Ungern-Sternberg, Ralph Winter, as well as seminar and conference participants at Bern, Cergy, Dusseldorf, Frankfurt, Leicester, Pompeu Fabra, X-CREST, CRC TR 224 (spring 2018), EARIE 2017, MaCCI Summer Institute 2016, MaCCI Annual Meeting 2017, and MaCCI IO Day 2017, Workshop on Competition and Bargaining in Vertical Chains (Dusseldorf), for helpful comments and suggestions. The first author gratefully acknowledges financial support from the Swiss National Science Foundation through grant no. 00018_176448. The second author gratefully acknowledges financial support from the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224.

[†]Faculty of Business and Economics, University of Lausanne and CEPR. Email: joao.montez@unil.ch.

[‡]University of Mannheim. Email: schutz@uni-mannheim.de.

Well-informed consumers visit the store with the lowest price and only if unable to purchase there (due to a stock-out) may they visit another store. However, it is hard to observe rivals' inventories, and historical data is unlikely to provide a reliable estimate—as inventories are transient by nature. Stores are therefore likely to make price and inventory choices without knowledge of rivals' choices.

It is natural to model such situations as a game where stores simultaneously choose a price-inventory pair, as such a model is formally equivalent to the arguably more realistic sequential game in which stores first make an inventory choice and then, without observing rivals' inventories, choose prices.

We introduce and study a class of such games, which we call all-pay oligopolies. This class can account for firm asymmetries, heterogeneity in consumer tastes, and the coexistence of informed and uninformed consumers. As unsold inventories often have some (partial) salvage value, we also allow a fraction of the unit cost to remain variable and be incurred only once a unit is sold (while the remainder is sunk). This variable fraction may also account for the cost of sales and post-sales services.

Beyond retail markets where our assumptions seem to apply best, it will become clear that a considerable variety of market situations can be usefully viewed through the conceptual lens of our model. It is for example possible to interpret it as describing a situation in which firms first secretly invest in capacity (for which the cost is sunk), and subsequent costly production takes place (for which the cost is variable) once prices have been set and consumers have made their purchase decisions.

We start by solving a constrained version of the model in which each store must source enough inventory to serve all its targeted demand at the price it chooses. This *constrained game* has the structure of an all-pay contest with outside options, non-monotonic winning and losing functions, and conditional investments. Based on this insight, we develop a method to obtain the closed-form characterization of the equilibrium set of such constrained games, and show that generically, there is a unique equilibrium.

Next, we show that this equilibrium is also the generically unique equilibrium of the *unconstrained game* (where firms freely choose inventory levels). In equilibrium, each store randomizes its price, ordering a low inventory when it sets a high price, and a high inventory when it holds a sale. Because each store holds enough inventory to serve all its targeted demand, the aggregate inventory level often exceeds demand, resulting in unsold inventories.

Large unsold inventories are observed in several industries, including grocery and apparel.¹ A common explanation is market demand uncertainty. In our model, market demand is deterministic, but *store-level* demand is stochastic due to the uncertain behavior of rivals. Stores' inability to anticipate the timing and depth of rivals' promotions makes it hard to adjust inventory purchases accordingly. Such strategic uncertainty provides another explanation for the persistence of unsold inventories.

This mechanism also offers an explanation as to why some products with a low variance

¹For example, U.S. supermarkets and grocery stores threw out \$46 billion worth of food in 2010 according to the USDA, and fashion retailer H&M has accumulated more than \$4 billion in unsold clothes in 2018.

in consumer demand exhibit a large variance in production, a well-documented fact known in the operations literature as the bullwhip effect. This term was coined by Procter and Gamble when it noticed that the volatility of diaper orders it received from retailers was quite high, even though it was (for obvious reasons) confident that end-consumer demand was reasonably stable (see, e.g., Lee, Padmanabhan, and Whang, 1997b). A similar effect has been found, for example, in orders for Barilla pasta (Hammond, 1994) and Hewlett-Packard printer cartridges (Lee, Padmanabhan, and Whang, 1997a). Our model suggests the bullwhip effect could, in part, be explained by retailers frequently offering discounts to attract price-conscious consumers to their stores.²

This model also provides a clear avenue to explore how oligopoly behavior under production in advance is affected by the extent to which inventory costs are sunk or recoverable. In particular, we find that as the fraction of the inventory cost that can be recovered tends to one, the equilibrium distribution of prices converges to an equilibrium of the associated Bertrand game, in which stores only choose prices and produce to order (i.e., source inventories to meet demand only after consumers made purchase decisions).

The equilibrium is thus said to be *Bertrand convergent*. Several oligopoly theory benchmarks (where production to order is typically assumed) can then be seen as the limiting outcome of similar situations with production in advance. Examples include standard Bertrand models with heterogeneous marginal or fixed costs (e.g., Blume, 2003; Anderson, Baik, and Larson, 2015), and Bertrand models where informed and uninformed consumers coexist, with or without advertising (e.g., Varian, 1980; Baye and Morgan, 2001).

This insight stands in contrast to Kreps and Scheinkman (1983)’s well-known result. Assuming that inventory choices become *observable* before the pricing stage, firms are symmetric, and rationing is efficient, they find that the Cournot outcome always arises in equilibrium, regardless of the salvage value of inventories. They conclude that situations where almost all the unit cost “is incurred subsequent to the realization of demand (situations that will look very Bertrand-like) will still give the Cournot outcome” (Kreps and Scheinkman, 1983, p. 337). The conventional interpretation of this result is that a Cournot analysis is appropriate if inventories (or capacities) must be chosen before prices, whereas a Bertrand analysis should be preferred when inventories can be sourced after prices have been set and consumers have decided where to purchase (see, e.g., Belleflamme and Peitz, 2010, pp. 66–67).

Our analysis suggests a more nuanced model selection, which explicitly takes into account whether inventory or capacity choices are *observed* by rivals or not. If inventory choices are observable, then the Cournot outcome is a reasonable benchmark. If instead inventory information is private, the Kreps-Scheinkman mechanism fails as a low inventory choice can no longer provide a commitment to soften price competition. In that case, our results suggest that a Bertrand approach is better justified (by continuity) if most of the unit cost

²Other explanations include adjustment to cost shocks (Blinder, 1986), increasing returns to scale upstream (Ramey, 1991), and stochastically-evolving demand (Kahn, 1987; Lee, Padmanabhan, and Whang, 1997b; Chen, Drezner, Ryan, and Simchi-Levi, 2000). At the macro level, the variance of production is also typically greater than that of demand (e.g., Blanchard, 1983).

can be salvaged. Yet, in many industries, the value of unsold inventories falls short of their acquisition value—for instance, when goods are perishable, or due to inventory holding costs (see also the dynamic extension in the conclusion).³ Above all, our closed-form equilibrium characterization generalizes the Bertrand-type analysis to those situations.

The fact that inventories may be unobservable also has implications for public policy. When inventories are observable, it follows for the model of Kreps and Scheinkman (1983) that an output subsidy alleviates the market power distortion and raises social welfare. However, when inventories are unobservable in that same symmetric model, a similar policy actually decreases social welfare. Indeed, despite the classic deadweight loss and the unsold inventories distortion, we show that the equilibrium outcome cannot be improved by standard taxation—such as combinations of linear and symmetric taxes or subsidies on sales, output, or unsold units. *Laissez-faire* is thus second-best efficient.

As the equilibrium involves mixed strategies, one may ask whether store managers should be expected to play dice. To address this concern, we add to a symmetric version of the model an idiosyncratic and privately-observed shock to firms’ per-unit costs. The resulting incomplete-information game has a unique equilibrium. As in the complete-information case, each store acts as a monopolist facing a stochastic residual demand. The novelty is that, for any cost realization, there is a unique price-inventory pair that solves a given store’s profit maximization problem, thus resulting in a strict pure-strategy equilibrium.

The incomplete-information and complete-information equilibria are qualitatively similar: Firms source enough inventory to serve their targeted demand, there is price dispersion (since higher cost-types set higher prices), and the equilibrium is Bertrand convergent. Moreover, the equilibrium price distribution of the incomplete-information game converges to the complete-information one as the cost distribution converges to a mass point—thereby providing a purification argument in the spirit of Harsanyi (1973).

The remainder of the paper is organized as follows. Section 1.1 relates our work to the existing literature. Section 2 builds intuition by solving a simple but non-generic example. Section 3 introduces and studies the rich class of all-pay oligopolies. Section 4 provides applications and relates them to classical Bertrand analysis. Section 5 analyzes an all-pay oligopoly with incomplete information. Section 6 studies efficiency and taxation. Section 7 discusses the N -firm case (studied formally in Online Appendix IV) and concludes.

1.1 Related Literature

Our analysis of the constrained game in Section 3.1, where firms must source enough inventories to supply their targeted demand, contributes to the literature on all-pay contests. In an all-pay contest with complete information, as thoroughly studied by Siegel (2009, 2010), there is a fixed number of prizes, each player submits a score, and prize winners are the players with the highest score. In much of the literature, including Baye, Kovenock, and

³Inventory holding costs include, among others, the cost of storage space, labor, insurance, and the opportunity cost of capital.

de Vries (1993), Che and Gale (1998, 2006), and Kaplan and Wettstein (2006), a player’s payoff conditional on winning or losing decreases continuously with his score, and the difference between the winning and losing payoffs is a constant—the value of the prize.

In our oligopoly setting, non-monotonic winning and losing payoffs arise as a direct consequence of market revenue concavity, and the difference between the winning and the losing payoff is not constant. All-pay contests with related properties have been studied in Kaplan, Luski, and Wettstein (2003), Siegel (2014a, b), and Chowdhury (2017). There are however differences. For instance, our losing functions are typically discontinuous in participation, due to fixed costs and to firms having the option to focus on their captive consumers only. Also, the weak and strong firms’ winning functions may cross in our model, as a firm may be advantaged in one dimension (e.g., have a lower unit cost), but disadvantaged in others (e.g., have a higher fixed cost)—the same holds for losing functions. These features, which arise naturally in oligopolies, affect the equilibrium structure. For example, the support of equilibrium prices may contain gaps, and a player may use multiple mass points.

Our Bertrand convergence result in Section 3.4 relates to Siegel (2010)’s finding that, in a complete-information all-pay auction, as payments become entirely conditional on winning, equilibrium play converges to the equilibrium of the limiting first-price auction—the Bertrand outcome of that model. In our model, as inventory costs become fully variable, the equilibrium converges to the equilibrium of the associated Bertrand game, where firms set prices and produce to order. Special cases of such games include asymmetric Bertrand models with affine costs (e.g., Marquez, 1997; Blume, 2003; Kartik, 2011; Anderson, Baik, and Larson, 2015) and clearinghouse models (e.g., Varian, 1980; Narasimhan, 1988; Baye, Kovenock, and de Vries, 1992; Baye and Morgan, 2001; Shelegia and Wilson, 2016). It has long been recognized that some such production-to-order games share characteristics with all-pay contests. Our work contributes to a better understanding of this connection.

Under incomplete information, in Section 5.1, we establish equilibrium existence and uniqueness in an all-pay contest with non-monotonic winning functions—a class of models that has received little attention in the literature.⁴ The limiting Bertrand game under incomplete information was studied by Spulber (1995), as an oligopoly game, and Hansen (1988), as a procurement auction with variable demand—which can also be viewed through the conceptual lens of our model. We show in Section 5.2 that the equilibrium of our model converges to Hansen and Spulber’s Bertrand equilibrium as inventory costs become fully recoverable.

Our equilibrium characterization thus provides a simple way to extend the analysis of production-to-order models (i.e., Bertrand models) to environments where inventories must be chosen in advance and costs are partially sunk, and does so under both complete and incomplete information.

An important difference relative to the literatures discussed above is that our main focus is not on the constrained game, but rather on the unconstrained one, where firms freely choose

⁴Kaplan, Luski, Sela, and Wettstein (2002) and Cohen, Kaplan, and Sela (2008) study related games, but do not prove existence and uniqueness. Incomplete-information contests with *monotonic* winning functions have been studied by, e.g., Weber (1985), Amann and Leininger (1996), and Moldovanu and Sela (2001, 2006).

inventories. Our results in Sections 3.2 and 5.1 that the equilibrium of the constrained game is the generically unique equilibrium of the unconstrained price-inventory game is novel, and involves considerations that are naturally absent in the all-pay contests literature and in the industrial organization literature on production to order.

A small number of papers has analyzed price-inventory models similar to ours. Maskin (1986) proves equilibrium existence in a class of price-inventory games with two firms. The non-generic example we present in Section 2 was previously studied by Levitan and Shubik (1978) with linear demand and Gertner (1986) under the assumption that inventory costs are completely sunk. Unfortunately, the proofs of equilibrium uniqueness they provide omit important non-trivial steps and/or contain several inaccuracies.⁵ For completeness, we provide a proof that addresses those shortcomings in Online Appendix I.

More importantly, the example of Section 2 is non-generic. Perturbing a game that is not generic leads to a generic game. Moreover, it is known that the equilibrium behavior in non-generic contests can be very different from that in generic ones (see, e.g., Section 3.2 in Siegel, 2009). These observations, and the fact that numerous applications in industrial organization do not become interesting unless firm asymmetries and heterogeneous consumer preferences for stores are allowed, also motivate our study of the rich class of all-pay oligopolies.

In a one-shot game, capacity and inventory choices are formally equivalent. A strand of literature studies oligopoly settings where firms first choose *observable* capacities, and then compete in prices. The leading reference is Kreps and Scheinkman (1983), discussed in detail above. A large literature (e.g., Davidson and Deneckere, 1986; Deneckere and Kovenock, 1996) explores the robustness of the Kreps-Scheinkman result. The general message conveyed in those papers is that production in advance provides a commitment to soften price competition. Yet, a counterfactual with unobservable inventories had so far not been properly investigated. Our work shows that, to soften price competition, production in advance must be combined with inventory observability. Indeed, a Bertrand-like, intense form of competition arises under production in advance if inventories remain unobservable.

Finally, Deneckere and Peck (1995) study a symmetric model with simultaneous price-inventory choices but stochastic demand. In the single pure-strategy equilibrium candidate, all stores choose the same price (i.e., there is no price dispersion), consumers are rationed when demand is high, and stores are unable to sell all their inventories when demand is low. In a pure-strategy equilibrium candidate, unsold inventories must however disappear as demand becomes certain (whereas in our model, some units remain unsold despite the absence of market demand uncertainty). More fundamentally, pure-strategy equilibria cease to exist when, close to our setting, demand uncertainty or the number of stores is low. Variations of this problem were subsequently studied by operations scholars as a branch of the large

⁵Levitan and Shubik (1978) confine attention to equilibria without mass points (except potentially at the choke price of demand). Gertner (1986)'s approach is generally correct. However, his proof, which was never peer-reviewed, contains a number of measure-theoretical inconsistencies, which our proof addresses. Tasnadi (2004, Section 4) assumes throughout that "any price above monopoly price is dominated," which, as we and the authors above show, is incorrect.

newsvendor literature (e.g., Bernstein and Federgruen, 2004, 2007; Zhao and Atkins, 2008).

2 A Non-Generic Example

Two symmetric firms, 1 and 2, sell a homogeneous product. The demand function D is non-increasing, strictly positive and continuous on $[0, p^0)$, and zero on $[p^0, \infty)$, with $p^0 \in (0, \infty)$. Each firm incurs a cost $c \in (0, p^0)$ for each unit it sells. For each unit that remains unsold, a fraction $\alpha \in [0, 1)$ of the unit cost c is recovered. Thus, $1 - \alpha$ captures the extent to which the inventory cost is sunk. If firm i chooses a price-inventory pair (p_i, q_i) and sells $s_i \leq q_i$ units, its profit is

$$(p_i - \alpha c)s_i - (1 - \alpha)cq_i.$$

The number of units sold, s_i , is determined by the vector of prices and inventories, the demand function, and the rationing rule. If firm i sets a price below its rival's ($p_i < p_j$), then it faces the entire market demand $D(p_i)$, and sells $s_i = \min(q_i, D(p_i))$. If instead $p_i > p_j$, then firm i only receives demand from those consumers who were unable to buy from firm j when $q_j < D(p_j)$. That residual demand is pinned down by the rationing rule, which we assume to be either efficient or random. Finally, if $p_i = p_j$, demand is allocated on a same-price fair-share basis.⁶

Firms compete by simultaneously choosing a price-inventory pair $(p_i, q_i) \in \mathbb{R}_+^2$ ($i = 1, 2$). This is formally equivalent to firms first choosing inventories, which remain unobservable, and then simultaneously setting prices. We refer to this game as the *unconstrained game*. We look for the set of Nash equilibria.

The Bertrand outcome ($p_1 = p_2 = c$, $q_1 = q_2 = D(c)/2$) is not an equilibrium: Firm i could profitably raise its price to $c + \varepsilon$ and supply the (strictly positive) residual demand at that price. The Cournot outcome is not an equilibrium either: Firm i could profitably undercut the Cournot price and supply the entire market demand. The same line of reasoning shows that any Nash equilibrium must involve non-degenerate mixing.

To introduce the approach later used to study more general all-pay oligopoly games, we first solve a constrained version of the model. In that *constrained game*, each firm i can only choose pairs (p_i, q_i) such that $q_i = D(p_i)$. We look for a Nash equilibrium in which both firms make zero profit—the Bertrand profits. Suppose firm j draws its price from a continuous (up to p^0) cumulative distribution function (CDF) F_j . If firm i sets a price p_i , it sells all its inventory with probability $1 - F_j(p_i)$, and nothing otherwise. From the zero profit condition,

$$\left((p_i - \alpha c)(1 - F_j(p_i)) - (1 - \alpha)c \right) D(p_i) = 0,$$

⁶Same-price fair-share ($s_i = \min(q_i, \max(D(p_i)/2, D(p_i) - q_j))$) is a standard assumption. The rationing and sharing rules are defined formally in Appendix B.1. See also Davidson and Deneckere (1986) for a discussion of rationing rules in pricing games.

we obtain

$$F^*(p) = \begin{cases} \frac{p-c}{p-\alpha c} & \text{if } p \in [c, p^0], \\ 1 & \text{if } p \geq p^0. \end{cases}$$

We thus obtain a Nash equilibrium where each firm i draws its price from the CDF F^* . Intuitively, firm i earns a higher margin on each unit it sells when it raises its price, but the probability that it ends up with costly unsold units increases as well. F^* exactly offsets these two effects. Note that F^* has a mass point at p^0 , i.e., on being inactive.

The above strategy profile is also a Nash equilibrium of the unconstrained game. Suppose firm j randomizes according to the proposed equilibrium strategy. If firm i chooses a price p_i and any inventory $q_i \in [0, D(p_i)]$, then it receives either the entire market demand, $D(p_i)$, or no residual demand at all. Hence, by the definition of F^* , firm i still makes zero profit.

The following proposition shows that there are no other equilibria:⁷

Proposition 1. *The constrained game and the unconstrained game have a unique Nash equilibrium, characterized above.*

Proof. See Online Appendix I.⁸ □

The equilibrium of Proposition 1 has the following properties: (i) it is unique; (ii) there is price dispersion; (iii) each firm always sources enough inventory to supply all demand at its chosen price; (iv) the equilibrium is non-strict in a very strong sense, as both firms are indifferent between all the price-inventory pairs that lie below the graph of D ; (v) some inventories remain unsold with positive probability; (vi) the distribution of prices is independent of the shape of D ; (vii) the market is not served with positive probability; (viii) firms earn Bertrand profits; and (ix) as inventory costs become fully recoverable, the equilibrium CDF of prices converges weakly to an atom at $p = c$, i.e., to the Bertrand outcome.

Note, however, that the model studied in this section is non-generic. Perturbations of this non-generic game will typically result in one or both of the firms making strictly positive profits, which breaks the equilibrium construction presented above. This is the case, e.g., if firms' costs are asymmetric, or firms' products are perceived to be differentiated by a subset of consumers. As such elements are natural and important features of oligopolistic competition, this raises several important questions. What is the equilibrium set of generic oligopoly games where firms choose prices and inventories simultaneously? What are the equilibrium properties of such games and, more specifically, which of the properties mentioned above are

⁷Uniqueness should be understood as being up to payoff-equivalent strategies, as all strategies with zero inventory give rise to the same payoffs.

⁸As we later prove a more general result (Theorem 1), the reader may wonder why the proof of Proposition 1 requires such a long development. Because the model studied in this section is non-generic, Proposition 4 does not apply, and so we need to use completely different techniques. Our approach for Proposition 1 essentially follows the one in Gertner (1986), which crucially builds on both firms making zero profit. The case of efficient rationing (Online Appendix I.6), which Gertner (1986) does not study, is significantly more involved than the random rationing one (Online Appendix I.5). Interestingly, the rationing rule plays no role in the proof of Theorem 1, so those complications become irrelevant in the generic case.

indeed generic? In this paper, we answer these and related questions by introducing a rich class of (generic) games, which we call all-pay oligopolies.

3 All-Pay Oligopolies

We now present a general model of all-pay oligopolies. There are two firms, 1 and 2.⁹ Firm i has constant unit cost $c_i > 0$, and can recover a fraction $\alpha_i \in [0, 1)$ of this cost on all unsold units. Total market demand is still given by the function D , assumed to be strictly positive, continuous, non-increasing, and log-concave on $[0, p^0)$, left-continuous at $p^0 \in (0, \infty)$, and identically equal to zero on (p^0, ∞) .¹⁰

In addition to cost asymmetries, consumers may also perceive the two firms as being differentiated, so that not all consumers necessarily wish to purchase from the firm setting the lowest price. We assume that a fraction of consumers $\mu_i \in [0, 1)$ only wish to buy from firm i , as in Varian (1980). Those consumers are referred to as firm i 's *captive segment*. The remaining fraction $1 - \mu_1 - \mu_2 \in (0, 1]$ of consumers, the *shoppers*, wish to buy from the firm with the lowest price. As a whole, we call those consumers the *contested segment*. Firm i can pay a fixed advertising cost of $A_i \geq 0$, as in, e.g., Baye and Morgan (2001), to make the shoppers aware of its product and price. This cost is akin to a fixed cost of production if there are no captive consumers, i.e., if $\mu_1 = \mu_2 = 0$.

Firms 1 and 2 simultaneously decide whether to pay the advertising cost, which price to set, and how many units to source. As in Section 2, we assume that rationing within a consumer segment is either random or efficient, and that demand is allocated in a same-price fair-share way in case of a tie.¹¹ We also need to specify how rationing works across segments: To fix ideas, we assume that each firm serves its captive consumers first. This last assumption does not affect the equilibrium characterization in generic cases.

We first study a constrained version of the model, where firms must source enough inventory to supply their targeted demand, in Section 3.1. We then study the unconstrained game, where firms can freely choose their inventories, in Section 3.2. In Section 3.3, we discuss the generic properties of the equilibrium. In Section 3.4, we show that equilibrium behavior converges to an equilibrium of the associated Bertrand model as production costs become fully recoverable.

3.1 The Constrained Game

As in Section 2, we first solve a constrained version of the model in which both firms must source enough inventories to supply their targeted demand. In that constrained game, firm i decides whether to target only its captive segment, or both its captive and the contested segment. In the former case, it does not to pay its advertising cost A_i and, conditional on

⁹We discuss the N -firm case in Section 7 and study it formally in Online Appendix IV.

¹⁰The potential discontinuity at p^0 allows, for example, to nest perfectly inelastic demand as a special case.

¹¹Ties play a limited role in the analysis. Rationing and sharing rules are defined formally in Appendix B.1.

choosing price p_i , it must source an inventory $\mu_i D(p_i)$. It thus finds it optimal to set its monopoly price p_i^m and receives its *outside option* payoff:¹²

$$o_i = \mu_i(p_i^m - c_i)D(p_i^m).$$

If instead firm i targets both its captive and the contested segment at price p_i , then it pays its advertising cost and sources an inventory $(1 - \mu_j)D(p_i)$. If firm j is not targeting the contested segment or $p_i < p_j$, firm i 's payoff is:

$$w_i(p_i) = (1 - \mu_j)(p_i - c_i)D(p_i) - A_i.$$

We call $w_i(\cdot)$ firm i 's *winning function*. If firm j is targeting the contested segment at a price $p_j < p_i$, then firm i receives its *losing payoff*:

$$l_i(p_i) = \left(\mu_i(p_i - c_i) - (1 - \mu_i - \mu_j)(1 - \alpha_i)c_i \right) D(p_i) - A_i.$$

Note that for every $p_i \in [c_i, p^0)$, winning is better than losing, i.e., $w_i(p_i) > l_i(p_i)$, and losing is worse than taking the outside option, i.e., $l_i(p_i) < o_i$. Payoffs in case of a tie play a limited role in the analysis and are therefore omitted here. The tuple $(w_i, l_i, o_i)_{i=1,2}$ defines an all-pay contest with outside options and, due to revenue log-concavity, potentially non-monotonic winning and losing functions.

We now proceed with the equilibrium characterization. If, for some firm i , o_i strictly exceeds $w_i(p_i^m)$, the highest profit firm i can receive if it targets its captive and the contested segment, then there is no scope for competition in the contested segment. In that case, the game is dominance solvable: Firm i focuses on its captive consumers with probability 1, and firm j plays a best response to that action.

Suppose that $w_i(p_i^m) > o_i$. We can then define firm i 's *reach* as

$$r_i = \min\{p \in [c_i, p^0] : w_i(p) = o_i\}.$$

By uni-modality of w_i , any price below r_i is a strictly dominated strategy for firm i . Let $r = \max\{r_1, r_2\}$ be the highest reach. We say that firm i is *strong* if $r_i < r$ and *weak* if $r_i = r$. As the weak firm will never price below r , this is the highest price at which the strong firm can be sure to capture the contested segment—the limit price. Naturally, the strong firm can price more aggressively than the weak one while still earning more than its outside option. If $p_i^m < r$, the game is also dominance solvable: Firm i targets its captive and the contested segment at its monopoly price, whereas firm j always focuses on its captive consumers and sets its own monopoly price.

¹²In the special case where $\mu_i = 0$, firm i can set any price when it takes its outside option since it produces nothing. It is of course without loss of generality to assume that firm i sets p_i^m .

In the remainder of the paper, we assume that $w_i(p_i^m) > o_i$ and $p_i^m > r$ for every i . We call games satisfying those conditions *non-trivial*. Standard arguments (see Section 2) imply that these games do not have a pure-strategy Nash equilibrium. We therefore look for mixed-strategy equilibria. An equilibrium is fully characterized by a pair of CDF's (F_1, F_2) over $[r, p^0]$. $F_i(p)$ is the probability that firm i targets its captive and the contested segment at a price less than or equal to p . The probability that firm i takes its outside option is therefore given by $1 - F_i(p^0)$.¹³

We look for an equilibrium in which firms mix continuously over some interval $[r, \bar{p})$ and distribute the remaining mass on higher prices and/or their outside option. Since r is the infimum of the support of both firms' strategies, the expected payoff of firm i in this putative equilibrium is given by $w_i(r)$. For firm i to be indifferent between all the prices in $[r, \bar{p})$, it has to be the case that for every $p \in [r, \bar{p})$,

$$(1 - F_j(p))w_i(p) + F_j(p)l_i(p) = w_i(r),$$

i.e., $F_j(p) = k_j(p)$, where

$$k_j(p) \equiv \frac{w_i(p) - w_i(r)}{w_i(p) - l_i(p)}, \quad \forall p \in [r, p^0).$$

The log-concavity of D implies that k_j is either single-peaked and achieves a global maximum at some $\bar{p}_j \in (r, p^0)$, or strictly increasing, in which case we set $\bar{p}_j = p^0$ (see Lemma C–(i) in the Appendix). If \bar{p}_j were strictly less than \bar{p} , then firm j 's CDF of prices would be strictly decreasing on (\bar{p}_j, \bar{p}) , which it cannot be. We therefore set $\bar{p} = \min\{\bar{p}_1, \bar{p}_2\}$. Then, k_j is continuous and strictly increasing on $[r, \bar{p})$. Moreover, the fact that $w_i(r) \geq o_i > l_i(p)$ for every $p \in (r, \bar{p})$ implies that $k_j(p) \in (0, 1)$. Hence, k_j has the properties of a CDF on the interval $[r, \bar{p})$. However, $\lim_{p \uparrow \bar{p}} k_j(\bar{p}) < 1$ (see Lemma C–(ii) in the Appendix), meaning that we have some mass left to distribute. How that mass is distributed depends crucially on whether \bar{p}_1 is higher or lower than \bar{p}_2 , as discussed next.

We focus on the generic case in which $r_i \neq r_j$. (We will later study the non-generic case $r_i = r_j$.) To fix ideas, suppose that $r_1 < r_2 (= r)$.

Suppose first that $\bar{p}_1 \geq \bar{p}_2$, so that $\bar{p} = \bar{p}_2$. The fact that k_2 is strictly decreasing on (\bar{p}, p^0) means that, given that firm 2 is already putting a total mass of $k_2(\bar{p})$ on the interval $[r, \bar{p}_2]$, firm 1 does not want to price anywhere in the interval $(\bar{p}, p^0]$. Moreover, since firm 1 is strong ($w_1(r) > o_1$), it does not want to take its outside option either. The only possibility is therefore that firm 1 puts the rest of its mass on \bar{p} . Firm 2 responds by putting the rest of its mass on its outside option. To summarize, each firm has a single mass point (the strong

¹³With this formulation, in the special case where $\mu_i = A_i = D(p^0) = 0$, firm i has two equivalent ways of being inactive: It can either pay the advertising cost (which, here, is equal to zero) and set $p_i = p^0$ and $q_i = (1 - \mu_j)D(p^0) = 0$, or not pay the advertising cost and set $p_i = p_i^m$ and $q_i = \mu_i D(p_i^m) = 0$. In that case, we assume without loss of generality that firm i chooses the latter way of being inactive, which ensures that $1 - F_i(p^0)$ is indeed the probability that firm i takes its outside option.

firm at \bar{p} , the weak firm on its outside option) and the CDF's are:

$$F_1(p) = \begin{cases} k_1(p) & \text{if } p \in [r, \bar{p}_2), \\ 1 & \text{if } p \in [\bar{p}_2, p^0], \end{cases} \quad \text{and} \quad F_2(p) = \begin{cases} k_2(p) & \text{if } p \in [r, \bar{p}_2), \\ k_2(\bar{p}_2) & \text{if } p \in [\bar{p}_2, p^0]. \end{cases} \quad (1)$$

It is readily verified that this pair of CDF's is a Nash equilibrium of the constrained game.

Next, suppose $\bar{p}_1 < \bar{p}_2$. Then, it is the weak firm that does not want to price anywhere in the interval $(\bar{p}, p^0]$. Hence, F_2 is constant on that interval. If $F_2(\bar{p}) < k_2(\bar{p}_2)$, then firm 1 can obtain strictly more than $w_1(r)$ by pricing at \bar{p}_2 (which cannot be). If instead $F_2(\bar{p}) > k_2(\bar{p}_2)$, then firm 1 does not want to price anywhere in $[\bar{p}, p^0]$ and must then take its outside option $o_1 < w_1(r)$ with positive probability (which also cannot be). It follows that $F_2(\bar{p}) = k_2(\bar{p}_2)$ and firm 2 puts its remaining mass on its outside option. Firm 1 responds by putting the rest of its mass on \bar{p}_2 . To summarize, firm 1 has a single mass point (at \bar{p}_2), firm 2 has two mass points (one at $\bar{p} < \bar{p}_2$ and the other one on its outside option), and CDF's are:

$$F_1(p) = \begin{cases} k_1(p) & \text{if } p \in [r, \bar{p}_1), \\ k_1(\bar{p}_1) & \text{if } p \in [\bar{p}_1, \bar{p}_2), \\ 1 & \text{if } p \in [\bar{p}_2, p^0], \end{cases} \quad \text{and} \quad F_2(p) = \begin{cases} k_2(p) & \text{if } p \in [r, \bar{p}_1), \\ k_2(\bar{p}_2) & \text{if } p \in [\bar{p}_1, p^0]. \end{cases} \quad (2)$$

It is readily verified that (F_1, F_2) is a Nash equilibrium.

In words, in both cases, firm 1 always targets both segments. With a strictly positive probability, it sources a low inventory and charges its reference price \bar{p}_2 . With complementary probability, it sources a high inventory and offers a discount, drawing its price from a continuous distribution over $[r, \bar{p}_2]$. Firm 2, with strictly positive probability, focuses exclusively on its captive segment at its monopoly price p_2^m . With complementary probability, it sources a high inventory to target both its captive and the contested segment. In that case, it draws its price from the segment $[r, \bar{p}]$, continuously if $\bar{p}_1 \geq \bar{p}_2$, and with a mass point at \bar{p}_1 if $\bar{p}_1 < \bar{p}_2$. Thus, in one case, firm 2 has a unique reference price (p_2^m), whereas it has two reference prices (p_2^m and \bar{p}_1) in the other case. (See Figure 1 in Section 4 for a graphical illustration of equilibrium behavior.)

In both cases, equilibrium uniqueness can be established using standard techniques:

Proposition 2. *Consider a non-trivial and generic ($r_1 \neq r_2$) all-pay oligopoly model, where firm 1 is the strong firm. The constrained game has a unique equilibrium. The equilibrium profile of CDF's of prices in the contested segment is described by equation (1) if $\bar{p}_2 \leq \bar{p}_1$, and by equation (2) otherwise. The strong firm targets the contested segment for sure, whereas the weak firm focuses exclusively on its captive consumers with probability $1 - F_2(p^0) > 0$. Equilibrium payoffs are $w_i(r)$ for $i = 1, 2$.*

Proof. See Appendix A. □

In contrast, the constrained game of a non-generic ($r_1 = r_2$) all-pay oligopoly model usually has a continuum of equilibria. The intuition for the equilibrium multiplicity is that in

the non-generic case, we have more leeway to allocate the mass that firms do not put on $[r, \bar{p}]$. Proposition A, stated and proven in the Appendix, provides a complete characterization of the set of equilibria, which we now describe informally. In any equilibrium, each firm earns an expected profit equal to its outside option; firms mix continuously on the interval $[r, \bar{p}]$ according to the CDF's k_1 and k_2 , and put the rest of their mass on \bar{p}_1 , \bar{p}_2 , and/or their outside option; at least one firm puts mass on its outside option.

3.2 The Unconstrained Game

We now study the unconstrained game, where firms can freely choose their inventories. We show that any equilibrium of the constrained game is also an equilibrium of the unconstrained game. Moreover, generically, the unique equilibrium of the constrained game is also the unique equilibrium of the unconstrained game. Finally, we briefly discuss the equilibrium characterization in non-generic cases.

Let (F_1, F_2) be an equilibrium of the constrained game. Suppose firm i pays the advertising cost and deviates to a price-inventory pair (p, q) such that it does not source enough inventory to supply its targeted demand, i.e., $q \in [\mu_i D(p), (1 - \mu_j)D(p)]$. If firm j does not have a mass point at p , then firm i earns a profit of

$$\tilde{\pi}_i(p, q) = \mu_i(p - c_i)D(p) - A_i + \left((p - \alpha_i c_i)(1 - F_j(p)) - (1 - \alpha_i)c_i \right) (q - \mu_i D(p)).^{14}$$

As $\tilde{\pi}_i(p, q)$ is linear in q , the optimal deviation given p is a corner solution, i.e., $q = \mu_i D(p)$ or $q = (1 - \mu_j)D(p)$. Since both corner solutions are permitted in the constrained game, and since (F_1, F_2) is an equilibrium of that game, the deviation is not profitable. We obtain:

Proposition 3. *In an all-pay oligopoly model, any equilibrium of the constrained game is also an equilibrium of the unconstrained game.*

In fact, provided that $A_i > 0$, a condition that holds generically, we have that for every price p , $\mu_i(p - c_i)D(p) - A_i < o_i$. The argument used to prove Proposition 3 then implies that $(p - \alpha_i c_i)(1 - F_j(p)) > (1 - \alpha_i)c_i$ for every price p in the support of F_i . Therefore, firm i strictly prefers sourcing enough inventory to supply the demand it targets. The key here is that firm i 's deviation profit $\tilde{\pi}_i$ is linear in q , with a strictly positive slope whenever $A_i > 0$. Further exploiting this argument, we obtain the following proposition:¹⁵

Proposition 4. *In a generic ($A_1, A_2 > 0$) all-pay oligopoly model, a strategy profile is an equilibrium of the constrained game if and only if it is an equilibrium of the unconstrained game.*

Proof. See Appendix B. □

¹⁴If firm j has a mass point at p , then either q is sufficiently low and the expression for $\tilde{\pi}_i$ is valid, or q is not sufficiently low and firm i would be strictly better off pricing just below p and avoiding the mass point.

¹⁵The proposition applies to both trivial and non-trivial games.

An immediate implication of Propositions 2 and 4 is that a generic $(A_1, A_2 > 0, r_1 \neq r_2)$ all-pay oligopoly game has a unique equilibrium.¹⁶

Proposition 4 is proven as follows. We assume for a contradiction that the unconstrained game has an equilibrium in which one (or both) of the firms does not always source enough inventory to supply its targeted demand. We define \hat{p} as the supremum of the set of prices below which both firms source enough inventory to supply their targeted demand, and show that this supremum is a maximum. Suppose firm i chooses a price-inventory pair (p, q) in the support of its equilibrium price-inventory distribution, with $p > \hat{p}$ and $\mu_i D(p) \leq q < (1 - \mu_j)D(p)$, and let $F_j(p)$ denote the probability that firm j chooses a pair (p_j, q_j) such that $p_j < p$ and $q_j > \mu_j D(p_j)$. Then, firm i makes an expected profit of

$$\tilde{\pi}_i(p, q) = \mu_i(p - c_i)D(p) - A_i + \left((p - \alpha_i c_i)(1 - F_j(p)) - (1 - \alpha_i)c_i \right) (q - \mu_i D(p)) + \varepsilon(p, q),$$

where we have ignored the possibility that firm j has a mass point at p to ease the exposition. The term $\varepsilon(p, q)$, which is non-negative and non-decreasing in q , captures the fact that firm i may still end up selling to some of the shoppers if firm j prices between the cutoff price \hat{p} and firm i 's price p . As p decreases to \hat{p} , the probability that firm j sets a price in (\hat{p}, p) converges to zero, and $\varepsilon(p, q)$ therefore tends to zero. Since $\tilde{\pi}_i(p, q) \geq o_i > \max_{p'} \mu_i(p' - c_i)D(p') - A_i$, this implies that $(p - \alpha_i c_i)(1 - F_j(p)) > (1 - \alpha_i)c_i$ for p sufficiently close to \hat{p} . It follows that $\tilde{\pi}_i(p, q)$ is strictly increasing in q for p sufficiently close to \hat{p} , a contradiction.

Combining Propositions 2–4 and Proposition A in the Appendix gives:

Theorem 1. *A non-trivial and generic $(A_1, A_2 > 0, r_1 \neq r_2)$ all-pay oligopoly game has a unique equilibrium. In that equilibrium, both firms source enough inventories to supply their targeted demand, and the equilibrium CDF's of prices are as characterized in Proposition 2. A non-trivial and non-generic all-pay oligopoly game may have multiple equilibria. Constrained equilibria, which are also unconstrained equilibria, are as characterized in Proposition A.*

We close this subsection by discussing the equilibrium multiplicity that can arise in *non-generic* cases. We do so in the context of a simple example with inelastic unit demand up to p^0 , symmetric firms, captive consumers, and no advertising cost. This model boils down to a production-in-advance version of Varian (1980)'s model of sales.

We first discuss constrained equilibria. By Proposition A in the Appendix, in any such equilibrium, both firms mix continuously between r and $\bar{p} = p^0$ according to the CDF F . Given that firm j puts mass $F(p^0)$ on $[r, p^0]$, firm i is indifferent between targeting both its captive and the contested segment at p^0 , and taking its outside option. This indifference gives rise to a continuum of equilibria in which firm i splits its remaining mass between its outside option and p^0 , whereas firm j puts all of its remaining mass on its outside option.

The unconstrained game also has equilibria that are not constrained equilibria. The proof of Proposition 4 can be adapted to show that in any equilibrium, conditional on pricing at

¹⁶Note that the genericity condition in Proposition 2 differs from that in Proposition 4. The intersection of those conditions ($r_1 \neq r_2$ and $A_1, A_2 > 0$) holds generically as well.

$p < p^0$ in the contested segment, firm i sources enough inventory so supply its targeted demand, i.e., $1 - \mu$ units. This implies that both firms still mix continuously over prices in $[r, p^0)$ with an inventory level of $1 - \mu$. Given that firm j puts mass $F(p^0)$ on $[r, p^0]$, firm i is therefore still indifferent between setting p^0 in the contested segment and taking its outside option. The fact that at $p = p^0$, firm i 's expected profit is linear in $q \in [\mu, 1 - \mu]$ implies that firm i is in fact indifferent between all the inventory levels in $[\mu, 1 - \mu]$. We therefore obtain a continuum of equilibria in which conditional on pricing at p^0 , each firm i draws its inventory from some probability measure λ_i over $[\mu, \bar{q}_i]$, with $\bar{q}_i \leq 1 - \mu$ and $\bar{q}_1 + \bar{q}_2 \leq 1$.

Recall however that the equilibrium multiplicity characterized above is non-generic: By Theorem 1, that multiplicity disappears when the game is slightly perturbed.

3.3 Generic Equilibrium Properties

We now discuss the qualitative features of the equilibrium of a generic all-pay oligopoly model, and compare them to what we found for the non-generic example of Section 2. Our analysis in Sections 3.1 and 3.2 reveals that in generic games, (I) the equilibrium is unique; (II) there is price dispersion; (III) firms source enough inventory to supply their targeted demand; (IV) firms *strictly prefer* sourcing enough inventory to supply their targeted demand; (V) some inventories remain unsold with positive probability; (VI) the distribution of prices depends on the shape of D ; (VII) all segments are served with probability 1 (as the strong firm always targets both its captive and the contested segment, and the weak firm always targets at least its captive consumers).

Properties (I), (II), (III), and (V) were also satisfied in the non-generic example of Section 2. Properties (IV), (VI), and (VII), however, were not: In that model, firms are indifferent between all the pure strategies that are not strictly dominated; the distribution of prices does not depend on the shape of D ; the market is not supplied with strictly positive probability. Our analysis thus reveals that those latter properties are non-generic: They disappear if the non-generic example of Section 2 is slightly perturbed.

Nevertheless, for small perturbations, the equilibrium of such perturbed games is typically quantitatively similar to that of the unperturbed game:¹⁷

Proposition 5. *Suppose D is continuous at p^0 , and let $(\gamma^n)_{n \geq 0}$ be a sequence of parameter vectors that converges to the non-generic parameter vector of Section 2. For every n , let (F_1^n, F_2^n) be a (constrained) equilibrium of the game with parameter vector γ^n . Then, $(F_1^n, F_2^n)_{n \geq 0}$ converges weakly to the mixed-strategy equilibrium of Proposition 1.*

Proof. See Appendix C.3. □

¹⁷ F_i^n was defined as the CDF of a measure over $[0, p^0]$. Since $F_i^n(p^0)$ may be strictly less than 1, F_i^n is not necessarily a *probability* measure, so the weak convergence of the sequence $(F_i^n)_{n \geq 0}$ may not be a well-defined concept. We circumvent this issue by studying an equivalent auxiliary game in which any mixed-strategy equilibrium can be described by a pair of *probability* measures over $[0, p^0]$. We then establish the weak convergence of the associated sequence of pairs of probability measures. See Appendix C.1 for details.

This result may seem surprising since in a generic all-pay oligopoly game, the strong firm always targets the contested segment, whereas, in the equilibrium of Section 2, both firms stay out of the market with strictly positive probability. The intuition is the following. In a generic game, the strong firm, firm 1, has a mass point at \bar{p}_2 . When parameters are close to those in Section 2, \bar{p}_2 is close to the choke price p^0 and, since D is continuous at p^0 , the strong firm is thus close to being inactive (i.e., it chooses an inventory close to zero).

We have not yet discussed the Bertrand-related properties—properties (viii) and (ix) of Section 2. This is a more involved exercise, which we undertake in the next subsection.

3.4 Bertrand Convergence

Under Bertrand competition, it is assumed that firms first choose prices, and must then satisfy all the demand directed to them. The constrained game studied in Section 3.1 is formally equivalent to a Bertrand game when costs are fully recoverable ($\alpha_1 = \alpha_2 = 1$), as there is then no downside to being left with unsold inventories.

We find that as α_1 and α_2 tend to 1, the equilibrium distribution of prices converges to an equilibrium of that Bertrand game. This verifies that property (ix) of Section 2 is generic:

Proposition 6. *Suppose D is continuous at p^0 , and let $(\gamma^n)_{n \geq 0}$ be a sequence of parameter vectors that converges to a parameter vector such that $\alpha_1 = \alpha_2 = 1$. For every n , let (F_1^n, F_2^n) be a (constrained) equilibrium of the game with parameter vector γ^n . Generically, $(F_1^n, F_2^n)_{n \geq 0}$ converges weakly to an equilibrium of the resulting Bertrand game.*

Proof. See Appendix C.4. □

The property that, regardless of the recoverability parameters, both firms earn their Bertrand profits (property (viii) of Section 2) is also generic, as we explain next. In a given all-pay oligopoly game, neither the winning payoff w_i nor the outside option o_i depends on the recoverability parameters. Hence, the reach of firm i , r_i , and the equilibrium profit of firm i , $w_i(r)$, remain the same regardless of (α_1, α_2) . By Proposition 6, as recoverability parameters tend to 1, equilibrium play tends to an equilibrium of the limiting Bertrand game. Hence, in the limit, the equilibrium profit of firm i is $w_i(r)$ —its Bertrand profit.

The following properties are therefore generic as well: (VIII) firms earn Bertrand profits; (IX) the equilibrium is Bertrand convergent.

As mentioned in the introduction, Proposition 6 is in contrast to what is known from the literature that studies the Kreps and Scheinkman (1983) framework (in which $\mu_i = A_i = 0$). Assuming *observable* inventory choices, that literature finds that if the sunk part of the unit cost is sufficiently high, then the equilibrium outcome is (close to) Cournot. Only when inventory costs are sufficiently recoverable can other outcomes arise.¹⁸ Those outcomes may

¹⁸Although closed-form solutions are unavailable for such cases, it is known that under efficient rationing, linear demand and costs, and unit cost asymmetry, the equilibrium may involve non-degenerate mixing, and

be more competitive than Cournot, yet they remain far from the intense competition arising in the Bertrand model.

One main takeaway from that literature is that observable inventory choices act as a commitment to soften price competition and protect margins (e.g., Tirole, 1987, p. 218). So far, a counterfactual with unobservable inventory choices had not been properly investigated. This subsection verifies that a Bertrand-like form of competition arises in the absence of the commitments provided by such observability: In terms of profits, this holds regardless of unit cost recoverability; in terms of prices, this holds if unit costs are sufficiently recoverable.

The Kreps-Scheinkman literature also suggests that mixed strategies, which result in unsold inventories and price dispersion, are not an unusual feature of oligopolies with production in advance—a feature that, as discussed in the introduction, seems to have a real-world counterpart. Due to the lack of explicit equilibrium characterizations, an in-depth exploration of the implications of this feature has understandably been all but inexistent. Our simple characterization provides an avenue for such exploration.

4 Applications

We now study specific all-pay oligopoly games and relate them to well-known Bertrand games.

Bertrand with fixed costs. In the Bertrand models ($\alpha_1 = \alpha_2 = 1$) studied by Sharkey and Sibley (1993), Marquez (1997), and Thomas (2002), firms have identical unit costs ($c_i = c$) and there are no captive consumers ($\mu_i = 0$). In the generic ($A_1 < A_2$) equilibrium, firms mix continuously over $[r, p^m)$, the strong firm (firm 1) puts the rest of its mass on p^m , and firm 2 puts the rest of its mass on inactivity (i.e., on staying out of the market).

Under production in advance ($\alpha_1, \alpha_2 < 1$), firm i mixes over $[r, \bar{p})$ according to

$$k_i(p) = \frac{(p - c)D(p) - A_2}{(p - \alpha_j c)D(p)}.$$

Note that $\bar{p} > p^m$, a property we will return to later on.¹⁹ How firm i distributes the rest of its mass depends on whether α_1 is larger or smaller than α_2 .

If $\alpha_1 \geq \alpha_2$, firm 1 is stronger in both dimensions, in the sense of having lower and more recoverable costs. This results in firm 1 being less willing to set high prices than firm 2 ($\bar{p}_1 \geq \bar{p}_2$). The equilibrium, represented graphically in the top panel of Figure 1, is thus described by equation (1), with the strong firm putting the rest of its mass on \bar{p} and the

that a similar result obtains under symmetric costs and *random* rationing—see, respectively, Deneckere and Kovenock (1996) and Davidson and Deneckere (1986). Yet, if the sunk part of the unit cost of both firms exceeds just 7.5% of the choke price, then the outcome is Cournot in the former case, and Cournot or close to it in the latter one.

¹⁹Here, we implicitly assume that $p^m < p^0$, which holds for instance if D is continuous at p^0 . To see why $\bar{p} > p^m$, note that starting from $p = p^m$, a small increase in p has no first-order impact on the numerator of k_i , but reduces the denominator. Hence, k_i is strictly increasing at p^m .

weak firm putting the rest of its mass on inactivity. The firms' mass points are therefore qualitatively similar to those under production to order, with the exception that $\bar{p} > p^m$.

If instead $\alpha_1 < \alpha_2$, firm 1 has a fixed cost advantage but a recoverability disadvantage. This results in firm 1 still being the strong firm ($r_1 < r_2 = r$), but being more willing to set high prices than firm 2 ($\bar{p}_1 < \bar{p}_2$). The equilibrium, represented graphically in the bottom panel of Figure 1, is thus described by equation (2), with the strong firm putting the rest of its mass on \bar{p}_2 and the weak firm having two mass points: one on $\bar{p} < \bar{p}_2$ and the other one on inactivity. This equilibrium differs significantly from the production-to-order one.

Thus, the extent to which inventories are recoverable affects equilibrium behavior in two ways. First, an increase in the α 's gives rise to a first-order stochastic dominance shift towards lower prices, i.e., firms price more aggressively. Second, as seen above, the firms' ranking in terms of unit cost recoverability determines the qualitative properties of the equilibrium.

Interestingly, in the special case where $\alpha_1 = \alpha_2$, which includes the limiting Bertrand case ($\alpha_i = 1$), we have that $k_1 = k_2$. Thus, despite having different fixed costs, firms 1 and 2 use the same pricing strategy up to \bar{p} —a feature that seems to have been overlooked in the production-to-order literature.

Regardless of how α_1 compares to α_2 , as $\bar{p} > p^m$, firms price above their monopoly price with positive probability. Proposition 6 shows that as costs become fully recoverable, the difference between our equilibrium and the production-to-order one becomes quantitatively small. Thus, in the limit, prices no longer exceed the monopoly level.

The intuition for why firms that produce in advance set prices above the monopoly level is the following. Consider a hypothetical monopolist facing a demand of $D(p)$ with probability λ and 0 otherwise. It maximizes


$$\left((p - \alpha c)\lambda - (1 - \alpha)c \right) D(p) = \lambda \left(p - \underbrace{\left(\alpha + \frac{1 - \alpha}{\lambda} \right) c}_{>1} \right) D(p),$$

and thus behaves as a monopolist with a unit cost that exceeds c . Hence, it prices above p^m . In our production-in-advance setting, given the strategy of its rival, a firm faces a demand of $D(p)$ with probability $1 - F(p)$, and 0 otherwise. A similar mechanism thereby raises the firm's perceived cost of production, which again rationalizes prices above the monopoly level.

Bertrand with heterogeneous unit costs. Suppose that $c_1 < c_2$, and there are neither captive consumers nor fixed costs ($\mu_i = A_i = 0$). When $\alpha_1 = \alpha_2 = 1$, the model boils down to the asymmetric Bertrand competition model. In any (undominated) equilibrium, the efficient firm serves the market at the unit cost of the inefficient firm, and firms earn $\bar{\pi}_1 = (c_2 - c_1)D(c_2)$ and $\bar{\pi}_2 = 0$ (Blume, 2003; Kartik, 2011; De Nijs, 2012).

Under production in advance ($\alpha_1, \alpha_2 < 1$), $r_1 = c_1$ and $r_2 = c_2$, so firm 1 is the strong firm. In equilibrium, firms earn their Bertrand profits as defined above, and mix continuously

Figure 1: Equilibrium CDF's for $r_1 < r_2$. Top panel: $\bar{p}_2 \leq \bar{p}_1$. Bottom panel: $\bar{p}_2 > \bar{p}_1$.



In both panels, $D(p) = 1 - p$, $c_1 = c_2 = 0.3$, $A_1 = 0.01$, $A_2 = 0.03$, and $\mu_1 = \mu_2 = 0$; in the top panel, $\alpha_1 = 0.9$ and $\alpha_2 = 0.1$; in the bottom panel, $\alpha_1 = 0.1$ and $\alpha_2 = 0.9$.

over $[r, \bar{p})$ according to

$$k_i(p) = \frac{(p - c_j)D(p) - \bar{\pi}_j}{(p - \alpha_j c_j)D(p)}.$$

As $\bar{\pi}_2 = 0$, the function $k_1(p) = (p - c_2)/(p - \alpha_2 c_2)$ is strictly increasing, and so $\bar{p}_1 = p^0$. The equilibrium is thus always described by equation (1), with the strong firm putting the rest of its mass on \bar{p} and the weak firm putting the rest of its mass on inactivity.

In contrast to the Bertrand outcome, under production in advance: (i) It is not necessarily the strong firm that ends up serving the market; (ii) the strong firm's inventory may remain unsold; (iii) there is non-trivial price dispersion; (iv) the strong firm prices above its monopoly price with positive probability (because, as in the case with fixed costs, the perceived unit cost exceeds the real one).

Despite these differences, the equilibrium converges to a Bertrand equilibrium as the α 's tend to 1. The Bertrand game has a continuum of equilibria (Blume, 2003). In Appendix C.5, we show that the limiting equilibrium corresponds to the Bertrand equilibrium in which firm 2 is the least aggressive in its randomization.

If, in addition to heterogeneous unit costs, firms have *identical* fixed costs and recoverability parameters ($A_1 = A_2$ and $\alpha_1 = \alpha_2$), then $r_1 < r_2$ and $\bar{p}_1 > \bar{p}_2$. This holds since firm 1 is advantaged in all dimensions (strictly so in one and weakly so in the other two). The production-in-advance equilibrium is therefore qualitatively unchanged as it is still given by equation (1). As the common recoverability parameter tends to 1, the game and equilibrium converge to the Bertrand game and equilibrium studied by Lang and Rosenthal (1991) and Anderson, Baik, and Larson (2015). (The latter reinterpret the Bertrand-with-fixed-costs model as a model of personalized pricing and advertising.)

If firm 1 is advantaged in some dimensions but disadvantaged in others, then the equilibrium may be described by equation (2) and thus be qualitatively different. If demand is continuous, such an equilibrium, where the weak firm has two mass points, emerges in the Bertrand limit if and only if the strong firm has a strictly higher unit cost than the weak firm.²⁰ By continuity (Proposition 6), when recoverability parameters are sufficiently high, the weak firm must also have two mass points under production in advance.

To obtain this type of equilibrium under production to order, it is necessary to have, in addition to interior monopoly prices, heterogeneous fixed and unit costs. Such equilibria have eluded the existing Bertrand literature, which has confined attention to either perfectly inelastic demand or heterogeneity in a single dimension.

Clearinghouse models. In Varian (1980)'s production-to-order model, firms are symmetric, demand is perfectly inelastic up to a choke price, and there are no advertising costs ($A = 0$). In equilibrium, firms randomize over prices according to a continuous probability measure. The fact that the CDF is continuous means that there is no reference price—as Narasimhan (1988) first pointed out, firms are always holding a sale.

²⁰To understand this condition, note that in the production-to-order limit, $\bar{p}_i = p_j^m$.

In our production-in-advance model, with the same parameters, firms mix over $[r, p^0)$ according to

$$k(p) = \frac{1 - \mu}{1 - 2\mu} \frac{p - r}{p - \alpha c}.$$

As discussed at the end of Section 3.2, in a generic equilibrium, one of the firms puts the rest of its mass on targeting both its captive and the contested segment at p^0 , whereas the other firm puts the rest of its mass on targeting only its captive segment at p^0 . Hence, both firms have a reference price, and so the frequency of promotions is no longer 1.

As inventory costs become fully recoverable, this equilibrium converges to an equilibrium of the resulting Bertrand game.²¹ Note that in the limit, both firms continue to have a mass point at the choke price—whereas there are no mass points in Varian’s equilibrium. The following seemingly innocuous difference explains the discrepancy: In our setting, firms choose not only their prices but also which segments to target, whereas in Varian’s model, both firms always target both their captive and the contested segment by assumption.

It is natural to expect that there are costs, fixed or variable, associated with targeting any given market segment—our model allows for both types of costs, whereas Varian’s accounts for neither. In the presence of either costs, firms need to actively decide not only what price to set but also which segments to target. In light of this, it seems fair to say that our model and equilibrium provide a more plausible description of sales.

Baye and Morgan (2001) consider a version of Varian’s symmetric model with advertising costs ($A > 0$) and downward-sloping demand. Due to those costs, the equilibrium of their Bertrand game features mass points at the monopoly price. However, each firm chooses not to advertise with strictly positive probability. Thus, the equilibrium of their game differs from the (generic) equilibrium of our limiting game, where one of the firms always advertises.²²

The following difference explains the discrepancy: In Baye and Morgan (2001), a firm that does not advertise, i.e., that does not list its price with the online gatekeeper, still receives demand from the contested segment provided its rival does not advertise either. Thus, just like in the Varian model, not only is each firm always targeting both its captive and the contested segment, but also the contested segment may be captured for free. In our setting, a consumer does not know that a product is available unless it is targeted, as in, e.g., Butters (1977), and Grossman and Shapiro (1984). Moreover, there is a cost associated with targeting the contested segment, and so the firms must again actively decide not only what price to set, but also which segments to target.

We believe that our setting provides a more accurate description of some industries, compared to Baye and Morgan (2001) and the literature that follows (e.g., Arnold, Li, Saliba, and Zhang, 2011; Shelegia and Wilson, 2016).²³ One example would be an industry where

²¹Because the game is non-generic, Proposition 6 cannot be applied to obtain Bertrand convergence. It is straightforward to adapt the proof to establish convergence manually.

²²Away from the limit, there are additional differences. For instance, firms price above the monopoly level.

²³The fundamental difference between these two approaches is whether the demand of a store that chooses not to advertise online is significantly affected by whether the rival advertises or not.

brick-and-mortar stores have some local captive demand, and there is a segment of remote online consumers who will never consider buying from a store that does not list its price on the gatekeeper’s platform. In contrast to Baye and Morgan (2001), our model predicts that in such an industry, at least one store must have an online presence.

As seen in the examples discussed above, the general analysis under production in advance in Sections 3.1 and 3.2 permits considerably more flexibility than existing similar models with production to order, allowing simultaneously for downward-sloping demand and firms that are heterogeneous in advertising costs, unit costs, share of captive consumers, and recoverability parameters. The convergence result in Proposition 6 also allows us to accommodate as limiting cases the Bertrand versions of such settings.

5 An All-Pay Oligopoly with Incomplete Information

We have shown that, in the generically unique equilibrium of an all-pay oligopoly model, both firms randomize over their price and inventory choices. As usual, firms must be indifferent between all the pure strategies in the support of their randomization. The model therefore predicts that managers play dice, despite them having no strict incentives to do so. To address this concern, we provide a purification argument in the spirit of Harsanyi (1973) in our framework with bi-dimensional continuous actions.

We study a model similar to the one in Section 2, in which firms have private information about their unit costs (Section 5.1). The model is therefore more restrictive in terms of asymmetries than the general class of all-pay oligopolies studied in Section 3, but it is richer in its information structure. We provide two convergence results in Section 5.2: (i) convergence to complete information as costs shocks become small (the purification result), and (ii) convergence to Bertrand as inventory costs become fully recoverable.

5.1 Framework and Equilibrium Analysis

Consider again the model of Section 2, but suppose that unit costs are drawn i.i.d. from a probability distribution G , which has strictly positive and continuous density over its support $[\underline{c}, \bar{c}]$, with $0 < \underline{c} < \bar{c} < p^0$. If the realizations of c_1 and c_2 were public information, then the analysis in Section 3 would apply to this game (with $A_i = \mu_i = 0$).²⁴ In this section, we study the case where those realizations remain private information. We assume that D is \mathcal{C}^2 , strictly decreasing and log-concave on $[0, p^0)$.

A pure strategy for firm i is a mapping $(p_i(\cdot), q_i(\cdot)) : [\underline{c}, \bar{c}] \rightarrow [0, p^0] \times \mathbb{R}_+$. A pure-strategy Bayesian equilibrium is a profile of strategies $(p_1(\cdot), q_1(\cdot), p_2(\cdot), q_2(\cdot))$ such that for every $c \in [\underline{c}, \bar{c}]$, the price-inventory pair $(p_i(c), q_i(c))$ chosen by type c of firm i maximizes that type’s expected profit, taking firm j ’s pure strategy as given.

²⁴Techniques similar to those applied in the present section could also be used to handle a setting closer to Section 3, with asymmetries, advertising costs, and captive consumers.

In line with what we did in the previous sections, we first focus on a constrained version of the model in which the firms must source enough inventories to supply their targeted demand, i.e., $q_i(\cdot) = D(p_i(\cdot))$. The analysis of this constrained game involves solving an all-pay contest with asymmetric information and non-monotonic winning functions.

For the time being, we restrict attention to symmetric Bayesian equilibria ($p_i(\cdot) = p_j(\cdot) = p^*(\cdot)$) with monotone and differentiable strategies. More precisely, we focus on symmetric equilibria with a cutoff type $c^0 \in (\underline{c}, \bar{c})$ such that p^* is continuous on $[\underline{c}, \bar{c}]$, differentiable with strictly positive derivative on $[\underline{c}, c^0)$, and equal to p^0 on $[c^0, \bar{c}]$. The pricing function p^* induces a CDF of prices $F^*(\cdot) \equiv G((p^*)^{-1}(\cdot))$ that is differentiable on $[\underline{p}, p^0] \equiv [p(\underline{c}), p^0]$.

The cutoff type c^0 is pinned down by the zero-profit condition

$$(p^0 - \alpha c^0)(1 - G(c^0)) = (1 - \alpha)c^0.$$

The expected profit of type $c < c^0$ when setting price p and facing the CDF of prices F^* is:

$$\pi(p, c) = \left((p - \alpha c)(1 - F^*(p)) - (1 - \alpha)c \right) D(p).$$

The first-order condition of profit maximization for that type c can be written as:

$$\frac{\partial \log \pi}{\partial p} = \frac{D'(p)}{D(p)} + \frac{1 - F^*(p) - (p - \alpha c)F^{*'}(p)}{(p - \alpha c)(1 - F^*(p)) - (1 - \alpha)c} = 0.$$

By definition of the putative equilibrium CDF F^* , pricing at p is optimal for type $c = G^{-1}(F^*(p))$. Combining this with the above optimality condition and rearranging terms, we obtain the ordinary differential equation $F^{*'}(p) = \Psi(p, F^*(p))$, where

$$\Psi(p, F) \equiv \frac{1}{p - \alpha G^{-1}(F)} \left(1 - F + \frac{D'(p)}{D(p)} \left((p - \alpha G^{-1}(F))(1 - F) - (1 - \alpha)G^{-1}(F) \right) \right).$$

F^* must also satisfy the boundary conditions $F^*(\underline{p}) = 0$ and $\lim_{p \uparrow p^0} F^*(p) = G(c^0) \equiv F^0$, and F^* must be increasing and satisfy the positive profit condition

$$(p - \alpha G^{-1}(F^*(p)))(1 - F^*(p)) > (1 - \alpha)G^{-1}(F^*(p)), \quad \forall p \in [\underline{p}, p^0].$$

Conversely, it is easy to show that, if F^* solves the above boundary-value problem and satisfies the monotonicity and positive profit conditions, then the induced pricing function $p^*(\cdot) = (F^*)^{-1} \circ G(\cdot)$ is a Bayesian equilibrium of the constrained game. In fact, for every $c \in [\underline{c}, \bar{c}]$, $p^*(c)$ is the unique global maximizer of $\pi(\cdot, c)$, i.e., the Bayesian equilibrium is strict. Much like in the analysis in Section 3.2 (see the discussion after Proposition 3), the fact that each type $c \in [\underline{c}, c^0)$ makes strictly positive profits implies that this equilibrium is also a strict Bayesian equilibrium of the unconstrained game, where firms can freely choose their inventory levels.

This analysis leaves a number of questions open. Does the boundary value problem have a

solution and if so, is the solution unique? Do equilibria of the unconstrained game necessarily satisfy the properties imposed above? The following theorem addresses these questions:

Theorem 2. *Consider an all-pay oligopoly model with incomplete information:*

- (a) *In any equilibrium, firms source enough inventories to supply the demand they target.*
- (b) *The boundary-value problem $F' = \Psi(p, F)$ with boundary conditions $F(\underline{p}) = 0$ for some $\underline{p} \in (0, p^0)$ and $\lim_{p \uparrow p^0} F(p) = F^0$ has a unique solution F^* that satisfies the positive profit and monotonicity conditions.*
- (c) *If G is \mathcal{C}^2 and convex, D is \mathcal{C}^3 on $[\underline{c}, p^0)$, and $DD''/(D')^2$ is non-decreasing, then the model has a unique equilibrium. That equilibrium is strict and the equilibrium CDF of prices is F^* .*
- (d) *There exists a unique equilibrium with continuous (resp. symmetric) pricing strategies. That equilibrium is strict and the equilibrium CDF of prices is F^* .*

Proof. See Online Appendix II. □

The equilibrium under incomplete information therefore also satisfies the generic properties (I)–(VI) of Section 3.3. Property (VII), however, is not satisfied as both firms are inactive with strictly positive probability. The Bertrand-related properties (properties (VIII) and (IX)) will be discussed in Section 5.2.

We close this subsection with an overview of the key steps of the proof. Part (a) shows that it is indeed without loss of generality to focus on equilibria of the constrained game. The proof follows a similar development as the proof of Proposition 4. Assuming for a contradiction that the property stated in part (a) does not hold and defining \hat{p} as the lowest price below which both firms source enough inventories to supply their targeted demand, we show that the expected profits of the types that price in the neighborhood of \hat{p} are bounded away from zero. A type that prices at, or just above, \hat{p} therefore makes a strictly positive margin on the last unit it sources, which implies that such a type must be sourcing enough inventories to supply its targeted demand, contradicting the definition of \hat{p} .

The proof of parts (b), (c), and (d) uses tools that were developed to study first-price auctions (e.g., Riley and Samuelson, 1981; Plum, 1992; Maskin and Riley, 2000, 2003; Lebrun, 1999, 2006) with some major additional difficulties, which all relate to the fact that a firm's winning function is non-monotonic in the price it sets.

The main technical difficulty in proving part (b) is that the differential equation $F' = \Psi(p, F)$ is singular at (p^0, F^0) , which prevents us from applying the Picard-Lindelöf theorem to establish the existence and uniqueness of a solution trajectory that goes through (p^0, F^0) . We circumvent this issue with a geometric argument similar in spirit to the one in Lebrun (2006). We show that the function $\underline{F} : p \mapsto \frac{p-c^0}{p-\alpha c^0}$ is a super solution of $F' = \Psi(p, F)$, whereas the zero-profit CDF of prices $\overline{F}(\cdot)$, implicitly defined as the unique solution of

$(p - \alpha G^{-1}(F))(1 - F) = (1 - \alpha)G^{-1}(F)$, is a sub solution.²⁵ Any solution trajectory that crosses \bar{F} (resp. \underline{F}) must remain above \bar{F} (resp. below \underline{F}) thereafter, implying that no such trajectory can solve the boundary value problem. Applying standard results for differential equations, we show that for some p^* , the solution trajectory that starts at $(p^*, 0)$ remains between \bar{F} and \underline{F} . The associated solution trajectory solves the boundary value problem. Finally, a similar argument is used to prove that the solution F^* is unique.

Parts (c) and (d) of Theorem 2 rule out equilibria with non-differentiable pricing functions. Supermodularity arguments imply that any equilibrium must involve strictly increasing pricing functions (possibly up to a cutoff type) and therefore continuous induced CDF's of prices (except at p^0). The main difficulty involves ruling out discontinuities in a firm's pricing function. Assume that $p_i(c^-) \equiv \lim_{c' \uparrow c} p_i(c') < \lim_{c' \downarrow c} p_i(c') \equiv p_i(c^+)$ for some type c . Standard arguments (see, e.g., Plum, 1992) imply that $p_i(c^-)$ and $p_i(c^+)$ are both optimal for that type. Moreover, by monotonicity of $p_i(\cdot)$, no type of firm i prices in $(p_i(c^-), p_i(c^+))$. If the winning function were monotonic in price, as in the auctions literature, then firm j would not be willing to price in that interval either. This would immediately imply that $p_i(c^-)$ and $p_i(c^+)$ cannot both be optimal for firm i , a contradiction.

In our oligopoly framework, however, a firm's winning function is hump-shaped in its price, and the above argument therefore does not apply. Some types of firm j may respond by setting a (modified) monopoly price in $(p_i(c^-), p_i(c^+))$: We show that if firm i 's pricing function is discontinuous at c , then there exists an interval of types $[\underline{c}_j, \bar{c}_j]$ such that

$$p_j(c_j) = p^m \left(\left(\alpha + \frac{1 - \alpha}{1 - F_i(p_i(c^-))} \right) c_j \right), \quad \forall c_j \in [\underline{c}_j, \bar{c}_j],$$

where $p^m(\tilde{c})$ is the monopoly price given unit cost \tilde{c} , $p_j(\underline{c}_j) = p_i(c^-)$, and $p_j(\bar{c}_j) = p_i(c^+)$. Firm j 's induced CDF of prices over $[p_i(c^-), p_i(c^+)]$ is therefore:

$$F_j(\cdot) = G \left((p^m)^{-1}(\cdot) / \left(\alpha + \frac{1 - \alpha}{1 - F_i(p_i(c^-))} \right) \right).$$

The assumptions made in part (c) of Theorem 2 imply that the monopoly pass-through of a cost increase is non-increasing in cost, i.e., $p^{m''} \leq 0$ (see, e.g., Fabinger and Weyl, 2012), and that F_j is convex on $[p_i(c^-), p_i(c^+)]$. In turn, the convexity of F_j implies that firm i 's expected payoff is single-peaked on $[p_i(c^-), p_i(c^+)]$, and that $p_i(c^-)$ and $p_i(c^+)$ cannot be simultaneously optimal for firm i , a contradiction. Once continuity has been established, the incentive compatibility constraints of nearby types can be used to show that CDF's of prices are symmetric and differentiable, and solve the boundary value problem.

The above complications and the assumption made in part (c) of the theorem can be avoided if one is willing to confine attention to Bayesian equilibria that satisfy certain properties, such as continuity or symmetry. Part (d) of the theorem shows that there is a unique

²⁵The function ϕ is a sub solution (resp., super solution) of $F' = \Psi(p, F)$ if $\phi'(p) < \Psi(p, \phi(p))$ (resp., $\phi'(p) > \Psi(p, \phi(p))$) for every p . See Teschl (2012) for a treatment of ordinary differential equations.

equilibrium satisfying such properties.

5.2 Purification and Bertrand Convergence

Purification. We now show that the equilibrium of Section 2 can be approximated by the equilibrium of a nearby game perturbed by incomplete information. As shown in Section 5.1, the equilibrium of the perturbed game is strict and in pure strategies, implying that managers need not play dice.

Proposition 7. *Let $(G^n)_{n \geq 0}$ be a sequence of probability measures over \mathbb{R}_+ . Assume that D and G^n ($n \geq 0$) satisfy the assumptions of Section 5.1, and let F^n be the equilibrium CDF of prices given G^n , as characterized in Theorem 2. Suppose that $(G^n)_{n \geq 0}$ converges weakly to the probability measure that puts full weight on $c \in (0, p^0)$. Then, $(F^n)_{n \geq 0}$ converges weakly to the equilibrium CDF of prices of Proposition 1.*

Proof. See Online Appendix III.1.²⁶ □

Non-degenerate mixed-strategy equilibria have a seemingly undesirable feature: Such equilibria are not regret-free, in the sense that, once firm i has observed the realization of firm j 's price, firm i no longer wants to mix in the way prescribed by its equilibrium strategy. The purification result does not address this concern.

In a retailing context, this concern may however be of limited relevance since it may be impossible (or too costly) for a firm to change and re-advertise its price in a reasonable time-frame. The fact that retailers do often end up with unsold inventories, without changing prices, seems to provide anecdotal support for this claim.

As discussed in Section 1.1, the model can also be seen as describing a procurement setting—the application that motivated Hansen (1988)'s analysis. In that case, the procurement agency is committed to its auction format and all submitted bids are final.

Bertrand convergence. We now show that the equilibrium of the production-in-advance game converges to a Bertrand equilibrium as costs become fully recoverable—the counterpart of generic property (IX) of Section 3.3.

The Bertrand version of our incomplete-information game was studied by Hansen (1988) and Spulber (1995). The Bertrand game has a unique equilibrium in which all but the highest types make strictly positive profits and all types price strictly below their monopoly price. The equilibrium CDF of prices solves the boundary value problem

$$F'(p) = \frac{1 - F}{p - G^{-1}(F)} \left(1 + \frac{D'(p)}{D(p)} (p - G^{-1}(F)) \right), \quad F(\underline{p}) = 0, \quad F(\bar{c}) = 1.$$

²⁶The proposition is proven by using the fact that for every n , F^n is bounded below by the super solution $\underline{F}^n : p \mapsto \frac{p - c^{0,n}}{p - \alpha c^{0,n}}$, where $c^{0,n}$ is the equilibrium cutoff type given G^n , and bounded above by the sub solution \bar{F}^n , the zero profit CDF of prices given G^n . Since $(\underline{F}^n)_{n \geq 0}$ and $(\bar{F}^n)_{n \geq 0}$ both converge pointwise to the equilibrium CDF of prices of Proposition 1, the result follows by the sandwich theorem.

Note that this differential equation coincides with the one studied in Section 5.1 for $\alpha = 1$, i.e., when inventory costs are fully recoverable. Showing that the equilibrium under production in advance converges to the Hansen-Spulber one as α tends to one does however require a few additional steps, which we undertake next:

Proposition 8. *Suppose that D and G satisfy the assumptions of Section 5.1, and let $(\alpha^n)_{n \geq 0}$ be a sequence over $[0, 1)$ such that $\alpha^n \xrightarrow[n \rightarrow \infty]{} 1$. For every n , let $F^*(\cdot, \alpha^n)$ be the equilibrium CDF of prices characterized in Theorem 2. Then, $(F^*(\cdot, \alpha^n))_{n \geq 0}$ converges weakly to $F^*(\cdot, 1)$, the Hansen-Spulber equilibrium CDF of prices.*

Proof. See Online Appendix III.2. □

Property (VIII), i.e., the fact that firms earn their Bertrand profits in equilibrium, can be interpreted in two ways under incomplete information. The *interim* interpretation is that each type earns its Bertrand profit: This property is clearly not satisfied, as high types make zero profit when α is small, and positive profits when α is high. In fact, in an example with linear demand and uniformly-distributed costs, each type's equilibrium profit is increasing in α (see Online Appendix II.7). The *ex-ante* interpretation is that each firm earns its Bertrand profit in expectation: That same example shows that this property does not hold either.

There are other qualitative differences between the equilibrium under production in advance and the Bertrand equilibrium: Under production in advance, some inventories remain unsold with positive probability; a positive mass of types sets prices above the monopoly level (for the same reason as in Section 4); a positive mass of types stays out of the market.

We close this subsection with an overview of the proof of Proposition 8. We first show that as n goes to infinity, the sequence of cutoff types $(c^{0,n})_{n \geq 0}$ tends to \bar{c} . Hence, the sequence of super solutions $(\underline{F}^n)_{n \geq 0}$ converges pointwise to 1 on $(\bar{c}, p^0]$. Since $F^*(\cdot, \alpha^n)$ is bounded below by \underline{F}^n for every n , this implies that $(F^*(\cdot, \alpha^n))_{n \geq 0}$ converges pointwise to 1 on that interval.

Next, we apply Helly's selection theorem to extract a subsequence $(F^*(\cdot, \alpha'^n))_{n \geq 0}$ that converges weakly to some probability measure \hat{F} (e.g., Theorem 23.9 in Billingsley, 2012). Using the fact that $F^*(\cdot, \alpha'^n)$ solves a differential equation, we argue that $(F^*(\cdot, \alpha'^n))_{n \geq 0}$ and $(F^{*'}(\cdot, \alpha'^n))_{n \geq 0}$ are equicontinuous on $[0, \bar{p}]$ for all $\bar{p} < \bar{c}$. The Arzelà-Ascoli theorem (e.g., Theorem 7.25 in Rudin, 1976) implies that $(F^*(\cdot, \alpha'^n))_{n \geq 0}$ and $(F^{*'}(\cdot, \alpha'^n))_{n \geq 0}$ converge uniformly on $[0, \bar{p}]$, so \hat{F} is \mathcal{C}^1 on that interval. We can then take limits in the differential equation to show that \hat{F} is a solution of the Hansen-Spulber differential equation. Finally, we show that \hat{F} satisfies the boundary conditions and conclude that $\hat{F} = F^*(\cdot, 1)$.

6 Efficiency and Taxation

The equilibrium outcome in an all-pay oligopoly game of complete information features three types of distortions: First, firms price above marginal cost with probability one (the classical deadweight loss); second, some inventories remain unsold with positive probability; third, it is

not necessarily the most efficient firm that ends up serving the contested segment. This raises the question of whether taxes or subsidies can alleviate those distortions, as they usually do in oligopoly models.

We find a surprisingly simple answer to this question in the context of the symmetric model seen in Section 2: Symmetric and linear taxation schemes can never improve efficiency.

Consider the oligopoly model of Section 2. Suppose that a firm faces a tax rate of t for each unsold unit, where t is common to both firms and can be positive or negative. The indifference condition for the mixed-strategy equilibrium becomes

$$p(1 - F(p)) - c + (\alpha c - t)F(p) = 0,$$

which gives rise to the CDF $F(p) = \frac{p-c}{p-\alpha c+t}$ for $p \in [c, p^0]$. Similarly, if firms face a tax rate of t per unit of output, then the indifference condition becomes

$$p(1 - F(p)) - (c + t) + \alpha c F(p) = 0,$$

which pins down the equilibrium CDF of prices $F(p) = \frac{p-(c+t)}{p-\alpha c}$. We could likewise envision combinations of symmetric per-unit and/or ad-valorem taxes on sales, output and unsold units. All those policies would give rise to a *symmetric* mixed-strategy equilibrium in which each firm i draws its price p_i from a common CDF, and sources $q_i = D(p_i)$.

Instead of separately examining the impact of each of those policies (or combinations thereof), we prove a stronger result. In a nutshell, suppose the social planner can choose any *symmetric* distribution of prices. Then, the planner would choose the equilibrium distribution of prices of Proposition 1. In this sense, the equilibrium is second-best efficient.

We now formalize this argument. Let \mathcal{F} be the set of CDF's over \mathbb{R}_+ . Suppose each firm i draws its price p_i according to $F \in \mathcal{F}$, and sources $D(p_i)$. Social welfare is then given by:

$$W(F) = \underbrace{\int_{[0, p^0]} \left(\int_p^{p^0} D(t) dt + pD(p) \right) dG(p)}_{\text{expected gross utility}} - \underbrace{2c \int_{[0, p^0]} D(p) dF(p)}_{\text{expected production costs}} + \underbrace{\alpha c \int_{[0, p^0]} D(p) dH(p)}_{\text{expected salvaged costs}},$$

where $G = 1 - (1 - F)^2$ and $H = F^2$ are the CDF's of the minimum and maximum price, respectively. We prove the following result:

Proposition 9. *The equilibrium policy F^* maximizes social welfare W . Moreover, if D is strictly decreasing, then F^* is the unique social welfare-maximizing policy.*

Proof. See Appendix D. □

Thus, it is not possible to raise social welfare by using a taxation scheme that gives rise to a Nash equilibrium in which each firm i draw its price p_i from a symmetric $F \in \mathcal{F}$, and sources $D(p_i)$. In particular, if D is strictly decreasing, any symmetric and linear taxation scheme must lower social welfare.

This result is best illustrated by the case of unit inelastic demand up to the choke price. In this case, social welfare only depends on the probability that a firm sets a price below p^0 and produces one unit. Call this probability β . With probability $1 - \beta$, the firm produces nothing. Expected social welfare is therefore given by:

$$W(\beta) = (1 - (1 - \beta)^2) p^0 - 2\beta c + \alpha\beta^2 c.$$

The first-order condition is $(1 - \beta)p^0 - c + \beta\alpha c = 0$, which coincides with the zero-profit condition at p^0 . Solving out for β , we obtain

$$\beta = \frac{p^0 - c}{p^0 - \alpha c},$$

which is also the probability that a firm is active (i.e., prices below p^0) in equilibrium.

The proof with an arbitrary demand function relies on an integration by parts argument for monotone but potentially discontinuous functions (see Border, 1996). We show that $W(F)$ can be rewritten as

$$W(F) = \Phi \left(p^0, \lim_{p \uparrow p^0} F(p) \right) \lim_{p \uparrow p^0} D(p) - \int_{[0, p^0)} \Phi(p, F(p)) dD(p),$$

where $\Phi(p, F) \equiv p(1 - (1 - F)^2) - 2cF + \alpha cF^2$. We then show that $F^*(p)$ is the unique maximizer of $\Phi(p, \cdot)$. This allows us to maximize $W(F)$ term by term and obtain Proposition 9.

Determining the optimal tax policy in the context of the general model studied in Section 3 would require completely different techniques, as we would no longer be able to choose directly the CDF's of prices. In the simple case where demand is inelastic up to p^0 , $c_1 < c_2$, and there are neither captive consumer nor fixed costs, it is easy to see that the optimal policy is to tax firm 2 out of the market by making its reach exceed p^0 . This can for example be done by imposing a tax rate of $t \in (p^0 - c_2, p^0 - c_1)$ per unit of output. This argument does not extend to the case of elastic demand, as such a policy would typically affect the deadweight loss as well. A thorough analysis of these issues is left for future research.

7 Concluding Remarks

We introduced and studied a class of games where stores source costly, unobservable inventories before simultaneously setting prices. Our framework allows for firm asymmetries, heterogeneous consumer tastes, endogenous consumer information through advertising, and salvage values for unsold units.

We first studied a constrained version of the model in which stores must source enough inventories to supply all their targeted demand. That constrained game is an all-pay contest with outside options, non-monotonic winning and losing functions, and conditional invest-

ments. The equilibrium is generically unique and in mixed strategies—thus, there is price dispersion and some inventories may remain unsold. Turning our attention to the unconstrained game, where firms freely choose inventories, we showed that the sets of constrained and unconstrained equilibria coincide generically. Thus, the equilibrium of the unconstrained game is generically unique and can be found by solving the constrained game.

We studied the limiting case where the per-unit inventory cost becomes fully recoverable. An important result is that equilibrium behavior converges to an equilibrium of the associated Bertrand game, in which stores only choose prices and produce to order. Several benchmark outcomes of oligopoly theory, where production to order is assumed, can thus also be seen as the limiting outcome of similar situations with production in advance. Away from that limit, our closed-form characterization generalizes the Bertrand-type analysis to those situations where the value of unsold inventories falls short of their acquisition value.

These results are robust to the introduction of incomplete information. Our analysis under incomplete information also purifies the mixed-strategy equilibria of all-pay oligopoly games of complete information.

Finally, we investigated whether taxes or subsidies alleviate the distortions that arise in all-pay oligopoly games, as is usually the case in oligopoly models. In a symmetric version of the model, we proved that the equilibrium cannot be improved by linear and symmetric taxation, i.e., in that sense, *laissez-faire* is second-best efficient.

To introduce product differentiation, we assumed that some consumers only wish to purchase the product of one store, and therefore only a fraction of consumers (the contested segment of shoppers) wishes to purchase from the store setting the lowest price. This specification of product differentiation ensures that the constrained game has the structure of an all-pay contest. That property would disappear if we were to use other more widely-used specifications of product differentiation that result in a smoother residual demand. The analysis of such models would therefore require completely different techniques.

In the main text, we focused on the two-store case. The general case with $N \geq 2$ stores is studied formally in Online Appendix IV. There, we use techniques similar to those in the main text to prove that the sets of constrained and unconstrained equilibria coincide generically (Propositions I and II). We establish equilibrium existence (Proposition IV), and provide a generic characterization of equilibrium payoffs (Proposition III) and a partial characterization of equilibrium behavior (Proposition V). If stores can be unambiguously ranked in terms of advantages (in a way that is made precise in the statements of Propositions VI and VII), we show that only the two most advantaged stores compete in the contested segment with the same strategies as in the duopoly case, while all the other stores target exclusively their captive segments.²⁷ Finally, in Online Appendix IV.6, we provide a preliminary analysis of the impact of entry on consumer surplus.

Throughout the paper, we confined attention to a static, one-shot setting. The equilibrium

²⁷More than two stores may target the contested segment with positive probability if the ranking assumption is not satisfied. See Proposition VIII for a worked-out example.

behavior in the one-shot game does however coincide with the per-period equilibrium behavior in a properly-specified dynamic game, as explained next. Suppose that time is discrete and runs for an infinite number of periods, goods never perish, and stores discount future profits with the same discount factor $\delta < 1$. Suppose also that, at the beginning of each period, each store can freely buy and sell inventories on a perfectly competitive wholesale market at an exogenously-given price of c .²⁸ We continue to assume that the stores do not observe each others' inventory holdings and set their prices simultaneously. The difference relative to the static model studied in the paper is that each unsold unit is now carried over to the next period.

Because stores can freely buy and sell in the wholesale market at the inventory-choice stage, a firm's opportunity cost of selling one unit of inventory to consumers is equal to a fraction δ of the unit cost c —regardless of how many units that firm has at the beginning of that period. For this reason, the dynamic game has a Markov-perfect equilibrium in which, in every period, stores play the Nash equilibrium of the one-shot game with cost c and recoverability parameter $\alpha = \delta$. Such dynamic versions of the model also provide micro-foundations for salvage values reflecting the time value of money, the time between periods, and the extent to which unsold inventories are perishable—the argument also holds for heterogeneous unit costs and discount factors.²⁹ Connecting with the Bertrand convergence results, this model predicts that if inventories remain unobservable and perish slowly, then an outcome close to Bertrand emerges when the time between periods is short.

Our Bertrand convergence results stand in contrast to Kreps and Scheinkman (1983)'s well-known result that under production in advance, if stores observe inventories before setting their prices, then the Cournot outcome should be expected. This suggests that the information stores have about rivals' inventories at the pricing stage affects the nature of competition significantly.

It is then natural to ask how inventory observability affects consumer surplus and social welfare. If demand is perfectly inelastic up to a choke price, any Cournot equilibrium results in both stores jointly serving all demand at the choke price. Thus, under inventory observability, consumers face high prices but the outcome is efficient. Under inventory unobservability, stores also serve all demand, but the mixed-strategy equilibrium results in lower prices and in some inventories being unsold. Thus, in the particular case of perfectly inelastic demand, inventory observability unambiguously reduces consumer surplus but raises social welfare.

However, this analysis becomes significantly more complex in the case of elastic demand. In general, it is no longer the case that inventory observability results in higher prices and social welfare. Moreover, in a dynamic setting, one may also worry that sharing inventory information could facilitate collusion. A more thorough analysis of the static and dynamic effects of inventory observability on market outcomes is left for future research.

²⁸In the one-shot setting studied in the paper, capacity and inventory choices are formally equivalent. This equivalence may break down in a dynamic setting, as it no longer seems reasonable to assume that firms can adjust their capacities at the beginning of each period.

²⁹Perishability can be captured by a fraction of the inventory being lost per unit of time.

Appendix

A Equilibrium Analysis in the Constrained Game

In this section, we state and prove a series of technical lemmas that jointly imply Proposition 2. We then provide a complete characterization of equilibria in the non-generic case (Proposition A below). Fix an all-pay oligopoly game satisfying the assumptions made at the beginning of Section 3.1. To fix ideas, assume $r_1 \leq r_2 = r$.

Lemma A. *In any equilibrium (F_1, F_2) of the constrained game, if F_i is discontinuous at $\hat{p} \in [0, p^0]$, then firm $j \neq i$ earns strictly less than its equilibrium payoff when it prices at \hat{p} .*

Proof. Put $F_i^-(\hat{p}) = \lim_{p \uparrow \hat{p}} F_i(p) < F_i(\hat{p})$, and let $\bar{\pi}_\iota$ denote firm ι 's equilibrium payoff ($\iota = 1, 2$). Using the same-price fair-share rule, firm ι 's payoff in case of a tie is:

$$t_\iota(p) = \mu_\iota(p - c_\iota)D(p) + (1 - \mu_1 - \mu_2) \left(\frac{1}{2}(p - \alpha_\iota c_\iota) - (1 - \alpha_\iota)c_\iota \right) D(p) - A_\iota.$$

Firm j 's expected payoff when it prices at \hat{p} is given by:

$$\tilde{\pi}_j = (1 - F_i(\hat{p}))w_j(\hat{p}) + (F_i(\hat{p}) - F_i^-(\hat{p}))t_j(\hat{p}) + F_i^-(\hat{p})l_j(\hat{p}).$$

Assume first that $t_j(\hat{p}) \geq w_j(\hat{p})$. Then, $(\hat{p} - \alpha_j c_j)D(\hat{p}) \leq 0$. Assume for a contradiction that $\hat{p} = p^0$. Since $p^0 > c_j$, this implies that $D(p^0) = 0$. Therefore, since firm i has a mass point at \hat{p} , $0 \leq o_i \leq \bar{\pi}_i = -A_i$, and $A_i = \mu_i = 0$. The convention we adopted in footnote 13 implies that F_i puts no mass on $p^0 = \hat{p}$, which is a contradiction. Hence, $\hat{p} < p^0$. It follows that $\hat{p} - \alpha_j c_j \leq 0$, and that $\tilde{\pi}_j < 0 \leq o_j \leq \bar{\pi}_j$, as in the statement of the lemma.

Assume instead that $t_j(\hat{p}) < w_j(\hat{p})$. Let $(p^n)_{n \geq 1}$ be a strictly increasing sequence such that $p^n \xrightarrow[n \rightarrow \infty]{} \hat{p}$ and F_i puts no mass on $\{p^n\}$ for every n . Then, for every n ,

$$\begin{aligned} \bar{\pi}_j &\geq (1 - F_i(p^n))w_j(p^n) + F_i(p^n)l_j(p^n), \\ &\xrightarrow[n \rightarrow \infty]{} (1 - F_i^-(\hat{p}))w_j(\hat{p}) + F_i^-(\hat{p})l_j(\hat{p}), \\ &> (1 - F_i(\hat{p}))w_j(\hat{p}) + (F_i(\hat{p}) - F_i^-(\hat{p}))t_j(\hat{p}) + F_i^-(\hat{p})l_j(\hat{p}), \\ &= \tilde{\pi}_j. \end{aligned} \quad \square$$

Lemma B. *In any equilibrium (F_1, F_2) of the constrained game, firm i 's expected profit is equal to $w_i(r)$, the infimum of the support of F_i is r , and $F_i(r) = 0$ ($i = 1, 2$), i.e., no firm has a mass point on r .*

Proof. Fix an equilibrium, and let $\bar{\pi}_i$ (resp. \underline{p}_i) denote firm i 's payoff (resp. the infimum of the support of F_i) in this equilibrium. Clearly, $\bar{\pi}_i \geq o_i$ for every firm i . Moreover, since every price $p < r$ is strictly dominated for firm 2, that firm puts not weight on $[0, r)$. Therefore,

$\bar{\pi}_1 \geq w_1(p)$ for every $p < r$, and $\bar{\pi}_1 \geq w_1(r)$. Hence, firm 1 puts no weight on $[0, r)$. To sum up, we have that, for every firm i , $\bar{\pi}_i \geq w_i(r)$ and $\underline{p}_i \geq r$.

Assume for a contradiction that $\bar{\pi}_i > w_i(r)$ for some firm i . Then, $\underline{p}_i > r$. Hence, if firm $j \neq i$ prices in the interval (r, \underline{p}_i) , then it wins the contest for sure. Since w_j is locally strictly increasing at r , this implies that $\bar{\pi}_j > w_j(r)$. Hence, $\bar{\pi}_i > o_i$ for $i = 1, 2$, and both firms participate in the contest for sure. Let \check{p}_i be the supremum of the support of F_i ($i = 1, 2$). If $\check{p}_i > \check{p}_j$, then there exists $p > \check{p}_j$ such that $\bar{\pi}_i = l_i(p) \leq o_i < \bar{\pi}_i$, which is a contradiction. Hence, $\check{p}_i = \check{p}_j \equiv \check{p}$. If firm i has a mass point at \check{p} but firm j does not, then $\bar{\pi}_i = l_i(\check{p})$, a contradiction. Therefore, by Lemma A, no firm has a mass point at \check{p} . There exists a strictly increasing sequence $(p^n)_{n \geq 1}$ such that $p^n \xrightarrow[n \rightarrow \infty]{} \check{p}$ and, for every n , $\bar{\pi}_i$ is equal to firm i 's expected profit when it prices at p^n . Lemma A implies that firm j puts no mass on $\{p^n\}$ for every n . Combining this with the continuity of F_j at \check{p} delivers a contradiction:

$$\bar{\pi}_i = (1 - F_j(p^n))w_i(p^n) + F_j(p^n)l_i(p^n) \xrightarrow[n \rightarrow \infty]{} l_i(\check{p}) \leq o_i < \bar{\pi}_i.$$

Hence, $\bar{\pi}_i = w_i(r)$ for $i = 1, 2$, which immediately implies that $\underline{p}_1 = \underline{p}_2 = r$.

Assume for a contradiction that firm i has a mass point at r . Then, by Lemma A, firm j cannot have a mass point at r . There exists a strictly decreasing sequence $(p^n)_{n \geq 0}$ such that $p^n \xrightarrow[n \rightarrow \infty]{} \check{p}$ and, for every n , $\bar{\pi}_j$ is equal to firm j 's expected profit when it prices at p^n . Lemma A implies that firm i puts no mass on $\{p^n\}$ for every n . Combining this with the right continuity of F_i delivers a contradiction:

$$\bar{\pi}_j = (1 - F_i(p^n))w_j(p^n) + F_i(p^n)l_j(p^n) \xrightarrow[n \rightarrow \infty]{} (1 - F_i(r))w_j(r) + F_i(r)l_j(r) < w_j(r). \quad \square$$

Recall from the analysis in the main text that $k_j(p) = (w_i(p) - w_i(r))/(w_i(p) - l_i(p))$ for every $p \in [r, p^0)$. We now establish some useful facts about k_j :

Lemma C. *The following holds:*

(i) k_j is strictly concave on $[r, p^0)$. Either k_j achieves a global maximum at some $\bar{p}_j \in (r, p^0)$, or it is strictly increasing on $[r, p^0)$. In the latter case, set $\bar{p}_j = p^0$.

(ii) $k_j(\bar{p}) (= \lim_{p \uparrow \bar{p}} k_j(p)) < 1$, where $\bar{p} = \min(\bar{p}_1, \bar{p}_2)$.

Proof. To prove the first part of the lemma, note that

$$\begin{aligned} \frac{1 - \mu_i - \mu_j}{1 - \mu_j} k_j(p) &= \frac{(p - c_i)D(p) - (r - c_i)D(r)}{(p - \alpha_i c_i)D(p)}, \\ &= \frac{p - c_i}{p - \alpha_i c_i} + (r - c_i)D(r)\Phi(\log((p - \alpha_i c_i)D(p))), \end{aligned}$$

where $\Phi(x) = -e^{-x}$. Since Φ is concave and increasing and $p \mapsto (p - \alpha_i c_i)D(p)$ is log-concave, it follows that $p \mapsto \Phi(\log((p - \alpha_i c_i)D(p)))$ is concave. Hence, k_j is the sum of a strictly concave function and a concave function. It follows that k_j is strictly concave.

We now turn to the second part of the lemma. If $\bar{p} < p^0$, the result follows immediately from the fact that $l_i(\bar{p}) < o_i \leq w_i(r)$ and $w_i(\bar{p}) > l_i(\bar{p})$. Suppose instead that $\bar{p} = p^0$. If $D(p^0) > 0$, then $\lim_{p \uparrow p^0} w_i(p) \geq o_i > \lim_{p \uparrow p^0} l_i(p)$, and therefore, $\lim_{p \uparrow p^0} k_j(p) < 1$. If instead $D(p^0) = 0$, then $w_i(r) = A_i = \mu_i = 0$ (for otherwise, k_j would start decreasing before p^0). Hence, $k_j(p) = \frac{p - c_i}{p - \alpha_i c_i}$, which is indeed bounded away from 1. \square

We now argue that the equilibrium F_1 and F_2 are uniquely pinned down on $[r, \bar{p}]$:

Lemma D. *In any equilibrium (F_1, F_2) of the constrained game, $F_i(p) = k_i(p)$ for every $p \in [r, \bar{p}]$ and $i \in \{1, 2\}$. Moreover, if $\bar{p}_i = \bar{p}$, then F_j is constant on $[\bar{p}, p^0]$ ($j \neq i$).*

Proof. Fix an equilibrium (F_1, F_2) . Let $\pi_i(p)$ denote firm i 's expected profit when it prices at p . Let $i \in \{1, 2\}$ and $p \in [r, p^0]$. If $F_i(p) < k_i(p)$, then firm j can price at (or just below) p and earn a profit strictly greater than $w_j(r)$, contradicting Lemma B. Hence, $F_i(p) \geq k_i(p)$ for every $p \in [r, p^0]$. Note also that $\pi_j(p) < w_j(r)$ whenever $F_i(p) > k_i(p)$. Moreover, if $D(p^0) > 0$, then $k_i(p^0)$ is well defined. Therefore, it is also the case that $\pi_j(p^0) < w_j(r)$ if $F_i(p^0) > k_i(p^0)$.

Suppose that $\bar{p}_i = \bar{p} < p^0$, and let $p \in (\bar{p}, p^0)$. Then,

$$F_i(p) \geq F_i(\bar{p}) \geq \lim_{p' \uparrow \bar{p}} F_i(p') \geq \lim_{p' \uparrow \bar{p}} k_i(p') = k_i(\bar{p}) = k_i(\bar{p}_i) > k_i(p).$$

Therefore, $\pi_j(p) < w_j(r)$ for every $p \in (\bar{p}, p^0)$, and F_j is constant on $[\bar{p}, p^0]$. We now show that F_j puts no mass on p^0 either. Since $\bar{p}_i < p^0$, we have that $w_j(r) > 0$. Hence, if $D(p^0) = 0$, then firm j clearly does not want to price at p^0 . If instead $D(p^0) > 0$, then $k_i(p^0)$ is well defined, and the above reasoning implies that $\pi_j(p^0) < w_j(r)$.

Assume for a contradiction that firm i puts strictly positive mass on some $\hat{p} \in (r, \bar{p})$. Since $F_i(p) \geq k_i(p)$ for every $p < \hat{p}$, $F_i(\hat{p}) > \lim_{p \uparrow \hat{p}} F_i(p) \geq k_i(\hat{p})$. By continuity of k_i and monotonicity of F_i , this implies that, for some $\varepsilon > 0$, $F_i(p) > k_i(p)$ for every $p \in [\hat{p}, \hat{p} + \varepsilon]$. Hence, $\pi_j(p) < w_j(r)$ for every $p \in [\hat{p}, \hat{p} + \varepsilon]$, and F_j is therefore constant on that interval. Hence, $F_j(\hat{p}) = F_j(\hat{p} + \varepsilon) \geq k_j(\hat{p} + \varepsilon) > k_j(\hat{p})$, and $\pi_i(\hat{p}) < w_i(r)$, contradicting the fact that firm i has a mass point at \hat{p} . We conclude that firm i has no mass points on $[r, \bar{p}]$, i.e., F_i is continuous on that interval ($i = 1, 2$). This implies in particular that π_i is continuous on $[r, \bar{p}]$. Hence, if $\pi_i(p) < w_i(r)$ at $p \in [r, \bar{p}]$, then F_i is constant on a neighborhood of p .

Assume for a contradiction that $F_j(\tilde{p}) > k_j(\tilde{p})$ for some $\tilde{p} \in (r, \bar{p})$. Then, F_i is constant on a neighborhood of \tilde{p} . Define $\hat{p} = \min\{p \in [r, \bar{p}] : F_i(p) = F_i(\tilde{p})\}$. (By continuity of F_i on $[r, \bar{p}]$, the minimum is well defined.) Then, $F_i(p) = F_i(\tilde{p}) \geq k_i(\tilde{p}) > k_i(p)$ for every $p \in [\hat{p}, \tilde{p})$. It follows that F_j is also constant on $[\hat{p}, \tilde{p})$. By continuity of F_j , this implies that $F_j(\hat{p}) = F_j(\tilde{p}) \geq k_j(\tilde{p}) > k_j(\hat{p})$. Hence, $\pi_i(\hat{p}) < w_i(r)$. Therefore, there exists $\eta > 0$ such that F_i is constant on $(\hat{p} - \eta, \hat{p} + \eta)$. This, however, contradicts the definition of \hat{p} . Hence, $F_j(p) = k_j(p)$ for every $j \in \{1, 2\}$ and $p \in [r, \bar{p}]$. \square

Combining Lemmas B and D and the analysis in the main text, we obtain Proposition 2. We now provide a complete characterization of the set of equilibria in the non-generic case:

Proposition A. Consider the constrained game of a non-generic ($r_1 = r_2$) all-pay oligopoly model. If $\bar{p}_1 = \bar{p}_2$, then:

- If $\bar{p} < p^0$ or $D(p^0) > 0$, then (F_1, F_2) is an equilibrium profile of CDF's if and only if there exists $(\bar{F}_1, \bar{F}_2) \in [k_1(\bar{p}), 1] \times \{k_2(\bar{p})\} \cup \{k_1(\bar{p})\} \times [k_2(\bar{p}), 1]$ such that, for $i = 1, 2$, $F_i(p) = k_i(p)$ if $p \in [r, \bar{p})$ and $F_i(p) = \bar{F}_i$ if $p \in [\bar{p}, p^0]$.
- If instead $\bar{p} = p^0$ and $D(p^0) = 0$, then the equilibrium is unique and given by $F_i(p) = k_i(p)$ for all $p \in [r, p^0]$ ($i = 1, 2$), where $k_i(p^0) \equiv \lim_{p \uparrow p^0} k_i(p)$.

If instead $\bar{p}_1 < \bar{p}_2$, then:

- If $\bar{p}_2 < p^0$ or $D(p^0) > 0$, then (F_1, F_2) is an equilibrium profile of CDF's if and only if there exists $(\bar{F}_1, \bar{F}_2) \in [k_1(\bar{p}_1), 1] \times \{k_2(\bar{p}_2)\} \cup \{k_1(\bar{p}_1)\} \times [k_2(\bar{p}_2), 1]$ such that

$$F_1(p) = \begin{cases} k_1(p) & \text{if } p \in [r, \bar{p}_1), \\ k_1(\bar{p}_1) & \text{if } p \in [\bar{p}_1, \bar{p}_2), \\ \bar{F}_1 & \text{if } p \in [\bar{p}_2, p^0], \end{cases} \quad \text{and } F_2(p) = \begin{cases} k_2(p) & \text{if } p \in [r, \bar{p}_1), \\ \bar{F}_2 & \text{if } p \in [\bar{p}_1, p^0]. \end{cases}$$

- If instead $\bar{p}_2 = p^0$ and $D(p^0) = 0$, then (F_1, F_2) is an equilibrium profile of CDF's if and only if there exists $\bar{F}_2 \in [\lim_{p \uparrow p^0} k_2(p), 1]$ such that

$$F_1(p) = \begin{cases} k_1(p) & \text{if } p \in [r, \bar{p}_1), \\ k_1(\bar{p}_1) & \text{if } p \in [\bar{p}_1, p^0], \end{cases} \quad \text{and } F_2(p) = \begin{cases} k_2(p) & \text{if } p \in [r, \bar{p}_1), \\ \bar{F}_2 & \text{if } p \in [\bar{p}_1, p^0]. \end{cases}$$

Proof. The proof follows the same development as the proof of Proposition 2. Lemma D pins down the equilibrium CDF's on $[r, \bar{p})$. The mass that remains can then be distributed over \bar{p}_1, \bar{p}_2 , and/or the firms' outside options as described in the main text. \square

B Equilibrium Inventory Choice in the Unconstrained Game

The goal of this section is to prove Proposition 4. We introduce notation and state preliminary lemmas in Section B.1. We then prove the proposition in Section B.2. In the following, we fix an all-pay oligopoly game satisfying the assumptions made at the beginning of Section 3.1.

B.1 Technical preliminaries

Let $i \neq j$ in $\{1, 2\}$. Let $Z_i(p_i, p_j, \tilde{q}_i, \tilde{q}_j)$ denote the demand for firm i 's product in the contested segment when prices are $(p_1, p_2) \in [0, p^0]^2$ and inventory levels are $(\tilde{q}_1, \tilde{q}_2) \in \mathbb{R}_+^2$.

(net of what firms i and j are selling in their captive segments). If $p_i < p_j$, then $Z_i = \min\{\tilde{q}_i, (1 - \mu_i - \mu_j)D(p_i)\}$. If instead $p_i > p_j$, then

$$Z_i = \begin{cases} \min(\tilde{q}_i, \max((1 - \mu_i - \mu_j)D(p_i) - \tilde{q}_j, 0)) & \text{under efficient rationing,} \\ \min\left(\tilde{q}_i, \max\left(\frac{D(p_i)}{D(p_j)}((1 - \mu_i - \mu_j)D(p_j) - \tilde{q}_j), 0\right)\right) & \text{under random rationing.} \end{cases}$$

Finally, if $p_i = p_j$, then, using the same-price fair-share rule,

$$Z_i = \min\left(\tilde{q}_i, \max\left(\frac{1}{2}(1 - \mu_i - \mu_j)D(p_i), (1 - \mu_i - \mu_j)D(p_i) - \tilde{q}_j\right)\right).$$

Note that, no matter whether rationing is random or efficient, Z_i is non-decreasing in q_i .

Next, we simplify the action sets by removing redundant and/or strictly dominated pure strategies. Note that, if firm i does not pay the advertising cost, then it is optimal for that firm to set $p_i = p_i^m$ and $q_i = \mu_i D(p_i^m)$.³⁰ Denote this strategy by $(p_i^m, \mu_i D(p_i^m))$. Next, we remove all the pure strategies in which firm i pays the advertising cost and chooses (p_i, q_i) such that $q_i \leq \mu_i D(p_i)$, because those strategies are either strictly dominated by $(p_i^m, \mu_i D(p_i^m))$, or outcome-equivalent to $(p_i^m, \mu_i D(p_i^m))$. Finally, we remove all the pure strategies in which firm i is pricing below cost or choosing (p_i, q_i) such that $q_i > (1 - \mu_j)D(p_i)$, as those strategies are strictly dominated.

This leaves us with the following set of pure strategies for firm i :

$$\mathcal{A}_i = \underbrace{\{(p_i, q_i) \in [c_i, p^0] \times \mathbb{R}_+ : \mu_i D(p_i) < q_i \leq (1 - \mu_j)D(p_i)\}}_{\equiv \mathcal{A}'_i} \cup \{(p_i^m, \mu_i D(p_i^m))\}.$$

A mixed strategy for player i is a probability measure σ_i over \mathcal{A}_i (\mathcal{A}_i is endowed with the σ -algebra of Borel sets). We decompose σ_i into σ'_i , a finite measure over \mathcal{A}'_i , and τ_i , a mass point on $(p_i^m, \mu_i D(p_i^m))$. We introduce the following notation: φ_i is the marginal on prices of σ'_i ; If $\varphi_i(\{p_i\}) > 0$, then we let $\chi_i(q_i|p_i)$ be the conditional probability distribution (over $(\mu_i D(p_i), (1 - \mu_j)D(p_i))$) of q_i given p_i .

Let $\pi_i(p_i, q_i, \sigma_j)$ be the expected profit received by firm i when it chooses a price-inventory pair $(p_i, q_i) \in \mathcal{A}'_i$ and firm j mixes according to σ_j . Let $\Delta_i(p_i, q_i, \sigma_j)$ denote the expected demand received by firm i in the contested segment given $(p_i, q_i) \in \mathcal{A}'_i$ and σ_j . In general, we have that

$$\Delta_i(p_i, q_i, \sigma_j) = \int_{\mathcal{A}_j} Z_i(p_i, p_j, q_i - \mu_i D(p_i), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j).$$

Note that, if $\varphi_j(\{p_i\}) = 0$, then

³⁰If $\mu_i = 0$, then it does not matter what price firm i sets, as long as $q_i = 0$. We assume without loss of generality that firm i sets $p_i = p_i^m$ in that case.

$$\begin{aligned}\Delta_i(p_i, q_i, \sigma_j) &= (\varphi_j([p_i, p^0]) + \tau_j) q_i \\ &\quad + \int_{\substack{c_j \leq p_j < p_i \\ \mu_j D(p_j) < q_j \leq (1-\mu_i)D(p_j)}} Z_i(p_i, p_j, q_i - \mu_i D(p_i), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j).\end{aligned}$$

Moreover, $\pi_i(p_i, q_i, \sigma_j) = (p_i - \alpha_i c_i) (\Delta_i(p_i, q_i, \sigma_j) + \mu_i D(p_i)) - (1 - \alpha_i) c_i q_i - A_i$.

Next, we define firm i 's expected demand when it sets a price "just below" p_i and an inventory of q_i (with $(p_i, q_i) \in \mathcal{A}'_i$):

$$\begin{aligned}\Delta_i^-(p_i, q_i, \sigma_j) &= (\varphi_j([p_i, p^0]) + \tau_j) q_i \\ &\quad + \int_{\substack{c_j \leq p_j < p_i \\ \mu_j D(p_j) < q_j \leq (1-\mu_i)D(p_j)}} Z_i(p_i, p_j, q_i - \mu_i D(p_i), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j).\end{aligned}$$

We now show that Δ_i^- is indeed firm i 's expected demand when it prices just below p_i :

Lemma E. *For every $(\hat{p}_i, \hat{q}_i) \in \mathcal{A}'_i$, for every mixed strategy σ_j for firm j , $\Delta_i(p_i, \hat{q}_i, \sigma_j) \xrightarrow[p_i \uparrow \hat{p}_i]{} \Delta_i^-(\hat{p}_i, \hat{q}_i, \sigma_j)$.*

Proof. Let $(p^n)_{n \geq 1}$ be a strictly increasing sequence such that $p^n \xrightarrow[n \rightarrow \infty]{} \hat{p}_i$. For every n ,

$$\Delta_i(p^n, \hat{q}_i, \sigma_j) = (\varphi_j([p^n, p^0]) + \tau_j) (\hat{q}_i - \mu_i D(p_i)) + \int_{\substack{c_j \leq p_j < \hat{p}_i \\ \mu_j D(p_j) < q_j \leq (1-\mu_i)D(p_j)}} Z^n(p_j, q_j) d\sigma'_j(p_j, q_j), \quad (3)$$

where

$$Z^n(p_j, q_j) = \mathbf{1}_{p_j \leq p^n} Z_i(p^n, p_j, \hat{q}_i - \mu_i D(p^n), q_j - \mu_j D(p_j))$$

for all $(p_j, q_j) \in \{(p'_j, q'_j) : 0 \leq p'_j < \hat{p}_i \text{ and } \mu_j D(p'_j) < q'_j \leq (1 - \mu_i)D(p'_j)\}$.

Note that, since the sequence of events $((p^n, p^0))_{n \geq 1}$ is non-increasing, we have that $\lim_{n \rightarrow \infty} \varphi_j([p^n, p^0]) = \varphi_j(\bigcap_{n \geq 1} [p^n, p^0]) = \varphi_j([\hat{p}_i, p^0])$.

Next, we turn our attention to the term in equation (3). The sequence of σ_j -integrable functions $(Z^n)_{n \geq 1}$ is non-negative and bounded above by the constant function $D(c_i)$, which is also σ_j -integrable. Moreover, $(Z^n)_{n \geq 1}$ converges pointwise to the function

$$\hat{Z}_i(p_j, q_j) = Z_i(\hat{p}_i, p_j, \hat{q}_i - \mu_i D(\hat{p}_i), q_j - \mu_j D(p_j)).$$

By Lebesgue's dominated convergence theorem, it follows that

$$\lim_{n \rightarrow \infty} \int_{\substack{c_j \leq p_j < \hat{p}_i \\ \mu_j D(p_j) < q_j \leq (1-\mu_i)D(p_j)}} Z^n(p_j, q_j) d\sigma_j(p_j, q_j) = \int_{\substack{c_j \leq p_j < \hat{p}_i \\ \mu_j D(p_j) < q_j \leq (1-\mu_i)D(p_j)}} \hat{Z}_i(p_j, q_j) d\sigma_j(p_j, q_j),$$

which proves the lemma. \square

Lemma E says that, no matter whether firm j has a mass point at \hat{p}_i , firm i can always secure a demand level arbitrarily close to $\Delta_i^-(\hat{p}_i, \hat{q}_i, \sigma_j)$ with a price arbitrarily close to \hat{p}_i .

For every $(p_i, q_i) \in \mathcal{A}'_i$, for every mixed strategy σ_j for firm j , define

$$\pi_i^- (p_i, q_i, \sigma_j) = (p_i - \alpha_i c_i) (\Delta_i^- (p_i, q_i, \sigma_j) + \mu_i D(p_i)) - (1 - \alpha_i) c_i q_i - A_i.$$

The following result is an immediate implication of Lemma E:

Lemma F. *Suppose that (σ_1, σ_2) is a mixed-strategy Nash equilibrium, and let $\bar{\pi}_i$ be firm i 's expected profit in that equilibrium. Then, for every $(\hat{p}_i, \hat{q}_i) \in \mathcal{A}'_i$, $\bar{\pi}_i \geq \pi_i^- (\hat{p}_i, \hat{q}_i, \sigma_j)$. Moreover, if $(\hat{p}_i, \hat{q}_i) \in \mathcal{A}'_i$ and $\bar{\pi}_i = \pi_i (\hat{p}_i, \hat{q}_i, \sigma_j)$, then $\pi_i (\hat{p}_i, \hat{q}_i, \sigma_j) = \pi_i^- (\hat{p}_i, \hat{q}_i, \sigma_j)$.*

B.2 Proof of Proposition 4

Proof. Suppose $A_1, A_2 > 0$. Let (σ_1, σ_2) be a Nash equilibrium of the all-pay oligopoly model. Let $\bar{\pi}_i$ denote firm i 's expected profit in that equilibrium. Clearly, for every firm i , $\bar{\pi}_i \geq o_i$.

For every $p \in [0, p^0]$ and $i \in \{1, 2\}$, define

$$S_i(p) = \{(p', q') \in [0, p] \times \mathbb{R}_+ : \mu_i D(p') < q' < (1 - \mu_j) D(p')\},$$

and $\phi_i(p) = \sigma_i(S_i(p))$. Clearly, ϕ_i is non-decreasing, and $\phi_i(p) = 0$ for p sufficiently low.

Assume for a contradiction that $\phi_i(p) > 0$ for some firm i and some price $p \in [0, p^0]$. Define

$$\hat{p} = \inf \{p \in [0, p^0] : \exists i \in \{1, 2\}, \phi_i(p) > 0\}.$$

We first argue that, for every $i \in \{1, 2\}$, $\phi_i(\hat{p}) = 0$. Assume for a contradiction that $\phi_i(\hat{p}) > 0$ for some firm i . We claim that $\varphi_i(\{\hat{p}\}) > 0$. To see this, let $(p^n)_{n \geq 1}$ be a strictly increasing sequence of prices that converges to \hat{p} . Note that

$$\{\hat{p}\} \times (\mu_i D(\hat{p}), (1 - \mu_j) D(\hat{p})) = S_i(\hat{p}) \setminus \bigcup_{n \geq 1} S_i(p^n).$$

Since, by definition of \hat{p} , $\sigma_i(S_i(p^n)) = 0$ for every n , it follows that

$$\varphi_i(\{\hat{p}\}) \geq \sigma_i(\{\hat{p}\} \times (\mu_i D(\hat{p}), (1 - \mu_j) D(\hat{p}))) = \sigma_i(S_i(\hat{p})) = \phi_i(\hat{p}) > 0.$$

Therefore, $\varphi_i(\{\hat{p}\}) > 0$, and $\chi_i(\cdot | \hat{p})$ is well defined, and does not put full weight on $q = (1 - \mu_j) D(\hat{p})$. In particular, there exists $\mu_i D(\hat{p}) < \hat{q} < (1 - \mu_j) D(\hat{p})$ such that $\bar{\pi}_i = \pi_i(\hat{p}, \hat{q}, \sigma_j)$. By Lemma F, it follows that $\pi_i(\hat{p}, \hat{q}, \sigma_j) = \pi_i^- (\hat{p}, \hat{q}, \sigma_j) = \bar{\pi}_i$. Therefore,

$$\bar{\pi}_i = \left((\hat{p} - \alpha_i c_i) (\varphi_j([\hat{p}, p^0]) + \tau_j) - (1 - \alpha_i) c_i \right) (\hat{q} - \mu_i D(\hat{p})) + \mu_i (\hat{p} - c_i) D(\hat{p}) - A_i,$$

where we have used the fact that firm i receives no residual demand in the contested segment when firm j prices strictly below \hat{p} . Since $\bar{\pi}_i \geq o_i = \mu_i (p_i^m - c_i) D(p_i^m)$, we have that

$$\begin{aligned} & \left((\hat{p} - \alpha_i c_i)(\varphi_j([\hat{p}, p^0]) + \tau_j) - (1 - \alpha_i)c_i \right) (\hat{q} - \mu_i D(\hat{p})) \\ & \geq \mu_i \left((p_i^m - c_i)D(p_i^m) - (\hat{p} - c_i)D(\hat{p}) \right) + A_i > 0, \end{aligned}$$

implying that $(\hat{p} - \alpha_i c_i)(\varphi_j([\hat{p}, p^0]) + \tau_j) - (1 - \alpha_i)c_i > 0$. Therefore, by Lemma F,

$$\begin{aligned} \bar{\pi}_i & \geq \pi_i^-(\hat{p}, (1 - \mu_j)D(\hat{p}), \sigma_j), \\ & = \left((\hat{p} - \alpha_i c_i)(\varphi_j([\hat{p}, p^0]) + \tau_j) - (1 - \alpha_i)c_i \right) (1 - \mu_j - \mu_i)D(\hat{p}) + \mu_i(\hat{p} - c_i)D(\hat{p}) - A_i, \\ & > \left((\hat{p} - \alpha_i c_i)(\varphi_j([\hat{p}, p^0]) + \tau_j) - (1 - \alpha_i)c_i \right) (\hat{q} - \mu_i D(\hat{p})) + \mu_i(\hat{p} - c_i)D(\hat{p}) - A_i, \\ & = \pi_i(\hat{p}, \hat{q}, \sigma_j) = \bar{\pi}_i, \end{aligned}$$

which is a contradiction. Hence, $\phi_i(\hat{p}) = 0$ for $i = 1, 2$.

By definition of \hat{p} , there exist a firm $i \in \{1, 2\}$ and a strictly decreasing sequence of prices $(p^n)_{n \geq 1}$ such that, $p^n \xrightarrow[n \rightarrow \infty]{} \hat{p}$, and, $\phi_i(p^n) > 0$ for every n . Since $\phi_i(\hat{p}) = 0$, this implies the existence of a sequence of price-inventory pairs $(p^n, q^n)_{n \geq 1}$ such that $p^n \xrightarrow[n \rightarrow \infty]{} \hat{p}$, and for every n , $p^n > \hat{p}$, $\mu_i D(p^n) < q^n < (1 - \mu_j)D(p^n)$, and $\bar{\pi}_i = \pi_i(p^n, q^n, \sigma_j)$. Moreover, by Lemma F, for every n ,

$$\begin{aligned} \bar{\pi}_i & = \pi_i^-(p^n, q^n, \sigma_j), \\ & = \mu_i(p^n - c_i)D(p^n) - A_i + \left((p^n - \alpha_i c_i)(\varphi_j([p^n, p^0]) + \tau_j) - (1 - \alpha_i)c_i \right) (q^n - \mu_i D(p^n)) \\ & \quad + (p^n - \alpha_i c_i) \int_{\substack{\hat{p} < p_j < p^n \\ \mu_j D(p_j) < q_j \leq (1 - \mu_i)D(p_j)}} Z_i(p^n, p_j, q^n - \mu_i D(p^n), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j).^{31} \end{aligned}$$

Lemma F also guarantees that, for every n and $q \in (\mu_i D(p^n), (1 - \mu_j)D(p^n)]$,

$$\begin{aligned} \bar{\pi}_i & \geq \pi_i^-(p^n, q, \sigma_j), \\ & = \mu_i(p^n - c_i)D(p^n) - A_i + \left((p^n - \alpha_i c_i)(\varphi_j([p^n, p^0]) + \tau_j) - (1 - \alpha_i)c_i \right) (q - \mu_i D(p^n)) \\ & \quad + (p^n - \alpha_i c_i) \int_{\substack{\hat{p} < p_j < p^n \\ \mu_j D(p_j) < q_j \leq (1 - \mu_i)D(p_j)}} Z_i(p^n, p_j, q - \mu_i D(p^n), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j). \end{aligned}$$

Note that the integral term in the above expression is non-decreasing in q . If $(p^n - \alpha_i c_i)(\varphi_j([p^n, p^0]) + \tau_j) - (1 - \alpha_i)c_i$ were strictly positive, then $\pi_i^-(p^n, q, \sigma_j)$ would be strictly increasing in q on the interval $(\mu_i D(p^n), (1 - \mu_j)D(p^n)]$. We would then obtain the following

³¹The reason why we can integrate over (\hat{p}, p^n) instead of $[\hat{p}, p^n)$ is the following. Either $\varphi_j(\{\hat{p}\}) = \sigma_j(\{\hat{p}\} \times (\mu_j D(\hat{p}), (1 - \mu_i)D(\hat{p}))) = 0$, and that set can be removed from the domain of integration. Or $\varphi_j(\{\hat{p}\}) > 0$, and the above analysis guarantees that $\chi_j(\cdot | \hat{p})$ puts full weight on $q_j = (1 - \mu_i)D(\hat{p})$.

contradiction:

$$\bar{\pi}_i = \pi_i^-(p^n, q^n, \sigma_j) < \pi_i^-(p^n, (1 - \mu_j)D(p^n), \sigma_j) \leq \bar{\pi}_i.$$

Therefore, $(p^n - \alpha_i c_i) (\varphi_j([p^n, p^0]) + \tau_j) - (1 - \alpha_i)c_i \leq 0$ for every n . Note, however, that

$$\begin{aligned} & \int_{\substack{\hat{p} < p_j < p^n \\ \mu_j D(p_j) < q_j \leq (1 - \mu_i) D(p_j)}} Z_i(p^n, p_j, q^n - \mu_i D(p^n), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j) \\ & \leq \int_{\substack{\hat{p} < p_j < p^n \\ \mu_j D(p_j) < q_j \leq (1 - \mu_i) D(p_j)}} D(\hat{p}) d\sigma_j(p_j, q_j) = \varphi_j(\hat{p}, p^n) D(\hat{p}) \xrightarrow{n \rightarrow \infty} 0. \end{aligned}$$

We obtain the following contradiction:

$$\begin{aligned} \bar{\pi}_i & \leq \mu_i(p^n - c_i)D(p^n) - A_i \\ & \quad + (p^n - \alpha_i c_i) \int_{\substack{\hat{p} < p_j < p^n \\ \mu_j D(p_j) < q_j \leq (1 - \mu_i) D(p_j)}} Z_i(p^n, p_j, q^n - \mu_i D(p^n), q_j - \mu_j D(p_j)) d\sigma_j(p_j, q_j), \\ & \xrightarrow{n \rightarrow \infty} \mu_i(\hat{p} - c_i)D(\hat{p}) - A_i < o_i \leq \bar{\pi}_i. \end{aligned} \quad \square$$

C Convergence Results Under Complete Information

Throughout this section, we assume that D is continuous at p^0 . In Section C.1, we provide an alternative formulation of the constrained game, which will be useful to derive our convergence results. Section C.2 establishes the continuity of p_i^m , r_i , \bar{p}_i , and k_i in the parameters of the model. Propositions 5 and 6 are proved in Sections C.3 and C.4, respectively. The Bertrand-without-fudge model is discussed in Section C.5.

C.1 An Alternative Formulation

Fix a vector of parameters $(c_1, c_2, \alpha_1, \alpha_2, \mu_1, \mu_2, A_1, A_2) \in (0, p^0)^2 \times [0, 1]^2 \times [0, 1]^2 \times \mathbb{R}_+^2$ such that $\mu_1 + \mu_2 < 1$. Note that we allow recoverability parameters to be equal to 1, which will be useful to prove Proposition 6. For every $i \in \{1, 2\}$, put $\mathcal{A}_i = \{0, 1\} \times \mathbb{R}_+^2$. A typical element of \mathcal{A}_i is (a_i, p_i, q_i) , where a_i is equal to 1 if firm i pays the advertising cost and to 0 otherwise, and (p_i, q_i) is the price-inventory pair chosen by firm i . Let $\pi_i(a_i, p_i, q_i, a_j, p_j, q_j)$ denote firm i 's payoff in the all-pay oligopoly model with parameters $(c_1, c_2, \alpha_1, \alpha_2, \mu_1, \mu_2, A_1, A_2)$, when firm ι chooses $(a_\iota, p_\iota, q_\iota) \in \mathcal{A}_\iota$ ($\iota \in \{1, 2\}$). The normal-form game associated with this all-pay oligopoly model is $\mathcal{G} = (\{1, 2\}, (\mathcal{A}_1, \mathcal{A}_2), (\pi_1, \pi_2))$.

The constrained game studied in Section 3.1 can be formally defined as follows. For $i = 1, 2$, let $\hat{\mathcal{A}}_i = [0, p^0] \cup \{\text{out}\}$ and

$$\hat{\psi}_i : \hat{p} \in \hat{\mathcal{A}}_i \mapsto \begin{cases} (1, \hat{p}, (1 - \mu_j)D(\hat{p})) & \text{if } \hat{p} \in [0, p^0], \\ (0, p_i^m, \mu_i D(p_i^m)) & \text{if } \hat{p} = \text{out}. \end{cases}$$

Define

$$\hat{\pi}_i(\hat{p}_i, \hat{p}_j) = \pi_i \left(\hat{\psi}_i(\hat{p}_i), \hat{\psi}_j(\hat{p}_j) \right), \quad i, j = 1, 2, \quad i \neq j, \quad (\hat{p}_i, \hat{p}_j) \in \hat{\mathcal{A}}_i \times \hat{\mathcal{A}}_j.$$

The constrained game is the normal-form game $\hat{\mathcal{G}} = \left(\{1, 2\}, \left(\hat{\mathcal{A}}_1, \hat{\mathcal{A}}_2 \right), (\hat{\pi}_1, \hat{\pi}_2) \right)$.

As discussed in footnote 17, our convergence results turn out to be easier to prove in an alternative formulation of the constrained game, which we now define formally. For $i = 1, 2$, let $\tilde{\mathcal{A}}_i = [0, p^0]$, and

$$\tilde{\psi}_i : \tilde{p} \in \tilde{\mathcal{A}}_i \mapsto \begin{cases} (1, \tilde{p}, (1 - \mu_j)D(\tilde{p})) & \text{if } \tilde{p} < p^0, \\ (0, p_i^m, \mu_i D(p_i^m)) & \text{if } \tilde{p} = p^0. \end{cases}$$

Define

$$\tilde{\pi}_i(\tilde{p}_i, \tilde{p}_j) = \pi_i \left(\tilde{\psi}_i(\tilde{p}_i), \tilde{\psi}_j(\tilde{p}_j) \right), \quad i, j = 1, 2, \quad i \neq j, \quad (\tilde{p}_i, \tilde{p}_j) \in [0, p^0]^2.$$

The auxiliary game is $\tilde{\mathcal{G}} = \left(\{1, 2\}, \left(\tilde{\mathcal{A}}_1, \tilde{\mathcal{A}}_2 \right), (\tilde{\pi}_1, \tilde{\pi}_2) \right)$.

The constrained game and the auxiliary game differ in only two ways: A firm's action set in the constrained game contains the additional element 'out'; In the constrained game, choosing $\hat{p}_i = p^0$ means "paying the advertising cost, setting a price of p^0 , and sourcing no inventory," whereas in the auxiliary game, such a strategy means "not paying the advertising cost, setting one's monopoly price, and sourcing enough inventory to supply one's captive consumers." Note, however, that in the constrained game, the pure strategy $\hat{p}_i = p^0$ is either payoff-equivalent to $\hat{p}_i = \text{out}$ (if $A_i = \mu_i = 0$), or strictly dominated by $\hat{p}_i = \text{out}$. Hence, firms put no mass on p^0 in equilibrium (recall the convention we adopted in footnote 13). For all intents and purposes, the auxiliary game is therefore equivalent to the constrained game.

Recall that a mixed-strategy equilibrium of the constrained game was defined as a pair of CDF's of finite measures (\hat{F}_1, \hat{F}_2) over $[0, p^0]$, with the understanding that $1 - \hat{F}_i(p^0)$ is the probability that firm i sets $\hat{p}_i = \text{out}$. Clearly, there is a one-to-one mapping between the equilibria of the constrained game and those of the auxiliary game. For a given equilibrium (\hat{F}_1, \hat{F}_2) of the constrained game, the associated pair of equilibrium CDF's $(\tilde{F}_1, \tilde{F}_2)$ in the auxiliary game is:

$$\tilde{F}_i(p) = \begin{cases} \hat{F}_i(p) & \text{if } p < p^0, \\ 1 & \text{if } p = p^0. \end{cases}$$

In the following, we prove our convergence results in the auxiliary game, and remove the tildes to ease notation.

C.2 Preliminaries

In this section, we show that p_i^m , r_i , \bar{p}_i , and k_i continue to be well-defined when α_1 and/or α_2 is equal to 1, and we study how these equilibrium objects are affected by small changes in the parameter vector.

The set of admissible parameter vectors is

$$\Gamma'' = \left\{ (c_1, c_2, \alpha_1, \alpha_2, \mu_1, \mu_2, A_1, A_2) \in (0, p^0)^2 \times [0, 1]^2 \times [0, 1]^2 \times \mathbb{R}_+^2 : \mu_1 + \mu_2 < 1 \right\}.$$

In the following, we denote a typical parameter vector by $\gamma \in \Gamma$, with the understanding that c_1 is the first component of γ , c_2 is the second component, etc.

We now make explicit the dependence of a firm's winning and losing functions on the parameters of the model by writing

$$\begin{aligned} W_i(p; \gamma) &= (1 - \mu_j)(p - c_i)D(p) - A_i, \\ L_i(p; \gamma) &= \mu_i(p - c_i)D(p) - (1 - \mu_i - \mu_j)(1 - \alpha_i)c_iD(p) - A_i, \end{aligned}$$

for every $i, j \in \{1, 2\}$ such that $i \neq j$, $p \in [0, p^0]$, and $\gamma \in \Gamma''$. Note that W_i and L_i are both continuous.

Monopoly prices and outside options. For every $\gamma \in \Gamma''$, let $P_i^m(\gamma)$ be the unique solution of the maximization problem $\max_{p \in [0, p^0]} (p - c_i)D(p)$. The theorem of the maximum guarantees that P_i^m is continuous. Firm i 's outside option is:

$$O_i(\gamma) = \mu_i(P_i^m(\gamma) - c_i)D(P_i^m(\gamma)),$$

which is also a continuous function.

As in Section 3.1, we restrict attention to parameter vectors that belong to the set

$$\Gamma' = \left\{ \gamma \in \Gamma'' : W_i(P_i^m(\gamma); \gamma) > O_i(\gamma), \forall i \in \{1, 2\} \right\}.$$

By continuity of W_i , P_i^m , and O_i , Γ' is open relative to Γ .

Reaches. For every $\gamma \in \Gamma'$ and $i \in \{1, 2\}$, define $R_i(\gamma)$ as the unique $p \in [0, P_i^m(\gamma)]$ such that $W_i(p; \gamma) = O_i(\gamma)$. The continuity of W_i and O_i implies that R_i is continuous.³² Therefore, $R = \max\{R_1, R_2\}$ is continuous as well.

As in Section 3.1, we further restrict attention to parameter vectors that belong to the set

$$\Gamma = \left\{ \gamma \in \Gamma' : R_i(\gamma) < P_j^m(\gamma), \forall i, j \in \{1, 2\} \text{ s.t. } i \neq j \right\}.$$

Again, the continuity of R_i and P_j^m implies that Γ is open, relative to Γ'' .

³²Assume for a contradiction that R_i is not continuous. There exist an $\varepsilon > 0$ and a sequence $(\gamma^n)_{n \geq 1}$ over Γ' such that $\gamma^n \xrightarrow[n \rightarrow \infty]{} \gamma \in \Gamma'$, but $|R_i(\gamma^n) - R_i(\gamma)| > \varepsilon$ for every n . Since $(R_i(\gamma^n))_{n \geq 1}$ is bounded, we can extract a subsequence $(R_i(\gamma'^n))_{n \geq 1}$ that converges to some $r \in [0, p^0]$. Clearly, $r \neq R_i(\gamma)$. Since $R_i(\gamma'^n) \leq P_i^m(\gamma'^n)$ for every n , the continuity of P_i^m implies that $r \leq P_i^m(\gamma)$. Moreover, since $W_i(R_i(\gamma'^n); \gamma'^n) = O_i(\gamma'^n)$, the continuity of W_i , R_i and O_i implies that $W_i(r; \gamma) = O_i(\gamma)$. By uniqueness of $R_i(\gamma)$, it follows that $r = R_i(\gamma)$, a contradiction.

The k functions and the \bar{p} cutoffs. For every $\gamma \in \Gamma$ such that $\alpha_j < 1$, define

$$K_i(p; \gamma) = \begin{cases} 0 & \text{if } p \in [0, R(\gamma)], \\ \frac{W_j(p; \gamma) - W_j(R(\gamma); \gamma)}{W_j(p; \gamma) - L_j(p; \gamma)} & \text{if } p \in (R(\gamma), p^0). \end{cases}$$

Note that, for every $p \in (R(\gamma), p^0)$,

$$K_i(p; \gamma) = \frac{1 - \mu_i}{1 - \mu_i - \mu_j} \left(\frac{p - c_j}{p - \alpha_j c_j} - \frac{(R(\gamma) - c_j)D(R(\gamma))}{(p - \alpha_j c_j)D(p)} \right). \quad (4)$$

If $R(\gamma) > c_j$, then $K_i(\cdot, \gamma)$ is single-peaked and achieves its global maximum at some $\bar{P}_i(\gamma) \in (R(\gamma), p^0)$, as shown in Lemma C. If instead $R(\gamma) = c_j$, then $\mu_j = 0$ and $K_i(p; \gamma) = \frac{p - c_j}{p - \alpha_j c_j}$ for all $p \in (R(\gamma), 1)$. Hence, either $\alpha_j < 1$ and $K_i(\cdot; \gamma)$ is strictly increasing on $(R(\gamma), p^0)$, or $\alpha_j = 1$ and $K_i(\cdot; \gamma)$ is constant and equal to 1 on $(R(\gamma), p^0)$. In the former case, we set $\bar{P}_i(\gamma) = p^0$. In the latter case, we do not define $\bar{P}_i(\gamma)$. The domain of \bar{P}_i is therefore

$$\bar{\Gamma}_i = \left\{ \gamma \in \Gamma : \alpha_j < 1 \text{ or } R(\gamma) > c_j \right\},$$

which is an open set. Note that $K_i(\cdot; \gamma)$ is continuous on $[0, p^0)$ whenever $\gamma \in \bar{\Gamma}_i$.

Convergence properties of K_i . Let $(\gamma^n)_{n \geq 1}$ be a sequence over Γ that converges to some $\gamma \in \Gamma$. We now argue that $(K_i(\cdot; \gamma^n))_{n \geq 1}$ converges pointwise to $K_i(\cdot; \gamma)$ on $[0, p^0) \setminus \{R(\gamma)\}$. To see this, let $p \in [0, p^0)$. Suppose first that $p < R(\gamma)$. Since R is continuous, we have that $p < R(\gamma^n)$ for n high enough. Hence, $K_i(p; \gamma^n) = 0$ for p high enough, and $\lim_{n \rightarrow \infty} K_i(p; \gamma^n) = 0 = K_i(p; \gamma)$. Next, suppose that $p > R(\gamma)$. Then, by continuity of R , $p > R(\gamma^n)$ for n high enough. Taking limits in equation (4), we obtain that $\lim_{n \rightarrow \infty} K_i(p; \gamma^n) = K_i(p; \gamma)$.

Continuity of \bar{P}_i . We now show that \bar{P}_i is continuous on its domain $\bar{\Gamma}_i$. Let $(\gamma^n)_{n \geq 1}$ be a sequence over $\bar{\Gamma}_i$ that converges to some $\gamma \in \bar{\Gamma}_i$. Let $R(\gamma) < \hat{p} < \bar{P}_i(\gamma)$. We show that $\bar{P}_i(\gamma^n) > \hat{p}$ for n sufficiently high. Let $\check{p} \in (\hat{p}, \bar{P}_i(\gamma))$. Then, $K_i(\check{p}; \gamma) > K_i(\hat{p}; \gamma)$. Since $\lim_{n \rightarrow \infty} K_i(\check{p}; \gamma^n) = K_i(\check{p}; \gamma)$ and $\lim_{n \rightarrow \infty} K_i(\hat{p}; \gamma^n) = K_i(\hat{p}; \gamma)$, it follows that $K_i(\check{p}; \gamma^n) > K_i(\hat{p}; \gamma^n)$ for n high enough. The uni-modality of $K_i(\cdot; \gamma^n)$ implies that $\hat{p} < \bar{P}_i(\gamma^n)$ for n high enough. The same line of reasoning implies that, for every $p > \bar{P}_i(\gamma)$, there exists $N \geq 1$ such that $p > \bar{P}_i(\gamma^n)$ for every $n \geq N$. It follows that $\bar{P}_i(\gamma^n) \xrightarrow[n \rightarrow \infty]{} \bar{P}_i(\gamma)$, and that \bar{P}_i is continuous.

More on the convergence properties of K_i . Let $(\gamma^n)_{n \geq 1}$ be a sequence over Γ that converges to some $\gamma \in \Gamma$. We now show that, if $\gamma \in \bar{\Gamma}_i$, then $\lim_{n \rightarrow \infty} K_i(R(\gamma); \gamma^n) = K_i(R(\gamma); \gamma) = 0$. Let $\varepsilon \in (0, K_i(\bar{P}_i(\gamma); \gamma))$. The continuity and monotonicity properties of $K_i(\cdot; \gamma)$ imply the existence of a price $p \in (R(\gamma), \bar{P}_i(\gamma))$ such that $K_i(p; \gamma) = \frac{\varepsilon}{2}$. Since $\lim_{n \rightarrow \infty} K_i(p; \gamma^n) = K_i(p; \gamma)$, we have that $K_i(p; \gamma^n) \in (0, \varepsilon)$ for n high enough. Moreover,

since $\lim_{n \rightarrow \infty} \bar{P}_i(\gamma^n) = \bar{P}_i(\gamma)$, we also have that $\bar{P}_i(\gamma^n) > p$ for n high enough. Therefore, by uni-modality of $K_i(\cdot; \gamma^n)$, $0 \leq K_i(R(\gamma); \gamma^n) < \varepsilon$ for n high enough. This proves that $\lim_{n \rightarrow \infty} K_i(R(\gamma); \gamma^n) = K_i(R(\gamma); \gamma) = 0$.

We summarize our findings in the following lemma:

Lemma G. *The following holds:*

- R , P_i^m , and R_i ($i \in \{1, 2\}$) are continuous on Γ . Moreover, \bar{P}_i ($i \in \{1, 2\}$) is continuous on $\bar{\Gamma}_i$.
- If the sequence $(\gamma^n)_{n \geq 1}$ converges to $\gamma \in \Gamma$, then, for $i = 1, 2$, $(K_i(\cdot; \gamma^n))_{n \geq 1}$ converges pointwise to $K_i(\cdot; \gamma)$ on $[0, p^0] \setminus \{R_i(\gamma)\}$. If, in addition, $\gamma \in \bar{\Gamma}_i$, then $(K_i(\cdot; \gamma^n))_{n \geq 1}$ converges pointwise to $K_i(\cdot; \gamma)$ on $[0, p^0]$.

C.3 Proof of Proposition 5

The proposition is proven using the auxiliary game of Section C.1 and the notation and results of Section C.2.

Proof. Let $(\gamma^n)_{n \geq 1}$ be a sequence that converges to $\gamma = (c, c, \alpha, \alpha, 0, 0, 0, 0)$, with $c \in (0, p^0)$ and $\alpha \in [0, 1)$. Let (F_1^*, F_2^*) be the equilibrium of the limiting game. Recall from Proposition 1 that $F_i^*(p) = K_i(p; \gamma)$ for every $p < p^0$. For every n , let $(F_1^n, F_2^n)_{n \geq 0}$ be a constrained equilibrium of the all-pay oligopoly game with parameter vector γ^n . By Propositions 2 and A, for $i \in \{1, 2\}$, $F_i^n(p) = K_i(p; \gamma^n)$ for every $p \in [0, \bar{P}_i(\gamma^n))$.

Let $p \in [0, p^0)$. Since $\gamma \in \bar{\Gamma}_i$, Lemma G implies that $\bar{P}_i(\gamma^n) \xrightarrow[n \rightarrow \infty]{} \bar{P}_i(\gamma) = p^0$. Therefore, $p < \bar{P}_i(\gamma^n)$ and $F_i^n(p) = K_i(p; \gamma^n)$ for n high enough. By Lemma G,

$$F_i^n(p) = K_i(p; \gamma^n) \xrightarrow[n \rightarrow \infty]{} K_i(p; \gamma) = F_i^*(p).$$

We have just shown that $(F_i^n)_{n \geq 0}$ converges pointwise to F_i^* at every point of continuity of F_i^* . It follows that $(F_i^n)_{n \geq 0}$ converges weakly to F_i^* . \square

C.4 Proof of Proposition 6

The proof relies on the auxiliary game of Section C.1, and uses the notation and results of Section C.2. Before proving the proposition, we first define genericity in this context: We say that a vector of parameters $\gamma \in \Gamma$ is generic if $c_1 \neq c_2$, $R_1(\gamma) \neq R_2(\gamma)$, and $R(\gamma) > c_i$ for $i = 1, 2$. (Note that this definition does not depend on (α_1, α_2) , as R_i does not depend on the value of the recoverability parameters.)

Proof. Let $(\gamma^n)_{n \geq 1}$ be a sequence that converges to a generic vector of parameters $\gamma = (c_1, c_2, 1, 1, \mu_1, \mu_2, A_1, A_2) \in \Gamma$. Suppose $R_1(\gamma) < R_2(\gamma)$ and $\alpha_1^n, \alpha_2^n < 1$ for every n . For every

n , let $(F_1^n, F_2^n)_{n \geq 0}$ be a constrained equilibrium of the all-pay oligopoly game with parameter vector γ^n .

Note that, for every $p \in (R(\gamma), p^0)$,

$$K_i(p; \gamma) = \frac{1 - \mu_i}{1 - \mu_i - \mu_j} \left(1 - \frac{(R(\gamma) - c_j)D(R(\gamma))}{(p - c_j)D(p)} \right).$$

Maximizing $K_i(\cdot; \gamma)$ is therefore equivalent to maximizing $(p - c_j)D(p)$. It follows that $\bar{P}_i(\gamma) = P_j^m(\gamma)$.

Assume first that $c_1 < c_2$, so that $\bar{P}_1(\gamma) = P_2^m(\gamma) > P_1^m(\gamma) = \bar{P}_2(\gamma)$. By Lemma G, $\lim_{n \rightarrow \infty} R_i(\gamma^n) = R_i(\gamma)$ and $\lim_{n \rightarrow \infty} \bar{P}_i(\gamma^n) = \bar{P}_i(\gamma)$ for $i = 1, 2$. Hence, for n high enough, we have that $R_1(\gamma^n) < R_2(\gamma^n)$ and $\bar{P}_2(\gamma^n) < \bar{P}_1(\gamma^n)$. By Proposition 2, for every $p < p^0$,

$$F_1^n(p) = \begin{cases} K_1(p; \gamma^n) & \text{if } p < \bar{P}_2(\gamma^n), \\ 1 & \text{otherwise,} \end{cases} \quad \text{and } F_2^n(p) = \begin{cases} K_2(p; \gamma^n) & \text{if } p < \bar{P}_2(\gamma^n), \\ K_2(\bar{P}_2(\gamma^n); \gamma^n) & \text{otherwise.} \end{cases}$$

Let $p < \bar{P}_2(\Gamma)$. Then, $p < \bar{P}_2(\gamma^n)$ for n high enough. Hence, using Lemma G, $F_1^n(p) = K_1(p; \gamma^n) \xrightarrow[n \rightarrow \infty]{} K_1(p; \gamma)$. The same line of reasoning implies that $F_1^n(p) \xrightarrow[n \rightarrow \infty]{} 1$ if $p > \bar{P}_2(\gamma)$. Hence, $(F_1^n)_{n \geq 1}$ converges pointwise to

$$F_1(p) = \begin{cases} K_1(p; \gamma) & \text{if } p < \bar{P}_2(\gamma), \\ 1 & \text{otherwise,} \end{cases}$$

at every point of continuity of F_1 . It follows that $(F_1^n)_{n \geq 1}$ converges weakly to F_1 .

Next, we turn our attention to the sequence $(F_2^n)_{n \geq 1}$. We first argue that K_2 is continuous on a neighborhood of $(\bar{P}_2(\gamma), \gamma)$. To see this, let $\varepsilon > 0$. By continuity of R , there exists a neighborhood V of γ such that $R(\tilde{\gamma}) < \bar{P}_2(\gamma) - \varepsilon$ for every $\tilde{\gamma} \in V$. Put $V' = (\bar{P}_2(\gamma) - \varepsilon, p^0) \times V$. Then, for every $(p, \tilde{\gamma}) \in V'$, $K_2(p; \tilde{\gamma})$ is given by equation (4), which is clearly continuous in $(p, \tilde{\gamma})$. Hence, K_2 is continuous on V' . Since $(\bar{P}_2(\gamma^n), \gamma^n) \xrightarrow[n \rightarrow \infty]{} (\bar{P}_2(\gamma), \gamma)$, it follows that $(\bar{P}_2(\gamma^n), \gamma^n) \in V'$ for n high enough. By continuity, it follows that $K_2(\bar{P}_2(\gamma^n); \gamma^n) \xrightarrow[n \rightarrow \infty]{} K_2(\bar{P}_2(\gamma), \gamma)$. Combining this with the argument used in the previous paragraph, we immediately obtain that $(F_2^n)_{n \geq 1}$ converges pointwise to

$$F_2(p) = \begin{cases} K_2(p; \gamma) & \text{if } p < \bar{P}_2(\gamma), \\ K_2(\bar{P}_2(\gamma), \gamma) & \text{otherwise} \end{cases}$$

on $[0, p^0] \setminus \{\bar{P}_2(\gamma)\}$.

All that is left to do now is show that $F_2^n(\bar{P}_2(\gamma)) \xrightarrow[n \rightarrow \infty]{} F_2(\bar{P}_2(\gamma))$. Partition the set of positive integers into $\mathcal{N} = \{n \geq 1 : \bar{P}_2(\gamma^n) \leq \bar{P}_2(\gamma)\}$ and $\mathcal{N}' = \{n \geq 1 : \bar{P}_2(\gamma^n) > \bar{P}_2(\gamma)\}$. If \mathcal{N} is infinite, then let ϕ be the strictly increasing bijection from $\{1, 2, \dots\}$ to \mathcal{N} . (If \mathcal{N} is

finite, there is nothing to prove.) For every $n \geq 1$,

$$F_2^{\phi(n)}(\bar{P}_2(\gamma)) = K_2(\bar{P}_2(\gamma^{\phi(n)}), \gamma^{\phi(n)}) \xrightarrow[n \rightarrow \infty]{} K_2(\bar{P}_2(\gamma), \gamma) = F_2(\bar{P}_2(\gamma)).$$

Similarly, if \mathcal{N}' is infinite, let ζ be the strictly increasing bijection from $\{1, 2, \dots\}$ to \mathcal{N}' . For every $n \geq 1$,

$$F_2^{\zeta(n)}(\bar{P}_2(\gamma)) = K_2(\bar{P}_2(\gamma), \gamma^{\zeta(n)}) \xrightarrow[n \rightarrow \infty]{} K_2(\bar{P}_2(\gamma), \gamma) = F_2(\bar{P}_2(\gamma)).$$

Therefore, $F_2^n(\bar{P}_2(\gamma)) \xrightarrow[n \rightarrow \infty]{} F_2(\bar{P}_2(\gamma))$, and $(F_2^n)_{n \geq 1}$ converges weakly to F_2 .

It is then straightforward to check that (F_1, F_2) is an equilibrium of the game with parameter vector γ .

Next, assume that $c_1 > c_2$, so that $\bar{P}_1(\gamma) = P_2^m(\gamma) < P_1^m(\gamma) = \bar{P}_2(\gamma)$. By Lemma G, $\lim_{n \rightarrow \infty} R_i(\gamma^n) = R_i(\gamma)$ and $\lim_{n \rightarrow \infty} \bar{P}_i(\gamma^n) = \bar{P}_i(\gamma)$ for $i = 1, 2$. Hence, for n high enough, we have that $R_1(\gamma^n) < R_2(\gamma^n)$ and $\bar{P}_2(\gamma^n) > \bar{P}_1(\gamma^n)$. By Proposition 2, for every $p < p^0$,

$$F_1^n(p) = \begin{cases} K_1(p; \gamma^n) & \text{if } p < \bar{P}_1(\gamma^n), \\ K_1(\bar{P}_1(\gamma^n); \gamma^n) & \text{if } p \in [\bar{P}_1(\gamma^n), \bar{P}_2(\gamma^n)), \\ 1 & \text{otherwise,} \end{cases}$$

and

$$F_2^n(p) = \begin{cases} K_2(p; \gamma^n) & \text{if } p < \bar{P}_1(\gamma^n), \\ K_2(\bar{P}_2(\gamma^n); \gamma^n) & \text{if } p \geq \bar{P}_1(\gamma^n). \end{cases}$$

Define, for every $p < p^0$,

$$F_1(p) = \begin{cases} K_1(p; \gamma) & \text{if } p < \bar{P}_1(\gamma), \\ K_1(\bar{P}_1(\gamma); \gamma) & \text{if } p \in [\bar{P}_1(\gamma), \bar{P}_2(\gamma)), \\ 1 & \text{otherwise,} \end{cases}$$

and

$$F_2(p) = \begin{cases} K_2(p; \gamma) & \text{if } p < \bar{P}_1(\gamma), \\ K_2(\bar{P}_2(\gamma); \gamma) & \text{if } p \geq \bar{P}_1(\gamma). \end{cases}$$

The techniques employed in the first part of the proof can be used to show that: K_i is continuous in a neighborhood of $(\bar{P}_i(\gamma), \gamma)$ ($i = 1, 2$); $K_i(\bar{P}_i(\gamma^n); \gamma^n) \xrightarrow[n \rightarrow \infty]{} K_i(\bar{P}_i(\gamma), \gamma)$ ($i = 1, 2$); $(F_1^n)_{n \geq 1}$ converges pointwise to F_1 on $[0, p^0) \setminus \{\bar{P}_2(\gamma)\}$; $(F_2^n)_{n \geq 1}$ converges pointwise to F_2 on $[0, p^0) \setminus \{\bar{P}_1(\gamma)\}$. It follows that $(F_i^n)_{n \geq 1}$ converges weakly to F_i for $i = 1, 2$.

It is then straightforward to check that (F_1, F_2) is an equilibrium of the game with parameter vector γ . \square

C.5 Bertrand Without Fudge

Since the standard model of Bertrand competition with heterogeneous marginal costs discussed in Section 3.3 is non-generic, we cannot apply Proposition 6 to study Bertrand convergence. It is, however, straightforward to adapt the argument in Section C.4 to establish convergence manually. This section relies on the alternative formulation of Section C.1, and use the notation and results of Section C.2.

Consider the following sequence of parameters: For every $n \geq 1$, $\gamma^n = (c_1, c_2, \alpha_1^n, \alpha_2^n, 0, 0, 0, 0)$ with $0 < c_1 < c_2 < p^0$ and $\alpha_1^n, \alpha_2^n < 1$. Suppose that $\alpha_i^n \xrightarrow[n \rightarrow \infty]{} 1$ for $i = 1, 2$, and let $\gamma = (c_1, c_2, 1, 1, 0, 0, 0, 0)$. Then, $R_i(\gamma^n) = R_i(\gamma) = c_i$ for every $n \geq 1$ and $i \in \{1, 2\}$. Moreover, for every $n \geq 1$, $i \in \{1, 2\}$, and $p \in (c_2, p^0)$,

$$K_1(p; \gamma^n) = \frac{p - c_2}{p - \alpha_2^n c_2}, \text{ and } K_2(p; \gamma^n) = \frac{p - c_1}{p - \alpha_1^n c_1} - \frac{(c_2 - c_1)D(c_2)}{(p - \alpha_1^n c_1)D(p)}.$$

It follows that $\bar{P}_1(\gamma^n) = p^0$ and $\bar{P}_2(\gamma^n) \in (c_2, p^0)$ for every n . Since $\gamma \in \bar{\Gamma}_2$, Lemma G implies that $\bar{P}_2(\gamma^n) \xrightarrow[n \rightarrow \infty]{} \bar{P}_2(\gamma) = P_1^m(\gamma)$.

By Proposition 2, the equilibrium profile of CDF's given the vector of parameters γ^n is given by:

$$F_1^n(p) = \begin{cases} 0 & \text{if } p < c_2, \\ \frac{p - c_2}{p - \alpha_2^n c_2} & \text{if } p \in [c_2, \bar{P}_2(\gamma^n)), \\ 1 & \text{if } p \in [\bar{P}_2(\gamma^n), p^0), \end{cases}$$

and

$$F_2^n(p) = \begin{cases} 0 & \text{if } p < c_2, \\ \frac{p - c_1}{p - \alpha_1^n c_1} - \frac{(c_2 - c_1)D(c_2)}{(p - \alpha_1^n c_1)D(p)} & \text{if } p \in [c_2, \bar{P}_2(\gamma^n)), \\ \frac{\bar{P}_2(\gamma^n) - c_1}{\bar{P}_2(\gamma^n) - \alpha_1^n c_1} - \frac{(c_2 - c_1)D(c_2)}{(\bar{P}_2(\gamma^n) - \alpha_1^n c_1)D(\bar{P}_2(\gamma^n))} & \text{if } p \in [\bar{P}_2(\gamma^n), p^0). \end{cases}$$

It is straightforward to adapt the techniques used in the previous subsections to show that, for $i = 1, 2$, $(F_i^n)_{n \geq 0}$ converges weakly to F_i , where

$$F_1(p) = \begin{cases} 0 & \text{if } p < c_2, \\ 1 & \text{if } p \in [c_2, p^0), \end{cases}$$

and

$$F_2(p) = \begin{cases} 0 & \text{if } p \leq c_2, \\ 1 - \frac{(c_2 - c_1)D(c_2)}{(p - c_1)D(p)} & \text{if } p \in [c_2, p_1^m), \\ \frac{p_1^m - c_1}{p_1^m - \alpha_1^n c_1} - \frac{(c_2 - c_1)D(c_2)}{(p_1^m - c_1)D(p_1^m)} & \text{if } p \in [p_1^m, p^0). \end{cases}$$

Moreover, (F_1, F_2) is a Nash equilibrium of the game with parameter vector γ .

Note that firm 1 is indifferent between all the prices in $[c_2, p_1^m)$. If firm 2 were to price less aggressively somewhere in that interval, then firm 1 would have a strictly profitable deviation.

Among the mixed-strategy equilibria identified by Blume (2003) and Kartik (2011), (F_1, F_2) is therefore the equilibrium in which firm 2 is the least aggressive in its randomization.

D Proof of Proposition 9

Proof. Let $F \in \mathcal{F}$. We first define a number of cumulative distribution functions, which will be useful to reexpress social welfare. Define

$$F^- : p \in [0, p^0] \mapsto F^-(p) = \begin{cases} F(p) & \text{if } p < p^0, \\ \lim_{p' \uparrow p^0} F(p') & \text{if } p = p^0, \end{cases}$$

and

$$D^- : p \in [0, p^0] \mapsto D^-(p) = \begin{cases} D(p) & \text{if } p < p^0, \\ \lim_{p' \uparrow p^0} D(p') & \text{if } p = p^0. \end{cases}$$

(The limits exist, as F and D are monotone.) Note that D^- is continuous on $[0, p^0]$. Let $G(p) = 1 - (1 - F(p))^2$ and $G^-(p) = 1 - (1 - F^-(p))^2$ be the cumulative distributions functions of the minimum price, based on the cumulative distribution functions F and F^- , respectively. Similarly, let $H(p) = F(p)^2$ and $H^-(p) = F^-(p)^2$ be the cumulative distributions functions of the maximum price, based on the cumulative distribution functions F and F^- , respectively. Finally, define

$$\Psi : p \in [0, p^0] \mapsto D(0) - D^-(p).$$

Ψ is continuous, bounded, non-decreasing and non-negative. Therefore, Ψ is the cumulative distribution function of some finite measure on $[0, p^0]$.

Expected social welfare is equal to expected consumer gross utility (U) plus expected total costs (C) minus expected recoverable costs (R). We first use Fubini's theorem to obtain a useful expression for expected total costs:

$$\begin{aligned} C &= 2c \int_{[0, p^0)} D(p) dF(p), \\ &= 2c \int_0^{p^0} D^-(p) dF^-(p), \\ &= 2c \left(\int_0^{p^0} (D(0) - \Psi(p)) dF^-(p) \right), \\ &= 2c \left(F^-(p^0) D(0) - \int_0^{p^0} \left(\int_0^p d\Psi(t) \right) dF^-(p) \right), \\ &= 2c \left(F^-(p^0) D(0) - \int_0^{p^0} \left(\int_t^{p^0} dF^-(p) \right) d\Psi(t) \right), \text{ by Fubini's theorem,} \end{aligned}$$

$$\begin{aligned}
&= 2c \left(F^-(p^0)D(0) - \int_0^{p^0} (F^-(p^0) - F^-(t)) d\Psi(t) \right),^{33} \\
&= 2c \left(F^-(p^0)(D(0) - \Psi(p^0)) + \int_0^{p^0} F^-(t)d\Psi(t) \right), \\
&= 2c \left(F^-(p^0)D^-(p^0) + \int_{[0,p^0)} F(t)d\Psi(t) \right).
\end{aligned}$$

Expected recoverable costs can be simplified in a similar way:

$$\begin{aligned}
R &= \alpha c \int_{[0,p^0)} D(p)dH(p), \\
&= \alpha c \left(H^-(p^0)D^-(p^0) + \int_{[0,p^0)} H(t)d\Psi(t) \right).
\end{aligned}$$

Next, we rewrite gross consumer utility at price p :

$$\begin{aligned}
U(p) &= \int_p^{p^0} D(t)dt + pD(p), \\
&= D(0)(p^0 - p) - \int_p^{p^0} \Psi(t)dt + p(D(0) - \Psi(p)), \\
&= p^0D(0) - p\Psi(p) - \int_p^{p^0} \left(\int_p^t d\Psi(x) + \Psi(p) \right) dt, \\
&= p^0D(0) - p\Psi(p) - (p^0 - p)\Psi(p) - \int_p^{p^0} \left(\int_x^{p^0} dt \right) d\Psi(x), \text{ by Fubini's theorem,} \\
&= p^0D(0) - p^0\Psi(p) - p^0(\Psi(p^0) - \Psi(p)) + \int_p^{p^0} xd\Psi(x), \\
&= p^0D(0) - p^0\Psi(p^0) + \int_p^{p^0} xd\Psi(x), \\
&= p^0D^-(p^0) + \int_p^{p^0} xd\Psi(x).
\end{aligned}$$

³³Note that

$$\int_t^{p^0} dF^-(p) = F^-(p^0) - \lim_{x \uparrow t} F(x).$$

Since F is monotone, the set of t 's such that $\lim_{x \uparrow t} F(x) \neq F(t)$ is at most countable. Since Ψ is continuous, the measure associated with Ψ assigns no weight to that set. Therefore,

$$\int_0^{p^0} (F^-(p^0) - F^-(t)) d\Psi(t) = \int_0^{p^0} \left(F^-(p^0) - \lim_{x \uparrow t} F(x) \right) d\Psi(t).$$

Therefore, expected consumer gross utility is given by:

$$\begin{aligned}
U &= \int_{[0, p^0)} U(p) dG(p), \\
&= \int_0^{p^0} \left(p^0 D^-(p^0) + \int_p^{p^0} x d\Psi(x) \right) dG^-(p), \\
&= p^0 D^-(p^0) G^-(p^0) + \int_0^{p^0} \left(\int_0^x dG^-(p) \right) x d\Psi(x), \\
&= p^0 D^-(p^0) G^-(p^0) + \int_0^{p^0} x G^-(x) d\Psi(x), \\
&= p^0 D^-(p^0) G^-(p^0) + \int_{[0, p^0)} x G(x) d\Psi(x).
\end{aligned}$$

Putting things together, we obtain expected social welfare:

$$\begin{aligned}
W(F) &= (p^0 G^-(p^0) - 2cF^-(p^0) + \alpha cH^-(p^0)) D^-(p^0) \\
&\quad + \int_{[0, p^0)} (pG(p) - 2cF(p) + \alpha cH(p)) d\Psi(p), \\
&= (p^0(1 - (1 - F^-(p^0))^2) - 2cF^-(p^0) + \alpha cF^-(p^0)^2) D^-(p^0) \\
&\quad + \int_{[0, p^0)} (p(1 - (1 - F(p))^2) - 2cF(p) + \alpha cF(p)^2) d\Psi(p), \\
&= \Phi(p^0, F^-(p^0)) D^-(p^0) + \int_{[0, p^0)} \Phi(p, F(p)) d\Psi(p),
\end{aligned}$$

where

$$\Phi(p, F) \equiv p(1 - (1 - F)^2) - 2cF + \alpha cF^2, \quad \forall (p, F) \in \mathbb{R}_+ \times [0, 1].$$

It is straightforward to show that, for every $p \in [0, p^0]$,

$$\arg \max_{F \in [0, 1]} \Phi(p, F) = \begin{cases} \{0\} & \text{if } p \leq c, \\ \left\{ \frac{p-c}{p-\alpha c} \right\} & \text{if } p \in (c, p^0]. \end{cases}$$

It follows that, for every policy $F \in \mathcal{F}$,

$$\begin{aligned}
W(F) &= \Phi(p^0, F^-(p^0)) D^-(p^0) + \int_{[0, p^0)} \Phi(p, F(p)) d\Psi(p), \\
&\leq \Phi \left(p^0, \frac{p^0 - c}{p^0 - \alpha c} \right) D^-(p^0) + \int_{[c, p^0)} \Phi \left(p, \frac{p - c}{p - \alpha c} \right) d\Psi(p), \\
&= \Phi(p^0, F^{*-}(p^0)) D^-(p^0) + \int_{[0, p^0)} \Phi(p, F^{*-}(p)) d\Psi(p), \\
&= W(F^*).
\end{aligned}$$

Next, we argue that F^* is the only optimal policy whenever D is strictly decreasing. Let $F \in \mathcal{F}$. Since D is strictly decreasing, $\Psi(p) < \Psi(p')$ for every $p < p' < p^0$. Therefore, the measure associated with Ψ assigns a positive weight to every non-degenerate interval. Assume that $F(\hat{p}) \neq F^*(\hat{p})$ for some $\hat{p} < p^0$. If $F(\hat{p}) > F^*(\hat{p})$, then, since F is non-decreasing and F^* is continuous on $[0, p^0)$, there exists $\varepsilon > 0$ such that $F(p) > F^*(p)$ for every $p \in (\hat{p}, \hat{p} + \varepsilon)$. It follows that $\Phi(p, F(p)) < \Phi(p, F^*(p))$ for every $p \in (\hat{p}, \hat{p} + \varepsilon)$. Since Ψ puts strictly positive weight on that interval, this implies that $W(F) < W(F^*)$, i.e., F is not optimal. Next, assume instead that $F(\hat{p}) < F^*(\hat{p})$. Then, by monotonicity of F and continuity of F^* , there exists $\varepsilon > 0$ such that $F(p) < F^*(p)$ for every $p \in (\hat{p} - \varepsilon, \hat{p})$. Again, this implies that F is not optimal. \square

References

- AMANN, E., AND W. LEININGER (1996): “Asymmetric All-Pay Auctions with Incomplete Information: The Two-Player Case,” *Games and Economic Behavior*, 14(1), 1–18.
- ANDERSON, S., A. BAIK, AND N. LARSON (2015): “Personalized pricing and advertising: An asymmetric equilibrium analysis,” *Games and Economic Behavior*, 92(C), 53–73.
- ARNOLD, M., C. LI, C. SALIBA, AND L. ZHANG (2011): “Asymmetric market shares, advertising and pricing: Equilibrium with an information gatekeeper,” *Journal of Industrial Economics*, 59(1), 63–84.
- BAYE, M., D. KOVENOCK, AND C. DE VRIES (1992): “It takes two to tango: Equilibria in a model of sales,” *Games and Economic Behavior*, 4(4), 493–510.
- (1993): “Rigging the Lobbying Process: An Application of the All-Pay Auction,” *American Economic Review*, 83(1), 289–94.
- BAYE, M. R., AND J. MORGAN (2001): “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91(3), 454–474.
- BELLEFLAMME, P., AND M. PEITZ (2010): *Industrial Organization: Markets and Strategies*. Cambridge University Press.
- BERNSTEIN, F., AND A. FEDERGRUEN (2004): “A General Equilibrium Model for Industries with Price and Service Competition,” *Operations Research*, 52(6), 868–886.
- (2007): “Coordination Mechanisms for Supply Chains Under Price and Service Competition,” *Manufacturing & Service Operations Management*, 9(3), 242–262.
- BILLINGSLEY, P. (2012): *Probability and Measure*. John Wiley & Sons, Hoboken, 4th edn.

- BLANCHARD, O. (1983): “The Production and Inventory Behavior of the American Automobile Industry,” *Journal of Political Economy*, 91(3), 365–400.
- BLINDER, A. S. (1986): “Can the Production Smoothing Model of Inventory Behavior be Saved?,” *The Quarterly Journal of Economics*, 101(3), 431–453.
- BLUME, A. (2003): “Bertrand without fudge,” *Economics Letters*, 78(2), 167–168.
- BORDER, K. C. (1996): “Integration and Differentiation,” Lecture notes.
- BUTTERS, G. R. (1977): “Equilibrium Distributions of Sales and Advertising Prices,” *Review of Economic Studies*, 44(3), 465–491.
- CHE, Y.-K., AND I. L. GALE (1998): “Caps on Political Lobbying,” *The American Economic Review*, 88(3), 643–651.
- (2006): “Caps on Political Lobbying: Reply,” *American Economic Review*, 96(4), 1355–1360.
- CHEN, F., Z. DREZNER, J. K. RYAN, AND D. SIMCHI-LEVI (2000): “Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times, and Information,” *Management Science*, 46(3), 436–443.
- CHOWDHURY, S. M. (2017): “The All-Pay Auction with Nonmonotonic Payoff,” *Southern Economic Journal*, 84(2), 375–390.
- COHEN, C., T. KAPLAN, AND A. SELA (2008): “Optimal rewards in contests,” *RAND Journal of Economics*, 39(2), 434–451.
- DAVIDSON, C., AND R. DENECKERE (1986): “Long-Run Competition in Capacity, Short-Run Competition in Price, and the Cournot Model,” *RAND Journal of Economics*, 17(3), 404–415.
- DE NIJS, R. (2012): “Further results on the Bertrand game with different marginal costs,” *Economics Letters*, 116(3), 502–503.
- DENECKERE, R., AND J. PECK (1995): “Competition Over Price and Service Rate When Demand is Stochastic: A Strategic Analysis,” *RAND Journal of Economics*, 26(1), 148–162.
- DENECKERE, R. J., AND D. KOVENOCK (1996): “Bertrand-Edgeworth duopoly with unit cost asymmetry,” *Economic Theory*, 8(1), 1–25.
- FABINGER, M., AND E. G. WEYL (2012): “Pass-Through and Demand Forms,” mimeo.

- GERTNER, R. H. (1986): “Simultaneous move price-quantity games and equilibrium without market clearing,” in “Essays in Theoretical Industrial Organization,” Ph.D. dissertation, Massachusetts Institute of Technology.
- GROSSMAN, G. M., AND C. SHAPIRO (1984): “Informative Advertising with Differentiated Products,” *Review of Economic Studies*, 51(1), 63–81.
- HAMMOND, J. H. (1994): “Barilla SpA (A),” Harvard Business School Case 694-046.
- HANSEN, R. G. (1988): “Auctions with Endogenous Quantity,” *RAND Journal of Economics*, 19(1), 44–58.
- HARSANYI, J. C. (1973): “Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points,” *International Journal of Game Theory*, 2(1), 1–23.
- KAHN, J. A. (1987): “Inventories and the Volatility of Production,” *American Economic Review*, 77(4), 667–679.
- KAPLAN, T., I. LUSKI, A. SELA, AND D. WETTSTEIN (2002): “All-Pay Auctions with Variable Rewards,” *Journal of Industrial Economics*, 50(4), 417–430.
- KAPLAN, T. R., I. LUSKI, AND D. WETTSTEIN (2003): “Innovative activity and sunk cost,” *International Journal of Industrial Organization*, 21(8), 1111 – 1133.
- KAPLAN, T. R., AND D. WETTSTEIN (2006): “Caps on Political Lobbying: Comment,” *American Economic Review*, 96(4), 1351–1354.
- KARTIK, N. (2011): “A note on undominated Bertrand equilibria,” *Economics Letters*, 111(2), 125–126.
- KREPS, D., AND J. SCHEINKMAN (1983): “Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes,” *Bell Journal of Economics*, 14(2), 326–337.
- KRISHNA, V., AND J. MORGAN (1997): “An Analysis of the War of Attrition and the All-Pay Auction,” *Journal of Economic Theory*, 72(2), 343–362.
- LANG, K., AND R. W. ROSENTHAL (1991): “The Contractors’ Game,” *The RAND Journal of Economics*, 22(3), 329–338.
- LEBRUN, B. (1999): “First Price Auctions in the Asymmetric N Bidder Case,” *International Economic Review*, 40(1), 125–42.
- (2006): “Uniqueness of the equilibrium in first-price auctions,” *Games and Economic Behavior*, 55(1), 131–151.
- LEE, H. L., V. PADMANABHAN, AND S. WHANG (1997a): “The Bullwhip Effect in Supply Chains,” *Sloan Management Review*, 38(3), 93–102.

- (1997b): “Information Distortion in a Supply Chain: The Bullwhip Effect,” *Management Science*, 43(4), 546–558.
- LEVITAN, R., AND M. SHUBIK (1978): “Duopoly with price and quantity as strategic variables,” *International Journal of Game Theory*, 7(1), 1–11.
- MARQUEZ, R. (1997): “A note on Bertrand competition with asymmetric fixed costs,” *Economics Letters*, 57(1), 87–96.
- MASKIN, E. (1986): “The Existence of Equilibrium with Price-Setting Firms,” *American Economic Review*, 76(2), 382–86.
- MASKIN, E., AND J. RILEY (2000): “Equilibrium in Sealed High Bid Auctions,” *Review of Economic Studies*, 67(3), 439–454.
- (2003): “Uniqueness of equilibrium in sealed high-bid auctions,” *Games and Economic Behavior*, 45(2), 395–409.
- MOLDOVANU, B., AND A. SELA (2001): “The Optimal Allocation of Prizes in Contests,” *American Economic Review*, 91(3), 542–558.
- (2006): “Contest architecture,” *Journal of Economic Theory*, 126(1), 70–96.
- NARASIMHAN, C. (1988): “Competitive Promotional Strategies,” *The Journal of Business*, 61(4), 427–449.
- PLUM, M. (1992): “Characterization and Computation of Nash-Equilibria for Auctions with Incomplete Information,” *International Journal of Game Theory*, 20(4), 393–418.
- RAMEY, V. A. (1991): “Nonconvex Costs and the Behavior of Inventories,” *Journal of Political Economy*, 99(2), 306–334.
- RILEY, J., AND W. F. SAMUELSON (1981): “Optimal Auctions,” *American Economic Review*, 71(3), 381–92.
- RUDIN, W. (1976): *Principles of mathematical analysis*. McGraw-Hill Book Co., New York, third edn., International Series in Pure and Applied Mathematics.
- SHARKEY, W., AND D. S. SIBLEY (1993): “A Bertrand model of pricing and entry,” *Economics Letters*, 41(2), 199–206.
- SHELEGIA, S., AND C. M. WILSON (2016): “A generalized model of sales,” Economics Working Papers 1541, Department of Economics and Business, Universitat Pompeu Fabra.
- SIEGEL, R. (2009): “All-Pay Contests,” *Econometrica*, 77(1), 71–92.

- (2010): “Asymmetric Contests with Conditional Investments,” *American Economic Review*, 100(5), 2230–2260.
- (2014a): “Asymmetric Contests with Head Starts and Nonmonotonic Costs,” *American Economic Journal: Microeconomics*, 6(3), 59–105.
- (2014b): “Contests with productive effort,” *International Journal of Game Theory*, 43(3), 515–523.
- SPULBER, D. (1995): “Bertrand Competition When Rivals’ Costs Are Unknown,” *Journal of Industrial Economics*, 43(1), 1–11.
- TASNADI, A. (2004): “Production in advance versus production to order,” *Journal of Economic Behavior & Organization*, 54(2), 191–204.
- TESCHL, G. (2012): *Ordinary Differential Equations and Dynamical Systems*, Graduate studies in mathematics. American Mathematical Society.
- THOMAS, C. J. (2002): “The effect of asymmetric entry costs on Bertrand competition,” *International Journal of Industrial Organization*, 20(5), 589 – 609.
- TIROLE, J. (1987): *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.
- VARIAN, H. R. (1980): “A Model of Sales,” *American Economic Review*, 70(4), 651–659.
- WEBER, R. (1985): “Auctions and competitive bidding,” in *Proceedings of Symposia in Applied Mathematics*, vol. 33, pp. 143–170.
- ZHAO, X., AND D. R. ATKINS (2008): “Newsvendors Under Simultaneous Price and Inventory Competition,” *Manufacturing & Service Operations Management*, 10(3), 539–546.