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

Joao Montez[†] Nicolas Schutz[‡]

November 12, 2018

Abstract

We generalize the classical Bertrand-type equilibrium to a production-in-advance setting, where firms source unobservable inventories first and then simultaneously set prices. In the limit, as inventory costs become fully recoverable, the equilibrium of our model converges to an equilibrium of the associated Bertrand game, where instead firms first choose prices and then produce to order. The analysis involves solving an all-pay contest with outside options, non-monotonic winning and losing functions, and conditional investments. We find that production in advance leads to each firm occasionally holding a sale relative to a reference price, which results in rivals being unable to sell all of their inventory. Our analysis suggests that in industries with production in advance, the choice between the Cournot equilibrium and our generalization of the Bertrand-type equilibrium should depend on whether or not firms observe their rivals' inventory choices before setting prices.

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

1 Introduction

Retail markets are typically characterized by production in advance, as each store chooses not only a price but also an inventory level—a quantity sourced from suppliers to be made readily available to consumers. It is however hard to observe rivals' inventories, and historical data is unlikely to provide a reliable estimate—as inventories are transient by nature. Stores therefore have to make price and inventory choices without knowledge of rivals' choices.

*We thank Simon Anderson, Pierre Boyer, Bruno Jullien, Sébastien Mitraille, Volker Nocke, Hans-Theo Normann, Martin Peitz, Régis Renault, Markus Reisinger, Patrick Rey, Armin Schmutzler, Andre Speit, Yossi Spiegel, Roland Strausz, Thomas Tröger, Thomas von Ungern-Sternberg, Ralph Winter, as well as participants at numerous seminars and conferences for helpful comments and suggestions. The authors 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.

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

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

We introduce and study a class of such games, which we call *all-pay oligopolies*. This class accounts for asymmetries between firms, heterogeneity in consumer tastes, and the coexistence of informed and uninformed consumers. As unsold inventories often have a salvage value, and sales or post-sales services may be costly, we allow a fraction of the unit cost to remain variable and incurred only once a unit is sold (while the remainder is sunk).[1]

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. We develop a method to obtain the closed-form characterization of the equilibrium set of such constrained games, and show that generically, there is a unique equilibrium. Next, we show that this equilibrium is also the generically unique equilibrium of the *unconstrained game* (where firms freely choose inventory levels).

In equilibrium, each store randomizes its price, ordering a low inventory when it sets a high price, and a high inventory when it holds a sale. Because each store holds enough inventory to serve all its targeted demand, the aggregate inventory level often exceeds 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 associated Bertrand game, in which stores only choose prices and produce to order (i.e., source inventories to meet demand only after consumers made purchase decisions). The equilibrium is thus said to be *Bertrand convergent*. 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 significantly short of their acquisition value.

This insight stands in contrast to Kreps and Scheinkman (1983)'s well-known result. Assuming that inventory choices become *observable* before the pricing stage, they conclude that situations where almost all the unit cost "is incurred subsequent to the realization of demand (situations that will look very Bertrand-like) will still give the Cournot outcome." The conventional interpretation of this result is that a Cournot analysis is appropriate under production in advance, whereas Bertrand should be preferred under production to order (see, e.g., Belleflamme and Peitz, 2010, pp. 66–67).

Our analysis suggests a more nuanced model selection for production-in-advance settings, which explicitly takes into account whether inventory choices are *observed* by rivals or *not*. In the former case, the Cournot outcome remains a reasonable benchmark. In the latter case, the Kreps-Scheinkman mechanism fails as a low inventory choice can no longer provide a commitment to soften price competition. Then, our Bertrand convergence result suggests

[1] Salvage value may arise from sales in outlet stores, buyback contracts with manufacturers (e.g., Montez, 2015), or future sales (see the dynamic extension discussed in the conclusion).

that a Bertrand approach is better justified if most of the unit cost can be salvaged, and our analysis characterizes the equilibrium behavior for lower salvage values.

Large unsold inventories, as predicted by our model, are observed in several industries, including grocery and apparel. The usual explanation relies on exogenous market demand uncertainty. By contrast, in our model, market demand is deterministic, but *store-level* demand becomes 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 another 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 which is 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 frequently offering discounts to attract price-conscious consumers to their stores.[2]

The equilibrium under complete information involves mixed strategies, reflecting stores' underlying incentives to be unpredictable. To provide a purification argument in the spirit of Harsanyi (1973), we add privately-observed shocks to unit costs in a symmetric version of the model.[3] In this incomplete-information game, we find that for any cost realization there is a unique price-inventory pair that solves a given store's profit maximization problem, thus resulting in a strict pure-strategy equilibrium. The equilibrium price distribution of the incomplete-information game converges to the complete-information one as the cost distribution converges to a mass point.

Finally, we provide a preliminary investigation of whether taxes or subsidies alleviate the distortions that arise in all-pay oligopoly games, as is usually the case in oligopoly models. In a symmetric version of the model, we find a surprising result: The equilibrium outcome cannot be improved by standard taxation. Laissez-faire is thus second-best efficient.

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 with complete information, as thoroughly studied by Siegel (2009, 2010), there is a fixed number of prizes, each player submits a score, and prize winners are the players

[2]Other explanations include adjustment to cost shocks (Blinder, 1986), increasing returns to scale upstream (Ramey, 1991), and stochastic demand (Kahn, 1987; Lee, Padmanabhan, and Whang, 1997b).

[3]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.

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 a player 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, under both complete and incomplete information.

Under complete information, as inventory costs become fully variable, the equilibrium converges to the 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 and Morgan, 2001; Iyer, Soberman, and Villas-Boas, 2005; Shelegia and Wilson, 2016). It has long 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 first-price auction—the Bertrand outcome of that model.

Under incomplete information, we establish equilibrium existence and uniqueness in an all-pay contest with non-monotonic winning functions—a class of models that has largely remained unstudied.[4] The limiting Bertrand game under incomplete information was studied by Spulber (1995), as an oligopoly game, and Hansen (1988), as a procurement auction with variable demand—which can also be viewed through the conceptual lens of our model. On the other hand, we study a production-in-advance model. We show that the equilibrium of our model converges to Hansen and Spulber's Bertrand equilibrium as inventory costs become fully recoverable.

An important difference relative to the literatures discussed above is that our objective

[4]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).

is not to understand the constrained game, but rather 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 that are naturally absent in the all-pay contests literature and in the industrial organization literature on production to order.

A small number of papers has analyzed price-inventory models similar to ours. Maskin (1986) proves equilibrium existence in a class of price-inventory games with two firms. 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 in industrial organization 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 of equilibrium uniqueness provided in those earlier papers omit important non-trivial steps or contain several inaccuracies.[5] 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. Yet, a counterfactual with unobservable inventories had not been properly investigated. Our work shows that, to soften price competition, production in advance must be combined with inventory observability: A Bertrand-like, intense form of competition arises under production in advance if inventories remain unobservable.

Finally, Deneckere and Peck (1995) study a symmetric model with simultaneous price-inventory choices but stochastic demand. In a pure-strategy equilibrium candidate, all stores choose the same price (i.e., no price dispersion), consumers are rationed when demand is high, and stores are unable to sell all inventories when demand is low. Unsold inventories must however disappear as demand becomes certain (whereas in our model, some units remain unsold despite the absence of market demand uncertainty). More fundamentally, pure-strategy equilibria 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).

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

The paper is organized as follows. Section 2 studies equilibrium behavior in all-pay oligopolies. Section 3 discusses equilibrium properties, including Bertrand convergence. Section 4 provides applications and relates them to classical Bertrand analysis. Section 5 analyzes an all-pay oligopoly with incomplete information. Section 6 studies efficiency and taxation. Section 7 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, ∞).[6] 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. We assume that a fraction of consumers $\mu_i \in [0, 1)$ only wishes to buy from firm i. Those consumers are referred to as firm i's *captive segment*. The remaining fraction $1 - \mu_1 - \mu_2 \in (0, 1]$ of consumers, the *shoppers*, wishes to buy from the firm with the lowest price. As a whole, we call those consumers the *contested segment*. 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 how rationing works across segments, we assume that each firm serves its captive consumers first.[7]

Firm i can pay a fixed advertising cost of $A_i \geq 0$ to make the shoppers aware of its product and price. This cost is akin to a fixed cost of production if there are no captive consumers, i.e., if $\mu_1 = \mu_2 = 0$.

Firms 1 and 2 simultaneously decide how many units to source, whether to pay the advertising cost, and which 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 the set of 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. If instead firm i sets a price $p_i > c_i$ and decides to serve

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

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

both its captive and the contested segment, then it captures at best the whole contested segment, in which case its payoff is

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

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 $o_i < w_i(p_i^m)$. Note that 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. We say that firm i is *strong* if $r_i < r$ and *weak* if $r_i = r$. As the weak firm will never price below r, this is the highest price at which the strong firm can be sure to capture the contested segment—the limit price. Naturally, the strong firm can price more aggressively than the weak one while still earning more than its outside option. If $p_i^m < r$, the game is again dominance solvable: Firm i serves its captive and the contested segment at its monopoly price, whereas firm j serves only its captive consumers at its own monopoly price.

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. Moreover, we are mainly 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 to obtain our results. When stating formal results, we mention the key (weaker) conditions required inside parentheses. We also provide some equilibrium characterization results for non-generic games. As we show, those games often have a continuum of equilibria, but this multiplicity disappears when the game is slightly perturbed and becomes generic.

The equilibrium characterization proceeds in two steps. In Section 2.1, 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. Next, in Section 2.2, we show that this equilibrium is also the generically unique equilibrium of the *unconstrained* game, where firms can freely choose their inventories. In Section 2.3, we generalize these steps to the N-firm case whenever possible.

2.1 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 log-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 $1-F_i(p^0)$.[8]

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. Since r is the infimum of the support of both firms' strategies, the expected payoff of firm i in this putative equilibrium is given by $w_i(r)$. For firm i to be indifferent between all the prices in $[r, \bar{p})$, it has to be the case that for every $p \in [r, \bar{p})$,

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

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

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

The log-concavity of D implies that either k_j is single-peaked and achieves a global 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

[8] 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.

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

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

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 $\bar{p}$.[9] 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 is the case where 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. In this setting, this means selecting the equilibrium where both firms put all of their remaining mass on their outside option.

The focus on the symmetric equilibrium can be unsatisfactory in a production-in-advance setting for the following reason. 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.

2.2 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:[10]

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

[9]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.

[10]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 (F_1, F_2) of the constrained game is also an equilibrium of the unconstrained game.*

The next question is whether the unconstrained game also has equilibria in which at least one firm does not always source enough inventory to supply its targeted demand. We find a negative answer for generic games. The following heuristic argument provides the intuition.

Assume for a contradiction that such an equilibrium exists. 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 $\mu_i D(p) \leq q < (1-\mu_j)D(p)$, and let $F_j(p)$ denote the probability that firm j chooses a pair (p_j, q_j) such that $p_j < p$ and $q_j > \mu_j D(p_j)$. Then, firm i makes an expected profit of:[11]

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

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 $A_i > 0$,

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

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:

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.*

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

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

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.

2.3 The N-Firm Case

We now discuss the case with $N \geq 3$ firms, providing a non-exhaustive summary of the formal treatment that can be found in Online Appendix I. (The remainder of the paper focuses on the two-firm case.)

As we did in Section 2.1, we can define the constrained game with N firms. Using techniques similar to those applied in Section 2.2, we prove that the sets of constrained and unconstrained equilibria coincide generically (see Propositions I and II). From now on, we focus on the constrained game. We establish equilibrium existence using an argument similar to the one in the proof of Corollary 1 in Siegel (2009) (see Proposition IV). Focusing on the generic case, where $r_1 < r_2 < \ldots < r_N$, and setting $r = r_2$, we show that in any equilibrium, the strongest firm earns $w_1(r)$ and every other firm i earns the value of its outside option o_i (see Proposition III).

Equilibrium behavior is hard to characterize in general. However, if firms can be unambiguously ranked in their willingness to target the contested segment, then the equilibrium is unique and the two highest-ranked firms compete in the contested segment with the same strategies as in the duopoly case, while all the other firms take their outside options.

We provide two such ranking conditions. First, if there are no captive consumers, then a ranking condition is that for any given pair of firms, the firm with the lower unit cost also has a lower fixed cost and more recoverable inventories (see Proposition VI). Second, if firms have identical unit costs, then a ranking condition is that for any given pair of firms, the firm that has fewer captive consumers also has lower advertising costs and more recoverable inventories (see Proposition VII).[12]

If such ranking conditions are not met, then there can exist equilibria in which three or more firms are actively targeting the contested segment. It is still the case that for sufficiently low prices, only firms 1 and 2 are targeting the contested segment, and do so with the same CDFs as in the duopoly case (see Proposition V). However, at higher prices, other firms may be targeting the contested segment.

To illustrate this point, we study a three-firm example where firms 1 and 2 have lower but less recoverable unit costs than firm 3, so that our ranking conditions are violated (see Online Appendix I.5). As long as α_3 is not too high, the unit-cost-advantage effect dominates

[12] The all-pay contests literature has also identified ranking conditions that guarantee non-participation of some players (see, e.g., Theorem 2 in Siegel, 2009).

and only firms 1 and 2 are active in equilibrium. However, once α_3 becomes sufficiently high, there is a unique equilibrium such that at low prices, only firms 1 and 2 mix continuously, while at higher prices, firm 2 becomes inactive and only firms 1 and 3 mix continuously (see Proposition VIII).

If firms are symmetric, it is very simple to characterize the unique symmetric equilibrium: All 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 2.1, it seems however questionable to focus on the symmetric equilibrium.[13]

3 Equilibrium Properties and Bertrand Convergence

The objective of this section is twofold. First, we list and discuss the qualitative features of the equilibrium. Second, we show that classical 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. Away from the limit, the analysis contained in this paper thus provides a toolbox to generalize the Bertrand-type equilibrium to markets with production in advance.

The analysis in Section 2 reveals that the following properties hold generically. (a) The equilibrium is unique. (b) There is price dispersion and each firm has at least one mass point (i.e., a reference price). The strong firm has a single reference price, 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. (c) Firms source enough inventory to supply their targeted demand (and strictly prefer to do so).[14] (d) All segments are served with probability 1: The strong firm always targets both its captive and the contested segment, and the weak firm always targets at least its captive consumers.

Combining the second and third properties, it follows that (e) 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 firms from adjusting inventory choices accordingly. As mentioned in the introduction, this mechanism provides an explanation as to why some products with low

[13]Moreover, the symmetric equilibrium has disputable comparative statics. For example, the expected price at which consumers purchase in the contested segment increases with the number of firms. This issue is discussed in detail in Online Appendix I.6.

[14]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.

variance in consumer demand feature high variance in inventory orders (the bullwhip effect).

The stochasticity of firm-level demand also drives the intuition behind the next property. (f) 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.[15] (For the same reason, the weak firm may also price above its monopoly price.) The intuition for prices above the monopoly level is the following. Consider a hypothetical monopolist facing a demand of $D(p)$ with probability λ, and 0 otherwise. It maximizes

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

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, prices can be higher in oligopoly than with a multi-segment monopolist—i.e., production in advance can lead to price-increasing competition.[16]

Finally, we establish two additional important properties: (g) The equilibrium is Bertrand convergent and (h) both firms earn their Bertrand profits regardless of the recoverability parameters.

Under Bertrand competition, it is assumed 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 2.1 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, thus establishing property (g):[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. □

[15] 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 global maximum at $\bar{p}_2$.

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

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

Our approach thus nests the classical Bertrand one as a limiting case. Away from that limit, the equilibrium described in Section 2 generalizes the Bertrand-type equilibrium to environments with production in advance when the salvage value of inventories falls significantly short of their acquisition value.

Property (h), that both firms earn their Bertrand profits regardless of the recoverability parameters, is 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 recoverability parameters tend to 1, equilibrium play tends to an equilibrium of the limiting Bertrand game. Hence, the Bertrand profit is also $w_i(r)$.

As mentioned in the introduction, Proposition 4 is in contrast to what is known from the literature that studies the Kreps and Scheinkman (1983) framework (in which $\mu_i = A_i = 0$). Assuming *observable* inventory choices, that literature finds that if the sunk part of the unit cost is sufficiently high, then the equilibrium outcome is (close to) Cournot. When inventory costs are sufficiently recoverable, other outcomes can arise, which may be more competitive than Cournot, but always remain far from the intense competition arising in Bertrand.[18]

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

4 Applications

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

4.1 Textbook Bertrand

In the *textbook Bertrand model* ($\alpha_1 = \alpha_2 = 1$), firms have identical unit costs ($c_i = c$), there are neither captive consumers ($\mu_i = 0$) nor advertising costs ($A_i = 0$), and demand is continuous at p^0. As is well known, firms price at marginal cost and earn zero profit in the unique equilibrium (Harrington, 1989).

Consider now a symmetric production-in-advance version of this model ($\alpha_i = \alpha < 1$).

[18]Deneckere and Kovenock (1996) and Davidson and Deneckere (1986) find that with linear demand, the equilibrium may involve mixing, but closed-form solutions are unavailable when costs are sufficiently recoverable. Yet, if the sunk part of the unit cost of both firms exceeds just 7.5% of the choke price, then the outcome is Cournot in the former paper, and Cournot or close to it in the latter one.

There is a unique equilibrium. Since k is increasing, both firms mix over $[c, p^0)$ according to

$$k(p) = \frac{p - c}{p - \alpha c}.$$

Each firm puts the rest of its mass $1 - k(p^0) > 0$ on its outside option, i.e., inactivity.

The equilibrium characterized above was previously uncovered by Levitan and Shubik (1978) and Gertner (1986) in variants of this model where inventory costs are fully sunk. 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 in fact, we need to use completely different techniques.[19]

This non-genericity also has implications in terms of equilibrium behavior. Note for example that the equilibrium CDF does not depend on the shape of $D(\cdot)$. Moreover, there is a positive probability that no firm serves the market—thus, property (d) is violated. Finally, both firms are indifferent between all the pure strategies that are not strictly dominated—thus, the equilibrium is non-strict in a very strong sense, and property (c) is violated.[20]

4.2 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). When $A_1 < A_2$, a condition that holds generically, firms mix continuously over $[r, p^m)$, the strong firm (firm 1) puts the rest of its mass on p^m, and firm 2 puts the rest of its mass on inactivity (i.e., on staying out of the market).

In the production-in-advance ($\alpha_1, \alpha_2 < 1$) framework explored in the present paper, 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 in turn determines how firms distribute the rest of their mass. This feature allows us to illustrate clearly the equilibrium characterization of Section 2.1.

If $\alpha_1 \geq \alpha_2$, firm 1 is stronger in both dimensions, in the sense of having lower and more

[19]Our approach essentially follows the one in Gertner (1986), 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.

[20]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.

recoverable costs. This results in firm 1 being less willing to set high prices than firm 2 ($\bar{p}_1 \geq \bar{p}_2$). The equilibrium, represented graphically in the top panel of Figure 1, is thus described by equation (1), with the strong firm putting the rest of its mass on $\bar{p}$ and the weak firm putting the rest of its mass on inactivity. The firms' mass points are therefore qualitatively similar to those under production to order, with the exception that $\bar{p} > p^m$, in line with equilibrium property (f).[21]

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

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

Proposition 4 shows that as costs become fully recoverable, the difference between our equilibrium and the production-to-order one becomes quantitatively small. Thus, for example, prices cease to exceed the monopoly level in the limit.

4.3 Bertrand with Heterogeneous Unit Costs

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

In our production-in-advance ($\alpha_1, \alpha_2 < 1$) framework, $r_1 = c_1$ and $r_2 = c_2$, so firm 1 is the strong firm. In equilibrium, firms still earn their Bertrand profits, but now 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 $\bar{\pi}_2 = 0$, the function $k_1(p) = (p - c_2)/(p - \alpha_2 c_2)$ is strictly increasing, and so $\bar{p}_1 = p^0$. The equilibrium is thus always described by equation (1), with the strong firm putting the rest of its mass on $\bar{p}$ and the weak firm putting the rest of its mass on inactivity.

Note that, in contrast to the Bertrand outcome: It is not necessarily the strong firm that ends up serving the market; the strong firm's inventory may remain unsold; there is price dispersion; and the strong firm prices above its monopoly price with positive probability. This equilibrium converges to the Bertrand equilibrium in which firm 2 is the least aggressive in

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

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

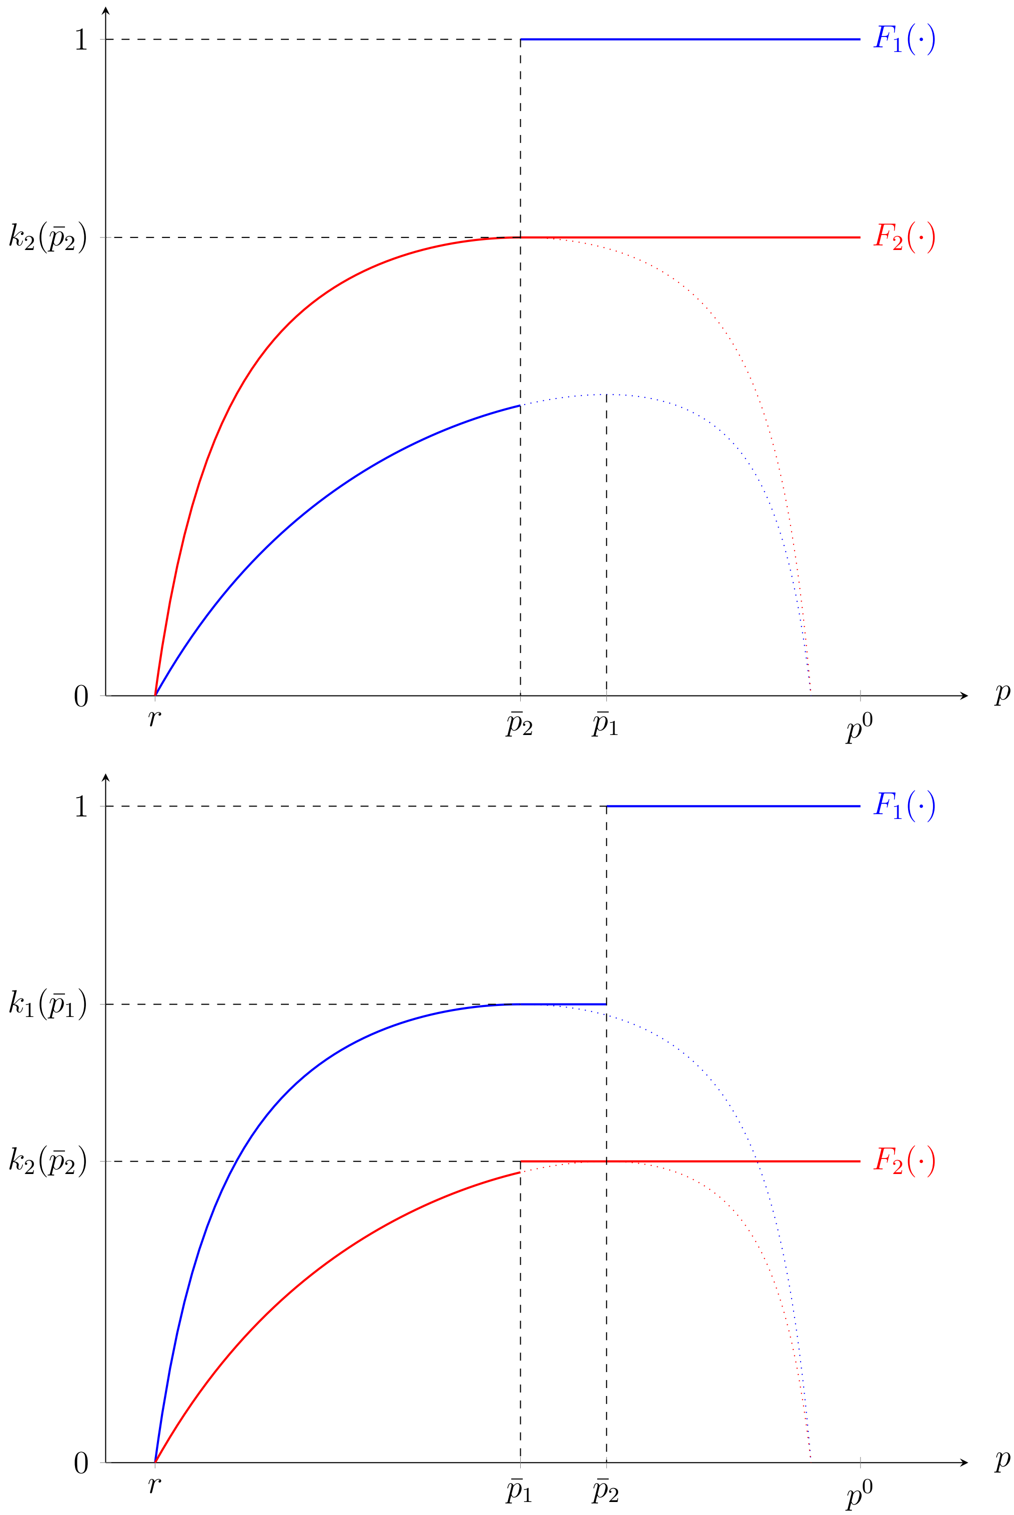

In both panels, $D(p) = 1 - p$, $c_1 = c_2 = 0.3$, $A_1 = 0.01$, $A_2 = 0.03$, and $\mu_1 = \mu_2 = 0$; in the top panel, $\alpha_1 = 0.9$ and $\alpha_2 = 0.1$; in the bottom panel, $\alpha_1 = 0.1$ and $\alpha_2 = 0.9$.

its randomization as the α's tend to 1 (see Appendix C.5).

Suppose next that in addition to heterogeneous unit costs, firms have *identical* fixed costs ($A_1 = A_2 > 0$). Bertrand versions of this model have been studied by Lang and Rosenthal (1991) and Anderson, Baik, and Larson (2015), with the latter reinterpreting it as a model of personalized pricing and advertising. In our production-in-advance framework, if $\alpha_1 = \alpha_2$, then $r_1 < r_2$ and $\bar{p}_1 > \bar{p}_2$. This holds since firm 1 is advantaged in all dimensions (strictly so in one and weakly so in the other two). The production-in-advance equilibrium is therefore given by equation (1), like in the case without fixed costs.

If instead firm 1 is advantaged in some dimensions but disadvantaged in others, then the weak firm may have two mass points as the equilibrium may instead be described by equation (2). In the Bertrand limit, such an equilibrium emerges if and only if the strong firm has a strictly higher unit cost than the weak firm.[22] To obtain this type of equilibrium under production to order, it is necessary to have, in addition to interior monopoly prices, heterogeneous fixed and unit costs. Such equilibria have eluded the existing Bertrand literature, which has confined attention to either perfectly inelastic demand or heterogeneity in a single dimension.

4.4 Clearinghouse Models

Consider next a textbook Bertrand model enriched with captive consumers. Suppose demand is perfectly inelastic. The symmetric version of this model has been 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—as Narasimhan (1988) first pointed out, firms are always holding a sale.

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

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

In a generic equilibrium, one of the firms puts the rest of its mass on targeting both its captive and the contested segment at p^0, whereas the other firm puts the rest of its mass on targeting only its captive segment at p^0. Hence, both firms have a reference price, and so firms are not always holding a sale.

As inventory costs become fully recoverable, this equilibrium converges to an equilibrium of the resulting Bertrand game.[23] 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 seemingly innocuous difference explains the discrepancy: In our setting, firms choose not only their prices but also which segments to target, whereas in Varian's model, firms always target both their captive and the contested segment by assumption.

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

[23] 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.

It is natural to expect that there are costs, fixed or variable, associated with targeting any given market segment—our model allows for both types of costs, whereas Varian's accounts for neither. In the presence of either costs, firms need to actively decide not only what price to set but also which segments to target.

The production-to-order literature has explored this avenue by adding advertising costs to Varian's model (e.g., Baye and Morgan, 2001; Iyer, Soberman, and Villas-Boas, 2005). Due to those costs, the equilibria of those Bertrand games feature mass points on monopoly prices. Yet, in the equilibria that were characterized, neither 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), 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).

A second source of discrepancy is the usual focus on the symmetric equilibrium, which as we argued above is non-generic. Indeed, Iyer, Soberman, and Villas-Boas (2005) study the symmetric equilibrium of a Bertrand game, assuming as we do that a consumer remains unaware of the product unless it is targeted. In this equilibrium, there is a positive probability that neither firm targets the contested segment. However, that model also has an asymmetric equilibrium that corresponds to the limit of our generic equilibrium as inventory costs become fully recoverable.

Consumer surplus and social welfare can vary significantly from one equilibrium to the other. Our genericity argument suggest that the asymmetric equilibrium is the robust one, and should therefore be used to address policy questions. Above all, our results also provide an avenue to extend such policy analysis to situations with production in advance.

5 An All-Pay Oligopoly with Incomplete Information

We study an *ex-ante* symmetric all-pay oligopoly in which firms have private information about their unit costs. We provide two convergence results: A purification result in the spirit of Harsanyi (1973) in our framework with bi-dimensional continuous actions, and a Bertrand convergence result.

5.1 Framework and Equilibrium Analysis

Consider an all-pay oligopoly with neither captive consumers nor advertising costs, with D strictly decreasing, log-concave, and $\mathcal{C}^2$ on $[0, p^0)$. Suppose that unit costs are drawn i.i.d. from a probability distribution G, which has strictly positive and continuous density over its support $[\underline{c}, \overline{c}]$, with $0 < \underline{c} < \overline{c} < p^0$. If the realizations of c_1 and c_2 were public information, then the analysis in Section 2 would apply. In this section, we study instead the case where those realizations remain private information. For simplicity, we assume that $\alpha_1 = \alpha_2$.

We look for pure-strategy Bayesian equilibria. A pure strategy for firm i is a mapping $(p_i(\cdot), q_i(\cdot)) : [\underline{c}, \overline{c}] \longrightarrow [0, p^0] \times \mathbb{R}_+$. As in Section 2, we first focus on the constrained game, in which firms source enough inventories to supply their targeted demand, i.e., $q_i(\cdot) = D(p_i(\cdot))$. We later show that the sets of constrained and unconstrained equilibria coincide.

We start by restricting our attention to symmetric Bayesian equilibria with monotone and differentiable strategies. More precisely, we focus on symmetric equilibria with a cutoff type $c^0 \in (\underline{c}, \overline{c})$ such that the common $p^*(\cdot)$ is continuous on $[\underline{c}, \overline{c}]$, differentiable with strictly positive derivative on $[\underline{c}, c^0)$, and equal to p^0 on $[c^0, \overline{c}]$. This p^* induces a CDF of prices $F^*(\cdot) \equiv G((p^*)^{-1}(\cdot))$ that is differentiable on $[\underline{p}, p^0) \equiv [p(\underline{c}), p^0)$.

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

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

The expected profit of type $c < c^0$ when setting price p is

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

The first-order condition for that type c can be written as

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

This condition and the fact that pricing at p must be optimal for type $c = G^{-1}(F^*(p))$ gives an ordinary differential equation $F^{*\prime}(p) = \Psi(p, F^*(p))$, where

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

F^* must be increasing, satisfy the boundary conditions $F^*(\underline{p}) = 0$ and $\lim_{p \uparrow p^0} F^*(p) = G(c^0)$, and be such that every type below c^0 makes strictly positive profits.

In Online Appendix III.5.1–III.5.2, we show that there exists a unique F^* that solves this boundary-value problem. Conversely, it is easy to show that this F^* induces a pricing function $p^*(\cdot) = (F^*)^{-1} \circ G(\cdot)$ which is a Bayesian equilibrium of the constrained game. In fact, for every c, $p^*(c)$ is the unique maximizer of $\pi(\cdot, c)$, i.e., the Bayesian equilibrium is strict. The fact that each type $c \in [\underline{c}, c^0)$ makes strictly positive profits implies that this equilibrium is also a strict Bayesian equilibrium of the unconstrained game.

We now show that, under some additional conditions, there are no other equilibria. The proof is non-trivial. It relies on tools that were developed to study first-price auctions (e.g., Riley and Samuelson, 1981; Plum, 1992; Maskin and Riley, 2000, 2003; Lebrun, 1999, 2006) with two major additional difficulties, both related to the fact that a firm's winning function is non-monotonic in the price it sets.

The first difficulty is that the differential equation $F' = \Psi(p, F)$ is singular at the boundary point $(p^0, G(c^0))$. We circumvent this by exploiting the properties of super- and sub-solutions, using geometric arguments similar in spirit to those in Lebrun (2006). The second difficulty arises when proving that any equilibrium pricing function is continuous. As a firm's winning function is hump-shaped in its price, standard arguments from the auctions literature do not apply. A restriction on the behavior of the curvature of D, which boils down to assuming that monopoly pass-through is non-increasing in cost (Fabinger and Weyl, 2012), together with a restriction on G allow us to establish continuity. These complications can alternatively be sidestepped if one is willing to confine attention to Bayesian equilibria that satisfy certain properties, such as continuity or symmetry. We have:

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

(a) The sets of constrained and unconstrained equilibria coincide.

(b) There exists a unique equilibrium with continuous (resp. symmetric) pricing strategies.

(c) There is a unique equilibrium if D is $\mathcal{C}^3$, $DD''/(D')^2$ is non-decreasing, and G is $\mathcal{C}^2$ and convex.

Proof. See Online Appendix III.1 for a sketch and Online Appendix III.2–III.6 for formal developments. □

Under the conditions of Theorem 2, the equilibrium under incomplete information therefore also satisfies the generic properties (a), (b), (c) and (e) of Section 3. Property (f) is satisfied as well: Types close to c^0 are pricing close to p^0, which strictly exceeds their monopoly prices.[24] Finally, property (d) is not satisfied as both firms are inactive with strictly positive probability. The Bertrand-related properties (g) and (h) will be discussed next.

5.2 Purification and Bertrand Convergence

Purification. We now show that, as the distribution of costs converges to a unit mass on c, the pure-strategy equilibrium converges to the mixed-strategy equilibrium of the resulting complete-information all-pay oligopoly, previously seen in Section 4.1:

Proposition 5. *Let $(G^n)_{n\geq 0}$ be a sequence of probability measures over $\mathbb{R}_+$. Assume that D and G^n $(n \geq 0)$ satisfy the assumptions of Section 5.1, and let F^n be the equilibrium CDF of prices given G^n. Suppose that $(G^n)_{n\geq 0}$ converges weakly to a unit mass on c. Then, $(F^n)_{n\geq 0}$ converges weakly to the equilibrium CDF of prices of the resulting complete-information game.*

Proof. See Online Appendix IV.1. □

[24] As in Section 3, we assume implicitly that monopoly prices are interior, which holds, e.g., if D is continuous at p^0.

A concern that purification cannot address is that mixed-strategy equilibria are not regret-free, in the sense that, once firm i has observed the realization of firm j's price, firm i no longer wants to mix in the way prescribed by its equilibrium strategy. 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.[25] The fact that retailers do often end up with unsold inventories, without changing prices, seems to provide anecdotal support for this claim.

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

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

This differential equation coincides with the one studied in Section 5.1 for $\alpha = 1$, i.e., when inventory costs are fully recoverable. Establishing Bertrand convergence, the counterpart of generic property (g) in Section 3, does however require a few additional non-trivial steps, which rely on Helly's selection theorem and the Arzelà-Ascoli theorem. We have:

Proposition 6. *Suppose that D and G satisfy the assumptions of Section 5.1, and let $(\alpha^n)_{n\geq 0}$ be a sequence over $[0,1)$ such that $\alpha^n \underset{n\to\infty}{\longrightarrow} 1$. For every n, let F^n be the equilibrium CDF of prices characterized in Theorem 2. Then, $(F^n)_{n\geq 0}$ converges weakly to the Hansen-Spulber equilibrium CDF of prices.*

Proof. See Online Appendix IV.2.1 for a sketch and Online Appendix IV.2.2–IV.2.4 for formal developments. □

Property (h), i.e., the fact that firms earn their Bertrand profits in equilibrium, can be interpreted in two ways under incomplete information. The *interim* interpretation is that each type earns its Bertrand profit: This property is clearly not satisfied, as high types make zero profit when α is small, and positive profits when α is high. In fact, in an example with linear demand and uniformly-distributed costs, each type's equilibrium profit is increasing in α (see Online Appendix III.8). The *ex-ante* interpretation is that each firm earns its Bertrand profit in expectation: That same example shows that this property does not hold either.

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

[25] The present model can also be seen as describing a procurement setting. In that case, the procurement agency is committed to its auction format and all submitted bids are final.

6 Efficiency and Taxation

The equilibrium outcome in an all-pay oligopoly of complete information features three types of distortions: First, firms price above marginal cost with probability one (the classical deadweight loss); second, some inventories remain unsold with positive probability; third, it is not necessarily the most efficient firm that ends up serving the contested segment. This raises the question of whether taxes or subsidies can alleviate those distortions, as they usually do in oligopoly models. We find a surprising answer in the context of the symmetric model seen in Section 4.1, as discussed next.

As will be explained below, for a large class of symmetric taxation schemes, the resulting equilibrium is symmetric, with each firm i drawing its price from a CDF F while sourcing $q_i = D(p_i)$.[26] Social welfare is then given by:

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

where $G = 1 - (1 - F)^2$ and $H = F^2$ are the CDFs of the minimum and maximum price, respectively. Let $\mathcal{F}$ be the set of CDFs over $\mathbb{R}_+$, and F^* be the equilibrium CDF under laissez-faire. We find that the equilibrium is second-best efficient in the following sense: If the social planner could choose *any* CDF of prices, he would still end up choosing the laissez-faire one.

Proposition 7. *The equilibrium CDF, F^*, maximizes social welfare. Moreover, if D is strictly decreasing, then F^* is the unique social welfare-maximizing CDF.*

Proof. See Appendix D. □

For example, suppose that firms face a tax or subsidy rate of t for each unsold unit. The equilibrium CDF of prices is $F(p) = (p - c)/(p - \alpha c + t)$. By Proposition 7, the optimal tax rate is $t = 0$. We could likewise envision combinations of symmetric per-unit and/or ad-valorem taxes on sales, output and unsold units. All those policies give rise to a mixed-strategy equilibrium in which each firm i draws its price p_i from a common CDF, and sources $q_i = D(p_i)$. By Proposition 7, the optimal tax rates are all zero as well.

Proposition 7 is best illustrated by the case of unit inelastic demand. In this case, social welfare only depends on the probability that a firm produces one unit. Call this probability β. The expected social welfare is then given by

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

[26] The problem becomes trivial if taxation can be asymmetric: It is optimal to tax one firm out of the market, and subsidize the remaining monopoly firm to eliminate the deadweight loss. Focusing instead on symmetric taxation schemes seems to be of higher practical relevance.

The first-order condition is $(1-\beta)p^0 - c + \beta\alpha c = 0$, and we obtain $\beta = (p^0 - c)/(p^0 - \alpha c)$, which is also the probability that a firm is active (i.e., prices below p^0) in equilibrium.

The proof with an arbitrary demand function is more involved and relies on an integration by parts argument for monotone but potentially discontinuous functions (see Border, 1996). In the appendix, we first show that $W(F)$ can be rewritten as

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

where $\Phi(p, F) \equiv p(1-(1-F)^2) - 2cF + \alpha c F^2$. We then show that $F^*(p)$ is the unique maximizer of $\Phi(p,\cdot)$, which allows us to maximize W term by term to obtain Proposition 7.

Determining the optimal tax policy in the context of the general model of Section 2 would be a different exercise for the following reason. In proving Proposition 7, we were able to ignore the specific mapping from taxation schemes to CDFs by letting the social planner choose any symmetric F. The restriction to symmetric CDFs was natural as it is satisfied for a wide class of taxation schemes. There is no natural counterpart to this requirement in the general model of Section 2, where CDFs are generically asymmetric, and the relationship between F_1 and F_2 depends on the details of the taxation scheme under consideration. A thorough analysis of these issues is left for future research.

7 Concluding Remarks

We introduced and studied a class of games where stores source costly unobservable inventories in advance and then simultaneously set prices. Our framework allows for asymmetries between firms, heterogeneous consumer tastes, endogenous consumer information through costly advertising, and salvage values for unsold units.

We first studied a constrained version of the model in which stores must source enough inventories to supply all their targeted demand. That constrained game is an all-pay contest with outside options, non-monotonic winning and losing functions, and conditional investments. The equilibrium is generically unique and in mixed strategies—thus, there is price dispersion and some inventories may remain unsold. Turning our attention to the unconstrained game, where firms freely choose inventories, we showed that the sets of constrained and unconstrained equilibria coincide generically. Thus, the equilibrium of the unconstrained game is generically unique and can be found by solving the constrained game. These results are robust to the introduction of incomplete information.

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

To introduce product differentiation, we assumed that some consumers only wish to purchase the product of one store, and therefore only a fraction of consumers (the contested segment of shoppers) wishes to purchase from the store setting the lowest price. This specification of product differentiation ensures that the constrained game has the structure of an all-pay contest. That property would disappear if we were to use other widely-used specifications of product differentiation that result in a smoother residual demand. The analysis of such models would therefore require completely different techniques.[27]

Throughout the paper, we confined attention to a static, one-shot setting. The equilibrium behavior in the one-shot game does however coincide with the per-period equilibrium behavior in a properly-specified dynamic game, as explained next. Time is discrete and runs for an infinite number of periods, goods do not perish, and stores discount future profits with a discount factor $\delta < 1$. At the beginning of each period, stores can freely buy and sell inventories on a perfectly competitive wholesale market at a given price of c.[28] Stores do not observe each others' inventory holdings and set prices simultaneously. The difference relative to the static model is that each unsold unit is now carried over to the next period.

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

This result stands in contrast to Kreps and Scheinkman (1983)'s well-known result that under production in advance, if stores observe inventories before setting their prices, then the Cournot outcome should be expected. This means that the information stores have about rivals' inventories at the pricing stage affects the nature of competition significantly. It is then natural to ask how inventory observability affects consumer surplus and social welfare. The analysis of the static and dynamic effects of inventory observability on market outcomes is left for future research.

[27]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.

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

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

Appendix

A Equilibrium Analysis in the Constrained Game

In this section, we state and prove a series of technical lemmas that jointly imply Proposition 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 8 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

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

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. $\square$

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

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

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

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

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

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

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

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

Combining Lemmas B and D and the analysis in the main text, we obtain Proposition 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. □

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 $q_i = \mu_i D(p_i^m)$.[30] Denote this strategy by $(p_i^m, \mu_i D(p_i^m))$. Next, we remove all the pure strategies in which firm i pays the advertising cost and chooses (p_i, q_i) such that $q_i \leq \mu_i D(p_i)$, because those strategies are either strictly dominated by $(p_i^m, \mu_i D(p_i^m))$, or outcome-equivalent to $(p_i^m, \mu_i D(p_i^m))$. Finally, we remove all the pure strategies in which firm i is pricing below cost or choosing (p_i, q_i) such that $q_i > (1-\mu_j)D(p_i)$, as those strategies are strictly dominated.

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

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

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, 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

[30] 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.

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

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 $(p_i, q_i) \in \mathcal{A}'_i$):

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

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

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

where

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

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

Note that, since the sequence of events $((p^n, p^0])_{n\ge 1}$ is non-increasing, we have that $\lim_{n\to\infty} \varphi_j\left((p^n, p^0]\right) = \varphi_j\left(\bigcap_{n\ge1}(p^n, p^0]\right) = \varphi_j\left([\hat{p}_i, p^0]\right)$.

Next, we turn our attention to the term in equation (5). The sequence of σ_j-integrable functions $(Z^n)_{n\ge1}$ is non-negative and bounded above by the constant function $D(c_i)$, which is also σ_j-integrable. Moreover, $(Z^n)_{n\ge1}$ 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 $(\hat{p}_i, \hat{q}_i) \in \mathcal{A}'_i$, $\bar{\pi}_i \geq \pi_i^-(\hat{p}_i, \hat{q}_i, \sigma_j)$. Moreover, if $(\hat{p}_i, \hat{q}_i) \in \mathcal{A}'_i$ and $\bar{\pi}_i = \pi_i(\hat{p}_i, \hat{q}_i, \sigma_j)$, then $\pi_i(\hat{p}_i, \hat{q}_i, \sigma_j) = \pi_i^-(\hat{p}_i, \hat{q}_i, \sigma_j)$.*

B.2 Proof of Proposition 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, $\bar{\pi}_i \geq o_i$.

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

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

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

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

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

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 $(p^n)_{n \geq 1}$ be a strictly increasing sequence of prices that converges to $\hat{p}$. Note that

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

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

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

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

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

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

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

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,

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

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 $\bar\pi_i=\pi_i\left(p^n,q^n,\sigma_j\right)$. Moreover, by Lemma F, for every n,

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

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

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

Note that the integral term in the above expression is non-decreasing in q. If $(p^n-\alpha_i c_i)$ $(\varphi_j\left([p^n,p^0]\right)+\tau_j)-(1-\alpha_i)c_i$ 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

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

contradiction:

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

Therefore, $(p^n - \alpha_i c_i)\left(\varphi_j\left([p^n, p^0]\right) + \tau_j\right) - (1-\alpha_i)c_i \leq 0$ 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 Section 4.4).

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 non-generic parameter vector of Section 4.1, the equilibrium converges weakly to the unique equilibrium of the limiting game. The "Bertrand-without-fudge" model, discussed in Section 4.3, is studied 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 2.1 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 8). 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

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

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

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

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

which is also a continuous function.

As in Section 2.1, 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.[32] Therefore, $R = \max\{R_1, R_2\}$ is continuous as well.

As in Section 2.1, 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 $p \in (R(\gamma), p^0)$,

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

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

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

[32] 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.

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 $\bar{\Gamma}_i$. Let $(\gamma^n)_{n\geq 1}$ be a sequence over $\bar{\Gamma}_i$ that converges to some $\gamma\in\bar{\Gamma}_i$. Let $R(\gamma)<\hat{p}<\bar{P}_i(\gamma)$. We show that $\bar{P}_i(\gamma^n)>\hat{p}$ for n sufficiently high. Let $\check{p}\in(\hat{p},\bar{P}_i(\gamma))$. Then, $K_i(\check{p};\gamma)>K_i(\hat{p};\gamma)$. Since $\lim_{n\to\infty}K_i(\check{p};\gamma^n)=K_i(\check{p};\gamma)$ 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 $V' = \big(\bar{P}_2(\gamma) - \varepsilon, p^0\big) \times V$. 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 Model of Section 4.1

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 non-generic parameter vector of Section 4.1. 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 Section 4.1.*

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 4.1 that $F_i^*(p) = K_i(p; \gamma)$ for every $p < p^0$. For every n, let $(F_1^n, F_2^n)_{n\geq 0}$ be a constrained equilibrium of the all-pay oligopoly 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 at the beginning of Section 4.3 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.

D Proof of Proposition 7

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

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

and

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

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

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

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

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

$$\begin{aligned}
C &= 2c \int_{[0,p^0)} D(p) dF(p), \\
&= 2c \int_0^{p^0} D^-(p) dF^-(p), \\
&= 2c \left(\int_0^{p^0} (D(0) - \Psi(p)) dF^-(p) \right), \\
&= 2c \left(F^-(p^0) D(0) - \int_0^{p^0} \left(\int_0^p d\Psi(t) \right) dF^-(p) \right), \\
&= 2c \left(F^-(p^0) D(0) - \int_0^{p^0} \left(\int_t^{p^0} dF^-(p) \right) d\Psi(t) \right), \text{ by Fubini's theorem,} \\
&= 2c \left(F^-(p^0) D(0) - \int_0^{p^0} \left(F^-(p^0) - F^-(t) \right) d\Psi(t) \right),^{33} \\
&= 2c \left(F^-(p^0)(D(0) - \Psi(p^0)) + \int_0^{p^0} F^-(t) d\Psi(t) \right),
\end{aligned}$$

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

Expected recoverable costs can be simplified in a similar way:

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

Next, we rewrite gross consumer utility at price p:

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

Therefore, expected consumer gross utility is given by:

$$U = \int_{[0,p^0)} U(p)dG(p),$$

[33] Note that

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

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

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

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

Putting things together, we obtain expected social welfare:

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

where

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

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

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

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

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

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

References

Amann, E., and W. Leininger (1996): "Asymmetric All-Pay Auctions with Incomplete Information: The Two-Player Case," *Games and Economic Behavior*, 14(1), 1–18.

Anderson, S., A. Baik, and N. Larson (2015): "Personalized pricing and advertising: An asymmetric equilibrium analysis," *Games and Economic Behavior*, 92(C), 53–73.

Arnold, M., C. Li, C. Saliba, and L. Zhang (2011): "Asymmetric market shares, advertising and pricing: Equilibrium with an information gatekeeper," *Journal of Industrial Economics*, 59(1), 63–84.

Baye, M., D. Kovenock, and C. de Vries (1993): "Rigging the Lobbying Process: An Application of the All-Pay Auction," *American Economic Review*, 83(1), 289–94.

Baye, M. R., and J. Morgan (2001): "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 91(3), 454–474.

Belleflamme, P., and M. Peitz (2010): *Industrial Organization: Markets and Strategies.* Cambridge University Press.

Bernstein, F., and A. Federgruen (2004): "A General Equilibrium Model for Industries with Price and Service Competition," *Operations Research*, 52(6), 868–886.

——— (2007): "Coordination Mechanisms for Supply Chains Under Price and Service Competition," *Manufacturing & Service Operations Management*, 9(3), 242–262.

Blanchard, O. (1983): "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy*, 91(3), 365–400.

Blinder, A. S. (1986): "Can the Production Smoothing Model of Inventory Behavior be Saved?," *The Quarterly Journal of Economics*, 101(3), 431–453.

Blume, A. (2003): "Bertrand without fudge," *Economics Letters*, 78(2), 167–168.

Border, K. C. (1996): "Integration and Differentiation," Lecture notes.

Butters, G. R. (1977): "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies*, 44(3), 465–491.

Che, Y.-K., and I. L. Gale (1998): "Caps on Political Lobbying," *The American Economic Review*, 88(3), 643–651.

——— (2006): "Caps on Political Lobbying: Reply," *American Economic Review*, 96(4), 1355–1360.

Chen, F., Z. Drezner, J. K. Ryan, and D. Simchi-Levi (2000): "Quantifying the Bullwhip Effect in a Simple Supply Chain: The Impact of Forecasting, Lead Times, and Information," *Management Science*, 46(3), 436–443.

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.

Davidson, C., and R. Deneckere (1986): "Long-Run Competition in Capacity, Short-Run Competition in Price, and the Cournot Model," *RAND Journal of Economics*, 17(3), 404–415.

De Nijs, R. (2012): "Further results on the Bertrand game with different marginal costs," *Economics Letters*, 116(3), 502–503.

Deneckere, R., and J. Peck (1995): "Competition Over Price and Service Rate When Demand is Stochastic: A Strategic Analysis," *RAND Journal of Economics*, 26(1), 148–162.

Deneckere, R. J., and D. Kovenock (1996): "Bertrand-Edgeworth duopoly with unit cost asymmetry," *Economic Theory*, 8(1), 1–25.

Fabinger, M., and E. G. Weyl (2012): "Pass-Through and Demand Forms," mimeo.

Gertner, R. H. (1986): "Simultaneous move price-quantity games and equilibrium without market clearing," in "Essays in Theoretical Industrial Organization," Ph.D. dissertation, Massachusetts Institute of Technology.

Grossman, G. M., and C. Shapiro (1984): "Informative Advertising with Differentiated Products," *Review of Economic Studies*, 51(1), 63–81.

Hammond, J. H. (1994): "Barilla SpA (A)," Harvard Business School Case 694-046.

Hansen, R. G. (1988): "Auctions with Endogenous Quantity," *RAND Journal of Economics*, 19(1), 44–58.

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.

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

Lang, K., and R. W. Rosenthal (1991): "The Contractors' Game," *The RAND Journal of Economics*, 22(3), 329–338.

Lebrun, B. (1999): "First Price Auctions in the Asymmetric N Bidder Case," *International Economic Review*, 40(1), 125–42.

——— (2006): "Uniqueness of the equilibrium in first-price auctions," *Games and Economic Behavior*, 55(1), 131–151.

Lee, H. L., V. Padmanabhan, and S. Whang (1997a): "The Bullwhip Effect in Supply Chains," *Sloan Management Review*, 38(3), 93–102.

——— (1997b): "Information Distortion in a Supply Chain: The Bullwhip Effect," *Management Science*, 43(4), 546–558.

Levitan, R., and M. Shubik (1978): "Duopoly with price and quantity as strategic variables," *International Journal of Game Theory*, 7(1), 1–11.

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.

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

Maskin, E., and J. Riley (2000): "Equilibrium in Sealed High Bid Auctions," *Review of Economic Studies*, 67(3), 439–454.

——— (2003): "Uniqueness of equilibrium in sealed high-bid auctions," *Games and Economic Behavior*, 45(2), 395–409.

Moldovanu, B., and A. Sela (2001): "The Optimal Allocation of Prizes in Contests," *American Economic Review*, 91(3), 542–558.

——— (2006): "Contest architecture," *Journal of Economic Theory*, 126(1), 70–96.

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

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

Plum, M. (1992): "Characterization and Computation of Nash-Equilibria for Auctions with Incomplete Information," *International Journal of Game Theory*, 20(4), 393–418.

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

Riley, J., and W. F. Samuelson (1981): "Optimal Auctions," *American Economic Review*, 71(3), 381–92.

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

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

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.

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.

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.