informs
http://pubsonline.informs.org/journal/mnsc/

MANAGEMENT SCIENCE
Vol. 64, No. 5, May 2018, pp. 2013–2030
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

We're Number 1: Price Wars for Market Share Leadership

Luís Cabral[a, b]

[a] Stern School of Business, New York University, New York, New York 10012; [b] Centre for Economic and Policy Research, London EC1V 0DX, United Kingdom

Contact: luis.cabral@nyu.edu (LC)

Received: March 30, 2015
Revised: April 28, 2016
Accepted: September 2, 2016
Published Online in Articles in Advance: May 9, 2017

https://doi.org/10.1287/mnsc.2017.2725

Abstract. I examine the dynamics of oligopolies when firms derive subjective value from market leadership. In equilibrium, prices alternate in tandem between high levels and occasional price wars, which take place when market leadership is at stake. The stationary distribution of market shares is typically multimodal; that is, much of the time, there is a stable market leader. Even though shareholders do not value market leadership per se, a corporate culture that values market leadership may increase shareholder value.

History: Accepted by Bruno Cassiman, business strategy.
Supplemental Material: The online appendix is available at https://doi.org/10.1287/mnsc.2017.2725.

Keywords: dynamic oligopoly • price wars • market shares • ordinal rankings

1. Introduction

In many industries, the prevailing managerial attitude places a disproportionate weight on being number one—the market share leader. For example, in May 2012, Airbus "accused Boeing Co. of trying to start a price war after the U.S. company pledged to work aggressively to regain a 50% share of the market" (Ostrower 2012). A February 2011 headline announced that "IBM reclaims server market share revenue crown in Q4," adding that "IBM and HP [Hewlett-Packard] will continue to duke it out" (Dignan 2011). According to CNN, "GM [General Motors] held onto its No. 1 rank by cutting prices on cars to the point where they were unprofitable" (Isidore 2012). And during a 2007 interview with a group of bloggers, SAP chief executive officer Henning Kagermann stated, "We are not arrogant, we are the market leader" (Farber 2007).[1]

In this paper, I examine the implications of ordinal comparisons, in particular the number one bias, for market competition. Specifically, I examine the behavior of managers who receive an extra utility kick from being market share leaders.[2] (I do not develop a theory to explain why managers derive utility from being market leaders, though I do discuss some rational and behavioral reasons for this pattern.)

I develop a model with two sellers and multiple buyers, all of whom live forever. Buyers reassess their choice of seller at random points in time. Buyers have preference for sellers and for money. Sellers have a preference for money and for being number one. The paper makes two central points, one normative and one positive. First, I show that a corporate culture that emphasizes the importance of market leadership may increase shareholder value even if shareholders do not care about market leadership (or market shares) per se. Second, I propose a simple model that leads to a rich theory of price wars and the evolution of market shares.

Specifically, I show that a firm's utility from begin market leader implies a price drop when market shares are close to 50%, and thus a lot is at stake. Moreover, I provide conditions such that, fearful of entering into a price war, competition is softened at states close to the price war region—so much so that shareholder value increases with respect to a situation where managers do not care about market share leadership. The softening of price competition also implies that the stationary distribution of market shares is bimodal; that is, most of the time, one firm is larger than the other one—and occasionally price wars for market share take place.

My paper also has implications for a central question in strategy and industrial organization: the persistence of differences across firms. Typically, these are explained by primitive differences across firms, such as unique resources (Wernerfelt 1984, Barney 1986, Dierickx and Cool 1989); endogenous difference due to increasing returns, such as learning curves or network effects (Cabral and Riordan 1994, Cabral 2011); or stickiness in market shares due to switching costs or related effects (Beggs and Klemperer 1992). My model features none of the above characteristics and still induces a stationary distribution of market shares that can be bimodal. In other words, for a "long" period of time, there is a large firm and a small firm (even though the model is symmetric); the only barrier to mobility that stops the small firm from becoming large is the price war it must go through to increase market share.

In this sense, my model also provides a new perspective on the concept of mobility barriers. In a seminar paper, Caves and Porter (1977) proposed an extension of the theory of entry barriers, one that goes beyond the movement of a firm from zero output to some positive level: for example, in some cases, established firms enter a new segment of a given industry. Exogenous or endogenous impediments to such segment entry are denoted mobility barriers. My theory of dynamic price competition suggests an additional instance of intraindustry mobility: a firm that is a market share follower becoming a market share leader. To the extent that the stationary distribution of market shares is multimodal (as I will show is frequently the case), this shift in relative positions is sufficiently "discontinuous" that the analogy of mobility barriers is meaningful. The barrier I will consider is endogenous and results from the market leader's aggressive price behavior when the laggard's market share becomes threateningly close to the leader's.

1.1. Related Literature and Contribution

The paper makes several contributions to the strategy and industrial organization literatures. First, it studies the implications of a fairly pervasive phenomenon—namely, firms' desire to be market share leaders. Baumol (1962) and others have developed models where firms follow objectives other than profit maximization. However, to the best of my knowledge this is the first paper in the industrial organization literature explicitly to consider pricing dynamics when number one effects are in place.

Second, I develop a realistic theory of price wars. For all of the richness of industrial organization theory, the core theory of price wars is still connected almost exclusively to collusion models. In Green and Porter (1984), price wars result from the breakdown of collusive equilibria during periods of (unobservable) low demand. Rotemberg and Saloner (1986) suggest that price wars correspond to firms refraining from collusion during periods of observable high demand.[3] By contrast, I assume that firms do not collude (they play Markov strategies). Instead of a repeated game, I assume firms play a dynamic game where the state is defined by each firm's market share. In this context, price wars emerge in states where a firms' value function is particularly steep—that is, during periods when a firm's gain from increasing market share is particularly high. In this sense, the pricing equilibrium in my model bears some resemblance to models with learning or network effects (Cabral and Riordan 1994, Besanko et al. 2010, Cabral 2011). However, the dynamics in these papers are driven by increasing returns, whereas I consider a setting with constant returns to scale.

Third, I provide an instance where corporate culture has a clear influence on the way firms compete. Specifically, I provide conditions such that a deviation from profit maximization may in effect lead to higher firm value. Vickers (1985), Fershtman and Judd (1987), and Sklivas (1987) have shown that profit-seeking shareholders may have an interest in delegating decisions to managers based on incentive mechanisms that differ from profit maximization.[4] Specifically, if the firms' decision variables are strategic complements (as is the case in my model), then equilibrium delegation contracts ask managers to pay less importance to profits than shareholders would: such contracts "soften" price competition and lead to overall higher profits than in the "normal" price competition game. My approach to delegation is very different, and so are the results. Specifically, number one effects ask firms to place more weight on market shares than shareholders would. This makes firms more, not less, aggressive.

1.2. Roadmap

The rest of this paper is structured as follows. In Section 2, as a prelude to the full-fledged framework, I develop a simple, two-stage model of foremarket and aftermarket competition. Although the model is rather stylized, it conveys the intuition that a "corporate culture" of striving to be number one may induce a credible commitment that is valuable to shareholders (who, by assumption, do not share the manager's "utility" from being market share leaders). In Sections 3–5, I develop an infinite-period, multistate model of price competition for market share. Besides confirming some of the qualitative features of Section 2, I also propose a novel theory of price wars for market shares.[5] Section 6 discusses robustness and extensions of the basic results. Section 7 concludes the paper.

2. A Two-Stage Model

Consider a two-stage model with a foremarket and an aftermarket. For example, the foremarket may correspond to a hardware purchase and the aftermarket to some consumable. Each consumer first chooses a seller in the foremarket, willing to pay v for at most one unit from either seller's product. Upon purchasing the basic product, a consumer is willing to pay μ in the aftermarket for at most one unit. However, if a consumer purchases in the aftermarket from a different seller than in the foremarket, then the consumer must pay an additional switching cost σ. The value of σ is each consumer's private information and has the following distribution:

$$\sigma = \begin{cases} \sigma_L & \text{with probability } 1-\lambda, \\ \sigma_H & \text{with probability } \lambda, \end{cases}$$

where $\sigma_H > \sigma_L > 0$ and $\lambda \in (0,1)$. Suppose that two firms, a and b, have equal-sized consumer installed bases, $n_a = n_b = n$, where

$$n > v/\mu. \tag{1}$$

The timing of the game is as follows. First, firms simultaneously set prices p_i $(i = a, b)$ for a new consumer coming on the market (foremarket). Next, the new consumer chooses to buy from firm a, firm b, or none. Next, firms simultaneously set prices q_i $(i = a, b)$ in the aftermarket. Finally, all consumers choose whether to purchase in the aftermarket from the same firm they purchased before or rather from a different firm (paying σ if they do so).

I now come to the central feature of the model: the benefit from leadership. I assume that firm a's corporate culture is such that its manager enjoys an additional payoff θ when it is the leader as measured by the aftermarket installed base; that is, when $n_a > n_b$.[6] When $\theta > 0$, I distinguish between firm a's shareholder value, which corresponds to firm a's profit (from foremarket and aftermarket sales) from firm a's manager's utility, which includes, *in addition*, the benefit from market leadership (if it applies).

The main result is that although shareholders do not benefit from market leadership, they may benefit from a culture that creates such a benefit in the manager's eyes.

Proposition 1. *There exists $\bar{\theta}$ such that if $\theta \geq \bar{\theta}$, then firm a's equilibrium shareholder value is greater when $\theta \geq \bar{\theta}$ than with $\theta = 0$.*

The proof of this and subsequent results may be found in the appendix. The intuition for Proposition 1 is that, by creating a corporate culture of market leadership, firm a effectively commits to becoming very aggressive should it lose its market leadership (which happens when firm b makes a sale in the foremarket). Aggressive pricing by firm a is harmful to firm b. In fact, it implies that firm a poaches all of firm b's customers in the aftermarket. Fearing such aggressive pricing, firm b "softens up" its foremarket behavior. Finally, such softening up by firm b increases firm a's foremarket profits—and, along the equilibrium path, the event of aggressive pricing by firm a never takes place, so firm a's shareholders are better off than they would be if firm a's managers were straight value maximizers.

Proposition 1 corresponds to the first of the two main points in the paper: the normative point that a corporate culture of seeking to be number one induces a credible threat of price aggression in states when market leadership is at stake. While the execution of the threat destroys shareholder value, the strategic commitment itself increases shareholder value. If, along the equilibrium path, the threat does not need to be followed through, then only the positive effect remains.

The idea that credible threats may improve a firm's position is hardly new (see, for example, the "top dog" strategy described in Fudenberg and Tirole 1984). The novel idea is that the pervasive culture of market leadership may do that job. In other words, one may think of corporate culture as a first-stage "investment" that influences the outcome of a second-stage market game.

The idea that contracts written between shareholders and managers may have strategic value is also not novel in of itself. Vickers (1985), Fershtman and Judd (1987), and Sklivas (1987) have shown that profit-seeking shareholders may have an interest in delegating decisions to managers based on incentive mechanisms that differ from profit maximization. Specifically, if the firms' decision variables are strategic complements (as is the case in my model), then equilibrium delegation contracts ask managers to pay less importance to profits than shareholders would: such contracts "soften" price competition and lead to overall higher profits than in the "normal" price competition game. My approach is very different, and so are the results, essentially because my approach is dynamic—that is, the game I consider evolves over a series of states—whereas that of Vickers (1985), Fershtman and Judd (1987), and Sklivas (1987) is essentially static. Specifically, number one effects ask firms to place more weight on market shares than shareholders would. This makes firms more, not less, aggressive. From a static point of view, this effect is bad news for shareholders, for excessively aggressive pricing means lower equilibrium profits. However, the price wars that follow from number one effects are rare (in the above example, they do not take place at all); the negative effect of overly aggressive pricing is more than compensated for by the deterrence effect implied by the threat of a price war.

2.1. Corporate Culture as Business Strategy

Proposition 1 begs the question of when a firm would want to have a number one culture. Formally, we can answer this question by considering a "corporate culture metagame" (see Figure 1) where, in a first stage, firms simultaneously choose the value of θ_i $(i = a, b)$; then, the firms having observed the choices of θ_i values, in a second stage, both play the game considered above. There are different versions that the first stage of this game could take. Here, I assume that each firm must choose between not having and having a number

Figure 1. Corporate Culture Metagame

		Firm b: 0	Firm b: θ
Firm a	0	$n\,\mu$ (a); $n\,\mu$ (b)	$n\,\mu$ (a); $n\,\mu + v$ (b)
	θ	$n\,\mu + v$ (a); $n\,\mu$ (b)	$-\theta$ (a); $-\theta$ (b)

one corporate culture, where the latter is defined by a specific value of θ.

Proposition 2. *Suppose that each firm must choose $\theta_i \in \{0, \theta\}$. There exists $\bar{\theta}$ such that, if $\theta \geq \bar{\theta}$, then (a) the choices θ_i are strategic substitutes, and (b) there exist two asymmetric equilibria in pure strategies: $(0, \theta)$ and $(\theta, 0)$.*

In essence, Proposition 2 states that the corporate culture metagame is akin to a game of chicken: both players prefer an aggressive culture given that the rival does not have one, but both players prefer not to have one given that the rival has one. In other words, the choices of an aggressive corporate culture are strategic substitutes.

Part of Proposition 2 is a corollary of Proposition 1: if my rival has no number 1 culture, then I want to have one. The additional part of Proposition 2 is that firms prefer to avoid a clash of aggressive number 1 cultures, in the same way that, in a game of chicken, the worst possible outcome is for both players to choose not to swerve. In the present context, a subgame where both firms have a number one culture induces a price war in the aftermarket, resulting in a loss for shareholders of the order of θ. In fact, that is approximately how far below cost managers are willing to take prices so as to guarantee market share leadership (from which managers, but not shareholders, derive utility).

Proposition 2 states that there exist two asymmetric equilibria in pure strategies. As often is the case in games of this type, there also exists a symmetric equilibrium in mixed strategies: in the first stage, each firm chooses a number one culture with probability p. In turn, this implies that all four possible outcomes—from no firm choosing a number one culture to both firms doing so—takes place with positive probability.

2.2. Summary

We may summarize this section by stating that a number one culture may be used as a value-enhancing commitment to be aggressive in pricing should a firm fall behind its rival. In a metagame where firms choose corporate culture before competing in the market, there exist asymmetric equilibria where one firm (and not the other) chooses a number one culture.

The specific extensive form considered in this section is highly stylized. My main purpose has been to illustrate one of the central points in the paper: the strategic value of committing to a number one corporate culture. In the next two sections, I propose a more general and realistic model, one that evolves over infinite periods and features many possible market share states. The result of higher shareholder value under aggressive number one management will reappear and so will an additional result: a positive theory of price wars—that is, the idea that price wars may result from the aggressive behavior of managers for whom market shares, in particular market leadership, are particularly important.

Specifically, in the next section, I lay down the basic model, including results that establish key properties of its stationary state. In Section 4, I show, by means of analytical results and numerical simulations, how the model leads to a natural theory of price wars. In Section 5, I return to the basic questions examined in the present section—namely, the extent to which commitment to an aggressive corporate culture may benefit shareholders who do not directly care about market share. The value added by Section 5 is that it is based on a more realistic model than the one considered in the present section. In Section 6, I consider a variety of robustness and extensions of the infinite-period model, including the possibility of demand-side number one effects and more than two competing firms.

3. An Infinite-Period Model

Consider a duopoly with two firms, a and b. I will use i and j to designate a firm generically; that is, $i, j = a, b$. Time is discrete and runs indefinitely: $t = 1, 2, \ldots$. The total number of consumers is given by η.

The model dynamics are given by the assumption that agents make "durable" decisions infrequently. Specifically, at random moments in time, a consumer is called to reassess its decision regarding the firm it buys from. One way to think about this is that each consumer's switching cost follows a stochastic process, alternating between the values of infinity (inactive consumer) and zero (active consumer). Alternatively, I may assume that consumers leave the market (death) and are replaced by new consumers in equal number (birth).[7]

Until later in the paper, I consider a symmetric equilibrium of a symmetric game. In particular, I assume both firms have the same corporate culture and play the same pricing strategy. I do this for several reasons. First, the main analytical results are cleaner in the symmetric case. Second, the symmetric model better highlights the distinction between model symmetry and outcome symmetry. Third, as mentioned in the previous section, the symmetric equilibrium of the corporate culture game admits an outcome whereby both firms choose a number one culture.

The timing of the game, as well as the state transition process, are described in Figure 2. Each period starts with each firm having a certain number of consumers, i and j, attached to it (where $i + j = \eta$). Firms set prices $p(i)$ and $p(j)$. I constrain prices to be a function of the state (i, j); that is, I restrict firms to playing Markov strategies. Since the total number of consumers is constant, the state space is one-dimensional and can be summarized by i.

After firms set prices, nature chooses a particular agent, whom I will call the "active" agent. Each agent

Figure 2. Timing

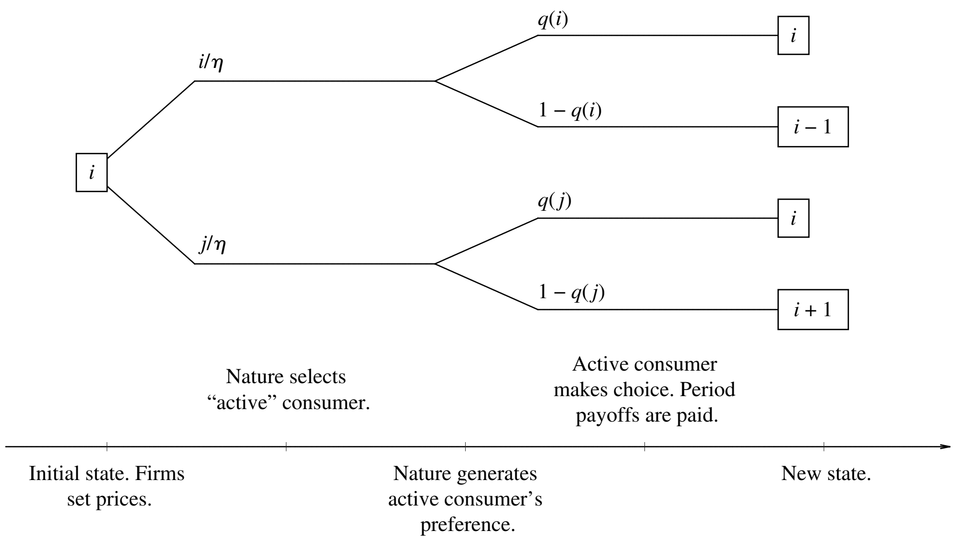

becomes active with equal probability. Then nature generates the active agent's preferences: values ζ_a and ζ_b, corresponding to consumer specific preference for each firm's product. I assume these values are independent and identically distributed (i.i.d.), drawn from a cumulative distribution function (cdf) $\Omega(\zeta)$, and that $\xi \equiv \zeta_a - \zeta_b$ is distributed according to cdf $\Phi(\cdot)$.[8] The active consumer then chooses one of the firms and period payoffs from sales are paid: the sale price to the firm that makes a sale and utility minus price to the consumer who makes a purchase.

In addition to sales revenues, I assume that firm i receives an extra benefit θ if it is the market leader—that is, if $i > j$. To preserve model symmetry, I also assume that if $i < j$, then firm i receives an extra negative benefit θ and that if $i = j$, then both firms receive zero extra benefit. This assumption guarantees that regardless of the state, the firms' joint payoff from market leadership is zero. Market leadership payoff may be summarized by $\theta\Delta(i)$, where $\Delta(i)$ is an indicator variable defined as follows:

$$\Delta(i) \equiv \operatorname{sgn}(i - j) = \begin{cases} +1 & \text{if } i > j, \\ 0 & \text{if } i = j, \\ -1 & \text{if } i < j, \end{cases}$$

where $j = \eta - i$. Recall that this term does not correspond to "real" value; rather, it is simply value perceived by firm i's managers.[9]

There are two sources of randomness in the model. One is that each period 1 consumer is selected by nature to be an active consumer. Second, nature generates utility shocks for the active agent such that the difference $\xi_i \equiv \zeta_i - \zeta_j$ is distributed according to cdf $\Phi(\xi)$. Many of the results below require relatively mild assumptions regarding Φ.

Assumption 1. (i) *The cdf* $\Phi(\xi)$ *is continuously differentiable,* (ii) $\phi(\xi) = \phi(-\xi)$, (iii) $\phi(\xi) > 0, \forall\, \xi$, *and* (iv) $\Phi(\xi)/\phi(\xi)$ *is strictly increasing.*

3.1. A Note on Model Assumptions

The model outlined above is fairly parsimonious. As is often the case, this begs the question of whether it captures reality appropriately. In particular, (a) I measure market shares by adding up previous purchase decisions, and (b) I assume there is one consumer only per period. Regarding the measurement of market shares, I have in mind the situation where consumers make both durable and nondurable purchases. For example, wireless consumers buy smartphones and commit to long-term plans occasionally (when they are "active" agents) and then buy usage on a monthly basis. To the extent that average nondurable purchases are relatively constant across firms, market shares in terms of nondurable purchases correspond to market shares in terms of my state space—that is, i/η. In other words, a firm that is a market leader in a given period is a firm with higher value of i. If variable profits from selling nondurables are zero, then the payoff function I consider is appropriate, for all profit is derived from selling durables. The case when nondurables produce variable profit is considered in Section 6.

Consider now the assumption of one sale per period. As I mentioned earlier, this may be interpreted as the reduced form of a continuous time model where each consumer becomes active as the result of a Poisson process. In that case, assuming the Markov state is given by i amounts to assuming that firms can adjust prices instantly after each purchase. This is obviously a simplifying assumption. In the online appendix, I argue that the main qualitative results are unlikely to change if I allow ν consumers to become active each period, which corresponds to the case when firms must commit to a price for a period of time.

Note that if I allow for a fraction of the total number of consumers to become active in each period, then, as $\eta \to \infty$, the model becomes deterministic, and all of the rich dynamics developed in Section 4 fall through. The reason is that, unless there are aggregate preference shocks, the law of large numbers kicks in. In this sense, my assumption that only one out of η consumers is active each period may also be interpreted as a reduced form of a world where a fraction $1/\eta$ of all consumers became active each period and are subject to a common preference shock. In that case, if γ were the churn rate per *calendar* time period, we would have $\gamma = \Delta/\eta$, where Δ is the number of model periods per period of calendar time.

3.2. Equilibrium

I will focus on symmetric Markov equilibria, which are characterized by a pricing strategy $p(i)$, where i is the number of living consumers who have purchased from firm i. In the remainder of the section, I first derive the determinants of consumer demand. Next, I derive the firm value functions and the resulting pricing strategy. Putting together demand and pricing, I derive a master equation that determines the evolution of market shares. The section concludes with two preliminary results: one regarding equilibrium existence and uniqueness and another regarding the stationary distribution of market shares.

3.3. Consumer Demand

At state i, an active consumer chooses firm i if and only if

$$\zeta_i - p(i) > \zeta_j - p(j), \tag{2}$$

or simply,

$$\xi_i \equiv \zeta_i - \zeta_j > x(i),$$

where

$$x(i) \equiv p(i) - p(j). \tag{3}$$

Firm i's demand function is simply given by

$$q(i) = 1 - \Phi(x(i)). \tag{4}$$

Notice that

$$\frac{\partial q(i)}{\partial p(i)} = -\phi(x(i)). \tag{5}$$

3.4. Pricing

Suppose that firms' costs are zero. Firm i's value function is then given by

$$\begin{aligned} v(i) = {} & q(i)\,p(i) \\ & + \frac{i}{\eta}\big(q(i)(\theta\Delta(i) + \delta v(i)) \\ & \qquad + (1 - q(i))(\theta\Delta(i-1) + \delta v(i-1))\big) \\ & + \frac{j}{\eta}\big(q(i)(\theta\Delta(i+1) + \delta v(i+1)) \\ & \qquad + (1 - q(i))(\theta\Delta(i) + \delta v(i))\big), \end{aligned} \tag{6}$$

where $i = 0, \ldots, \eta$ and $j = \eta - i$.[10] The various terms in (6) correspond to various possibilities regarding consumer "death" and "birth." Suppose, for example, that the active consumer is a firm j consumer, something that happens with probability j/η. Suppose, moreover, that this consumer chooses firm i, which happens with probability $q(i)$. Then firm i receives sales revenue $q(i)p(i)$ (first row), current extra payoff $\theta\Delta(i+1)$, and continuation payoff $\delta v(i+1)$.[11]

Note that, with some abuse of notation, (6) corresponds both to firm i's Bellman equation and to the recursive system that determines the value function. As a Bellman equation, the $v(\cdot)$ on the right-hand side should be treated as $v_c(i)$—that is, continuation value. This is important when deriving first-order conditions, to the extent that the terms on the right-hand side should be treated as constant in the firm's optimization problem.

Define

$$w(i) \equiv \theta(\Delta(i+1) - \Delta(i)) + \delta(v(i+1) - v(i)). \tag{7}$$

Put into words, this denotes firm i's value from poaching a customer from firm j. This is divided into two different components: the immediate value in terms of market leadership, $\theta\Delta(i+1)$ if firm i makes the sale and minus $\theta\Delta(i)$ if it does not; and the discounted future value, $\delta v(i+1)$ if firm i makes the sale and minus $\delta v(i)$ if it does not.

Using (7), the first-order condition for maximizing the right-hand side of (6) with respect to $p(i)$ is given by

$$q(i) + \frac{\partial q(i)}{\partial p(i)} p(i) + \frac{i}{\eta}\frac{\partial q(i)}{\partial p(i)} w(i-1) + \frac{j}{\eta}\frac{\partial q(i)}{\partial p(i)} w(i) = 0,$$

or simply,

$$p(i) = \frac{1 - \Phi(x(i))}{\phi(x(i))} - \frac{i}{\eta} w(i-1) - \frac{j}{\eta} w(i), \tag{8}$$

where I substitute (4) for $q(i)$ and (5) for $\partial q(i)/\partial p(i)$.

If $\theta = 0$, then there are no number one effects: $v(i) = v(i+1)$, $w(i-1) = w(i) = 0$, and we have a standard static product differentiation model. Specifically, only the first term on the right-hand side of (8) matters, where $x(i) = p(i) - p(j)$. By contrast, if $\theta > 0$, then $w(i) \neq 0$, and firms lower their price to the extent of what they have to gain from making the next sale, which is given by $(i/\eta)w(i-1) + (j/\eta)w(i)$. From firm i's perspective, with probability i/η, the next sale is a battle for keeping one of its customers; that is, it is the difference between the continuation value of state i and the continuation value of state $i-1$. With probability j/η, the next sale is a battle for attracting a rival customer; that is, it is the difference between the continuation value of state $i+1$ and the continuation value of state i.

Plugging this back into the value function (6) yields

$$v(i) = \frac{(1-\Phi(x(i)))^2}{\phi(x(i))} + \frac{i}{\eta}(\theta\Delta(i-1) + \delta v(i-1)) + \frac{j}{\eta}(\theta\Delta(i) + \delta v(i)). \tag{9}$$

Under a static oligopoly, we would only have the first term on the right-hand side. The additional terms suggest that a firm's value corresponds to the value in case it *loses* the challenge for the next consumer: either losing the battle for keeping one of its consumers (a battle that takes place with probability i/η) or losing the battle for capturing one of the rival's consumers (a battle that takes place with probability j/η). This is the intuition underlying the Bertrand paradox (also known as the Bertrand trap; see Cabral and Villas-Boas 2005): to the extent that firms lower their price by the value of winning a sale, their expected value is the value corresponding to losing the sale (zero in the standard symmetric Bertrand model and the first term on the right-hand side if there is product differentiation). In other words, price competition implies rent dissipation; in the present case, the $w(i)$ rent.

System (9) can be solved sequentially:

$$v(i) = \left(1 - \frac{j}{\eta}\delta\right)^{-1} \left(\frac{(1-\Phi(x(i)))^2}{\phi(x(i))} + \frac{i}{\eta}(\theta\Delta(i-1) + \delta v(i-1)) + \frac{j}{\eta}\theta\Delta(i)\right). \tag{10}$$

Finally, I am also interested in distinguishing firm value (the function that firm decision makers maximize) from shareholder value (the firm's financial gain). The latter is given by

$$s(i) = q(i)p(i) + \frac{i}{\eta}(q(i)\delta s(i) + (1-q(i))\delta s(i-1)) + \frac{j}{\eta}(q(i)\delta s(i+1) + (1-q(i))\delta s(i)). \tag{11}$$

In other words, (11) corresponds to (6) with the difference that it excludes number one effects; that is, $\theta = 0$.

3.5. Market Shares

Recalling that $x(i) = p(i) - p(j)$ and subtracting (8) from the corresponding $p(j)$ equation, we get

$$p(i) - p(j) = \frac{1-\Phi(x(i))}{\phi(x(i))} - \frac{i}{\eta}w(i-1) - \frac{j}{\eta}w(i) - \frac{1-\Phi(x(j))}{\phi(x(j))} + \frac{j}{\eta}w(j-1) + \frac{i}{\eta}w(j), \tag{12}$$

or simply,

$$x(i) = \frac{1-2\Phi(x(i))}{\phi(x(i))} - \frac{i}{\eta}(w(i-1) - w(j)) - \frac{j}{\eta}(w(i) - w(j-1)), \tag{13}$$

where I use the fact that $1 - \Phi(x(j)) = \Phi(x(i))$.

Equation (13) is the "master equation" determining the evolution of market shares (in expected value). Recall that $q(i) = 1 - \Phi(x(i))$, so a higher $x(i)$ implies a lower probability that firm i makes the next sale. If $\theta = 0$, so that $w(i) = 0$ for all i, then we have a standard static product differentiation model: all terms on the right-hand side except the first one are zero, and as a result $x(i) = 0$, too: each firm makes a sale with the same probability.

More generally, what factors influence the value of $x(i)$? Essentially, the difference across firms is the value of winning the sale: as shown before, firms lower their prices to the extent of their incremental value of winning a sale; the firm that has the most to win will be the most aggressive, thus increasing the likelihood of a sale. The value of winning a sale may be decomposed into (a) the immediate benefit from an increment in market share, $\theta\Delta(i+1) - \theta\Delta(i)$ or $\theta\Delta(i) - \theta\Delta(i-1)$ as the case may be; and (b) the discounted future value from market share, $v(i+1) - v(i)$ or $v(i) - v(i-1)$, as the case may be.

3.6. Equilibrium

Equations (10) and (13) define a Markov equilibrium, where I note that $w(i)$ is given by (7). Given the values of $v(i)$ and $x(i)$, prices $p(i)$ and sales probabilities $q(i)$ are given by (8) and (4), respectively. Many of the results in the next sections pertain to the limit case when $\delta \to 0$. These results are based on the following existence and uniqueness result.

Lemma 1. *There exists a unique equilibrium in the neighborhood of $\delta = 0$. Moreover, equilibrium values are continuous in δ.*

3.7. Stationary Distribution of Market Shares

Given the assumption that $\Phi(\cdot)$ has full support (part (iii) of Assumption 1), $q(i) \in (0,1), \forall\, i$; that is, there are no corner solutions in the pricing stage. It follows that the Markov process of market shares is ergodic, and I can compute the stationary distribution over states. This is given by the (transposed) vector m that solves $mM = m$. Since the process in question is a "birth-and-death" process, whereby the state only moves to adjacent states, I can directly compute the stationary distribution of market shares.

Lemma 2. *The stationary distribution $m(i)$ is recursively determined by*

$$m(i) = m(0)\prod_{k=1}^{i}\frac{q(i-1)}{1-q(i)}\cdot\frac{\eta-i+1}{i},$$

where

$$m(0) = \left(1+\sum_{i=1}^{\eta}\prod_{k=1}^{i}\frac{q(i-1)}{1-q(i)}\cdot\frac{\eta-i+1}{i}\right)^{-1}.$$

Lemmas 1 and 2 allow for a partial analytical characterization of equilibrium. I will develop two types of analytical results: one corresponds to taking limits as $\delta\to 0$; the second, to taking derivatives with respect to δ at $\delta = 0$ (that is, linearizing the model). I complement these analytical results with numerical simulations for higher values of δ. These numerical simulations confirm the analytical results for small δ but also uncover additional features not present in the small δ case.

4. A Theory of Price Wars

I cannot find a general analytical closed-form solution for the model's equilibrium. However, I can characterize the equilibrium when $\delta = 0$, and, by Lemma 1, in the neighborhood of $\delta = 0$, the equilibrium values take on values close to the limit case $\delta = 0$. In the following results, I assume for simplicity that η is even, and I denote the symmetric state by $i^* \equiv \eta/2$.

Proposition 3. *There exists a unique equilibrium in the neighborhood of $\delta = 0$. Moreover,*

$$\lim_{\delta\to 0} p(i) = \begin{cases} \dfrac{1}{2\phi(0)} - \theta & \text{if } i = i^*, \\ \dfrac{1}{2\phi(0)} - \dfrac{\eta+1}{\eta}\theta & \text{if } i = i^*\pm 1, \\ \dfrac{1}{2\phi(0)} & \text{otherwise,} \end{cases}$$

$$\lim_{\delta\to 0} q(i) = \frac{1}{2}$$

$$\lim_{\delta\to 0} v(i) = \begin{cases} \dfrac{1}{4\phi(0)} - \theta & \text{if } i \le i^*-1, \\ \dfrac{1}{4\phi(0)} - \dfrac{i^*-1}{\eta}\theta & \text{if } i = i^*, \\ \dfrac{1}{4\phi(0)} + \dfrac{i^*-1}{\eta}\theta & \text{if } i = i^*+1, \\ \dfrac{1}{4\phi(0)} + \theta & \text{if } i \ge i^*+2, \end{cases}$$

$$\lim_{\delta\to 0} m(i) = \frac{\eta!}{i!\,(\eta-i)!\,2^{\eta}}.$$

The limiting stationary distribution is maximal at i^*.

Put into words, when firm market shares are close to each other, firms engage in a price war for market leadership, whereby both firms decrease price by up to θ from the static Hotelling price level $1/(2\phi(0))$. This is similar to the idea underlying the Bertrand paradox: the potential gain from being a market leader is competed away through pricing. Specifically, I define the "price war region" of the state space as the set $\{i^*-1, i^*, i^*+1\}$. Proposition 3 then states that, in the limit as $\delta\to 0$, prices are set lower than $1/(2\phi(0))$ (price war) when $i \in \{i^*-1, i^*, i^*+1\}$ and are equal to $1/(2\phi(0))$ (peace) when $i \notin \{i^*-1, i^*, i^*+1\}$.

Note that, in the limit as $\delta\to 0$, $p(i) = p(j)$. As a result, the probability of making a sale is uniform at $\frac{1}{2}$. This implies that market share dynamics follow a straightforward reversion to the mean process: smaller firms increase their market share on average, whereas larger firms decrease their market share on average. This is particularly bad for profits because it implies a constant tendency to engage in a price war.

The dashed lines in Figure 3 illustrate this situation. (In this and in the remaining figures in the paper, I assume $\eta = 100$, so that i is both the state and firm i's market share.[12] I also assume that ξ is distributed according to a standardized normal.[13]) The top left panel depicts the equilibrium price function, whereas the top right panel shows the stationary distribution of market shares. (Note that, since the equilibrium is symmetric, $p(i)$ and $m(i)$ are only a function of the state, not of the firm's identity.) The bottom panels show the value functions for firm managers (left) and shareholders (right).

Beginning with the price mapping, we see that prices are set at a constant level (the static equilibrium level) when the state is outside the price war region. Inside the price war region, firm prices drop by up to θ, which is the change in firm value from moving up one unit in the state space. Since the price mapping is symmetric about i^*, each firm's sale probability is flat at $\frac{1}{2}$. It follows that the stationary distribution of market shares is a simple multinomial centered around i^* (that is, around 50% market share).

The bottom right panel shows that shareholder value drops sharply when i is near i^*—that is, in the price war region. This follows form the fact that prices are lower near the symmetric state and also the fact that shareholders do not receive any benefit from being number one. In other words, since shareholders do not care for market leadership per se, number one effects are only bad news: they lead to price wars, which in turn destroy shareholder value.

With respect to firm value, the bottom left panel indicates that, in the limit as $\delta\to 0$, $v(i)$ is increasing in i. In particular, if $i > i^*$, then firm i receives utility θ in addition to expected revenues. This benefit from leadership is balanced out by the negative utility suffered by the laggard.

Figure 3. Equilibrium When $\theta = 1$ and $\delta = 0$ (Dashed Lines) and $\delta = \frac{3}{4}$ (Solid Lines)

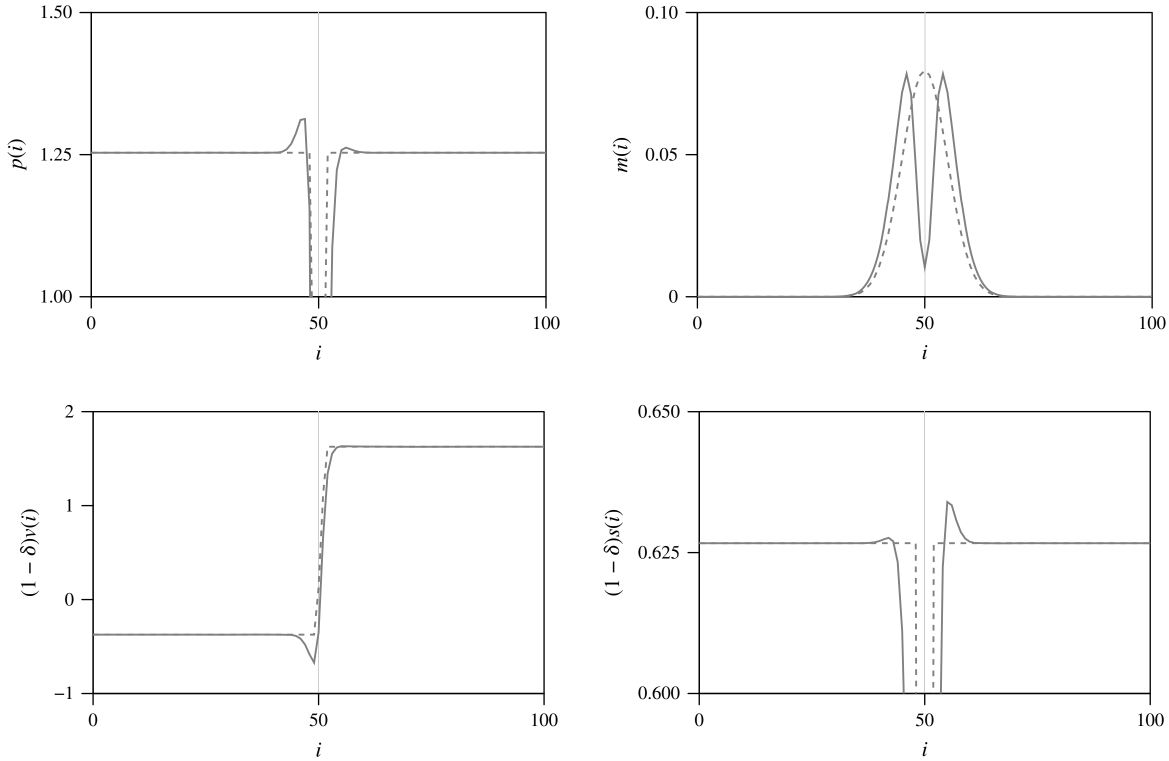

Finally, although not obvious from Figure 3, industry joint value, $v(i) + v(j)$, at states near i^* is actually lower when $\theta > 0$ than when $\theta = 0$. This follows from Proposition 3, as the next result attests.

Corollary 1. *In the limit as $\delta \to 0$, joint industry value $v(i) + v(j)$ is strictly decreasing in θ if $i \in \{i^* - 1, i^*, i^* + 1\}$ and constant otherwise.*

This is an important point, one that warrants further elaboration. The idea is akin to the Bertrand paradox. In a first-price auction where the payoff from winning is given by $+\pi$ and the payoff from losing is given by $-\pi$, the greater the value of π, the lower the equilibrium value by both bidders: the winner gets π from winning minus 2π, the equilibrium bid, whereas the loser gets $-\pi$. In the present context, an increase in θ increases the payoff from winning a sale and decreases the payoff of losing it. Although the total payoff from market leadership is constant (specifically, $\theta\Delta(i) + \theta\Delta(j) = 0$), the equilibrium value received by each firm is decreasing in θ: in equilibrium, each firm fares as well as when it loses the sale.[14]

An additional implication of Proposition 3, similar to Corollary 1, is that industry joint value is higher at asymmetric states than at symmetric states, so that, at symmetric or near-symmetric states, the leader has more to gain from increasing its lead than the laggard has to lose from falling farther behind.

Corollary 2. *At $\delta = 0$, $v(i) + v(j)$ is strictly increasing in $|i - j|$ if $|i - j| \le 2$. Moreover,*

$$v(i^* + 1) - v(i^*) > v(i^*) - v(i^* - 1),$$
$$v(i^* + 2) - v(i^* + 1) > v(i^* - 1) - v(i^* - 2).$$

Put into words, the second part of Corollary 2 states that, at state $i^* + 1$, what the leader has to lose by moving down one step is more than what the laggard has to gain by moving up one step, and what the leader has to gain by moving up one step is more than what the laggard has to lose by moving down one step. This is the dynamic equivalent of the "joint-profit effect" of Gilbert and Newbery (1982). In their paper, the effect results from the convexity of the profit function; in my paper, it results from the convexity of the value function.[15]

Notice that the two parts of Corollary 2 are equivalent: both stem from the value function being "convex." In fact, $v(i^* + 1) - v(i^*) > v(i^*) - v(i^* - 1)$ is equivalent to $v(i^* + 1) + v(i^* - 1) > v(i^*) + v(i^*)$, and $v(i^* + 2) - v(i^* + 1) > v(i^* - 1) - v(i^* - 2)$ is equivalent to $v(i^* + 2) + v(i^* - 2) > v(i^* - 1) + v(i^* + 1)$. In other words, if the value function is convex, then its "slope" is greater for the leader than for the laggard. Similarly, by a discrete analog of Jensen's inequality, joint profit increases when the state becomes more asymmetric.

Figure 4. Market Leadership Benefit (Left) and Value Function (Right) at $\delta = 0$ for $\theta = 0$ (Light Lines) and $\theta > 0$ (Dark Lines), Where i^* Is the Symmetric State

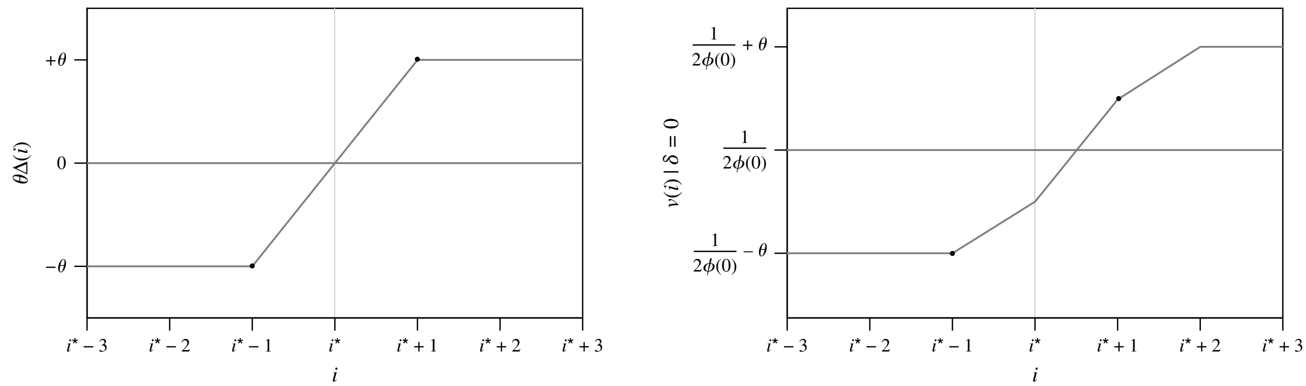

Figure 4 illustrates Corollaries 1 and 2. The left-hand panel depicts the market leadership mapping. As can be seen, the mapping is symmetric about $(i^*, 0)$, so that the sum $\Delta(i) + \Delta(j)$ is equal to zero at every state. The same is not true, however, regarding value functions, as can be seen from the right-hand panel. For example, at state i^*, each firm's payoff when $\theta > 0$ is lower than it would be if $\theta = 0$ (Corollary 1). Moreover, $v(i)$ is "convex." At $i = i^*$, this corresponds to the fact that $v(i^*) - v(i^* - 1) < v(i^* + 1) - v(i^*)$; at $i = i^* - 1$, it corresponds to the additional fact that $v(i^* + 2) - v(i^* + 1) > v(i^* - 1) - v(i^* - 2)$ (Corollary 2).

Corollaries 1 and 2 have important implications for system dynamics in the neighborhood of $\delta = 0$, as I will show next.

4.1. Positive, Small Values of δ

Proposition 3 considers the limit when $\delta \to 0$. From Lemma 1, I know that the system's behavior is continuous around $\delta = 0$; that is, the limit $\delta \to 0$ is a good indication of what happens for low values of δ. Additional information can be obtained by linearizing the system around $\delta = 0$ and thus determining the direction in which equilibrium values change as δ moves away from zero.[16]

Recall that in the limit, as $\delta \to 0$, $p(i) = p(j)$ and $q(i) = q(j)$. My next result shows that, in the near symmetric states $i^* - 1$ and $i^* + 1$, the market leader sets a low price and sells with higher probability. Moreover, the laggard is strictly worse off by increasing its market share.

Proposition 4. *There exists a $\delta' > 0$ such that if $0 < \delta \le \delta'$, then $\theta > 0$ implies*

$$p(i^* + 1) < p(i^* - 1),$$
$$q(i^* + 1) > q(i^* - 1),$$
$$v(i^* - 1) < v(i^* - 2).$$

(Notice that, given the demand curve (4), the first two inequalities are equivalent.)

As mentioned earlier—and as shown by (8)—firm i's first-order condition includes the value of winning a sale, either the value of keeping an existing customer, $w(i - 1)$, or the value of poaching a rival's customer, $w(i)$. When $\delta = 0$, the value of winning a customer is based on the mapping $\theta\Delta(i)$, as illustrated in the left panel of Figure 4. Consider, for example, a firm with $i^* - 1$ customers. If this firm gains one customer, its payoff increases by θ, whereas its rival, by moving from $i^* + 1$ to i^*, decreases by θ. Conversely, if the firm at $i^* - 1$ loses one customer, then its leadership payoff remains the same, whereas its rival, by moving from $i^* + 1$ to $i^* + 2$, also sees its payoff remain constant. In sum, for $\delta = 0$, what the leader has to gain (respectively, lose) from making a sale is the same as the laggard has to lose (respectively, gain). As a result, both firms apply the same "subsidy" to their price level, and $q(i) = 1/2$ for all i, as stated in Proposition 3.

Consider now the case when δ is positive but infinitesimal. Given that the active consumer is a j consumer, firm i's value from winning a sale is given by $w(i) \equiv \theta\Delta(i+1) - \theta\Delta(i) + \delta v(i+1) - \delta v(i)$. At $\delta = 0$, as we have seen, the values of $w(i)$ for leader and laggard balance out exactly. As we increase δ infinitesimally, the value of $w(i)$ increases at the rate $v(i+1) - v(i)$, where the value functions are evaluated at $\delta = 0$. Proposition 4 exploits the fact that, while the values of $\theta\Delta(i)$ add up to a constant, so that leader and laggard have the same to win or lose, the same is not true for $v(i+1) - v(i)$, as Corollary 2 states.

Specifically, consider the near-symmetric state $(i^* - 1, i^* + 1)$. As Corollary 2 shows, the lagging firm has less to gain from moving up the value function than the leader has to lose from losing to the laggard. Moreover, the laggard has less to lose from falling farther behind than the leader has to gain from moving further ahead. In other words, the value function is "convex." Given the intuition underlying the first-order conditions (8), this implies that the leader prices more aggressively, which results in it making a sale with a higher probability than the laggard.

4.2. Higher Values of δ

For high values of δ, I cannot find a closed-form analytical solution or linear expansion approximation. However, I can solve the model numerically. The dark lines in Figure 3 show the model's solution for $\delta = \frac{3}{4}$.[17] The solution looks qualitatively similar to $\delta = 0$ in various respects—namely, in the property that prices drop when firms' market shares are close to each other. However, upon closer inspection, important differences become apparent as well. First, as suggested by Proposition 4, when $\delta > 0$, the pricing function is no longer symmetric around i^*. In particular, just outside the price war region, the large firm's price is lower, whereas the smaller firm's price is higher. This implies that the probability of a sale by a leader increases when the leader's market share drops to close to i^*. As the top right panel in Figure 3 shows, this (perhaps) also implies that the stationary distribution of market shares is bimodal.[18] That is, most of the time, the system lies at an asymmetric state, where one firm is larger and the other firm smaller.

4.3. Price and Market Share Dynamics

Proposition 3 shows that firms engage in price wars when the state space is close to the symmetric state, whereas Proposition 4 suggests that market shares tend to remain stable around asymmetric outcomes. I now examine the implications of these properties. Figure 5 illustrates the dynamics of price and market shares by showing the results of a model simulation when $\delta = \frac{3}{4}$ and $\theta = 1$ (the parameter values corresponding to Figure 3).[19] In the top panel, the solid line represents firm i's price and the dashed line firm j's price. A horizontal line marks the equilibrium price level when $\theta = 0$.[20] The bottom panel depicts firm a's market share.

Figure 5. Price and Market Share Dynamics ($\delta = \frac{3}{4}$, $\theta = 1$)

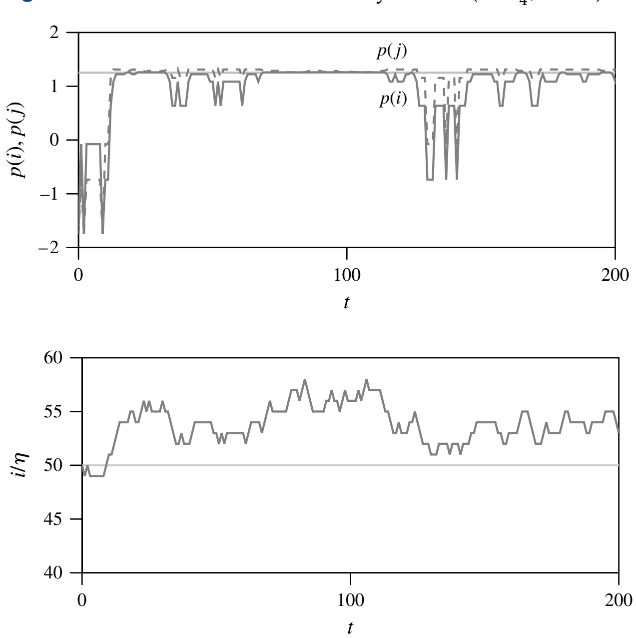

According to my model, a price war is a period of significantly lower prices that takes place when the firms' market shares are close enough (if $\delta = 0$, when $|i - j| \le 2$). Figure 5 shows some of these price wars. The deepest price war takes place right from the start, which is not surprising since I started the simulation at $i = i^*$. Mostly out of sheer luck, firm a makes most of the sales during this period; that is, firm a "wins" the price war. As a result, firm a's market share increases to over 50%, as can be seen from the bottom panel. Now that there is a "clear" market share leader, prices increase to a high level, at or slightly above the static equilibrium price level.

Whenever firm a's market shares decrease to close to 50%, a price war begins again. Notice that firm a, the market share leader, is also the price leader going into a price war. This is consistent with the idea that price wars are a defense against threats to market share leadership. Finally, just as price wars are triggered by a drop in the leader's market share, so the end of a price war is determined by an increase in the leader's market share.

In other words, firm i (the market leader in the particular simulation on Figure 3) has more to lose from lowering its market share than firm j has to gain from increasing its market share. This results in more aggressive pricing by the market share leader. Most of the time (and in the first 200 periods of the simulation I considered), the market share leader remains so; that is, my model implies persistence in leadership. However, with probability 1, market leadership changes in finite time.

My paper is by no means the first paper to feature symmetric equilibria with asymmetric outcomes and price wars near symmetry states. Besanko et al. (2010), for example, show that learning curves lead to "trenchy" price equilibria whereby prices drop when competitors' market shares are close to each other.[21] My model differs from the previous literature in that it does not feature increasing returns to scale. In fact, by construction, $\theta\Delta(i) + \theta\Delta(j)$ is equal to zero. Specifically, if prices were set at a constant level, then my model would imply that industry joint value $v(i) + v(j)$ is constant across states, whereas Besanko et al. (2010) or Cabral (2011), for example, would imply that $v(i) + v(j)$ is increasing in $|i - j|$.

Moreover, while the stationary distribution of market shares is multimodal, it still places significant mass on symmetric or near-symmetric states. (If $\delta = 0$, the stationary distribution of market shares is a binomial centered around 50%.) As a result, price wars are relatively frequent, whereas in models with increasing returns to scale they are rare: once one of the firms becomes dominant, it takes a long time for tipping to take place. This is an important distinction, one that warrants further discussion. In dynamic market share models, there is a natural reversion-to-the-mean force: consumer death (a firm with 100% of the market can only decrease its market share). Against this force pushing toward market share balance, there may be various forces pushing the system away from symmetry. Increasing returns (learning curves or network effects) represent one such force. In my model, the force that pushes away from symmetry is price wars. However, to the extent that price wars only kick in at states close to symmetry, the effect of price wars is only felt at states close to symmetry. As a result, we have a stationary distribution where much of the weight is on states close to the threshold of the price war region. This results in frequent movements inside the price war region. In other words, unlike models with increasing returns, price wars are observed cyclically along the equilibrium path.

In Section 1, I mentioned the server market as an example where market share leadership is considered important. For several years, IBM and HP fought for the number one position (as measured by market share of dollar sales). In 2014, IBM sold a considerable portion of its server assets to Lenovo. This essentially ended a period of fight for market share leadership. My model would predict an increase in HP's price to ensue, basically, the result of a shift to a "peace" phase.

The data, depicted in Figure 6, are broadly consistent with the model's predictions—though admittedly, it is also consistent with other models. In both cases, we observe a price that is relatively constant until 2013 and gradually increases as the market share gap between HP and IBM increases—that is, as HP becomes a clear market share leader. (Considering that there may be many spurious factors affecting price, I also plot the price index: HP's price divided by industry average price, measured in the right scale of the upper panel. The pattern is similar.)

Figure 6. Server Market: Prices and Market Shares

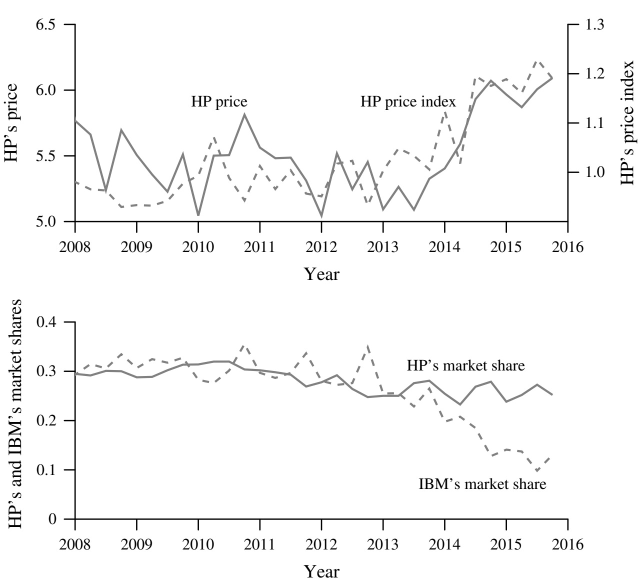

5. Corporate Culture and Shareholder Value

The bottom right panel in Figure 3 shows an important difference between the equilibrium with $\delta = 0$ and the equilibrium with $\delta = \frac{3}{4}$. In the former case, number one effects are unambiguously detrimental to shareholder value. This is fairly intuitive: number one effects lead firms (symmetrically) to lower prices when in state $i \in \{i^* - 1, i^*, i^* + 1\}$. Lower prices lower shareholder value; moreover, number one effects accrue no shareholder utility. All in all, wanting to be number one is bad for shareholders.

However, if δ is sufficiently high (e.g., $\delta = \frac{3}{4}$), then there are states when shareholder value is greater with $\theta > 0$ than with $\theta = 0$. To understand this, it helps to notice that, as shown in Proposition 4, $v(i)$ is *decreasing* for values of i lower than, but close to, i^*. In other words, a laggard becomes worse off as its market share approaches the leader's. The reason is that the increase in market share induces very aggressive pricing behavior by the leader, which in turn reduces the laggard's value: the laggard receives no benefit from market leadership but pays the cost of a leader eager to defend its benefit from market leadership.

As seen earlier, the first-order condition for optimal pricing includes a "subsidy" in the amount of the expected continuation gain from making a sale, either the value of keeping an own contested consumer, $w(i-1)$, or the value of poaching a consumer from the rival firm, $w(i)$. If the value function is decreasing (and the payoff from market leadership does not change), then a declining $v(i)$ implies a negative $w(i)$, which in turn implies that the price "subsidy" becomes a "tax." In other words, the "threat" of entering a price war with the leader softens the laggard. This effect may be so strong as to increase the leader's shareholder value (in the states where the laggard softens up). In other words, even though shareholders do not care about market leadership per se, shareholder value may increase when managers care for market leadership. (I should stress the word "may" as this is a possibility result, not a general analytical result.)

Although the Markov equilibrium I consider differs greatly from a repeated game (where, by definition, there is no state space such as market share), there is an interesting similarity between the above effect and the so-called topsy-turvy principle in collusive repeated game equilibria (Shapiro 1989). Consider a repeated game where each period 1 consumer buys one unit from one of two firms. Consider a class of grim-strategy equilibria whereby price is $\bar{p}$ along the equilibrium path and strategies are such that, if any player sets $p \neq \bar{p}$, then play reverts to $\underline{p}$ forever (for simplicity, I ignore issues of subgame perfection or renegotiation proofness). Suppose that buyers choose the firm with the lowest price and that willingness to pay is sufficiently high that it is not binding. Then, for a given discount factor, the lower the value of $\underline{p}$, the higher the maximum $\bar{p}$ that is sustainable as a Nash equilibrium of the repeated game.

In my model, equilibrium play moves between states in the "price war" region and outside the price war region. If we think of price war states as similar to punishment periods in the repeated game, then the corresponding topsy-turvy principle is that the deeper the price cuts in the price war region, the higher the prices once outside of the price war region. The reason is that deeper price cuts imply a bigger drop in $v(i)$ for the laggard when close to i^* and, consequently, a higher "tax" on price.

5.1. Asymmetric Number One Effects

So far in this section, I have considered the case when both firms benefit from being market leaders; that is, both adopt a number culture. To consider the possibility of a firm creating its own culture—in the sense of how much it values market leadership—it is helpful to consider the case when the value of θ is firm specific (as was the case in Section 2).

Figure 7 shows equilibrium prices in the case when $\theta_a > 0$, whereas $\theta_b = 0$. Since the game is now asymmetric, I must distinguish between firm a's pricing strategy, $p_a(i)$, and firm b's pricing strategy, $p_b(j)$. To best view price levels at a given state, I plot firm b's values against a left-pointing axis measuring $j = \eta - i$. For example, suppose the current state is $i = 53$; that is, firm i's market share is given by 53%. Then I plot at the same horizontal position firm b's price $p_b(47)$, where $47 = \eta - i$.

Figure 7. Asymmetric Game: $\theta_a = 1, \theta_b = 0$, and $\delta = \frac{3}{4}$ (Horizontal Line: Symmetric $\theta_a = \theta_b = 0$ Case)

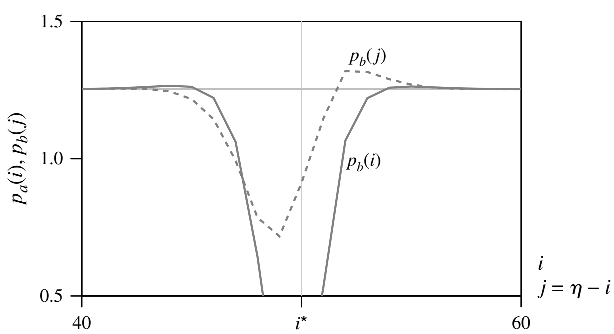

The qualitative features of firm a's pricing function are similar to the symmetric case: when firm a's market share is close to 50%, its prices are lower as its value function is very steep. Notice that, while firm b gains nothing from being the market leader, it too lowers its price when market shares are similar. This results from strategic complementarity in pricing as well as from the fact that $v_b(i)$, too, is steeper when $i \approx i^*$. In particular, when firm b is a market leader with a short lead, it knows that a decrease in market share implies a significant decrease in value: it implies entering a value-destroying price war with a rival who cares about market share leadership. In fact, as Figure 7 suggests, firm b leads firm a in cutting prices when firm b's market share drops toward i^* (which, in Figure 7, corresponds to an increase in i toward i^*).

Figure 7 also shows that when firm a has a moderate lead, firm b prices above the static equilibrium level. This is the "softening" effect first mentioned in Section 2: fearing that firm a will become very aggressive should it (firm a) lose market share, firm b sets a higher price. Strategic complementarity implies that firm a also sets a higher price. We conclude that, in these states, the number one corporate culture benefits firm a in terms of both a higher margin and a higher probability of making a sale—and thus further increasing market share.

One important difference between the symmetric and asymmetric cases is that the stationary distribution of market shares is no longer symmetric. In fact, consistent with the fact that firm a places extra value on being a market leader—whereas firm b does not—most of the time firm a is effectively the market leader. (The stationary distribution of markets shares may be unimodal or bimodal, but the higher mode is always to the right of 50%.)

5.2. The Corporate Culture Metagame

Considering various possible asymmetric games, we can now examine the nature of a corporate culture metagame such as the one presented in Section 2. There is one important difference with respect to the previous two-stage game: Before, I essentially assumed an initial state (namely, 50-50—state i^*). Now, I consider the impact of corporate culture choices at various possible states.

One possible conceptual experiment is as follows: Suppose that a Martian shareholder randomly appears on planet Earth and acquires firm a. Would such a shareholder choose a number one corporate culture? Figure 8 plots the difference, in terms of discounted shareholder value, between a game with $\theta_a = 1$ and a game with $\theta_a = 0$, given a value of $\theta_b \in \{0, 1\}$ and given a current state i. If the new shareholder could make a one-time irreversible corporate culture decision, and if firm b corporate culture were fixed, then we conclude

Figure 8. Corporate Culture Best Response: Change in $s_a(i)$ When θ_a Switches from 0 to 1 Given that $\theta_b = 0$ (Solid Line) or $\theta_b = 1$ (Dotted Line)

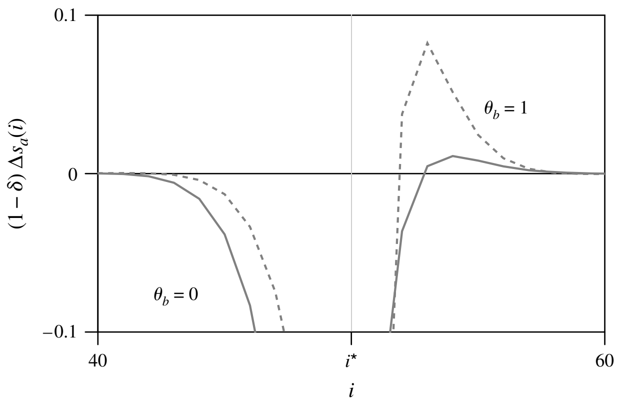

that the shareholder would opt for a number one culture if firm a is currently the market share leader.

I tried multiple parameter values for the above game and metagame, and this qualitative feature seems robust. Intuitively, if firm a is a market share leader, then a number one culture best corresponds to the intuition provided by the two-stage game in Section 2: shareholders get the benefits from a number one culture (softening the rival, who fears entering into a price war) without incurring the costs of number one leadership (actual price wars).

Interestingly, this is true even if the rival already has a number one culture (as shown by the dashed line in Figure 8). But this then raises an interesting problem: If firm a's market share were to drop below 50%, then firm b, if faced with the choice of a corporate culture, would be faced with a similar trade-off—and would choose a number one culture.

However, the previous analysis was predicated on the assumption that the rival's corporate culture choice is given. An alternative approach to the corporate culture metagame would be to assume that shareholders can make an irreversible choice at any time t (that is, at any state i or j). A forward-looking shareholder would then consider the possibility that, at a future date, the rival firm would also adopt a number one culture. The bottom right panel of Figure 3 shows that, even if both players choose a number one culture, shareholder value increases with respect to the $\theta_a = \theta_b = 0$ case. This suggests that even a forward-looking shareholder would opt for a number one culture if it ever found itself in a leadership position.

An alternative conceptual experiment is to consider a metagame whereby, before the pricing game is played, shareholders must choose a corporate culture; then each firm is thrown into an initial state according to the stationary distribution of market shares.[22] While I do not have any analytical results, numerical simulations suggest that $\theta = 0$ is a dominant strategy in this metagame.

6. Robustness and Extensions

There are several directions along which the basic model can be extended. In the online appendix, I detail the following.

6.1. Sensitivity Analysis

I perform comparative dynamics with respect to various model parameters: θ, the intensity of the number one effect; δ, the discount factor; and μ and σ, the mean and standard deviation of preference shocks. I find that the results vary rather "smoothly" with respect to these parameters.

Moreover, the qualitative results described in the previous sections remain valid for a vast set of possible parameter values. Naturally, the quantitative results change. In particular, the length and depth and frequency of price wars depend on parameter values in expected ways.[23]

6.2. Demand-Driven Number One Effects

In the previous sections I considered the case when managers derive extra utility from being market leaders; but an equally compelling observation is that consumers enjoy purchasing from a market leader. A natural way to extend the model is to assume that consumer net utility from buying firm i's product is given by

$$\zeta_i + \lambda\Delta(i) - p(i).$$

It can be shown that, under consumer number one effects, the market leader is able to price higher than the laggard and sell with a higher probability than the laggard. Moreover, just like the manager's utility case, if θ and δ are high enough, then equilibrium pricing results in a bimodal stationary distribution of market shares.

6.3. Alternative Number One Effect Mappings

How much do the results depend on the fact that market leadership is such a "discontinuous" mapping; that is, it switches from $-\theta$ when $i < j$ to $+\theta$ when $i > j$? To address this question, I consider alternative mappings where market share leadership utility ramps up from $-\theta$ to $+\theta$ gradually as market share increases. The qualitative results are very similar, with one quantitative difference: the more gradual the number one utility mapping, the longer and shallower prices wars are.

6.4. Multiple Simultaneous Active Agents

For analytical tractability, I assume that one consumer makes a decision at a time. I conjecture that the model's main qualitative features remain valid if we were to allow for multiple active consumers each period. In the appendix, I discuss this extension in greater detail.

6.5. N Firms

So far, all of my analysis have centered on the duopoly case. Some of the examples I motivated the paper with are indeed duopolies (e.g., Boeing, Airbus), but many feature more than two firms (e.g., IBM, HP). Compared with the $N = 2$ case, the general N case creates an extra layer of computational burden: even for given value functions, I need to solve the first-order conditions numerically. Moreover, the stochastic process is no longer a birth-and-death process, so the stationary distribution of market shares cannot be computed analytically. Other than that, the model is similar to what I developed in Section 3, although some new wrinkles are created by the presence of more than two firms (e.g., is the leadership benefit—or lack thereof—different for a number two and a number three firm?).

7. Conclusion

Many firms seem to place a disproportionate weight on the goal of being market share leaders. In this paper, I develop a positive analysis of market competition when firms have such preferences. Whereas the standard explanation for price wars is associated with the idea of collusive equilibria, I present a theory of price wars that is entirely based on battles for market share.

Various real-world examples seem to fit my theory of price wars better than the collusion theory. Consider, for example, the events following Rupert Murdoch's acquisition of the *London Times*. The *Times*, which started from a low i state (daily circulation of 360K), initially slashed its prices (from 48p to 30p, then to 20p). These low prices were followed by some competitors (e.g., *Independent, Daily Telegraph*) but not by all (e.g., not by the *Guardian*). Even the newspapers that cut prices did so to a less extent than the *Times* did. As a result, the *Times*' market share gradually increased, reaching a daily circulation of 860K after three years of price war. Eventually, prices were brought back to the initial levels. These events are roughly consistent with my model and a shock to θ_a corresponding to the *Times* acquisition by Murdoch.

A related question suggested by my model is whether it makes business sense for firms to aim for market leadership. Consider, for example, the case of General Motors, who, in 2007, lost the position of global market share leader and later recovered it—although for one year only. According to CNN,

> GM had held onto market share and its No. 1 rank by cutting prices on cars to the point where they were unprofitable. Bob Lutz, former vice chairman of GM, said worrying about its market share rank did the company more harm than good. "There is absolutely nothing to be gained by being the world's biggest," he said. "I tried to tell them to say, no, it's not our objective to be No. 1. But they just couldn't do it." (Isidore 2012)

In other words, Lutz suggested that rank is irrelevant as far as firm value goes. By contrast, my analysis shows that, even if rank is not directly relevant in terms of shareholder value, it may be so by the behavior that it induces; it may, in fact, increase shareholder value.

I am by no means the first to suggest that committing to a course of action that departs from profit maximization may increase a firm's payoff. In these situations, a crucial issue is whether players have the power to commit to an ex post suboptimal course of action. For example, complex contracts may be difficult to observe or verify—and are subject to renegotiation. In this sense, the goal of being number one seems compelling because it is simple—and simplicity is an important condition for credibility.

Acknowledgments

The author is grateful to Glenn Ellison, Michael Katz, Dmitri Kuksov, Francisco Ruiz-Aliseda, Paulo Somaini, Brian Wu, and participants at various seminars and conferences for helpful comments. The usual disclaimer applies.

Appendix

Proof of Proposition 1. Consider first the case when $\theta = 0$. Since $\sigma_L > \mu$, the equilibrium in the aftermarket is for firms to set $q = \mu$ and serve the customers attached to them (that is, it would require a negative price to poach customers from the rival). Given that, competition for the new customer (foremarket) leads to equilibrium prices $p = -\mu$.

Now consider the case when $\theta_a = \theta$ is very high, whereas $\theta_b = 0$. If firm a serves the new consumer in the foremarket, then the continuation subgame is isomorphic to the case when $\theta = 0$: firms set $q = \mu$ and serve the customers in their installed base. Suppose, however, that firm b serves the new customer in the foremarket. Then, by increasing the number of customers by two, firm a's manager increases its utility by θ. If the value of θ is high enough, then firm a's manager is better off by setting $q_a = \mu - \sigma_H$, thus attracting all of firm b's customers. In fact, if $q_a > \mu - \sigma_H$, then there is a strictly positive probability that no firm b customer switches from firm b, thus forgoing the extra θ benefit.

Now consider the equilibrium price level $p_a = p_b = \nu$ and suppose that the new consumer picks firm a.[24] In equilibrium, firm b receives a payoff of $n\mu$. By slightly undercutting its rival, firm b gets a payoff of ν in the foremarket and 0 in the aftermarket. The condition that firm b has no profitable deviation is the equivalent to (1). Regarding firm a, $p_a = \nu$ gives it the highest possible payoff in the foremarket. (Any higher price, and the new customer would not make a purchase.) Since in equilibrium firm a is the market leader (n versus $n - 2$ customers), poaching additional customers from firm b would cost firm a a loss ($\sigma_L < \mu$) and would not increase leadership benefit. □

Proof of Proposition 2. Proposition 1 implies that there exists a $\bar{\theta}$ such that, if $\theta \geq \bar{\theta}$, then the best response to $\theta_j = 0$ is $\theta_i = \theta$. What if $\theta_j = \theta$? If firm i chooses $\theta_i = 0$, then, by Proposition 1, its subgame equilibrium payoff is given by $n\,\mu$. If instead firm i chooses $\theta_i = \theta$, then in the ensuing game both firms' payoffs are zero: one of the firms prices sufficiently low such that it attracts all consumers with probability 1 but does so by pricing so low that its net payoff is 0; the other firm prices σ_H higher, loses all of its consumers, and receives an equilibrium payoff of zero as well. Together with the fact that choosing $\theta_i = 0$ implies a positive payoff, we conclude that the best response to $\theta_j = \theta$ is $\theta_i = 0$. □

Proof of Lemma 1. As mentioned in the text, equilibrium is (essentially) defined by the sets of equations (10) and (13), which determine the values of $v(i)$ and $x(i)$, respectively. Given $x(i)$ and $v(i)$, the values of $p(i)$ and $q(i)$ are determined uniquely. Let

$$\mathbf{z} \equiv (x(0), \ldots, x(\eta), v(0), \ldots, v(\eta)).$$

We thus have a system of $2(\eta+1)$ equations and $2(\eta+1)$ unknowns. Represent this system as $f_i(\mathbf{z};\delta) = 0$, where $i = 0, \ldots, 2\eta$.

Setting $\delta = 0$, (13) may be rewritten as

$$x(i) + \frac{2\Phi(x(i)) - 1}{\phi(x(i))} = \Gamma(i), \tag{A.1}$$

where $\Gamma(i)$ is bounded and exogenously given. Part (iv) of Assumption 1 implies that the left-hand side is strictly increasing in $x(i)$, ranging continuously from $-\infty$ to $+\infty$ as $x(i)$ varies from $-\infty$ to $+\infty$. It follows that there exists a unique value $x(i)$. Uniqueness of $x(i)$ in turn implies uniqueness of $v(i)$, $p(i)$, and $q(i)$.

Continuity implies that, for each δ in the neighborhood of $\delta = 0$, I can find an $\epsilon \in (0,\infty)$ such that equilibrium $x(i)$ and $v(i)$ must be in $[-\epsilon, \epsilon]$. I thus henceforth restrict to this compact set of $x(i)$ and $v(i)$ values.

Setting $\delta = 0$, (9) implies

$$v(i) - \frac{(1-\Phi(x(i)))^2}{\phi(x(i))} = \Omega(i). \tag{A.2}$$

Taking into account that $\delta = 0$ implies $x(i) = \frac{1}{2}$, (A.1) and (A.2) imply that the matrix of partial derivatives of $f_i(\mathbf{z};\delta)$ (derivatives of f_i with respect to z_j) at $\delta = 0$, ∇f, is a block matrix (where I is the identity matrix):

$$\nabla f = \begin{bmatrix} 3I & 0 \\ I & I \end{bmatrix},$$

which has full rank. By part (i) of Assumption 1, all values of f_i are continuously differentiable. Therefore, the implicit function theorem implies that there exists a unique equilibrium in the neighborhood of $\mathbf{z}^*$ and $\delta = 0$, where $\mathbf{z}^*$ is the (unique) equilibrium at $\delta = 0$. By continuity and the assumption that all elements of $\mathbf{z}$ belong to a compact set, there exists no other equilibrium in the neighborhood of $\mathbf{z}^*$, which finally implies that there exists a unique equilibrium in the neighborhood of $\delta = 0$. □

Proof of Lemma 2. The game I consider has the structure of a "birth-and-death" Markov process (Kelly 1979, Section 1.3); that is, from any given state i, the only transitions to consider are to the neighboring states: $M(i,k) = 0$ if $|i-k| > 1$. These processes are recursive (Kelly 1979, Lemma 1.5). It follows that they are also stationary. Recursiveness also implies that detailed balance holds (Kelly 1979, Theorem 1.3); namely,

$$M(i-1,i)m(i-1) = M(i,i-1)m(i). \tag{A.3}$$

The value of $M(i-1,i)$ corresponds to nature's selecting as an active consumer one of the consumers with the firm that currently has $\eta - i + 1$ consumers, and that agent switching to the other firm (that is, the firm currently having $i-1$ consumers). This happens with probability

$$M(i-1,i) = \frac{\eta - i + 1}{\eta} q(i-1).$$

Similarly,

$$M(i,i-1) = \frac{i}{\eta}(1 - q(i)).$$

Equation (A.3) allows me to compute the stationary distribution recursively. Given $m(0)$, we have

$$m(i) = m(0)\prod_{k=1}^{i} \frac{M(i-1,i)}{M(i,i-1)} = m(0)\prod_{k=1}^{i} \frac{q(i-1)}{1-q(i)} \cdot \frac{\eta - i + 1}{i}.$$

Since $\sum_{k=0}^{\eta} m(k) = 1$,

$$m(0) = \left(1 + \sum_{i=1}^{\eta}\prod_{k=1}^{i} \frac{q(i-1)}{1-q(i)} \cdot \frac{\eta - i + 1}{i}\right)^{-1}.$$

Equation (A.3) also implies that $m(i) > m(i-1)$ if and only if

$$\frac{\eta - i + 1}{\eta} q(i-1) > \frac{i}{\eta}(1 - q(i)).$$

By a similar argument, $m(i) > m(i+1)$ if and only if

$$\frac{\eta - i}{\eta} q(i) < \frac{i+1}{\eta}(1 - q(i+1)),$$

which concludes the proof. □

Proof of Proposition 3. Suppose that $\delta = 0$. Then (7) becomes

$$w(i) = \theta(\Delta(i+1) - \Delta(i)) = \theta \mathbb{1}_{i \in \{i^*-1, i^*\}},$$

where $\mathbb{1}_x$ is an indicator variable that equals 1 if x is true and 0 otherwise:

$$w(j-1) = \theta \mathbb{1}_{j-1 \in \{i^*-1, i^*\}} = \theta \mathbb{1}_{j \in \{i^*, i^*+1\}} = \theta \mathbb{1}_{i \in \{i^*, i^*-1\}} = w(i).$$

Substituting in (13), this implies $x(i) = 0$ (see the proof of Lemma 1). From (8), we get

$$p(i) = \frac{1}{2\phi(0)} - \frac{i}{\eta}\theta \mathbb{1}_{i \in \{i^*, i^*+1\}} - \frac{j}{\eta}\theta \mathbb{1}_{i \in \{i^*-1, i^*\}},$$

which implies the first expression in the result.

Substituting $\delta = 0$ in (10), I get

$$\mathring{v}(i) = \frac{1}{4\phi(0)} + \frac{i}{\eta}\theta\Delta(i-1) + \frac{j}{\eta}\theta\Delta(i),$$

which in turn implies the expression in the result.

I next turn to the stationary distribution of market shares. Since $\lim_{\delta \to 0} q(i) = \frac{1}{2}$, Lemma 2 implies that, in the limit as $\delta \to 0$,

$$m(i) = m(0)\prod_{k=0}^{i} \frac{\eta - i + 1}{i} = \frac{\eta!}{i!(\eta - i)!},$$

where

$$m(0) = \left(1 + \sum_{i=1}^{\eta} \frac{\eta!}{i!(\eta-i)!}\right)^{-1} = \left(\sum_{i=0}^{\eta} \frac{\eta!}{i!(\eta-i)!}\right)^{-1} = 2^{-\eta},$$

which leads to the expression for $m(i)$ in the result. Finally, setting $q(i-1) = q(i) = q(i+1) = \frac{1}{2}$, the Lemma 2 conditions for $m(\eta/2)$ to be greater than its neighbors become

$$\left(\eta - \frac{\eta}{2} + 1\right)\frac{1}{2} > \frac{\eta}{2}\left(1 - \frac{1}{2}\right) \quad \text{and} \quad \left(\eta - \frac{\eta}{2}\right)\frac{1}{2} < \left(\frac{\eta}{2} + 1\right)\left(1 - \frac{1}{2}\right),$$

both of which are equivalent to $\eta + 2 > \eta$. □

Proof of Corollary 1. From the expression for $v(i)$ in Proposition 3, we get

$$\lim_{\delta \to 0} v(i) + v(j) = \begin{cases} \dfrac{1}{2\phi(0)} - 2\dfrac{i^* - 1}{\eta}\theta & \text{if } i = i^*, \\ \dfrac{1}{2\phi(0)} - \dfrac{\eta - i^* - 1}{\eta}\theta & \text{if } i = i^* \pm 1, \\ \dfrac{1}{2\phi(0)} & \text{otherwise,} \end{cases}$$

which is decreasing in θ, strictly if $i \in \{i^* - 1, i^*, i^* + 1\}$. □

Proof of Proposition 4. Define, for a generic variable x,

$$\mathring{x} \equiv x|_{\delta=0} \qquad \dot{x} \equiv \left.\frac{\partial x}{\partial \delta}\right|_{\delta=0}.$$

Differentiating (7), I get

$$\mathring{w}(i) = \mathring{v}(i+1) - \mathring{v}(i). \tag{A.4}$$

Differentiating (13), I get

$$\begin{aligned} -3\dot{x}(i) &= \frac{i}{\eta}(\mathring{w}(i-1) - \mathring{w}(j)) + \frac{j}{\eta}(\mathring{w}(i) - \mathring{w}(j-1)) \\ &= \frac{i}{\eta}(\mathring{v}(i) - \mathring{v}(i-1) - \mathring{v}(j+1) + \mathring{v}(j)) \\ &\quad + \frac{j}{\eta}(\mathring{v}(i+1) - \mathring{v}(i) - \mathring{v}(j) + \mathring{v}(j-1)), \end{aligned} \tag{A.5}$$

where the second equality follows from (A.4). Considering that $\delta = 0$ implies $x(i) = 0$ for all i (see Proposition 3), by substituting $i^* + 1$ for i in (A.5), we can state that $x(i^* + 1) < 0$ if and only if

$$\begin{aligned} &\frac{i^* + 1}{\eta}(\mathring{v}(i^* + 1) - \mathring{v}(i^*) - \mathring{v}(i^*) + \mathring{v}(i^* - 1)) \\ &\quad + \frac{i^* - 1}{\eta}(\mathring{v}(i^* + 2) - \mathring{v}(i^* + 1) - \mathring{v}(i^* - 1) + \mathring{v}(i^* - 2)) > 0. \end{aligned}$$

(Recall that, if $i = i^* + 1$, then $j = \eta - i = i^* - 1$.) Dividing throughout by θ/η (a positive constant), and substituting the equilibrium values of $v(i)$ in Proposition 3 for the various $\mathring{v}$, I get

$$(i^* + 1)\left(3\frac{i^* - 1}{\eta} - 1\right) + (i^* - 1)\left(-\frac{i^* - 1}{\eta} + 1\right) > 0,$$

which, recalling that $\eta = 2i^*$, is equivalent to $i > 2$. This proves the first two inequalities.

Differentiating (9) with respect to δ at $\delta = 0$, I get

$$\dot{v}(i) = -\dot{x}(i) + \frac{i}{\eta}\mathring{v}(i-1) + \frac{j}{\eta}\mathring{v}(i). \tag{A.6}$$

From Proposition 3, $\mathring{v}(i)$ is constant for $i < i^*$; that is, $i < i^*$ implies $\mathring{v}(i) = \mathring{v}(i^* - 1)$. Substituting into (A.6), we get that $i < i^*$ implies

$$\dot{v}(i) = -\dot{x}(i) + \mathring{v}(i^* - 1). \tag{A.7}$$

Also substituting $\mathring{v}(i^* - 1)$ for $\mathring{v}(i)$ in (A.5), I get $\dot{x}(i^* - 2) = 0$. Finally, (A.7) implies

$$\begin{aligned} \dot{v}(i^* - 1) - \dot{v}(i^* - 2) &= -\dot{x}(i^* - 1) + \mathring{v}(i^* - 1) + \dot{x}(i^* - 2) - \mathring{v}(i^* - 1) \\ &= -\dot{x}(i^* - 1), \end{aligned}$$

which, as shown above, is negative. □

Endnotes

[1] Ferrier et al. (1999) quote a series of *Wall Street Journal* headlines (though no formal cites are supplied), including "Alex Trotman's goal: To make Ford No. 1 in world auto sales," "Kellogg's cutting prices... to check loss of market share," and "Amoco scrambles to remain king of the polyester hill."

[2] In Section 6, I also examine the behavior of managers who sell to consumers who get an extra utility kick from buying from a market share leader.

[3] Maskin and Tirole (1988) consider an alternating-move game and show the existence of equilibria featuring Edgeworth cycles, which are similar to, but different from, price wars.

[4] The idea goes back to (at least) Schelling's (1960, pp.142–143) observation that "the use of thugs or sadists for the collection of extortion or the guarding of prisoners, or the conspicuous delegation of authority to a military commander of known motivation, exemplifies a common means of making credible a response pattern that the original source of decision might have been thought to shrink from or to find profitless, once the threat had failed."

[5] The cost of a more general model is that some of the results are derived by numerical computation.

[6] The model presented in this section and the model presented in the next section are sufficient general to encompass two interpretations of the value of θ. One, which I will use throughout the paper, is that the value of θ corresponds to corporate culture. A second is that θ reflects individual psychological characteristics—namely, a manager's "big ego."

[7] Similar to Cabral (2011), the assumption of discrete time with exactly one consumer being "active" in each period may be interpreted as the reduced form of a continuous time model where each consumer becomes "active" with a constant hazard rate ν. The relevant discount factor is then computed as $\delta \equiv \eta\nu/(r + \eta\nu)$, where r is the continuous time interest rate.

[8] Under the model interpretation that consumers are born and die, the i.i.d. assumption seems reasonable. Under the active/inactive consumer assumption, this assumption has the unreasonable implication that the preferences of an active consumer are independent of its previous preferences. In this sense, my model may be seen as an approximation or as assuming that consumers and firms do not take this time correlation into account when computing value functions.

[9] In Section 6, I consider the possibility that consumers derive utility from purchasing from a market share leader; that is, consumers derive utility $\lambda\Delta(i)$ in addition to the $\zeta_i - p(i)$ term considered above.

[10] Notice that, for the extreme case $i = 0$, (6) calls for values of $v(\cdot)$, which are not defined. However, these values are multiplied by zero.

[11] The reason why the index in the various components differs—i for $p(i)$ and $i+1$ for $\theta\Delta(i+1)$ and $v(i+1)$—is that strategy $p(i)$ is defined over the initial state, i, whereas payoff $\theta\Delta(i')$ and continuation value $v(i')$ are defined over the new state i' resulting from the current active consumer's decision.

[12] The qualitative features of the results remain the same for different values of η. However, in the limit when $\eta \to \infty$, aggregate noise vanishes and the model becomes deterministic.

[13] The assumption that ξ follows a standardized normal implies no additional loss of generality with respect to ξ being normal, on account of my symmetry assumption and an appropriate change of units.

[14] Cabral and Villas-Boas (2005) denote by *Bertrand supertrap* the situation (as is the present case) when the strategic effect of an exogenous change is greater in absolute value and opposite in sign to the direct effect.

[15] The joint value effect corresponds vaguely to the principle of least action in classical mechanics; dynamic pricing implies that, in expected terms, the state space moves in the direction that joint value is maximized.

[16] This is similar to the approach followed by Budd et al. (1993) and Cabral and Riordan (1994).

[17] I also considered higher values of δ (e.g., $\delta = 0.9$). For higher values of δ, convergence to equilibrium is more difficult to obtain numerically. Given that, either one lowers the precision of calculations or the time spent computing equilibrium values becomes prohibitively high. In the numerical simulations I consider, I set a tolerance level of 1E-6 for the percent difference in the total value of value functions at every state between two consecutive iterations. See also Doraszelski and Pakes (2007).

[18] Simulations show that this requires the value of δ to be sufficiently high. In fact, all curves vary smoothly with δ; for $\delta = 0$, the stationary distribution is unimodal, as we saw earlier.

[19] The simulation starts with $i = 50$ and is based on the default random seed.

[20] Specifically, $\theta = 0 \Rightarrow p(i) = 1/\phi(0) = 1.2533$, given my assumption that $\xi \sim \mathrm{N}(0, 1)$.

[21] In fact, I adapt the term "trenchy" from Besanko et al. (2010).

[22] I am grateful to a referee for this suggestion.

[23] See also Section 6.3.

[24] As in the asymmetric cost Bertrand game, we assume that while prices are the same, the firm with higher "margin" makes the sale with probability 1. We can think of this equilibrium as the limit of a pricing game on a finite grid, in which case, firm a prices just above v. See Tirole 1988.

References

Barney JB (1986) Strategic factor markets: Expectations, luck, and business strategy. *Management Sci.* 32(10):1231–1241.
Baumol WJ (1962) On the theory of expansion of the firm. *Amer. Econom. Rev.* 52(5):1078–1087.
Beggs AW, Klemperer P (1992) Multi-period competition with switching costs. *Econometrica* 60(3):651–666.
Besanko D, Doraszelski U, Kryukov Y, Satterthwaite M (2010) Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica* 78(2):453–508.
Budd C, Harris C, Vickers J (1993) A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease? *Rev. Econom. Stud.* 60(3):543–573.
Cabral L (2011) Dynamic price competition with network effects. *Rev. Econom. Stud.* 78(1):83–111.
Cabral L, Riordan MH (1994) The learning curve, market dominance, and predatory pricing. *Econometrica* 62(5):1115–1140.
Cabral L, Villas-Boas M (2005) Bertrand supertraps. *Management Sci.* 51(4):599–613.
Caves RE, Porter ME (1977) From entry barriers to mobility barriers: Conjectural decisions and contrived deterrence to new competition. *Quart. J. Econom.* 91(2):241–261.
Dierickx I, Cool K (1989) Asset stock accumulation and sustainability of competitive advantage. *Management Sci.* 35(12):1504–1511.
Dignan L (2011) IBM reclaims server market share revenue crown in Q4, says Gartner. *ZDNet* (February 24), http://www.zdnet.com/article/ibm-reclaims-server-market-share-revenue-crown-in-q4-says-gartner/.
Doraszelski U, Pakes A (2007) A framework for applied dynamic analysis in IO. *Handbook of Industrial Organization*, Vol. 3 (Elsevier, Oxford, UK), 1887–1966.
Farber D (2007) SAP CEO: We are not arrogant, we are the market leader. *ZDNet* (April 25), http://www.zdnet.com/article/sap-ceo-we-are-not-arrogant-we-are-the-market-leader/.
Ferrier WJ, Smith KG, Grimm CM (1999) The role of competitive action in market share erosion and industry dethronement: A study of industry leaders and challengers. *Acad. Management J.* 42(4):372–388.
Fershtman C, Judd KL (1987) Equilibrium incentives in oligopoly. *Amer. Econom. Rev.* 77(5):927–940.
Fudenberg D, Tirole J (1984) The fat-cat effect, the puppy-dog ploy, and the lean and hungry look. *Amer. Econom. Rev.* 74(2):361–366.
Gilbert RJ, Newbery DMG (1982) Preemptive patenting and the persistence of monopoly. *Amer. Econom. Rev.* 72(3):514–526.
Green EJ, Porter RH (1984) Noncooperative collusion under imperfect price information. *Econometrica* 52(1):87–100.
Isidore C (2012) GM back on top in global sales race. *CNNMoney* (January 19), http://money.cnn.com/2012/01/19/autos/gm_global_sales_leader/.
Kelly FP (1979) *Reversibility and Stochastic Networks* (Wiley, New York).
Maskin E, Tirole J (1988) A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica* 56(3):571–599.
Ostrower J (2012) Airbus accuses Boeing of price war. *Wall Street J.* (May 15), https://www.wsj.com/articles/SB10001424052702304371504577406830992897136.
Rotemberg JJ, Saloner G (1986) A supergame-theoretic model of price wars during booms. *Amer. Econom. Rev.* 76(3):390–407.
Schelling TC (1960) *The Strategy of Conflict* (Harvard University Press, Cambridge, MA).
Shapiro C (1989) Theories of oligopoly behavior. Schmalensee R, Willig R, eds. *Handbook of Industrial Organization*, Vol. 1, Chap. 6 (Elsevier B.V., Amsterdam), 329–414.
Sklivas SD (1987) The strategic choice of managerial incentives. *RAND J. Econom.* 18(3):452–458.
Tirole J (1988) Short-un price competition. *The Theory of Industrial Organization*, Chap. 5 (MIT Press, Cambridge, MA).
Vickers J (1985) Delegation and the theory of the firm. *Econom. J.* 95(380a):138–147.
Wernerfelt B (1984) A resource based view of the firm. *Strategic Management J.* 5(2):171–180.