MARKETS FOR PARTIALLY CONTRACTIBLE KNOWLEDGE: BOOTSTRAPPING VERSUS BUNDLING

James J. Anton
Duke University

Dennis A. Yao
Harvard University and
University of Pennsylvania

Abstract
We discuss how a seller can appropriate rents when selling knowledge that lacks legal property rights by solving either an expropriation or a valuation problem and then analyze how seller rents increase when a portion of the intellectual property (IP) can be protected. The analysis shows that a sequential strategy in which the protected portion of the IP is sold prior to selling the unprotected IP is superior to selling both portions of the IP as a bundle. (JEL: D23, L14, O31)

1. Introduction: Selling Ideas without Legal Property Right Protection

A fundamental problem in market transactions involving intellectual property (IP) is that a buyer's willingness to pay for IP depends on knowledge about the property's value and that value may be difficult to determine absent an expropriable disclosure of the IP (Arrow 1962). Legal institutions such as patents, copyright, and trade secret laws provide some protection against expropriation, but protection is minimal in many settings.[1] This paper reviews ways in which a seller can appropriate rents in those no-property-right settings and then discusses how these insights illuminate our understanding of partial-property-right settings.

Consider the sale of an idea in the absence of legal property rights. Establishing the value of the idea may be impossible without revealing the idea and then being expropriated without compensation. What can a seller do to appropriate rents? Analyses of this problem can be categorized into approaches that focus on expropriation incentives and on the valuation problem.

Acknowledgments: We thank Patrick Rey for helpful comments and the Fuqua Business Associates Fund, the Reginald Jones Center at Wharton, and the GE Fund for financial support.

E-mail addresses: Anton: james.anton@duke.edu; Yao: dyao@hbs.edu

1. See Arora, Fosfuri, and Gambardella (2001) for a general discussion of technology sales.

Journal of the European Economic Association April–May 2005 3(2–3):745–754

1.1. *Expropriation Incentives and the Blackmail Model*

Expropriation-incentive approaches begin by exchanging the underlying valuation problem for an expropriation problem: the seller gives up knowledge and relies on the receiver's incentives not to expropriate. Expropriation incentives are the focus of Anton and Yao (1994), Baccara and Razin (2003), and Biais and Perotti (2003).

As an example, consider our blackmail model (Anton and Yao 1994) in which an inventor would like to sell IP to one (or both) of two buyers who compete in an output market. The value of the IP is assumed to be binary and is private information of the seller. No demonstration possibilities exist that would not also give away the IP and the seller has no legal property rights. Profits associated with the market use of the invention are observable and contractible.[2]

We demonstrate that a seller can appropriate a significant share of an invention's underlying value by freely revealing the invention to a buyer. Free revelation solves the valuation (adverse selection) problem but introduces an expropriation problem. Although the buyer can now costlessly expropriate the invention, the seller can still leverage its knowledge by credibly threatening to reveal the invention to a competitor. The buyer then offers an inventive payment structured to prevent revelation to a competitor, thereby avoiding having a "monopoly" situation turn into a "duopoly" situation.

The structure of off-equilibrium-path incentives that discourage expropriation is key to the "blackmail" analysis. An inventor, having been expropriated, might simply reveal to the second firm out of spite. While this action has a certain biblical appeal ("neither mine nor thine") it provides no economic basis for understanding what would happen if, instead of outright expropriation, the first firm offers a contract with a minimally attractive expected payment. The economic credibility of our blackmail threat can be understood in terms of gains to trade between the seller and the second buyer. Whenever the first buyer offers a contract with an expected payment that falls short of duopoly profits (approximately), the seller and the second firm can design a contract that captures the remaining positive gains to trade. This contract employs the structure of the prior contract (if any) to lead the seller to reveal to the second firm while also screening out the bad seller type.

In related work, Baccara and Razin (2003) model development as requiring cooperation of at least two parties who bargain over the split of the expected development rents. Expropriation opportunities determine bargaining threats and the proposer (seller) joins with a previously uninformed party (buyer) rather than with

2. The underlying contracts are based on the market outcomes with innovation but not on the amount of IP transferred. With payments contingent on the IP transferred, the parties can effectively create private property rights.

an informed party. Choosing an uniformed party increases information diffusion, weakens the expropriation threat, and gives the proposer a greater payoff.

In Biais and Perotti (2003) an inventor undertakes a project only with a positive signal from each of two experts. Expert evaluation requires disclosure of the underlying idea, so the problem is to avoid expropriation of the idea by an expert. Biais and Perotti show that expropriation incentives are mitigated by the difficulty one expert has in avoiding expropriation by the second expert who may infer the first expert's signal by the request for an evaluation.[3]

1.2. Valuation and the Partial Disclosure of Enabling Information

Expropriation can be mitigated by withholding some or all of the underlying knowledge and relying on signaling to establish value.[4] Under this valuation-focused approach the buyer pays to obtain the withheld knowledge whereas under the expropriation-incentive approach the buyer pays the seller not to reveal the information to others. We focus on solutions that signal value by disclosing a portion of the IP that gives enabling knowledge to the potential buyer.[5] This line of research was pioneered by Bhattacharya and Ritter (1983), who modeled the impact of an enabling disclosure on the capital cost paid by a firm to develop the invention.[6]

Anton and Yao (2002) applies the disclosure approach to the sale of IP. In that model the amount of wealth a seller can put at risk determines the division of available rents. Larger wealth allows for a greater payment wedge between the high-profit (market success) and no-profit payoffs offered in the buyers' contracts and relaxes the disclosure incentive constraints between higher and lower seller types. Thus, sellers with more wealth gain a larger share of the rent, while sellers with no wealth are pushed to zero rent and full disclosure. By contrast, the blackmail approach allows a seller with near-zero wealth to obtain a significant rent. In fact, seller wealth can even be used against the seller.

3. Finally, models that examine imitation made possible in the course of licensing negotiations such as Gallini and Wright (1990) and Gans and Stern (2000) also can be interpreted as having important expropriation incentive elements.

4. A related problem involves getting the attention of the buyer. See Dewatripont and Tirole (2003).

5. Enabling disclosures are commonly used in practice, in part, because the signal is direct and less open to misunderstanding. It is possible to avoid enabling disclosures by structuring an ex ante contract with payments contingent on market outcomes. Ex ante contracting exploits common knowledge about the proportion of sellers offering positive value and is somewhat artificial given relatively free entry of sellers with no value. Such contracts result in meager payments when there is a high proportion of sellers offering no value and the seller has little capital that can be risked to help assure performance.

6. See also d'Apresmont, Bhattacharya, and Gerard-Varet (2000).

2. Partial Contractibility, Bundling, and Bootstrapping

Consider a case in which one portion of the IP may be protectible (contractible) while another is not. For example, an idea may consist of a portion that is patentable (e.g., in biotechnology, a host cell and vector) and a portion that is not (e.g., the process technology).

Given this unevenness in protection, should an inventor sell the IP sequentially in two parts or bundle the parts together? One possible advantage with a sequential sale is that proceeds from the prior sale may relax the wealth constraint in the subsequent sale. A disadvantage is that the initial sale will create asymmetric bidders, thereby weakening the competitive pressure in the downstream auction. We find that a sequential sale leads to a higher seller payoff than a bundled sale.

2.1. The Bootstrapping Model

We extend the partial disclosure model of Anton and Yao (2002) by allowing the seller to have protected IP as well as unprotected IP. In this bootstrapping model, a seller, S, has IP which is indexed by a probability of success, θ, and two buyers, $i = 1, 2$, have IP of $\alpha > 0$. This IP is an input to a potential innovation by the buyers. For simplicity, we assume an IP structure in which the probability of success from combinations of IP is equal to the largest of the components (e.g., if a buyer with IP α learns $\theta > \alpha$, then the buyer's probability of success equals θ). A unique success (state M) has a payoff of Π and 0 to the successful and unsuccessful firms, respectively. Otherwise (state 0), both earn 0.

The game form for the base model is as follows:

(1) The seller privately observes $\theta \in [\underline{\theta}, \bar{\theta}]$ according to a c.d.f. F with mean μ.
(2) The seller chooses an IP disclosure $r \leq \theta$. This disclosure is observed and can be freely used by both buyers.
(3) After the disclosure, each buyer offers the seller a contract (R^i_M, R^i_0). The payoffs must satisfy $R^i_M \geq -L$ and $R^i_0 \geq -L$ for $i = 1, 2$, where L is the wealth of the seller.
(4) The seller accepts up to one contract and then chooses how much remaining IP, t_i, to reveal where $r \leq t_i \leq \theta$, for $i = 1, 2$.
(5) The innovation outcome for each i is realized given the underlying IP input, $\max\{\alpha, t_i\}$, and is independent for $i = 1, 2$.

Now also suppose the seller has a fully protected know-how advantage of $\beta > 0$ relative to the buyers (i.e., no buyer can use this know-how without a license from the seller) so that $\theta \in [\alpha + \beta, \bar{\theta}]$ where $\beta < \bar{\theta} - \alpha$, and that this is common knowledge. Effectively, $\underline{\theta} = \alpha + \beta$. Further, assume that β does not block the firms from using non-β know-how that they have (α) or later acquire ($\theta - \beta$) and that β is only available in a discrete chunk.

The bootstrapping game has two stages: an initial auction in which the protected IP of β is sold and a subsequent (asymmetric) auction in which the remaining IP is sold. In Stage 1 the seller offers β to buyers in a lump-sum, first-price auction. P_β denotes the equilibrium price. The winning buyer, henceforth designated by i, gains exclusive use of β, while the losing buyer, j, is precluded from using β.

In Stage 2, where the remaining θ is sold, the wealth of the seller changes from initial wealth L_0 to $L_1 = L_0 + P_\beta$ as a result of the initial β auction, and the subsequent auction is conducted between asymmetric bidders i and j. We solve for a separating Perfect Bayesian Equilibrium.[7]

2.2. Analysis of the Bootstrapping Model

We begin with Stage 2 and work backwards to Stage 1. The analysis of Stage 2 follows Anton and Yao (2002), except that the auction is held between asymmetrically positioned bidders i (can use β) and j (cannot use β). In equilibrium, the bidder with the lower willingness to pay (WTP) for the incremental unprotected IP is indifferent between winning and losing the Stage 2 auction. Although contract offers differ, each bidder offers the same expected value to the seller. Bidder i has the higher WTP and is always selected.

The disclosure function is set to eliminate the incentives of any θ type to misrepresent itself through its disclosure signal and receive a contract offer for a different type. Let φ_{bs} denote the bootstrapping disclosure strategy. Firm j's offer is structured to eliminate the incentive of the seller to deviate from the equilibrium disclosure and accept a contract offer from j instead of from i.

Suppose that $r = \varphi_{bs}(\theta) > \alpha + \beta$ has been disclosed and θ inferred. If i offers a contract with $R^i_M \geq R^i_0$ and S accepts, then i earns $\theta[1 - (r - \beta)](\Pi - R^i_M) - (1 - \theta[1 - (r - \beta)])R^i_0$ while j earns $(r - \beta)(1 - \theta)\Pi$, as S will reveal θ fully and exclusively to i. Similarly, if j offers a contract with $R^j_M \geq R^j_0$ and S accepts, then j earns $(\theta - \beta)(1 - r)(\Pi - R^j_M) - [1 - (\theta - \beta)(1 - r)]R^j_0$ while i earns $r[1 - (\theta - \beta)]\Pi$. Then, calculating the WTP for each buyer, we have $\mathrm{WTP}^i = (1 + \beta)(\theta - r)\Pi > (1 - \beta)(\theta - r)\Pi = \mathrm{WTP}^j$ and, thus, the buyer who acquired β has the greater incentive to acquire the seller's remaining know-how

7. Here buyers cannot infer the seller's type from the decision of the seller to opt for bootstrapping versus bundling. While the Proposition 2 result that all seller types prefer bootstrapping over bundling is suggestive that a pooling equilibrium exists, analysis of the signaling possibility requires an examination of underlying beliefs.

in Stage 2. $\text{WTP}^i > \text{WTP}^j$ holds because β provides i with a head start for the state M payoff. We then have

LEMMA 1. *Consider a disclosure equilibrium and suppose $r = \varphi_{bs}(\theta)$ has been disclosed by the seller. Then, for $\theta > r > \alpha + \beta$, the contracting stage satisfies (i) the buyers i and j offer contracts such that $R_M^i \geq R_0^i$ and $R_M^j \geq R_0^j$ and*

$$\theta[1-(r-\beta)](R_M^i - R_0^i) + R_0^i = (\theta - r)(1-\beta)\Pi$$
$$= (\theta - \beta)(1-r)(R_M^j - R_0^j) + R_0^j;$$

(ii) the seller accepts (R_M^i, R_0^i) and reveals fully and exclusively to i; and (iii) the payoffs are $\Pi^S = (1-\beta)(\theta - r)\Pi$ for S, $\Pi^i = \theta[1-(r-\beta)]\Pi - \Pi^S$ for i, and $\Pi^j = (r-\beta)(1-\theta)\Pi$ for j.

Buyer i contracts with the seller while Buyer j sets the price as $\Pi^S = \text{WTP}^j$.

We now construct the disclosure equilibrium and solve for the disclosure function. Set the state 0 payment from i to the minimum value, $R_0^i = -L_1$, as this corresponds to the Pareto-dominant disclosure equilibrium for the seller. (See Anton and Yao 2002 for details.) In equilibrium, a seller of type θ earns $\Pi^S(\theta) = (1-\beta)[\theta - \varphi_{bs}(\theta)]\Pi$. At a feasible deviation disclosure $\hat{r} = \varphi_{bs}(\hat{\theta})$, where $\theta \geq \hat{r}$, the seller can accept the offer of $(\hat{R}_M^i, -L_1)$ from i and earn

$$U(\theta, \hat{\theta}, \hat{r}, -L_1) = \theta(1-(\hat{r}-\beta))\hat{R}_M^i - (1-\theta(1-(\hat{r}-\beta)))L_1$$
$$= \frac{\theta}{\hat{\theta}}(1-\beta)(\hat{\theta}-\hat{r})\Pi - L_1\left(1 - \frac{\theta}{\hat{\theta}}\right)$$

upon substitution with Lemma 1 (i). Incentive compatibility, $\Pi^S(\theta) \geq U(\theta, \hat{\theta}, \hat{r}, -L_1)$, then implies a differential equation for φ_{bs}: $d(\varphi_{bs}(\theta)/\theta)/d\theta = -L_1/[(1-\beta)\Pi\theta^2]$ and $\varphi_{bs}(\alpha+\beta) = \alpha+\beta$. The bigger the wedge between R_M^i and $-L_1$, the easier it is for incentive compatibility to be met as higher types have a greater probability weight on the R_M^i term than do lower types. Note also that increasing r for a given θ lowers the WTP of the bidders and lowers seller rents. This dissipation feature of r is critical for achieving incentive compatibility. Solving the differential equation gives Proposition 1.

PROPOSITION 1. *Suppose $L_1 < (\alpha+\beta)(1-\beta)\Pi$. An equilibrium with partial disclosure exists and is given by*

$$\varphi_{bs}(\theta) = \frac{L_1}{(1-\beta)\Pi} + \left[1 - \frac{L_1}{(\alpha+\beta)(1-\beta)\Pi}\right]\theta$$

for $\alpha + \beta \leq \theta \leq \bar{\theta}$, contract offers from i and j of $R_0^i = -L_1$ and $R_0^j = -\alpha L_1/(\alpha+\beta)$ with R_M^i and R_M^j as implied by Lemma 1, and the seller accepting the offer from i.

As β rises (or as L_1 falls), φ_{bs} shifts up; however, L_1 is endogenously related to P_β, via $L_1 = L_0 + P_\beta$, due to the proceeds from the Stage-1 sale of β to i. The associated equilibrium payoffs are $\Pi^S(\theta) = (\theta/(\alpha + \beta) - 1)L_1$ for S, $\Pi^i(\theta) = \theta[1-(\varphi_{bs}(\theta)-\beta)]\Pi - \Pi^S(\theta)$ for i, and $\Pi^j(\theta) = (\varphi_{bs}(\theta)-\beta)(1-\beta)\Pi$ for j.

In Stage 1, buyers bid to acquire the β know-how. The price must make the (ex ante symmetric) buyers indifferent between becoming bidder i or j. Thus, $P_\beta = E_\theta\{\Pi^i(\theta) - \Pi^j(\theta)\}$, where E_θ is the expectation over θ. Substituting for $\Pi^i(\theta)$ and $\Pi^j(\theta)$ and simplifying yield $P_\beta = \beta\Pi[1+(L_1/((1-\beta)\Pi))(\mu/(\alpha+\beta)-1)]$. Then $L_1 = (L_0+\beta\Pi)/[1-(\beta/1-\beta)(\mu/(\alpha+\beta)-1)]$ upon using $L_1 = L_0+P_\beta$. To satisfy the existence condition in Proposition 1 we assume $L_0 < (\alpha - \beta\mu)\Pi$.

2.3. Comparison of the Bootstrap Model to a Bundled Model

We now adapt the base model to a bundled sale in the partial contractibility setting. The payoffs to the winner and loser of the bundled auction are $\theta[1-(r-\beta)](\Pi - R_M) + [1-\theta[1-(r-\beta)]]L_0$ and $(r-\beta)(1-\theta)\Pi$, respectively. Now the buyers are symmetric at the contracting stage. In contrast to the sequential sale of IP, the revelation of incremental IP following a disclosure is bundled with the protected IP. Proceeding as before, the equilibrium seller payoff and disclosure strategy, respectively, are given by $\Pi^s(\theta) = [\theta - (\varphi_{bu}(\theta) - \beta)]\Pi$ and $\varphi_{bu}(\theta) = \beta + L_0/\Pi + [(\alpha - L_0/\Pi)/(\alpha+\beta)]\theta$. A direct comparison of payoffs leads to Proposition 2.

PROPOSITION 2. *Each seller type θ receives a higher payoff from the sequential sale of IP in the bootstrap model than from the bundled sale of IP. Further, the payoff difference is strictly increasing in θ and strictly positive for even the lowest seller type.*

The lowest type, $\underline{\theta} = \alpha+\beta$, has a strict preference for the sequential sale. This type has no incremental unprotected IP to sell. In the bundled sale, the disclosure equilibrium separates the lowest type and then buyers bid only for the protected β. In the bootstrap case when buyers bid initially for β they also bid to assume the stronger role of buyer i in the subsequent stage. Since P_β is an average value across types, this effect benefits even the lowest seller type.

An earlier sale of β improves the wealth position of the seller for the subsequent sale which reduces the amount of disclosure needed to satisfy incentive compatibility. Bidding pressure, however, is weaker with the asymmetric buyers. This weaker pressure narrows the contract payment spread and makes separation more difficult, pushing toward a larger disclosure (compare φ_{bs} and φ_{bu} at a common L). In equilibrium, the wealth effect dominates.[8]

3. Applications Involving the Arrow Problem

A critical concern of knowledge-driven firms is the loss of valuable IP through employee departure. Two important categories related to this problem involve employees departing with private knowledge (e.g., an invention they have discovered but have not yet disclosed) to, say, a start-up firm and employees departing with a firm's private knowledge. The former category is a direct application of the basic seller-buyer problem with the employee as seller with an option for (less efficient) self-production and the employer as an efficient "buyer." This problem has been analyzed by Anton and Yao (1995), Anand and Galetovic (2000), and Hellmann (2004), among others. The other category includes Rajan and Zingales (2001) and Hellmann and Perotti (2004), where employee expropriation concerns are modeled as drivers of organizational structure. These papers can be viewed as exploring how different expropriation-incentive environments affect employee or employer payoffs.

Another interesting topic is the impact of intermediaries in the facilitation of knowledge development and innovation. In the IP sale context, intermediaries such as a venture capitalist can serve the valuation function while having a weaker incentive to expropriate than a direct buyer. Biais and Perotti's (2003) model of experts has some elements of such a model, though there is no role in that model for signals of the underlying know-how as a means to avoid expropriation. This is a fruitful area for further work.

The existence of some property right protection raises the question of how endogenous choice involving the use of property rights affects expropriation incentives or valuation. Gallini and Wright (1990) can be interpreted along these lines. Their three-period model involves an ex ante license contract that creates a contractual reward structure which alters the future expropriation incentives of the licensee. Also, Bhattacharya and Guriev (2004) analyze how the choice of selling with and without patents affects expropriation incentives which in turn affects seller payoffs and R&D incentives and Anton and Yao (2004) explore how eschewing confidentiality rights signals a seller's underlying know-how.

8. The dominance of the wealth effect also implies that the seller benefits from an inability to extract value from the stronger buyer in the bidding. If the seller could commit to a handicap in the asymmetric auction, each seller type would earn strictly less because such a commitment reduces each buyer's willingness to pay for β, thereby undermining the relative value of the sequential sale approach.

Appendix

Proof of Lemma 1. This is similar to the proof of Lemma 1 in Anton and Yao (2002), except that we must account for the asymmetric payoffs to buyers i and j. This implies that if the seller is indifferent between two contract offers, then the seller must choose i with probability one. Otherwise, i can profitably offer a slight payment increase to attract the seller. □

Proof of Proposition 1. This is analogous to the proof of Proposition 1 in Anton and Yao (2002), in that we simply verify the equilibrium. The only significant difference is that the contract offer from j, (R_M^j, R_0^j), is off the equilibrium path since the seller accepts the offer from i with probability one. To analyze the incentives, suppose θ discloses (a feasible) $\hat{r} = \varphi(\hat{\theta})$, accepts an offer from j where $\hat{R}_M^j \geq \hat{R}_0^j$, and reveals fully and exclusively to j. This yields

$$U^j(\theta, \hat{\theta}, \varphi(\hat{\theta}), \hat{R}_0^j) = (\theta - \hat{\theta})\left[\frac{\Pi^S(\hat{\theta}) - \hat{R}_0^j}{\hat{\theta} - \beta}\right] + \Pi^S(\hat{\theta}),$$

and then incentive compatibility requires that $\Pi^S(\theta) \geq U^j(\theta, \hat{\theta}, \varphi(\hat{\theta}), \hat{R}_0^j)$, which implies $(\theta - \hat{\theta})\left[-\alpha L_1 - (\alpha + \beta)\hat{R}_0^j\right] \leq 0$. Then, as this must also hold for $\hat{\theta}$ deviations to $r = \varphi(\theta)$, where, $\varphi(\theta) \leq \hat{\theta} < \theta$, we conclude that the above expression must be zero, which yields $\hat{R}_0^j$ as in Proposition 1. □

Proof of Proposition 2. We compare the two equilibrium payoffs at each θ. In the bootstrap equilibrium, a seller of type θ has payoffs across the M and 0 states of $\hat{R}_M^i(\theta) + L_1$ with probability $\theta[1 - (\varphi(\theta) - \beta)]$ and 0 with the complementary probability for an expected payoff of $\Pi_{bs}^S(\theta) = L_1\theta/(\alpha + \beta)$. For the benchmark case, we have $R_M(\theta) + L_0$ with probability $\theta[1 - (\varphi_{bu}(\theta) - \beta)]$ and 0 with the complementary probability, and expected payoff of $\Pi_{bu}^S(\theta) = (\beta\Pi + L_0)\theta/(\alpha + \beta)$. The comparison reduces to $L_1 \gtreqless \beta\Pi + L_0$. Since $L_1 = (\beta\Pi + L_0)[1 - (\beta/(1 - \beta))(\mu/(\alpha + \beta) - 1)]^{-1}$ and the denominator lies between zero and 1, we have $L_1 > \beta\Pi + L_0$ and the result is established with strict inequality for all $\theta \in [\alpha + \beta, \bar{\theta}]$. □

References

Anand, Bharat N. and Alexander Galetovic (2000). "Weak Property Rights and Holdup in R&D." *Journal of Economics and Management Strategy*, 9, 615–642.

Anton, James J. and Dennis A. Yao (1994). "Expropriation and Inventions: Appropriable Rents in the Absence of Property Rights." *American Economic Review*, 84, 190–209.

Anton, James J. and Dennis A. Yao (1995). "Start-ups, Spin-offs, and Internal Projects." *Journal of Law, Economics & Organization*, 11, 362–378.

Anton, James J. and Dennis A. Yao (2002). "The Sale of Ideas: Disclosure, Property Rights, and Incomplete Contracts." *Review of Economic Studies*, 69, 513–531.

Anton, James J. and Dennis A. Yao (2004). "Attracting Skeptical Buyers." Working paper, Duke University.

Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella (2001). *Markets for Technology: The Economics of Innovation and Corporate Strategy.* MIT Press.

Arrow, Kenneth J. (1962). "Economic Welfare and the Allocation of Resources for Inventions." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by Richard R. Nelson. Princeton University Press.

d'Aspremont, Claude, Sudipto Bhattacharya and L-A Gerard-Varet (2000). "Bargaining and Sharing Innovative Knowledge." *Review of Economic Studies*, 67, 255–271.

Baccara, Mariagiovanna and Ronny Razin (2003). "From Thought to Practice: Appropriation and Endogenous Market Structure with Imperfect Intellectual Property Rights." Working paper, New York University.

Bhattacharya, Sudipto and Sergei Guriev (2004). "Knowledge Disclosure, Patents, and Optimal Organization of Research and Development." Working paper, Princeton University.

Bhattacharya, Sudipto and Jay R. Ritter (1983). "Innovation and Communication: Signalling with Partial Disclosure." *Review of Economic Studies,* 50, 331–346.

Biais, Bruno and Enrico Perotti (2003). "Entrepreneurs and New Ideas." Working paper, CEPR.

Dewartripoint, Mathias and Jean Tirole (2003). "Modes of Communication." Working paper, ECARES.

Gallini, Nancy T. and Brian D. Wright (1990). "Technology Transfer Under Asymmetric Information." *RAND Journal of Economics,* 21, 147–160.

Gans, Joshua S. and Scott Stern (2000). "Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction." *Journal of Economics and Management Strategy,* 9, 485–511.

Hellmann, Thomas (2004). "When Do Employees Become Entrepreneurs?" Working paper, Stanford Graduate School of Business.

Hellmann, Thomas and Enrico Perotti (2004). "Circulation of Ideas: Firms versus Markets." Working paper, Stanford Graduate School of Business.

Rajan, Raghuram G. and Luigi Zingales (2001). "The Firm as a Dedicated Hierarchy: A Theory of the Origin and Growth of Firms." *Quarterly Journal of Economics,* 116, 805–851.