

WHEN DOES PREDATION DOMINATE COLLUSION?

THOMAS WISEMAN

ABSTRACT. I study repeated competition among oligopolists. The only novelty is that firms may go bankrupt and permanently exit: the probability that a firm survives a price war depends on its financial strength, which varies stochastically over time. Under some conditions including no entry, an anti-folk theorem holds: when firms are patient, so that strength levels change relatively fast, every Nash equilibrium involves an immediate price war that lasts until at most one firm remains. Surprisingly, the possibility of entry may facilitate collusion, as may impatience. The model can explain some observed patterns of collusion and predation.

1. INTRODUCTION

In this paper, I study dynamic competition among multiple firms. The model differs from the usual repeated game setting only in that firms may go bankrupt and irrevocably exit the market. In each period, each active firm has a state variable corresponding to its financial strength. That state, which is publicly observed, moves stochastically over time, with Markov transitions that depend on per-period profits. A firm goes bankrupt with positive probability after a string of low profits. Thus, any active firm can start a price war that ends only when it or all its rivals are driven into bankruptcy. A firm in a strong financial position is more likely to survive a price war, all else equal, than a weak firm. The main result (Theorem 1) establishes a condition under which when firms are patient, then in any Nash equilibrium a price war begins very quickly (before much discounted time has passed) and lasts until a single active firm is left as a monopolist. That is, collusion is impossible for patient firms.

Date: September 15, 2016.

I thank John Asker, Allan Collard-Wexler, Joseph Harrington, Johannes Hörner, Colin Krainin, David McAdams, Marcin Peški, William Sudderth, Caroline Thomas, Nicolas Vieille, and seminar and conference participants for helpful comments, and the referees and editor for useful suggestions.

Department of Economics, University of Texas at Austin, wiseman@austin.utexas.edu.

To fix ideas, consider a market with symmetric Bertrand competition between two firms. An active firm always has the option of starting a price war by setting a price of zero until either it or its rival goes bankrupt. The firm can thus guarantee itself an expected long-run payoff close to the monopoly payoff times the firm's probability of winning a price war. If a firm is stronger financially than its rival, then it has a better than even chance of winning a price war. When the firms are very patient, then, in equilibrium the stronger firm must get over half of the available profits. A firm with an initial advantage, therefore, must immediately start a price war. Otherwise, the stochastic movement of financial strength levels means that its rival will quickly gain the advantage. At that point, the rival can ensure itself over half of the profit, and the long-run share remaining for the first firm is less than its expected payoff would have been from an immediate price war. The no-collusion result (Theorem 1) follows.

More generally, suppose that there are N firms competing, whether in Bertrand fashion or otherwise. As patience increases, in equilibrium each firm gets an expected payoff no less than the expected long-run payoff from starting a price war, equal to the monopoly payoff π^M times the firm's probability of winning the price war. If firms collude and no bankruptcies occur in equilibrium, then a firm can wait until it is in a strong financial position relative to its rivals, start a price war when its winning probability is high, and thereby get an expected payoff equal to that high probability multiplied by π^M . But if that expected payoff exceeds a $1/N$ share of collusive profits (Condition NC in Section 3), then for each firm to get such an expected payoff is infeasible, since the sum of payoffs would exceed the total collusive profits available. Thus, bankruptcies must occur quickly, before the random process driving financial state transitions gives every firm a turn being stronger than its rivals.

Intuitively, a firm initially in a strong position prefers starting a price war right away to risking a deterioration in its position. To avert a price war, firms that are initially weak would be willing to promise a strong rival a large share of future profits. However, when those weak firms eventually become strong they are no longer willing to make the promised transfers; they prefer to take their chances with a price war. Since the strong firm can foresee that outcome, an initial price war cannot be averted.

The prediction of fighting till bankruptcy is consistent with historical examples, as detailed below. For fixed levels of patience, the model can also be used to explain patterns of collusion in markets. Dividing market demand asymmetrically may make collusion easier. The effect of cartel size on the ability to collude is nonmonotonic: collusion between n firms may be sustainable when collusion between $n - 1$ or $n + 1$ firms is not. Similarly, more patient firms may be more or less able to collude than less patient firms – increasing patience has the usual repeated-game effect of making a one-time gain less attractive, but it also reduces the weight placed on the temporary losses during a price war relative to the potential long-run stream of monopoly profits for the winner. Subgame perfect equilibria (SPEs) exist where firms collude when they have equal financial strengths and fight only when one is stronger. A final, counterintuitive result is that the possibility of entry (at a cost) facilitates collusion: a potential entrant might risk a price war against a single firm, but the chances of winning a price war against multiple firms are low enough to discourage entry.

The assumption that firms may be forced to exit after strings of low profits is based on models of capital market imperfections driven by moral hazard, such as Holmstrom and Tirole (1997). Tirole's (2006) Chapter 3 summarizes those models and their predictions that firms with low net worth may face credit rationing. His Chapter 2 describes some stylized facts about corporate finance, in particular the finding that “firms with more cash on hand and less debt invest more, controlling for investment opportunities.” Many papers in the literature on the “deep pockets” or “long purse” model of predatory pricing make similar assumptions about financially weak firms facing exit as a result of borrowing constraints. See, for example, Benoit (1984), Fudenberg and Tirole (1985, 1986), Poitevin (1989), Bolton and Scharfstein (1990), Kawakami (2010) and Motta's (2004, Chapter 7) summary.

Historical experience. There are many historical instances of oligopolistic firms competing to drive each other out of the market instead of colluding. At one point during the 1824-1826 price war for New Jersey-Manhattan steamboat passenger traffic, for example, the Union Line, run by Cornelius Vanderbilt, offered a fare of zero, and in addition would provide passengers with a free dinner (Lane, 1942, p.47). In the end, the competing Exchange Line, “whose resources were small compared to those

of” the Union Line, collapsed (Lane, 1942, p.47). During the price war, “the well-financed Union Line simply took some planned temporary losses, marking them down for what they were: an investment in future prices” (Renehan, 2007, p. 108). There were similar episodes throughout Vanderbilt’s career. Weiman and Levin (1994) describe Southern Bell Telephone’s aggressive pricing policy toward competitors around 1900: SBT endured persistent “operating losses with no tendency toward accommodation” and “rejected the offer of a competitor in Lynchburg, Virginia, to ‘raise rates ... , each company agreeing not to take the subscribers of the other’ ” (p.114). Lamoreaux (1985) describes repeated failures to collude on high prices in the steel wire nail industry (pp.62-76) and the newsprint industry (pp.42-45).

Exit need not occur through bankruptcy. A failing firm might sell out to a rival at a low price. If there is some friction preventing mergers between financially strong firms, that interpretation of “forced exit” is also consistent with the model here. Genesove and Mullin (2006) argue that predatory pricing by the American Sugar Refining Company around 1900 allowed it to buy out competitors at prices that were “simply too low to be consistent with competitive conditions” (p.62). Similarly, Burns (1986) finds that the American Tobacco Company significantly reduced buyout costs through predatory pricing. SBT sometimes pursued this strategy as well: “In 1903, SBT purchased the independent [competitor] in Richmond, Virginia, for only one-third of the original asking price ... but only after five years of heavy losses” (Weiman and Levin, 1994, pp.119-120). More recently, a sequence of bankruptcies and mergers has reduced the number of major airlines in the United States from ten in 2000 to just four. Overall, Borenstein (2011) estimates that the industry lost \$59 billion (in 2009 dollars) in the thirty years after deregulation in 1978.

There also are examples where small groups of incumbent firms coordinate price cuts to drive out entrants. Scott Morton (1997) describes how British shipping cartels around 1900 jointly lowered prices in response to entry, especially entry by firms with fewer financial resources. She argues that if the entrant was in a weak financial position, then the cartel “was more likely to think it could drive the entrant into bankruptcy or exit” (p.700). Lamoreaux (1985 pp.80-82) documents how the Rail Association, a cartel of steel rail producers, facing multiple new entrants, cut prices

in 1897 with the result that “firms that had entered the industry under the pool’s pricing umbrella quickly retreated.” Lerner (1995) examines the disk drive industry from 1980-1988. He finds that in the second half of that period, when financing was difficult to obtain, “drives located adjacent to those sold by thinly capitalized undiversified rivals were priced lower than other drives” (p.585). The analysis of collusive action to deter entry, or “parallel exclusion,” in Section 4.2.1 argues for an increased focus on such behavior in antitrust enforcement.

In contrast, there are of course many more industries characterized by a stable market structure with multiple firms. The model here can explain that outcome either through some impatience on the part of the firms or through the possibility of entry, as discussed above. More broadly, the no-collusion result here applies when the prospect of forced exit in the relatively near future is relevant, as might be the case in a new industry where entry requires financing a large capital investment (as in many of the examples above) or in an industry facing declining demand.

Relation to literature. Theorem 1 is in contrast with folk theorems for repeated games such as Fudenberg and Maskin (1986). Here, a firm cannot be rewarded or punished in the future by a rival firm that is about to exit, and so the logic of the folk theorem does not apply. Folk theorems for stochastic games (Dutta, 1995, Fudenberg and Yamamoto 2011, Hörner *et al.* 2011) require that the set of equilibrium payoffs be independent of the initial state in the limit as players become patient. That assumption fails here because bankruptcy is an absorbing state for each firm. More precisely, when Condition **NC** holds, the intersection of the feasible and individually rational payoff sets across all different combinations of strength levels for active firms is empty. Similarly, the results of Green and Porter (1984) and Rotemberg and Saloner (1986) on collusion when prices are imperfectly observed and when demand fluctuates, respectively, do not hold in the model of this paper.

Besides the literature on deep-pockets predation, three other papers are closely related. Kawakami and Yoshida (1997) study repeated duopoly with exit. In their model, whichever firm deviates first from collusion by undercutting its rival is guaranteed to win the resulting price war. As a consequence, if firms are patient relative to the deterministic length of a price war, then the temptation to preempt makes

collusion impossible. Fershtman and Pakes (2000) numerically analyze outcomes in a model where entry and exit are voluntary and firms' quality levels change stochastically as a function of investment. Powell (2004) examines a bargaining model of war. In each period countries can either divide surplus peacefully or go to war, after which the winner gets all future surpluses. Extending the work in Powell (1999) and Fearon (2004), Powell (2004) focuses on understanding when war will occur.

Finally, there is also a small literature (including Milnor and Shapley, 1957, Rosenthal and Rubinstein, 1984, and Maitra and Sudderth, 1996, Section 7.16) on “games of survival,” where players, who have an initial stock of capital, repeat a stage game, and the resulting positive or negative payoffs are added to their stocks. A player whose stock becomes negative is “ruined” and exits the game. The goal is to force opponents into ruin. Shubik and Thompson (1959) and Shubik (1959, Chapters 10 and 11) generalize such games to allow players either to consume their stage-game payoffs or to add them to their capital stock.

The structure of the paper is as follows: Sections 2 and 3 present the model and the no-collusion result. Section 4 illustrates the result for the cases of Cournot and Bertrand competition. Section 5 describes behavior for fixed levels of patience. Section 6 examines the robustness of Theorem 1, and Section 7 describes how the theorem can be extended to other dynamic games. Section 8 concludes.

2. MODEL

There are $N > 1$ expected-profit maximizing firms interacting in an infinite-horizon stochastic game. There is a bounded set $\hat{S} = S \cup \{0\}$ of states of the world for each firm, where an element s of $S \subseteq (0, K]$, $K > 0$, represents the strength of the firm's financial position, and a firm in state 0 is bankrupt. States in S are referred to as active states. Let $\mathbf{s} \in \hat{S}^N$ denote a vector of states for all firms, let $I(\mathbf{s}) \equiv \{i : s_i \neq 0\}$ denote the set of active (that is, not bankrupt) firms at state vector \mathbf{s} , and let $\mathbf{S}(n)$ denote the set of state vectors where n firms are active. In each period, all active firms compete in the following symmetric stage game: each firm i chooses an action from a compact set A . Each bankrupt firm does not participate and must play the “inactive action” $a^\emptyset \in A$, which yields zero profit. Given an action $a_i \in A$ for firm

i and a vector of actions $a_{-i} \in A^{N-1}$ for the rival firms, firm i 's profit in a period is given by the measurable function $\pi(a_i, a_{-i})$. Payoffs to mixed actions are defined in the usual way.¹

Let π^M denote monopoly profit. That is,

$$\pi^M \equiv \sup_{a_i \in A} \pi(a_i, a^\emptyset, \dots, a^\emptyset).$$

More generally, for each number of active firms $n \in \{2, \dots, N\}$, let $\pi^C(n)$ denote the maximum total profits that the firms could earn by cooperating:

$$\pi^C(n) \equiv \sup_{(a_1, \dots, a_n) \in A^n} \sum_{i=1}^n \pi(a_i, a_{-i}, a^\emptyset, \dots, a^\emptyset).$$

Firms discount future payoffs at the rate $\delta < 1$ per period; the total payoff to a firm that earns stream of profits $\{\pi_t\}_{t \geq 1}$ is $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \pi_t$. Outcomes of previous periods are publicly observed, as is the realized state.

Here are a few examples of the stage game to illustrate the definitions above:

Example 1. *Bertrand competition.* Each firm i chooses a price $p_i \in [0, \bar{p}]$. Firm i 's quantity demanded is given by the following function:

$$x_i(p_i, \underline{p}, n^*) = \begin{cases} Q(p_i) & \text{if } p_i < \underline{p} \\ \frac{Q(p_i)}{n^*+1} & \text{if } p_i = \underline{p} \\ 0 & \text{if } p_i > \underline{p}, \end{cases}$$

where \underline{p} is the minimum price set by firm i 's rivals, n^* is the number of rival firms that set that price, and the function $Q(\cdot)$ is market demand. Each firm then produces, at a constant marginal cost normalized to zero, to meet its quantity demanded. Market demand $Q(\cdot)$ is differentiable and strictly decreasing below some strictly positive choke price $\hat{p} < \bar{p}$ such that $Q(p) = 0$ for all $p \geq \hat{p}$. A bankrupt firm must set price \bar{p} . Firm i 's profit, then, is given by $\pi(p_i, \underline{p}, n^*) \equiv p_i x_i(p_i, \underline{p}, n^*)$. The monopoly profit equals the maximum collusive profit and is given by

$$\pi^M = \pi^C(n) = \max_{p \geq 0} pQ(p),$$

with the associated monopoly price p^M , which is assumed to be unique.

¹Profits are deterministic given actions. All results generalize if random noise is added to payoffs, given appropriate adjustments of the assumptions on transitions in Section 2.1.

Example 2. *Linear Cournot competition.* Each firm i chooses a quantity $q_i \in [0, \bar{q}]$ at cost cq_i , where $c > 0$. A bankrupt firm must set quantity 0. When the total market quantity is Q , the market price is given by $\max\{a - bQ, 0\}$, where $a > c$, $b > 0$, and $\bar{q} \geq a/b$. Firm i 's profit is then $(a - bQ - c)q_i$. The monopoly profit and the maximum collusive profit coincide and are given by $\pi^M = \pi^C(n) = (a - c)^2/4b$, achieved by a total quantity of $(a - c)/2b$.

Example 3. *Cournot competition with quadratic costs.* The environment is the same as in Example 2, except that now the cost to a firm of producing quantity q is cq^2 . The monopoly profit is $\pi^M = a^2/4(b + c)$, achieved by quantity $a/2(b + c)$. The maximum collusive profit is

$$\pi^C(n) = n \frac{a^2}{4(nb + c)},$$

achieved when each firm produces quantity $a/[2(nb + c)]$.

In Example 3, for any number of firms $n > 1$ the maximum collusive profit $\pi^C(n)$ exceeds monopoly profit π^M because of the diseconomies of scale: multiple firms can produce more cheaply than a single firm.

Example 4. *Hotelling competition.* The number of firms is $N = 2$, and each firm i chooses a price $p_i \in [0, 1]$. A bankrupt firm must set price 1. Firm 1 is located at $x = 0$ and firm 2 is located at $x = 1$. There is a unit mass of consumers distributed uniformly along the interval $[0, 1]$. Each consumer has unit demand, and a consumer of type x has value $v^1(x) = 1 - x$ for firm 1's product and value $v^2(x) = 1 - (1 - x) = x$ for firm 2's product. A consumer of type x will purchase from firm i if $v^i(x) - p_i > \max\{0, v^j(x) - p_j\}$. Thus, the quantity demanded for firm i is

$$q(p_i, p_j) = \min \left\{ 1 - p_i, \frac{1 - (p_i - p_j)}{2} \right\}.$$

The cost of production is 0, so firm i 's profit is $p_i q(p_i, p_j)$. The monopoly profit is $\pi^M = \frac{1}{4}$, achieved when the firm sets price $\frac{1}{2}$ and sells to the half of the market closest to its own position. The maximum collusive profit is $\pi^C(2) = \frac{1}{2}$, achieved when each firm sets price $\frac{1}{2}$ and sells to the closer half of the market.

In Example 4, product differentiation means that the two firms jointly can make more profit than either on its own. In fact, the maximum collusive profit is twice the monopoly profit: each firm acts as a monopolist in its half of the market.

2.1. States and state transitions. Bankruptcy is an absorbing state. Recall that a bankrupt firm earns a continuation payoff of 0: $\pi(a^\emptyset, a_{-i}) = 0$ for any a_{-i} . If and when every firm but at most one becomes bankrupt, the surviving firm, if any, becomes a monopolist and earns the corresponding flow of profits $(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \pi^M = \pi^M$.

Otherwise, states do not affect payoffs directly. Instead, a firm in a weaker financial position is closer to bankruptcy. In particular, after each period an active firm's state for the next period is determined stochastically according to a transition rule Γ that is a measurable function of the firm's current state and profit: $\Gamma(\pi, s)[s']$ denotes the probability that tomorrow's state is less than or equal to s' given that today's state is s and today's profit is π .² Tomorrow's state is stochastically increasing in today's state. Formally, for any profit π and any two states s^H, s^L such that $s^H > s^L$, the distribution $\Gamma(\pi, s^H)$ strictly first-order stochastically dominates the distribution $\Gamma(\pi, s^L)$. That is, $\Gamma(\pi, s^H)[s] \leq \Gamma(\pi, s^L)[s]$ for all s , and the inequality is strict for some s . State transitions are independent across firms.

To derive the main results, I will assume that bankruptcy is a threat but is avoidable if firms collude, and that firms' financial strengths vary significantly over time. The three formal conditions on the transition function Γ are as follows.

First, say that *bankruptcy is avoidable* if firms sharing the maximum feasible profit will not go bankrupt: there exists $\epsilon > 0$ such that for any active state $s \in S$, $\Gamma(\pi, s)[0] = 0$ whenever

$$\pi \geq \min \left\{ \pi^M, \min_{n \in \{2, \dots, N\}} \pi^C(n)/n \right\} - \epsilon.$$

In particular, if bankruptcy is avoidable, then a monopolist will not go bankrupt.

Second, say that *bankruptcy is achievable* if stage-game competition is such that a firm can, without unavoidably going bankrupt itself, impose low enough profits on its rivals that a stream of such profits will eventually lead to bankruptcy. For each

²Section 6 relaxes the assumptions that transition and payoff functions are symmetric across firms and that payoffs are independent of financial strength.

number of active firms $n \leq N$, let $\underline{\pi}(a, n)$ be the upper bound on a rival's profit when a firm plays action a :

$$\underline{\pi}(a, n) \equiv \sup_{(a_2, \dots, a_n) \in A^{n-1}} \pi(a_2, \dots, a_n, a, a^\emptyset, \dots, a^\emptyset).$$

Then *bankruptcy is achievable* if there exist $\epsilon > 0$ and a positive integer τ such that the following two conditions hold: first, for each number of active firms $n \leq N$, there exists a (possibly mixed) action $\underline{a}(n)$ such that for any state $s \in S$ and any sequence of profits $\pi^\tau = (\pi_1, \dots, \pi_{\tau-1})$ satisfying $\pi_t \leq \underline{\pi}(\underline{a}(n), n)$ for each $t \in \{1, \dots, \tau - 1\}$,

$$\Pr(s_{t_0+\tau} = 0 | s_{t_0} = s, \pi^\tau) \geq \epsilon;$$

and second,

$$\Pr(s_{t_0+1} = 0 | s_{t_0} = s, \pi) \leq 1 - \epsilon$$

for any profit π .³

Finally, say that *transitions are uniformly irreducible* if whenever a firm's profit is high enough to avoid bankruptcy, then its state moves stochastically throughout the state space S . Specifically, *transitions are uniformly irreducible* if for any $\epsilon_0 > 0$, there exist $\epsilon_1 > 0$ and a positive integer τ such that the following holds: for any two active states $s, s' \in S$ and any sequence of profits $\pi^\tau = (\pi_1, \dots, \pi_{\tau-1})$, if

$$\Pr(s_{t_0+\tau} = 0 | s_{t_0} = s, \pi^\tau) < \epsilon_1, \text{ then}$$

$$\Pr(|s_{t_0+\tau} - s'| \leq \epsilon_0 | s_{t_0} = s, \pi^\tau) \geq \epsilon_1.$$

Uniform irreducibility implies that there is a bound on the expected waiting time to transition from any active state to the neighborhood of any other active state, conditional on not going bankrupt, whenever a firm earns profits that make bankruptcy unlikely, and that bound is independent of the exact level of profits.

The assumptions that bankruptcy is avoidable, that bankruptcy is achievable, and that transitions are uniformly irreducible are maintained throughout:

³One well-known oligopoly where bankruptcy was *not* achievable is the Joint Executive Committee, the 1880s railroad cartel. Firms faced a “no-exit” constraint,” as Porter (1983, p.303) puts it: “bankrupt railroads were relieved by the courts of most of their fixed costs and instructed to cut prices to increase business” (Porter, 1983, p.303, citing Ulen, 1978, pp.70-74).

Assumption. *Bankruptcy is both avoidable and achievable, and transitions are uniformly irreducible.*

Those conditions rule out cases where in the long run all firms inevitably go bankrupt, where the threat of bankruptcy is irrelevant, and where a firm's financial strength is stable and thus plays no strategic role.

It would be natural to assume that financial strength is worsening on average if profits are low and increasing if profits are high, but that assumption is not necessary for the no-collusion result (Theorem 1). In the case of Bertrand competition, adding such a monotonicity condition yields the even stronger result that equilibrium profits and prices must be low in every period with two or more active firms (Theorem 2).

In this setting, a strategy is a mapping from the current state vector and the history of past state vectors and action profiles to an action for an active firm. Recall that a bankrupt firm must play a^\emptyset every period.

Let $E^\delta(\mathbf{s})$ be the set of payoffs obtained in Nash equilibria, given initial state vector \mathbf{s} and discount factor δ . Fink (1964), Mertens and Parthasarathy (1987, 1991), and Solan (1998) give sufficient conditions for $E^\delta(\mathbf{s})$ to be nonempty. In particular, a Nash equilibrium exists for any initial state vector if the set of stage-game actions A and the state space \hat{S} are finite, or if the profit function π and transition function Γ are continuous in actions and profits, respectively. (Recall that A is compact.) The stage games in Examples 2-3 satisfy those conditions. Examples 1 and 4 do not satisfy them, but nevertheless equilibria exist in those settings, as will be shown.

2.2. Price wars. One strategy available to any active firm i is to start a price war; that is, to try and drive all its rivals into bankruptcy by playing the action $\underline{a}(n)$ described in the definition of achievable bankruptcy in the previous section. Denote by σ^{PW} , for “price war,” the strategy of playing $\underline{a}(n)$ whenever there are $n - 1$ active rivals and earning the monopoly profit if all other firms exit.⁴ The expected payoff from σ^{PW} provides a lower bound on a firm's continuation payoff in equilibrium at any state vector. That lower bound will play a key role in the proof of the no-collusion

⁴If multiple actions $\underline{a}(n)$ satisfy the definition of achievable bankruptcy, select any one.

result (Theorem 1), specifically in Lemma 2 below. To analyze that bound, it will be useful to introduce the following notation.

Given any state vector \mathbf{s} and strategy profile σ , define for each active firm $i \in I(\mathbf{s})$ the possibly infinite random variable

$$T_i^{PW}(\mathbf{s}, \sigma) \equiv \min \{ \tau > 0 : s_{i,t+\tau} = 0 \mid \mathbf{s}_t = \mathbf{s}, \sigma \}.$$

$T_i^{PW}(\mathbf{s}, \sigma)$ denotes the number of periods until firm i goes bankrupt, starting from state vector \mathbf{s} , when firms use strategy σ . For each $n \in \{2, \dots, N\}$, initial state vector $\mathbf{s} \in \mathbf{S}(n)$, strategy profile σ , and active firm $i \in I(\mathbf{s})$, define the probability

$$\alpha_i(\mathbf{s}, \sigma) \equiv \Pr \left[T_i^{PW}(\mathbf{s}, \sigma) > \max_{j \in I(\mathbf{s}) \setminus \{i\}} T_j^{PW}(\mathbf{s}, \sigma) \right];$$

$\alpha_i(\mathbf{s}, \sigma)$ is the probability that firm i “wins” – that is, the probability that each other active firm goes bankrupt first. For each firm j that is bankrupt at state vector \mathbf{s} , define $\alpha_j(\mathbf{s}, \sigma) \equiv 0$.

The relevant lower bound on continuation payoffs in Lemma 2 depends on the probability of winning a price war that a firm can guarantee itself by playing the price-war strategy σ^{PW} :

$$\alpha_i(\mathbf{s}) \equiv \inf_{\sigma_{-i}} \alpha_i(\mathbf{s}, (\sigma^{PW}, \sigma_{-i})).$$

The assumption that bankruptcy is achievable implies that $\alpha_i(\mathbf{s})$ is strictly between 0 and 1 for any active firm i . The following lemma shows, straightforwardly, that $\alpha_i(\mathbf{s})$ increases when firm i 's financial position improves or when its rivals' financial positions worsen. That result follows because given any current profit level, the distribution of a firm's strength tomorrow is strictly increasing, in the sense of first-order stochastic dominance, in its strength today. Note that Lemma 1 uses the assumption that state transitions are independent across firms, conditional on profits. All formal proofs are in the appendix.

Lemma 1. *For each $n \in \{2, \dots, N\}$ and each pair of active firms i, j , $j \neq i$, $\alpha_i(\mathbf{s})$ is strictly increasing in s_i and strictly decreasing in s_j .*

The next lemma, Lemma 2, says that when firms are patient, in equilibrium each firm gets an expected payoff at least equal to the monopoly profit π^M times the

firm's guaranteed probability of winning a price war $\alpha_i(\mathbf{s})$. The idea is that given state vector \mathbf{s} , any active firm i can, by starting a price war, attain an expected payoff of $\alpha_i(\mathbf{s})\pi^M$ discounted by $E\delta^T$, where T is the random variable denoting the length of a price war won by firm i . The expected length of the price war is bounded independently of δ , so for δ close to 1 each firm can guarantee itself a payoff arbitrarily close to $\alpha_i(\mathbf{s})\pi^M$.

Lemma 2. *For any $\epsilon > 0$, there exist $\delta(\epsilon) < 1$ such that for any $\delta > \delta(\epsilon)$, each $n \in \{2, \dots, N\}$, each initial state vector $\mathbf{s} \in \mathbf{S}(n)$, and each firm i , firm i 's expected payoff in any Nash equilibrium is at least $\alpha_i(\mathbf{s})\pi^M - \epsilon$.*

In the next section, Lemmas 1 and 2 are used to provide a condition under which firms cannot collude for long in equilibrium.

3. IMMEDIATE PRICE WAR

Suppose first that the set of active states S is finite, so that $S = \{1, \dots, K\}$. For each number of active firms n , let $\mathbf{s}^i(n) \in \mathbf{S}(n)$ denote the state vector where firm i 's financial strength is $s_i = K$, the highest state, and $s_j = 1$, the lowest non-bankruptcy state, for each rival firm $j \in I(\mathbf{s}^i(n) \setminus \{i\})$. By Lemma 1, $\alpha_i(\mathbf{s}^i(n))$ represents a firm's highest guaranteed probability of winning a price war across all state vectors with $n - 1$ rivals. If $\alpha_i(\mathbf{s}^i(n))$ multiplied by the monopoly profit π^M exceeds a $1/n$ share of the maximum collusive profit $\pi^C(n)$, then patient firms will be unable to sustain collusion, given that bankruptcy is achievable and avoidable and that transitions are uniformly irreducible.

The intuition for the no-collusion result is that in the absence of bankruptcies, the randomness of state transitions implies that a patient firm i can get a payoff close to $\alpha_i(\mathbf{s}^i(n))\pi^M$ by waiting until its most favorable state vector $\mathbf{s}^i(n)$ is reached and then starting a price war. If $n\alpha_i(\mathbf{s}^i(n))\pi^M > \pi^C(n)$, though, the sum of those payoffs exceeds the maximum total payoff available. Thus, bankruptcies must occur in equilibrium.

For a general state space, the same argument applies as long as a firm's winning probability is high enough in a neighborhood of the most favorable state vector $\mathbf{s}^i(n)$,

where $s_i = K$ and $s_j = \inf \{s \in S\}$ for each rival firm j . For $\epsilon > 0$, define

$$\alpha^\epsilon(n) \equiv \inf \left\{ \alpha_i(\mathbf{s}) : \mathbf{s} \in \mathbf{S}(n), \|\mathbf{s} - \mathbf{s}^i(n)\| \leq \epsilon \right\},$$

and let

$$\alpha^*(n) \equiv \sup_{\epsilon > 0} \{\alpha^\epsilon(n)\}.$$

Then a sufficient condition for collusion to be impossible for patient firms is the following:

Condition.

$$\alpha^*(n)\pi^M > \frac{1}{n}\pi^C(n). \quad (\mathbf{NC})$$

Note that when the set of states S is finite, $\alpha^*(n) = \alpha_i(\mathbf{s}^i(n))$.

When Condition **NC** holds, then combining the payoff restrictions from Lemma 2 with uniformly irreducible state transitions yields the following result, stated formally as Lemma 4 in the appendix: for any Nash equilibrium strategy profile and any state vector with at least two active firms, the expected discounted time that elapses before the first bankruptcy shrinks to zero as the discount factor δ approaches 1. Iterating that result immediately yields the no-collusion result, Theorem 1. As firms become patient, in any Nash equilibrium one firm quickly drives its rivals into bankruptcy. Lemma 2's payoff bounds imply that each firm i wins the price war with probability at least $\alpha_i(\mathbf{s})$, where \mathbf{s} is the initial state vector.

Theorem 1. *If Condition **NC** holds for each $n \in \{2, \dots, N\}$, then for any $\epsilon > 0$, there exists $\delta(\epsilon) < 1$ such that for any $\delta > \delta(\epsilon)$, any initial state vector $\mathbf{s} \in \hat{S}^N$, and any Nash equilibrium σ^* , the following hold:*

- *the expected discounted time until at most one firm remains active is less than ϵ ,*
- *there exist probabilities $\{x_i(\mathbf{s}|\sigma^*)\}_i$, $\sum_{i=1}^N x_i(\mathbf{s}|\sigma^*) \leq 1$, such that for each firm i , $x_i(\mathbf{s}|\sigma^*) \in [\alpha_i(\mathbf{s}), 1 - \sum_{j \neq i} \alpha_j(\mathbf{s})]$ and firm i has probability within ϵ of $x_i(\mathbf{s}|\sigma^*)$ of being the single survivor.*

Thus, in the limit each firm i 's realized discounted average profit is π^M with probability $x_i(\mathbf{s}|\sigma^)$ and 0 with probability $1 - x_i(\mathbf{s}|\sigma^*)$.*

Theorem 1 provides a straightforward sufficient condition for predation to dominate collusion. When Condition **NC** holds for each number of active firms $n \leq N$, then long-run behavior is predictable: a winner-take-all price war must occur.⁵ Note that the theorem describes behavior in any *Nash* equilibrium, without the stricter requirement of subgame perfection. Section 4 below illustrates Condition **NC** in the examples from Section 2.

3.1. Necessity of Condition NC. Condition **NC**, the sufficient condition – given the maintained assumptions of uniform irreducibility and achievability and avoidability of bankruptcy – for long-run collusion to be unsustainable by patient firms is in fact close to necessary. If at any state vector a $1/n$ share of the collusive profit $\pi^C(n)$ is greater than the monopoly profit π^M times the probability that a firm could win a price war that it starts unilaterally, then equal sharing of $\pi^C(n)$ is sustainable in equilibrium by patient firms. The distance between that condition and Condition **NC** comes from two sources. The first is the difference between an anticipated and an unanticipated price war. Recall that $\alpha_i(\mathbf{s})$ is the probability that firm i wins a price war starting from state vector \mathbf{s} when its rivals try their best to stop it. An unexpected deviation to start a price war would give the firm a one-period head start and thus might yield a probability of winning slightly higher than $\alpha_i(\mathbf{s})$. Second, a firm’s optimal strategy in a price war may be more complicated than σ^{PW} . A strategy that depended on the discount factor and the current state vector rather than just the number of active rivals might generate a higher expected payoff.

The maintained assumptions also are close to necessary. If there were some state vector \mathbf{s} that was absorbing when at most one firm deviates from equal sharing of the collusive profit $\pi^C(n)$, then patient firms could collude, so something like irreducibility is necessary for Theorem 1. Similarly, if bankruptcy were not achievable, then standard repeated games constructions of collusive equilibria would apply. Lastly, if bankruptcy were not avoidable, then under any profile of strategies in the long run all firms would go bankrupt.

⁵In this case, the set of feasible and individually rational dynamic payoffs has full dimension, but as $\delta \rightarrow 1$ the limiting set of equilibrium payoffs is a singleton. See Sorin (1986) for another example of a stochastic game with that property.

4. EXAMPLES

Recall that in the Bertrand and linear Cournot settings, Examples 1 and 2, the maximum collusive profit $\pi^C(n)$ is the same as the monopoly profit π^M for any number of active firms n . In that case, Condition **NC** simplifies to $\alpha^*(n) > 1/n$: if a firm in a very favorable position has a better than even chance of winning a price war, then a price war is inevitable.

On the other hand, in Cournot competition with quadratic costs and in Hotelling competition, Examples 3 and 4, $\pi^C(n)$ is strictly higher than π^M , and so a maximum winning probability $\alpha^*(n)$ higher than $1/n$ is needed to rule out collusion. In fact, in the Hotelling example with two firms, the collusive profit $\pi^C(2)$ is twice the monopoly profit π^M , so Condition **NC** does not hold for any probability $\alpha^*(2)$. In particular, the strategy where both firms price at the profit-maximizing level $p = \frac{1}{2}$ after every history and neither goes bankrupt is an equilibrium.

The next subsection explores when Condition **NC** holds in the Cournot setting with quadratic costs. The Bertrand case is the focus of Section 4.2.

4.1. Cournot competition with quadratic costs. To illustrate Condition **NC**, I focus on Example 3, Cournot competition with quadratic costs, with two firms and a discrete state space, so that the set of active states is $S = \{1, \dots, K\}$. Assume that a firm goes bankrupt with positive probability only if its profits are zero or lower. In that case, if a firm – without loss of generality, Firm 1 – wants to drive its rivals into bankruptcy, it must produce quantity at least a/b in order to push the market price down to zero. Otherwise, a rival firm could make strictly positive profit by choosing a small enough quantity. But then Firm 1's profit is no more than $-c(a/b)^2 < 0$. If the probability of bankruptcy is strictly decreasing in profit, then Firm 1 is at a disadvantage.

Specifically, transitions are such that when Firm 1 starts a price war, then 1) after each period each firm either moves down one state or stays at the current state, 2) the probability $\gamma^D > 0$ that Firm 2, earning profit 0, moves down is constant across active states, and 3) the probability that Firm 1, earning $-c(a/b)^2$, moves down is $\theta\gamma^D$, where $\theta \in [1, 1/\gamma^D)$. The parameter θ represents Firm 1's disadvantage.

Besides the parameters γ^D and θ , it will be useful to keep track of the ratio of collusive profit $\pi^C(2)$ to monopoly profit π^M . That ratio, $R \equiv 2(b+c)/(2b+c)$, is strictly between 1 and 2: colluding firms get more than π^M but not twice as much. Recall that $K \geq 2$ is the number of non-bankruptcy levels of financial strength.

In this setting, Condition **NC** can be rewritten as

$$\alpha^*(2) > \frac{R}{2}. \quad (1)$$

When Condition **1** holds, collusion is impossible for patient firms. The question, then, is which parameter values satisfy Condition **1**. In particular, first, for fixed values of R , γ^D , and θ , how many states higher must Firm 1 be than Firm 2 to have a better than $R/2$ chance of winning a price war that it starts? That is, how large must K be for Condition **1** to hold? Second, when there are exactly two non-bankruptcy states, how large can Firm 1's disadvantage θ grow before it no longer has a better than $R/2$ chance of winning when it is strong and Firm 2 is weak?

The assumptions on the transition functions make it straightforward to calculate the maximum guaranteed probability of winning a price war, $\alpha^*(2) = \alpha_i(K, 1)$, and answer those questions. The details of the computation are in the online appendix. The results are presented in Claim **1**.

Claim 1. In the example above, Condition **NC** holds if the number of financial strength levels K exceeds $\underline{K}(R, \theta)$, where

$$\underline{K}(R, \theta) \equiv \frac{\theta}{\left(1 - \frac{R}{2}\right) \cdot e} + 1.$$

Condition **NC** holds for any $K \geq 2$ if the transition disadvantage θ is less than $\bar{\theta}(R) \equiv 4 - 2R$.

The required number of states $\underline{K}(R, \theta)$ is increasing in both its arguments. When the ratio $R = \pi^C(2)/\pi^M$ is high, a firm needs a large strength advantage to make the expected payoff from starting a price war more attractive than half the collusive profit. When θ is higher, then the firm starting a price war moves toward bankruptcy more quickly, so it needs a greater initial advantage to yield the same probability of winning. For the same reasons, the upper bound $\bar{\theta}(R)$ is decreasing.




FIGURE 4.1. Minimum number of active states to ensure that Condition **NC** holds, as a function of $R = \pi^C(2)/\pi^M$, at transition disadvantage $\theta = 1.5$.

Numerically, as R increases from its lower bound $R = 1$, the threshold $\bar{\theta}(R)$ decreases linearly from 2 until it hits its lower bound $\theta = 1$ at $R = \frac{3}{2}$. That is, a one-state advantage gives Firm 1 a better than even chance of winning the price war as long as its state deteriorates no more than twice as quickly as Firm 2's. At $R = \frac{3}{2}$, where collusive profit is 50 percent higher than monopoly profit, the threshold $\underline{K}(\frac{3}{2}, \theta)$ increases at a linear rate from 2.4 at $\theta = 1$, to 3.9 at $\theta = 2$, to 14.7 at $\theta = 10$: collusion is unsustainable when there are at least 3, 4, or 15 non-bankruptcy strength levels, respectively, when Firm 1's state deteriorates at the same rate as Firm 2's, twice as fast, or ten times as fast. Figure 4.1 shows how large K must be, as a function of R , to rule out collusion when $\theta = 1.5$.

4.2. Bertrand competition. This section focuses on Example 1, Bertrand competition with linear costs and no capacity constraints, a setting where the maximum collusive profit $\pi^C(n)$ equals the monopoly profit π^M . Again, a firm can go bankrupt only if it earns zero profit or lower. Here, the price-war strategy σ^{PW} for firm i is to set price $p_i = 0$ in every period until either firm i or all its rivals go bankrupt. In this Bertrand setting, the set of equilibrium payoffs $E^\delta(\mathbf{s})$ is nonempty for any discount factor. The strategy profile where each active firm plays the price-war strategy σ^{PW} constitutes a Nash equilibrium and is in fact subgame perfect.

When firm i plays σ^{PW} , then every firm, regardless of the pricing decisions of firms other than i , makes a profit of zero in every period until either firm i or all its rivals go bankrupt. It follows that two firms with the same financial strength have the same chance of winning a price war, regardless of who starts it. Then for any state vector \mathbf{s} the sum across firms of the winning probabilities $\alpha_i(\mathbf{s})$ is less than 1 only because of the possibility that the last two active firms go bankrupt in the same period – that

is, because of the positive probability that everyone loses the price war. When the per-period probability of transitioning to bankruptcy is small, then the likelihood of such a tie also is small. The following lemma formalizes that result.

Lemma 3. *In the Bertrand case, if $\Gamma(\pi, s)[0] < \bar{\gamma}$ for all π and all $s > 0$, then $\sum_i \alpha_i(\mathbf{s}) \geq 1 - \bar{\gamma}$.*

Lemma 3 has two implications. First, it means that Condition **NC** always holds in the Bertrand case when the per-period probability of bankruptcy is low. Lemma 1 shows that α_i is strictly monotonic in strength levels, so when the sum of the α_i 's is close to 1, it follows that a firm in a stronger state than its rivals has a strictly better than $1/n$ chance of winning a price war: $\alpha^*(n) > 1/n$. Theorem 1 then implies that a price war is inevitable when firms are very patient.

Left open is the possibility that firms may collude for a large number of periods before a price war starts. The logic of Theorem 1 uses the fact that a firm facing the most advantageous vector of states prefers an immediate price war because it knows that its position can only deteriorate. If all firms are in intermediate states, though, then they may be willing to wait and see whether their positions will improve. (See an example with a fixed δ in Section 5.4.) However, the second implication of Lemma 3 is that under a mild monotonicity condition on state transitions, that possibility can be ruled out in the Bertrand case when the per-period probability of bankruptcy is low. The monotonicity condition specifies that a higher profit for a firm today strictly improves the distribution over the next period's financial strength:

Definition 1. State transitions are *strictly monotonic* if for any profit levels $\pi, \pi' \in [0, \pi^M]$ such that $\pi' > \pi$ and any non-bankruptcy state $s \in S$, $\Gamma(\pi', s)$ strictly first-order stochastically dominates $\Gamma(\pi, s)$.

Strict monotonicity means that a firm would improve its chance of winning a price war by being the first to undercut its rivals. When the sum of the α_i 's is close to 1, Theorem 1 implies that for high δ , there is very little scope in equilibrium to punish such a deviation. The vector of financial states pins down payoffs almost completely, and so if undercutting improves a firm's relative position even a little, then it would be a profitable deviation for a patient firm. Under the monotonicity condition, Theorem

2 shows that when the per-period bankruptcy probability is small and firms are patient, then in any period when more than one firm is active, the market price (that is, the lowest price set by an active firm) is low with high probability and expected profits are close to zero.

Theorem 2. *Suppose that state transitions are strictly monotonic, and choose any $\epsilon > 0$. In the Bertrand case there exists $\bar{\gamma}(\epsilon) > 0$ such that whenever $\Gamma(\pi, s)[0] < \bar{\gamma}$ for all π and all $s > 0$, there is a $\delta(\epsilon) < 1$ with the following property: for any $\delta > \delta(\epsilon)$, any initial state vector $\mathbf{s} \in \hat{S}^N$, and any Nash equilibrium σ^* , in any period on the equilibrium path where at least two firms are active, both the total expected profit and the probability that the market price exceeds ϵ are less than ϵ .*

Theorem 2 implies that in subgame perfect equilibrium prices and profits are low at any history, on or off the equilibrium path, with at least two active firms.

The proof is similar to the proof that profits are zero in a one-shot Bertrand equilibrium – if total profit were high, then some firm could increase its profit by undercutting. The argument is slightly more complicated here because firms maximize the discounted stream of payoffs rather than just today’s profit. The additional pieces of the argument are as follows: monotonicity of state transitions implies that a firm can improve its expected financial strength tomorrow and weaken its rivals by increasing its profit today at its rivals’ expense. When the sum of the α_i ’s is close to 1, Lemma 1 and Theorem 1 show that a firm’s equilibrium payoff is increasing in its financial strength and decreasing in its rivals’. Thus, in equilibrium undercutting must not yield a one-shot boost in profit – total profit must therefore be low. In contrast to the usual folk-theorem argument for repeated games, where a one-shot deviation yields a gain of order $1 - \delta$, here a deviation can have long-lasting effects.

4.2.1. *Entry and Joint Monopolization.* In the model in Section 2, the number of active firms cannot increase, although it may decrease through bankruptcies. In this section, I consider the effect of allowing new firms to enter the market in the Bertrand case. As Genesove and Mullin (2006, p.47) point out, “The continued exercise of market power depends upon deterring entry.” In their survey of empirical studies of cartels, Levenstein and Suslow (2006, Sections 4.4.4 and 5) find that entry was

a major cause of cartel breakdowns. Harrington (2006, pp.68-69), similarly, argues that entry in the international vitamin C market “most likely caused the collapse of the cartel” in the 1990s. The standard model of competition with free entry (in the absence of increasing returns to scale) predicts that entry will drive industry profits to zero – collusion and entry are incompatible.⁶

Here, however, that effect is reversed. Theorem 1 shows that in the absence of entry, patient firms cannot collude when Condition NC holds. If potential new firms can choose to enter at a small cost, however, then collusion may become sustainable. A potential competitor that sees a single incumbent firm may be willing to pay the entry cost because it has a chance to survive the resulting price war and earn positive profits. If there are multiple incumbent firms, then the entrant’s odds of winning are much lower and may not justify the entry cost. That is, joint monopolization may be profitable when monopolization by a single firm is not, because the ability of the cartel to jointly deter entry (“parallel exclusion”) is greater than a single firm’s.⁷

Formally, say that *entry is possible* if the following holds: there is a countably infinite sequence of potential entrants. At the start of each period, with probability $\rho_e > 0$ (independent across time) the next potential entrant i gets a one-time opportunity to enter the market by paying cost $c_e > 0$. If firm i decides not to enter, then it gets payoff zero. If firm i enters, then it immediately becomes active with a financial state drawn from a commonly known distribution F_e over the set of active states S . Firm i observes its realized financial state before making its entry decision. As before, history is publicly observed.

Theorem 3 shows that when entry is possible, there is an equilibrium in which a cartel of $n^* > 1$ firms colludes on the monopoly price. The firms in the cartel deter entry by threatening to jointly wage a price war against any new entrant.

Theorem 3. *When entry is possible in the Bertrand setting, then there is an SPE in which in the long run i) there are n^* active firms, ii) firms collude (set the monopoly*

⁶See Harrington (1989) for an exception.

⁷A single active firm can deter entry if the cost of entry is high enough relative to an entrant’s chance of winning a price war, as might be the case if a new firm is financially weak.

price p^M in each period), and iii) no entry occurs, as long as n^* exceeds a finite lower bound \underline{n} and δ exceeds a lower bound $\delta(n^*) < 1$.

The proof is constructive: firms price collusively when there are no more than n^* firms and price at marginal cost otherwise. If a firm deviates, then other firms price at marginal cost as long as the deviating firm is still active. Entry is allowed when there are fewer than n^* firms, but any further entry is treated as a deviation and punished. The details of the proof are straightforward but somewhat tedious.⁸

Together, Theorems 1 and 3 show that the possibility of entry can make the market *less* competitive, at least in the short run – instead of an initial price war followed by eventually monopolization, consumers face joint monopolization from the start.⁹ That finding suggests that, as Hemphill and Wu (2013) argue, the practice of parallel exclusion merits more attention in antitrust analysis.

Theorem 3 deals with the limiting case of perfect patience and shows the existence of a collusive equilibrium where entry is deterred completely. For a high but fixed level of patience, it is possible to construct similar equilibria where collusion is again sustainable, but occasional entry does occur, especially among stronger potential competitors – parallel exclusion discourages entry better than a monopolist could, but still not perfectly. Such equilibria match qualitatively the behavior of the shipping cartels in Scott Morton (1997) toward entrants.

5. PATTERNS OF COLLUSION FOR FIXED δ

Theorem 1 shows that in the limit as firms' discount factor δ approaches 1, collusion cannot be sustained in equilibrium when Condition **NC** holds. The focus of this section is to consider patterns of predation and collusion that may occur for less than perfect levels of patience.

⁸Theorem 3 is the only place where the assumption in the Bertrand example that firms have zero fixed costs matters. If firms faced a positive fixed cost, then there would be an upper bound, possibly below \underline{n} , on the number of active firms in the long run: firms not making enough profit to cover their fixed costs would go bankrupt.

⁹Besides the collusive equilibrium described in Theorem 3, when entry is possible there are also other equilibria that may feature lower prices.

5.1. The role of patience. Increasing firms' patience has conflicting effects on the ability to collude. There is the usual repeated-game effect: a patient firm is less willing to sacrifice future collusive profits for a one-time gain from undercutting its rivals, so increasing δ makes collusion easier to sustain. On the other hand, increasing δ also makes a firm value a potential future stream of monopoly profits more highly relative to the short-run losses of a price war, so collusion becomes harder to sustain. Proposition 1 in the online appendix describes examples where in one case raising δ facilitates collusion, but in another case it destroys the possibility of collusion.

Note that the effect of adding a positive trend to demand for the firms' product (so that the monopoly profit π^M increases over time) is qualitatively similar to that of an increase in δ : the opportunity cost of a price war today is lower relative to the value of future monopoly profits, but the one-shot gain from undercutting today is similarly smaller relative to the value of future collusive profits.

5.2. Cartel size. Changes in the number of firms have a similarly ambiguous effect. The more rivals a firm has, the smaller its share of the collusive profit and the greater the one-time benefit from undercutting the monopoly price. However, a firm's chances of winning a price war also decrease with the number of other firms. In the standard model of repeated Bertrand competition, only the first effect is present, and so increasing the size of a cartel always reduces the incentive to collude.

Proposition 2 in the online appendix describes one example where collusion on the monopoly price is sustainable with two or three firms but not with four, and another example where three firms cannot collude but two or four can. Those patterns suggest that when the initial number of firms is high, there may be a price war that lasts not until only one firm is left, but until the number of active firms is consistent with collusion. In that case, long-run cartel size may depend on initial conditions.

5.3. Unequal market shares. For a fixed level of patience, collusion with unequal sharing of profits may be sustainable when symmetric sharing is not. In particular, giving a larger share to firms in strong financial states can help sustain collusion. The difficulty with symmetric sharing is that a strong firm's advantage in the probability of winning a price war is not matched by an advantage in profits. Transferring profit

from weak firms to a strong firm can reduce the strong firm’s incentive to start a price war in two ways. First, it obviously increases the firm’s instantaneous profit. Second, if higher profits today lead to higher expected financial states tomorrow, then the transfers increase the average number of periods that the strong firm will maintain its advantage and corresponding high profits.¹⁰

Proposition 3 in the online appendix describes a two-firm example in which symmetric collusion, where each firm gets half the total collusive profit each, period is not sustainable. There is, though, an SPE such that when the firms’ financial strengths are uneven, the stronger firm gets 90 percent of the profit.¹¹

To share profit unequally, firms could, for example, divide the market into regions of different sizes and assign each region to a specific firm. Griffin (2000, p.11) presents evidence that “members of most cartels recognize that price-fixing schemes are more effective if the cartel also allocates sales volume among the firms,” and Harrington (2006, p.25) in his study of European cartels in the late twentieth century finds that “[a]ll but a few cartels had an explicit market sharing arrangement.” Harrington’s (2006, pp.31-33) descriptions of (i) bargaining over sales quotas between large and small firms in the lysine and calcium chloride cartels and (ii) variations over time in the assigned cartel market shares for vitamins B2 and B5 are consistent with the phenomenon in Proposition 3. Alternatively, firms might simply transfer money to each other. Levenstein and Suslow (2006, section 4.4.2) give examples.

5.4. Recurring collusion. The logic behind Theorem 1 is that when a firm is at its strongest and all its rivals are at their weakest, then the firm’s position can only worsen, so it has an incentive to start a price war right away. When firms are at even strength, though, each may be willing to wait and see whether its position will improve

¹⁰The nonmonotonic effects of patience and cartel size on the sustainability of collusion in Propositions 1 and 2 also apply to asymmetric collusion, although the examples become more complicated.

¹¹Transfers play a different role here than in the collusive equilibrium of Harrington and Skrzypacz (2007). There, transfers help with the incentive constraint: a firm that exceeds its quota must make a transfer to its rivals in the next period. (Harrington and Skrzypacz, 2011, list examples of such schemes.) Here, transfers help with the individual rationality constraint by encouraging the strong firm to participate in the cartel. In Athey and Bagwell (2001) and Athey *et al.* (2004), where firms’ costs vary over time, dividing quantity allows firms to allocate production efficiently.

while in the meantime earning collusive profits. In the online appendix, Proposition 4 presents a Bertrand example where there is no equilibrium in which the two firms always collude on the monopoly price p^M . There is, however, an SPE in which they price at p^M when they have equal strength, and they set price 0 otherwise.¹²

6. ROBUSTNESS

The availability of a public randomization device does not affect any of the results. If firms could commit to future prices, then a price war need not result, because weak firms could credibly promise to give a currently strong firm a high share of profits forever. However, Lemma 2 would still hold – initial financial strengths would still influence long-run payoffs, which are restricted by the threat of a price war.

The appendix proves a stronger version of Theorem 1 that allows both stage-game payoff functions and state transition functions to vary both by firm and by state. For example, a firm with a lower marginal cost than its rivals or a lower probability of transitioning to bankruptcy would have an advantage in a price war for any vector of financial strengths. Each firm’s probability of winning a price war would still be increasing in its own strength and decreasing in its rivals strengths, though, and so the argument behind Theorem 1 goes through with no substantial changes. Even a firm with a permanent advantage would want to “lock in” that advantage at its highest point by starting a price war.

Private states. In the model of Section 2, firms’ financial strengths are publicly observed, but I speculate that the no-collusion result extends to the case of asymmetric information. Suppose that each firm observes its own state but not the states of other firms. Instead, a firm knows only which of its rivals are still active. In every period (except possibly the first, if firms start at the strongest state), the uniform irreducibility condition implies that each firm assigns probability bounded away from zero to its’ rivals being in states below the strongest state K . When a firm is in or very close to state K , then, it expects to have an advantage over its rivals in a price war. Thus, the argument behind Theorem 1 likely still holds.

¹²Fershtman and Pakes (2000) present a somewhat similar example, in which collusion cannot be sustained between one firm in a high-quality state and two other firms in very low-quality states.

More precisely, uniform irreducibility implies the following: there exists $\underline{\epsilon} > 0$ such that for any vector of states \mathbf{s} , after each period, regardless of the action profile a played, either 1) the probability that one or more firms go bankrupt is at least $\underline{\epsilon}$ or 2) for each active firm i , the probability that firm i moves to a new state $s'_i \notin [s_i - \underline{\epsilon}, s_i + \underline{\epsilon}]$ is at least $\underline{\epsilon}$. If neither condition held, then the outcome from repeatedly playing a starting from state vector \mathbf{s} would contradict the definition of uniform irreducibility. Then the probability that a firm in state K moves to a state below $K - \underline{\epsilon}$ is at least $\underline{\epsilon}$ whenever an action unlikely to lead to bankruptcy is played. Since tomorrow's state is stochastically increasing in today's state, the probability that a firm is in a state below $K - \underline{\epsilon}$ after such an action is at least $\underline{\epsilon}$ *regardless* of its previous state. Thus, no matter what information firm i has about firm j 's state in period t – even if firm j 's action choice in period t perfectly reveals its period- t state – firm i must assign probability no higher than $1 - \underline{\epsilon}$ to firm j 's state in period $t + 1$ being above $K - \underline{\epsilon}$, as long as the action profile in period t was unlikely to cause bankruptcy.

For any sequence of actions, with high probability either a bankruptcy occurs or each firm approaches the strongest state K , and so the logic of Theorem 1 still applies. Condition **NC** extends to a sufficient condition for the no-collusion result with private states: just replace $\alpha^*(n)$ with a firm's probability of winning if it starts a price war when it is close to state K and each of its $n - 1$ rivals is below state $K - \underline{\epsilon}$ with probability at least $\underline{\epsilon}$. Note that in the Bertrand case the set of equilibria is nonempty, because the price-war strategy profile is still a Nash equilibrium with private states. In fact, it is a sequential equilibrium when combined with appropriate beliefs. Existence of equilibrium in the general case is not guaranteed, though.

Unbounded state space. In the model so far, there is an upper bound K on the space of financial strengths. However, the argument behind the no-collusion result does not rely on the state space being bounded. Rather, the proof requires the existence of a state vector at which a firm 1) has a high probability of winning a price war, so that Condition **NC** holds, and 2) expects that its probability of winning will very likely decrease in the future. Those conditions are satisfied if conditional on the action profile a firm's state follows a stationary autoregressive process, for example. The key is that the process be mean-reverting, although the mean and the rate of

reversion may depend on actions. I speculate that that property plus irreducibility is enough for the no-collusion result.

The same holds true if states are private information. When a firm reaches a state high enough above the action-conditional mean, then the independence across firms and the mean-reverting nature of the state transitions make it very likely that its rivals are weaker. Thus, the firm expects to have an advantage in a price war. Mean reversion will erode that advantage, so the firm prefers an immediate price war.

7. OTHER DYNAMIC GAMES

The logic behind the no-collusion result extends to a broad class of dynamic games. Suppose that there are a finite set of players I and a metric space S representing the set of states. State transitions may depend on actions. In each state s , each player in a subset $I(s) \subseteq I$ can choose to end the game, in which case players receive a vector of termination payoffs, $w(s)$, that depends on the state. In each period that the game continues, players receive stage-game payoffs as a function of their actions and the state. For a game that ends in period T , then, players' overall payoffs are given by

$$(1 - \delta) \left(\sum_{t=1}^{T-1} \delta^{t-1} u(a_t, s_t) \right) + \delta^{T-1} w(s_T), \quad (2)$$

where s_t and a_t are the realized state and the chosen action profile, respectively, in period t .¹³ For any strategy profile σ and initial state s , let $U^\delta(\sigma, s)$ be the expectation over actions, states, and the stopping time of the payoff vector in Expression 2.

For each state s , let $\Sigma^\infty(s)$ denote the set of strategy profiles under which no player ever stops the game, given initial state s . Define

$$\bar{u}(s) \equiv \limsup_{\delta \rightarrow 1} \left[\sup_{\sigma \in \Sigma^\infty(s)} \sum_{i \in I} U_i^\delta(\sigma, s) \right]$$

as the upper bound on the sum of payoffs for patient firms achievable by such strategies. (In the model of Section 2, that upper bound is the maximum collusive payoff.)

¹³Such games are related to stopping games, although the literature on stopping games typically assumes that payoffs are undiscounted, that a player's only action choices within a period are to stop or continue, and that there is a fixed vector of payoffs that players receive if no one ever stops.

Denote by $B(s, \epsilon)$ the closed ball of radius ϵ centered at state s . Let $S(i, \epsilon) \equiv \{s \in S : i \in I(s') \forall s' \in B(s, \epsilon)\}$, and let $I^* \subseteq I$ denote the set of players for whom $S(i, \epsilon)$ is nonempty for some $\epsilon > 0$: players in I^* can stop the game in some neighborhood of states. Next, for each player $i \in I^*$, $\epsilon > 0$, and state $s \in S(i, \epsilon)$, let

$$\bar{w}_i^\epsilon(s) \equiv \inf \{w_i(s') : s' \in B(s, \epsilon)\},$$

and define

$$\bar{w}_i \equiv \sup_{\epsilon > 0, s \in S(i, \epsilon)} \{\bar{w}_i^\epsilon(s)\}$$

as player i 's highest termination payoff across all neighborhoods where player i has the chance to end the game. (That value corresponds to the monopoly profit π^M times $\alpha^*(N)$.) The no-collusion result extends as follows:

Theorem 4. *In the class of dynamic games described above, if i) for any $\epsilon > 0$, there exists a finite $T(\epsilon)$ such that the expected transition time from s to $B(s', \epsilon)$ is less than $T(\epsilon)$ for any two states s and s' and any strategy $\sigma \in \Sigma^\infty(s)$, and ii) $\sum_{i \in I^*} \bar{w}_i > \sup_{s \in S} \bar{u}(s)$, then as $\delta \rightarrow 1$, in any equilibrium the expected discounted time until some player ends the game shrinks to zero.*

The proof of Theorem 4 is essentially identical to that of Theorem 1 and is omitted. The intuition is as before: if the sum of payoffs that players have a chance to guarantee themselves by ending the game exceeds the total available payoffs, then players cannot commit to continue the game even if doing so is efficient. Examples might include deciding when to call a vote or election, or an employer trying to break up a strike by making expiring offers of higher wages to individual workers in sequence.

The model in Section 2 differs from the class of stochastic games considered here in two ways. First, when firms compete the “stop” decision is not a one-time choice: firms can conduct a price war for a while and then return to colluding. Second, the value of stopping depends on the discount factor, which determines the relative weight on the low profits during a price war in total discounted payoffs. Neither difference is important as firms become patient.

8. SUMMARY AND DISCUSSION

This paper models oligopolistic competition between patient firms. After a string of low profits, firms can be driven into bankruptcy and permanent exit. The expected time until bankruptcy during a price war depends on the firm's financial strength, which varies stochastically over time. That feature changes the range of equilibrium outcomes for patient firms dramatically: under Condition **NC** collusion is unsustainable, and every Nash equilibrium involves an immediate price war that lasts until at most one firm is left. That result is robust to several natural extensions of the model, and it can be generalized to a broad class of stochastic games. The possibility of entry at a small cost may facilitate collusion. The model also characterizes patterns of firm behavior for fixed levels of patience.

The discussion of entry and joint monopolization in Section 4.2.1 suggests that antitrust policies toward parallel exclusion are important. (Hemphill and Wu, 2013, argue that “we lack a ... robust understanding of parallel exclusion” (p.1253), and that there has been a “neglect of [its] importance” in legal analysis (p.1182).) In the simple model here, exclusion occurs only through the threat of price wars, but vertical integration and control of inputs, tying contracts and rebates, and patents could have similar effects.

Finally, note that the model in Section 2 can be applied very naturally to international relations. Instead of price wars, there are actual wars where countries with stochastically varying military strength levels compete for productive resources. In that context, Condition **NC** represents a substantial weakening of Powell's (2004) sufficient condition for war, Condition (2).¹⁴ While Powell (2004, p.238) uses Condition (2) to argue that “rapid shifts in the distribution of military power lead to war,” he also says that such “shifts ... due to differential rates of economic growth are empirically too small to account for war through this mechanism.” The analysis here suggests that the possibility of large gradual shifts in the future can cause war in the present. Krainin and Wiseman (forthcoming) examine this issue in more detail.

¹⁴Krainin (forthcoming) presents a different generalization, for the case of two countries facing a deterministic sequence of vectors of minmax values.

APPENDIX A. APPENDIX: PROOFS FROM SECTION 2.2

Proof of Lemma 1. Pick an $n \in \{2, \dots, N\}$, a state vector $\mathbf{s} \in \mathbf{S}(n)$, and an active firm $i \in I(\mathbf{s})$. Choose another state vector $\mathbf{s}' \in \mathbf{S}(n)$ such that $s'_i > s_i$ and $s'_j = s_j$ for each firm $j \neq i$. Then the fact that the distribution $\Gamma(\pi, s'_i)$ strictly first-order stochastically dominates $\Gamma(\pi, s_i)$ for any profit level π means that after any finite sequence of actions, the distribution over firm i 's state conditional on starting at state s'_i strictly first-order stochastically dominates the distribution conditional on starting at state s_i . It follows that $T_i^{PW}(\mathbf{s}', \sigma)$ strictly first-order stochastically dominates $T_i^{PW}(\mathbf{s}, \sigma)$ for any strategy profile σ . Therefore, $\alpha_i(\mathbf{s}') > \alpha_i(\mathbf{s})$, and $\alpha_j(\mathbf{s}') < \alpha_j(\mathbf{s})$ for each $j \in I(\mathbf{s}) \setminus \{i\}$, since the condition that bankruptcy is achievable implies that each of those values is strictly between 0 and 1. \square

Proof of Lemma 2. First recall that a bankrupt firm j gets a continuation payoff of zero, and that $\alpha_j(\mathbf{s}) = 0$, as required. Recall also that at any state vector where firm i is a monopolist (that is, where only firm i is active), firm i receives continuation payoff π^M . To see that an active firm i must get at least a payoff of approximately $\alpha_i(\mathbf{s}) \pi^M$ in any Nash equilibrium when the firms are patient, define $\tau_i(\mathbf{s})$ as the upper bound on the expected time in a price war started by firm i until all of firm i 's rivals are bankrupt, conditional on firm i winning the price war:

$$\tau_i(\mathbf{s}) \equiv \sup_{\sigma_{-i}} E \left[\max_{j \in I(\mathbf{s}) \setminus \{i\}} T_j^{PW}(\mathbf{s}, (\sigma^{PW}, \sigma_{-i})) \mid \max_{j \in I(\mathbf{s}) \setminus \{i\}} T_j^{PW}(\mathbf{s}) < T_i^{PW}(\mathbf{s}, (\sigma^{PW}, \sigma_{-i})) \right].$$

Let

$$\tau^* \equiv \sup \{ \tau_i(\mathbf{s}') : \mathbf{s}' \in \mathbf{S}(n), i \in I(\mathbf{s}') \}.$$

The assumption that bankruptcy is achievable ensures that these expectations exist and that the supremum is finite. Then regardless of the strategy of rival firms, the expected payoff to firm i from starting a price war in state \mathbf{s} is at least $\alpha_i(\mathbf{s}) \delta^{\tau^*} \pi^M$, which converges to $\alpha_i(\mathbf{s}) \pi^M$ as δ converges to 1. (Note that δ^T is convex in T , so $E\delta^T \geq \delta^{ET}$.) For high enough δ , then, firm i can guarantee itself a payoff arbitrarily close to $\alpha_i(\mathbf{s}) \pi^M$, and so its equilibrium payoff must be at least that high. \square

APPENDIX B. PROOF FROM SECTION 3

Here, I prove a more general version of Theorem 1 that holds without the assumptions that firms are symmetric and that payoffs are independent of the state. Now firms may have different action spaces, payoff functions, state spaces, and transition functions, and payoff functions may depend on the (active) state, as long as the assumptions on those objects from Section 2 still hold for each firm. A firm i 's chance of winning a price war now depends not only on the number of active firms but on their identities, and so does the maximum collusive profit.

Let I denote a subset of the N players and $\mathbf{S}(I) \subseteq \mathbf{S}$ the set of state vectors at which exactly the firms in I are active: $\mathbf{S}(I) \equiv \{\mathbf{s} : I(\mathbf{s}) = I\}$. Define $\pi^C(I, \mathbf{s})$ as the maximum collusive stage-game profit that the firms in I can achieve together at state vector $\mathbf{s} \in \mathbf{S}(I)$, and let

$$\pi^C(I) \equiv \sup_{\mathbf{s} \in \mathbf{S}(I)} \pi^C(I, \mathbf{s}).$$

Similarly, let $\pi_i^M(s)$ be the monopoly profit for an active firm in state s , and let

$$\pi_i^M \equiv \inf_{s \in S} \pi_i^M(s).$$

Lemma 1 goes through as before, as does Lemma 2 once π^M is replaced with π_i^M .

Let $\mathbf{s}^i(I) \in \mathbf{S}(I)$ denote the state vector where $s_i = K_i$, the highest state for firm i , and $s_j = \inf \{s \in S\}$, the lower bound on non-bankruptcy states, for each active rival firm $j \in I \setminus \{i\}$. For $\epsilon > 0$, define

$$\alpha_i^\epsilon(I) \equiv \inf \left\{ \alpha_i(\mathbf{s}) : \mathbf{s} \in \mathbf{S}(I), \|\mathbf{s} - \mathbf{s}^i(I)\| \leq \epsilon \right\},$$

and let

$$\alpha_i^*(I) \equiv \sup_{\epsilon > 0} \{\alpha_i^\epsilon(I)\}.$$

The generalization of Condition **NC** to the general setting is the following:

$$\sum_{i \in I} \alpha_i^*(I) \pi_i^M > \pi^C(I). \quad (\mathbf{NC}^*)$$

Theorem 5 extends the no-collusion result, Theorem 1, to the general case.

Theorem 5. *If Condition **NC*** holds for every subset of at least two players I , then for any $\epsilon > 0$, there exists $\delta(\epsilon) < 1$ such that for any $\delta > \delta(\epsilon)$, any initial state vector $\mathbf{s} \in \hat{S}^N$, and any Nash equilibrium σ^* , the following hold:*

- *the expected discounted time until at most one firm remains active is less than ϵ ,*
- *there exist probabilities $\{x_i(\mathbf{s}|\sigma^*)\}_i$, $\sum_{i=1}^N x_i(\mathbf{s}|\sigma^*) \leq 1$, such that for each firm i , $x_i(\mathbf{s}|\sigma^*) \in [\alpha_i(\mathbf{s}), 1 - \sum_{j \neq i} \alpha_j(\mathbf{s})]$ and firm i has probability within ϵ of $x_i(\mathbf{s}|\sigma^*)$ of being the single survivor.*

The first part of Theorem 5 follows from iterating the following lemma. Lemma 2 then yields the second part. Let T^b be the random variable denoting the number of periods until the first bankruptcy; the distribution of T^b depends on the strategy profile and the initial state vector.

Lemma 4. *Choose any subset I of at least two active firms. If Condition **NC*** holds, then for any $\epsilon > 0$, there exist $\delta(\epsilon) < 1$ and a finite integer $\bar{T}(\epsilon)$ such that for any $\delta \geq \delta(\epsilon)$, any initial state vector $\mathbf{s} \in \mathbf{S}(I)$, and any Nash equilibrium σ^* , $\Pr [T^b > \bar{T}(\epsilon) | \mathbf{s}, \sigma^*] < \epsilon$. It follows that $\lim_{\delta \rightarrow 1} E [\delta^{T^b} | \mathbf{s}, \sigma^*] = 1$.*

Proof of Lemma 4. For each $\delta \in (0, 1)$, let σ^δ be a Nash equilibrium of the game with discount factor δ and initial state vector \mathbf{s} , and let $v(\sigma^\delta)$ denote the corresponding vector of equilibrium expected payoffs. By Condition **NC*** and the definition of $\alpha_i^*(I)$, there exists $\epsilon_0 > 0$ such that $\sum_{i \in I} \alpha_i^{\epsilon_0}(I) \pi_i^M > \pi^C(I)$. Denote the number of active firms by $n \equiv \#(I)$, and choose a positive $\epsilon^* < [\sum_{i \in I} \alpha_i^{\epsilon_0}(I) \pi_i^M - \pi^C(I)] / 2n$. Define

$$\underline{P}(I) \equiv \frac{\pi^C(I)}{-2n\epsilon^* + \sum_{i \in I} \alpha_i^{\epsilon_0}(I) \pi_i^M} < 1,$$

and choose an integer $m^* > \ln(\frac{1}{2}\epsilon) / \ln(\underline{P}(I))$, so that $(\underline{P}(I))^{m^*} < \frac{1}{2}\epsilon$.

For each firm $i \in I$, let T^i be the random variable denoting the (possibly infinite) number of periods until a state vector \mathbf{s}' within ϵ_0 of firm i 's most favorable state vector $\mathbf{s}^i(I)$, so that $\alpha_i(\mathbf{s}') \geq \alpha_i^{\epsilon_0}(I)$, is reached. Let $\underline{T} \equiv \max_{j \in I} T^j$, the time at which the ϵ_0 -neighborhood of each active firm's most favorable state vector has been reached. Define \underline{T}^{m^*} as the random variable denoting the (possibly infinite) number of periods until m^* such collections of arrivals occur.

Choose a positive η such that

$$\eta < \min \left\{ \frac{\epsilon}{2}, \frac{n\epsilon^*}{\sum_{i \in I} (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*)} \right\}.$$

Uniform irreducibility implies that for some finite threshold $\hat{T}(\eta)$, either

$$\Pr [T^b \leq \hat{T}(\eta) | \mathbf{s}, \sigma] \geq 1 - \eta \text{ or } \Pr [\underline{T}^{m^*} \leq \hat{T}(\eta) | \mathbf{s}, \sigma] \geq 1 - \eta$$

for any strategy σ . For any δ , if the first inequality holds, then $\Pr [T^b > \hat{T}(\eta) | \mathbf{s}, \sigma^\delta] < \epsilon$ by the definition of η . Otherwise, we must have $\Pr [\underline{T}^{m^*} \leq \hat{T}(\eta) | \mathbf{s}, \sigma^\delta] \geq 1 - \eta$. In that case, by Lemma 2, we have that for all δ above some $\hat{\delta}(\epsilon^*) < 1$ and for each firm i , firm i 's equilibrium continuation value if state vector $\mathbf{s}^i(I)$ is reached is at least $\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*$. It follows that for such δ ,

$$v_i(\sigma^\delta) \geq \Pr [T^b > T^i | \mathbf{s}, \sigma^\delta] E [\delta^{T^i} | T^b > T^i, \mathbf{s}, \sigma^\delta] (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*).$$

Since the sum of equilibrium expected profits cannot exceed the maximum collusive profit (recall that any active firm can mimic the “inactive action” a^\emptyset of a bankrupt firm), summing across active firms yields

$$\begin{aligned} \pi^C(I) &\geq \sum_{i \in I} v_i(\sigma^\delta) \\ &\geq \sum_{i \in I} \Pr [T^b > T^i | \mathbf{s}, \sigma^\delta] E [\delta^{T^i} | T^b > T^i, \mathbf{s}, \sigma^\delta] (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*) \\ &\geq \sum_{i \in I} \Pr [T^b > \underline{T} | \mathbf{s}, \sigma^\delta] E [\delta^{T^i} | T^b > T^i, \mathbf{s}, \sigma^\delta] (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*) \\ &\geq \sum_{i \in I} \Pr [T^b > \underline{T} | \mathbf{s}, \sigma^\delta] E [\delta^{T^i} | \mathbf{s}, \sigma^\delta] (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*) \\ &\geq \sum_{i \in I} \Pr [T^b > \underline{T} | \mathbf{s}, \sigma^\delta] E [\delta^{\underline{T}^{m^*}} | \mathbf{s}, \sigma^\delta] (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*). \end{aligned} \quad (3)$$

The fourth inequality follows from the fact that when $T^b < T^i$, then T^i is infinite: after a firm in I exits, the neighborhood of $\mathbf{s}^i(I)$ is never reached.

Because

$$E [\delta^{\underline{T}^{m^*}} | \mathbf{s}, \sigma^\delta] \geq \Pr [\underline{T}^{m^*} \leq \hat{T}(\eta) | \mathbf{s}, \sigma^\delta] \delta^{\hat{T}(\eta)} \geq (1 - \eta) \delta^{\hat{T}(\eta)},$$

the definition of η implies that there exists a threshold $\tilde{\delta}(\epsilon^*) < 1$ such that for all $\delta > \tilde{\delta}(\epsilon^*)$,

$$\sum_{i \in I} E [\delta^{\underline{T}^{m^*}} | \mathbf{s}, \sigma^\delta] (\alpha_i^{\epsilon_0}(I)\pi_i^M - \epsilon^*) \geq \sum_{i \in I} (\alpha_i^{\epsilon_0}(I)\pi_i^M - 2\epsilon^*). \quad (4)$$

Plugging Expression 4 into Expression 3 yields

$$\pi^C(I) \geq \sum_{i \in I} \Pr \left[T^b > \underline{T} | \mathbf{s}, \sigma^\delta \right] \left(\alpha_i^{\epsilon^0}(I) \pi_i^M - 2\epsilon^* \right),$$

and rearranging gives

$$\Pr \left[T^b > \underline{T} | \mathbf{s}, \sigma^\delta \right] \leq \frac{\pi^C(I)}{-2n\epsilon^* + \sum_{i \in I} \alpha_i^{\epsilon^0}(I) \pi_i^M} = \underline{P}(I)$$

for all $\delta \geq \delta^*(\epsilon^*) \equiv \max \{ \hat{\delta}(\epsilon^*), \tilde{\delta}(\epsilon^*) \}$.

That is, the probability that the first bankruptcy occurs after the ϵ^0 -neighborhood of each active firm i 's most favorable state vector $\mathbf{s}^i(I)$ is reached is at most $\underline{P}(I)$ when δ is high enough. Once each neighborhood has been reached, repeat the argument. Thus, the probability that no firm goes bankrupt before m^* such collections of arrivals of most favorable neighborhoods is at most $(\underline{P}(I))^{m^*} < \frac{1}{2}\epsilon$. Then, since we are considering the case $\Pr \left[\underline{T}^{m^*} \leq \hat{T}(\eta) | \mathbf{s}, \sigma^\delta \right] \geq 1 - \eta$, we have

$$\begin{aligned} \Pr \left[T^b > \hat{T}(\eta) | \mathbf{s}, \sigma^\delta \right] &= \Pr \left[T^b > \hat{T}(\eta) \text{ and } T^b > \underline{T}^{m^*} | \mathbf{s}, \sigma^\delta \right] \\ &\quad + \Pr \left[T^b > \hat{T}(\eta) \text{ and } \underline{T}^{m^*} \geq T^b | \mathbf{s}, \sigma^\delta \right] \\ &\leq \Pr \left[T^b > \underline{T}^{m^*} | \mathbf{s}, \sigma^\delta \right] + \Pr \left[\underline{T}^{m^*} > \hat{T}(\eta) | \mathbf{s}, \sigma^\delta \right] \\ &\leq (\underline{P}(I))^{m^*} + \eta \\ &\leq \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

To complete the proof, set $\bar{T}(\epsilon) = \hat{T}(\eta)$, and choose a threshold $\delta(\epsilon) \in (\delta^*(\epsilon^*), 1)$ such that $(1 - \epsilon)\delta^{\bar{T}(\epsilon)} > 1 - 2\epsilon$ for all $\delta \geq \delta(\epsilon)$. Then for all $\delta \geq \delta(\epsilon)$ and any Nash equilibrium σ^δ , $\Pr \left[T^b > \bar{T}(\epsilon) | \mathbf{s}, \sigma^\delta \right] < \epsilon$ and $E \left[\delta^{T^b} | \mathbf{s}, \sigma^\delta \right] \geq (1 - \epsilon)\delta^{\bar{T}(\epsilon)} + \epsilon \cdot 0 > 1 - 2\epsilon$. \square

APPENDIX C. PROOFS FROM SECTION 4.2

Proof of Lemma 3. In the Bertrand example, the price-war strategy σ^{PW} gives each firm a profit of zero, so

$$\alpha_i(\mathbf{s}) = \Pr \left[T_i^{PW}(s_i) > \max_{j \in I(\mathbf{s}) \setminus \{i\}} T_j^{PW}(s_j) \right],$$

where

$$T_i^{PW}(s) \equiv \min \{ \tau > 0 : s_{i,t+\tau} = 0 | s_t = s, \pi_{i,t'} = 0 \forall t' \geq t \}$$

is the number of periods until firm i goes bankrupt, starting from state s , when it earns zero profit in every period. Let $T_{(1)}$ and $T_{(2)}$ be the highest and second-highest realized values of $T_i^{PW}(s_i), i \in I(\mathbf{s})$. Then the probability that no firm wins the price war, $1 - \sum_i \alpha_i(\mathbf{s})$, is the probability that $T_{(1)} = T_{(2)}$. Because state transitions are independent across firms, the probability that the last firm goes bankrupt in the same period as the second-to-last firm (that outcome includes the case that the last $n > 2$ firms all go bankrupt at the same time) is no greater than $\bar{\gamma}$, the bound on the per-period probability of going bankrupt. Thus, $1 - \sum_i \alpha_i(\mathbf{s}) \leq \bar{\gamma}$. \square

Proof of Theorem 2. First, note that for any $\eta > 0$, strategy profile σ^* can be a Nash equilibrium for small $\bar{\gamma}$ and large δ only if no firm i has a deviation on the path of play such that i) for any vector of other firms' prices in the support of σ_{-i}^* , firm i 's profit in the period is weakly higher and every other firm j 's profit is weakly lower, and ii) firm i 's expected profit in the period increases by at least η . (Call such a deviation an “ η -deviation.”) The reason is the following. Lemma 3 shows that for any $\epsilon' > 0$, $\sum_i \alpha_i(\mathbf{s}) \geq 1 - \epsilon'/2$ whenever $\Gamma(\pi, s)[0] < \bar{\gamma}(\epsilon') \equiv \epsilon'/2$ for all π and all $s > 0$. In that case, it follows that $1 - \sum_{j \neq i} \alpha_j(\mathbf{s}) \leq \alpha_i(\mathbf{s}) + \epsilon'/2$ for each firm i , so Theorem 1 implies that there exists $\delta(\epsilon'/2) < 1$ such that firm i 's equilibrium payoff in any state vector \mathbf{s} , $v_i(\sigma^*, \mathbf{s})$, is within ϵ' of $\alpha_i(\mathbf{s}) \pi^M$ whenever $\delta \geq \delta(\epsilon'/2)$. That is, $v_i(\sigma^*, \mathbf{s})$ is close to the payoff that results from all firms earning zero profits in every period until only one survives:

$$v_i(\sigma^*, \mathbf{s}) \leq \epsilon' + (1 - \delta)0 + \delta E[\alpha_i(\mathbf{s}') | \mathbf{s}, \pi = \mathbf{0}, s'_i \neq 0] \pi^M.$$

Because state transitions are strictly monotonic (by assumption), and $\alpha_i(\mathbf{s})$ is strictly increasing in s_i and strictly decreasing in \mathbf{s}_{-i} (by Lemma 2),

$$E[\alpha_i(\mathbf{s}') | \mathbf{s}, \pi_i = \eta, \pi_{-i} = \mathbf{0}, s'_i \neq 0] > E[\alpha_i(\mathbf{s}') | \mathbf{s}, \pi = \mathbf{0}, s'_i \neq 0].$$

Thus, for small enough ϵ' and $\bar{\gamma}$, and large enough δ , an η -deviation by firm i would give firm i a total expected payoff strictly greater than $v_i(\sigma^*, \mathbf{s})$. In Nash equilibrium, then, no firm has an η -deviation on the path of play.

The next step is to show that if total expected profit (in the current period) is at least ϵ , then some firm has an η -deviation. Suppose that total expected profit is at least ϵ , and choose $\eta > 0$ such that $\epsilon > 2N^2\eta/(N - 1)$.

For each active firm i , let

$$\bar{p}_{-i} \equiv \min_{j \in I(\mathbf{s}), j \neq i} \max \left\{ p : p \in \text{supp } \sigma_j^* \right\},$$

the highest price in the support of the pricing strategy of all of firm i 's rivals; \bar{p}_{-i} is the supremum of the range of prices that firm i can set to make sales with positive probability. Observe that a price $p_i > \bar{p}_{-i}$ cannot be in the support of firm i 's equilibrium strategy, because if it were, then firm i would have an η -deviation. Setting price $p_i > \bar{p}_{-i}$ yields a profit of 0. Since total expected profit is at least ϵ , there must be some firm j (possibly $j = i$) that expects profit at least ϵ/N . Instead of setting $p_i > \bar{p}_{-i}$, firm i can mimic the part of firm j 's strategy below \bar{p}_{-i} ; that is, firm i can play according to $\hat{\sigma}_i(\sigma_j^*, \bar{p}_{-i})$, where

$$\hat{\sigma}_i(\sigma_j^*, \bar{p}_{-i})[p] \equiv \begin{cases} \frac{\sigma_j^*(p)}{\int_{p' \leq \bar{p}_{-i}} \sigma_j^*(p')} & \text{if } p \leq \bar{p}_{-i} \\ 0 & \text{otherwise.} \end{cases}$$

Strategy $\hat{\sigma}_i$ yields expected profit at least $\frac{1}{2}\epsilon/N$ (sharing firm j 's profit), and by assumption $\frac{1}{2}\epsilon/N > N\eta/(N-1) > \eta$. Further, $\hat{\sigma}_i$ always yields profit at least 0 (the profit from setting $p_i > \bar{p}_{-i}$), and yields each competing firm a weakly lower profit than would setting $p_i > \bar{p}_{-i}$. Therefore, $\hat{\sigma}_i$ would be an η -deviation from $p_i > \bar{p}_{-i}$.

Thus, when δ is high, in equilibrium no firm sets a price above the support of any other firm's pricing strategy: the upper bound on equilibrium prices is the same for all firms. Let \bar{p} denote that common upper bound. For $x > 0$, let the decreasing, right-continuous function $q_{-i}^*(x) \equiv \text{Pr}[p_j \geq \bar{p} - x \forall j \neq i | \sigma^*]$ denote the probability that each firm $j \neq i$ sets price $\bar{p} - x$ or higher. If $q_{-i}^*(0) = 0$, then for small enough x pricing above $\bar{p} - x$ cannot be part of firm i 's equilibrium strategy: setting price $\bar{p} - x$ yields a profit close to 0, and playing $\hat{\sigma}_i(\sigma_j^*, \bar{p} - x)$ (that is, mimicking the part of the most profitable firm j 's strategy below $\bar{p} - x$) would be an η -deviation.

It must be, then, that $q_{-i}^*(0) > 0$ for each firm i : firm i sets price $p_i = \bar{p}$ with positive probability. Setting $p_i = \bar{p}$ can be part of firm i 's equilibrium strategy only if $q_{-i}^*(0)$ satisfies $q_{-i}^*(0)\pi(\bar{p})/N > \frac{1}{2}\epsilon/N - \eta$; otherwise, the mimicking strategy $\hat{\sigma}_i(\sigma_j^*, \bar{p})$ would be an η -deviation from $p_i = \bar{p}$. On the other hand, setting a price just below \bar{p} is an η -deviation from $p_i = \bar{p}$ unless $q_{-i}^*(0)\pi(\bar{p})/N > q_{-i}^*(0)\pi(\bar{p}) - \eta$. That is, it must

be the case that

$$\frac{1}{2}\epsilon - N\eta < q_{-i}^*(0)\pi(\bar{p}) < \frac{N}{N-1}\eta,$$

implying that

$$\epsilon < \frac{2N}{N-1}\eta + 2N\eta = \frac{2N + 2N(N-1)}{N-1}\eta = \frac{2N^2}{N-1}\eta,$$

but by assumption $\epsilon > 2N^2\eta/(N-1)$. This is a contradiction, so total expected profit in equilibrium must be less than ϵ .

Finally, note that if the probability that the market price exceeds ϵ is at least ϵ , then by setting price ϵ a firm would make expected profit at least $\epsilon\pi(\epsilon) > 0$. Given any $\eta > 0$, that strategy would be an η -deviation unless the equilibrium expected profit is at least $\epsilon\pi(\epsilon) + \eta$. As was just shown, though, the discount factor δ can be chosen high enough that equilibrium profit is below that level. Thus, it must be that for high enough δ , the probability that the market price exceeds ϵ is less than ϵ . \square

Proof of Theorem 3. The equilibrium strategy σ^* can be represented by a countable-state automaton: one on-path state $\omega^{<n^*}$ for the case where there are fewer than n^* active firms, one on-path state $\omega^{=n^*}$ when there are n^* active firms, one on-path state $\omega^{>n^*}$ when there are more than n^* active firms, and two “punishment states” for each firm j (both those initially active and potential entrants): $\omega^{j,\geq n^*}$ for the case where there are at least n^* active firms, and $\omega^{j,<n^*}$ when there are fewer than n^* active firms. In state $\omega^{>n^*}$, σ^* specifies that any potential entrant stays out, and that active firms set a price of 0. In state $\omega^{=n^*}$, entrants stay out, and firms set the monopoly price p^M . In state $\omega^{<n^*}$, a potential entrant enters, and firms set the monopoly price p^M . In each punishment state $\omega^{j,\geq n^*}$, entrants stay out, firm j (the firm being punished) sets price p^M , and any other active firms set a price of 0. In each punishment state $\omega^{j,<n^*}$, entrants enter, firm j sets price p^M , and any other active firms set a price of 0.

In a slight abuse of terminology, define a “unilateral deviation by firm i ” in a period as the case where either a single firm i deviates from the specified actions or a potential entrant deviates from the specified entry decision and then a single firm i deviates from the specified price. (Using this definition will mean that only the *last* deviation in a period will be punished.)

Transitions between state of the automaton are as follows: from state $\omega_t = \omega^{j, \geq n^*}$ or $\omega_t = \omega^{j, < n^*}$, if there are no unilateral deviations and firm j is active at the end of the period, then the next state is again a punishment state for firm j : either $\omega^{j, \geq n^*}$ (if the number of active firms at the end of the period, n_{t+1} , is at least n^*) or $\omega^{j, < n^*}$ (if $n_{t+1} < n^*$). If firm i deviates unilaterally and is still active at the end of the period, then either $\omega_{t+1} = \omega^{i, \geq n^*}$ or $\omega_{t+1} = \omega^{i, < n^*}$, according to the value of n_{t+1} . If either i) there are no unilateral deviations and firm j is not active at the end of the period, or ii) firm i deviates unilaterally and is not active at the end of the period, then the next state is one of the on-path states, according to the value of n_{t+1} : $\omega_{t+1} = \omega^{> n^*}$ if $n_{t+1} > n^*$, $\omega_{t+1} = \omega^{= n^*}$ if $n_{t+1} = n^*$, and $\omega_{t+1} = \omega^{< n^*}$ if $n_{t+1} < n^*$.

From state $\omega_t \in \{\omega^{> n^*}, \omega^{= n^*}, \omega^{< n^*}\}$, transitions are similar: if firm i deviates unilaterally and is still active at the end of the period, then $\omega_{t+1} = \omega^{i, \geq n^*}$ or $\omega_{t+1} = \omega^{i, < n^*}$, according to the value of n_{t+1} . If either i) there are no unilateral deviations, or ii) firm i deviates unilaterally and is not active at the end of the period, then the next state is one of the on-path states, according to the value of n_{t+1} .

The initial state also is determined by the number of active firms at the start of the game: $\omega_1 = \omega^{> n^*}$ if $n_1 > n^*$, $\omega_1 = \omega^{= n^*}$ if $n_1 = n^*$, and $\omega_1 = \omega^{< n^*}$ if $n_1 < n^*$. Note that if firms follow σ^* , then the outcome will be as specified in the theorem. It remains to show that the strategy is a best response.

First, consider the incentive of an active firm to deviate from the specified price. In any state, the expected discounted average payoff v^{dev} to a firm i from deviating converges to zero as δ approaches 1: after the period of the deviation, the expected time τ until firm i goes bankrupt is finite. (Even if firm i outlasts its current competitors, entrants will continue the price war and eventually drive firm i into bankruptcy.) Thus,

$$\lim_{\delta \rightarrow 1} v^{dev} \leq \lim_{\delta \rightarrow 1} (1 - E\delta^\tau) \pi^M = 0.$$

(Since δ^τ is convex in τ , $E\delta^\tau \geq \delta^{E\tau}$.) In state $\omega^{= n^*}$ or $\omega^{< n^*}$, following the equilibrium strategy σ^* yields an active firm a payoff of π^M/n^* . Thus, for high enough δ , no active firm gains from deviating from price p^M in state $\omega^{= n^*}$ or $\omega^{< n^*}$.

Neither does an active firm i have an incentive to deviate from p^M in state $\omega^{> n^*}$, $\omega^{j, < n^*}$, or $\omega^{j, \geq n^*}$, $j \neq i$: there is no short-run gain from deviating (since other firms

are setting price 0, if firm i sets a higher price it will make no sales), and deviating reduces the continuation value (since then firm i will be punished). Similarly, in state $\omega^{i,<n^*}$ or $\omega^{i,\geq n^*}$, firm i cannot gain from deviating: its choice of price does not affect its continuation value, and setting price p^M is optimal in the short run (whether or not there are any active rivals).

Finally, consider the incentives of a potential entrant to deviate from the prescribed entry-no entry choice. Staying out yields a payoff of 0. In state $\omega^{<n^*}$, following σ^* and entering yields a payoff of $\pi^M/n^* - (1 - \delta)c_e$, which is positive for large δ . Entering in state $\omega^{j,<n^*}$ gives a payoff of close to at least $(\pi^M/n^*)\underline{\alpha} - (1 - \delta)c_e$, where

$$\underline{\alpha} \equiv \inf_{s \in S} \alpha_1(s, K) > 0$$

is the lower bound on firm i 's probability of outlasting the punished firm j (at which point the price war ends). Again, that value is positive for large δ .

In state $\omega^{>n^*}$, $\omega^{=n^*}$ or $\omega^{j,\geq n^*}$, $j \neq i$, σ^* specifies staying out and getting payoff 0. Entering can yield a positive profit only if the firm outlasts all currently active rivals in the subsequent price war, which occurs with probability no higher than $\bar{\alpha}^{n^*}$, where

$$\bar{\alpha}^{n^*} \equiv \max_{\mathbf{s} \in \mathbf{S}(n^*)} \alpha_i(\mathbf{s}).$$

Even if the firm succeeds in becoming the only active firm, eventually entrants will drive it into bankruptcy; as above, let $E\tau$ denote the (finite) expected number of periods until that bankruptcy. Then the expected payoff from entering is no more than $\bar{\alpha}^{n^*} E\tau(1 - \delta)\pi^M - (1 - \delta)c_e$; since

$$\lim_{n^* \rightarrow \infty} \bar{\alpha}^{n^*} = 0,$$

that expected payoff is negative for large enough n^* .

Thus, if n^* is large enough, and δ is large enough relative to n^* , then σ^* is an SPE. □

REFERENCES

- [1] Athey, S. and Bagwell, K. (2001) “Optimal Collusion with Private Information,” *RAND Journal of Economics*, **32**, pp. 428–465.
- [2] Athey, S., Bagwell, K. and Sanchirico, C. (2004). “Collusion and Price Rigidity,” *Review of Economic Studies*, **71**, pp. 317–349.
- [3] Benoit, J-P. (1984). “Financially Constrained Entry in a Game with Incomplete Information,” *RAND Journal of Economics*, **15**, pp. 490–499.
- [4] Bolton, G. and Scharfstein, D. (1990). “A Theory of Predation Based on Agency Problems in Financial Contracting,” *American Economic Review*, **80**, pp. 93–106.
- [5] Borenstein, S. (2011). “Why Can’t US Airlines Make Money?,” *American Economic Review: Papers & Proceedings*, **101**, pp. 233–237.
- [6] Burns, M. (1986). “Predatory Pricing and the Acquisition Cost of Competitors,” *Journal of Political Economy*, **94**, pp. 266-296.
- [7] Dutta, P.K. (1995). “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, **66**, pp. 1-32.
- [8] Fearon, J. (1995). “Rationalist Explanations for War,” *International Organization* **49**, pp. 379-414.
- [9] Fearon, J. (2004). “Why Do Some Civil Wars Last So Much Longer than Others?,” *Journal of Peace Research*, **41**. pp. 275-301.
- [10] Fershtman, C. and Pakes, A. (2000). “A Dynamic Oligopoly with Collusion and Price Wars,” *RAND Journal of Economics*, **31**, pp. 207-236.
- [11] Fink, A. (1964). “Equilibrium in a Stochastic n -person Game,” *Journal of Science of the Hiroshima University, Series A-I*, **28**, pp. 89-93.
- [12] Fudenberg, D. and Maskin, E. (1986). “The Folk Theorem for Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, **54**, pp. 533-554.
- [13] Fudenberg, D. and Tirole, J. (1985). “Predation without Reputation,” MIT Department of Economics working paper #377.
- [14] Fudenberg, D. and Tirole, J. (1986). “A ‘Signal-Jamming’ Theory of Predation,” *RAND Journal of Economics*, **17**, pp. 366-376.
- [15] Fudenberg, D. and Yamamoto, Y. (2011). “The Folk Theorem for Irreducible Stochastic Games with Imperfect Public Monitoring,” *Journal of Economic Theory*, **146**, pp. 1664-1683.
- [16] Genesove, D. and Mullin, W. (2006). “Predation and its Rate of Return: The Sugar Industry, 1887–1914,” *RAND Journal of Economics*, **37**, pp. 47–69.
- [17] Green, E. and Porter, R. (1984). “Noncooperative Collusion under Imperfect Price Information,” *Econometrica*, **52**, pp. 87–100.

- [18] Griffin, J. (2000). "An Inside Look at the Cartel at Work: Common Characteristics of International Cartels." Address before the American Bar Association, Washington, D.C. (www.usdoj.gov/atr/public/speeches/4489.htm).
- [19] Harrington, Jr., J. (1989). "Collusion and Predation under (Almost) Free Entry," *International Journal of Industrial Organization*, **7**, pp. 381-401.
- [20] Harrington, Jr., J. (2006). "How Do Cartels Operate?," *Foundations and Trends in Microeconomics*, **2**, pp. 1-105.
- [21] Harrington, Jr., J. and Skrzypacz, A. (2011). "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *American Economic Review*, **101**, pp. 1-25.
- [22] Harrington, Jr., J. and Skrzypacz, A. (2007). "Collusion under Monitoring of Sales," *RAND Journal of Economics*, **38**, pp. 314-31.
- [23] Hemphill, C. and Wu, T. (2013). "Parallel Exclusion," *Yale Law Journal*, **122**, pp. 1182-1253.
- [24] Holmstrom, B. and Tirole, J. (1997). "Financial Intermediation, Loanable Funds and the Real Sector," *Quarterly Journal of Economics*, **62**, pp. 663-691.
- [25] Hörner, J., Sugaya, T., Takahashi, S., and Vieille, N. (2011). "Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem," *Econometrica*, **79**, pp. 1277-1318.
- [26] Kawakami, T. (2010). "Collusion and Predation under the Condition of Stochastic Bankruptcy," *Japanese Economic Review*, **61**, pp. 408-426.
- [27] Kawakami, T. and Yoshida, Y. (1997). "Collusion under Financial Constraints: Collusion or Predation When the Discount Factor is near One?" *Economics Letters*, **54**, pp. 175-178.
- [28] Krainin, C. (forthcoming). "Preventive War as a Result of Long Term Shifts in Power," *Political Science Research and Methods*.
- [29] Krainin, C. and Wiseman, T. (forthcoming). "War and Stability in Dynamic International Systems," *Journal of Politics*.
- [30] Lamoreaux, N. (1985). *The Great Merger Movement in American Business, 1895-1904*. Cambridge, UK: Cambridge University Press.
- [31] Lane, W. (1942). *Commodore Vanderbilt: An Epic of the Steam Age*, New York: Alfred A. Knopf.
- [32] Lerner, J. (1995). "Pricing and Financial Resources: An Analysis of the Disk Drive Industry, 1980-88," *Review of Economics and Statistics*, **77**, pp. 585-598.
- [33] Levenstein, M. and Suslow, V. (2006). "What Determines Cartel Success?," *Journal of Economic Literature*, **44**, pp. 43-95.
- [34] Maitra, A. and Sudderth, W. (1996). *Discrete Gambling and Stochastic Games*. New York: Springer-Verlag.

- [35] Mertens, J.-F. and Parthasarathy, T. (1987). "Equilibria for Discounted Stochastic Games," C.O.R.E. Discussion Paper 8750.
- [36] Mertens, J.-F. and Parthasarathy, T. (1991). "Nonzero-sum Stochastic Games," in T. Raghavan, T. Ferguson, T. Parthasarathy, and O. Vrieze, eds., *Stochastic Games and Related Topics*, Boston: Kluwer Academic Publishers.
- [37] Milnor, J., and Shapley, L. (1957). "On Games of Survival," in M. Dresher, A.W. Tucker, and P. Wolfe, eds., *Contributions to the Theory of Games III*, Princeton, NJ: Princeton University Press.
- [38] Motta, M. (2004). *Competition Policy: Theory and Practice*. Cambridge, UK: Cambridge University Press.
- [39] Poitevin, M. (1989). "Financial Signalling and the 'Deep-Pocket' Argument," *RAND Journal of Economics*, **20**, pp. 26-40.
- [40] Porter, R. (1983) "A Study of Cartel Stability: The Joint Executive Committee, 1880-1886," *Bell Journal of Economics*, **14**, pp. 301-314.
- [41] Powell, R. (1999). *In the Shadow of Power*. Princeton, NJ: Princeton University Press.
- [42] Powell, R. (2004). "The Inefficient Use of Power: Costly Conflict with Complete Information," *American Political Science Review*, **98**, pp. 231-241.
- [43] Ranehan, Jr., E. (2007). *Commodore: The Life of Cornelius Vanderbilt*. New York: Basic Books.
- [44] Rosenthal, R. and Rubinstein, A. (1984). "Repeated Two-Player Games with Ruin," *International Journal of Game Theory*, **13**, pp. 155-177.
- [45] Rotemberg, J. and Saloner, G. (1986). "A Supergame-Theoretic Model of Price Wars during Booms," *American Economic Review*, **76**, pp. 390-407.
- [46] Scott Morton, F. (1997). "Entry and Predation: British Shipping Cartels 1879-1929," *Journal of Economics & Management Strategy*, **6**, pp. 679-724.
- [47] Shubik, M. (1959). *Strategy and Market Structure: Competition, Oligopoly, and the Theory of Games*. New York: John Wiley and Sons, Inc.
- [48] Shubik, M. and Thompson, G. (1959). "Games of Economic Survival," *Naval Research Logistics Quarterly*, **6**, pp. 111-123.
- [49] Solan, E. (1998). "Discounted Stochastic Games," *Mathematics of Operations Research*, **23**, pp. 1010-1021.
- [50] Sorin, S. (1986). "Asymptotic Properties of a Non-Zero Sum Stochastic Game," *International Journal of Game Theory*, **15**, pp. 101-107.
- [51] Tirole, J. (2006). *The Theory of Corporate Finance*. Princeton, NJ: Princeton University Press.
- [52] Ulen, T. (1978). *Cartels and Regulation*. Ph.D. dissertation, Stanford University.
- [53] Weiman, D. and Levin, R. (1994). "Preying for Monopoly? The Case of Southern Bell Telephone Company, 1894-1912," *Journal of Political Economy*, **102**, pp. 103-126.

Online Appendix

APPENDIX D. PROOF FROM SECTION 4.1

Proof of Claim 1. Let $\alpha(k)$ denote the value of $\alpha_1(k, 1)$, the guaranteed probability of winning a price war for Firm 1 when it is in state k and Firm 2 is in state 1. Observe that $\alpha(K) = \alpha^*(2)$, so Condition 1 is equivalent to $\alpha(K) > R/2$.

The function $\alpha(k)$ satisfies the following system of equations:

$$\begin{aligned}\alpha(1) &= \gamma^D(1 - \theta\gamma^D) + (1 - \gamma^D)(1 - \theta\gamma^D)\alpha(1) \\ \alpha(k) &= \gamma^D + (1 - \gamma^D) \left[(1 - \theta\gamma^D)\alpha(k) + \theta\gamma^D\alpha(k - 1) \right] \text{ for } k > 1.\end{aligned}$$

That is, when both firms are in state 1, then Firm 1 wins the price war if Firm 2's state deteriorates and Firm 1's does not. If neither firm's state changes, then Firm 1 has again probability $\alpha(1)$ of winning. Similarly, when Firm 1 starts at state $k > 1$, then it wins for sure if Firm 2's state moves down, it wins with probability $\alpha(k - 1)$ if only Firm 1's state moves down, and it wins with probability $\alpha(k)$ if neither state changes. Recall that in the example a firm's state cannot move up during a price war.

Rearranging yields

$$\alpha(1) = \frac{1 - \theta\gamma^D}{1 + \theta(1 - \gamma^D)}, \quad \alpha(k) = 1 - \beta + \beta\alpha(k - 1) \text{ for } k > 1,$$

where

$$\beta \equiv \frac{\theta(1 - \gamma^D)}{1 + \theta(1 - \gamma^D)} \in (0, 1).$$

The dependence of β on θ and γ^D is suppressed for readability. Taking differences yields

$$\Delta(2) = (1 - \beta)(1 - \alpha(1)), \quad \Delta(k) = \beta\Delta(k - 1) \text{ for } k > 2,$$

where $\Delta(k) \equiv \alpha(k) - \alpha(k - 1)$. Thus,

$$\begin{aligned}\alpha(k) &= \alpha(1) + \Delta(2) + \Delta(3) + \dots + \Delta(k) \\ &= \alpha(1) + \Delta(2) + \beta\Delta(2) + \dots + \beta^{k-2}\Delta(2) \\ &= \alpha(1) + \frac{1 - \beta^{k-1}}{1 - \beta}\Delta(2) \\ &= \alpha(1) + (1 - \alpha(1))(1 - \beta^{k-1}),\end{aligned}$$

so $\alpha(K) > R/2$ if

$$\begin{aligned} \frac{R}{2} &< \alpha(1) + (1 - \alpha(1))(1 - \beta^{K-1}) \\ &\Leftrightarrow \\ \frac{\theta^K(1-\gamma^D)^{K-1}}{[1+\theta(1-\gamma^D)]^K} &= \beta^{K-1}(1 - \alpha(1)) < 1 - \frac{R}{2}. \end{aligned}$$

The left-hand side is maximized at $\gamma^D = (\theta - K + 1)/\theta$, so plugging in that value gives the sufficient condition

$$\frac{\theta/(K-1)}{[K/(K-1)]^K} < 1 - \frac{R}{2}.$$

At $K = 2$, that condition becomes

$$\frac{\theta}{4} < 1 - \frac{R}{2}.$$

Rearranging yields the second part of the claim.

For the first part of the claim, note that $[K/(K-1)]^K > e$, so $\alpha(K) > R/2$ if

$$\begin{aligned} \frac{\theta/(K-1)}{[K/(K-1)]^K} &< \frac{\theta/(K-1)}{e} < 1 - \frac{R}{2} \\ &\Leftrightarrow \\ K &> \frac{\theta}{(1-\frac{R}{2})e} + 1, \end{aligned}$$

as desired. □

APPENDIX E. PROOFS FROM SECTION 5

All the examples in this section are from the Bertrand case, and in all the following propositions, “sustainable” means that there is an SPE that generates the specified outcome, and “unsustainable” means that there is no Nash equilibrium that yields that outcome.

E.1. Patience and cartel size.

Example 5. There are two, three, or four firms ($N \in \{2, 3, 4\}$) and two non-bankruptcy states ($S = \{1, 2\}$). The discount factor is either $\delta = 0.8$ or $\delta = 0.82$. Transition rates satisfy the following:

- $\Gamma(\pi, 2)[1] = 0.1$ and $\Gamma(\pi, 2)[0] = 0$ for all $\pi \geq 0$;
- $\Gamma(\pi, 1)[1] > 0.1$ for all $\pi \geq 0$ and $\Gamma(0, 1)[1] = 1$; and

- $\Gamma(0, 1)[0] = 0.09$ and $\Gamma(\pi, 1)[0] = 0$ for all $\pi > 0$.

That is, the probability that a strong firm becomes weak is 0.1 for any level of profits. The probability that a weak firm becomes strong is less than 0.9 for any profit and 0 given no profit. Only a weak firm earning zero profit can go bankrupt, and the probability that such a firm goes bankrupt is 0.09.

Example 6. There are two, three, or four firms ($N \in \{2, 3, 4\}$) and two non-bankruptcy states ($S = \{1, 2\}$). The discount factor is either $\delta = 0.9$ or $\delta = 0.95$. Transition rates satisfy the following:

- $\Gamma(\pi, 2)[1] = 0.15$ and $\Gamma(\pi, 2)[0] = 0$ for all $\pi \geq 0$;
- $\Gamma(\pi, 1)[1] > 0.15$ for all $\pi \geq 0$ and $\Gamma(0, 1)[1] = 1$; and
- $\Gamma(0, 1)[0] = 0.1$ and $\Gamma(\pi, 1)[0] = 0$ for all $\pi > 0$.

Proposition 1. *In Example 5, when $N = 4$, collusion on the monopoly price π^M is unsustainable when $\delta = 0.8$ but sustainable when $\delta = 0.82$. In Example 6, when $N = 3$, collusion on the monopoly price π^M is sustainable when $\delta = 0.9$ but unsustainable when $\delta = 0.95$.*

Proposition 2. *In Example 5, when $\delta = 0.8$, collusion on the monopoly price π^M is sustainable when $N = 2$ or $N = 3$ but unsustainable when $N = 4$. In Example 6, when $\delta = 0.9$, collusion on the monopoly price π^M is unsustainable when $N = 2$ or $N = 4$ but sustainable when $N = 3$.*

Proof of Propositions 1 and 2. Note that a price war (where all firms price at 0 until only one is left) is always an SPE, and it gives all players their minmax payoffs. Thus, a given on-path strategy profile is sustainable (in either Nash or subgame perfect equilibrium) if and only if no firm can gain by a one-shot deviation followed by a price war.

First consider Example 5, and suppose that there are two active firms ($n = 2$). Let $v_2^{ss'}(\delta)$ denote the expected continuation payoff after a deviation (that is, the payoff from a price war) to a firm in state s whose rival is in state s' . In particular,

$$\begin{aligned} v_2^{21}(\delta) &= (1 - \delta)0 + \delta \left[.09\pi^M + (1 - .09) \cdot .1v_2^{11}(\delta) + (1 - .09)(1 - .1)v_2^{21}(\delta) \right] \\ v_2^{11}(\delta) &= (1 - \delta)0 + \delta \left[.09 \cdot 0 + (1 - .09) \cdot .09\pi^M + (1 - .09)^2 v_2^{11}(\delta) \right]. \end{aligned}$$

Collusion on the monopoly price p^M is part of an SPE if a firm prefers collusive profits ($\pi^M/2$) to the profit from an optimal deviation: undercutting its rival and starting a price war. Note that if a firm prefers not to deviate when it is strong (state 2) and its rival is weak (state 1), then it cannot gain from deviating in any other state vector. The reason is that the equilibrium payoff is the same for any state vector, and so is the one-shot profit from undercutting, while the expected continuation payoff in a price war $\underline{v}_2^{ss'}(\delta)$ is increasing in the firm's state s tomorrow and decreasing in its rival's s' ; the condition in both examples that $\Gamma(\pi, 1)[1] > \Gamma(\pi, 2)[1]$ for all π implies that tomorrow's state is positively correlated with today's for each firm.

Let $\Delta_2(\delta)$ denote the payoff from that optimal deviation:

$$\Delta_2(\delta) \equiv (1 - \delta)\pi^M + \delta \left[.09\pi^M + (1 - .09) \cdot 1\underline{v}_2^{11}(\delta) + (1 - .09)(1 - .1)\underline{v}_2^{21}(\delta) \right].$$

At $\delta = 0.8$, solving yields $\underline{v}_2^{21}(0.8) \approx 0.250\pi^M$, $\underline{v}_2^{11}(0.8) \approx 0.194\pi^M$, and $\Delta_2(0.8) \approx 0.450\pi^M$. Since $\Delta_2(0.8) < \frac{1}{2}\pi^M$, no deviation is profitable, and collusion is sustainable.

When there are three active firms, again the best time to deviate is when a firm is strong and its rivals are weak. The relevant price war payoffs in that case are

$$\begin{aligned} \underline{v}_3^{211}(\delta) &= (1 - \delta)0 + \\ \delta \left[(.09)^2 \pi^M + 2 \cdot .09(.91) \cdot 1\underline{v}_2^{11}(\delta) + .9\underline{v}_2^{21}(\delta) \right] &+ (.91)^2 \left[.1\underline{v}_3^{111}(\delta) + .9\underline{v}_3^{211}(\delta) \right], \\ \underline{v}_3^{111}(\delta) &= (1 - \delta)0 + \\ \delta \left[.09 \cdot 0 + (.09)^2 (.91)\pi^M + 2 \cdot .09(.91)^2 \underline{v}_2^{11}(\delta) + (.91)^3 \underline{v}_3^{111}(\delta) \right]. \end{aligned}$$

The payoff from the optimal deviation, $\Delta_3(\delta)$, is

$$(1 - \delta)\pi^M + \delta \left[(.09)^2 \pi^M + 2 \cdot .09(.91) \cdot 1\underline{v}_2^{11}(\delta) + .9\underline{v}_2^{21}(\delta) \right] + (.91)^2 \left[.1\underline{v}_3^{111}(\delta) + .9\underline{v}_3^{211}(\delta) \right].$$

At $\delta = 0.8$, solving yields $\underline{v}_3^{211}(0.8) \approx 0.107\pi^M$, $\underline{v}_3^{111}(0.8) \approx 0.073\pi^M$, and $\Delta_3(0.8) \approx 0.307\pi^M$. Since $\Delta_3(0.8) < \frac{1}{3}\pi^M$, no deviation is profitable, and collusion is again sustainable. Repeating that analysis for the case of four active firms, however, yields $\underline{v}_4^{2111}(0.8) \approx 0.056\pi^M$, $\underline{v}_{34}^{1111}(0.8) \approx 0.034\pi^M$, and $\Delta_4(0.8) \approx 0.256\pi^M$. Since $\Delta_4(0.8) > \frac{1}{4}\pi^M$, a firm gains by undercutting when it is strong and its rivals are all weak. Thus, collusion is not sustainable.

In summary: in Example 5, collusion on the monopoly price is sustainable at $\delta = 0.8$ when there are two or three active firms, but not when there are four. Applying the same analysis for $\delta = 0.82$ yields $\Delta_4(0.82) \approx 0.247\pi^M$. Since $\Delta_4(0.82) < \frac{1}{4}\pi^M$, in that setting collusion is sustainable.

In Example 6, similar calculations yield $\Delta_2(0.9) \approx 0.506\pi^M$, $\Delta_3(0.9) \approx 0.3328\pi^M$, $\Delta_4(0.9) \approx 0.252\pi^M$, and $\Delta_3(0.95) \approx 0.395\pi^M$. Thus, at $\delta = 0.9$, collusion on the monopoly price is sustainable by three firms but not by two or four. At $\delta = 0.95$, such collusion is not sustainable by three firms, either. \square

E.2. Unequal market shares.

Example 7. There are two firms ($N = 2$) and two non-bankruptcy states ($S = \{1, 2\}$). The discount factor is $\delta = 0.95$. Transition rates satisfy the following:

- $\Gamma(\pi, 2)[1] = 0.17$ for $\pi \in [0, 0.75\pi^M]$, $\Gamma(\pi, 2)[1] = 0.05$ for $\pi > 0.75\pi^M$, and $\Gamma(\pi, 2)[0] = 0$ for all $\pi \geq 0$;
- $\Gamma(\pi, 1)[2] = 0.05$ for $\pi \in (0, 0.25\pi^M)$, $\Gamma(\pi, 1)[2] = 0.1$ for $\pi \geq 0.25\pi^M$, and $\Gamma(0, 1)[2] = 0$; and
- $\Gamma(0, 1)[0] = 0.1$ and $\Gamma(\pi, 1)[0] = 0$ for all $\pi > 0$.

Proposition 3. *In Example 7, collusion on the monopoly price π^M where the firms split the monopoly profit π^M each period is unsustainable. If the firms can divide market demand (or, equivalently, if they can transfer money to each other), then there is a collusive SPE in which (i) the firms set the monopoly price p^M each period, (ii) when $s_1 = s_2$, both firms get profit $\pi^M/2$, and (iii) when $s_i > s_j$, firm i gets profit $0.9\pi^M$ and firm j gets profit $0.1\pi^M$.*

Proof. The proof is similar to the proofs of Propositions 1 and 2. First, calculating the value of \underline{v}^{21} , the expected continuation payoff from a price war to a firm in state 2 whose rival is in state 1, yields $\underline{v}^{21} \approx 0.513\pi^M > \frac{1}{2}\pi^M$, so equal sharing is not sustainable: the strong firm would prefer to start a price war.

As before, a price war is always a SPE. Therefore, again as before, to prove that the specified unequal sharing strategy is a SPE, it is sufficient to show that no firm can gain by a one-shot, on-path deviation followed by a price war. Let $\hat{v}^{ss'}$ denote the

expected continuation payoff from the specified strategy to a firm in state s whose rival is in state s' . In particular,

$$\begin{aligned}\hat{v}^{21} &= (1 - \delta).9\pi^M + \delta [(.05)^2\hat{v}^{12} + (.05)(.95)(\hat{v}^{22} + \hat{v}^{11}) + (.95)^2\hat{v}^{21}] \\ \hat{v}^{12} &= (1 - \delta).1\pi^M + \delta [(.05)^2\hat{v}^{21} + (.05)(.95)(\hat{v}^{22} + \hat{v}^{11}) + (.95)^2\hat{v}^{12}] \\ \hat{v}^{22} &= (1 - \delta).5\pi^M + \delta [(.17)^2\hat{v}^{11} + (.17)(.83)(\hat{v}^{21} + \hat{v}^{12}) + (.83)^2\hat{v}^{22}] \\ \hat{v}^{11} &= (1 - \delta).5\pi^M + \delta [(.1)^2\hat{v}^{22} + (.1)(.9)(\hat{v}^{21} + \hat{v}^{12}) + (.9)^2\hat{v}^{11}].\end{aligned}$$

Solving yields $\hat{v}^{21} \approx 0.638\pi^M$, $\hat{v}^{12} \approx 0.362\pi^M$, and $\hat{v}^{22} = \hat{v}^{11} = 0.5\pi^M$. Next, let $\Delta^{ss'}$ denote the expected continuation payoff from the optimal deviation (undercutting the price p^M and starting a price war) to a firm in state s whose rival is in state s' :

$$\begin{aligned}\Delta^{21} &= (1 - \delta)\pi^M + \delta [.1\pi^M + .9(.05\underline{v}^{11} + .95\underline{v}^{21})] \\ \Delta^{12} &= (1 - \delta)\pi^M + \delta [(.1)(.17)\underline{v}^{21} + (.1)(.83)\underline{v}^{22} + (.9)(.17)\underline{v}^{11} + (.9)(.83)\underline{v}^{12}] \\ \Delta^{22} &= (1 - \delta)\pi^M + \delta [(.05)(.17)\underline{v}^{11} + (.05)(.83)\underline{v}^{12} + (.95)(.17)\underline{v}^{21} + (.95)(.83)\underline{v}^{22}] \\ \Delta^{11} &= (1 - \delta)\pi^M + \delta [.1\pi^M + .9(.1\underline{v}^{21} + .9\underline{v}^{11})].\end{aligned}$$

Solving yields $\Delta^{21} \approx 0.577\pi^M$, $\Delta^{12} \approx 0.268\pi^M$, $\Delta^{22} \approx 0.364\pi^M$, and $\Delta^{11} \approx 0.474\pi^M$. Since $\Delta^{ss'} < \hat{v}^{ss'}$ for every state vector ss' , deviating is never profitable, and the specified strategy is an SPE. \square

E.3. Recurring collusion.

Example 8. There are two firms ($N = 2$) and two non-bankruptcy states ($S = \{1, 2\}$). The discount factor is $\delta = 0.95$. Transition rates satisfy the following:

- $\Gamma(\pi, 2)[1] = 0.05$ and $\Gamma(\pi, 2)[0] = 0$ for all $\pi \geq 0$;
- $\Gamma(\pi, 1)[1] = 0.99$ for all $\pi \geq \pi^M/2$ and $\Gamma(0, 1)[1] = 1$; and
- $\Gamma(0, 1)[0] = 0.06$ and $\Gamma(\pi, 1)[0] = 0$ for all $\pi > 0$.

Proposition 4. *In Example 8, collusion on the monopoly price π^M is unsustainable. There is an SPE in which when both firms are active, they set the monopoly price p^M in any period in which $s_1 = s_2$, and set price 0 otherwise.*

Proof. Define the strategy σ^* as follows: in the first period, and after any history in which neither firm has deviated, actions are as specified in the proposition. After a deviation, both firms set price 0 until one firm goes bankrupt. For states $s, s' \in$

$\{1, 2\}^2$, let $v^{ss'}$ denote the expected payoff to a firm in state s whose rival is in state s' if both follow σ^* . Those values satisfy the following equations:

$$\begin{aligned} v^{21} &= (1 - \delta)0 + \delta \left[.06\pi^M + (1 - .06).05v^{11} + (1 - .06)(1 - .05)v^{21} \right] \\ v^{12} &= (1 - \delta)0 + \delta \left[.06 \cdot 0 + (1 - .06).05v^{11} + (1 - .06)(1 - .05)v^{12} \right] \\ v^{22} &= (1 - \delta)\frac{1}{2}\pi^M + \delta \left[(1 - .05).05 [v^{21} + v^{12}] + .05^2v^{11} + (1 - .05)^2v^{22} \right] \\ v^{11} &= (1 - \delta)\frac{1}{2}\pi^M + \delta \left[(1 - .01).01 [v^{21} + v^{12}] + .01^2v^{22} + (1 - .01)^2v^{11} \right]. \end{aligned}$$

Solving yields $v^{21} \approx 0.509\pi^M$, $v^{12} \approx 0.133\pi^M$, $v^{22} \approx 0.386\pi^M$, and $v^{11} \approx 0.451\pi^M$. Similarly, let $\underline{v}^{ss'}$ denote the expected continuation payoff after a deviation (that is, the payoff from a price war) to a firm in state s whose rival is in state s' :

$$\begin{aligned} \underline{v}^{21} &= (1 - \delta)0 + \delta \left[.06\pi^M + (1 - .06).05\underline{v}^{11} + (1 - .06)(1 - .05)\underline{v}^{21} \right] \\ \underline{v}^{12} &= (1 - \delta)0 + \delta \left[.06 \cdot 0 + (1 - .06).05\underline{v}^{11} + (1 - .06)(1 - .05)\underline{v}^{12} \right] \\ \underline{v}^{22} &= (1 - \delta)0 + \delta \left[(1 - .05).05 [\underline{v}^{21} + \underline{v}^{12}] + .05^2\underline{v}^{11} + (1 - .05)^2\underline{v}^{22} \right] \\ \underline{v}^{11} &= (1 - \delta)0 + \delta \left[.06 \cdot 0 + (1 - .06).06\pi^M + (1 - .06)^2\underline{v}^{11} \right]. \end{aligned}$$

Solving yields $\underline{v}^{21} \approx 0.474\pi^M$, $\underline{v}^{12} \approx 0.098\pi^M$, $\underline{v}^{22} \approx 0.187\pi^M$, and $\underline{v}^{11} \approx 0.334\pi^M$.

To verify that σ^* is an SPE, it is necessary to check only that neither firm wants to deviate first. (After the first deviation, the price that a firm sets does not affect its payoff.) When both firms are in state 2, a firm's best deviation would be to a price just below p^M , yielding an expected payoff of

$$(1 - \delta)\pi^M + \delta \left[(1 - .05).05 [\underline{v}^{21} + \underline{v}^{12}] + .05^2\underline{v}^{11} + (1 - .05)^2\underline{v}^{22} \right] \approx 0.237\pi^M.$$

Since 0.239 is less than the payoff from σ^* , $v^{22} \approx 0.386\pi^M$, a firm cannot gain from deviating. Similarly, when both firms are in state 1, undercutting yields

$$(1 - \delta)\pi^M + \delta \left[.06\pi^M + (1 - .06).01\underline{v}^{21} + (1 - .06)(1 - .01)\underline{v}^{11} \right] \approx 0.406\pi^M.$$

The payoff from σ^* , $v^{11} \approx 0.451\pi^M$, is higher, so again the deviation is not profitable.

Finally, note that when the firms have asymmetric strengths (and are setting price 0 under σ^*), deviating to a different price has no effect on today's profit and lowers continuation payoffs (since $\underline{v}^{ss'} < v^{ss'}$ for all $s, s' \in \{1, 2\}^2$). Thus, σ^* is an SPE.

To see that there is no Nash equilibrium in which the two firms always collude on the monopoly price, observe that in that case a strong firm would gain by undercutting

a weak rival: the resulting payoff,

$$(1 - \delta)\pi^M + \delta \left[.06\pi^M + (1 - .06).05\underline{v}^{11} + (1 - .06)(1 - .05)\underline{v}^{21} \right] \approx 0.524\pi^M,$$

strictly exceeds $\frac{1}{2}\pi^M$. □

In the equilibrium of Proposition 4, on average a price war lasts for 9.3 periods. It ends either when the weaker firm goes bankrupt (probability 0.56) or when the stronger firm's state declines and the now-evenly matched firms start to collude again (probability 0.44). On average, then, there are 1.8 distinct price wars before one firm goes bankrupt. When both firms are initially strong, the interval of collusion lasts on average for 11.5 periods; when both are weak, collusion lasts for 50.3 periods on average. Thus, from an initial state vector where both firms are weak, the expected time until one firm goes bankrupt is 106 periods. By comparison, if both firms priced at 0 until one went bankrupt, the expected time till bankruptcy would be only 8.6 periods. Significant collusion can occur, although that collusion is only temporary.