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

Joao Montez† Nicolas Schutz‡

August 28, 2019

Abstract

We study production-in-advance in a setting where firms first source inventories that remain unobservable to rivals, and then simultaneously set prices. In the unique equilibrium, each firm occasionally holds a sale relative to its reference price, resulting in firms sometimes being left with unsold inventory. In the limit as inventory costs become fully recoverable, the equilibrium converges to an equilibrium of the game where firms only choose prices and produce to order—the associated Bertrand game (examples of such games include fully-asymmetric clearinghouse models). Thus, away from that limit, our work generalizes Bertrand-type equilibria to production in advance, and challenges the commonly-held view associating production in advance with Cournot outcomes. The analysis involves, as an intermediate step, mapping the price-inventory game into an asymmetric all-pay contest with outside options and non-monotonic winning and losing functions. We lay out applications to taxation, merger analysis, information sharing, ex-ante investments, and vertical relations.

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

---

*We thank the Editor, three anonymous referees, Simon Anderson, Heski Bar-Isaac, Pierre Boyer, Bruno Jullien, Sébastien Mitraille, Volker Nocke, Hans-Theo Normann, Martin Peitz, Alexander Rasch, Régis Renault, Markus Reisinger, Patrick Rey, Armin Schmutzler, Sandro Shelegia, Andre Speit, Yossi Spiegel, Roland Strausz, Thomas Tröger, Thomas von Ungern-Sternberg, Ralph Winter, as well as seminar and conference participants at Bergen, Bern, Católica Porto Business School, Cergy Pontoise, CREST-Ecole Polytechnique, Dusseldorf, East Anglia, Goethe University Frankfurt and Frankfurt School of Finance & Management, Leicester, Nova School of Business & Economics, Pompeu Fabra, Toulouse, BECCLE Law & Economics Conference, Bergamo Workshop on Advances in IO 2018, Berlin IO Day 2019, CEPR Applied IO Conference 2019, EARIE 2018, Lisbon Meetings in Game Theory 2018, MaCCI Annual Meeting 2017, MaCCI Summer Institute 2018, Meeting of the IO Committee of the German Economic Association 2019, Swiss Theory Day 2018, Workshop on Competition and Bargaining in Vertical Chains for helpful comments and suggestions. We gratefully acknowledge financial support from the Swiss National Science Foundation through grant no. 00018_176448 and from the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224.

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

‡Department of Economics and MaCCI, University of Mannheim and CEPR. Email: schutz@uni-mannheim.de.

1 Introduction

Most retail markets are characterized by production in advance, as each store chooses not only its price but also its inventory level, sourced from suppliers and made readily available to consumers. As rivals' inventories are hard to observe, stores have 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—or, equivalently, a 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 allows for asymmetries between firms, heterogeneity in consumer tastes, and the coexistence of informed and uninformed consumers. Importantly, we 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 fraction captures the salvage value of unsold inventories (for instance through outlets stores sales, buyback contracts with manufacturers, or future sales) and the cost of sales and post-sales services.

We start by solving a constrained version of the model in which each store chooses a price and must source enough inventory to serve all its targeted demand. This *constrained game* has the structure of an all-pay contest with outside options, non-monotonic winning and losing functions, and conditional investments. Adapting tools from the all-pay contests literature Siegel (2009, 2010), we develop a method to obtain the closed-form characterization of the equilibria of such constrained games, and show that there is a generically unique equilibrium. Next, we show that this 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 total demand, resulting in unsold inventories.

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 game in which firms only choose prices and produce to order—the associated Bertrand game.[1] The equilibrium is thus *Bertrand convergent*. Away from that limit, our closed-form equilibrium characterization thus generalizes the Bertrand-type equilibrium to production-in-advance industries where the value of unsold inventories falls short of their acquisition value.

This insight stands in contrast to Kreps and Scheinkman (1983)'s well-known result for production-in-advance industries. Assuming that inventory choices become *observable* before the pricing stage, those authors conclude that the equilibrium outcome is Cournot, even when almost all the unit cost "is incurred subsequent to the realization of demand (situations that will look very Bertrand-like)." The conventional interpretation of that result is that a Cournot analysis is appropriate under production in advance, whereas Bertrand is preferable under

[1] Under production to order, inventories are sourced to meet demand *after* consumers' purchase decisions.

production to order (see, e.g., Belleflamme and Peitz, 2010, pp. 66–67).

Our analysis suggests that the model selection for production-in-advance settings should explicitly take into account whether inventory choices are *observed* by rivals or *not*. In the former case, the Cournot outcome is a reasonable benchmark. In the latter case, a low inventory choice cannot provide a commitment to soften price competition, and so the Kreps-Scheinkman mechanism fails. Then, as our Bertrand convergence result suggests, a Bertrand approach is better justified when most of the unit cost can be salvaged, and our analysis characterizes the equilibrium behavior for lower salvage values.

Also in contrast to the observable-inventories case, with unobservable inventories the extent to which unit costs are sunk or variable fundamentally determines price levels, consumer surplus, and social welfare. We find that less recoverable costs induce businesses to carry lower levels of inventories, which ultimately results in higher markups and prices.

In a dynamic version of our model, a lower discount factor (e.g., a higher interest rate) increases the cost of holding excess inventories, and thus, like in our static model, leads to higher prices. This mechanism suggests a countervailing effect of increasing interest rates on price levels—which is absent if inventories are observable, or in production-to-order and perfect-competition versions of the same model. In the industrial organization literature, the relationship between firms' discount factors and price levels is typically viewed through the lens of collusion models, and thus perceived to be positive. Our analysis uncovers that, in production-in-advance industries, those two variables may in fact be negatively related.

Large unsold inventories are a well-documented feature of several industries, including grocery and apparel.[2] Common explanations rely on exogenous market demand uncertainty. By contrast, in our model market demand is deterministic, but *store-level* demand becomes endogenously 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 endogenous strategic uncertainty, arising from the stores' need to be unpredictable, provides a new explanation for the persistence of unsold inventories.

This mechanism also offers an explanation as to why some products with a low variance 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). Also at the macro level, the variance of production is typically greater than that of demand (e.g., Blanchard, 1983). Our model suggests the bullwhip effect could be explained by retailers sourcing inventories in advance and frequently offering discounts to their consumers.[3]

[2]For 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 accumulated more than $4 billion in unsold clothes in 2018.

[3]Previous explanations include adjustment to cost shocks (Blinder, 1986), increasing returns to scale

Our framework and closed-form equilibrium characterization can be used as a building block to uncover novel economic insights. In this vein, we lay out applications to taxation, merger analysis, horizontal information sharing, pre-competition investments, and vertical relations. Among other results, we find that: In a symmetric version of the model, the equilibrium is second-best efficient (in that it cannot be improved by standard taxation); a merger to monopoly tends to raise social welfare and consumer surplus if unit costs are high and mostly sunk; and finally, in a setup with ex-ante investments, the industry endogenously gravitates towards asymmetry.

Related literature. Our analysis of the constrained game, where firms must source enough inventories to supply their targeted demand, contributes to the literature on all-pay contests. In an all-pay contest, 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 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 (2014b, c), and Chowdhury (2017). There are however differences. For instance, our losing functions are typically discontinuous in participation, due to fixed costs or to firms having the option to focus on their captive consumers only. Also, in our model, the weak and strong firms' winning functions may cross, 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 firms may use multiple mass points.

Our equilibrium characterization extends the analysis of production-to-order models (i.e., Bertrand models) to environments where inventories are chosen in advance. As inventory costs become fully variable, the equilibrium of the production-in-advance game converges to an equilibrium of the associated Bertrand game. 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; Iyer, Soberman, and Villas-Boas, 2005; Kocas and Kiyak, 2006; Shelegia and Wilson, 2016, 2019). It has been recognized that production-to-order games share characteristics with all-pay contests. Our work contributes to a better understanding of this connection. Our Bertrand convergence result also relates to Siegel (2010)'s finding that, in an all-pay auction, as payments become entirely conditional on winning, equilibrium play converges to the equilibrium of the limiting

upstream (Ramey, 1991), and stochastic demand (Kahn, 1987; Lee, Padmanabhan, and Whang, 1997b).

first-price auction—the Bertrand outcome of that model.

In contrast to the all-pay contests and production-to-order literatures, our main contribution is the study of the unconstrained game, where firms freely choose inventories. Our result that the equilibrium of the constrained game is the generically unique equilibrium of the unconstrained price-inventory game is novel, and involves considerations and technical challenges that are naturally absent from those literatures.

A small number of papers has analyzed simultaneous price-inventory models related to ours. Maskin (1986) proves equilibrium existence for a class of duopoly games. A non-generic version of our model was previously studied by Levitan and Shubik (1978) with linear demand and Gertner (1986) when inventory costs are completely sunk. Perturbing a game that is not generic leads to a generic game and it is known that the equilibrium behavior in non-generic contests can be very different from that in generic ones (see, e.g., Siegel, 2009). These observations, and the fact that numerous applications do not become interesting unless firm asymmetries and heterogeneous consumer preferences for stores are allowed, motivate the study of the rich class of all-pay oligopolies we introduce in this paper. Moreover, the proofs in those earlier papers omit important non-trivial steps or contain several inaccuracies.[4] For completeness, we provide a proof that addresses those shortcomings in Online Appendix II.

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

Finally, Deneckere and Peck (1995) study pure-strategy equilibria in a symmetric, simultaneous price-inventory model with stochastic demand. In equilibrium, all stores choose the same price (i.e., no price dispersion), consumers are rationed when demand is high, and some inventories remain unsold when demand is low. However, pure-strategy equilibria do not exist when, close to our setting, demand uncertainty or the number of stores is low. Variations of this problem were subsequently studied in the newsvendor literature by operations scholars (e.g., Bernstein and Federgruen, 2004, 2007; Zhao and Atkins, 2008).

The paper is organized as follows. Section 2 introduces the class of all-pay oligopolies. Section 3 characterizes the generic equilibrium and identifies its properties. Section 4 studies the limiting case as inventory costs become fully recoverable as well as its connection to classic

[4] Levitan and Shubik (1978) confine attention to equilibria without mass points (except potentially at the choke price). Gertner (1986)'s approach is generally correct, but his proof, which was never peer-reviewed, contains measure-theoretical inconsistencies. In Gertner's setting, Tasnadi (2004, 2018) shows that firms make zero profits in any symmetric equilibrium, but does not characterize equilibrium behavior.

Bertrand models, and provides a dynamic formulation where such a limit is obtained as firms become patient. Section 5 discusses extensions and applications. Section 6 concludes.

2 All-Pay Oligopolies

There are two firms, 1 and 2. Firm i incurs a constant unit cost $c_i > 0$ for each unit it sells. For each unit that is sourced but remains unsold, a fraction $\alpha_i \in [0, 1)$ of the unit cost c_i is recovered. Thus, $1 - \alpha_i$ captures the extent to which the inventory cost is sunk.

Total market demand is given by the function D, assumed to be strictly positive, continuous, non-increasing, and log-concave on $[0, p^0)$, left-continuous at the choke price $p^0 \in (0, \infty)$, and identically equal to zero on (p^0, ∞).[5] Consumers may perceive the two firms as being differentiated, so that not all consumers necessarily wish to purchase from the firm setting the lowest price. A fraction of consumers $\mu_i \in [0, 1)$, firm i's *captive segment*, only wishes to buy from firm i. The remaining fraction $1 - \mu_1 - \mu_2 \in (0, 1]$ of consumers, the *contested segment* of *shoppers*, wishes to buy from the firm with the lowest price. Firm i can pay a fixed advertising cost of $A_i \geq 0$ to make the shoppers aware of its product and price. This becomes a fixed or entry cost if there are no captive consumers, i.e., $\mu_1 = \mu_2 = 0$. As commonly assumed in oligopoly theory, rationing within a consumer segment is either random or efficient, and demand is allocated in a same-price fair-share way in case of a tie. To specify rationing across segments, we assume that each firm serves its captive consumers first.[6]

Firms 1 and 2 simultaneously decide how many units to source, whether to pay the advertising cost, and what price to set. This is formally equivalent to firms first choosing inventories, which remain unobservable, and subsequently making advertising and pricing decisions. We look for Nash equilibria.

We now introduce terminology and additional assumptions that will be used throughout the paper. Much of this terminology follows closely the one used in the all-pay contests literature (see Siegel, 2009). Note that firm i can always guarantee itself a payoff of

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

by not advertising and serving its captive consumers only at its monopoly price p_i^m. We call o_i the *outside option* payoff of firm i—which is zero if $\mu_i = 0$. If instead firm i sets a price $p_i > c_i$ and decides to serve both its captive and the contested segment, then it captures at best the whole contested segment, in which case its payoff is

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

[5] A discontinuity at p^0 can capture the fact that consumers may have access to an alternative source of supply at price p^0, so that the demand faced by the oligopoly players is truncated at that price. It also allows us to nest perfectly inelastic demand as a special case.

[6] Rationing and sharing rules, formally defined in Appendix B.1, do not affect the equilibrium characterization of generic cases. In contrast to other settings (Kreps and Scheinkman, 1983; Davidson and Deneckere, 1986), all we need is that sales are weakly increasing in own inventories.

Thus, the most firm i receives when serving its captive and the contested segment is $w_i(p_i^m)$.

If $o_i > w_i(p_i^m)$ for some firm i, 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 $o_i < w_i(p_i^m)$. Then, firm i will never use a price below

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

Call this price the *reach* of firm i and let $r = \max\{r_1, r_2\}$ be the highest reach, known as the threshold in the all-pay contests literature (Siegel, 2009). 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 again dominance solvable: Firm i serves its captive and the contested segment at its monopoly price, whereas firm j serves its captive consumers.

In the remainder of the paper, we focus on the non-trivial case where $w_i(p_i^m) > o_i$ and $p_i^m > r$ for every i. In Online Appendix VII.1, we show that under the convention that $c_i \leq c_j$, these conditions are equivalent to $w_i(p_i^m) > o_i$ and $w_j(p_i^m) > o_j$.

Moreover, we are interested in characterizing equilibrium behavior in generic games. A game is generic if $r_1 \neq r_2$ and all parameters are interior and differ across firms. Genericity is sometimes stronger than necessary for our results. When stating formal results, we mention the key (weaker) conditions required inside parentheses. We also provide equilibrium characterization results for non-generic games. Those games often have a continuum of equilibria, but this multiplicity disappears when the game is slightly perturbed and becomes generic.

3 Equilibrium Analysis

The equilibrium analysis proceeds in two steps. In Section 3.1, we study a *constrained* version of the model, where firms must source enough inventory to supply their targeted demand, and show that the equilibrium is generically unique. In Section 3.2, we show that this equilibrium is also the generically unique equilibrium of the *unconstrained* game, where firms can freely choose their inventories. In Section 3.3, we discuss the equilibrium's qualitative features.

3.1 The Constrained Game

In the constrained game, each firm must source enough inventories to supply its targeted demand. Firm i must therefore decide whether to target only its captive segment, or both its captive and the contested segment. In the former case, it receives its outside option payoff o_i. If instead firm i targets both its captive and the contested segment at price p_i, then it pays its advertising cost and sources $(1 - \mu_j)D(p_i)$ units. If firm j is not targeting the contested segment or $p_i < p_j$, then firm i sells all the units it sources and receives $w_i(p_i)$, which we call firm i's *winning payoff*. If firm j is targeting the contested segment at a price $p_j < p_i$, then $(1 - \mu_i - \mu_j)D(p_i)$ units remain unsold and firm i receives its *losing payoff*:

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

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, conditional investments and, due to revenue concavity, potentially non-monotonic winning and losing functions.

We now proceed with the equilibrium characterization. Standard arguments imply that there is no pure-strategy Nash equilibrium. A mixed-strategy equilibrium is fully characterized by a pair (F_1, F_2) of cumulative distribution functions (CDF) 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 $1 - F_i(p^0)$.[7]

In Appendix A (see Lemma B), we show that the firms' payoffs are $w_i(r)$ in any equilibrium. This payoff characterization is to be expected from the all-pay contests literature (Baye, Kovenock, and de Vries, 1992; Siegel, 2009, 2014c). On the other hand, the properties of the equilibrium can be quite different from those found in that literature.

We look for an equilibrium in which firms mix continuously over some interval $[r, \bar{p})$ and distribute the remaining mass on higher prices or their outside option. 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

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

By log-concavity of D, either k_j is single-peaked and achieves its maximum at some $\bar{p}_j \in (r, p^0)$, or it is strictly increasing, in which case we set $\bar{p}_j = p^0$ (see Lemma C in the Appendix).

Loosely speaking, $\bar{p}_j$ can be viewed as the highest price firm i is willing to set, in that firm j's CDF would need to decrease to induce firm i to set higher prices. 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 in the Appendix), meaning that we have some mass left to distribute. As discussed next, how that mass is distributed depends crucially on whether $\bar{p}_2$ is higher or lower than $\bar{p}_1$, i.e., on

[7] In the special case where $\mu_i = A_i = D(p^0) = 0$, firm i can be inactive either by advertising and setting $p_i = p^0$ and $q_i = 0$, or equivalently by not advertising and setting $p_i = p_i^m$ and $q_i = 0$. The latter ensures that $1 - F_i(p^0)$ is indeed the probability that firm i takes its outside option.

whether firm 1 is more or less willing to set high prices than firm 2.

We use the convention that firm 1 is the strong firm, i.e., $r_1 < r_2 = r$. We will later study the non-generic case $r_1 = r_2$. Suppose first that $\bar{p}_1 \geq \bar{p}_2$, so that $\bar{p} = \bar{p}_2$. The fact that k_2 is decreasing on $(\bar{p}, p^0)$ means that, given that firm 2 already puts a 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 firm at $\bar{p}$, the weak firm on its outside option) and the CDFs are:

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

It is readily verified that this pair of CDFs 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 the CDFs are:

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

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 $[r, \bar{p}]$. 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 4.2 for a graphical illustration of equilibrium behavior in a simple example.)

Equilibrium uniqueness is then established using standard techniques. Summarizing:

Proposition 1. *The constrained game of a generic ($r_1 \neq r_2$) all-pay oligopoly has a unique equilibrium. The equilibrium profile of CDFs of prices in the contested segment is described by equation* (1) *if the weak firm is willing to set higher prices than the strong firm (i.e., $\bar{p}_2 \leq \bar{p}_1$), 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 a strictly positive probability. Equilibrium payoffs are $w_i(r)$ for $i = 1, 2$.*

Proof. See Appendix A. □

By contrast, the constrained game of a non-generic ($r_1 = r_2$) all-pay oligopoly usually has a continuum of equilibria. Intuitively, in the non-generic case, there is more leeway to allocate the mass that firms do not put on $[r, \bar{p})$, so that the equilibria differ only in the probability that players take their outside option or price at $\bar{p}$.[8] Proposition A, stated and proven in the Appendix, provides a complete characterization of the set of equilibria.

One particular instance where firms have identical reaches arises when they are symmetric in all dimensions. In this case, the set of equilibria can be described as follows: Firms mix continuously over $[r, \bar{p})$ according to k; at most one firm has a mass point at $\bar{p}$; both firms put the rest of their mass on outside options. In symmetric settings, it is almost a convention to select the symmetric equilibrium, which here means selecting the equilibrium where both firms put all of their remaining mass on their outside option.

Proposition 1 shows that this focus on the symmetric equilibrium can be unsatisfactory in a production-in-advance setting. We have found that in a generic game, the strong firm targets the contested segment with probability 1, whereas in the symmetric equilibrium, both firms do not target the contested segment with a strictly positive probability. Thus, the symmetric equilibrium cannot be approached by a sequence of equilibria of generic games. This implies that if the symmetric equilibrium is selected, then any slight perturbation of the game must result in completely different behavior—whereas the asymmetric equilibrium in which one of the firms puts no mass on its outside option survives such perturbations.

3.2 The Unconstrained Game

We now study the unconstrained game, where firms can freely choose their inventories. Let (F_1, F_2) be an equilibrium of the constrained game. Suppose firm i deviates to a price-inventory pair (p, q) with $q \in [\mu_i D(p), (1 - \mu_j) D(p))$, i.e., such that it does not source enough inventory to supply its targeted demand. Unless firm j has a mass point at p, firm i earns:[9]

$$\tilde{\pi}_i(p, q) = \mu_i (p - c_i) D(p) - A_i + \Big((p - \alpha_i c_i)(1 - F_j(p)) - (1 - \alpha_i) c_i \Big) (q - \mu_i D(p)). \qquad (3)$$

[8] This multiplicity is similar in spirit to the one in Siegel (2014a), except that in his paper the analogue of $\bar{p}$ (a reserve price) is exogenous, whereas in our case it arises endogenously from non-monotonicities.

[9] If 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.

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

The next question is whether the converse of Proposition 2 also holds, i.e., whether any equilibrium of the unconstrained game is also an equilibrium of the constrained game. We find a positive answer for generic games. The proof for this direction is technically more challenging and, therefore, we only provide here a heuristic argument.

Assume for a contradiction that an equilibrium exists, in which at least one firm does not always source enough inventories to supply its targeted demand. Define $\hat{p}$ as the supremum of the set of prices below which both firms always source enough inventory to supply their targeted demand. It can be shown that this supremum is a maximum. Suppose next that firm i chooses a pair (p,q) in the support of its equilibrium strategy, with $p > \hat{p}$ and $\mu_i D(p) \le q < (1-\mu_j)D(p)$, 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:[10]

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

where $\varepsilon(p,q)$ 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. Since sales are non-decreasing in own inventories, ε is non-decreasing in q. Moreover, as p decreases to $\hat{p}$, the probability that firm j prices in $(\hat{p}, p)$ converges to zero, and $\varepsilon(p,q)$ therefore tends to zero. When $A_i > 0$,

$$\tilde{\pi}_i(p,q) \ge o_i > \max_p \mu_i(p-c_i)D(p) - A_i. \tag{4}$$

This implies that in generic games, for any p sufficiently close to $\hat{p}$ and in the support of firm i's marginal on prices,

$$(p-\alpha_i c_i)(1-F_j(p)) > (1-\alpha_i)c_i.$$

It follows that $\tilde{\pi}_i(p,q)$ is strictly increasing in q for every such p, a contradiction. Formalizing this argument, we obtain:

Proposition 3. *In a generic* ($A_1, A_2 > 0$) *all-pay oligopoly, 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. □

Combining Propositions 1–3 and Proposition A in the Appendix gives:

[10] We have ignored the possibility that firm j has a mass point at p for ease of exposition.

Theorem 1. *A generic ($A_1, A_2 > 0$, $r_1 \neq r_2$) all-pay oligopoly has a unique equilibrium. Both firms source enough inventories to supply their targeted demand and the CDFs of prices are as characterized in Proposition 1.*

A non-generic all-pay oligopoly may have multiple equilibria. Constrained equilibria, which are also unconstrained equilibria, are as characterized in Proposition A in the Appendix.

The equilibrium multiplicity that can arise in *non-generic* cases is discussed in Appendix B.3. There, we show that non-generic all-pay oligopolies can have not only a continuum of constrained equilibria, but also a continuum of equilibria that are not constrained equilibria. Both sources of multiplicity disappear if the game is slightly perturbed.[11]

3.3 Equilibrium Properties

We now list and discuss the generic qualitative features of the equilibrium.

(a) There is price dispersion and each firm has at least one mass point (i.e., a reference price). The strong firm has a single mass point, whereas the weak firm has one or two depending on whether it is more or less willing to set high prices than the strong firm.

We provide sufficient conditions under which either case of Proposition 1 arises in Online Appendix VII.2. In a nutshell, the weak firm has one mass point if it has higher and less recoverable costs than the strong firm. Conversely, it has two mass points if it has lower and more recoverable costs. A firm with such a variable cost advantage can still be weak if it has more captive consumers or higher advertising costs since either of these makes it less interesting to target the contested segment.[12]

(b) Firms source enough inventory to supply their targeted demand, and strictly prefer to do so. The latter follows from combining equation (3) and the second inequality in equation (4), implying that $\tilde{\pi}_i(p, q)$ is linear and strictly increasing in q for every p in the support of firm i's equilibrium strategy.

(c) All segments are served with probability one: The strong firm always targets both its captive and the contested segment, and the weak firm always targets at least its captive consumers. Combining this with the previous property, it follows that stock-outs never occur.

(d) Some inventories remain unsold with positive probability. Although market demand is deterministic, the strategic uncertainty resulting from the rival's need to be unpredictable implies that firm-level demand is stochastic. Firms' inability to anticipate the timing and depth of rivals' promotions therefore prevents them from adjusting inventory choices accordingly.[13] As mentioned in the introduction, this mechanism provides an explanation as to why

[11] Positive advertising costs thus give rise to a selection criterion in the case with zero advertising costs.

[12] As we show in the Online Appendix (Corollary II), if $A_1 < A_2$, $\mu_1 < \mu_2$ and $c_1 = c_2$, then firm 2 has one mass point if $\alpha_1 \geq \alpha_2$ and two mass points otherwise. For a worked-out example, see also Section 4.2.

[13] Mixed-strategy equilibria are not regret-free: Once firm i has observed the realization of firm j's price, it no longer wants to mix. 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.

some products with low variance in consumer demand feature high variance in inventory orders (the bullwhip effect).

(e) The strong firm prices above its monopoly level with positive probability. This follows as the strong firm always has a mass point on $\bar{p}_2$, which strictly exceeds p_1^m.[14] (For the same reason, the weak firm may also price above its monopoly price.) The intuition for prices above the monopoly level is also driven by the stochasticity of firm-level demand induced by strategic uncertainty. Consider a hypothetical monopolist facing a demand of $D(p)$ with probability λ, and 0 otherwise. It maximizes

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

and thus behaves as a monopolist with a unit cost that exceeds c. Hence, it prices above p^m. Likewise, in our setting, given the strategy of its rival, firm i faces demand from the contested segment with probability $1 - F_j(p)$, and no demand from that segment otherwise. A similar mechanism thereby raises the firm's perceived cost of production, rationalizing again prices above the monopoly level. Thus, with a strictly positive probability, prices can be higher in oligopoly than with a multi-segment monopolist, i.e., production in advance can lead to price-increasing competition—see our application to merger analysis in Section 5.[15]

(f) When inventory costs are more recoverable, firms price more aggressively in the contested segment in the first-order stochastic dominance sense. (See Online Appendix VII.3 for formal statements and proofs.) Thus, the price at which the shoppers purchase decreases with inventory recoverability, and likewise for the strong firm's captive consumers. A corollary is that if the weak firm has no captive consumers, then consumers are better off when inventory costs are more recoverable.[16]

4 Bertrand Convergence and Dynamics

We show in Section 4.1 that classic Bertrand-type analysis is nested as a limiting case: As inventory costs becomes fully recoverable, the equilibrium converges to an equilibrium of the resulting Bertrand game, where firms set prices and produce to order. Conversely, away from the limit, our analysis provides a toolbox to generalize the Bertrand-type equilibrium to markets with production in advance. We provide illustrations for this in the context of

[14] Here, we implicitly assume that $p_1^m < p^0$, which holds, e.g., if D is continuous at the choke price p^0. To see why $\bar{p}_2 > p_1^m$, note that starting from $p = p_1^m$, a small increase in p has no first-order impact on the numerator of k_2 (i.e., $w_1(p) - w_1(r)$), but strictly reduces the denominator (i.e., $w_1(p) - l_1(p)$) as $(p - \alpha_1 c_1)D(p)$ is locally decreasing. The result follows as k_2 is single-peaked and achieves its maximum at $\bar{p}_2$.

[15] See Chen and Riordan (2008) for another instance of price-increasing competition.

[16] The weak firm's captive consumers are also better off when α_2 increases, but we cannot always guarantee they are better off when α_1 increases: Increasing the latter transfers some mass the weak firm was putting on its outside option towards prices that may be above its monopoly level. However, if $c_1 < c_2$ and α_1 is high, then increasing either α enables all consumers to purchase at a lower price (see the Online Appendix).

several classic Bertrand games in Section 4.2. Finally, in Section 4.3, we introduce a simple dynamic version of the model, which provides microfoundations for salvage values reflecting firms' impatience. Connecting this with our earlier results, this model implies that, when goods are non-perishable, firms that are more patient tend to set lower prices, and firms that are very patient price close to Bertrand.

4.1 The Convergence Result

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

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:[17]

Proposition 4. *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.3. □

Proposition 4 delivers two additional properties:

(g) The equilibrium is Bertrand convergent. Our approach thus nests the classic Bertrand one as a limiting case, and generalizes it to production in advance away from that limit.

(h) Firms earn their Bertrand profits regardless of the recoverability parameters, as explained next. In an all-pay oligopoly, 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 4, as those parameters tend to 1, equilibrium play tends to an equilibrium of the limiting Bertrand game. Hence, the Bertrand profit is also $w_i(r)$. This property implies that firms have no incentives to make inventories more recoverable, since any such advantage is competed away through more aggressive pricing. It is reminiscent of Siegel (2010)'s Corollary 1 that equilibrium payoffs in a simple contest do not depend on the extent to which investment costs are sunk.

As mentioned in the introduction, Proposition 4 is in contrast to what is known from the literature that studies production in advance in the Kreps and Scheinkman (1983) framework, in which $\mu_i = A_i = 0$ (e.g., Davidson and Deneckere, 1986; Deneckere and Kovenock, 1996). Assuming *observable* inventory choices, that literature finds that if the sunk part of the unit

[17] F_i^n was defined as the CDF of a measure over $[0, p^0]$. Since F_i^n is not necessarily a *probability* measure, the weak convergence of $(F_i^n)_{n\geq 0}$ may not be well defined. We circumvent this by studying an equivalent auxiliary game in which any equilibrium is a pair of *probability* measures over $[0, p^0]$, and proving the weak convergence of the associated sequences of probability measures. See Appendix C.1 for details.

cost is sufficiently high, then the equilibrium outcome is (close to) Cournot. When inventory costs are sufficiently recoverable, other outcomes can arise, which remain close to Cournot, and thus far from the intense competition arising in Bertrand.

A main takeaway from that literature is that observable inventory choices act as a commitment to soften price competition and protect firms' profits (e.g., Tirole, 1987, p. 218). So far, a counterfactual with unobservable inventory choices had not been properly investigated. We have shown that in the absence of the commitments provided by such observability, firms earn their Bertrand profits. Pricing behavior is however not Bertrand-like, unless unit costs are sufficiently recoverable. In general, prices will be higher than under Bertrand, and may even exceed the Cournot level in expectation. Consumers may therefore be worse off when inventory choices are unobservable, as discussed in greater detail in Section 5.

4.2 Illustrations

We now study specific all-pay oligopolies and relate them to well-known Bertrand games using the convergence result in Proposition 4.

Textbook Bertrand. In the *textbook Bertrand model* ($\alpha_i = 1$), units costs are identical ($c_i = c$), there are neither captive consumers nor advertising costs ($A_i = \mu_i = 0$), and demand is continuous. In the unique equilibrium, firms price at marginal cost (Harrington, 1989).

Consider the production-in-advance version of this model with $\alpha_i = \alpha < 1$. There is a unique equilibrium: Each firm mixes over $[c, p^0)$ according to $k(p) = (p - c)/(p - \alpha c)$ and puts the rest of its mass $1 - k(p^0) > 0$ on inactivity. As α increases from 0 to 1, prices decrease (in a first-order stochastic dominance sense) and converge to a mass point at c—Bertrand convergence. Each firm's inventory increases and converges to a mass point at $D(c)$ (and likewise for unsold inventories). The total incurred cost of production, $cD(\min(p_1, p_2)) + (1 - \alpha)cD(\max(p_1, p_2))$, also converges to the Bertrand one, $cD(c)$.

The equilibrium characterized above was previously uncovered by Levitan and Shubik (1978) and Gertner (1986) in variants of this model with fully sunk inventory costs. The proofs of equilibrium uniqueness they provide are however incomplete. In Online Appendix II, we provide a uniqueness proof that addresses these shortcomings. In view of Theorem 1, the reader may wonder why the proof requires such a long development. The reason is that the present model is non-generic in two crucial ways: Advertising costs are zero and firms are identical in every dimension, implying that firms have identical reaches. For this reason, Proposition 3 does not apply and we need to use completely different techniques.[18]

This non-genericity also has implications for equilibrium behavior. For example, the equilibrium CDF does not depend on the shape of $D(\cdot)$, and there is a positive probability

[18] Our approach in this case essentially follows Gertner (1986)'s, which crucially builds on both firms making zero profit. The case of efficient rationing (Online Appendix II.6), which Gertner (1986) does not study, is significantly more involved than the random rationing one (Online Appendix II.5). Interestingly, the rationing rule plays no role in the proof of Theorem 1, so those complications are generically irrelevant.

that no firm serves the market, violating property (c). Also, both firms are indifferent between all the pure strategies that are not strictly dominated, violating property (b).[19]

Consider next a textbook Bertrand model with heterogeneous unit costs, $c_1 < c_2$. In undominated equilibria, the efficient firm serves the market at the unit cost of the inefficient firm, with profits $\bar{\pi}_1 = (c_2 - c_1)D(c_2)$ and $\bar{\pi}_2 = 0$ (Blume, 2003; Kartik, 2011; De Nijs, 2012).

In our production-in-advance framework, $r_1 = c_1$ and $r_2 = c_2$, so firm 1 is the strong firm. Firms still earn their Bertrand profits, but mix continuously 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 $\bar{\pi}_2 = 0$, k_1 is strictly increasing, and so $\bar{p}_1 = p^0$. The equilibrium is thus described by equation (1): The strong firm puts the rest of its mass on $\bar{p}$ and the weak firm on inactivity.

In contrast to the Bertrand outcome, with production in advance there is price dispersion, the strong firm may not serve the market, its inventory may remain unsold, and it sometimes prices above its monopoly price. As the α's tend to 1, the equilibrium does however converge to the Bertrand equilibrium in which firm 2 prices least aggressively (see Appendix C.5).

Bertrand with fixed costs. Consider next a textbook Bertrand model enriched with fixed costs ($A_i > 0$). These games have been studied by Sharkey and Sibley (1993), Marquez (1997), and Thomas (2002). Such games have also been reinterpreted as models of personalized pricing and advertising (Anderson, Baik, and Larson, 2015, 2019).

With (generically) $A_1 < A_2$, in the Bertrand game firms mix continuously on $[r, p^m)$, the strong firm (firm 1) puts the rest of its mass on p^m, and the weak firm on inactivity.

In our production-in-advance framework, firm i mixes over $[r, \bar{p})$ according to

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

Whether α_1 is larger or smaller than α_2 determines whether $\bar{p}_1$ is larger or smaller than $\bar{p}_2$, which then determines how firms distribute the rest of their mass. This feature provides a sharp illustration of the equilibrium characterization of Section 3.1.

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 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 weak firm on inactivity. The mass points are qualitatively similar to those under production to order with the exception that $\bar{p} > p^m$, in line with equilibrium property (e).

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

[19] Despite these qualitative differences, the quantitative features of the equilibrium are similar to those of nearby generic games. See Appendix C.4 for a formal statement of this result.

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

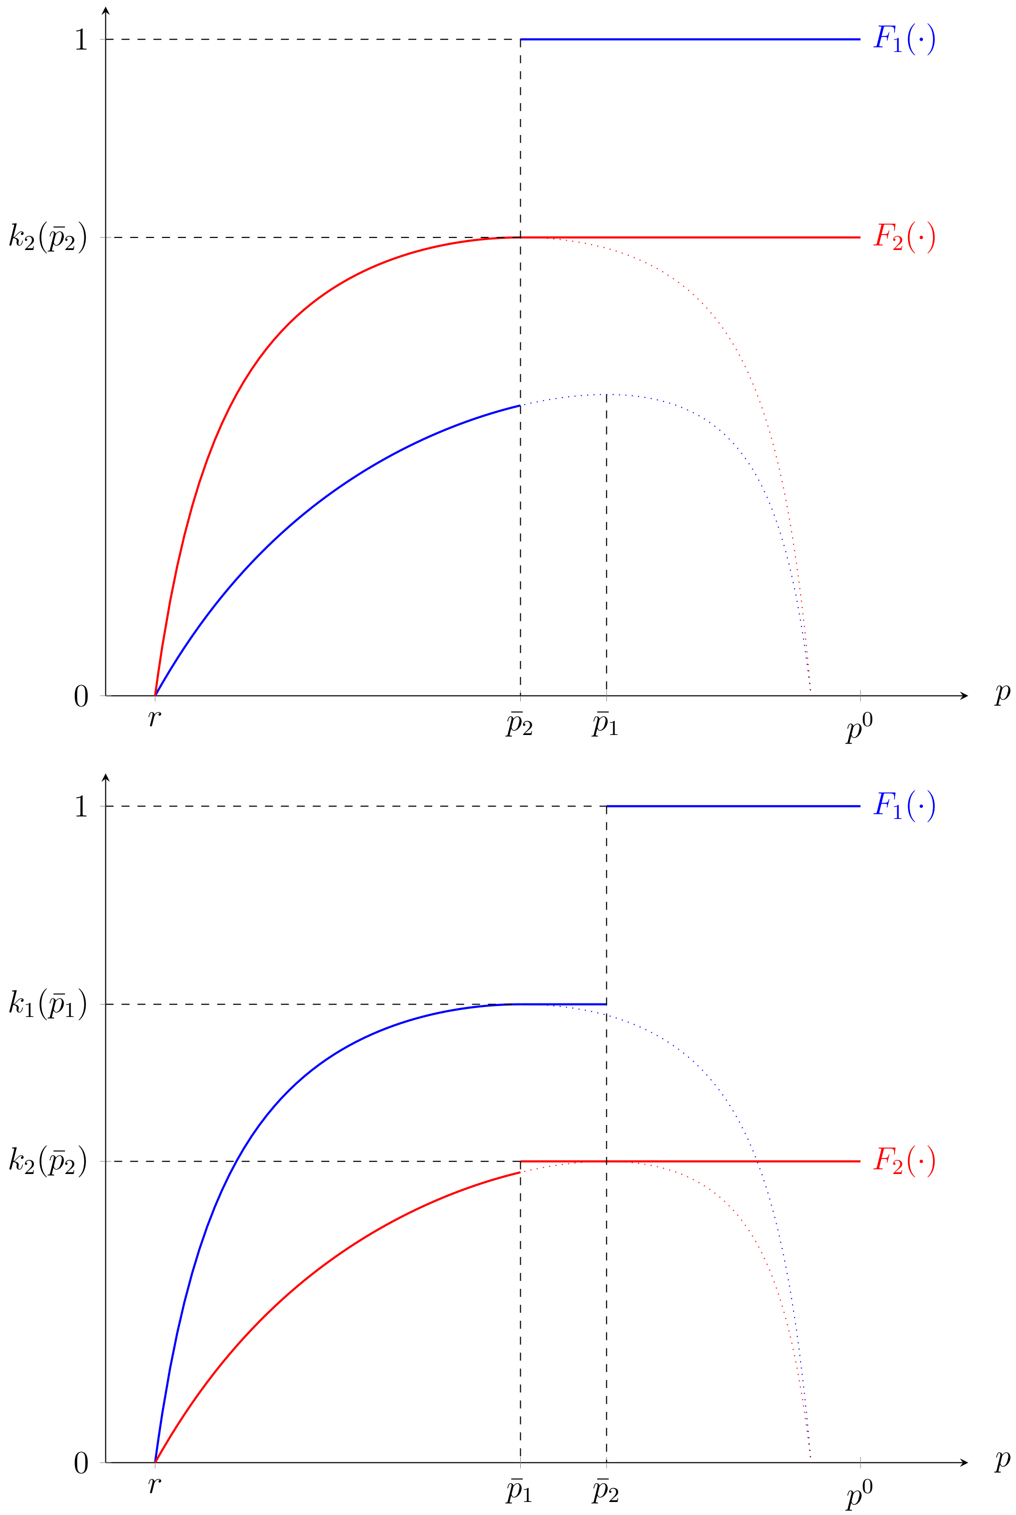

In both panels, $D(p) = 1 - p$, $c_1 = c_2 = 0.3$, $A_1 = 0.01$, $A_2 = 0.03$, and $\mu_1 = \mu_2 = 0$; in the top panel, $\alpha_1 = 0.9$ and $\alpha_2 = 0.1$; in the bottom panel, $\alpha_1 = 0.1$ and $\alpha_2 = 0.9$.

set high prices than firm 2 ($\bar{p}_1 < \bar{p}_2$). The equilibrium, represented 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, but the quantitative differences disappear in the limit.

Clearinghouse models. Consider next a symmetric textbook Bertrand model enriched with captive consumers. This model, with perfectly inelastic demand, was studied by Varian (1980). 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.

In that model but instead with production in advance, firms mix over $[r, p^0)$ according to

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

In a generic equilibrium, one firm puts the rest of its mass on targeting both its captive and the contested segment at p^0 and the other firm on targeting only its captive segment at p^0. Hence, both have reference prices. As inventory costs become fully recoverable, this equilibrium converges to an equilibrium of the resulting Bertrand game.[20] Importantly, both firms continue to have a mass point at the choke price in the limit—whereas there are no mass points in Varian's equilibrium. The following difference explains the discrepancy: In our setting, firms explicitly choose which segments to target, whereas in Varian's model, firms always target both their captive and the contested segment by assumption.

The production-to-order literature has previously incorporated the idea that firms may need to actively decide which segments to target by adding advertising costs to Varian's model (e.g., Baye and Morgan, 2001; Iyer, Soberman, and Villas-Boas, 2005). The equilibria that were characterized already feature mass points, but have the property that no firm advertises in the contested segment with a strictly positive probability. Thus, those equilibria differ from the (generic) equilibrium of our limiting game, where one of the firms always advertises. The sources of this discrepancy are twofold.

First, in Baye and Morgan (2001) and the literature that follows (e.g., Arnold, Li, Saliba, and Zhang, 2011; Shelegia and Wilson, 2016, 2019), a firm that does not advertise still receives demand from the contested segment provided its rival does not advertise either—so the contested segment may be captured for free. Instead, 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). The second source of discrepancy is the usual focus on the symmetric equilibrium. Iyer, Soberman, and Villas-Boas (2005) study that equilibrium in a Bertrand game, assuming as we do that a consumer remains unaware of the product unless it is targeted, and find that neither firm targets the contested segment with positive probability. However, they did not study the asymmetric equilibrium, which is the generic one.

[20] Because the game is non-generic, Proposition 4 cannot be applied to obtain Bertrand convergence. It is straightforward to adapt the proof to establish convergence manually.

Consumer surplus and social welfare can vary significantly from one equilibrium to the other. Genericity suggests that the asymmetric equilibrium is the robust one, and should therefore be used to address policy questions. Our work provides an avenue for such policy analysis under both production in advance and production to order.

4.3 Dynamics and Convergence

So far, we have studied a static, one-shot setting. We now show that the equilibrium behavior in the one-shot game coincides with the per-period equilibrium behavior in a dynamic discrete-time infinite-horizon game of complete information. We will see that the recoverability parameter α_i in the static game reflects firms' impatience in the dynamic game.

At the beginning of every period t, each firm i buys or sells inventory in a perfectly competitive wholesale market at a price $c > 0$, decides whether to pay the advertising cost A_i, and chooses its price p_i. Next, demand (as in Section 2) is realized and firm i incurs an additional per-unit cost $\widetilde{c}_i$ for each unit it sells. Unsold units are then carried over to the next period.[21] (Perishability can be easily introduced by a fraction of the inventory being lost per unit of time.) Firms maximize discounted profits, using a common discount factor $\delta \in (0, 1)$. To avoid collusive effects, we focus on Markov-perfect equilibria.[22]

This dynamic game has an equilibrium in which, in every period, firms choose their actions as if in the one-shot game with unit cost $c_i = c + \widetilde{c}_i$ and recoverability parameter $\alpha_i = (\widetilde{c}_i + \delta c)/(\widetilde{c}_i + c)$, buying or selling in the wholesale market the quantity needed to match their desired inventory. To see this, note that a firm's cost of selling one unit to consumers is equal to the additional cost $\widetilde{c}_i$ plus an opportunity cost equal to a fraction δ of the wholesale price c—the discounted value of one unit of inventory bought or sold in the wholesale market at the beginning of the next period.

Writing $\delta = e^{-r\Delta}$, this model provides microfoundations for salvage values reflecting the time value of money (r) and the time between periods of competition (Δ). Note that α_i decreases with r and Δ, and tends to one as these tend to zero. Connecting with property (f), this model predicts that in industries where inventory holdings are unobservable at the pricing stage, more patient firms tend to set lower prices. Moreover, connecting with Proposition 4, this model also predicts that an outcome close to Bertrand emerges if firms are patient enough or the time between periods of competition is short.

The relationship between discount factors and price levels is typically viewed through the lens of collusion models, and thus a positive relationship is expected. Our analysis uncovers that, in the absence of collusion, in production-in-advance industries those two variables are negatively related. This mechanism also suggests a possible countervailing link between

[21] In Online Appendix VII.4, we set out a simple model where consumers can stockpile: Stockpiling occurs in equilibrium, but the equilibrium structure is similar to the one in the main text.

[22] Complete information, which implies that inventories are observed at the end of each period, is assumed for simplicity. Our equilibrium remains a Markov-perfect equilibrium of an incomplete-information game in which firms never observe rivals' inventories.

interest rates and price levels—which is absent if inventories are observable, or in production-to-order and perfect-competition versions of the same model.

5 Extensions and Applications

In this section, we extend our analysis to the N-firm case and to incomplete information. We then provide applications of our framework that offer novel economic insights into taxation, merger analysis, horizontal information sharing, ex-ante investments (e.g., in persuasive advertising) and endogenous asymmetries, and vertical relations.

The N-firm case. The formal treatment can be found in Online Appendix I. As in Sections 3.1 and 3.2, we define the constrained game with N firms and prove that the sets of constrained and unconstrained equilibria coincide generically (Propositions I and II). We establish equilibrium existence in the constrained game using an argument similar to the proof of Siegel (2009)'s Corollary 1 (Proposition IV). Next, we focus on the generic case, where firms can be strictly ordered by their reach, i.e., $r_1 < r_2 < ... < r_N$, and set $r = r_2$. As in Siegel (2014c)'s Theorem 1, we show that, in any equilibrium, the strongest firm earns its winning payoff at r and all other firms earn their outside options (Proposition III).

If firms can be unambiguously ranked in their willingness to target the contested segment, then the equilibrium is unique: The two highest-ranked firms compete in the contested segment with the same strategies as in duopoly, while the other firms take their outside options. We provide two such ranking conditions. First, if there are no captive consumers, the condition is that for any given pair of firms, the firm with the lower unit cost has a lower fixed cost and more recoverable inventories. Second, if firms have identical unit costs, the condition is that for any given pair of firms, the firm with fewer captive consumers has lower advertising costs and more recoverable inventories.[23] (Propositions VI and VII.)

If such ranking conditions are not met, then equilibria are harder to characterize. At sufficiently low prices, it is still the case that only firms 1 and 2 are targeting the contested segment, and do so with the same CDFs as in duopoly (Proposition V). However, other firms may be targeting the contested segment at higher prices. We illustrate this with a three-firm example that violates our ranking conditions: Firms 1 and 2 have lower but less recoverable unit costs than firm 3. If firm 3's unit costs are not too recoverable, the unit-cost-advantage effect dominates and only firms 1 and 2 are active in the unique equilibrium. Otherwise, there is also a unique equilibrium, such that at sufficiently high prices, firm 2 becomes inactive and only firms 1 and 3 mix continuously (Proposition VIII).

If firms are symmetric, it is easy to characterize the unique symmetric equilibrium: All

[23] The all-pay contests literature has also identified ranking conditions that guarantee non-participation of some players. See, e.g., Baye, Kovenock, and de Vries (1992) and Theorem 2 in Siegel (2009).

firms randomize continuously according to the CDF

$$F(p) = 1 - \left(1 - \frac{w(p) - w(r)}{w(p) - l(p)}\right)^{\frac{1}{N-1}}$$

over $[r, \bar{p}]$, where $\bar{p}$ is the unique maximizer of the right-hand side. The remaining mass is put on the outside option. For the reason discussed at the end of Section 3.1, it seems however questionable to focus on the symmetric equilibrium.[24]

Incomplete information. The equilibrium under complete information involves mixed strategies, reflecting firms' underlying incentives to be unpredictable. To provide a purification argument in the spirit of Harsanyi (1973), we study an incomplete-information model. This is non-trivial as our framework involves bi-dimensional continuous actions, whereas Harsanyi's theorem applies to finite games and we are not aware of existing results that would apply in our setting. Formal statements and proofs can be found in Online Appendix III–V.

Consider an all-pay oligopoly with neither captive consumers nor advertising costs. Suppose that unit costs are drawn i.i.d. from a well-behaved probability distribution. If the realizations of c_1 and c_2 were public information, then the analysis in Section 3 would apply. We consider instead the case where those realizations remain private information.

We first study the constrained game, which involves analyzing an all-pay contest of incomplete information with non-monotonic winning functions—a class of models that has largely remained unstudied.[25] The constrained game has a pure-strategy Bayes-Nash equilibrium, which is unique under certain conditions; moreover, the sets of constrained and unconstrained equilibria coincide (Theorem I). Purification follows: The equilibrium price distribution of the incomplete-information game converges to the complete-information one as the cost distribution converges to a mass point (Proposition X).

Finally, we provide a Bertrand convergence result. The limiting Bertrand game under incomplete information was studied by Hansen (1988) and Spulber (1995). We show that the equilibrium of our model converges to the Bertrand equilibrium as inventory costs become fully recoverable (Proposition XI).

Efficiency and taxation. The equilibrium in an all-pay oligopoly features three types of distortions: First, firms price above marginal cost (the classic deadweight loss); second, some inventories may remain unsold; third, it is not always the most efficient firm that serves the contested segment. This raises the question of whether taxes or subsidies can alleviate those distortions, as they usually do in oligopoly models.

In Online Appendix VI, we study this question in the symmetric production-in-advance version of the textbook Bertrand model. We observe that any linear and symmetric taxation

[24] The symmetric equilibrium has disputable comparative statics: The expected price at which consumers purchase in the contested segment increases with the number of firms, as discussed in Online Appendix I.6.

[25] Kaplan, 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).

scheme results in an equilibrium in which each firm draws its price from a common CDF. We then find a surprising answer: If the social planner could choose any common CDF, he would choose the laissez-faire one (Proposition XII)—the equilibrium is second-best efficient.

For example, if firms face a tax or subsidy t for each unsold unit, then the equilibrium CDF of prices is $F(p) = (p - c)/(p - \alpha c + t)$. The result described above implies that the optimal t is zero. This remains true for combinations of symmetric per-unit and/or ad-valorem taxes on sales, output and unsold units, as all such policies give rise to a mixed-strategy equilibrium in which each firm draws its price from a common CDF.

While the proof with an arbitrary demand function is involved, the result can be easily illustrated in the case of unit inelastic demand. In that case, social welfare only depends on the probability β that a firm produces one unit. The expected social welfare is then given by

$$\left(1 - (1 - \beta)^2\right) p^0 - 2\beta c + \alpha \beta^2 c,$$

where the first term is expected gross utility, the second is expected production costs, and the third is expected salvaged costs. The β that maximizes social welfare is $(p^0 - c)/(p^0 - \alpha c)$, the probability that a firm is active in the laissez-faire equilibrium.

Determining the optimal tax policy in the context of the general model of Section 2 would be a different exercise for reasons that are outlined in the Online Appendix. A thorough analysis of this question is left for future research.

Welfare-increasing mergers. We have seen that production in advance gives rise to uncoordinated competition, which results in costly unsold inventories. This inefficiency disappears under monopoly, thus providing a novel source of efficiency gains from mergers. Those gains need to be weighted against the impact of mergers on the classic deadweight loss.

Consider the symmetric production-in-advance version of the textbook Bertrand model. If inventory costs are almost fully recoverable, the merger must reduce social welfare since the pre-merger outcome is close to efficient (as prices are near marginal cost and the loss from unsold inventories is close to zero). As inventory costs become less recoverable, pre-merger prices increase, and consumer surplus decreases. Since firms make zero profits, pre-merger social welfare decreases as inventories become less recoverable. Post-merger social welfare is on the other hand independent of α. A merger is therefore more likely to increase social welfare if inventory costs are more sunk. For instance, if demand is linear and costs are fully sunk, the merger raises social welfare if and only if the unit cost exceeds circa 16% of the choke price, and that threshold increases with α. (See Online Appendix VII.5 for a treatment of the linear-demand case.)

For the same reasons, a merger harms consumers if inventory costs are almost fully recoverable, and becomes more likely to benefit consumers as costs become less recoverable. In the linear-demand case, if inventory costs are fully sunk, the merger raises consumer surplus if and only if the unit cost exceeds circa 66% of the choke price, and that threshold increases with α. Intuitively, when inventory costs are sunk and high, firms are wary of being left with costly unsold inventories, thus source little, and often set prices above the monopoly level.

The discussion above suggests that a merger to monopoly reduces social welfare and consumer surplus if costs are low or salvage values are high, but otherwise may, and sometimes will, raise social welfare and consumer surplus. A better understanding of these effects in a general all-pay oligopoly is left for future research.

Inventory information. A classic antitrust question is whether horizontal information sharing should be discouraged. Earlier work has for example focused on information exchange on demand and costs (e.g., Novshek and Sonnenschein, 1982; Gal-Or, 1986; Raith, 1996; Vives, 2002). The analysis carried out in this paper allows us to investigate the effects of inventory or capacity information being made available to rivals before pricing decisions.

In the production-in-advance version of the textbook Bertrand model, when such information is available we are in the Kreps and Scheinkman (1983) framework. Firms then earn their Cournot profits in equilibrium (assuming efficient rationing). On the other hand, in the absence of that information, firms make zero profit, i.e., their Bertrand profits. Hence, firms unambiguously benefit from this information.

The consumer surplus and social welfare effects of inventory information are however unclear. We study this issue in a companion paper (Montez and Schutz, 2019a). We find that public information also benefits consumers if costs are high and mostly sunk, and market demand is not too 'curved.' Moreover it is more likely to raise social welfare if inventory costs are high and mostly sunk, and demand is sufficiently 'curved.'

In that paper, we also explore the additional concern that horizontal information sharing on inventories may facilitate collusion. We find a surprising answer: Inventory information may hinder collusion since, despite making short-run deviations less attractive, it weakens punishments. With public information, sustaining collusion may however require carrying excess inventories to deter deviations—whereas such excess inventories are not needed under private information. Thus, public information, despite eliminating inventory waste in a static setting, can actually encourage it in a dynamic one.

Ex-ante investments and endogenous asymmetries. An important aspect of our analysis is that it allows us to solve fully asymmetric duopoly models with production in advance, including the limiting case of production-to-order clearinghouse models. This is an essential building block in the study of models of competition with ex-ante investments.

Consider a non-trivial two-stage game in which firms are initially symmetric. In stage 1, each firm i makes a costly investment to render a fraction of the shoppers' segment captive to its product—this can be thought of as persuasive advertising. In the second stage, firms play the resulting all-pay oligopoly game.[26] (See Online Appendix VII.6 for a formal analysis and description of the game.)

This game has no subgame-perfect equilibrium in which both firms make symmetric and pure investment choices. To see this, note that, assuming $\mu_i \geq \mu_j$ and letting $C(\cdot)$ be the

[26] Our choice of investment variable is made for expositional simplicity. Similar insights would obtain for, e.g., investment in unit cost reductions.

investment cost, firms' profits in the continuation subgame are given by

$$\pi_i = \mu_i(p^m - c)D(p^m) - C(\mu_i) \quad \text{and} \quad \pi_j = (1-\mu_i)(r_i - c)D(r_i) - A - C(\mu_j).$$

Although those payoffs are continuous, we show in the Online Appendix that they are kinked at $\mu_i = \mu_j$ since, as can be seen from the expressions above, the weak firm (firm i) benefits more from investing than the strong firm. This kink breaks the quasi-concavity of the payoffs and rules out symmetric equilibria in pure strategies.

This argument implies that any subgame-perfect equilibrium in which the investment choices are pure must be asymmetric. We provide conditions under which such equilibria exist in the Online Appendix. In specifications where such equilibria do not exist, the Glicksberg theorem ensures the existence of an equilibrium in which the investment choice is mixed. Such mixing must also result in ex-post asymmetries.

There are thus forces in all-pay oligopolies, including the limiting case of pure clearinghouse models, that push them towards asymmetric outcomes in the presence of ex-ante investments.[27] This argument provides an additional motivation for the study of asymmetric models that goes beyond the genericity considerations discussed throughout the paper: Once firms' technologies are endogenized, the industry naturally gravitates towards asymmetry.

Vertical relations. Manufacturers often use buyback policies, whereby retailers can return unsold goods for a fraction of the wholesale price. An obvious explanation is aggregate demand uncertainty (e.g., Deneckere, Marvel, and Peck, 1996; Krishnan and Winter, 2007).

In a companion paper (Montez and Schutz, 2019b), we study a model in which a manufacturer offers linear and non-discriminatory retail contracts consisting of a wholesale and a buyback price. We use the equilibrium characterization of the present paper to show that buybacks allow the manufacturer to intensify retail competition and squeeze retail margins. On the other hand, buybacks induce retailers to source more inventories, some of which eventually remain unsold. Such overproduction is ultimately costly to the manufacturer.[28]

The optimal contract solves the trade-off between double marginalization and overproduction. We find that the optimal wholesale price increases with the upstream marginal cost, while the optimal return price (as a fraction of the wholesale price) decreases with that cost. Consistent with this result, generous return policies are common in the distribution of books and CDs, products that have a low cost relative to the retail price.

[27]This is surprising since symmetric outcomes are the norm in symmetric oligopoly models with ex-ante investments (d'Aspremont and Jacquemin, 1988; Suzumura, 1992; Leahy and Neary, 1997; Baye and Morgan, 2009). Münster (2007) finds asymmetric outcomes in an all-pay auction model with pre-bidding investments.

[28]The idea that returns intensify retail competition was first proposed by Padmanabhan and Png (1997) in a model with observable inventories. However, Wang (2004) showed that in such a model, retailers still use their inventory choice as a commitment to soften price competition regardless of the buyback price, as in Kreps and Scheinkman (1983). Wang concluded that returns do not affect the manufacturer's equilibrium profit. See Montez (2015) for a study of buybacks in a private-contracts setting with manufacturer opportunism.

6 Concluding Remarks

We introduced and studied a class of games where firms source unobservable inventories before simultaneously setting prices. Solving these games requires first studying a model where firms must source enough inventories to supply all their targeted demand. The equilibrium of that constrained game is generically unique and in mixed strategies—thus, there is price dispersion and some inventories may remain unsold. This is also the generically unique equilibrium of the unconstrained game, where firms freely choose inventories.

As inventory costs become fully recoverable, equilibrium behavior converges to an equilibrium of the game in which firms only choose prices and produce to order—the Bertrand game. Several benchmark outcomes of oligopoly theory, where production to order is assumed, can thus 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 equilibrium to situations where the value of unsold inventories falls short of their acquisition value.

We have assumed throughout that aggregate demand is deterministic. Under this assumption, the firms' strategies can be collapsed into a single dimension: As firms always source enough inventories to supply their targeted demand, the strategies reduce to a price distribution and potentially a mass point on the outside option. Introducing aggregate demand uncertainty often breaks this property, and the tools developed in this paper thus no longer apply.[29] The analysis of such models would require completely different techniques, the development of which is a challenging avenue for future research.

We introduced product differentiation by assuming 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 again that the firms' strategies can be collapsed into a single dimension. That property would disappear if we were to use other widely-used specifications of product differentiation that result in a smoother residual demand, which again would require completely different techniques.[30]

An advantage of this specification is that it nests clearinghouse models as limiting cases. To the best of our knowledge, our work is the first to provide a method for solving such models under full asymmetry. This is an essential building block in solving various applications, such as models with ex-ante investments. As we showed, such models seem to naturally gravitate towards asymmetric outcomes.

An interesting property of all-pay oligopolies is that equilibrium profits do not depend on the extent to which inventory costs are recoverable, since any such advantage is competed away through more aggressive pricing. Therefore, and perhaps surprisingly, firms have no

[29]As pointed out by a referee, similar issues can arise if there are N firms and some consumer segments only shop between partially-overlapping subsets of those firms.

[30]Smooth product differentiation does not result in a pure-strategy equilibrium. Indeed, in such a model with linear demand à la Shubik and Levitan (1980), there is no pure-strategy equilibrium. (A proof is available from the authors.) The reason is that inventory constraints introduce kinks in residual demand functions, which break the quasi-concavity of profit functions.

incentives to increase the salvage value of their inventories. As more aggressive pricing gives rise to gains in consumer surplus and thus social welfare, this provides a justification for public policies that result in more recoverable and less perishable inventories, even before accounting for environmental concerns.

Appendix

A Equilibrium Analysis in the Constrained Game

In this section, we state and prove a series of technical lemmas that jointly imply Proposition 1. Then, in Proposition A below, we provide a complete characterization of equilibria in the non-generic case. Consider an all-pay oligopoly satisfying the assumptions made at the beginning of Section 2. 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 7 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 $p^n \underset{n \to \infty}{\longrightarrow} \hat{p}$ and F_i puts no mass on $\{p^n\}$ for every n. Then, for every n,

$$\begin{aligned}
\bar{\pi}_j &\geq (1 - F_i(p^n))w_j(p^n) + F_i(p^n)l_j(p^n), \\
&\underset{n \to \infty}{\longrightarrow} (1 - F_i^-(\hat{p}))w_j(\hat{p}) + F_i^-(\hat{p})l_j(\hat{p}), \\
&> (1 - F_i(\hat{p}))w_j(\hat{p}) + (F_i(\hat{p}) - F_i^-(\hat{p}))t_j(\hat{p}) + F_i^-(\hat{p})l_j(\hat{p}), \\
&= \tilde{\pi}_j.
\end{aligned}$$

□

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 $p^n \underset{n\to\infty}{\longrightarrow} \check{p}$ 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:

$$\bar{\pi}_i = (1 - F_j(p^n))w_i(p^n) + F_j(p^n)l_i(p^n) \underset{n\to\infty}{\longrightarrow} l_i(\check{p}) \leq o_i < \bar{\pi}_i.$$

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 $(p^n)_{\geq 0}$ such that $p^n \underset{n\to\infty}{\longrightarrow} \check{p}$ 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:

$$\bar{\pi}_j = (1 - F_i(p^n))w_j(p^n) + F_i(p^n)l_j(p^n) \underset{n\to\infty}{\longrightarrow} (1 - F_i(r))w_j(r) + F_i(r)l_j(r) < w_j(r). \qquad \square$$

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

$$\frac{1 - \mu_i - \mu_j}{1 - \mu_j} k_j(p) = \frac{(p - c_i)D(p) - (r - c_i)D(r)}{(p - \alpha_i c_i)D(p)},$$

$$= \frac{p - c_i}{p - \alpha_i c_i} + (r - c_i)D(r)\Phi\left(\log\left((p - \alpha_i c_i)D(p)\right)\right),$$

where $\Phi(x) = -\mathrm{e}^{-x}$. Since Φ is concave and increasing and $p \mapsto (p - \alpha_i c_i)D(p)$ is log-concave, it follows that $p \mapsto \Phi\left(\log\left((p - \alpha_i c_i)D(p)\right)\right)$ 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. □

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 1. 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. 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 CDFs 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 CDFs 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 CDFs 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 1. Lemma D pins down the equilibrium CDFs on $[r, \bar{p})$. The mass that remains can then be distributed over $\bar{p}_1$, $\bar{p}_2$, or the firms' outside options as described in the statement of the proposition. $\square$

B Equilibrium Behavior in the Unconstrained Game

The goal of this section is to prove Proposition 3. We introduce notation and state preliminary lemmas in Section B.1. We then prove the proposition in Section B.2. We also briefly discuss the equilibrium multiplicity that can arise in a non-generic all-pay oligopoly in Section B.3. In the following, we consider an all-pay oligopoly satisfying the assumptions made at the beginning of Section 2.

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

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

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

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

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

Next, we simplify the action sets by removing redundant 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 $q_i = \mu_i D(p_i^m)$.[31] 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:

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

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,

[31] If $\mu_i = 0$, 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.

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

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

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

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

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

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, $\Delta_i(p_i, \hat{q}_i, \sigma_j) \underset{p_i \uparrow \hat{p}_i}{\longrightarrow} \Delta_i^-(\hat{p}_i, \hat{q}_i, \sigma_j)$.*

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

$$\Delta_i\left(p^n, \hat{q}_i, \sigma_j\right) = \left(\varphi_j\left((p^n, p^0]\right) + \tau_j\right)\left(\hat{q}_i - \mu_i D(p_i)\right) + \int_{\substack{c_j \leq p_j < \hat{p}_i \\ \mu_j D(p_j) < q_j \leq (1-\mu_i) D(p_j)}} Z^n(p_j, q_j) d\sigma'_j(p_j, q_j), \tag{5}$$

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

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

Next, we turn our attention to the term in equation (5). 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

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

which proves the lemma. □

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

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

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 $(\hat{p}_i,\hat{q}_i)\in\mathcal{A}_i'$, $\bar{\pi}_i\ge\pi_i^-(\hat{p}_i,\hat{q}_i,\sigma_j)$. 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 3

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

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

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

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

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

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 $(p^n)_{n\ge1}$ be a strictly increasing sequence of prices that converges to $\hat{p}$. Note that

$$\{\hat{p}\}\times(\mu_i D(\hat{p}),(1-\mu_j)D(\hat{p}))=S_i(\hat{p})\setminus\bigcup_{n\ge1}S_i(p^n).$$

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

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

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,

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

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

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

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,

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

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, $p'^n \underset{n\to\infty}{\longrightarrow} \hat{p}$, 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 $p^n \underset{n\to\infty}{\longrightarrow} \hat{p}$, 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,

$$\begin{aligned}\bar{\pi}_i &= \pi_i^-(p^n,q^n,\sigma_j),\\ &= \mu_i(p^n-c_i)D(p^n) - A_i + \Big((p^n-\alpha_i c_i)\left(\varphi_j\left([p^n,p^0]\right)+\tau_j\right) - (1-\alpha_i)c_i\Big)(q^n-\mu_i D(p^n))\\ &\quad + (p^n-\alpha_i c_i)\int_{\substack{\hat{p}<p_j<p^n \\ \mu_j D(p_j)<q_j\leq(1-\mu_i)D(p_j)}} Z_i\left(p^n,p_j,q^n-\mu_i D(p^n),q_j-\mu_j D(p_j)\right)d\sigma_j(p_j,q_j).\end{aligned}$$ [32]

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

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

[32] The 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 $q_j = (1-\mu_i)D(\hat{p})$.

$$+ (p^n - \alpha c_i) \int_{\substack{\hat{p}<p_j<p^n \\ \mu_j D(p_j)<q_j \leq (1-\mu_i)D(p_j)}} Z_i\left(p^n, p_j, q - \mu_i D(p^n), q_j - \mu_j D(p_j)\right) d\sigma_j(p_j, q_j).$$

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 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

$$\int_{\substack{\hat{p}<p_j<p^n \\ \mu_j D(p_j)<q_j \leq (1-\mu_i)D(p_j)}} Z_i\left(p^n, p_j, q^n - \mu_i D(p^n), q_j - \mu_j D(p_j)\right) d\sigma_j(p_j, q_j)$$
$$\leq \int_{\substack{\hat{p}<p_j<p^n \\ \mu_j D(p_j)<q_j \leq (1-\mu_i)D(p_j)}} D(\hat{p}) d\sigma_j(p_j, q_j) = \varphi_j((\hat{p}, p^n))D(\hat{p}) \underset{n\to\infty}{\longrightarrow} 0.$$

We obtain the following contradiction:

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

□

B.3 Equilibrium Multiplicity in Non-Generic Cases

In this subsection, we discuss 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 (see the end of Section 4.2).

We first discuss constrained equilibria. By Proposition A, 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 3 can be adapted to show that in any equilibrium, conditional on pricing at $p < p^0$ 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.

C Convergence Results

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. In Section C.2, we establish the continuity of p_i^m, r_i, $\bar{p}_i$, and k_i in the parameters of the model. We prove Proposition 4 in Section C.3. In Section C.4, we state and prove Proposition B, which says that, as parameters converge to the parameter vector of the symmetric production-in-advance version of the textbook Bertrand model, the equilibrium converges weakly to the unique equilibrium of the limiting game. Bertrand convergence in the production-in-advance version of the asymmetric Bertrand model is established 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 4. 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 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 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

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

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

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

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 7). 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 CDFs 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 CDFs $(\tilde{F}_1, \tilde{F}_2)$ 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 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

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

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

$$W_i(p;\gamma) = (1-\mu_j)(p-c_i)D(p) - A_i,$$
$$L_i(p;\gamma) = \mu_i(p-c_i)D(p) - (1-\mu_i-\mu_j)(1-\alpha_i)c_iD(p) - A_i,$$

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

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

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.[33] 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

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

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

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 $p \in (R(\gamma), p^0)$,

$$K_i(p;\gamma) = \frac{1-\mu_i}{1-\mu_i-\mu_j}\left(\frac{p-c_j}{p-\alpha_jc_j} - \frac{(R(\gamma)-c_j)D(R(\gamma))}{(p-\alpha_jc_j)D(p)}\right). \tag{6}$$

[33] Assume 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 $\gamma^n \underset{n\to\infty}{\longrightarrow} \gamma \in \Gamma'$, 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.

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

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

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 $\lim_{n \to \infty} K_i(p; \gamma^n) = 0 = K_i(p; \gamma)$. Next, suppose that $p > R(\gamma)$. Then, by continuity of R, $p > R(\gamma^n)$ for n high enough. Taking limits in equation (6), we obtain that $\lim_{n \to \infty} K_i(p; \gamma^n) = K_i(p; \gamma)$.

Continuity of $\bar{P}_i$. We now show that $\bar{P}_i$ is continuous on its domain $\overline{\Gamma}_i$. 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 $\lim_{n \to \infty} K_i(\check{p}; \gamma^n) = K_i(\check{p}; \gamma)$ and $\lim_{n \to \infty} K_i(\hat{p}; \gamma^n) = K_i(\hat{p}; \gamma)$, 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 $\bar{P}_i(\gamma^n) \underset{n \to \infty}{\longrightarrow} \bar{P}_i(\gamma)$, 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 $\lim_{n \to \infty} K_i(R(\gamma); \gamma^n) = K_i(R(\gamma); \gamma) = 0$. Let $\varepsilon \in \big(0, K_i(\bar{P}_i(\gamma); \gamma)\big)$. 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 $\lim_{n \to \infty} K_i(p; \gamma^n) = K_i(p; \gamma)$, we have that $K_i(p; \gamma^n) \in (0, \varepsilon)$ for n high enough. Moreover, since $\lim_{n \to \infty} \bar{P}_i(\gamma^n) = \bar{P}_i(\gamma)$, 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 $\lim_{n \to \infty} K_i(R(\gamma); \gamma^n) = K_i(R(\gamma); \gamma) = 0$.

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 4

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 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, $\lim_{n\to\infty} R_i(\gamma^n) = R_i(\gamma)$ and $\lim_{n\to\infty} \bar{P}_i(\gamma^n) = \bar{P}_i(\gamma)$ 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 1, 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, $F_1^n(p) = K_1(p;\gamma^n) \underset{n\to\infty}{\longrightarrow} K_1(p;\gamma)$. The same line of reasoning implies that $F_1^n(p) \underset{n\to\infty}{\longrightarrow} 1$ 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 $V' = \left(\bar{P}_2(\gamma) - \varepsilon, p^0\right) \times V$. Then, for every $(p, \tilde{\gamma}) \in V'$, $K_2(p;\tilde{\gamma})$ is given by equation (6), which is clearly continuous in $(p, \tilde{\gamma})$. Hence, K_2 is continuous on V'. Since $(\bar{P}_2(\gamma^n), \gamma^n) \underset{n\to\infty}{\longrightarrow} (\bar{P}_2(\gamma), \gamma)$, it follows that $(\bar{P}_2(\gamma^n), \gamma^n) \in V'$ for n high enough. By continuity, it follows that $K_2(\bar{P}_2(\gamma^n);\gamma^n) \underset{n\to\infty}{\longrightarrow} K_2(\bar{P}_2(\gamma), \gamma)$. 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 $F_2^n(\bar{P}_2(\gamma)) \underset{n\to\infty}{\longrightarrow} F_2(\bar{P}_2(\gamma))$. 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 $\{1,2,\ldots\}$ to $\mathcal{N}$. (If $\mathcal{N}$ is finite, there is nothing to prove.) For every $n \geq 1$,

$$F_2^{\phi(n)}(\bar{P}_2(\gamma)) = K_2(\bar{P}_2(\gamma^{\phi(n)}),\gamma^{\phi(n)}) \underset{n\to\infty}{\longrightarrow} K_2(\bar{P}_2(\gamma),\gamma) = F_2(\bar{P}_2(\gamma)).$$

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

$$F_2^{\zeta(n)}(\bar{P}_2(\gamma)) = K_2(\bar{P}_2(\gamma),\gamma^{\zeta(n)}) \underset{n\to\infty}{\longrightarrow} K_2(\bar{P}_2(\gamma),\gamma) = F_2(\bar{P}_2(\gamma)).$$

Therefore, $F_2^n(\bar{P}_2(\gamma)) \underset{n\to\infty}{\longrightarrow} F_2(\bar{P}_2(\gamma))$, 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, $\lim_{n\to\infty} R_i(\gamma^n) = R_i(\gamma)$ and $\lim_{n\to\infty} \bar{P}_i(\gamma^n) = \bar{P}_i(\gamma)$ 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 1, 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 $(\bar{P}_i(\gamma), \gamma)$ $(i = 1, 2)$; $K_i(\bar{P}_i(\gamma^n); \gamma^n) \underset{n\to\infty}{\longrightarrow} K_i(\bar{P}_i(\gamma), \gamma)$ $(i = 1, 2)$; $(F_1^n)_{n\geq 1}$ 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. □

C.4 Convergence to the Production-in-Advance Version of the Textbook Bertrand Model

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

Proposition B. *Suppose D is continuous at p^0, and let $(\gamma^n)_{n\geq 0}$ be a sequence of parameter vectors that converges to the symmetric parameter vector of the production-in-advance version of the textbook Bertrand model. 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 the production-in-advance version of the textbook Bertrand model.*

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 Section 4.2 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 with parameter vector γ^n. By Propositions 1 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 $\bar{P}_i(\gamma^n) \underset{n\to\infty}{\longrightarrow} \bar{P}_i(\gamma) = p^0$. Therefore, $p < \bar{P}_i(\gamma^n)$ and $F_i^n(p) = K_i(p; \gamma^n)$ for n high enough. By Lemma G,

$$F_i^n(p) = K_i(p; \gamma^n) \underset{n\to\infty}{\longrightarrow} K_i(p; \gamma) = F_i^*(p).$$

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^*. □

C.5 Bertrand Without Fudge

Since the standard model of Bertrand competition with heterogeneous marginal costs discussed in Section 4.2 is non-generic, we cannot apply Proposition 4 to study Bertrand convergence. It is, however, straightforward to adapt the argument in Section C.3 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 $\alpha_i^n \underset{n\to\infty}{\longrightarrow} 1$ 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 $\bar{P}_2(\gamma^n) \underset{n\to\infty}{\longrightarrow} \bar{P}_2(\gamma) = P_1^m(\gamma)$.

By Proposition 1, the equilibrium profile of CDFs 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.

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.

Anderson, S. P., A. Baik, and N. Larson (2019): "Price Discrimination in the Information Age: Prices, Poaching, and Privacy with Personalized Targeted Discounts," CEPR Discussion Papers 13793, C.E.P.R. Discussion Papers.

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.

——— (2009): "Brand and Price Advertising in Online Markets," *Management Science*, 55(7), 1139–1151.

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.

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.

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, Y., and M. H. Riordan (2008): “Price-increasing competition,” *RAND Journal of Economics*, 39(4), 1042–1058.

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.

d'Aspremont, C., and A. Jacquemin (1988): “Cooperative and Noncooperative R&D in Duopoly with Spillovers,” *American Economic Review*, 78(5), 1133–37.

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., H. P. Marvel, and J. Peck (1996): “Demand Uncertainty, Inventories, and Resale Price Maintenance,” *The Quarterly Journal of Economics*, 111(3), 885–913.

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.

Gal-Or, E. (1986): “Information Transmission—Cournot and Bertrand Equilibria,” *Review of Economic Studies*, 53(1), 85–92.

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.

Harrington, J. E. (1989): "A re-evaluation of perfect competition as the solution to the Bertrand price game," *Mathematical Social Sciences*, 17(3), 315 – 328.

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.

Iyer, G., D. Soberman, and J. M. Villas-Boas (2005): "The Targeting of Advertising," *Marketing Science*, 24(3), 461–476.

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.

Kocas, C., and T. Kiyak (2006): "Theory and evidence on pricing by asymmetric oligopolies," *International Journal of Industrial Organization*, 24(1), 83–105.

Kreps, D., and J. Scheinkman (1983): "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *Bell Journal of Economics*, 14(2), 326–337.

Krishnan, H., and R. A. Winter (2007): "Vertical Control of Price and Inventory," *American Economic Review*, 97(5), 1840–1857.

Leahy, D., and J. P. Neary (1997): "Public Policy towards R&D in Oligopolistic Industries," *American Economic Review*, 87(4), 642–62.

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.

Maggi, G. (1996): "Strategic Trade Policies with Endogenous Mode of Competition," *American Economic Review*, 86(1), 237–258.

Marquez, R. (1997): "A note on Bertrand competition with asymmetric fixed costs," *Economics Letters*, 57(1), 87–96.

Marvel, H. P., and J. Peck (1995): "Demand Uncertainty and Returns Policies," *International Economic Review*, 36(3), 691–714.

Maskin, E. (1986): "The Existence of Equilibrium with Price-Setting Firms," *American Economic Review*, 76(2), 382–86.

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.

Montez, J. (2015): "Controlling opportunism in vertical contracting when production precedes sales," *RAND Journal of Economics*, 46(3), 650–670.

Montez, J., and N. Schutz (2019a): "Private versus Public Inventories," mimeo.

——— (2019b): "When Returns Intensify Retail Competition," mimeo.

Münster, J. (2007): "Contests with investment," *Managerial and Decision Economics*, 28(8), 849–862.

Narasimhan, C. (1988): "Competitive Promotional Strategies," *The Journal of Business*, 61(4), 427–449.

Novshek, W., and H. Sonnenschein (1982): "Fulfilled Expectations Cournot Duopoly with Information Acquisition and Release," *Bell Journal of Economics*, 13(1), 214–218.

Padmanabhan, V., and I. P. L. Png (1997): "Manufacturer's Return Policies and Retail Competition," *Marketing Science*, 16(1), 81–94.

Raith, M. (1996): "A General Model of Information Sharing in Oligopoly," *Journal of Economic Theory*, 71(1), 260–288.

Ramey, V. A. (1991): "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy*, 99(2), 306–334.

Sharkey, W., and D. S. Sibley (1993): "A Bertrand model of pricing and entry," *Economics Letters*, 41(2), 199–206.

Shelegia, S., and C. Wilson (2019): "All-Pay Contests with Costly Participation and Default Allocations," mimeo.

Shelegia, S., and C. M. Wilson (2016): "A generalized model of sales," Economics Working Papers 1541, Department of Economics and Business, Universitat Pompeu Fabra.

Shubik, M., and R. Levitan (1980): *Market Structure and Behavior*. Harvard University Press, Cambridge.

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 all-pay auctions with interdependent valuations," *Journal of Economic Theory*, 153(C), 684–702.

——— (2014b): "Asymmetric Contests with Head Starts and Nonmonotonic Costs," *American Economic Journal: Microeconomics*, 6(3), 59–105.

——— (2014c): "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.

Suzumura, K. (1992): "Cooperative and Noncooperative R&D in an Oligopoly with Spillovers," *American Economic Review*, 82(5), 1307–20.

Tasnadi, A. (2004): "Production in advance versus production to order," *Journal of Economic Behavior & Organization*, 54(2), 191–204.

——— (2018): "Corrigendum to "Production in advance versus production to order"," Unpublished manuscript.

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.

Vives, X. (2002): "Private Information, Strategic Behavior, and Efficiency in Cournot Markets," *RAND Journal of Economics*, 33(3), 361–376.

Wang, H. (2004): "Do Returns Policies Intensify Retail Competition?," *Marketing Science*, 23(4), 611–613.

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.