DYNAMIC LIMIT PRICING: ONLINE APPENDIX

Flavio Toxvaerd*

December 1, 2016

Abstract. This appendix offers a detailed and self-contained analysis of the benchmark single-round version of the dynamic model presented in DYNAMIC LIMIT PRICING. In addition, the appendix offers a worked example and detailed discussion of different aspects of the dynamic extension omitted from the main article, such as equilibrium selection and post-entry collusion. Note that there is some overlap between this text and that in the article and that the equation numbering is independent from that of the main manuscript.

1. The Static Model

The following model is a simple version of the model of Milgrom and Roberts (1982). Consider an incumbent monopolist I and a potential entrant E. The monopolist serves a market with demand $Q(p)$ and the entrant can enter the market at cost $F>0$ to compete with the incumbent. The monopolist can be one of two types, high cost (H) or low cost (L), with probability μ and $(1-\mu)$ respectively. The incumbent knows his type, but his type is unknown to the entrant (who only knows the probability μ). Let $C_H(q)$ and $C_L(q)$ be the cost functions of H and L respectively. Denote by $\pi_i(p)$ the profit function of the incumbent of type $i=H,L$ when he sets price p. These profits are given by

$$\pi_i(p)=pQ(p)-C_i(Q(p)),\quad i=H,L \tag{1}$$

Let D_i be the duopoly profit of the incumbent of type $i=H,L$ when competing against E and let $D_E(i)$ be the duopoly profits of E when competing against the incumbent of type $i=H,L$. Denote by p_H^M and p_L^M the monopoly prices under the technologies $C_H(\cdot)$ and $C_L(\cdot)$ respectively.

Throughout, I make the following assumptions:

Assumptions

1 $C_i(q)$, $i=H,L$ and $Q(p)$ are differentiable, for $q>0$ and $p>0$ respectively.

2 $C_H'(q)>C_L'(q),\forall q\in\mathbb{R}_+$, with $C_H(0)\geq C_L(0)$.

3 $Q'(p)<0,\forall p\geq 0$.

4 $D_E(L)-F<0$.

5 $D_E(H)-F>0$.

6 $\pi_i(p)$ is strictly increasing for $p<p_i^M$ and strictly decreasing for $p>p_i^M$, $i=H,L$.

7 $\pi_i(p_i^M)>D_i$, $i=H,L$.

8 $\mu D_E(H)+(1-\mu)D_E(L)-F<0$.

*Faculty of Economics, University of Cambridge.

Assumption 2 makes precise the sense in which type L is more efficient than type H. Assumption 3 simply states that demand is downward sloping. Assumptions (4)-(5) imply that E will not enter if he knows that I is of type L, whereas he will enter if he knows that I is of type H. Thus these conditions are necessary for a separating limit price equilibrium to exist. Assumption 6 means that the incumbent's profit function is single peaked, whereas Assumption 7 ensures that entry deterrence is desirable for the incumbent, ceteris paribus. Under Assumption 8, the entrant expects to make negative profits against the incumbent if he cannot distinguish between the two types and thus stays out. This is a necessary condition for a pooling limit price equilibrium to exist.

The game between I and E is played in three stages. At the first stage, I sets a price that will serve as a signal for E of I's type. After observing the price set by I, E decides at the second stage whether or not to enter (incurring the entry fee F). Denote E's entry decision by $s_E \in \{0,1\}$, where $s_E = 0$ stands for *stay out* and $s_E = 1$ stands for *enter*. At the third stage, if E enters he will learn I's type and compete against him in complete information fashion. Both incumbent and entrant discount the future by a factor $\delta \in [0,1]$. The payoff to E is given by

$$\Pi_E(p) \equiv \begin{cases} 0 & \text{if } s_E = 0 \\ D_E(H) - F & \text{if } s_E = 1, \quad i = H \\ D_E(L) - F & \text{if } s_E = 1, \quad i = L \end{cases} \tag{2}$$

A strategy for I is a price for each of his two types, p_H or p_L, at the first stage, a price at the second stage if the entrant stays out and a quantity or price to set at the third stage if the entrant enters (depending on the mode of competition), both as functions of his type and the decisions made at the first stage. A strategy for E is a decision rule to enter or not as a function of the price set by I at the first stage and a quantity or price to set at the third stage in case he enters (again, depending on the mode of competition).

If E enters at the second stage, then at the third stage I and E play a duopoly game of complete information. Hence in any subgame perfect equilibria of the game after E's entry, I's equilibrium payoffs in the third stage are D_H or D_L. If E stays out, then I's equilibrium payoffs at the third stage are $\pi_H(p_H^M)$ or $\pi_L(p_L^M)$, depending on his type.[1] That is, the payoffs to the incumbent of type $i = H, L$ are given by

$$\Pi_i(p) \equiv \begin{cases} \pi_H(p) + \delta\pi_H(p_H^M) & \text{if } s_E = 0, \quad i = H \\ \pi_H(p) + \delta D_H & \text{if } s_E = 1, \quad i = H \\ \pi_L(p) + \delta\pi_L(p_L^M) & \text{if } s_E = 0, \quad i = L \\ \pi_L(p) + \delta D_L & \text{if } s_E = 1, \quad i = L \end{cases} \tag{3}$$

Next, I state some key definitions that will be used throughout this section. Let $\sigma \equiv (p_L, p_H, \overline{p})$ denote a triple of pure strategies of the game, i.e. a price charged by each type of I and a threshold price governing E's entry decision (details are given below). Throughout this article, attention will be restricted to pure strategy perfect Bayesian equilibria. Denote by p_H^* and p_L^* the equilibrium prices charged by the H type and the L type respectively.

Definition 1. *σ is a separating equilibrium if $p_H^* \neq p_L^*$ and a pooling equilibrium if $p_H^* = p_L^*$. σ is a limit price equilibrium if $p_H^* < p_H^M$ or $p_L^* < p_L^M$ or both.*

The aim of the analysis that follows is to characterize separating and pooling limit price equilibria of the game. Note that under the maintained assumptions, ceteris paribus, the high cost incumbent will wish to set a higher monopoly price than the low cost incumbent. Formally, the following result obtains:

[1] In the dynamic version of the model, stages one and two will together constitute a period and stage three will be a separate period.

Lemma 2. *(relative efficiency of types)*
(i) $\pi_L(p) - \pi_H(p)$ *is strictly decreasing in* p *and (ii)* $p_H^M > p_L^M$.

Proof. (i) First, note that $\pi_L(p) - \pi_H(p) = C_H(Q(p)) - C_L(Q(p))$ and thus

$$\frac{\partial}{\partial p}\left[\pi_L(p) - \pi_H(p)\right] = Q'(p)\left[C_H'(Q(p)) - C_L'(Q(p))\right] \tag{4}$$

By Assumption 3, $Q'(p) < 0$. Thus, by Assumption 2 it follows that $\frac{\partial}{\partial p}\left[\pi_L(p) - \pi_H(p)\right] < 0$.
(ii) By the definition of monopoly prices and Assumption 6, it follows that

$$p_L^M Q_L^M - C_L(Q_L^M) \;>\; p_H^M Q_H^M - C_L(Q_H^M) \tag{5}$$
$$p_H^M Q_H^M - C_H(Q_H^M) \;>\; p_L^M Q_L^M - C_H(Q_L^M) \tag{6}$$

Adding these inequalities, I obtain

$$C_H(Q_L^M) - C_L(Q_L^M) > C_H(Q_H^M) - C_L(Q_H^M) \tag{7}$$

Hence, by Assumption 2, $Q_L^M > Q_H^M$ and by Assumption 3, $p_L^M < p_H^M$ ■

An implication of this fact is that an inefficient incumbent would only set lower prices than an inefficient incumbent's monopoly price, in order to convince the entrant that it is in fact an efficient incumbent.

Perfect Bayesian equilibrium requires that beliefs be derived from Bayes' rule whenever possible. This means that one must assign beliefs after out of equilibrium (i.e. probability zero) events have been observed. For simplicity, the out of equilibrium beliefs of E will be assumed to have the following monotone structure:

$$\mu'(p) = \begin{cases} 1 \text{ if } & p \le p' \\ 0 \text{ if } & p > p' \end{cases}$$

where μ' is the probability assigned to the incumbent being of type L and p' is the L type's equilibrium strategy (i.e. either the separating price in a separating equilibrium or the common price in a pooling equilibrium).[2] That is, for any observed price above the L type's equilibrium price, the entrant will assign probability one to the incumbent being of type H. For prices below the L type's equilibrium price, the entrant will assign probability one to the incumbent being of type L.

This structure on beliefs is equivalent to a monotone decision rule for the entrant of the form

$$s_E(p) = \begin{cases} 1 \text{ if } & p > \overline{p} \\ 0 \text{ if } & p \le \overline{p} \end{cases} \tag{8}$$

for some appropriately chosen threshold price $\overline{p}$ (determined by the entrant). The equivalence is straightforward (it follows from Assumptions 3 and 4) and is shown below for each of the two types of equilibria respectively. The restriction to such monotone entry rules is routine in the literature.

1.1. Separating Limit Price Equilibria.

Characterization. In a separating equilibrium, the entrant can, by definition, infer the incumbent's type merely by observing its chosen equilibrium price. Hence assume that $p_H^* \ne p_L^*$. The best reply strategy of E in this case is to enter if $p = p_H^*$ and to stay out if $p = p_L^*$, i.e. $s_E(p_H^*) = 1$ and $s_E(p_L^*) = 0$. Therefore the H type incumbent is best off setting $p = p_H^M$,

[2] This is for simplicity only. Any off equilibrium beliefs that favor entry would do.

knowing that entry will occur in the second period, so $s_E(p_H^M) = 1$. Hence the high cost incumbent's equilibrium price is given by

$$p_H^* = p_H^M \tag{9}$$

To obtain a limit price equilibrium, it is thus required that

$$p_L^* \neq p_L^M \tag{10}$$

Next, the entrant's cutoff price can be characterized as follows:

Lemma 3. *(entrant's optimal decision)*
$\overline{p} = p_L^*$ *and* $\overline{p} < p_L^M$.

Proof. Suppose to the contrary that $\overline{p} \geq p_L^M$. Then $s_E(p_L^M) = 0$ and L is therefore best off switching from p_L^* to p_L^M, contradicting (10). Next, observe that in a separating equilibrium, $s_E(p_L^*) = 0$ (E knows that L set p_L^*) and hence $p_L^* \leq \overline{p}$. Suppose that $p_L^* < \overline{p}$. Since $\overline{p} < p_L^M$, it follows by Assumption 6 that L is better off by increasing his price from p_L^* to $\overline{p}$, which is a contradiction. Thus, $\overline{p} = p_L^*$ ■

The characterization so far of the separating equilibrium prices may be summarized in the following way:

Corollary 4. *(characterization of separating equilibria)*
(i) In any separating limit price equilibrium, $p_L^* < p_L^M$ *and (ii) in any separating equilibrium, either* $p_L^* = p_L^M \leq \overline{p}$ *or* $p_L^* = \overline{p} < p_L^M$.

These results completely characterize the entrant's equilibrium behavior. I proceed by further analyzing the incumbent's equilibrium strategy.

The Incentive Compatibility Constraints. Since $p_H^* = p_H^M$, the following incentive compatibility constraint for H should hold:

$$\Pi_H(p_H^M) \geq \Pi_H(p), \quad \forall p \tag{11}$$

This simply means that the H type's equilibrium strategy is globally optimal. Clearly, (11) holds for $p > \overline{p}$ because in this case, E enters and I can do no better than to set the monopoly price. Consider p such that $p \leq \overline{p}$. By Lemma 2, in a separating limit pricing equilibrium $\overline{p} = p_L^*$ and hence by Assumption 6 and Lemma 1 (ii) (which can be found in Appendix), it follows that $p \leq \overline{p} = p_L^* < p_L^M < p_H^M$ and thus it is sufficient to consider the following inequality:

$$\Pi_H(p_H^M) \geq \Pi_H(p_L^*) \tag{12}$$

By the definition of Π_H given in (3), (12) is equivalent to

$$\pi_H(p_L^*) \leq (1-\delta)\pi_H(p_H^M) + \delta D_H \tag{13}$$

For later reference, note that the right-hand side of (13) is strictly positive. This means that for the incentive compatibility constraint (13) to be satisfied, it is not *necessarily* the case that the H type's profits from mimicking the L type are negative. As shall be shown in Section 3, this result does not carry over to the dynamic setting.

To write (13) in terms of prices, first define the set

$$A_H \equiv \left\{p : \pi_H(p) = (1-\delta)\pi_H(p_H^M) + \delta D_H\right\} \tag{14}$$

This set is simply the set of prices for which the H type's incentive compatibility constraint is binding. Since $D_H = \pi_H(p)$ for some p, then by Assumptions 6 and 7 the set A_H is non-empty and contains at most two points. Next, define

$$\widehat{\alpha} \equiv \min A_H, \quad \widehat{\beta} \equiv \max A_H \tag{15}$$

where $\widehat{\alpha} < \infty$ and $\widehat{\beta} \leq \infty$. Hence, according to (13), p_L^* must satisfy

$$p_L^* \notin [\widehat{\alpha}, \widehat{\beta}] \tag{16}$$

For later use, note that by definition,

$$\pi_H(\widehat{\alpha}) = \delta D_H + (1-\delta)\pi_H(p_H^M) = \pi_H(\widehat{\beta}) \tag{17}$$

Observe that $p_L^* < p_L^M < p_H^M < \widehat{\beta}$. In conclusion, for the H type's incentive compatibility constraint to hold, it must be that

$$p_L^* \leq \widehat{\alpha} \tag{18}$$

In other words, in order for the high cost incumbent to be willing to tell the truth, the low cost incumbent's strategy must be sufficiently low.

I now turn to the L type. The incentive compatibility constraint for L is given by

$$\Pi_L(p_L^*) \geq \Pi_L(p), \quad \forall p \tag{19}$$

Again, this inequality simply states that the L type's equilibrium strategy is globally optimal. But the relevant alternative strategy p is only $p = p_L^M$ (because deterring entry is only optimal if it yields higher payoffs than setting the monopoly price in the first period and accommodating entry). Hence (19) becomes

$$\Pi_L(p_L^*) \geq \Pi_L(p_L^M) \tag{20}$$

By the definition of Π_L given by (3), inequality (20) is equivalent to

$$\pi_L(p_L^*) + \delta\pi_L(p_L^M) \geq \pi_L(p_L^M) + \delta D_L \tag{21}$$

Consequently,

$$\pi_L(p_L^*) \geq (1-\delta)\pi_L(p_L^M) + \delta D_L \tag{22}$$

is the relevant incentive compatibility constraint for L. Define the set

$$A_L \equiv \left\{p : \pi_L(p) = (1-\delta)\pi_L(p_L^M) + \delta D_L\right\} \tag{23}$$

Again, this set is the set of prices for which the L type's incentive compatibility constraint is binding. Since $D_L = \pi_L(p)$ for some p, then by Assumptions 6 and 7 the set A_L is non-empty and contains at most two points. Let

$$\alpha_0 \equiv \min A_L, \quad \beta_0 \equiv \max A_L \tag{24}$$

where $\alpha_0 < \infty$ and $\beta_0 \leq \infty$.

In terms of prices, the L type's incentive compatibility (22) can then be written as

$$p_L^* \in [\alpha_0, \beta_0] \tag{25}$$

where, by definition, it is the case that

$$\pi_L(\alpha_0) = \delta D_L + (1-\delta)\pi_L(p_L^M) = \pi_L(\beta_0) \tag{26}$$

This means that for the low cost incumbent to be willing to engage in costly signaling, the separating equilibrium price must be high enough. The previous results can be summarized as follows:

Proposition 5. *(characterization of separating limit price equilibria)*
Any separating limit price equilibrium is a triple $(p_H^*, p_L^*, \overline{p})$ *such that (i)* $p_H^* = p_H^M$*, (ii)* $\overline{p} = p_L^*$*, (iii)* $\alpha_0 \leq p_L^* \leq \widehat{\alpha}$ *and (iv)* $p_L^* < p_L^M$*.*

Hence, to show existence of a separating limit price equilibrium, I need to show that $\alpha_0 < \widehat{\alpha}$.[3]

Existence of Separating Limit Price Equilibria. The existence of separating limit price equilibria is now considered. Fortuitously, existence is secured under very mild conditions on the primitives of the model, as the following result shows:

Proposition 6. *(existence of separating limit price equilibria)*
Suppose that

$$\pi_L(p_L^M) - D_L > \pi_H(p_H^M) - D_H \tag{27}$$

Then $\widehat{\alpha} > \alpha_0$ *and the set of separating limit pricing equilibria is non-empty.*

Proof. From (27), (17) and (26), it follows that

$$\pi_L(p_L^M) - \pi_H(p_H^M) > \pi_L(\alpha_0) - \pi_H(\widehat{\alpha}) \tag{28}$$

Adding and subtracting $\pi_H(p_L^M)$ yields

$$\pi_L(p_L^M) + \pi_H(p_L^M) + \left[\pi_H(p_L^M) - \pi_H(p_H^M)\right] > \pi_L(\alpha_0) - \pi_H(\widehat{\alpha}) \tag{29}$$

By the definition of p_H^M, it follows that $\pi_H(p_L^M) - \pi_H(p_H^M) \leq 0$. It thus follows from (29) that

$$\pi_L(p_L^M) + \pi_H(p_L^M) > \pi_L(\alpha_0) - \pi_H(\widehat{\alpha}) \tag{30}$$

Since $\alpha_0 \leq p_L^M$, it follows by Lemma 1 that

$$\pi_L(\alpha_0) - \pi_H(\alpha_0) \geq \pi_L(p_L^M) - \pi_H(p_L^M) \tag{31}$$

Combined with (30), this implies that $\pi_H(\alpha_0) < \pi_H(\widehat{\alpha})$. Finally, $\alpha_0 \leq p_H^M$ and $\widehat{\alpha} \leq p_H^M$ and therefore it follows by Assumption 6 that $\alpha_0 < \widehat{\alpha}$ ■

Condition (27) holds for the cases of Cournot competition with linear demand and fixed marginal costs (for sufficiently high demand intercept) and Bertrand competition with or without product differentiation (see Tirole, 1988).

Equilibrium Selection. As seen above, the solution concept perfect Bayesian equilibrium fails to uniquely determine L's equilibrium price p_L^*. The reason for this lies in the arbitrariness of off-equilibrium path beliefs, which are not pinned down by Bayes' rule and on which the notion of perfectness imposes no restrictions. To get a sharp characterization of

[3] If one eliminates the limit pricing requirement, then in addition to the set $\{(p_H^*, p_L^*, \overline{p}) :$ (i), (ii), (iii), (iv) satisfied$\}$ there are other separating equilibrium points if $\widehat{\alpha} \geq p_L^M$. In particular, any separating equilibrium is a triple $(p_H^*, p_L^*, \overline{p})$ such that (i) $p_H^* = p_H^M$, (ii) $\overline{p} = p_L^*$ and (iii) p_L^* satisfies the inequality $p_0 \leq p_L^* \leq \min\{\widehat{\alpha}, p_L^M\}$.

equilibrium behavior, I therefore make use of the notion of *equilibrium dominance.* This entails using notions of both backward and forward induction. Specifically, it requires that the incumbent's equilibrium strategy (at the signaling stage) form part of a perfect Bayesian equilibrium of the game obtained after deletion of strategies that are not a weak best response to any of the entrant's possible equilibrium strategies (at the entry stage). In other words, a deviation from an equilibrium price will be interpreted as coming from the L type whenever the H type cannot possibly benefit from such a deviation (for any best response of the entrant) whereas the L type incumbent would stand to benefit from such a deviation. The reasonableness of this criterion lies in the fact that it requires the entrant to assign probability zero to a type of incumbent who would find the observed action to be dominated by the equilibrium action, irrespective of the entrant's response to such a deviation from equilibrium play. In other words, if irrespective of the entrant's response to a non-equilibrium price, one type of incumbent could not possibly benefit from such a deviation and thus earn lower payoff than by setting its equilibrium price, the entrant will disregard the possibility that the incumbent is of that type. As the following proposition shows, this refinement yields a unique equilibrium:

Proposition 7. *(uniqueness of separating limit price equilibrium satisfying dominance)*
(i) Suppose that $\alpha_0 < \widehat{\alpha} \leq p_L^M$. Then only $p_L^* = \widehat{\alpha}$ satisfies equilibrium dominance. (ii) Suppose that $\alpha_0 < p_L^M \leq \widehat{\alpha}$. Then only $p_L^* = p_L^M$ satisfies equilibrium dominance.

Proof. (i) Suppose that $\alpha_0 < \widehat{\alpha} \leq p_L^M$ and let p' satisfy $p_L^* < p' < \widehat{\alpha}$. Whichever strategy E picks, it is a strictly dominated strategy for H to choose p'. If $s_E(p') = 1$, then because $p' < \widehat{\alpha} \leq p_L^M \leq p_H^M$ it follows that

$$\pi_H(p') + \delta D_H < \pi_H(\widehat{\alpha}) + \delta D_H \tag{32}$$

If in turn $s_E(p') = 0$, then

$$\pi_H(p') + \delta\pi_H(p_H^M) < \pi_H(\widehat{\alpha}) + \delta\pi_H(p_H^M) = \pi_H(p_H^M) + \delta D_H \tag{33}$$

Hence, even if H fools E to believe that he is L, he will obtain less than $\pi_H(p_H^M)+\delta D_H$ which he would obtain under the equilibrium strategy $p_H^* = p_H^M$. In the game obtained after eliminating the strategy p' from H's strategy set, E must play $s_E(p') = 0$ because p' can have been set only by L and thus by backward induction staying out at the price p' is a best response for E. But in the new reduced game, L can profitably deviate from p_L^* to p' and obtain $\pi_L(p') - \pi_L(p_L^*) > 0$, which follows from Assumption 6 and the fact that $p' \leq p_L^M$. For completeness, note that no type of incumbent can benefit from deviations to prices such that $p' \in [\alpha_0, p_L^*]$. The proof of (ii) follows similar steps as that of (i) ■

The price selected by the equilibrium dominance approach is known as the least-cost separating equilibrium price, as it is the equilibrium price which involves the lowest possible cost for the L type in terms of foregone profits. In other words, it is the highest price (lower than the monopoly price) consistent with the incentive compatibility constraints. This outcome is known in the literature as the *Riley outcome.*

1.2. Pooling Limit Price Equilibria. In a pooling equilibrium, it is by definition the case that

$$p_L^* = p_H^* = p^* \tag{34}$$

This means that the entrant cannot infer the incumbent's type merely by observing his chosen equilibrium price. Observe first that if

$$\mu D_E(H) + (1-\mu)D_E(L) - F > 0 \tag{35}$$

then pooling equilibria cannot exist, because the expected profit of E when he cannot distinguish between the incumbent's types is positive and he thus enters regardless of p^*. By backward induction, each type of incumbent is better off setting his monopoly price. Since $p_H^M > p_L^M$, I thus have that $p_L^* \neq p_H^*$, contradicting the supposition that the two types pool. Assumption 8 rules out this case, thus ensuring that a pooling equilibrium is feasible.

Characterization. Before characterizing the incumbent's equilibrium price, the entrant's cutoff rule can be characterized in the following way:

Lemma 8. *(entrant's optimal decision)*
$\overline{p} = p^*$ *and* $p^* \leq p_L^M$.

Proof. Clearly $\overline{p} \geq p^*$. Otherwise, E's decision rule dictates entry if p^* is charged. That is, $s_E(p^*) = 1$ if $p^* > \overline{p}$ and thus each type of incumbent would benefit from deviating to their respective monopoly prices, contradicting (34). Next observe that if $p^* > p_L^M$, then L is best off setting the price p_L^M and entry will still be deterred (i.e. $s_E(p_L^M) = 0$). Consequently, $p^* \leq p_L^M$ as claimed. Finally, suppose to the contrary that $\overline{p} > p^*$. Since $p^* \leq p_L^M < p_H^M$, it follows by Assumption 6 that the H type is better off increasing his price slightly above p^* to increase profits while still deterring E's entry. Therefore $\overline{p} = p^*$ must hold as claimed ■

The Incentive Compatibility Constraints. The incentive compatibility constraints for the H type and the L type are given by

$$\pi_H(p^*) + \delta\pi_H(p_H^M) \geq \pi_H(p_H^M) + \delta D_H \quad (36)$$
$$\pi_L(p^*) + \delta\pi_L(p_L^M) \geq \pi_L(p_L^M) + \delta D_L, \quad p^* < p_L^M \quad (37)$$

Note that for each type, the best alternative strategy to choosing the entry deterring pooling price is to set the monopoly price and inviting entry. Also note that if $p^* = p_L^M$, then there is no incentive compatibility constraint for the L type.[4] The two incentive compatibility constraints (36)-(37) can be rewritten as

$$\pi_H(p^*) \geq (1-\delta)\pi_H(p_H^M) + \delta D_H \quad (38)$$
$$\pi_L(p^*) \geq (1-\delta)\pi_L(p_L^M) + \delta D_L, \quad p^* < p_L^M \quad (39)$$

Using (15) and (24), inequality (38) holds if and only if $\widehat{\alpha} \leq p^* \leq \widehat{\beta}$ whereas (39) holds for $p^* \geq \alpha_0$ as long as $p^* < p_L^M$. Combining these constraints, I obtain:

Proposition 9. *(characterization of pooling limit price equilibria)*
Any pooling equilibrium is a tuple $(p^*, \overline{p})$ *such that (i)* $\overline{p} = p^*$ *and (ii)* p^* *satisfies*

$$\max\{\alpha_0, \widehat{\alpha}\} \leq p^* \leq p_L^M < p_H^M \quad (40)$$

It should be noted that a pooling equilibrium necessarily involves limit pricing, because at least the H type (and potentially the L type) charges below his monopoly price.

Equilibrium Selection. As was the case with the set of separating limit price equilibria, there is a continuum of pooling limit price equilibria. Again, equilibrium dominance can be used to select a unique equilibrium satisfying equilibrium dominance as follows:

[4]Throughout the article, the qualifier $p^* < p_L^M$ will reappear in connection with constraints on pooling prices. It will henceforth be implicit that if $p_t^* = p_L^M$ in some period t, then there is no incentive compatibility constraint for the L type in that period.

Proposition 10. *(uniqueness of pooling limit price equilibrium satisfying dominance)*
The only pooling equilibrium limit price satisfying equilibrium dominance is $p^* = p_L^M$.

Proof. The set of pooling equilibrium prices is the set

$$\{p^* : \max\{\alpha_0, \widehat{\alpha}\} \leq p^* \leq p_L^M\} \tag{41}$$

Suppose that $p^* < p_L^M$. First note that $s_E(p_L^M) = 1$, for otherwise the L type is better off switching from p^* to p_L^M. Thus it is a strictly inferior strategy for H to select p_L^M or p_H^M. Indeed, by H's incentive compatibility constraint (36) I have

$$\pi_H(p^*) + \delta\pi_H(p_H^M) \geq \pi_H(p_H^M) + \delta D_H > \pi_H(p_L^M) + \delta D_H \tag{42}$$

Consider the new reduced game, which is obtained from the original game by eliminating p_L^M from H's strategy set. In the equilibrium of the new game, $s_E(p_L^M) = 0$, because this price can only have been set by the L type. Hence L, in the new game, is better off deviating from p^* to p_L^M ■

As was the case with the selected separating limit price equilibrium, equilibrium dominance selects the least-cost pooling limit price equilibrium.

Last, note the following result, which further reduces the set of pooling limit price equilibria satisfying equilibrium dominance:

Proposition 11. *(possible non-existence of pooling limit price equilibrium satisfying dominance)*
If $\widehat{\alpha} > p_L^M$, *then no pooling equilibrium satisfying equilibrium dominance exists.*

Proof: First, note that $\widehat{\alpha} \geq p_L^M$ if and only if

$$(1-\delta)\pi_H(p_H^M) + \delta D_H \geq \pi_H(p_L^M) \tag{43}$$

To see this, note that from (17) it follows that

$$\delta D_H = \pi_H(\widehat{\alpha}) - (1-\delta)\pi_H(p_H^M) \tag{44}$$

Substituting this in (43) yields

$$\pi_H(\widehat{\alpha}) \leq \pi_H(p_L^M) \tag{45}$$

Since $\widehat{\alpha} < p_H^M$ and $p_L^M < p_H^M$, the result then follows from Assumption 6. Next, recall that for pooling on $p^* = p_L^M$ to be incentive compatible, inequality (38) must hold, i.e.

$$\pi_H(p_L^M) \geq (1-\delta)\pi_H(p_H^M) + \delta D_H \tag{46}$$

The result then follows immediately. For the knife's edge case $\pi_H(p_L^M) = (1-\delta)\pi_H(p_H^M) + \delta D_H$, pooling on $p^* = p_L^M$ is incentive compatible ■

1.3. Existence of Limit Price Equilibria Satisfying Equilibrium Dominance. Before continuing the analysis, some comments on the existence of limit price equilibria satisfying equilibrium dominance are in order. Note that the above existence result concerns itself only with the existence of limit price equilibria and *not* with the existence of limit price equilibria satisfying equilibrium dominance. After performing equilibrium selection, the set of equilibria can, if non-empty, be divided into two distinct regimes, namely a *limit price regime* and a *monopoly price regime*. The former obtains if $\widehat{\alpha} < p_L^M$ and the latter if $\widehat{\alpha} \geq p_L^M$. These regimes will reappear in an important way in the dynamic game. In the monopoly price regime, the

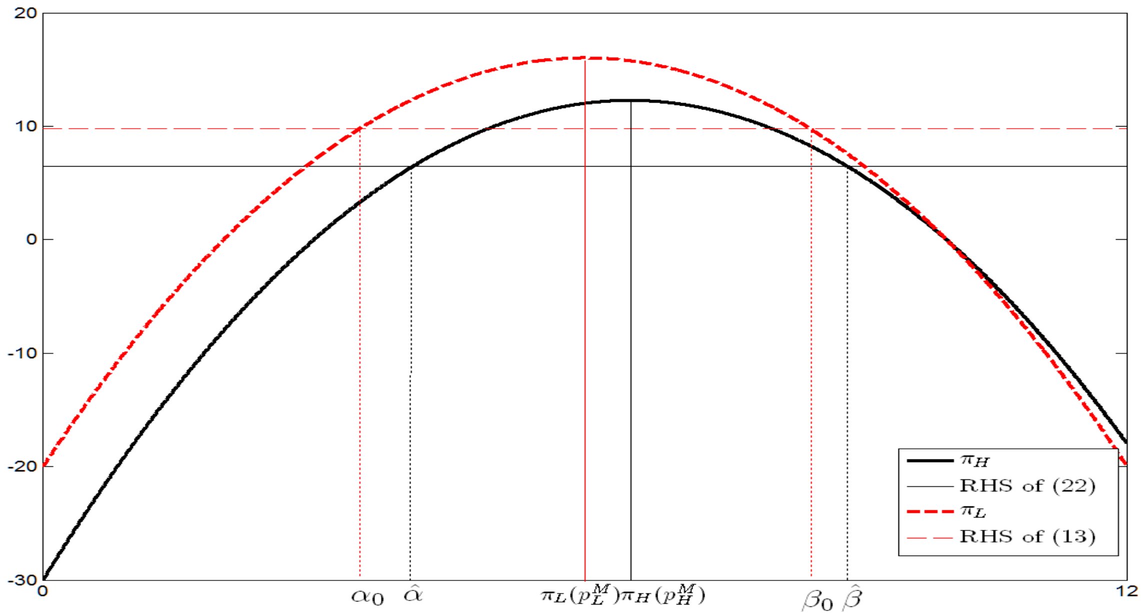

Figure 1: Profits and Incentive Compatibility Constraints in Static Setting.

unique selected equilibrium is characterized by firms separating by setting their respective monopoly prices, whereas in the limit price regime, both pooling and separating limit price equilibria satisfying equilibrium dominance coexist. Which regime obtains, depends on the parameter constellation and on the specifics of the mode of competition.

For later reference, it should be reiterated that the condition determining the regimes is given by (43). That is, the monopoly price regime obtains if and only if

$$\pi_H(p_H^M) + \frac{\delta D_H}{1-\delta} \geq \frac{\pi_H(p_L^M)}{1-\delta} \tag{47}$$

This inequality has an interesting interpretation. The left-hand side is the profit for the H type of revealing his type by earning monopoly profits in the first period and then earning discounted duopoly profits in perpetuity thereafter. The right-hand side is the discounted profit stream for the H type from mimicking the L type's monopoly price in perpetuity.

The profit functions and the incentive compatibility constraints are illustrated in Figure 1.[5]

1.4. Worked Example. In order to illustrate the analysis so far, I will next analyze a concrete functional form example. Consider the case in which the post-entry game takes the form of homogeneous goods Cournot competition. Production involves incurring fixed marginal costs but no fixed costs, so $C'_H(q) = c_H > c_L = C'_L(q)$ and $C_H(0) = C_L(0) = 0$. The entrant has fixed marginal costs $C'_E(q) = c_E$, but no fixed costs (other than the entry fee F). Market demand is given by the function $p = a - bQ$, where $a, b > 0$ and $Q = q_I + q_E$ is total quantity

[5]The graphs are based on the worked example below.

produced. With this specification, it is straightforward to verify that

$$D_E(L) = \frac{(a + c_L - 2c_E)^2}{9b} \tag{48}$$

$$D_E(H) = \frac{(a + c_H - 2c_E)^2}{9b} \tag{49}$$

$$D_L = \frac{(a + c_E - 2c_L)^2}{9b} \tag{50}$$

$$D_H = \frac{(a + c_E - 2c_H)^2}{9b} \tag{51}$$

$$\pi_L(p_L^M) = \frac{(a - c_L)^2}{4b} \tag{52}$$

$$\pi_H(p_H^M) = \frac{(a - c_H)^2}{4b} \tag{53}$$

$$\pi_L(p_H^M) = \frac{(a - c_H)(a + c_H - 2c_L)}{4b} \tag{54}$$

$$\pi_H(p_L^M) = \frac{(a - c_L)(a + c_L - 2c_H)}{4b} \tag{55}$$

The sufficient condition for a separating limit price equilibrium to exist, i.e. inequality (23), reduces to the requirement that

$$2a + 7c_H + 7c_L - 16c_E \geq 0 \tag{56}$$

This condition holds provided that the demand intercept is sufficiently high or that the entrant is not too inefficient relative to the incumbent.

Parametric Example. Consider the following benchmark parameter constellation:

a	b	c_H	c_L	c_E	F	μ	δ
10	1	3	2	2	8	0.5	0.7

It is easily verified that for this choice of parameters, both separating and pooling limit price equilibria exist (i.e. assumptions A4 and A5 and conditions (13), (22) and (35) are satisfied). Furthermore, the relevant regime is the limit price regime (because condition (43) is violated). The monopoly prices are given by $p_L^M = 6$ and $p_H^M = 6.5$ respectively. The incentive compatibility constraints are given by

$$IC_H : \quad p_L^* \notin [4.1, 8.9] \tag{57}$$

$$IC_L : \quad p_L^* \in [3.5, 8.5] \tag{58}$$

It follows that the set of separating limit price equilibria are given by pairs of prices (p_L^*, p_H^*) with

$$p_H^* = 6.5 \tag{59}$$

$$p_L^* \in [3.5, 4.1] \tag{60}$$

The least cost separating equilibrium limit price is given by $p_L^* = 4.1$. Last, the least cost pooling limit price equilibrium is given by $p^* = 6$.

2. Worked Example in Dynamic Model

To illustrate the results of the dynamic analysis, I now return to the example considered earlier. For ease of reference, the two relevant incentive compatibility constraints for the dynamic setting are reproduced below:

$$\pi_L(p^*_{t,L}) \geq \left(1 - \frac{\delta - \delta^{T-t+1}}{1-\delta}\right)\pi_L(p^M_L) + \left(\frac{\delta - \delta^{T-t+1}}{1-\delta}\right)D_L \quad (61)$$

$$\pi_H(p^*_{t,L}) \leq (1-\delta^{T-t})\pi_H(p^M_H) + \left(\frac{\delta - \delta^{T-t+1}}{1-\delta}\right)D_H - \left(\frac{\delta - \delta^{T-t}}{1-\delta}\right)\pi_H(p^M_L) \quad (62)$$

Now consider the case $T = 4$, i.e. where the static game is repeated once (as long as entry has not occurred). In this case, the incentive compatibility constraints (62) and (61) for separation at $t = 1$ are given by

$$IC_H : \quad p^*_L \notin [3, 10] \quad (63)$$

$$IC_L : \quad p^*_L \in [2.3, 9.7] \quad (64)$$

The set of separating equilibria (with separation in the first period) is in this case given by (p^*_L, p^*_H) with

$$p^*_H = 6.5 \quad (65)$$

$$p^*_L \in [2.3, 3] \quad (66)$$

It follows that when moving from $T = 2$ to $T = 4$, the range of possible limit price equilibria increases.[6] More importantly, the least cost separating equilibrium has decreased, from $p = 4.1$ to $p = 3$, in accordance with the results in the previous section.

For the purpose of illustration, I will also consider the case $T = 100$. For this long horizon, the set of separating equilibria is given by (p^*_L, p^*_H) with

$$p^*_H = 6.5 \quad (67)$$

$$p^*_L \in [-2.9, -2] \quad (68)$$

so even the least cost separating equilibrium (if it exists), will involve negative prices.

I next turn to the necessary conditions in Assumptions A4' and A8' for limit price equilibria to exist. As already noted above, these conditions hold at $t = 2$, at which point the remainder of the game coincides with the static game. But at $t = 1$ (when $T = 4$), the constraints in Assumptions A4' and A8' are in fact *violated*. This means that the entrant's expected, discounted post-entry profits are now so large, that *no amount of signaling can make entry unprofitable* and neither pooling nor separating limit price equilibria exist. Thus in this case, the unique equilibrium involves immediate separation on the monopoly prices $(p^*_L, p^*_H) = (p^M_L, p^M_H) = (6, 6.5)$ and it is in fact immaterial whether the prevailing regime is the limit price regime or the monopoly price regime. The important point to take from this example is that with this specification of post-entry competition and for this parameterization, limit pricing is feasible in the static setting, whereas under a minimal dynamic extension, it is not.

From the above analysis, it is clear that for a given set of product market parameters (in the parametric example, a, b, c_H, c_L, c_E) and parameters F, δ, μ, each of the incentive compatibility constraints and each of the many feasibility constraints deliver values of T for which the constraints hold (or do not hold, as the case may be). It is therefore tempting to try to rank such

[6] To be precise, the measure of the interval of possible prices increases, but neither interval is contained in the other.

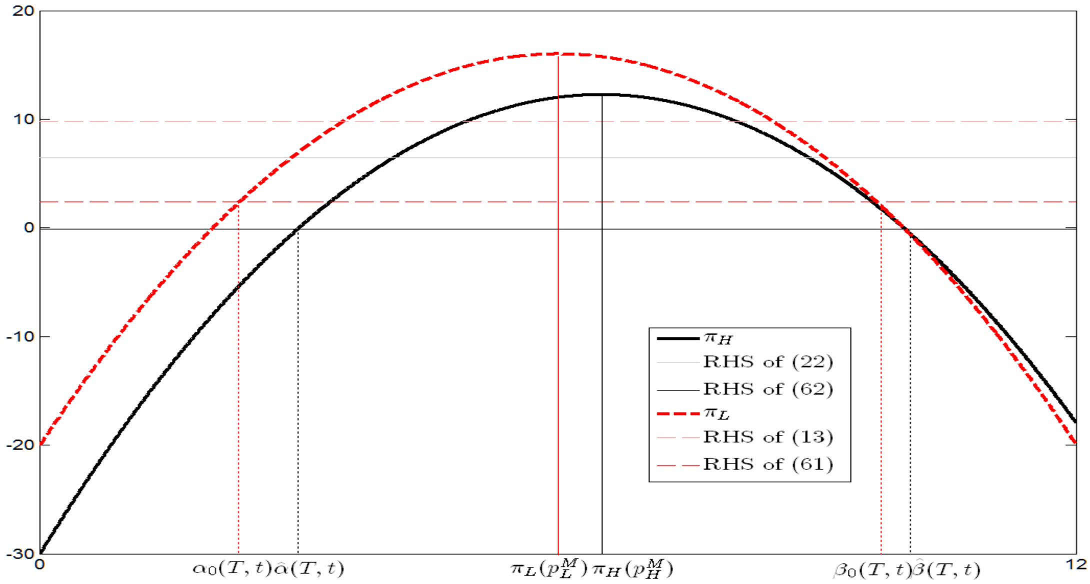

Figure 2: Profits and Incentive Compatibility Constraints in Dynamic Setting.

thresholds in a clean manner and draw general conclusions about existence of different types of equilibria for different horizons T. Unfortunately, doing this is an uphill battle, because (i) the ranking of thresholds may change depending on the choice of parameters $a, b, c_H, c_L, c_E, F, \delta, \mu$ and (ii) the constraints depend on these parameters in different and complicated ways. But as a practical matter, for a given choice of parameters, it is relatively straightforward to evaluate the different constraints and determine the existence and nature of the equilibria of the dynamic model.

3. Discussion on Equilibrium Selection in Dynamic Model

Rather than characterizing the entire set of equilibria of the dynamic game, I will argue that only a subset of the equilibria are 'reasonable' in a specific sense to be developed in further detail below. Because the literature on signaling in dynamic settings is still in its infancy, I will start by giving a brief review of it and emphasize the distinct contribution of the present analysis to that literature. In the vast majority of signaling models, there is only one instance of signaling, even if the model is otherwise dynamic. When there are multiple opportunities for the informed party to engage in signaling, the details of how (and if) private information changes over time and its observability by the uninformed party become crucial. The most conventional analyses are those of models in which private information is regenerated each period or in which the uninformed party's observations are imperfect signals of the informed party's actions. In either case, updating on the equilibrium path can always be achieved by application of Bayes' rule. articles of this type include Mester (1992) and Vincent (1998) as well as Saloner (1984), Roddie (2010) and Gedge et al. (2013) in the context of limit pricing.

When private information is perfectly persistent over time and the informed party's actions are perfectly observable, then the modeler must confront the issue of assigning out of equilibrium beliefs. There have been two different approaches to deal with such beliefs in the existing literature, namely (i) *support restrictions* and (ii) *belief resetting*. Both approaches rely on the fact that the solution concept perfect Bayesian equilibrium does not impose any restrictions on beliefs after probability zero events. In the former approach, once posterior beliefs are concentrated entirely on some state of nature, no possible observation will prompt a shift of probability towards alternative states of nature. In other words, once posterior beliefs are degenerate, the game is treated as one of perfect information regardless of how it subsequently unfolds. In the latter approach, posteriors are allowed to fluctuate over time. In particular, this approach allows beliefs that assign positive probability to events that previously were assigned zero probability.

Support restrictions have been used in different contexts by Rubinstein (1985), Grossman and Perry (1986) and LeBlanc (1992) and in a limit pricing context by Gryglewicz (2009). Although such restrictions may be perfectly appropriate for some analyses, there are instances in which they are clearly inappropriate. Madrigal et al. (1987), Noldeke and van Damme (1990a) and Vincent (1998) discuss the treatment of degenerate posteriors in depth and show that such support restricted equilibria may fail to exist.

As an alternative to support restrictions, some authors have resorted to repeated resetting of beliefs. In actual fact, beliefs are degenerate along the entire equilibrium path when employing this approach, but the equilibria are constructed *as if* beliefs are reset to their prior values.[7] This is the avenue taken by Cho (1990), Noldeke and van Damme (1990b), by Kaya (2009) in a limit pricing context and discussed by Vincent (1998). The equilibria studied in Kaya (2009) and Noldeke and van Damme (1990b) exploit the fact that the players may simply disregard the public information contained in past play and proceed "as if" they had not observed past play at all. The point here is not that the equilibria studied by these authors are not equilibria (which they clearly are). Rather, I argue that the reliance of such equilibria on the players ignoring past evidence can serve as a useful feature to help choose between different kinds of equilibria in this type of setting.

In terms of applications to economic modeling such as limit pricing, these two approaches differ radically in their predictions in that support restrictions effectively make repeated signaling impossible (by definition), whereas belief resetting allows for a potentially very rich set of equilibria, in which signaling occurs repeatedly. The precise assumptions adopted by the modeler therefore have profound implications for the analysis at hand and therefore merit scrutiny.

What support restrictions and belief resetting have in common, is that with neither approach does the observation of out of equilibrium play prompt the uninformed party to make sense of the deviation. This is at odds with the way that static signaling models are habitually analyzed. In such settings, out of equilibrium beliefs are not all treated equally, some being deemed more reasonable than others. In this way, equilibrium selection techniques are useful in that they reduce the equilibrium set significantly, sometimes even to a unique equilibrium. Hitherto, equilibrium selection techniques have not been widely applied to dynamic settings of signaling. This is unfortunate, because equilibrium selection obviates the need to choose between support restrictions and belief resetting. Furthermore, it is entirely consistent with the way that static signaling models are analyzed. The approach I adopt in the present article, is to make use of equilibrium selection reasoning in the dynamic game. Specifically, I make use of reasoning along the lines of criterion D1 against a natural benchmark equilibrium in which post separation beliefs are degenerate.

Last, it should be mentioned that some articles do feature repeated signaling without relying on the arbitrariness of out of equilibrium beliefs. These articles include Noldeke and van Damme (1990b), Bar-Isaac (2003) and Sorenson (2004). In these articles, the informed party is unable to effectively separate in a single period and is hence forced to distribute costly signaling over several periods.[8,9]

[7] This distinction is immaterial since under this approach, beliefs are not fully incorporated into subsequent behavior.

[8] In the present article, all signaling takes place in (at most) a single period, for reasons mirroring those articulated by Weiss (1983) and Noldeke and Van Damme (1990b).

[9] Note that whereas many equilibrium selection approaches rely on the arbitrariness of off-equilibrium path

To be more specific about the differences between the approaches taken in the existing literature, suppose that the incumbent's past actions have convinced the entrant that its type is L and that the incumbent subsequently charges price p_H^M. This sequence of events should confound the entrant, because an L type could have set the preferred price p_L^M without suffering adverse consequences. There are different ways to interpret the situation. One is to insist on the informational content of past behavior and to simply ascribe p_H^M to a "mistake" by the L type incumbent (which echoes the approach prescribed by the notion of subgame perfection in complete information games). This type of obstinacy in updating is the essence of support restrictions. A second way to proceed is to suppose that past behavior was in fact "mistaken" and to infer from the observation of p_H^M that the incumbent is in fact not an L type after all. But if such past behavior is ignored, then the L type incumbent should set his price such as to (again) credibly convey the information that he is in fact an L type incumbent and thus deter entry, despite the fact that it is already "known" (or has already been inferred) that he is an L type. This is exactly the way in which belief resetting makes repeated signaling possible. In the former approach, the informational content of past actions is given all weight whereas in the latter, the informational content in the incumbent's current action is given all weight.

A third approach, is to consider the two pieces of conflicting evidence together and to make sense of the conflict by using heuristics familiar to the equilibrium refinement literature. This approach consists of asking which type of incumbent, *given the belief that he is type* L, would stand to gain from setting price p_H^M? It turns out that answering this question gives a very natural prediction in this game. The key is to observe that *given* that the entrant already assigns probability one to the incumbent being an L type, the L type *cannot possibly benefit from setting any price different from* p_L^M, as long as observing p_L^M does not prompt the entrant to revise his belief that the incumbent is of type L. On the other hand, an H type incumbent *would* benefit from this price if E disregards this piece of confounding evidence (which he is entitled to do as out of equilibrium beliefs are arbitrary in a perfect Bayesian equilibrium).

Extending this reasoning to the dynamic game, the natural benchmark equilibrium price sequence after separation has occurred is $(p_L^M, p_L^M, ..., p_L^M)$. Next, *given* this benchmark equilibrium price sequence, all deviations can be dealt with by using reasoning similar to that inherent in criterion D1 of Cho and Kreps (1987). The D1 criterion works as follows. Fix some perfect Bayesian equilibrium of the game under consideration and consider a deviation by the informed party from its equilibrium strategy. Criterion D1 then dictates that if the set of responses by the entrant that makes the type $i = H, L$ incumbent willing to deviate to the observed deviation price is strictly smaller than for type $j \neq i$, then the entrant should assign infinitely larger probability to incumbent j having deviated than to incumbent i. In the reasoning above, the benchmark equilibrium was simply that in which after separation, the L type sets its monopoly price p_L^M whereas the H type, off the equilibrium path, chooses to mimic the behavior of the L type by also setting the price p_L^M in periods after separation.[10]

Note that the procedure I make use of is not quite a direct application of D1, as I do not consider an arbitrary equilibrium. Rather, the present approach accords special significance to the equilibrium in which the uninformed party at each information set makes full use of all available information (i.e. acts without ignoring available evidence). The reason that this is sensible is that in the static setting, there is no sense in which prior information "favors" any equilibrium over the other. In the dynamic setting however, prior information "suggests" or "indicates" one particular equilibrium over all other equilibria. In contrast, belief resetting amounts to *actively disregarding* the most focal equilibrium, on which beliefs are naturally

---

beliefs, some, such as those introduced in Fudenberg, Kreps and Levine (1988) are explicitly constructed to avoid this issue.

[10]That this is indeed the optimal deviation for the inefficient incumbent in the limit price regime, will be verified below.

anchored because of past play and application of Bayes' rule.

It should be noted that the equilibrium chosen by the anchored D1 criterion has a very nice property, namely that it is the equilibrium satisfying D1 which is preferred by both the incumbent and the entrant. That it satisfies D1 and is preferred by the incumbent follows from the arguments above. That it is also preferred by the entrant, can be seen by considering an alternative equilibrium which also satisfies D1. Assume that the game is in the limit price regime. In the benchmark equilibrium, the equilibrium strategies are given by (p_L^M, p_L^M) in the period after separation has occurred. If these prices do not form part of the highest payoff equilibrium satisfying D1, then there exists some other D1 equilibrium such that the H type would have strictly higher payoffs from setting an alternative price $p \neq p_L^M$. But by D1, in this alternative equilibrium it must be that $p_L^* = p_L^M$ (or else it would be better for the L type to switch to this price, as he would still be believed to be type L). But by D1, upon observing $p \neq p_L^* = p_L^M$, the entrant must assign probability one to the incumbent being type H, and thereby enter. As will be shown below, in the limit price regime, the optimal deviation has the H type incumbent set p_L^M and avoid entry (rather than setting the next best alternative p_H^M and inviting entry). In conclusion, there is no equilibrium satisfying D1 which is preferred by either the incumbent or the entrant, a feature which lends added support to the one selected by the anchored D1 criterion.

4. Infinite Horizon and the Effects of Post-Entry Collusion

In the infinite horizon setting, the possibility of post-entry cooperation (which is absent in the static setting) affects the desirability and feasibility of limit pricing in interesting ways. Before exploring these effects more systematically, it is worth recalling the tradeoffs involved for the incumbent and the entrant respectively. The reason that the incumbent may wish to engage in costly limit pricing in order to deter entry, is that post-entry competition reduces its profits. Similarly, the logic of limit pricing is that for a sufficiently low price, the incumbent convinces the entrant that its post-entry profits would be too low to offset the entry fee F. In short, the magnitude of post-entry profits for the two market participants directly affect the different incentive constraints necessary for limit pricing to be viable.

In any dynamic setting in which the final period T is finite, the post-entry payoffs $(D_i, D_E(i))$, $i = H, L$ to the incumbent and the entrant are given by the Nash equilibrium duopoly payoffs. This fact turns out to make entry deterrence through limit pricing much easier to achieve in the static setting than in the a setting with an infinite horizon. To see this, note that in the infinite horizon game, the set of payoffs that E and I can achieve in equilibrium is bounded below by the payoffs $(D_i, D_E(i))$, $i = H, L$ in the static setting. From the Folk Theorem, it is known that for a sufficiently high discount factor, any feasible and individually rational payoff vector can be sustained in equilibrium. In other words, the competing firms may be able to coordinate on payoffs that are above the levels achieved in the static setting. In turn, this directly influences the two firms' incentives to deter entry and to stay out respectively. To be specific, the higher the incumbent's (per-period) post-entry payoff D_i, the weaker is the incumbent's incentive to engage in costly signaling in the short term in order to maintain incumbency. This is seen most clearly in the L type's incentive compatibility constraint (61). As D_L increases, it becomes increasingly difficult to satisfy the constraint.[11]

Turning to the entrant, recall that the basic tradeoff influencing the entry decision is that of entry fees versus post-entry payoffs $D_E(i)$ from duopoly competition. But the higher these latter payoffs are, the more difficult is it to discourage entry, even under complete information. To see this, consider the constraint A4', which ensures that the entrant will choose to stay out against an efficient incumbent. As $D_E(L)$ increases, the constraint becomes increasingly difficult to satisfy.

[11] When D_i reaches $\pi_L(p_L^M)$, the incentive compatibility constraint can only be satisfied with equality.

In conclusion, the prospect of less than cut-throat post-entry competition, makes the incumbent more reluctant to engage in costly entry deterrence. By the same token, the prospect of higher post-entry duopoly profits enjoyed by engaging in collusion with the incumbent makes it more difficult to discourage the potential entrant from entering the industry and competing with the incumbent.

REFERENCES

[1] BAR-ISAAC, H. "Reputation and Survival: Learning in a Dynamic Signalling Model", *Review of Economic Studies*, 70 (2003), pp. 231-251.

[2] CHO, I.-K. "Uncertainty and Delay in Bargaining", *Review of Economic Studies*, 57 (1990), pp. 575-595.

[3] CHO, I.-K. AND D. M. KREPS. "Signaling Games and Stable Equilibria", *Quarterly Journal of Economics*, 102 (1987), pp. 179-221.

[4] FUDENBERG, D., D. KREPS AND D. LEVINE. "On the Robustness of Equilibrium Refinements", *Journal of Economic Theory*, 44 (1988), pp. 354-380.

[5] GEDGE, C., J. W. ROBERTS AND A. SWEETING. "A Model of Dynamic Limit Pricing with an Application to the Airline Industry", *mimeo* (2013).

[6] GROSSMAN, S. AND M. PERRY. "Sequential Bargaining under Asymmetric Information", *Journal of Economic Theory*, 39 (1986), pp. 120-154.

[7] GRYGLEWICZ, S. "Signaling in a Stochastic Environment and Dynamic Limit Pricing", *mimeo* (2009).

[8] KAYA, A. "Repeated Signaling Games", *Games and Economic Behavior*, 66 (2009), pp. 841-854.

[9] LEBLANC, G. "Signalling Strength: Limit Pricing and Predatory Pricing", *RAND Journal of Economics*, 23 (1992), pp. 493-506.

[10] MADRIGAL, V., T. C. C. TAN AND S. R. DA C. WERLANG. "Support Restrictions and Sequential Equilibria", *Journal of Economic Theory*, 43 (1987), pp. 329-334.

[11] MESTER, L. J. "Perpetual Signalling with Imperfectly Correlated Costs", *Rand Journal of Economics*, 23 (1992), pp. 548-563.

[12] MILGROM, P. AND J. ROBERTS. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis", *Econometrica*, 50 (1982), pp. 443-459.

[13] NOLDEKE, G. AND E. VAN DAMME. "Switching Away from Probability One Beliefs", *mimeo* (1990a).

[14] NOLDEKE, G. AND E. VAN DAMME. "Signalling in a Dynamic Labour Market", *Review of Economic Studies*, 57 (1990b), pp. 1-23.

[15] RODDIE, C. "Signaling and Reputation in Repeated Games", *mimeo* (2010).

[16] RUBINSTEIN, A. "A Bargaining Model with Incomplete Information about Preferences", *Econometrica*, 53 (1985), pp. 1151-1172.

[17] SALONER, G. "Dynamic Equilibrium Limit-Pricing in an Uncertain Environment", *MIT Working Paper #342* (1984).

[18] Sorenson, T. L. "Limit Pricing with Incomplete Information: Reassuring Answers to Frequently Asked Questions", *Journal of Economic Education*, 35 (2004), pp. 62-78.

[19] Tirole, J. *The Theory of Industrial Organization*, MIT Press, 1988.

[20] Vincent, D. R. "Repeated Signalling Games and Dynamic Trading Relationships", *International Economic Review*, 39 (1998), pp. 275-293.

[21] Weiss, A. "A Sorting-cum-Learning Model of Education", *Journal of Political Economy*, 91 (1983), pp. 420-442.