

How to Curb the Price-Competition Paradox¹

Hans Carlsson

Department of Economics

Lund University

June 30, 2010

Abstract

A well-known result says that a static duopoly where both firms have constant unit costs and compete in prices is characterized by a unique Nash-equilibrium outcome in which the low-cost firm serves the whole market at a price equal to the other firm's unit cost. In the symmetric-cost case the market is split and both firms earn zero profits. We argue that the result is based on an inappropriate use of the Nash concept because the equilibria involve the use of weak best replies. To obtain a more satisfactory solution we look for strategy sets which are closed under best replies (so-called curb sets, see Basu and Weibull, 1991). We show that in the asymmetric-cost case any curb set is compatible with market prices ranging from the high-cost firm's unit cost to the low-cost firm's monopoly price. This result remains valid in settings with increasing or slightly decreasing returns to scale.

1 Introduction

A well-known text-book story of price competition goes like this: two firms which produce a homogenous good at constant unit costs choose prices simultaneously in a one-shot interaction. The entire market demand goes to the firm charging the lowest price; in case of a tie market demand is split equally between the firms. The predicted outcome is that the firm with the lowest unit cost will get the entire demand at a price equal to the unit cost of the other firm. In the special case where firms have the same unit cost, price will be equal to this common unit cost and both firms will make zero profits. This

¹ Financial support from the Swedish Competition Authority, Jan Wallander's and Tom Hedelius' Foundation and Tore Browaldh's Foundation is gratefully acknowledged.

outcome appears almost paradoxical; it is highly unintuitive that in a market with only two firms none of them will be able to make a profit. The paradox is barely mitigated in the more realistic case where firms have slightly different costs.²

Notwithstanding its lack of intuitiveness the outcome is typically regarded as *the* solution of the price-competition model. The usual text-book approach is to look for more sensible outcomes in variants involving repeated interaction, differentiated goods or capacity constraints. No-one, however, questions the correctness of the solution to the original model. This paper will argue that such a questioning is legitimate.

The solution to the price-competition model is supported by the fact that it is the unique Nash-equilibrium outcome of the corresponding strategic game. We will argue that the notion of Nash equilibrium, even when associated with a unique outcome, is not necessarily an adequate basis for predicting an outcome. In this particular instance, the problem is that the Nash equilibrium is *weak*: at least one firm is indifferent between its equilibrium price and a large range of other prices. Thus it fails to meet the rather compelling requirement that the solution to a model should include for each agent all the choices that are rational given the other agents' possible choices. This requirement naturally leads to consider set-valued solution concepts such as Basu and Weibull's (1991) notions of *curb sets* and *curb* sets*: A curb set is a product set of strategies which includes all the best replies for each player to all strategy combinations included for the other players. A curb* set is similar but excludes the use of dominated strategies.

In the following we analyse price competition models using the notion of curb. The conclusions yielded by this analysis differ radically from the traditional ones. The solution is typically compatible with market prices ranging from the high-cost firm's unit cost to the low-cost firm's monopoly price. The result is valid in settings with increasing, constant or slightly decreasing returns to scale. The only setting where the paradox can be supported by the curb concept is the very special case in which costs are symmetric and proportional.

² As has been observed by several scholars the attribution of the model—and, thus, the paradox—to Joseph Bertrand is not well-founded (see Bertrand, , and, e.g., Dimand and Dore, 1999).

Section 2 introduces and motivates notions of curb sets and curb* sets. Section 3 presents the price-competition model with constant and increasing returns to scale and provides the main results which are extended to models with slightly decreasing returns to scale and slightly differentiated goods in Section 4. Section 5 concludes.

2 Desiderata for solutions

The price-competition model is conveniently viewed as a static or one-shot game. In the present section we define curb sets and present two desiderata for solutions to static games which lead to the notion of *tight* curb sets. To start we provide some basic definitions.³

An n -person game in the strategic form is defined by its pure strategy sets S_i , and payoff functions $u_i: S \rightarrow \Re$ where $S \equiv S_1 \times \dots \times S_n$, $i = 1, \dots, n$. We allow players to use mixed strategies and extend the payoff functions in the usual way; $M(S_i)$ denotes the set of mixed strategies for player i and $M(S) \equiv M(S_1) \times \dots \times M(S_n)$ is the set of mixed strategy profiles.

A *retract* is a product set of strategies $T = T_1 \times \dots \times T_n$, $T_i \subseteq M(S_i)$, $i = 1, \dots, n$. A retract T is *pure* if each T_i contains only pure strategies. We extend the above notation so that, for any subset T_i of (pure or mixed) strategies, $M(T_i)$ is the set of all mixtures of strategies in T_i , i.e. the convex hull of T_i ; moreover for any retract $T = T_1 \times \dots \times T_n$, $M(T) \equiv M(T_1) \times \dots \times M(T_n)$. If T is a pure retract, $M(T)$ is the *mixed extension* of T .

For a mixed strategy combination $\sigma_{-i} \in \prod_{j \neq i} M(S_j)$, we let $B_i(\sigma_{-i})$ denote i 's set of pure-strategy best replies to σ_{-i} :

$$B_i(\sigma_{-i}) := \arg \max_{s_i \in S_i} u_i(s_i, \sigma_{-i}).$$

Given any retract T , we let $B_i(T)$ be the set of pure strategies for i which are best replies to some mixed strategy combination $\sigma_{-i} \in \prod_{j \neq i} M(T_j)$:

$$B_i(T) := \bigcup_{\sigma_{-i} \in \prod_{j \neq i} M(T_j)} B_i(\sigma_{-i}).$$

³ This section borrows from Carlsson (2010).

Notice that $B(T) \equiv B_1(T) \times \dots \times B_n(T)$ is a pure retract and that, by definition

$$(2.1) \quad B(T) \equiv B(M(T)).$$

The corresponding sets of *mixed* best replies will be denoted $\beta_i(\sigma_{-i})$ and $\beta_i(T)$. Recall the elementary property

$$(2.2) \quad \beta_i(\sigma_{-i}) = M(B_i(\sigma_{-i})).$$

Definition: A pure retract T is a *curb set* if it is closed under pure best replies, i.e.

$$(2.3) \quad B(T) \subseteq T.$$

A curb set is *tight* if $B(T) = T$ and *minimal* if it does not contain a smaller curb set.

Using (2.2) it is easily seen that the corresponding requirement for mixed strategies, viz. $\beta(T) \subseteq M(T)$, is automatically met if (2.3) holds: *the mixed extension of a curb set is closed under mixed best replies*. Also notice that any finite game has a minimal curb set and that minimal curb sets are tight.

We consider a solution to a game both as a set of recommendations about strategies that players should use and as a set of conjectures that each player can hold about the other players' behavior. In order to explain the rationale for curb as a solution concept it is useful to make this division explicit. Thus we will start from the assumption that a solution always consists of two product sets of mixed strategies $X = X_1 \times \dots \times X_n$, $X_i \subseteq M(S_i)$, and $Y = Y_1 \times \dots \times Y_n$, $Y_i \subseteq M(S_i)$, $i = 1, \dots, n$. Each X_i represents the behavior recommended for player i and each Y_i is the set of conjectures which the other players may have about i 's behavior. Naturally these two aspects of a solution should be closely connected: X_i should contain strategies that can be rationalized given that i 's conjectures about the others' behavior lie in Y_{-i} and Y_i should represent the set of conjectures which the other players can entertain about i 's behavior given that i will choose a strategy in X_i . We propose the following two assumptions.

A1. For each i , $Y_i = M(X_i)$.

A2. For each i , $X_i = \beta_i(Y)$.

Regarding A1 it seems natural to postulate that the other players can have conjectures corresponding to any strategy in X_i : if σ_i is rational for i under the given solution, the other players are justified in conjecturing that i will play σ_i . Granted this postulate it is also reasonable to allow conjectures corresponding to probability distributions on X_i , whence the condition. As for A2 this assumption posits that a strategy is rational for i if and only if it is the best reply to some conjecture allowed by the solution. This requirement serves to guarantee that the solution is stable in the sense that conjectures and behaviors reinforce each other. Indeed if A2 is not satisfied, then $\beta_i(Y) \not\subset X_i$ or $X_i \not\subset \beta_i(Y)$. In the first case there exists for some player i a conjecture $\sigma_{\cdot i}$ and a strategy σ_i such that σ_i is a best reply to $\sigma_{\cdot i}$ but not an element of X_i . As σ_i is optimal under some conjecture allowed by the proposed solution it seems questionable to exclude it from the solution. By contrast, in the second case, for some player i , X_i contains some element that cannot be rationalized by any conjecture in $Y_{\cdot i}$ and, thus, need not be included in the solution.

The following result establishes a close connection between the above assumptions and the notion of a tight curb set.

Proposition 2.1 (Carlsson, 2010): *If X and Y satisfy A1 and A2, then Y is the mixed extension of a tight curb set. Conversely, if T is a tight curb set, then $X = \beta(T)$ and $Y = M(T)$ satisfy A1 and A2.*

This result shows that if one thinks that Assumptions A1 and A2 are attractive properties of a solution, then one should look for tight curb sets.⁴ In some games, however, the curb condition may seem too restrictive. In the following example from Basu and Weibull (1991) the game has a unique curb set which is simply the whole game. Yet (T, L) is, arguably, the only really rational strategy profile as both B and R are weakly dominated. (Recall that a strategy σ_i is weakly dominated if there exists σ_i' which

⁴ For axiomatic treatments and properties of various curb notions, see Dufwenberg et al. (2001) and Voorneveld et al. (2005). In addition to their appeal as solutions to games played by rational players curb sets are interesting in adaptive learning models. For instance, Hurkens (1995) shows that players with bounded memory who play best replies against their beliefs will end up playing strategies from a minimal curb set.

gives i at least the same payoff against any σ_{-i} and strictly higher payoff against some σ_{-i} .)

	L	R
T	1, 1	0, 1
B	1, 0	0, 0

By excluding the use of (weakly) dominated strategies we can get sharper predictions. We let $S_i^* \subseteq S_i$ be the set of undominated pure strategies for player i . B_i^* is the undominated best-reply correspondence for player i :

$$B_i^*(\sigma_{-i}) := B_i(\sigma_{-i}) \cap S_i^*$$

and

$$B_i^*(T_{-i}) := B_i(T_{-i}) \cap S_i^*.$$

Definition: A pure retract T is a *curb** set if it is closed under undominated best replies, i.e. $B^*(T) \subseteq T$. A curb* set T is *tight* if $B^*(T) = T$ and *minimal* if it does not contain a smaller curb* set.

Clearly we can modify Proposition 1 to get a case for tight curb* sets by replacing assumption A1 by the following:

A2*. For each i , $X_i = \beta_i^*(Y_{-i})$.

Whether curb or curb* is the most appropriate solution concept thus partly boils down to whether one prefers A2 or A2*. In the context of the price-competition model the two notions do not yield very different conclusions but the curb sets are easier to characterize and will be the focus of our analysis.

It is interesting to contrast these notions with that of Nash equilibrium which is the standard solution concept for static games. Recall that a strategy profile σ is a *Nash equilibrium* if, for each player i , σ_i is a best reply to σ_{-i} , i.e.

$$\sigma_i \in M(B_i(\sigma_{-i})).$$

A Nash equilibrium σ is *strict* if, for each i , $B_i(\sigma_{-i}) = \{\sigma_i\}$ ⁵; otherwise it is *weak*.

From our point of view the main shortcoming of Nash equilibrium as a solution concept is that it may fail to satisfy assumption A2 (and its variant A2*): Nash equilibria do not necessarily contain all the best replies to the strategies they prescribe. This will be the case only when the equilibrium is strict: a set $\{\sigma\}$ consisting of a single strategy profile is a curb set if and only if σ is a strict Nash equilibrium. Thus, in models that do not have strict equilibria an analysis in terms of some curb concept is necessary to insure a solution which is supported by adequate incentives. Indeed, as noted by Basu and Weibull (1991), the curb requirement may be seen as a set-valued generalization of the notion of strict equilibrium.

The following Proposition summarizes well-known properties of curb sets and is valid for curb* sets as well.

Proposition 2.2:

- (i) *Every finite game has at least one tight curb set.*
- (ii) *Any curb set contains the support of at least one Nash equilibrium.*

The game in Table 2.2 shows that curb sets and Nash equilibrium may yield quite different conclusions. The Nash equilibria of the game are all of the form (M, σ_2) where $\sigma_2(L) = x$, $\sigma_2(C) = 1 - x$, $\sigma_2(R) = 0$, and $x \in [1/3, 2/3]$ and all result in the same payoffs, viz. (2, 1). For the curb analysis we can use Proposition 2.2 (ii) to conclude that any curb set contains $(\{M\}, \{L, R\})$. Any curb set must also contain T and B which are 1's best replies to L and R . Thus, the game has a unique curb set, which consists of the full strategy space.

	L	R
T	3, 3	0, 4
M	2, 1	2, 1
B	0, 4	3, 3

Table 2.2

⁵ As a strict Nash equilibrium is necessarily in pure strategies there is no need to write $M(B_i(\sigma_{-i})) = \{\sigma_i\}$.

As we will see next a similar phenomenon—albeit on a larger scale—appears in price-competition games. In such games, generally, every Nash equilibrium has the property that one firm wins the whole market by charging a price which is strictly lower than the other firm's price. The losing firm cannot compete profitably and will thus earn zero. As a result every curb set will contain all strategies for the loser which earn zero against the winner's equilibrium price as well as the winner's best replies to these strategies. Such a curb set typically consists of a wide range of prices for each firm implying a high degree of market price indeterminacy.

3 Price Competition with Constant or Increasing Returns to Scale

Studies of the price-competition model have mostly focused on the case with proportional costs. However, both the Nash-equilibrium and the curb analyses are very similar whether there are constant or increasing returns to scale. Thus it is convenient to use a model which covers both these cases.

The model has two firms, 1 and 2, with increasing and concave cost functions C_1 and C_2 such that $C_1(0) = C_2(0) = 0$. Total demand is given by a continuous decreasing function $D: \mathfrak{R}_+ \rightarrow \mathfrak{R}_+$; we assume there exists $\bar{p} > 0$ such that $D(p) = 0$ if and only if $p \geq \bar{p}$. In the duopoly game firms choose prices p_1 and p_2 from $P \equiv [0, \bar{p}]$ simultaneously and total demand goes to the firm with the lowest price. In case of a tie, each firm gets half the total demand.⁶ Let $\Pi_i^m(p)$ denote the profit firm i would get as a monopoly charging price p :

$$\Pi_i^m(p) := pD(p) - C_i(D(p))$$

and let $\Pi_i^d(p)$ be firm i 's *duopoly profit*, i.e. its profit from selling $D(p)/2$ at price p :

$$\Pi_i^d(p) := pD(p)/2 - C_i(D(p)/2).$$

Firm i 's profit when prices p_i and p_j are chosen, $\pi_i(p_i, p_j)$, can then be expressed follows

⁶ None of the results to follow depends critically on the equal-split assumption. If instead we assume that, for some $\lambda \in (0, 1)$, firm 1 gets $\lambda D(p)$ and firm 2 $(1 - \lambda)D(p)$, all the results remain valid.

$$\pi_i(p_i, p_j) = \begin{cases} \Pi_i^m(p) & \text{if } p_i < p_j \\ \Pi_i^d(p) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j. \end{cases}^7$$

To ensure non-triviality, we make the following assumption:

A3: *There exists p such that $\Pi_i^m(p) > 0$.*

Notice that if $\Pi_i^m(p) > 0$, then $\Pi_i^m(p) > \Pi_i^d(p)$.

We assume each firm has a unique monopoly price to be denoted p_i^m (p_i^m maximizes $\Pi_i^m(p)$). We let m_i denote firm i 's *reservation price*, i.e. the highest lower bound on the prices at which firm i may make a profit:

$$m_i := \inf \{p \in P; \Pi_i^m(p) > 0\}$$

(Notice that, in the constant-returns case where, for some c_i , $C_i(q) = c_i q$, the reservation price coincides with the constant unit cost c_i .) Obviously $m_i < p_i^m$ and $\Pi_i^m(m_i) = 0$; in the constant-returns case we have $\Pi_i^d(m_i) = 0$ but whenever returns are strictly increasing $\Pi_i^d(m_i) < 0$. We assume:

A4: $m_1 \leq m_2$.

A5: $m_2 < p_1^m$.

A4 entails no loss of generality, while A5 serves to avoid the trivial case where firm 2 cannot compete profitably at firm 1's monopoly price.

We let d_i denote the highest lower bound on the prices at which firm i 's duopoly profit is positive if there exist such prices:

$$d_i := \inf \{p \in P; \Pi_i^d(p) > 0\};$$

otherwise we set $d_i := p_i^m$. Notice that $d_i \geq m_i$.

⁷ Here and in the sequel it is assumed that $i, j = 1, 2$ and $i \neq j$.

A well-known feature of the model—due to the payoff discontinuities—is that best replies do not exist for relevant ranges when prices can be chosen from a real interval. This makes the definition of curb sets problematic (What does it mean for a strategy set to be closed under best replies when best replies do not exist?) and in some cases entail nonexistence of Nash equilibrium. To circumvent this complication we assume prices have to be integer multiples of some monetary unit $k > 0$ and study the model's properties for small k . Thanks to this assumption the game is finite so Nash equilibrium and tight curb sets exist.

The game with monetary unit k will be denoted $G(k)$ and $P(k) := \{nk \in P; n \in N\}$ is the set of feasible prices—and, thus, each firm's set of pure strategies—in $G(k)$. We will assume k is small enough so that each firm could make a profit if it were alone on the market and let $p_i^m(k)$ be i 's monopoly price in $G(k)$ (in case it is not unique, we let $p_i^m(k)$ be the lowest monopoly price); notice that $p_i^m(k)$ converges to p_i^m as k goes to zero. To avoid an uninteresting complication we impose the following restriction:

A6: *There is no $p \in P(k)$ such that $\Pi_i^m(p) = 0$ or $\Pi_i^d(p) = 0$.*

We let $m_i(k)$ denote firm i 's reservation price in $G(k)$, $m_i(k) := \inf \{p \in P(k); \Pi_i^m(p(k)) > 0\}$ and assume k is small enough so that $m_i(k) < p_i^m(k)$. Notice that $m_i(k)$ converges to m_i as k goes to zero.

For a mixed strategy $\sigma_i \in M(P(k))$ with support $\text{supp } \sigma_i$ we use the notation

$$\sigma_i^- = \min \text{supp } \sigma_i$$

$$\sigma_i^+ = \max \text{supp } \sigma_i.$$

A strategy pair (σ_1, σ_2) is *overlapping* if $[\sigma_1^-, \sigma_1^+] \cap [\sigma_2^-, \sigma_2^+] \neq \emptyset$. Notice that if and only if (σ_1, σ_2) is *nonoverlapping*, a single firm—the *winner*—will serve the whole market with probability one.

In view of the following analysis it is useful to partition the model into three cases:

Case A: $m_1 < m_2$.

Case B: $m_1 = m_2 = m$ and $d_1 > m$ or $d_2 > m$.

Case C: $m_1 = m_2 = d_1 = d_2 = m$.

Remark 3.1: In Case C we have $\Pi_i^d(m) = \Pi_i^p(m) = 0$. Thus, clearly, $C_i(D(m)/2) = C_i(D(m))/2$ so, by the concavity assumption, costs are proportional and symmetric for quantities up to $D(m)$.

Proposition 3.1 summarizes Nash-equilibrium properties of the three cases.

Proposition 3.1:

(i) *In Case A, if k is small enough, every Nash equilibrium is nonoverlapping, firm 1 is the winner and market price lies in $[m_1(k), m_2(k)]$ or slightly above $m_2(k)$.*

(ii) *In Case B, if k is small enough, every Nash equilibrium is nonoverlapping and market price is slightly above m .*

(iii) *In Case C, if k is small enough, every Nash equilibrium is overlapping and market price is slightly above m .*⁸

The results in Proposition 3.1 to a large extent confirm well-known results for the corresponding model with continuous strategy spaces (see, e.g., Deneckere and Kovenock (1996) and Hoernig (2007)). The proof is given in Appendix A. Here we confine ourselves to providing the main intuitions. In all three cases the low market prices result from the usual underbidding logic. In Case A, the asymmetric outcome is driven by the cost asymmetry; in Case B it is a consequence of scale economies: for prices slightly above m a single firm will make a profit but if two firms charge the same price at least one of them will incur a loss. (Incidentally, if in this case both $d_1 > m$ and $d_2 > m$, no mixed-strategy Nash equilibrium exists in the continuous-strategy model; see Hoernig, 2007.) By contrast in Case C costs are symmetric and returns constant at the

⁸ We say that the (possibly random) market price *is slightly above* x if for any $\varepsilon > 0$ there exists $k^* > 0$ such that every point y in the support of the market price satisfies $x < y < x + \varepsilon$ in any game $G(k)$ with $k < k^*$.

relevant prices so two firms charging a price slightly above m will make small profits. As a result all Nash equilibria will be overlapping.⁹

Proposition 3.2 states the main properties of curb sets for Case A and B and of tight curb sets for Case C.

Proposition 3.2:

(i) *In Case A, if k is small enough and $V = V_1 \times V_2$ is a curb set, then $\min V_1 \leq m_2(k)$, $\max V_1 \geq p_1^m(k)$, and V_2 contains the set $[m_2(k) + k, \bar{p}] \cap P(k)$.*

(ii) *In Case B, if k is small enough and $V = V_1 \times V_2$ is a curb set, then, for some i , $\min V_i \leq m_j(k)$ and $\max V_i \geq p_i^m(k)$; moreover V_j contains the set $[m_j(k) + k, \bar{p}] \cap P(k)$.*

(iii) *In Case C, if k is small enough and $V = V_1 \times V_2$ is a tight curb set, then each V_i contains only prices slightly above m .*

The proof of Proposition 3.2 is quite straight-forward. It exploits Propositions 2.2 and 3.1 as well as the following Lemma which states a fundamental underbidding property of firm i 's best-reply correspondence in $G(k)$ which we denote B_i^k .

Lemma 3.1: *If $p' > m_i$, k is small enough, and σ_j is such that $\sigma_j^+ \in [p', p_i^m(k) + k] \cap P(k)$, then $\max B_i^k(\sigma_j) < \sigma_j^+$.*

Proof: If i chooses $p_i \geq \sigma_j^+$, its payoff will be at most $\max \{\sigma_j(\sigma_j^+) \Pi_i^d(\sigma_j^+), 0\}$ where $\sigma_j(\sigma_j^+)$ is the probability σ_j assigns to σ_j^+ . Define $F(\sigma_j^+) := \max \{\Pi_i^m(p); p \in [0, \sigma_j^+]\}$ and notice that $F(\sigma_j^+) > \max \{\Pi_i^d(\sigma_j^+), 0\}$ as $\sigma_j^+ > m_i$ and $\Pi_i^m(p) > \Pi_i^d(p)$ if $\Pi_i^m(p) > 0$. The result now follows from the observation that, if k is small, i can get at least approximately $\sigma_j(\sigma_j^+)F(\sigma_j^+)$ against σ_j in $G(k)$ by playing a uniformly mixed strategy with support on a large number of prices close to one associated with $F(\sigma_j^+)$. \square

Proof of Proposition 3.2: (i) Using Propositions 2.2 and 3.1 we can conclude that V contains the support of a nonoverlapping Nash equilibrium (σ_1, σ_2) where firm 1 is the winner. Thus $\min V_1 \leq m_2(k)$; otherwise we would have $\sigma_1^- > m_2(k)$ and firm 2 could make a profit by playing $m_2(k)$. As firm 2 earns zero in any Nash equilibrium, V_2 must

⁹ This is a point where Assumption A6 is needed: if m ($= m_1 = m_2$) belongs to $P(k)$ there are non-overlapping Nash equilibria where one firm plays m and the other $m + k$ or some mixed strategy with support above m , both firms earning zero.

contain the set X_2 of all strategies which earn zero against $\min V_1$. Clearly $X_2 \supseteq [m_2(k) + k, \bar{p}] \cap P(k)$. V_1 contains all best replies for firm 1 against X_2 , in particular $p_1^m(k)$.

(ii) It suffices to consider the case where $d_1 > m$. By Propositions 2.2 and 3.1 V contains the support of a nonoverlapping Nash equilibrium (σ_1, σ_2) . Thus either $\sigma_1^+ < \sigma_2^-$ or $\sigma_2^+ < \sigma_1^-$. If $\sigma_1^+ < \sigma_2^-$, firm 2 earns zero in the equilibrium so we must have $\sigma_1^- \leq m_2(k)$ (otherwise firm 2 could make a profit by playing $m_2(k)$). Thus $\min V_1 \leq m_2(k)$ and $\pi_2(\min V_1, p_2) \leq 0$ for any p_2 . Consequently V_2 contains the set Y_2 of all strategies which earn zero against $\min V_1$ and $Y_2 \supseteq [m_2(k) + k, \bar{p}] \cap P(k)$. V_1 contains all best replies for firm 1 against Y_2 , in particular $p_1^m(k)$. The case where $\sigma_2^+ < \sigma_1^-$ is shown in a similar way.

(iii) By Propositions 2.2 and 3.1 V contains the support of an overlapping Nash equilibrium (σ_1, σ_2) with prices slightly above m . Thus $\min V_i < \min p_i^m(k)$. Assume $\min V_1 \leq \min V_2$. Then $\min V_1 \geq m_1(k)$ as V is tight and any $p_1 < m_1(k)$ entails a loss for firm 1 against any $p_2 \geq p_1$ and thus cannot be a best reply. If we fix p' and k as in Lemma 3.1 and assume $\max V_1 \geq \max V_2$, then we can conclude from the tightness of V that $\max V_1 < p'$. \square

Corollary 3.1:

(i) *In the limit as k goes to zero any curb set in Case A and B allows any market price between the highest reservation price and the lowest monopoly price.*

(ii) *In Case C any market prices corresponding to a tight curb set converge to the reservation price as k goes to zero.*

The Corollary shows that from a curb point of view the price-competition paradox occurs only in what may be described as a special case of a special case: unless costs are both proportional and symmetric any curb set is compatible with a very wide range of market prices.

As curb* sets in some cases yield sharper predictions than curb sets it makes sense to study whether the use of curb* might alter the above conclusion. The following proposition and corollary show that there are no major changes. Let $P_i^*(k)$ denote the set of undominated strategies for firm i in $G(k)$ and notice that $m_i(k)$ and $p_i^m(k)$ belong to

$P_i^*(k)$. We assume Π_i^m has a unique maximum on $[0, p_j^m]$ and denote it b_i . Notice that $b_i = p_i^m$ if $p_i^m \leq p_j^m$; moreover $b_i = p_j^m$ if $p_i^m > p_j^m$ and Π_i^m is strictly quasi-concave. We let $b_i(k)$ be i 's lowest best reply to $p_j^m(k)$. Clearly $b_i(k)$ converges to b_i as k goes to zero.

Proposition 3.3:

(i) *In Case A, if k is small enough and $V = V_1 \times V_2$ is a curb* set, then $\min V_1 \leq m_2(k)$, $\max V_1 \geq b_1(k)$, and V_2 contains $p_2^m(k)$.*

(ii) *In Case B, if k is small enough and $V = V_1 \times V_2$ is a curb* set, then, for some i , $\min V_i \leq m_j(k)$ and $\max V_i \geq b_i(k)$; moreover V_j contains $p_j^m(k)$.*¹⁰

The proofs are straightforward and therefore omitted: curb* sets of a game can be found by looking for curb sets of the game obtained by removing all weakly dominated strategies from the original game.

Corollary 3.2: *In the limit as k goes to zero any curb* set in Case A and B allows market prices ranging from the highest reservation price to $\min b_i$.*

The Corollary shows that the substitution of curb* for curb will never alter the lowest market price allowed; the highest market price allowed will be lower under curb* only if $\min b_i < p_i^m$. In the more general Case A this requires both that $p_2^m < p_1^m$ (although $m_1 < m_2$) and that Π_1^m not be strictly quasi-concave. Moreover even if $\min b_i < p_i^m$, there may still remain a large degree of price indeterminacy. Hence the use of curb* does not really change the overall picture.

4 Two Variants

The results of the curb analysis show that the price-competition paradox only occurs in a very special case. A possible line of defence of the paradox might be to point to the existence of nearby models with well-behaved strict Nash equilibria close to the paradoxical outcome. In the present section we examine this argument by considering two variants of the above model, one with decreasing returns to scale and another with differentiated products. We study their behavior as the proportional-cost case of the

¹⁰ As the tight curb sets for Case C only contain undominated strategies they coincide with the tight curb* sets.

model of the previous section is approximated. In each case the above results are confirmed: the paradox occurs only if costs are symmetric; otherwise we get a highly indeterminate solution.

4.1 Decreasing Returns to Scale

We consider a model where firm i 's total cost is

$$C_i(q) = c_i q + t q^2$$

for some $c_i \in [0, 1)$ and $t > 0$. Total demand is given by

$$D(p) = \max\{1 - p, 0\}.$$

We assume a firm must meet the whole demand it is facing at a given price.¹¹ $\Pi_i^m(p)$ and $\Pi_i^d(p)$ can now be derived:

$$\Pi_i^m(p) = (1 - p)(p - c_i - t(1 - p))$$

$$\Pi_i^d(p) = \left(\frac{1 - p}{2}\right) \left(p - c_i - \frac{t(1 - p)}{2}\right).$$

We assume c_1 and c_2 , $c_1 \leq c_2$, are given and analyse the model's properties as t goes to zero so that a setup with proportional costs is approximated. Firm i 's monopoly price is given by

$$p_i^m = \frac{1 + c_i + 2t}{2 + 2t}.$$

m_i and d_i are defined as before and are now given by

¹¹ On this model see, e.g., Dastidar (1995) and Vives (1999). An alternative and, arguably, more interesting approach (e.g. Tirole, 1988) has each firm supply as much as it wants and allows for rationing. Such models, which are quite hard to analyse, are typically characterized by mixed-strategy Nash equilibria (in the non-trivial sense) which are overlapping and weak. Hence one cannot apply the line of argument used in the preceding section where the large size of curb sets in the general case could be derived from the fact that all Nash equilibria were nonoverlapping. Nor can one apply the argument in the special case which was based on the correspondence between (tight) curb sets and strict Nash equilibria. Nevertheless, I strongly conjecture that the results will be similar to the ones obtained for the approach studied here albeit by a more convoluted argument.

$$m_i = \frac{c_i + t}{1 + t}$$

$$d_i = \frac{2c_i + t}{2 + t}.$$

Clearly $\Pi_i^m(p) < 0$ if $p < m_i$, $\Pi_i^m(p) = 0$ if $p = m_i$, $\Pi_i^m(p) > 0$ if $p > m_i$, and similarly for $\Pi_i^d(p)$ and d_i .

e_i denotes the unique price at which $\Pi_i^m(p)$ and $\Pi_i^d(p)$ are equal:

$$e_i = \frac{2c_i + 3t}{2 + 3t}.$$

We have $\Pi_i^m(p) < \Pi_i^d(p)$ for $p < e_i$ and $\Pi_i^m(p) > \Pi_i^d(p)$ for $p > e_i$; hence it is profitable to undercut your competitor's price above, but not below e_i . Moreover

$$c_i < d_i < m_i < e_i < p_i^m,$$

which, i. a., shows that $\Pi_i^d(e_i) = \Pi_i^m(e_i) > 0$. We will assume $c_2 < \frac{1}{2}$ to insure that $d_2 < p_1^m$ and, thus, exclude the trivial case where firm 2 cannot compete profitably when firm 1 chooses its monopoly price.

Let us first consider the symmetric case where $c_1 = c_2 = c$, $d_1 = d_2 = d$, and $e_1 = e_2 = e$. It is easily seen that, in this case, any price p in the interval (d, e) corresponds to a strict and symmetric Nash equilibrium where each firm earns $\Pi_i^d(p) > 0$: if a firm deviates by choosing $p' < p$, it will earn less than $\Pi_i^m(p)$ since Π_i^m is strictly increasing in (d, e) ; as, moreover, $\Pi_i^m(p) < \Pi_i^d(p)$ the deviator will be worse off; the same conclusion is obvious for deviations to $p' > p$ which lead to zero profit. It is clear that the conclusion applies to a discrete model as well: in such a model, any $p \in (d, e) \cap P(k)$ defines a symmetric and strict Nash equilibrium. As any such equilibrium corresponds to a tight curb set there is no conflict with the curb analysis in this case. Furthermore, as both d and e converge to c as t goes to zero, market prices converge to the marginal cost of the canonical price-competition model so we get a confirmation of the paradox.

The precariousness of this result, however, becomes apparent once you consider the asymmetric case where $c_1 < c_2$. As long as $d_2 < e_1$ each price between these numbers

corresponds to a symmetric and strict Nash equilibrium. For t small enough though (the critical value being $2(c_2 - c_1)/(2 + c_1 - 3c_2)$), we have $e_1 < d_2$. In such case firm 1 has an incentive to undercut firm 2's price below d_2 where the 2 cannot even match 1's price without incurring a loss. As a result all Nash equilibria are nonoverlapping and of the form (p_1, σ_2) where $p_1 \in (e_1, d_2]$ and σ_2 is a mixed strategy with the probability mass concentrated just above p_1 in order to discourage firm 1 from choosing a price above p_1 . Again the conclusion applies to the discrete model as well. As firm 2 earns zero in any equilibrium we may use Proposition 2.2 to infer that, in any curb set, firm 2's component V_2 contains all strategies which earn zero against any $p_1 \in (e_1, d_2]$. Thus V_2 contains the set $[d_2, 1] \cap P(k)$ and V_1 contains all the best replies to this set, in particular V_1 contains $[d_2, p_1^m] \cap P(k)$. In the limit as t goes to zero we get the same indeterminate result regarding market prices as as in the asymmetric case of our main model.

4.2 Product differentiation

We consider a spatial model of product differentiation usually attributed to Hotelling; the present version borrows from Tirole (1988, pp. 279-280). Consumers are uniformly distributed on the unit interval and the two firms are located at its end points. A consumer's net utility is $v - p_i - td_i$ if he buys the good from firm i where $v > 0$ is his valuation of the good, p_i is firm i 's price, t is a transportation cost parameter, and d_i is the distance to firm i . If the net utility is positive and maximal for one firm only, he buys one unit of the good from that firm; if the net utility is negative for both firms, he does not buy. As, clearly, demand will be zero for a firm charging a price above v we will assume firms are restricted to choose prices in $[0, v]$.

The firms have constant unit costs c_1 and c_2 . We assume $c_1 \leq c_2 < v$ and that t is small so that a model with non-differentiated goods is approximated; in particular we assume $t < (v - c_2)/4$. Given these assumptions the best-reply correspondences can be shown to be as follows:

$$B_i(p_j) = \begin{cases} \{p_j - t\} & \text{if } p_j > c_i + 3t \\ \left\{ \frac{p_j + c_i + t}{2} \right\} & \text{if } c_i - t < p_j \leq c_i + 3t \\ [p_j + t, v] & \text{if } p_j \leq c_i - t. \end{cases}$$

The derivation is given in Appendix B; here we provide the main intuitions. Against sufficiently high p_j ($p_j > c_i + 3t$) it is optimal for firm i to choose $p_i = p_j - t$ which is the highest price at which i captures the whole market: for p_i above $p_j - t$ the loss in market share dominates the gains from a higher price in this case. Against intermediate p_j ($p_j \in (c_i - t, c_i + 3t)$) the optimal p_i is determined by a first-order condition where these two effects balance each other. For $p_j \leq c_i - t$, clearly, firm i cannot make a profit so $B_i(p_j)$ consists of all p_i for which $\pi_i(p_i, p_j) = 0$, viz. $[p_j + t, v]$.

From the best-reply correspondences it is straightforward to derive a strict Nash equilibrium in which the market is divided between the firms and prices are given by:

$$p_i = \frac{2c_i + c_j + 3t}{3}.$$

Such an equilibrium exists if $c_2 < c_1 + 3t$. If $c_2 > c_1 + 3t$, there exist only equilibria in which firm 1 serves the whole market. These equilibria are of the form (p_1, σ_2) where $p_1 \in (c_1, c_2 - t]$ and σ_2 is a mixed strategy with the probability mass sufficiently concentrated just above $p_1 + t$ to discourage firm 1 from choosing a price above p_1 .

It is easily seen that the strict Nash equilibrium will survive in the limit as t goes to zero only if $c_1 = c_2$; once more we get support for the paradox only in this special case. For $c_1 < c_2$ only equilibria in which firm 1 serves the whole market survive. As a result in any curb set $V = V_1 \times V_2$, V_2 must contain the set of prices $[c_2, v]$ at which firm 2 cannot incur a loss and V_1 contains all the best replies against this set, viz. $[c_2 - t, v - t]$. Thus, in the limit as t goes to zero market price can be anything in the interval $[c_2, v]$: we get the same kind of indeterminacy as in our main model.¹²

¹² As best replies are well-defined in this model we need not have recourse to a discrete variant.

5 Discussion

The above analysis has demonstrated that one reaches radically different conclusions depending on whether one uses Nash equilibrium or tight curb sets as solution concepts for price-competition games. The only case where the solutions coincide is so special that it may well be disregarded. The overall picture is that Nash equilibrium predicts the paradoxical outcome where market price will not exceed the highest reservation price while any solution based on curb is compatible with any price between the highest reservation price and the lowest monopoly price.

The point of confronting the equilibrium and curb analyses is not mainly an empirical one, it is rather a conceptual and pedagogical issue. From an empirical point of view both solution concepts perform quite poorly in the price-competition model. Nash equilibrium leads to an unintuitive prediction which is contradicted by experimental evidence¹³ while curb sets yield highly indeterminate predictions. In favor of curb one may note that it has somewhat more plausible comparative statics: While the Nash prediction only depends on the firms' reservation prices the set of market prices that are compatible with tight curb also reacts to changes in demand that affect monopoly prices. Still both concepts fail to produce the rather obvious prediction that market price should be negatively correlated with the number of firms.¹⁴ If one wants more interesting predictions one should either look for a behavioral solution concept or for a richer model including features such as repetition, capacity constraints or differentiated goods.

This raises the question why we should bother about the traditional price-competition model at all. The answer is that it is very difficult to get around: it is the simplest possible model of oligopolistic competition and it would be an embarrassment if game theory could not say anything sensible about it. That, however, turns out hard if one insists on basing the solution on Nash equilibrium as many textbooks witness. Consider for instance the following discussion of a price-competition duopoly where both firms have a marginal cost of 10:

¹³ See, e.g., Dufwenberg and Gneezy (2000).

¹⁴ We have restricted the above analysis to the case of two firms for which the Nash prediction is most paradoxical. The extension to more players, however, is straightforward: under Nash equilibrium market price will not exceed the second lowest reservation price while any curb set is compatible with any price between the second lowest reservation price and the lowest monopoly price.

“The Bertrand model is easy to interpret as a one-period simultaneous move game. Both firms reason: ‘If I set price at any $p > 10$, then my opponent will set price at $p - \varepsilon$, and I will sell nothing. But if I set price at $p = 10$, then I either capture the entire market or split the market 50-50. I should therefore set price equal to 10.’” (Waldman and Jensen, 2001, p. 212)

This quotation shows how difficult it is to defend the paradoxical Nash equilibrium, in particular when, as in this example, it involves weakly dominated strategies. Curb sets provide an appealing alternative to trying to justify the unjustifiable.

Appendix A

Proof of Proposition 3.1: (i) Let (σ_1, σ_2) be a Nash equilibrium. Setting $p' = m_2$ in Lemma 3.1 shows that $\sigma_1^+ < \sigma_2^+$ or $\sigma_2^+ < m_2$. If, however, $\sigma_2^+ < m_2$ and $\sigma_2^+ \leq \sigma_1^+$, then firm 2 will make a loss which cannot happen in equilibrium. Therefore $\sigma_1^+ < \sigma_2^+$ and, as a result we must have $\pi_2(\sigma_1, p_2) = 0$ for every $p_2 \in \text{supp } \sigma_2$. Hence the equilibrium is non-overlapping by Assumption A6. It is obvious that $\sigma_1^- \geq m_1(k)$. It is also clear that σ_1^+ is at most slightly above $m_2(k)$; otherwise Lemma 3.1 shows that firm 2 could make a profit by playing an appropriate mixed strategy with support in $[m_2(k), \sigma_1^+ - k]$.

(ii) It suffices to show the case where $d_1 > m$. If the equilibrium (σ_1, σ_2) is overlapping, then both firms will make a profit by Assumption A6. Thus $\sigma_2^+ = \sigma_1^+$ and, moreover, $\sigma_1^+ \geq d_1$; otherwise firm 1 would make a loss. If, however, $\sigma_1^+ \geq d_1$, then $\sigma_2^+ < \sigma_1^+$ by Lemma 3.1. Thus (σ_1, σ_2) is non-overlapping. It is also clear that the support of the winner’s strategy is at most slightly above m ; otherwise the other firm could make a profit by Lemma 3.1.

(iii) Remark 3.1 shows that, for prices above m , $\Pi_1^d(p)$, $\Pi_2^d(p)$, $\Pi_1^m(p)$, and $\Pi_2^m(p)$ all have the same sign. In view of Assumption A6, if there is a winning Nash-equilibrium strategy, it can have support only on prices where the sign is positive. Then, however, the losing firm could make a profit by imitating the winner. Consequently, the equilibrium is overlapping. Using Lemma 3.1 it is easily seen that market price is slightly above m . \square

Appendix B

Let firm 1 be located at $x = 0$ and firm 2 at $x = 1$. A consumer located at $x \in [0, 1]$ will buy from firm 1 if $v - p_1 - tx > \max \{v - p_2 - t(1 - x), 0\}$ and from firm 2 if $v - p_2 - t(1 - x) > \max \{v - p_1 - tx, 0\}$; if the net utilities from buying are negative for both firms, he will not buy. Assuming a unit-size population of consumers firm i 's demand given prices p_i and p_j then becomes:

$$D_i(p_i, p_j) = \begin{cases} 1 & \text{if } p_i \leq p_j - t \\ \frac{p_j - p_i + t}{2t} & \text{if } p_j - t < p_i < p_j + t \text{ and } p_i + p_j \leq 2v - t \\ \frac{v - p_i}{2t} & \text{if } p_j - t < p_i < p_j + t \text{ and } p_i + p_j > 2v - t \\ 0 & \text{if } p_i \geq p_j + t. \end{cases}$$

Firm i 's profit given prices p_i and p_j is given by $\pi_i(p_i, p_j) = (p_i - c_i)D_i(p_i, p_j)$. It is easily checked that π_i is strictly decreasing in p_i for $p_i \in (p_j - t, p_j + t)$ as long as $p_j > c_i + 3t$. ($d\pi_i/dp_i$ is negative for such p_i if $p_i > (p_j + c_i + t)/2$ for $p_i \in (p_j - t, p_j + t)$ which implies $p_j - t > (p_j + c_i + t)/2$ and, thus, $p_j > c_i + 3t$.) Thus $B_i(p_j) = \{p_j - t\}$ for such p_j . For $p_j \in (c_i - t, c_i + 3t)$ firm i 's unique best reply is given by the first-order condition $d\pi_i/dp_i = 0$ which implies $p_i = (p_j + c_i + t)/2$. For $p_j \leq c_i - t$, it is obvious that $B_i(p_j)$ consists of all p_i for which $\pi_i(p_i, p_j) = 0$.

References

- Basu, K. and J.W. Weibull (1991), "Strategy subsets closed under rational behavior," *Economics Letters* **36**, pp. 141–146.
- Bertrand, J. (1883), "Revue de la Théorie Mathématique de la Richesse Sociale et des Recherches sur les Principes Mathématiques de la Théorie des Richesses," *Journal des Savants* 1883, pp. 499-508.
- Blume, A., (2003), "Bertrand without fudge," *Economics Letters* **78**, pp. 167-168.

- Carlsson, H. (2010), "Desiderata for Solutions to Static Games," ...
- Dastidar, K.D. (1995), "On the existence of pure strategy Bertrand equilibrium," *Economic Theory* 5, pp. 19-32.
- Deneckere, R. J. and D. Kovenock (1996), "Bertrand-Edgeworth duopoly with unit cost asymmetry," *Economic Theory* 8, pp. 1-25.
- Dimand, R. W. and M. H. I. Dore (1999), "Cournot, Bertrand, and game theory: A further note", *Atlantic Economic Journal* 27, pp. 325-333.
- Dufwenberg, M. and U. Gneezy, (2000), "Price competition and market concentration: an experimental study," *International Journal of Industrial Organization* 18, pp. 7-22.
- Dufwenberg, M., M. H. Norde, H. Reijnierse, and S. Tijs, (2001), "The consistency principle for set-valued solutions and a new direction for normative game theory," *Mathematical Methods of Operations Research* 54, 119-131.
- Hoernig, S. H. (2007), "Bertrand Games and Sharing Rules," *Economic Theory* 31, pp. 573-585.
- Hurkens, S. (1995), "Learning by Forgetful Players," *Games and Economic Behavior* 11, pp. 304-329
- Tirole, J. (1988), *The Theory of Industrial Organization*, The MIT Press.
- Vives, X. (1999), *Oligopoly Pricing: Old Ideas and New Tools*. The MIT Press.
- Mark Voorneveld, M., W. Kets, H. Norde (2005), "An axiomatization of minimal curb sets," *International Journal of Game Theory* 33, pp. 479-490.
- Waldman, D.E. and E.J. Jensen (2001), *Industrial Organization: Theory and Practice*, Addison Wesley Longman.