Patent Licensing to Bertrand Competitors*

Benny Moldovanu† and Aner Sela‡

October 2001

Abstract

A cost-reducing process innovation protected by a patent is sold to one of several firms engaged in price competition. Incomplete information about production costs yields an auction model with both private and common value components. Our main result is that standard auction mechanisms lead to inefficient allocations. This sharply contrasts with the result obtained under complete information. The inefficiency result extends to patent race frameworks which resemble all-pay auctions. An auction where the lowest bidder gets the patent is shown to have several equilibria, one of which is efficient.

*We wish to thank Simon Anderson and two anonymous referees for many helpful editorial comments. We also wish to thank Mort Kamien, Thomas Kittsteiner and Asher Wolinsky for insightful remarks.

†Department of Economics, University of Mannheim, L7, 3-5, 68131 Mannheim. mold@pool.uni-mannheim.de

‡Department of Economics, Ben-Gurion University of the Negev, P.O.B. 653, Beer-Sheva 84105, Israel. anersela@bgumail.bgu.ac.il

1. Introduction

In a classical paper, Arrow (1962) defined the value of a patent on a cost-reducing innovation as the revenue which the innovator could achieve by licensing the innovation to producing firms. Arrow also pointed out that the value depends on the downstream market structure. Extensive literature can be found on patent races, patent licensing and the value of patents. This literature is surveyed in Tirole (1988), Reinganum (1989) and Kamien (1992). In particular, Kamien's survey compares licensing procedures in various oligopolistic market structures. An important feature of most of the relevant papers is the fact that information is complete: all relevant parameters (e.g., production costs before and after an innovation is introduced) and hence downstream profits are common knowledge.

Our main goal is to investigate the effects of incomplete information[1] in a simple model of patent auctions conducted in an oligopolistic industry where firms compete in prices. We will show that this feature drastically changes the conclusions we obtain in the same model under complete information.

We consider several firms engaged in price competition under conditions of asymmetric information about production costs: this will be the "status-quo" situation. An innovation, which is sold through an auction, enables a reduction in marginal cost. The main questions we address are: Which firm will acquire the patent ? Is the allocation efficient ? What is the revenue to the seller of the patent, i.e., what is the "value" of the patent ?

The above questions are often addressed in industry studies. We now briefly

[1] Kamien (1992) concludes his survey with the following: "Obviously much remains to be done in bringing the models of patent licensing closer to reality. For example, introducing uncertainty regarding the magnitude of the cost reduction provided by an invention...."

describe two cases which differ with respect to the information available to the competing firms.

1) Competitors in the mature steel industry know each other well, and engineers often visit competitors' plants (see Ghemawat, 1997, and von Hippel, 1988). In 1983, a West-German company called SMS discovered a new process that enabled thinner casting of flat-rolled steel sheets.[2] Nucor of Charlotte, NC (who was a minion at the time) adopted the new technology in 1987. It was the first firm to do so in the U.S., and until 1995 it faced no competitors using that technology.[3] Ghemawat thinks that Nucor's first adoption of the new technology (instead of adoption by one of the large integrated mills) "appears anomalous". He writes: "Why was Nucor, rather than other competitors, the one that adopted first ?..."The most obvious reason for Nucor's adoption of thin-slab casting before other mills is that it seems to have been more efficient than the average."

2) Stobaugh (1988) reviews innovation patterns in the petrochemical industry, and focuses on nine main compounds used in fibers, plastics and rubbers. A key factor for innovation is the amount of capital available for research. There are many and frequent process innovations - Stobaugh calculates an average of one innovation every two years for each one of the nine compounds, up to 1974.[4] An important feature of the industry is its secrecy: petrochemical firms go to great

[2] All main process innovations in the steel industry were concerned with thinner casting.

[3] Subsequently, Nucor became one of the largest US steel makers.

[4] About 30% of all innovations involve radically new raw materials or operational steps. For example, acrylonitrile was first commercialized by BASF in 1933. Standard Oil of Ohio introduced a radically new production process in 1960, cut prices from 26 to 18 cents per pound, and caused widespread shutdown of existing plants. Its subsequent profits were estimated at $700 million.

length to keep proprietary information from leaking to competitors.[5]

Stobaugh finds that product innovators were almost always among the largest firms in the industry:[6] the innovating firm had, on average, 75% of the size of the largest firm in the specific line of business. In marked contrast, firms of various sizes have participated in process innovation. The corresponding averages are 60% for early process innovators, 25% for late process innovators and 10% for technology purchasers.[7] Most telling, only four out of the nine studied product innovators (which initially monopolized the market for an average period of 5.7 years) still manufacture the respective product, and none is a world leader. The reason seems to be that only one of the product innovators developed a new process for the product it originally commercialized. As we will see, this pattern agrees well with our funding below.

Our paper is organized as follows: In Section 2 we present the model with complete information, and we show that the firm with the ex-ante *lowest* cost (i.e., the ex-ante most efficient firm) attaches the highest value to the patent, and consequently bids highest for it at standard auctions. In particular, standard auction mechanisms such as the second-price and first-price sealed-bid auctions are efficient.

[5] For example, Union Carbide had a policy of not even filling process patents in an attempt to keep operations secret. Units were called by code names (a practice also used at Monsanto), and non-technical personnel did not even know what product was produced where, and under which conditions.

[6] Stobaugh compares the list of product innovators with the "Who's Who" of the petrochemical industry.

[7] Interestingly, technology sellers in the industry succeed to reap up to 75% of the estimated savings for the purchaser.

In Section 3 we introduce incomplete information: we assume that the ex-ante production costs are private information, while the percentage reduction in marginal cost due to the innovation is common knowledge.[8] This yields a model with *interdependent values*. In subsection 3.1 we show that standard auctions which award the patent to the highest bidder are not efficient:[9] in any equilibrium of such an auction there will be instances where the patent is not bought by the firm with the ex-ante lowest cost, contradicting the efficiency criterion for the complete information case. We compute inefficient equilibria for the simpler case where the post-innovation cost is equal for all firms. These equilibria usually involve pooling. In subsection 3.2 we construct an efficient auction. Interestingly, such an auction awards the patent to the lowest bidder! In subsection 3.3 we briefly discuss R&D races for patents. In Section 4 we gather several concluding comments. In particular, we address a possible application of our model to takeover contests.

2. Licensing under complete information

Consider first a Bertrand duopoly in a market with inelastic demand, normalized to be equal to one unit. There are 2 firms producing a homogenous good. Firm i produces with constant marginal cost c_i. The firms simultaneously choose prices

[8] For example, imagine a new device that performs twice as fast as older models.

[9] Jehiel and Moldovanu (2001) discuss various types of inefficiencies arising for **any** allocation procedure if signals are multidimensional. In the present model signals are one-dimensional, and efficient allocation procedures do exist (see below).

p_1 and p_2. Sales for firm i are given by

$$x_i(p_i, p_j) = \begin{cases} 1, \text{ if } p_i < p_j \\ \frac{1}{2}, \text{ if } p_i = p_j \\ 0, \text{ if } p_i > p_j \end{cases}$$

and profits are given by

$$\pi_i(p_i, p_j) = x_i(p_i, p_j)(p_i - c_i).$$

Assume now that an inventor uses an auction[10] to sell a cost-reducing technical innovation protected by a patent. We assume that only one firm will be licensed.[11] The technical innovation allows firm i to produce with a marginal cost equal to $\min(c_i, \alpha c_i + c)$, where $0 \leq \alpha \leq 1$, and $c < \min(c_1, c_2)$. This parametrization incorporates the cases of a percentage reduction in marginal costs $(c = 0)$, and a common cost $(\alpha = 0)$.

We are interested in two main questions: Which firm will buy the patent? Which price will the buyer pay, i.e., what is the "value" of the patent?

Definition 2.1. *The value of the patent for a firm is the difference between the profit it makes in case it acquires the patent, and the profit in case it does not. An auction is said to be efficient if it awards the patent to the firm with the highest value.*

Proposition 2.2. *The firm with the ex-ante lowest cost attaches the highest value to the patent and obtains it at a standard auction.*

[10] We focus here on standard first- and second-price sealed-bid auctions. Other mechanisms such as the Dutch or English auctions yield here the same results.

[11] This is, in fact, optimal for an inventor facing Bertrand competitors (see Kamien, 1992).

Proof. Assume without loss of generality that $c_2 < c_1$. We offer here the proof for the case where $\alpha < \frac{c_2 - c}{c_1}$, so that both firms attach a positive value to the patent. In the other case, firm 1 attaches a zero value to the patent, and the proof is simpler.

The losing firm always makes zero profits. If firm 1 gets the patent, then it competes with cost $\alpha c_1 + c$ against firm 2 that has cost $c_2 > \alpha c_1 + c$. This implies that firm 1 will serve the market by charging a price p_1 which is slightly lower than c_2.[12] Hence, the value of the patent for firm 1 is simply $c_2 - \alpha c_1 - c$. Similarly, the value of the patent for firm 2 is given by $c_1 - \alpha c_2 - c$. Since $c_1 - \alpha c_2 - c > c_2 - \alpha c_1 - c$, firm 2 (i.e., the firm with ex-ante lower costs) attaches a higher value to winning, and it will consequently win the patent by bidding $c_2 - \alpha c_1 - c$ in a sealed-bid first-price auction or $c_1 - \alpha c_2 - c$ in a sealed-bid second-price auction.[13] In both cases, the revenue to the seller is equal to $c_2 - \alpha c_1 - c$. *Q.E.D.*

It is straightforward to show that Proposition 2.2 generalizes to the case of more than 2 bidders: the firm that stands to gain most from acquiring the patent is always the firm that will face in the post-innovation world the least efficient best competitor. Thus, the firm that bids most is the ex-ante most efficient firm, and we have the following:

Corollary 2.3. *Under complete information, standard auctions are efficient.*

[12] Asymmetric Bertrand games (and first-price auctions) have no equilibria in pure strategies here, but introducing a smallest money unit immediately yields the intuitive solution. Other modeling approaches use sharing rules for the case where both firms charge the same price. For example, Lederer and Hurter (1986)use a "lowest cost" sharing rule that could be applied here. This technical problem usually disappears under conditions of incomplete information with continuous types.

[13] The arguments for English or Dutch auctions are completely analogous.

An explanation about the use of the term *efficiency* is necessary: since consumers do not participate at the auction, the first natural step is to judge the allocation procedure according to its properties for the involved agents (firms and patent holders). Of course, a procedure that creates a high value for those agents may be less desirable for consumers. To see this, consider again the set-up of Proposition 2.2: if $c_2 < c_1$, the innovation does not affect the market price, which remains fixed at c_1. But consumers prefer the alternative where the less efficient firm 1 gets the patent, in which case the price goes down to c_2. In our model demand is inelastic. Assuming that consumers value the product at v, the price is a linear transfer from consumers to firms, and it can immediately be seen that awarding the patent to the ex-ante most efficient firm also maximizes total welfare. In more general models of oligopoly our measure of efficiency (which does not include consumers) need not coincide with maximization of total welfare. But, as long as the market price depends on the distribution of costs in the oligopoly, an allocation procedure that awards the patent to an arbitrary firm (as is the case in the pooling equilibria displayed below) cannot maximize total welfare.

3. Licensing under incomplete information

We now allow for incomplete information about ex-ante production costs. While the complete-information treatment fits well mature industries, the framework below is better suited for emerging or very dynamic and secretive industries where there is still considerable uncertainty about competitors. In this context, we show that efficiency-based explanations offer no indication about the identity of the firm that is likely to acquire the innovation.

The model is as follows: There are n firms competing in a Bertrand oligopoly. Firm i's marginal cost in the status-quo, c_i , is *private information* to that firm. All firms j, $j \neq i$, believe that c_i is distributed on the interval $[c^L, c^H]$ according to density $f > 0$, and independently of other costs.

A cost-reducing technical innovation allows production with marginal cost equal to min $(c_i, \alpha c_i + c)$,where, $0 \leq \alpha \leq 1$ and $c < c^L$ are common knowledge. An independent inventor sells the innovation through an auction. Note that the value of the patent for firm i is determined by two factors: 1) i's ex-ante marginal cost (through the term αc_i); 2) the ex-ante cost of the most efficient competitor (which is not known to i at the time of the auction). Hence, our model is not a *private values model* but rather displays a feature called *interdependent values* in the auction literature.

In order to avoid signalling issues, we look at the case where the true production costs of the bidding firms are revealed after the auction, i.e., the posterior Bertrand game is conducted under complete information. It is known that signalling through bids in order to "manipulate" competitors' beliefs in a downstream interaction game may cause inefficiencies (see Das Varma (2000) and Goeree (2000)).[14]

[14]Das Varma studies a private values model where the ex-ante costs are common knowledge and equal (hence the firms are symmetric). If firm $i, i = 1, 2$ wins the patent, it produces with cost $c - \theta_i$, where θ_i is private information. Goeree (2000), expanding on Jehiel and Moldovanu (1996, 2000), studies an auction setting where the downstream interaction is modeled in reduced form.

3.1. Inefficiency in standard auctions

In many basic auction settings the efficiency properties that hold for auction procedures under complete information continue to hold under incomplete information, at least as long as the bidders are ex-ante symmetric[15] (see Myerson, 1981 or Milgrom and Weber, 1982). Our model is ex-ante symmetric, and, given the result for the complete information case, the efficiency of standard auctions seems intuitive. Two conditions must be fulfilled in order to obtain equilibrium efficiency (i.e., a patent award to the ex-ante most efficient firm) for any realization of firms' marginal costs[16]:

1. The equilibrium strategies must be symmetric.

2. The equilibrium strategies must be strictly monotonically decreasing in cost.

The following result shows that the second-price auction does not possess an equilibrium with the above properties. The proofs for other standard auctions (e.g., first-price, all-pay, English, Dutch, etc...) are completely analogous and use the respective first-order conditions.

Proposition 3.1. *The second-price auction is not efficient.*

[15] Ex-ante symmetry requires here that all firms have the same utility function, and all costs are drawn from the same distribution. It is well known that ex-ante asymmetries lead to inefficiencies (see Myerson, 1981).

[16] Note also that any two procedures that yield the same physical allocation are, up to a constant, revenue-equivalent. Revenue equivalence is well known in settings with private values (i.e., in situations where the valuation of one agent does not depend on information available to other agents). Our model has interdependent values, but revenue equivalence continues to hold as long as types are independent (for a formal result, see for example Fieseler, Kittsteiner and Moldovanu, 2000)

Proof. Assume that $\alpha < \frac{c^L - c}{c^H}$, so that the winning firm is sure to have the lowest cost. Assume that a symmetric equilibrium in strictly monotonically decreasing strategies exists, and denote the common bidding function by b.[17]

Assume that firm i with cost c_i bids x. The maximization problem for firm i reads:

$$\max_x \int_{b^{-1}(x)}^{c^H} \left((c_j - \alpha c_i - c) - b(c_j)\right)(n-1)(1-F(c_j))^{n-2} f(c_j) dc_j$$

where b^{-1}denotes the inverse function of b. The first-order condition is:

$$-(b^{-1})'(x)(n-1)(1-F(b^{-1}(x))^{n-2} f(b^{-1}(x)) \left(b^{-1}(x) - \alpha c_i - c - b(b^{-1}(x)\right) = 0.$$

Since in a symmetric equilibrium we must have $x = b(c_i)$, the unique solution to the above equation is $b(c_i) = (1-\alpha)c_i - c$. Observe that the candidate equilibrium bidding function $(1-\alpha)c_i - c$ is strictly monotonically **increasing,** which is a contradiction to our assumption. Therefore, the second-price auction does not have a symmetric equilibrium in strictly monotonically decreasing strategies, and thus it cannot be efficient. *Q.E.D.*

The intuition behind the above result is as follows:[18] Consider for simplicity the case $\alpha = 0$. If a symmetric separating equilibrium is played, then, on the margin, a low-cost firm expects to outbid a rival who also has a low cost. But this means that, conditional on winning, the price she will be able to charge in the market (and hence her market profit) will be low. Therefore, on the margin, a low cost firm has a lower willingness to pay than a firm with high cost, which can expect to charge a high market price. Note that such an argument cannot

[17] Note that any strictly monotonic function is differentiable almost everywhere. In order to shorten the argument we assume below that b is differentiable everywhere.

[18] We wish to thank a referee for describing this intuitive explanation.

work in a private values model where the value conditional on winning does not depend on competitors' characteristics.

Since standard auctions (or other similarly structured methods where the highest bidder wins) are nevertheless likely to be employed in practice, it is of interest to compute their equilibria. The task is complex since the first-order approach is not useful here, and the structure of equilibria must be guessed. We display below inefficient equilibria for the simpler model where $\alpha = 0$, i.e., where the ex-ante costs are private information, but the winner's post-innovation cost is given by the constant c.

Under the assumption that a separating symmetric equilibrium exists in the second-price auction, we have shown in the proof of Proposition 3.1 (see also the intuition following it) that the only candidate for a separating equilibrium has to be an increasing bidding function, namely $b_i(c_i) = (1-\alpha)c_i - c$. For $\alpha = 0$ this becomes $b_i(c_i) = c_i - c$.

Proposition 3.2. *The strategy profile where each firm bids according to $b_i(c_i) = c_i - c$ constitutes an equilibrium of the second-price auction*[19] *only for the case of two bidders.*[20]

[19] For $n = 2$, we can easily use revenue equivalence in order to compute an inefficient separating equilibrium for the first-price sealed-bid auction: $b_i(c_i) = E[c_{-i} - c \mid c_{-i} \leq c_i]$, $i = 1, 2$, where $-i$ is the firm other than i, and where E denotes the expectation according to f.

[20] This construction works also for the case where $c^L \leq c \leq c^H$. The strategy profile $b = (b_1(c_1), b_2(c_2))$ where $\forall i, i = 1, 2,$ $b_i(c_i) = c_i - c$ for $c_i \geq c$, and $b_i(c_i) = 0$ for $c_i < c$, is an (inefficient) equilibrium

Proof. For a given realization of cost parameters, let $i^* = \arg\max_i c_i$. Assume that all firms $i \neq i^*$, bid according to $b_i(c_i) = c_i - c$. By bidding $b = c_{i^*} - c$, firm i^* wins the patent and pays $\max_{i \neq i^*} c_i - c$. In the subsequent Bertrand competition this firm makes a profit of $\min_{i \neq i^*} c_i - c$. Hence the overall payoff for i^* is given by $(\min_{i \neq i^*} c_i - c) - (\max_{i \neq i^*} c_i - c) = \min_{i \neq i^*} c_i - \max_{i \neq i^*} c_i$. It is clear that for $n > 2$ the last term is negative for almost all realizations. Hence, a bid $b = 0$ is more advantageous for i^* than the bid $b = c_{i^*} - c$. For $n = 2$ we have $\min_{i \neq i^*} c_i - \max_{i \neq i^*} c_i = 0$, hence firm i^* is indifferent to whether it wins or loses and the bid $c_{i^*} - c$ is optimal. It is easy to see then that the proposed strategy profile is optimal also if a firm has the higher cost. *Q.E.D.*

Proposition 3.3. *Consider the strategy profile where* $\forall i, b_i(c_i) = b^* = E[\min_{j \neq i} c_j - c]$. *This strategy profile constitutes an equilibrium*[21] *for both a first-price and a second-price auction. For* $n > 2$, *this is the unique symmetric equilibrium.*

Proof. Assume that all firms $i, i \neq i^*$, use this strategy. If firm i^* bids b^*, then its overall expected payoff is $\frac{1}{n}(E[\min_{j \neq i} c_j - c] - b^*) = 0$. Bidding $b < b^*$ yields a zero expected payoff, while bidding $b > b^*$ yields $E[\min_{j \neq i} c_j - c] - b < 0$. Hence a bid of b^* is optimal.

Pooling cannot be sustained at a level $\widetilde{b} \neq b^*$ since: 1) If $\widetilde{b} > b^*$, the firms make negative profits and prefer to bid zero; 2) If $\widetilde{b} < b^*$, each firm makes an expected profit of $\frac{1}{n}(E[\min_{j \neq i} c_j - c] - \widetilde{b}) = \frac{1}{n}(b^* - \widetilde{b}) > 0$, and each firm has an incentive to always win the patent by bidding slightly higher.

[21] Since ties appear in equilibrium with positive probability (in fact with probability 1), we assume below that each of the tied winners gets the patent with the same probability.

The proof that for $n > 2$ a symmetric equilibrium cannot have strictly monotonic intervals follows by combining arguments similar to those in Propositions 3.1 and 3.2. *Q.E.D.*

To summarize, the main conclusion of this section is that in the presence of incomplete information it is hard to make a determinate prediction about the identity of the firm that will acquire an innovation through standard auctions. Efficiency-based explanations lose their predictive power.

3.2. An efficient equilibrium

The above results reveal that standard auctions will generally not award the patent to the firm that values it most. Is there any mechanism that achieves this goal ?

Consider the sealed-bid auction where the bidder with the **lowest bid** obtains the patent and pays the next higher bid. The strategy profile where all firms bid zero (and hence no revenue is created) is an obvious pooling equilibrium. But, we also have the following rather intriguing result:

Proposition 3.4. *Assume that* $\alpha < \frac{c^L - c}{c^H}$ *, and consider the sealed-bid auction where the bidder with the* ***lowest bid*** *obtains the patent and pays the next higher bid. The strategy profile where,* $\forall i$*,* $b_i(c_i) = (1-\alpha)c_i - c$ *constitutes an equilibrium.*[22] *This equilibrium leads to an efficient allocation.*

Proof. Assume that all bidders besides firm j use the above strategy. Assume first that $j = \arg\min_i c_i$. Then, by bidding $(1-\alpha)c_j - c$, j gets the patent, pays $(1-\alpha)\min_{i\neq j}(c_i) - c$, produces with cost $\alpha c_j + c$ and charges $\min_{i\neq j}(c_i)$. This

[22] The equilibrium is not in dominant strategies. Nevertheless, the equilibrium does not depend on the function governing the distribution of costs.

yields for j a payoff of

$$\min_{i\neq j}(c_i) - (\alpha c_j + c) - ((1-\alpha)\min_{i\neq j}(c_i) - c) = \alpha(\min_{i\neq j}(c_i) - c_j) \geq 0$$

It is clear that j cannot improve her payoff by bidding less. If j bids more, then either the outcome does not change, or j loses the patent, yielding utility of zero.

Assume now that $j \neq \arg\min_i c_i$. Then, by bidding $(1-\alpha)c_j - c$, firm j does not win, and has a payoff of zero. Obviously, j cannot improve her payoff by bidding more. In order to change the outcome, j must bid $b \leq (1-\alpha)\min_{i\neq j} c_i - c$. In that case, j gets the patent, pays $(1-\alpha)\min_{i\neq j} c_i - c$, produces with cost $\alpha c_j + c$, and charges $\min_{i\neq j} c_i$. This yields a payoff of

$$\min_{i\neq j} c_i - (\alpha c_j + c) - ((1-\alpha)\min_{i\neq j} c_i - c) = \alpha(\min_{i\neq j} c_i - c_j) \leq 0$$

Hence bidding $(1-\alpha)c_j - c$ is also optimal in this case.

Because equilibrium strategies are increasing and because the patent goes to the lowest bidder, the firm with the ex-ante lowest cost gets the patent, as required. *Q.E.D.*

Since in our model marginal costs are drawn independently of each other, the above result, together with revenue equivalence, can be used to compute efficient equilibria of other mechanisms such as the one where the lowest bidder gets the patent, but pays her own bid. An efficient equilibrium also exists for the procedure where a price-clock goes down until a unique firm remains active, in which case the clock stops and the remaining firm wins the patent and pays the price where the clock stopped.

For the case where $\alpha = 0$ we can compare the seller's expected revenue in the efficient equilibrium of Proposition 3.4 with her expected revenue in the standard auctions. Since the respective physical allocations are different, revenue equivalence cannot be applied, and the result is not immediate.

Proposition 3.5. *The seller's revenue in the efficient equilibrium of Proposition 3.4 is higher than the expected revenue in standard auctions.*

Proof. Denote by E_l^k the expectation of the $k-$th order statistic out of l identical and independently distributed random variables which govern the firms' cost distributions. The seller's expected revenue in the efficient auction[23] is given by $E_n^2 - c$. The seller's expected revenue in the pooling equilibrium of a standard auction[24] is given by $E_{n-1}^1 - c$. The result follows by observing that $E_n^2 > E_{n-1}^1$.[25] For the case $n = 2$ we can also consider the inefficient separating equilibrium of Proposition 3.2, which yields an expected revenue of $E_2^1 - c$. Since $E_2^1 < E_2^2$, the result follows for this case as well. *Q.E.D.*

Note that both the pooling equilibrium where all firms bid zero (and where revenue is also zero), and the separating, efficient equilibrium involve very risky strategies: the winner's bid is only a lower bound for the price she will pay.[26] For these reasons, the above revenue comparison cannot constitute a practical recommendation for choosing among auction formats.

[23] Recall that each firm i bids $b_i(c_i) = c_i - c$, and the lowest bidder pays the next higher bid.

[24] Recall that each firm i bids $b^* = E[\min_{j \neq i} c_j - c]$.

[25] This follows by stochastic dominance arguments (see Theorem 1.B.16 in Shaked and Shantihikumar, 1994.)

[26] We wish to thank a referee for mentioning the zero-bid equilibrium and the risks associated with the present auction format.

3.3. R&D Races

Another model that can be treated within our framework is that where several firms are engaged in an R&D race. Each firm bears the cost of investment in R&D, but only one of them succeeds to get a patent on a cost reducing innovation. In real-life settings, some random process usually influences the identity of the innovator (besides the firms' characteristics and the respective amounts spent on R&D). If the influence of this random process is strong enough, one should expect some variation with regard to the innovator's identity, and theoretical predictions may be problematic. But even if one simply assumes that the firm spending the largest amount on R&D gets to innovate, our model predicts that variations will occur. Indeed, such a situation is isomorphic to a first-price sealed-bid all pay auction in which every agent submits and pays a bid for the item being sold, while only the highest bidder receives the item.

By an argument which is completely analogous to the one in Proposition 3.1, in equilibrium the innovator cannot always be the ex-ante most efficient firm. For the special case where $\alpha = 0$, the strategy profile where $b_i(c_i) = b^* = \frac{1}{n}E[\min_{j \neq i} c_j - c]$ constitutes an inefficient pooling equilibrium. The argument is analogous to the one in Proposition 3.3.[27] Thus, even without a random element it is impossible to make a determinate prediction about the identity of the firm that will emerge as winner of the R&D race: all firms spend the same amount, independently of their ex-ante productivity.

Recall here Stobaugh's findings about the petrochemical industry. The large variance of the types of firms that participated in process innovation seems to fit

[27] If we interpret this as an all-pay auction organized by a patent holder, note how revenue equivalence holds among the present procedure and the second-price (or first-price) auction.

well with our above results: when costs are private information, pooling equilibria may prevail, and the identity of the innovator cannot be a-priori determined.

4. Concluding Remarks

We have studied competition over a patent whose acquisition allows a decrease in production cost. We have shown that, in general, standard auction mechanisms are inefficient under conditions of incomplete information. Consequently, efficiency-based arguments cannot determine the identity of the innovating firm. Our result was obtained in a model[28] that abstracted from the possibility that bids at the auction may serve as signalling devices.

The large body of work on vertical relations has emphasized the allocative externalities arising in situations where a principal contracts with downstream competing firms. Our simple model draws attention to informational externalities. Since the value of a specific bilateral contract between the principal and one of the firms will necessarily depend on the other firms' characteristics, a departure from the *private values* paradigm is needed. The impact of incomplete information on welfare, and the interplay between allocative and information aspects seem to us interesting topics in many vertical relations.

[28] One can extend the model in order to allow for an elastic demand function (for example by introducing some form of imperfect substitutability among the firms' products). Besides the issues discussed in this paper, existence of efficient auction equilibria in such models will also depend on whether the firms' actions are strategic complements or substitutes.

5. References

Arrow, K., 1962. Economic Welfare and the Allocation of Resources for Invention. In: the Rate and Direction of Inventive Activity: Economic and Social Factors, NBER Conference no. 13, Princeton: Princeton University Press.

Das Varma, G., 2000. Bidding for a Process Innovation under Alternative Modes of Competition. Discussion paper, Duke University.

Fieseler, K. , Kittsteiner, T. and Moldovanu, B., 2000. Lemons, Partnerships, and Efficient Trade. Discussion Paper, University of Mannheim.

Goeree, J., 2000. Bidding for the Future. Discussion paper, University of Virginia.

von Hippel, E., 1988. *The Sources of Innovation.* New York: Oxford University Press.

Jehiel, P. and Moldovanu, B., 1996. Strategic Non-Participation. *RAND Journal of Economics* 27, 84-98.

Jehiel, P. and Moldovanu, B., 2000. Auctions with Downstream Interaction among Agents. *RAND Journal of Economics* 31, 768-791.

Jehiel, P. and Moldovanu, B., 2001. Efficient Design with Interdependent Valuations. *Econometrica* 69, 1237-1259.

Kamien, M., 1992. Patent Licensing. In: R. Aumann and S. Hart (eds.), *Handbook of Game Theory.* Amsterdam: North-Holland.

Lederer, P. and A. Hurter, Jr., 1986, Competition of Firms: Discriminatory Pricing and Location. *Econometrica* 54, 623-640.

Milgrom P., Weber, R.J., 1982. A Theory of Auctions and Competitive Bidding. *Econometrica* 50, 1089-1122.

Myerson, R., 1981. Optimal Auction Design. *Mathematics of Operations Research* 6, 58-73.

Reinganum, J., 1989. The Timing of Innovation. In: R. Schmalensee and R. Willig (eds.), *Handbook of Industrial Organization.* Amsterdam: North-Holland.

Shaked, M. and Shanthikumar, G., 1994. *Stochastic Orders and Their Applications.* San Diego: Academic Press