NBER WORKING PAPER SERIES

SHIRKING, SHARING RISK, AND SHELVING:
THE ROLE OF UNIVERSITY LICENSE CONTRACTS

Marie Thursby
Jerry Thursby
Emmanuel Dechenaux

Working Paper 11128
http://www.nber.org/papers/w11128

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2005

We thank Ajay Agrawal, Ashish Arora, Irwin Feller, Josh Lerner, Yann Meniere, Fiona Murray, Richard Nelson, Arti Rai, and participants of workshops at Emory University, Harvard Business School, the University of California at Los Angeles, the NBER Entrepreneurship and Higher Education Meetings, the International Industrial Organization Conference at Northwestern University, and the Georgia Tech Roundtable for Research on Engineering Entrepreneurship for comments. We gratefully acknowledge financial support from the Alan and Mildred Peterson Foundation through the Purdue University Technology Transfer Initiative. Thursby and Thursby acknowledge support from the National Science Foundation (SES0094573) and Kauffman Foundation. Corresponding author. Department of Economics, Kent State University, Kent, OH 44242. Phone: (330) 672-1106. Fax: (330) 672-9808. E-mail: edechena@kent.edu.

Shirking, Sharing Risk, and Shelving: The Role of University License Contracts
Marie Thursby, Jerry Thursby, and Emmanuel Dechenaux
NBER Working Paper No. 11128
February 2005, Revised March 2007
JEL No. D82,L14,O3

ABSTRACT

In this paper, we develop a theoretical model of university licensing to explain why university license contracts often include payment types that differ from the fixed fees and royalties typically examined by economists. Our findings suggest that milestone payments, annual payments, and consulting are common because moral hazard, risk sharing, and adverse selection all play a role when embryonic inventions are licensed. Milestones address inventor moral hazard without the inefficiency inherent in royalties. Royalties are optimal only when the licensee is risk averse. The potential for a licensee to shelve inventions is an adverse selection problem which can be addressed by annual fees if shelving is unintentional, but may require an upfront fee if the firm licenses an invention with the intention to shelve it. Whether the licensing contract prevents shelving depends in part on the university credibly threatening to take the license back from a shelving firm. This supports the rationale for Bayh-Dole march-in rights but also shows the need for the exercise of these rights can be obviated by contracts.

Marie Thursby
College of Management
Georgia Institute of Technology
800 West Peachtree street, NW
Atlanta, GA 30332-0520
and NBER
marie.thursby@mgt.gatech.edu

Jerry Thursby
Dept. of Economics
Emory University
Atlanta, GA 30322
jthursb@emory.edu

Emmanuel Dechenaux
Kent State University
Department of Economics
Kent, OH 44242-0001
edechena@kent.edu

1 Introduction

In this paper we address the disparity between existing theories of licensing which focus on simple upfront fees and royalties and recent empirical evidence on contracts. Several surveys show that additional fees, such as those paid annually as well as fees paid when technical or business milestones are reached, are common [Edwards *et al.* 2003, Thursby *et al.* 2001, and Agrawal and Cockburn 2006]. For example, in the Thursby *et al.* survey of 62 US university technology transfer offices, 97% of the respondents reported that royalties were included in license contracts either "almost always" or "often," 92% reported the same for upfront fees, 89% for annual payments, and 72% for milestone payments. Existing theories cannot explain either the prevalence of annual payments and milestone payments or the multiple payment types observed.

We construct a model of university patent licensing in order examine when, if ever, it is optimal to include annual payments and/or milestone payments in a license contract. This is an ideal setting in which to examine these contract terms, as commercialization of university inventions often requires further effort by both the inventor and licensing firm, neither of which is observable. The need for inventor effort presents a moral hazard problem because inventors may "shirk" if they prefer research to development. The need for licensee effort suggests a problem with adverse selection since firms may "shelve" inventions either because their intent in licensing is simply to block other firms from developing them or, more innocently, because by the time development is completed expected profits are less than originally anticipated. While moral hazard has been studied, albeit in the context of royalties or equity [Jensen and Thursby 2001], shelving by licensing firms has not.[1] Since the intent of laws allowing university ownership of inventions, such the US Bayh-Dole Act and Bayh-Dole-inspired legislation in Europe [Verspagen 2006], is to promote commercialization, this is an important oversight.

We consider the problem of a Technology Licensing Office (TLO) which has the responsibility for designing and offering an exclusive license contract to a firm that has expressed interest in developing an invention owned by the university. The invention requires further inventor and licensee collaboration in technical development, as well as licensee investment in commercialization. Importantly, we consider contract terms with possible payment types based on events or quantities the TLO can observe (contract acceptance, technical and commercial success, output). In this context, we analyze the role that different payment types play in extracting rents as well as providing appropriate incentives to the inventor and the firm.

Assuming risk neutrality, if incentives provision is not a concern for the TLO, we

[1] Macho-Stadler *et al.* [2006] examine shelving by the TLO as a device to improve the quality of inventions licensed.

show that the contract that maximizes the TLO's payoff consists of a simple upfront fee that extracts the firm's rent. On the other hand, if inventor effort is not observable, but absent of any concern with shelving, the TLO's optimal offer includes a simple milestone payment, which serves both to provide incentives to the inventor and extract rents from the firm. Hence, when the licensee is genuinely interested in developing the invention, risk aversion on the part of the TLO and the firm are necessary for optimal contracts to include an upfront fee and a royalty, although it is not sufficient. With risk aversion, an optimal contract may thus comprise an upfront fee, a milestone payment and an output royalty.

We then extend the basic model to account for the fact that in the course of development, a firm may learn that while the invention may be useful for other firms, it is no longer useful to it. We show that in equilibrium, annual fees are included in contracts to induce these firms to return the license. We also introduce the possibility of a firm licensing the invention with the intent to shelve rather than develop it. Hence, upfront fees may serve a different purpose here than they do in an environment in which shelving is not regarded as a problem (where they merely extract rents or spread risk). It follows from our analysis that even when all parties are risk neutral, shelving concerns prompt the need for including an upfront fee, a milestone payment and an annual fee together in the same contract.

These results contribute to the extensive theoretical literature on licensing which has focused primarily on simple contracts involving fixed fees and royalties, with little attention paid to milestones (see Kamien [1992] for a review). Exceptions are Arora [1995] and Bousquet *et al.* [1998] discuss the role of state-contingent fees in the transfer of tacit knowledge and risk sharing, respectively. By largely ignoring issues related to licensing of inventions that require further development, this literature is unable to explain the complications that arise in university licensing. For example, Bousquet *et al.* [1998] who consider milestones impractical for risk sharing since there are no development milestones in their model. By contrast, when early stage inventions are licensed, milestones are not only feasible but also may be easier to define than royalties.[2] Furthermore, while earlier work showed that royalties can address inventor moral hazard, in our model milestone payments serve this function without the inefficiency of royalties [Jensen and Thursby 2001, Macho-Stadler *et al.* 1996 and Choi 2001].

Perhaps our most novel results pertain to the adverse selection problem that results from shelvers having private information about their intent. Prior studies find that licensor and/or licensee private information can be used to justify royalties (see Gallini and Wright [1990] and Beggs [1992]). We show that with shelving, fees paid upfront, annually, or with achievement of milestones address the problem while royalties

[2] Milestones are often as simple as the licensee having a business plan. Other milestones include development benchmarks such as clinical trials.

exaggerate it.

Our work is also related to the literature on the organization of R&D with incomplete contracts. The work closest to ours is that of Aghion and Tirole [1994a, b] which examines conditions under which an invention should be owned by the research unit, final customer, or some combination. They derive conditions under which ownership is irrelevant for efficiency. One of the conditions is whether the invention could be developed independently by the research unit or the customer. The types of inventions that we model are those which cannot be independently developed by either the university inventor or the firm. Moreover, in their model the final customer has no incentive to prevent development of the invention. We contribute to this literature by showing that contingent ownership combined with appropriate contract terms is important when shelving is a possibility.

Finally, we contribute to the literature on university licensing and associated public policy concerns (see Agrawal [2001] and Thursby and Thursby [2003] for reviews). The Bayh-Dole Act includes "march-in rights" which allow the US government to take ownership from the university in the absence of reasonable commercialization efforts, making university ownership state-contingent. The fact that the federal government has never exercised these rights has contributed to the view that the law should be strengthened [Rai and Eisenberg 2003]. We contribute to the debate by showing that contract terms and a willingness of universities to terminate licenses may well provide a market mechanism to minimize shelving, thus obviating the need for exercise of government march-in.

Section 2 provides survey results to motivate our consideration of multiple development stages as well as shelving concerns. Section 3 presents the basic model. In Section 4, we focus on the moral hazard problem when shelving by the licensee is not a concern. In Section 5, we consider the adverse selection problem that arises when the firm's opportunity cost of commercializing the invention is not observable by the TLO. Section 6 discusses the robustness of our results on shelving to risk aversion by one or more of the parties involved. We also briefly discuss the role a consulting contract may play in increasing inventor effort. Section 7 concludes.

2 University license contracts

Our focus on moral hazard and adverse selection, as well as risk and the types of contracts we consider, is motivated largely by the results of two surveys. One is the previously mentioned survey of 62 US university TLO's and the other is a follow-on survey of 112 firms that license-in university inventions. Details on survey design and sample frame are in Appendix A.

As shown in Table 1, respondents to both surveys characterize the majority of

Stage	TLO	Industry
Proof of concept (no prototype)	45 %	38 %
Prototype (only lab scale)	37 %	36 %
Preclinical stage	26 %	15 %
Clinical stage	10 %	5 %
Manufacturing feasibility known	15 %	9 %
Ready for practical or commercial use	13 %	7 %
Inventor collaboration needed	71 %	40 %
Inventions that fail	NA	46 %

Table 1: Stage of development at the time of license.

inventions licensed as embryonic. They estimate that three fourths or more of the inventions licensed are no more than a lab scale prototype at the time of license. For the subset of inventions requiring animal or clinical trials prior to use, less than one fourth have data available at the time of license. Manufacturing feasibility is known for no more than 15% of the inventions. No more than 13% are considered ready for commercial use, so that at least 87% of the inventions licensed require further development by the licensee, making shelving a potential issue for most inventions licensed. Although the estimates differ, both sets of professionals think inventor cooperation in further development is necessary for commercial application of many of the inventions licensed. Thus moral hazard with respect to inventor effort is also an issue. Finally, developing university inventions is quite risky as respondents to the business survey report 46% of the university inventions they license fail. Of those, 47% failed for purely technical reasons.

Our consideration of annual payments comes from the university survey, which included a number of open ended questions on the circumstances in which TLO professionals include various payment types in contracts. With regard to annual payments, 19% volunteered that they include annual payments as due diligence to ensure the licensee makes "reasonable efforts to commercialize." Interestingly, one respondent mentioned using, not annual fees, but an upfront payment when the "feasibility of march-in to recover the license was low." Whether or not these payments indeed deter such actions is an open question, but an additional 18% of the respondents said they had problems with firms shelving despite their best attempts at due diligence.[3] In general, due diligence to ensure licensee development appears to be a thorny issue for TLO personnel as 78% of the respondents noted that when they had to terminate licenses, the reason was failure to meet due diligence requirements or payments.

[3]Our question was "Have you had problems with companies despite proper due diligence terms acquiring a technology and shelving it to prevent its commercialization?"

Both surveys include information on milestone payments. Examples of milestones mentioned include reaching animal or clinical trials (as shown in Table 1, the majority of inventions requiring such trials were licensed prior to trials). Other common milestones are development of a business plan or prototypes, as well as milestones based on earnings. When respondents to the business survey were asked the importance of various payment terms (milestone payments, royalties, up-front fees and equity) when inventor cooperation was critical to development of the technology, they noted that milestone payments were by far the most important payment term. While the university survey did not contain information linking milestones with inventor cooperation, respondents emphasized that one of their most difficult issues is maintaining the needed inventor involvement.

When, if ever, the use of annual and milestone payments is optimal is not clear a priori, nor is it clear that the complex contracts noted in university licenses are optimal. For these reasons, in the following sections, we set up a stylized model of university licensing in which a TLO must design and offer an exclusive license contract to a firm that has expressed interest in developing a university invention. The invention requires further technical development whose completion is impossible without the inventor's cooperation. Commercialization of the technology also requires further investments by the firm, but the TLO has no information on the firm's true intent.[4] Importantly, we consider contract terms with possible payment types based on events or quantities the TLO can observe (contract acceptance, technical and commercial success, output).

3 A model of university licensing

Consider the problem faced by a university TLO with the responsibility for licensing an invention that requires further development before it can be commercialized. The model is similar to Jensen and Thursby [2001] except that we exploit the fact that, as shown in Table 1, many inventions must go through multiple stages of development before they can be successfully commercialized. In the first stage, inventor effort and firm investment are needed to determine technical success. In the second stage, which may or may not be reached, the firm invests in commercialization. The probability of technical success in the first stage is given by $p(e, X)$ where e is inventor effort and X is the firm's investment. We assume $p(0, X) = p(e, 0) = 0$ and $p(e, X) \in [0, 1)$. $p(e, X)$ is increasing in both arguments and strictly concave in e. For simplicity, we assume a fixed investment, $X > 0$, by the firm is necessary if technical success is to be determined (e.g., the cost of testing equipment). Since $p(e, 0) = 0$, we can write

[4]This is in contrast to recent work in which the TLO serves an intermediary purpose by being better informed than other players in the technology transfer process (Hoppe and Ozdenoren [2005], Hellmann [2006], and Macho-Stadler [2006]).

$p(e) = p(e, X)$ for $X > 0$. If the invention is a technical success, the firm then invests a fixed amount C in commercialization to earn profit $\pi > 0$ with probability z and zero with probability $1 - z$, where $z \in (0, 1]$. We assume that $z\pi - C \geq 0$ holds. The TLO, inventor, and the firm are all expected utility maximizers, and except in Section 4.2 we assume each is risk neutral.

The university owns and can exclusively license the invention but, in accordance with Bayh-Dole, this property right is contingent on the licensee making "reasonable" efforts to commercialize it. The TLO acts on behalf of the university.[5] One of the issues addressed by the university survey was TLO objectives. All respondents reported having multiple objectives including earning revenue and sponsored research, as well as simply executing licenses and encouraging commercialization. The most important of these objectives was earning revenue, with 98% of the respondents noting that revenue was either moderately or extremely important to them. Next in importance was the number of inventions commercialized, with 92% noting commercialization was either moderately or extremely important to them [Thursby *et al.* 2001]. Accordingly, we assume the TLO maximizes utility, $\mathbf{U}_A(\tilde{R}; L)$ given by

$$\mathbf{U}_A(\tilde{R}; L) = EU_A((1 - \alpha)\tilde{R}) - \mathcal{L},$$

where $\tilde{R}$ is random total revenue from licensing, $(1 - \alpha)$ is the TLO's share of revenue, and $\mathcal{L}$ is a function representing the expected loss associated with enforcement of the Bayh-Dole Act. Specifically, we assume that $\mathcal{L} = 0$ if the firm invests in technical development and the invention fails, or if the firm invests in technical development, finds it successful and invests in commercial development. In all other cases or if no firm accepts the contract offered by the TLO, $\mathcal{L} = L \geq 0$.[6] Thus even if licensing revenue is certain, the TLO strictly prefers an outcome in which the firm invests in commercialization.

The inventor's utility is given by $U_I(\alpha\tilde{R}) = \alpha\tilde{R}$, where α is her share of license revenue. She incurs strictly positive and increasing disutility of effort represented by the continuous function $V(e)$, with $V(0) = 0$. We also assume increasing marginal disutility ($V''(e) > 0$).

With risk neutrality, the firm's full expected utility is given by its (random) profit net of license payments. If it invests X in Stage 1 and C in Stage 2, the firm's expected

[5] We abstract from any agency problems between the TLO and administration. See Jensen *et al.* [2003] regarding the alignment of TLO and administration objectives.

[6] Implicitly, we are assuming that if either a technical failure or a technical success followed by a failure to commercialize are observed, the federal agency that funded the research may request an investigation into the causes of the failure. Importantly, we are assuming that as a result of the investigation, the federal agency knows whether the inventor and the firm invested in development.

profits gross of payments to the TLO are

$$p(e)(z\pi - C) - X.$$

The timing is as follows. The TLO offers the firm an exclusive license contract that specifies all payment terms. We denote a contract by $O = \{T_1, T_2, \ldots, T_n\}$, where T_i is the amount associated with payment type i if this payment is included in the contract. We restrict attention to payment types that are contingent on events and quantities the TLO can observe. In this baseline model, if the firm rejects the contract offer, the game ends. If it accepts, the inventor and the firm simultaneously choose effort and investment, respectively, to determine technical success of the invention. If development results in a technical failure, the game ends. If the invention is successful technically, the firm decides whether or not to invest in commercialization.

In general, inventor effort and firm investment (e, X or C) are neither observable nor contractible. However, as a benchmark we consider the TLO's problem if they are observable and contractible. In this case, the TLO can offer an enforceable contract specifying the amount of effort expected from the inventor, and the optimal payment by the licensee is an upfront fee contract, $O^* = \{F^*\}$, that extracts all of the firm's profits [Kamien 1972]. With observable effort, the TLO also offers a contract to the inventor, which stipulates the amount of effort e^*. The TLO's problem is to choose effort to maximize its utility subject to the firm's and the inventor's participation constraints.

$$\begin{aligned} &\text{Maximize} && \mathbf{U}_A(F; L) \\ &\text{with respect to} && e \geq 0 \text{ and } F \geq 0 \\ &\text{subject to} && \alpha F - V(e) \geq 0 \text{ (Inventor's participation constraint)} \\ &\text{and} && p(e)(z\pi - C) - X - F \geq 0 \text{ (Firm's participation constraint).} \end{aligned} \tag{1}$$

Since the firm's expected payoff is increasing in the level of effort, so is the TLO's. Hence, the TLO will pick the maximum level of effort consistent with the inventor's participation. That is, if effort is contractible, for a given upfront fixed fee F, the TLO will require an amount of effort $e(F)$ from the inventor given by

$$e(F) = \max\{e \geq 0 | \alpha F - V(e) = 0\}.$$

Under our continuity and monotonicity assumptions, $e(F)$ exists and is unique. Moreover, $e'(F) > 0$ and $F > 0$ implies $e(F) > 0$. Finally, since the contract extracts all of the firm's rent, F^* is defined as follows

$$F^* = \max\{F \geq 0 | p(e(F))(z\pi - C) - X - F = 0\}.$$

We assume that F^* is strictly positive, otherwise the return from investing in the invention is not sufficient for the firm and the TLO to agree on mutually beneficial contracts that do not involve a positive transfer of money from the TLO to the firm.

4 Non-contractible inventor effort

Now suppose there is no problem with firm shelving, but that inventor effort is subject to moral hazard since her effort is unobservable. We show that when development milestones are feasible, a payment tied to their achievement can solve the problem of inventor moral hazard. However, since effort is not observable, the inventor's rent is strictly positive and effort is generally less than the level $e^* \equiv e(F^*)$ characterized above.

4.1 Moral hazard: Is there a role for milestone payments?

Our university survey revealed that TLO personnel view obtaining faculty participation as one of the more challenging parts of their jobs [Thursby *et al.* 2001 and Jensen *et al.* 2003]. Jensen and Thursby [2001] show that obtaining inventor effort requires some type of payment tied to commercial success, such as royalties or equity. In their model, there is a single development stage so there is no role for milestone payments. In this section, we show that when development milestones are feasible, a payment tied to their achievement can solve the problem of inventor moral hazard.

In the context of our model, the inventor solves:

$$\text{Maximize } E[U_I(\alpha\tilde{R})|e] - V(e) \text{ with respect to } e. \tag{2}$$

It is clear that the expected utility term will not depend on e if the reward $\alpha\tilde{R}$ is the same whether or not the invention works. We therefore consider a contract $O = \{F, M\}$, which includes a payment contingent on technical success, M, in addition to any upfront fee F. The *milestone* payment is made if technical success occurs, but not otherwise. The firm's expected profit from accepting O is then given by

$$p(e)(z\pi - M - C) - X - F.$$

In the remainder of the paper we denote by $\hat{e}$, the inventor's optimal level of effort given the TLO's contract and often specify the payment types as arguments in the function $\hat{e}$. The TLO maximizes $E\mathbf{U}_A(\tilde{R}; L)$ with respect to M and F, subject to the firm's and the inventor's participation constraints and assuming that inventor effort is optimal, that is $e = \hat{e}(F, M)$.[7] Again, our assumptions on the marginal disutility of

[7]The optimal effort level at an interior solution $\hat{e}(F, M)$ is implicitly given by the first order

effort imply that $\hat{e}(F, M)$ is strictly positive only if the milestone payment is strictly positive.[8] The contract $O^{**} = \{F^{**}, M^{**}\}$ that induces positive inventor effort and maximizes the TLO's expected payoff is such that the firm's participation constraint is binding for effort given by $\hat{e}(F^{**}, M^{**}) > 0$. Hence, the upfront fee F^{**} is set equal to the firm expected profit from investing in technical development. Since we assume that all parties are risk neutral, a simple argument then implies that at the optimum, $M^{**} > 0$ and $F^{**} = 0$.[9] Therefore, M^{**} is defined by

$$M^{**} = \max\{M \geq 0 | p(\hat{e}(0, M))[z\pi - C - M] - X = 0\}. \tag{3}$$

4.2 Risk aversion and risk-sharing: Is there a role for upfront fees and royalties?

The natural question to ask is how the milestone M^{**} that solves the moral hazard problem compares with a royalty? Relying on the same arguments made by Jensen and Thursby (2001), it is simple to show that when all players are risk neutral, the milestone payment dominates a royalty. That is, denote the firm's optimal profit as $\pi(x)$, where x is output, and let $x(r)$ denote the firm's profit maximizing level of output when the royalty rate is r. Assume that $x(r)$ is unique and $x'(r) < 0$ and that royalty revenue $rx(r)$ is strictly concave so that its maximum is also unique. When all parties are risk neutral, the only use of a royalty is to induce inventor effort. As in Jensen and Thursby [2001] inventor effort is increasing in the royalty rate if royalty revenue is increasing in the rate, so that one could find a royalty that induces strictly positive effort, but the royalty introduces an output distortion. Thus, for the same reason that equity dominates royalties in Jensen and Thursby, the milestone dominates a royalty when the royalty reduces the firm's equilibrium profits. To see this, recall that when the inventor is risk neutral, other things equal, the TLO will set a fixed fee F that extracts all rents from the firm. Then the firm's expected profit is $p(e)[z(\pi(x(r)) - rx(r)) - M - C] - X - F$. By setting a milestone payment equal to $M' \equiv zrx((r)) + M$ and the royalty rate to zero, all else equal, the TLO provides

condition to the inventor's problem $p'(\hat{e}(F, M))\alpha M - V'(\hat{e}(F, M)) = 0$ and does not depend on F since the inventor is risk neutral.

[8]Given the complementarity of the inventor's effort and the firm's investment, there is another equilibrium in the technical development subgame in which the firm does not invest and the inventor spends no effort so that the project fails with probability one. As in Jensen and Thursby's [2001] analysis of moral hazard with royalties and sponsored research, it is straightforward to show that this equilibrium is unstable with standard assumptions on the firm and inventor's problems. We therefore restrict our attention to the equilibrium with positive effort.

[9]This follows from the fact that after substituting for the firm's net expected rent in place of F into the TLO's expected payoff, the TLO's problem is to maximize a share of the firm's gross expected rent, $(1 - \alpha)[p(\hat{e}(M, F))(\pi - C) - X]$, an expression that is strictly increasing in M. Hence, the optimal level for the milestone is the maximum level consistent with the firm's participation.

an incentive for the firm to raise its profit maximizing output and can thus raise the milestone payment so as to strictly increase its own payoff.

Why then do we observe contracts that include royalties and upfront fees as well as milestones? As in Bousquet *el al.* (1998), risk aversion is a natural explanation given the embryonic nature of university inventions.

Consider risk aversion on the part of the inventor. Although, as in Jensen and Thursby [2001], effort is independent of the upfront fee F if the inventor is risk neutral, optimal effort is decreasing in the fee if the inventor is risk averse. Furthermore, optimal effort is increasing in the milestone payment regardless of the inventor's attitude toward risk.[10] Thus the result from the previous section, that the TLO's contract will only include a milestone payment, is robust to allowing for inventor risk aversion. It follows that inventor risk aversion alone does not explain the presence of upfront fees in observed contracts. On the other hand, when either the TLO or the firm are risk averse, the TLO may find it optimal to include a positive upfront fee in order to reduce the variance of payments across states of nature (technical success or failure). Hence, in the absence of a problem with licensee shelving, TLO or firm risk aversion are necessary for upfront fees to be included in licensing contracts.

While a sufficiently high level of risk aversion exhibited by either the TLO, the firm or both can explain upfront fees, interestingly, uncertainty related to commercialization and risk aversion on the part of the firm are both necessary to explain the presence of royalties. In our model, once the firm invests C, it obtains $\pi(x)$ with probability $z < 1$, so that the worst state of nature for the firm is one in which the invention is a technical success but commercial success is not realized. In this case, the firm has to pay the fixed fee plus the milestone payment but earns no revenue from the invention. A positive royalty rate will reduce the variance in the distribution of profits and result in more equal payments across states. It is straightforward to show that with firm risk aversion, for a given upfront fee F, if z is sufficiently small, a simultaneous marginal increase in r and a marginal decrease in M that keep the inventor's incentives (and hence effort) constant increase the firm's expected payoff. Therefore, the TLO can increase the upfront fee and increase its own payoff.

The following proposition summarizes the main insight from this section.

Proposition 1 *When all parties are risk neutral, the optimal licensing contract with unobservable inventor effort contains only a positive milestone payment tied to technical success. Moreover, (i) TLO or firm risk aversion are necessary, but not sufficient for the optimal contract to include an upfront fee and (ii) firm risk aversion and the*

[10]Comparative statics results also show that if the TLO's contract contains a milestone payment, inventor effort is increasing in the inventor's share of revenue α if she is risk neutral or not too risk averse, a result consistent with empirical studies of inventor response to economic incentives [Goldfarb and Colyvas 2003, Lach and Schankerman 2003].

existence of uncertainty at the commercialization stage are necessary, but not sufficient for the optimal contract to include an output royalty.

5 Shelving

In this section, we examine the problems that arise because the firm's commercialization effort is not contractible. We consider situations in which the firm may be better off shelving the invention than commercializing it. We allow for two types of shelving, one in which a firm licenses with the intent to prevent commercialization of the invention and another in which a firm licenses intending to commercialize the invention, but decides not to commercialize it after spending time on development. For simplicity, we refer to the first type of shelving as *intentional* and to the second as *unintentional.*

The motivation for intentional shelving is similar to that of "sleeping patents" examined by Gilbert and Newbery [1982] in which a monopolist patents substitutes for its product to keep others from producing it. Well known examples include DuPont's patenting of 200 substitutes for Nylon. More recently, Cohen *et al.* [2000] find that when firms in their survey patent inventions, 82% (64%) patent them in order to block rivals in the case of product (process) inventions.

The most publicized example of unintentional shelving was involved in the CellPro-Johns Hopkins march-in dispute. In 1997 Bayh and Cutler petitioned the National Institutes of Health to take back the Johns Hopkins license for the My-10 antibody originally licensed to Becton Dickinson. While the company had invested in development, over time it decided to withdraw from the therapeutic business, so that developing the antibody had no economic value to the company.

Suppose that a firm has expressed interest in the invention. We assume that at the time the TLO offers a license contract, the firm is of one of two types. With probability s the firm is interested in licensing the invention to prevent development (either by itself or a rival). It is natural to think of this firm either producing or trying to develop a substitute for the invention. The firm, which we call an *intentional shelver* earns profit π^m when it obtains the license for the invention but does not commercialize it, and $\pi^c \leq \pi^m$, if it obtains the license and commercializes it. The shelver earns $\pi^d \leq \pi^m$ if another firm, holding the exclusive license, commercializes the invention. Therefore, a shelver would never invest in commercial development since $\pi^c - C < \pi^m$ by assumption. Moreover, a shelver saves an amount $D \equiv \pi^m - \pi^d$ when it obtains and shelves the license, preventing commercialization of the innovation.

With probability $1 - s$, at the time the TLO makes a contract offer, the firm is interested in licensing the invention to develop and commercialize a product. However, we assume that after technical success is determined but before the decision to incur C is made, new information becomes available which reveals whether the commercial potential for the invention is good or bad. With probability q, commercial potential is good and the firm's expected profit from commercializing is equal to π (in this section, without loss of generality, we assume $z = 1$). With probability $1 - q$, the firm does not expect to earn any profit from commercializing the invention. Hence, with probability $(1 - s)(1 - q)$, the firm is an *unintentional shelver*, which, ex-post, will decide not to pursue commercialization even though it knows the invention works. We assume that $q(\pi - C) > X$ holds, so that there exists a range of probabilities of technical success and milestone payments for which such a firm is willing to invest in technical development. The probability that the firm is not a shelver and invests in commercial development following in technical success is equal to $(1 - s)q$.

The timing in this extension of our baseline game is as follows. Nature picks a firm to which the TLO offers a contract. The firm accepts or rejects the contract. If it rejects, the TLO must decide whether or not to search for a second firm at a cost $K > 0$. For simplicity, we assume that with probability one, the firm the TLO finds as a result of this search is not an intentional shelver. The second firm may accept or reject the contract. If it rejects the contract, we assume that the project is abandoned. If the first firm accepts, the firm and the inventor play the simultaneous development game. If the outcome of the development game is a failure, the TLO may take the license back from the firm. Note that failure can occur because the firm did not invest or the inventor spent no effort, or simply because the invention does not work. If the TLO takes the license back from the firm, it decides whether or not to search for a second firm at cost K and the outcome of the search process is as described above. Again, we assume the TLO can only search once after taking a license back. Moreover, we assume that the opportunity to find a second firm vanishes by the time the commercialization stage starts. Therefore, the threat of taking the license back to license to a second firm is only credible early in the development process. Throughout the analysis, we focus on pure strategy equilibria.

There are three critical dates at which the TLO may want to search for a different firm: immediately after the initial contract offer is made (if that offer was rejected), immediately after a technical failure is observed and, finally, after commercial success if the TLO is able to determine that the firm is a shelver.

The information available to a potential second firm about investments made by the first firm plays an important role in this context. To abstract from potential agency problems between the TLO and the second firm, we assume that the firm can examine the technology prior to making a decision. Furthermore, the firm has the ability to determine whether the technology is a sure technical failure (as would be the case if technical development resulted in a failure even though both the first firm and the inventor invested in development) or if it may potentially work (as would be the case if the first firm shirked).[11]

To characterize the optimal contract offered to the first firm, it is necessary to analyze the TLO's behavior when the first firm rejects the TLO's original contract offer. In this case, the TLO may try to license the invention to another firm. If the TLO searches for a second firm, then, it will optimally offer a milestone only contract similar to O^{**} to that firm. This contract $O^{**'} = \{M^{**'}\}$ maximizes the TLO's payoff and is such that the firm's participation constraint is binding. That is, $M^{**'}$ is defined by

$$p(\hat{e}(M^{**'}))[q(\pi - C) - M^{**'}] - X = 0.$$

When $q = 1$, $O^{**'} = O^{**}$. Of course, it is optimal to search for a second firm if and only if

$$-L \leq (1 - \alpha)p(\hat{e}(M^{**'}))M^{**'} - K \tag{4}$$

holds. Since our premise is that the TLO views shelving as a problem we make the following assumption on the parameters:

A1: The loss from failing to provide adequate incentives for development satisfies $-L \leq (1 - \alpha)p(\hat{e}(M^{**'}))(M^{**'} - qD) - K$.

A1 ensures that the TLO would not find it profitable to sell the technology to an intentional shelver via an upfront-fee only contract. To see this, suppose the TLO knows the firm's type and is faced with an intentional shelver. The assumption implies that the TLO would prefer to turn that firm down and search for another firm as opposed to extracting all rents from the shelver and committing not to take the license back in case of a technical failure.[12] Clearly, **A1** is more likely to be satisfied if the probability of commercial success with the second firm, $p(\hat{e}(M^{**'}))q$, is high. Note that the loss expected from the enforcement of the march-in provision included in the Bayh-Dole act helps prevent the TLO from writing an enforceable contract stating that the TLO will not take the license back from the firm even if a technical failure is observed.

[11]We also assume that the information the second firm is able to gather by examining the technology is not verifiable, so that the TLO cannot use it in court against a shelver.

[12]Suppose that the TLO knows the firm's type. To derive the inequality in **A1**, note that in general, if the TLO searches and finds another firm, it will offer $O^{**'}$ to that firm. Hence, the TLO will search if and only if (4) is satisfied. Suppose that this is the case. Shelvers can anticipate the TLO's optimal response and a simple argument shows that they benefit from licensing the invention if and only if the fee F set by the TLO satisfies $p(\hat{e}(M^{**'}))qD \geq F$. Thus, the maximum rent the TLO can extract from a shelver is equal to $p(\hat{e}(M^{**'}))qD$. Hence, the TLO's maximum utility from licensing to a shelver is given by $(1 - \alpha)p(\hat{e}(M^{**'}))qD - L$. Comparing the latter to the net utility obtained from searching for another firm, we obtain the inequality in **A1**. Finally, it is straightforward to show that **A1** is sufficient for the TLO to indeed find it profitable to search for a second firm.

5.1 The optimal contract in the absence of intentional shelvers

In this section, let $s = 0$ so intentional shelving is not a concern. Since, ex-ante, the firm does not know the value of the licensed invention, it expects to earn a gross expected profit equal to $q(\pi - C)$ from investing X and achieving technical success. We have shown in the previous section that when additionally, $q = 1$ holds, then it is optimal for the TLO to offer a milestone only contract. The problem facing the TLO when $q < 1$ is that it may want to take the license back from an unintentional shelver even though the firm achieved technical success. Thus, the TLO may have an incentive to include a second milestone payment to be paid at the commercialization stage, when the firm's knowledge of the technology is sufficient for it to decide whether it wants to continue working or abandon the project. The TLO's incentive to take the license back from an unintentional shelver stems from the opportunity to license to a second firm. It is thus important to determine when it is optimal for the TLO to license to a second firm after it has observed a technical success with the first firm.

Formally, suppose that technical development yields a success, but that the firm finds that it does not want to develop it further. Then, the firm's expected profit from keeping the license is equal to 0. Hence, if the contract includes a milestone payment $A > 0$ to be paid at the commercialization stage, the firm would prefer to return the license to the TLO to avoid having to make the payment. Suppose that the contract also contains a technical milestone payment equal to M and that the second milestone, A, satisfies $A < \pi - C$. The firm anticipates that it will return the license upon finding out that the project has no commercial value, but that it will keep the license and commercialize the project if it has value. It then follows that prior to the technical development stage, the firm's expected payoff is given by

$$p(\hat{e}(M, A))[q(\pi - C - A) - M] - X \tag{5}$$

where $\hat{e}(M, A))$ is the investor's optimal effort level given the TLO's contract and the fact that the TLO will take the license back from an unintentional shelver in order to license to a second firm.

To determine the optimal level of effort, it is thus necessary to characterize the optimal contract offered to the second firm in case the first firm turns out to be an unintentional shelver. At this point, technical development has been completed so that the TLO maximizes his payoff by offering an upfront-fee contract $O_2 = \{F_2\}$ to the second firm. The optimal upfront fee extracts all rents from the firm and is equal to its expected profit, $q(\pi - C)$. Therefore, the TLO sets

$$F_2 = q(\pi - C)$$

and optimal inventor effort is given by

$$\hat{e}(M, A) = \arg\max_{e}\{p(e)\alpha[M + qA + (1 - q)q(\pi - C)] - V(e)\}.$$

Finally, for the TLO to find it optimal to incur the search cost K to find a second firm upon taking the license back from an unintentional shelver, the following has to be satisfied

$$-L \leq (1 - \alpha)q(\pi - C) - K.$$

It is simple to check that the above inequality is implied by **A1**.

Since the second milestone payment, A, is constrained to be strictly positive, an optimal solution to the TLO's problem does not exist because of an open-endedness problem. Indeed, the optimal contract is such that A is set to its lowest possible value. To avoid this technical problem, we let ϵ be the smallest feasible strictly positive fee. Summarizing the above characterization of the TLO's optimal contract, we have the following result.

Proposition 2 *Assume* ***A1****. If the pool of firms interested in obtaining the license from the TLO does not include intentional shelvers, then the TLO's optimal contract is* $O^{us} = \{M^{us}, A^{us}\}$*, where* $A^{us} = \epsilon$ *and* M^{us} *is set so that (5) is binding when* $A = A^{us}$ *and* $M = M^{us}$*. That is, the contract includes two milestone payments. One payment is conditional upon technical success and the other payment is made at the commercialization stage.*

The second milestone payment A^{us} can be interpreted as an annual fee. Such fees are commonly observed in contracts, even though, in the absence of the potential for unintentional shelving, it is not clear what purpose they serve. Absent unintentional shelving, the TLO would not include the annual fee in order to increase the technical milestone payment.

Our result is derived for the special case in which unintentional shelvers obtain no profit from commercializing the license. Therefore, an infinitesimally small annual fee is sufficient for such firms to strictly prefer returning the license to the TLO (as opposed to keeping it, but not investing in development). Another possibility for modeling unintentional shelving is to assume that firms differ in their opportunity cost of investing in commercialization. Since this opportunity cost includes the profit that could potentially be earned from devoting resources to another project, it is likely to remain uncertain even after technical success has occurred. Suppose that ex-ante all firms are identical and genuinely interested in developing the license (and have expected commercialization costs of C), but that each recognizes that upon technical success, it will learn the true distribution of the commercialization cost. Then, with ex-post heterogeneity in the distribution of development costs, firms differ in their probability of commercialization. In this case, a strictly positive, and possibly substantial annual fee will be optimal in order to prevent firms with high expected opportunity costs of commercialization from holding on to the license until they decide whether or not to invest.

5.2 Equilibrium contracts with intentional shelvers

When the firm expressing interest in licensing the technology might be an intentional shelver, the TLO faces a trade-off between letting such firms obtain a license (but shirk in development) and preventing them from accepting the contract, which implies distorting incentives for non-shelvers. Since the unintentional and intentional shelving problems are essentially independent in our model, in this section we assume that $q = 1$. Of course, if $q < 1$ the results are qualitatively unchanged except for the fact that, when unintentional shelving is a possibility, the TLO's optimal contract includes a positive annual fee as we showed in Proposition 2.

To analyze the TLO's problem in the presence of intentional shelvers, note first that if the TLO does not incur the search cost to find a second firm after observing a technical failure with the first firm, then the contract O^{**} does not provide sufficient incentives for a non-shelver to invest in technical development. Indeed, suppose that the TLO offers this contract to the first firm and assume that both types of firms accept it. Observing that the TLO does not search for a second firm upon being faced with a technical failure, a shelver does not invest X and an invention licensed by a shelver fails with probability one. In this case, the inventor's optimal effort is given by

$$\arg\max_{e}\{p(e)\alpha(1-s)M^{**} - V(e)\}$$

and is strictly less than $\hat{e}(M^{**})$. Therefore a non-shelver would incur an expected loss from accepting O^{**} and investing X. It thus follows that a non-shelver is indifferent between accepting the contract and shirking, and rejecting O^{**}. Therefore, the TLO's payoff will be equal to $-L$. Our question is then: do there exist contracts that satisfy all relevant incentive compatibility constraints as well as the additional requirement that, upon accepting the contract, a non-shelver is willing to invest in development?

5.2.1 A separating contract with an upfront fee

In this section, we construct a contract that can be supported in a (constrained) Perfect Bayesian equilibrium in the class of contracts that only non-shelvers accept. We refer to contracts belonging to that class as separating contracts. The equilibrium we construct has the following features. In equilibrium, shelvers reject the TLO's offer so that, assuming **A1** holds, the TLO searches for a second firm upon being turned down by a shelver. Importantly, since shelvers turn down the contract offer, if a firm accepts the contract, then the TLO believes that the firm is a non-shelver with probability one. This implies that upon observing a technical failure following the development stage, the TLO knows that the invention is worthless. Therefore, in equilibrium, the TLO does not search for a second firm upon observing a technical failure with the first firm. To begin the characterization of the equilibrium contract, note that it follows from the latter observations that if the TLO's contract includes a fixed fee F, then by deviating and accepting the TLO's contract, a shelver can earn a certain profit equal to

$$\pi^m - F. \tag{6}$$

Therefore, the incentive compatibility constraint that guarantees that a shelver prefers to reject the contract is one such that its expected payoff from doing so is greater than (6). To determine a shelver's expected payoff from rejecting the contract, observe that upon being turned down by the first firm, the TLO offers O^{**} to the second firm. Hence, a shelver's expected payoff from rejecting the TLO's offer is equal to

$$p(\hat{e}(M^{**}))\pi^d + [1 - p(\hat{e}(M^{**}))]\pi^m. \tag{7}$$

Hence, a shelver's incentive compatibility constraint, which states that the expected payoff from rejecting the contract must be greater than the payoff from accepting it, is given by

$$F \geq p(\hat{e}(M^{**}))D > 0. \tag{8}$$

It thus follows immediately that a separating contract must include a strictly positive upfront fee.

Suppose for now that the TLO's contract is $O = \{F, M\}$. Clearly, a non-shelver's participation constraint must be satisfied as well. That is, a non-shelver's expected payoff from accepting the contract must be greater than zero. Hence

$$p(\hat{e}(F, M))[\pi - C - M] - X - F \geq 0$$

must hold, where

$$\hat{e}(F, M) = \arg\max_e \{\alpha[F + p(e)M] - V(e)\}.$$

Note that the probability that the firm is a non-shelver, $1 - s$, does not enter the inventor's expected payoff since, in equilibrium, the inventor's updated belief that a firm that accepted the contract is a non-shelver must be equal to one.

We know that increasing the upfront fee does not raise effort, but forces the TLO to decrease the milestone payment. Hence, the TLO will want to offer the smallest upfront fee capable of deterring shelvers, from which it follows that F will be set so that (8) is binding, i.e. $F = p(\hat{e}(M^{**}))D$. In equilibrium, a non-shelver's participation constraint is therefore given by

$$p(\hat{e}(F, M))[\pi - C - M] - X \geq p(\hat{e}(M^{**}))D. \tag{9}$$

The upfront fee being the only instrument the TLO can use to sort shelvers from non-shelvers, if a milestone payment such that (9) holds does not exist, then an equilibrium separating contract does not exist. We thus have the following result.

Proposition 3 *Assume **A1**. If there does not exist a milestone payment M for which (9) holds, a separating contract cannot be supported in a Perfect Bayesian equilibrium. If there exists an M for which (9) holds, a separating contract can be supported in a (constrained) Perfect Bayesian equilibrium. The equilibrium contract is $O^{sep} = \{F^{sep}, M^{sep}\}$ where $F^{sep} = p(\hat{e}(M^{**}))D$ and M^{sep} is set such that a non-shelver's participation constraint given by (9) is binding. In equilibrium, the TLO searches for a second firm if and only if the first firm rejects O^{sep}. Furthermore, as D goes to 0, O^{sep} converges to O^{**}.*

Separation therefore results in an effort distortion that stems from the need for raising the upfront fee to keep shelvers out. A separating contract is feasible only if the expected net value of the invention to non-shelvers $\pi - C$ is high enough for such firms to be willing to pay the milestone payment.

5.2.2 Pooling contracts and the role of the threat to take the license back

We now characterize a contract that forms a (constrained) Perfect Bayesian equilibrium in the class of contracts that both types of firms accept. We refer to contracts in that class as pooling contracts.

The critical feature of an equilibrium pooling contract in pure strategies is that shelvers must shirk at the technical development stage. If in equilibrium, both types of firms invest X in technical development, then, upon observing a technical failure, the TLO's updated probability that the invention is worthless is equal to one. Hence, the TLO expects that the second firm will find the technology worthless so that engaging in search is itself worthless. However, if the TLO does not search, then a shelver is better off shirking at the technical development stage and will not work in the first place. We show that this important aspect of incentives provision with intentional shelving manifests itself in the form of an additional necessary incentive compatibility constraint.

To characterize the equilibrium pooling contract when it exists, let $O^{po} = \{F^{po}, M^{po}\}$ be the contract the TLO offers to the first firm. Suppose for now that shelvers accept the contract, but shirk at the technical development stage as is required in equilibrium. Let $\hat{e}^{po}$ denote the inventor's optimal effort level and note that $\hat{e}^{po}$ only depends on the original contract terms, but also on revenue that may be obtained with a second firm in the event that the license is taken back from the first firm.

Under which conditions will the TLO find it optimal to search for a second firm when technical development yields a failure? Critical to the TLO's optimal decision is its updated belief that the firm is a shelver when it observes a technical failure. Taking $\hat{e}^{po}$ as given and using Bayes rule, in equilibrium, this belief is given by

$$\mu(s) = \frac{s}{s + (1-s)(1-p(\hat{e}^{po}))}.$$

If the TLO decides to search for another firm, then we assume for now that the TLO offers a milestone only contract to that firm (even though incentives for inventor effort have already been provided). Let this milestone payment be equal to M_2. If the TLO decides to search for a second firm, its expected payoff in the continuation of the game is equal to

$$\mu(s)p(\hat{e}^{po})(1-\alpha)M_2 - K, \tag{10}$$

where M_2 extracts the firm's rent and is thus the solution to the firm's zero-profit condition

$$p(\hat{e}^{po})[\pi - C - M_2] - X = 0.$$

It is optimal for the TLO to take the license back after a technical failure and search for a second firm if and only if (10) is greater than $-\mu(s)L$.

Lemma 1 *Consider a pooling contract containing a milestone payment only in which, (i) shelvers shirk in development, (ii) non-shelvers invest X in development and (iii) the TLO searches for a second firm upon observing a technical failure. If such a contract is part of a (constrained) PBE, then $M^{po} = M_2^{po} = M^{**}$.*

Proof. To show that $M^{po} = M_2^{po}$, note that for given effort level e and milestone payment M, a non-shelver's participation constraint is given by

$$p(e)[z\pi - C - M] - X \geq 0, \tag{11}$$

while, with milestone payment M_2, the second firm's participation constraint is given by

$$p(e)[z\pi - C - M_2] - X \geq 0. \tag{12}$$

By assumption, the TLO searches for a second firm upon observing a technical failure and the second firm accepts the contract and works in development if and only if the first firm shirked. Hence, from the inventor's point of view, the probability that a second firm accepts the TLO's contract after a technical failure is s, the probability that the first firm was a shelver. Therefore, the inventor's expected payoff is given by

$$p(e)\alpha[(1-s)M + sM_2] - V(e).$$

Finally, the TLO's expected payoff is equal to

$$p(e)(1-\alpha)[(1-s)M + sM_2] - [(1-s)(1-p(e)) + s]K.$$

Hence, the TLO's payoff is strictly increasing in M, M_2 and e. It follows that the TLO will set the maximum values of M and M_2 consistent with a non-shelver and a second firm accepting their respective contract. Therefore if for a given e, M^{po} solves (11) satisfied with equality and M_2^{po} solves (12) satisfied with equality, $M^{po} = M_2^{po}$. To show that $M^{po} = M_2^{po} = M^{**}$, we note that in equilibrium, $M^{po} = M_2^{po}$ implies that the inventor's expected payoff is equal to $p(e)\alpha M^{po} - V(e)$, while the TLO's expected payoff is equal to $p(e)(1-\alpha)M^{po} - [(1-s)(1-p(e)) + s]K$. Therefore, setting $M^{po} = M^{**}$ is optimal from the TLO's point of view since it induces effort level $e^{**} = \hat{e}(M^{**})$ and implies that (11) and (12) are binding. □

Lemma 1 implies that one of the relevant constraints to determine whether a constrained PBE with successful ex-post licensing exists is given by (10) evaluated at $\hat{e}(M^{**})$ and M^{**}, and set greater than or equal to $-\mu(s)L$. That is,

$$\frac{sp(\hat{e}(M^{**}))}{s + (1-s)(1-p(\hat{e}(M^{**}))}(1-\alpha)M^{**} - K \geq -\frac{sL}{s + (1-s)(1-p(\hat{e}(M^{**}))}. \tag{13}$$

When (13) holds, the TLO takes the license back from the firm upon observing a technical failure, searches for a second firm and offers a contract to that firm. For beliefs to be correct in equilibrium, i.e., for the TLO to believe that there is a chance the invention failed because the firm shirked, it must be the case that a shelver prefers shirking to investing. Hence, the following inequality must hold

$$p(\hat{e}(M^{**}))\pi^d + [1 - p(\hat{e}(M^{**}))]\pi^m \geq \pi^m - p(\hat{e}(M^{**}))M^{**} - X$$

or, after re-arranging terms,

$$p(\hat{e}(M^{**}))[D - M^{**}] - X \leq 0. \tag{14}$$

However, note that by the definition of M^{**}, combining (3) and (14), the condition simplifies to

$$\pi - C \geq D. \tag{15}$$

We assume that the above inequality holds and maintain the assumption for all of the results below. The following proposition describes the equilibrium when the contract is constrained to be a pooling contract. Before stating the proposition, we require two additional definitions. First, let $\hat{e}^s(M)$ be given by

$$\hat{e}^s(M) = \arg\max_e \{p(e)\alpha(1-s)M - V(e)\}.$$

Effort level $\hat{e}^s(M)$ is optimal from the inventor's point of view if the inventor believes that the probability the firm invests X is equal to $1-s$. Now define M^s,

$$M^s = \max\{M | p(e^s(M))(\pi - C - M) - X = 0\}.$$

Proposition 4 describes the form taken by pooling contracts that can be supported in a (constrained) PBE.

Proposition 4 *Assume **A1**. Then,*

(i) if (13) holds, a pooling contract can be supported in a (constrained) Perfect Bayesian equilibrium. In equilibrium, the TLO offers O^{**} to the first firm. Upon observing a technical failure, the TLO incurs the search cost K to find a second firm. If the second firm is willing to purchase the license, the TLO offers O^{**} to that firm.

(ii) If (13) does not hold, then a pooling contract can be supported in a (constrained) Perfect Bayesian equilibrium if and only if the probability that the firm is a shelver, s, is sufficiently low. In equilibrium, the TLO offers contract $O^{pool} = \{F^{pool}, M^{pool}\}$ to the first firm where $F^{pool} > 0$ is possible. If the contract is such that $F^{pool} = 0$, then $M^{pool} = M^s$. In equilibrium, both types of firms accept the contract. Furthermore, at the technical development stage, shelvers shirk, but non-shelvers invest X. The TLO does not search for a second firm upon observing a technical failure. Finally, as s goes to zero, O^{pool} converges to O^{**}.

Proposition 4(i) shows that if the TLO's threat of taking the license back from the firm to license it again is credible, then the optimal contract is O^{**}, the same contract the TLO would offer in the absence of shelving concerns. The result follows from the fact that by searching for a second firm upon technical failure the TLO is given a second chance. Hence, it does not have to distort incentives in order to offer a contract that is incentive compatible for all parties involved. From Proposition 4(ii), note that if the threat to license the invention a second time is not credible, the TLO may find it optimal to include an upfront fee in the contract because of the trade-off between revenue raised through the milestone and revenue raised through the upfront fee. This trade-off does not exist in the absence of intentional shelvers. More specifically, an increase in the milestone payment will increase effort and thus the amount of revenue raised through the milestone payment. However, with probability s, the firm is a shelver who fails to meet the milestone with probability one. Hence, the TLO may find it profitable to increase the upfront fee in order to obtain some revenue from the shelvers.

Proposition 4(i) depends in part on our assumption that the TLO finds a firm with probability one upon searching. When this assumption is relaxed, the milestone payment M^{**} is no longer feasible because inventor effort decreases as the probability that a second, back-up firm exists decreases. It is important to note that if both the probability of finding a second firm and L, the loss associated with the enforcement of the Bayh-Dole act, are low, then the TLO may find it optimal to include a small upfront fee to the contract offered to the first firm (so as to obtain some revenue from shelvers). On the other hand, if L is sufficiently high, the TLO will find it in its interest to provide as great an incentive for effort as possible in order to increase the probability of success. In the latter case, generating revenue is not the TLO's primary concern. As a result, even though the probability of finding a second firm is not equal to one, the TLO offers a contract including a milestone payment only.

5.2.3 When does the TLO offer a separating contract?

In this section, we ask when the TLO will want to offer a separating contract with a high upfront fee rather than a pooling contract. It is interesting to note that in our model, the probability that the firm is a shelver (s) does not affect the existence of a separating contract, while it is critical in guaranteeing the existence of a pooling contract. On the other hand, while the difference between duopoly and monopoly profit does not affect the existence of a pooling contract (at least as long as (15) holds), it clearly is critical in determining whether a separating contract exists.

How s and D affect existence for the two types of contracts is intuitive. A separating contract is designed so that the original contract offer sorts between shelvers and non-shelvers. Hence, in equilibrium, the contract itself will not depend on s since shelvers are eliminated at the contract offer stage. However, in order to provide adequate incentives for shelvers to reject the contract, the TLO must make it sufficiently costly for them to obtain the license. Hence, the upfront fee depends on D, the amount a shelver gains by blocking commercialization. Now turning to the pooling contract, in equilibrium, both types of firms accept the contract, but each type behaves differently at the development stage. Hence, the probability that the firm is of a given type matters. However, since the payment scheme is not type dependent, as long as shelvers prefer to shirk rather than invest, the contract does not depend on how much a shelver gains by accepting the license contract and thus does not depend on D.

Figure 1 partitions the (s, D) space into regions where the different types of contracts exist, holding all other parameters constant. The critical values of the probability that the first firm is a shelver, s_1 and s_2, are defined as follows. The lower bound s_1 is the minimum value of the probability above which (13) is satisfied. The existence of $s_1 < 1$ is guaranteed by assumption **A1**. The upper bound s_2 is the maximum value of s below which M^s exists. If $s > s_2$, then there does not exist a value of the milestone payment for which, given that the inventor spends $\hat{e}(M)$ on development effort, a non-shelver would make a non-negative profit by working on development and paying the milestone. Note that $s_2 \geq s_1$ is possible. Finally, $\tilde{D}$ is the maximum value of D below which there exists a value of the milestone payment for which (9) holds.

Figure 1 shows that there are values of s and D where two different types of contracts can be supported in a (constrained) Perfect Bayesian equilibrium. However, out of the set of feasible contracts, the TLO will offer the contract that maximizes its expected payoff. Suppose a (constrained) PBE separating contract exists and consider the TLO's expected payoff from offering such a contract. Using Proposition 3, the TLO's expected payoff is equal to

$$(1-\alpha)[(1-s)(p(\hat{e}^{sep})M^{sep} + F^{sep}) + sp(\hat{e}^{**})M^{**}] - sK. \tag{16}$$

In (16), with probability $1-s$, the firm is not a shelver, and the TLO obtains the expected licensing revenue from licensing to the firm. With probability s, the firm is a shelver that turns down the contract offer. In this case, the TLO searches for a second firm at cost K and obtains expected revenue $p(\hat{e}^{**})M^{**}$ from licensing to that firm.

Now suppose that s is large enough so that a (constrained) PBE pooling contract with the possibility for licensing to a second firm exists, i.e. $s \geq s_1$. Then, the TLO's payoff from offering such a contract is

$$(1-\alpha)p(\hat{e}^{**})M^{**} - [s + (1-s)(1-p(\hat{e}^{**}))]K. \tag{17}$$

Equation (17) states the TLO can guarantee itself an expected licensing revenue of $p(\hat{e}^{**})M^{**}$. This follows from (i) in Proposition 4. Both a non-shelving first firm and a second firm are offered the same contract, which results in the same inventor effort. However, the TLO searches for a second firm at cost K whenever it observes a technical failure. The probability of this happening is equal to $s + (1-s)(1-p(\hat{e}^{**}))$.

Finally, if s is low, $s \leq s_2$, then in a (constrained) PBE pooling contract, the TLO does not license to a second firm upon observing a technical failure and its expected payoff from a pooling contract is equal to

$$(1-\alpha)(1-s)p(\hat{e}^{pool})M^{pool} + F^{pool} - sL. \tag{18}$$

Following (ii) in Proposition 4, (18) states that the TLO's expected payoff consists of the expected milestone payment only if the firm is a non-shelver (probability $1 - s$) and the upfront fee if one is included in the contract. With probability s, the firm is a shelver, who shirks with probability one. In this case, a technical failure results. The TLO does not search for a second firm and consequently, incurs the loss associated with enforcement of the Bayh-Dole act.

Suppose s and D are such that both a separating contract and a pooling contract with re-licensing can be supported in a (constrained) PBE (the southeast corner of the rectangle in Figure 1). Comparing (16) to (17) yields a difference in the TLO's expected payoff between the two types of contracts equal to

$$(1 - s)[(1 - \alpha)(p(\hat{e}^{sep})M^{sep} + F^{sep} - p(\hat{e}^{**})M^{**}) + (1 - p(\hat{e}^{**}))K].$$

For the TLO, the nature of trade-off between offering a separating contract as opposed to a pooling contract when both exist is as follows. The pooling contract provides better incentives to the inventor because the possibility of search for a second firm makes it possible for the TLO to use a contract with a high milestone payment (and no upfront fee). However, the pooling contract can be supported in equilibrium only if the TLO incurs the search cost upon technical failure. On the one hand, by offering a separating contract instead, the TLO creates an effort distortion by increasing the upfront fee and decreasing the milestone payment. On the other hand, it also eliminates shelvers at the contract offer stage and does not have to search for a second firm, thereby saving on the cost of search. The level at which the upfront fee has to be raised to deter shelving, F^{sep}, increases with the incentive to shelve, D, and so does the effort distortion. Therefore the higher D, the more likely it is that the TLO will offer a pooling contract.

Similarly, using (16) and (18), the difference between the TLO's expected payoff from a separating contract and the expected payoff from a pooling contract with no re-licensing is equal to

$$(1 - \alpha)[(1 - s)(p(\hat{e}^{sep})M^{sep} + F^{sep} - p(\hat{e}^{pool})M^{pool} - F^{pool}) + sp(\hat{e}^{**})M^{**}] - s(K - L).$$

It is clear that $L \geq K$ is sufficient for the above expression to be strictly positive and thus, for the separating contract to dominate the pooling contract with no re-licensing from the TLO's point of view. The separating contract dominates the pooling contract with no re-licensing both in terms of revenue and in terms of incentive provisions. Hence, offering a separating contract also allows the TLO to avoid raising suspicion by the federal agency which funded the research as to whether or not its contracts provide adequate incentives for development.

6 Extensions and Discussion

6.1 Shelving and risk aversion

The analysis of optimal contracts in the presence of shelvers in the previous section assumes all players are risk neutral. However, as shown in Section 4, risk aversion may play an important role in shaping contracts. We thus discuss the robustness of our results to relaxing the risk neutrality assumptions.

A result similar to Proposition 2 holds independently of the parties' risk attitudes because the annual fee is necessary in order to prevent unintentional shelvers from holding on to the license. Interestingly, when the firm is risk averse, the annual fee used to separate unintentional shelvers from other firms also serves to spread the risk ex-ante. In fact, the annual fee may act as a substitute for an output royalty in order to spread the risk, but avoids the output distortion implied by a royalty.

Regarding intentional shelving, when the TLO is risk averse, an upfront fee continues to be optimal to separate shelvers from non-shelvers. In fact, as compared to an environment without shelvers, the distortion caused by intentional shelving may be less important when the TLO is risk averse than when the TLO is risk neutral. This results from the fact that a risk averse TLO may find it optimal to include an upfront fee in a contract simply to reduce the risk inherent in a milestone-only contract. It follows that to keep shelvers out, the TLO might only have to raise the upfront fee slightly above the level that is optimal in the absence of shelvers. Note that with firm risk aversion, as we argued in the text preceding Proposition 1, the TLO's optimal contract in the absence of shelving may include an upfront fee. Thus, the distortion resulting from the TLO's attempts to separate intentional shelvers from non-shelvers may also be mitigated when the firm is risk averse. However, it is clear that risk sharing via a royalty is more difficult to achieve when intentional shelving is a concern. Indeed, shelvers benefit from contracts that require payments such as royalties, based on the realization of events they wish to prevent from happening.

The main impact of inventor risk aversion comes from the strict disincentive for effort that an upfront fee creates. With inventor risk aversion, raising the upfront fee to keep shelvers out is more costly to the TLO than when the inventor is risk neutral. Hence, the separating contract, if it exists, is not as attractive when the inventor is risk averse.

The features of the pooling contracts described in Proposition 4 may also be affected by risk aversion. With pooling contracts, inventor risk aversion reinforces the TLO's incentive to offer a milestone-only contract since an upfront fee would decrease the inventor's effort. For reasons similar to those discussed above, the role played by risk aversion on the part of the TLO and the firm is more complex since an upfront fee can be used for risk sharing.

6.2 The role of consulting

The remaining puzzle with regard to inventor effort is that respondents to our business survey indicated that when faculty effort is critical for further development, consulting contracts are quite common. On average, they indicated using consulting 58.7% of the time. This is curious if, as we suggest, milestones or royalties address inventor moral hazard. In this section, we analyze incentives for the firm and inventor to agree on a consulting contract after the firm has accepted the TLO's contract but before they start technical development. We rule out potential shelving by the firm and we do not model the offer process by which the inventor and the firm determine the terms of the consulting contract. Rather, we study the conditions under which both parties have an incentive to engage in consulting. We assume that consulting makes the effort the inventor spends in the context of the consulting contract observable to the firm and that the consulting contract is enforceable.[13]

To analyze the incentives for consulting, note that after the firm has accepted the contract and paid the upfront fee, its continuation expected payoff before the inventor spends any effort is given by:

$$p(e^{**})[z\pi - C - M^{**}] - X \geq 0.$$

This expression is strictly increasing in the effort level (unless $X = 0$ and $p(e^{**})[q\pi(x) - m^{**} - C] = 0$), so that a firm's expected profit increases if it can increase inventor effort above e^{**}. For any $e > e^{**}$, we define $c(e)$ as this increase in expected profit

$$c(e) = (p(e) - p(e^{**}))[z\pi - C - M^{**}] \tag{19}$$

and consider the conditions under which the inventor would accept a contract offering a share of $c(e)$ if she exerts additional effort. A necessary condition for the inventor to accept a share of $c(e)$ is

$$c(e) = (p(e) - p(e^{**}))[z\pi - C - M^{**}] > (p(e^{**}) - p(e))\alpha M^{**} - (V(e^{**}) - V(e)). \tag{20}$$

That is, the increase in the firm's expected profit must more than compensate for the loss in expected payoff that occurs because effort is not optimal from the inventor's point of view.

In general, given O^{**}, the maximum increase in surplus or "gains from trade" from a consulting contract is achieved by maximizing $c(e)$ (indeed, $p(e)$) with respect to e subject to (20). If a consulting contract is feasible given O^{**}, the TLO maximizes its

[13]Note that in reality, the contract may not be enforceable in court. However most consulting relationships between a licensing firm and an inventor are not one-shot as in our model. In this case, enforceability may still occur through the potential loss in reputation from shirking.

profit by letting the firm and the inventor agree on such a contract since it increases the probability of success above $p(e^{**})$ without affecting the TLO's contract terms.[14] Thus, if consulting is allowed and there exists $e > e^{**}$ satisfying (20), it must be part of the equilibrium since it increases all three players' payoff.

Equation (20) provides insight into how consulting relates to the characteristics of the license. Since $p(e)$ is strictly concave and $V(e)$ is strictly convex, (20) is more likely to be satisfied when e^{**} is low.

7 Conclusion

University-industry technology transfer is an important part of national innovation systems and one fraught with incentive problems, largely because of the informational asymmetries and investment needed for industrial application of many university inventions. In this paper, we focus on the role of contracts, and in particular the form of payment in overcoming these distortions and argue that commonly observed license contracts can be explained by the presence of multiple distortions. Table 2 summarizes our results. We show that in the presence of moral hazard and adverse selection, simple contracts are unlikely to be optimal. Interestingly, milestones, which have largely been ignored in the literature, are the only form of payment which can work to mitigate problems with inventor and licensee effort as well as risk. is milestones. When milestones are feasible, then the only use for royalties is to mitigate risk. In our context, upfront fees serve not only to mitigate risk or extract rents, but also problems with licensee shelving.

Notice that the contracting problems we examine are predicated on the split ownership implicit in Bayh-Dole, that is, the university owns the invention but the government reserves the right to take it back in the absence of reasonable commercialization effort. We argue that this march-in provision provides the incentive for the university to execute separating contracts, so that in equilibrium actual march in would not occur.

It is the university ownership of the invention that makes our contracting problems fundamentally different from those of Aghion and Tirole [1994a]. In our model the researcher (inventor) has a moral hazard problem that does not exist in their framework where either the researcher or the customer (licensee in our case) owns the invention. However, it is well understood from principal-agent theory that if the agent is risk neutral and faces no limited liability constraints, the principal can usually fully solve the moral hazard problem by "selling" the project to the agent and extracting rent with

[14] However, since the TLO is the first-mover in the game, if it anticipates that a consulting contract will be signed between the firm and the inventor, it has an incentive to slightly increase the milestone payment (see Appendix B).

	Moral Hazard only	Moral Hazard with risk aversion *Firm*	*TLO*	Moral Hazard with Unintentional Shelving	Intentional Shelving
Fixed fees					
Upfront		+	+		+
Annual				+	
Milestone	+	+	+	+	+
Royalty		+			
Consulting*	+				

* We do not discuss consulting in the presence of shelvers or when the players are risk averse.

Table 2: Theoretical results. What should we expect to see in a contract? In each column, a "+" indicates that according to our model, the payment type is likely to be observed if the problem in the column header is a serious concern.

a fixed fee (see, for instance, Laffont [1989]). This solution is reminiscent of a commonly observed practice in university licensing, which consists of letting the inventor start up her own firm to develop and commercialize the invention. An interesting question for further research, particularly given increasing commercialization through inventor startup companies, is when it would be optimal for the university to transfer ownership to the inventor. This question has also been the topic of debate among a number of European countries where traditionally ownership has resided with the inventor [OECD 2003]. Another question, currently a point of contention between some firms and universities, is when the firm funds the research, whether firm ownership is optimal.

References

[1] Aghion, Philippe and Tirole, Jean. "The Management of Innovation," *Quarterly Journal of Economics*, 1994a, 109(4), pp. 1185-1209.

[2] Aghion, Philippe and Tirole, Jean. "Opening the black box of innovation," *European Economic Review*, 1994b, 38, pp. 701-710.

[3] Agrawal, Ajay. "Research on University to Industry Knowledge Transfer: Literature Review and Unanswered Questions," *International Journal of Management Reviews,* 2001, 3, pp. 285-302.

[4] Agrawal, Ajay and Cockburn, Iain. *LES Licensing Survey*. 2006.

[5] Arora, Ashish. "Licensing Tacit Knowledge: Intellectual Property Rights and the Market for Know-How," *Economics of Innovation and New Technology*, 1995, 4, pp. 41-59.

[6] Beggs, Alan. "The Licensing of Patents under Asymmetric Information." *International Journal of Industrial Organization*, 1992, (10), pp. 171-191.

[7] Bousquet, Alain, Cremer, Helmuth, Ivaldi, Marc, and Wolkowiz, Michel. "Risk Sharing in Licensing," *International Journal of Industrial Organization,* 1998, 16, pp. 535-554.

[8] Choi, Jai Pil. "Technology Transfer with Moral Hazard," *International Journal of Industrial Organization*, 2001, 19, pp. 249-266.

[9] Cohen, Wesley, Nelson, Richard, and Walsh, John. "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)," 2000, NBER Working Paper No. w7552.

[10] Edwards, Mark, Murray, Fiona, and Yu, Robert. "Value Creation and Sharing Among Universities, Biotechnology and Pharma," *Nature Biotechnology*, 2003, 21(6), pp. 618-624.

[11] Gallini, Nancy and Wright, Brian. "Technology Transfer under Asymmetric Information," *RAND Journal of Economics*, 1990, (21), pp. 147-160.

[12] Gilbert, Richard and Newbery, David. "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review,* 1982, 72(3), pp. 514-526.

[13] Goldfarb, Brent, and Colyvas, Jeannette. "Do Strong Property Rights Inspire Academic Entrepreneurship," 2003, mimeo.

[14] Hellmann, Thomas. "The Role of Patents for Bridging the Science to Market Gap," *Journal of Economic Behavior & Organization,* 2006, forthcoming.

[15] Hoppe, Heidrun and Emre Ozdenoren. "Intermediation in Innovation." *International Journal of Industrial Organization*, 2005, (23), pp. 483-503.

[16] Jensen, Richard and Thursby, Marie. "Proofs and Prototypes for Sale: The Licensing of University Inventions," *American Economic Review*, 2001, 91(1), pp. 240-259.

[17] Jensen, Richard, Thursby, Jerry, and Thursby, Marie. "The Disclosure and Licensing of University Inventions: The best we can do with the s**t we get to work with," *International Journal of Industrial Organization*, 2003, 21(9), pp. 1271-1300.

[18] Kamien, Morton. "Patent Licensing." in Robert Aumann and Sergiu Hart, *Handbook of Game Theory,* 1992, Amsterdam: North Holland.

[19] Lach, Saul, Schankerman, Mark. "Royalty Sharing and Technology Licensing in Universities.," *Journal of the European Economic Association*, 2004, 2:2-3, pp. 252-264.

[20] Laffont, Jean Jacques. *The Economics of Uncertainty and Information.* Cambridge, MA: The MIT Press, 1989.

[21] Macho-Stadler, Ines, Martinez-Giralt, Xavier and Pérez-Castrillo, David. "The Role of Information in Licensing Contract Design." *Research Policy*, 1996, (25), pp. 43-57.

[22] Macho-Stadler, Ines, Pérez-Castrillo, David and Veugelers, Reinhilde. "Licensing of University Inventions: The Role of Technology Transfer Office," *International Journal of Industrial Organization*, 2006, Forthcoming.

[23] Organization for Economic Co-operation and Development, *Turning Science into Business: Patenting and Licensing at Public Research Organizations,* 2003.

[24] Rai, Arti and Eisenberg, Rebecca. "Bayh-Dole Reform and the Progress of Biomedicine," *Law and Contemporary Problems,* 2003, 66, pp. 288-313.

[25] Thursby, Jerry, Jensen, Richard, and Thursby, Marie. "Objectives, Characteristics and Outcomes of University Licensing: A Survey of Major U.S. Universities," *Journal of Technology Transfer*, 2001, 26, pp. 59-72.

[26] Thursby, Jerry and Thursby, Marie. "Who is Selling the Ivory Tower: The Sources of Growth in University Licensing," *Management Science*, 2002, 48(1), pp. 90-104.

[27] Thursby, Jerry and Thursby, Marie. "University Licensing and the Bayh-Dole Act," *Science*, 2003, 301, p. 1052.

[28] Thursby, Jerry and Thursby, Marie. "Are Faculty Critical? Their Role in University Licensing," *Contemporary Economic Policy,* 2004, 22(2), pp. 162-178.

[29] Verspagen, Bart. "University Research, Intellectual Property Rights and European Innovation Systems," *Journal of Economic Surveys*, 2006, 20(4), pp. 607-632.

8 Appendix A: Survey Data

Our data come from two sources. The first is a survey of university based technology transfer professionals and the second is a survey of business executives who actively license-in from universities. The university survey was sent to the top 135 U.S. universities in terms of licensing revenue as reported in the 1996 AUTM Survey and 62 responded. The majority of universities responding were public, and of the public universities responding, 62% were land-grant institutions. Private universities accounted for 37% of the responses. Average industry sponsored research for universities in the sample was $16.9 million in 1996, and federally sponsored research was $149.6 million. The average technology office in the sample reported 26.3 licenses executed, 92.3 inventions disclosures, 30.1 new patent applications and $4.2 million in income for 1996. Compared to the 131 U.S. universities who responded to the 1996 AUTM survey, the respondents to our survey represent 68% of industry sponsored research, 75% of federally sponsored research, 71% of royalty income, 74% of the licenses executed, 70% of the invention disclosures and 48% of the new patent applications. For further details see Jensen and Thursby [2001].

The business survey was designed to be answered by individuals actively engaged in executing licenses, options, and/or sponsored research agreements with universities between 1993-1997. We received responses from 112 business units that had licensed-in university inventions. Firms in our sample accounted for at least 15% of the license agreements and 17% of sponsored research agreements reported by AUTM in 1997. Seventy-nine firms in the sample responded to a question on the top five universities with whom they had contractual agreements. The 85 universities mentioned include 35 of the top 50 universities in terms of industry sponsored research and 40 of the top 50 licensing universities in the 1997 AUTM Survey. The majority of respondents were employed by small firms, with 46% answering for firms with less than one-hundred employees and 17% for firms with more than one hundred but less than five hundred employees. In terms of industry segments, 31% of the respondents identified pharmaceuticals as the main industry in which their firm operated, 36% indicated biotechnology and medical devices as their main industry, and 33% indicated other industries. 91% of the sample conducted some R&D in-house. On average, 37% of the R&D conducted in-house was basic or discovery research, 44% was new product development, and 18% was process improvement. Finally, many of the firms in the business survey are not publicly traded, which precludes the usual tests for selectivity bias. As reported in Thursby and Thursby [2004], we used an alternative approach of comparing data on respondents with that of the general population reported in the AUTM survey as well as our earlier university survey. Other details of survey design can be found in Thursby and Thursby [2002, 2004].

9 Appendix B: Proof of Proposition 4

The proof of (i) follows from Lemma 1, the fact that (13) is necessary and sufficient for the TLO to find it profitable to search for a second firm and arguments made in the text.

To prove (ii), note first that the fact that (13) does not hold by assumption implies that the TLO cannot credibly commit to searching for a second firm upon observing a technical failure. Therefore, if a constrained PBE exists, on the equilibrium path, the TLO does not search for a second firm. Suppose then that the TLO offers a contract $O = \{M, F\}$ to the first firm. If the inventor only expects non-shelvers to work on development, then his optimal effort is given by $e^s(M)$, where $e^s(M)$ is defined in the text. It is then optimal for a non-shelver to accept the contract and work on development if and only if

$$p(e^s(M))(\pi - C - M) - X - F \geq 0. \tag{A1}$$

Since upon a rejection by the first firm, the TLO offers O^{**} to a second firm, a shelver will prefer to accept the TLO's contract and shirk if and only if

$$\pi^m - F \geq p(\hat{e}(M^{**}))\pi^d + [1 - p(\hat{e}(M^{**}))]\pi^m \iff$$

$$F \leq p(\hat{e}(M^{**}))D. \tag{A2}$$

Thus, in a pooling contract, F must satisfy the above inequality. Suppose that (A1) holds. If both types of firms accept the contract, the TLO's payoff from offering O is given by

$$(1 - \alpha)[(1 - s)p(e^s(M)) + F] - sL. \tag{A3}$$

It follows that in equilibrium (A1) is binding and (A2) may or may not bind (since M and F can always be adjusted so that (A1) binds as long as (A2) is satisfied). Solving for F from (A1) satisfied with equality and substituting in (A3) yields that the TLO's expected payoff is equal to

$$(1 - \alpha)[p(e^s(M))(\pi - C) - sp(e^s(M))M - X] - sL.$$

Differentiating with respect to M yields

$$(1 - \alpha)[p'(e^s(M))\frac{\partial e^s(M)}{\partial M}(\pi - C - sM) - sp(e^s(M))]. \tag{A4}$$

Hence, for sufficiently large s, the TLO's payoff is not monotonically increasing in the milestone payment, which implies that $F > 0$ is possible. Setting (A4) equal to zero characterizes the optimal milestone payment M^{pool} and $F^{pool} = p(e^s(M^{pool}))(\pi - C -$

$M^{pool}) - X$.

Finally, noting that in the limit as s goes to zero, the terms sM and $sp(e^s(M))$ vanish from (A4), it follows that M^{pool} converges to M^{**}, which implies that F^{pool} goes to 0 and thus, O^{pool} converges to O^{**}.

10 Appendix C: A possible consulting contract

In the absence of a consulting contract, the inventor's effort level is given by $\hat{e}(M)$, as long as M is large enough that the firm finds it profitable to invest. Hence, the benefit of a consulting contract that enforces effort level $e \geq \hat{e}(M)$ is equal to

$$c(e, \hat{e}) = (p(e) - p(\hat{e}))[z\pi - C - M]$$

A possible consulting contract between the inventor and the firm that enforces effort level e^c is as follows (where e^c will generally depend on M). The inventor receives a transfer payment equal to $c_I(e^c, \hat{e}) \in (0, c(e^c, \hat{e}))$ if and only if the firm observes effort level $e \geq e^c$. If the inventor's effort is $e < e^c$, the inventor does not receive any payment from the firm. The effort level e^c and the inventor's consulting fee $c_I(e^c, \hat{e})$ may be the outcome of a bargaining game between the inventor and the firm. In general, the size of this fee will depend on each party's respective bargaining power.

The inventor's expected payoff from accepting the contract is given by

$$\alpha p(e^c)M + c_I(e^c, \hat{e}) - V(e^c).$$

Therefore, the inventor will accept the consulting contract if and only if

$$c_I(e^c, \hat{e}) \geq V(e^c) - V(\hat{e}) - [p(e^c) - p(\hat{e})]\alpha M.$$

Moreover, if the inventor accepts the consulting contract, the firm's expected payoff is equal to

$$p(e^c)[z\pi - C - M] - c_I(e, \hat{e}) - X.$$

Hence, the firm will find it profitable to offer the consulting contract if and only if

$$c_I(e^c, \hat{e}) \leq [p(e^c) - p(\hat{e})][z\pi - C - M].$$

Finally, conditional on the inventor accepting the consulting contract, ex-ante, the TLO's expected payoff is given by

$$(1 - \alpha)p(e^c(M))M$$

which, for every given M, satisfies $(1-\alpha)p(e^c(M))M \geq (1-\alpha)p(\hat{e}(M))M$ and $(1-\alpha)p(e^c(M))M > (1-\alpha)p(\hat{e}(M))M$ if and only if $e^c(M) > \hat{e}(M)$.

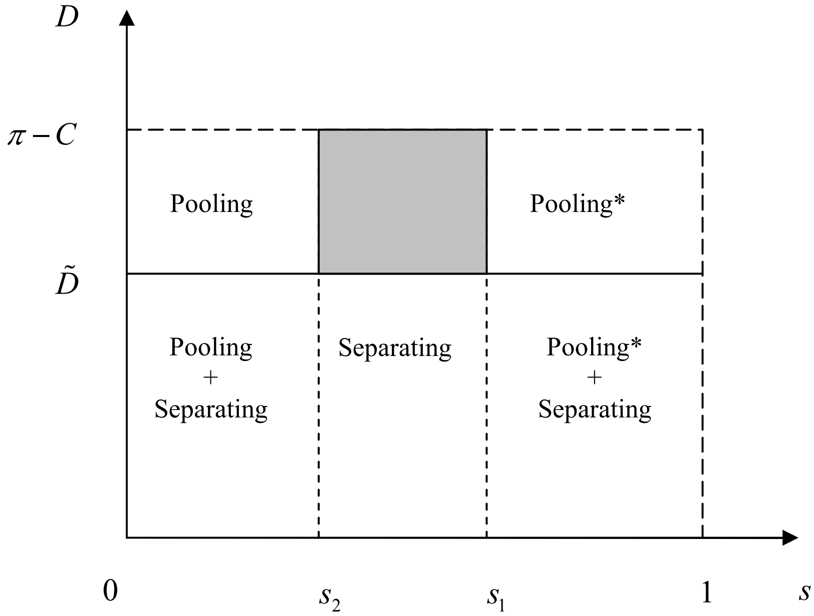

Figure 1: Values of s and D for which different types of contracts can be supported in a (constrained) PBE. Pooling* refers to an equilibrium in which the TLO takes the license back from the firm and searches for a second firm upon observing a technical failure with the first firm.