The Market for R&D Failures[*]

Roy Shalem[†] and Manuel Trajtenberg[‡]

First draft: September 2008. This version: September 2009.

[†]Berglas School of Economics, Tel-Aviv University; email: shalemroy@gmail.com.

[‡]Berglas School of Economics - Tel-Aviv University, NBER and CEPR; email: Manuel@post.tau.ac.il.

Abstract

The informational structure of R&D is characterized by the fact that firms are exposed to more information related to R&D successes of their competitors than to one that reveals their failures. This paper theoretically explores the feasibility of trading knowledge of R&D failures. The implications of constructing a market for R&D failures are substantive, in that they may decrease the effective cost of R&D, shorten the average time to discovery and increase the number of active R&D projects. We present the conditions required for a mutually beneficial trading of R&D failures information between competitors, as well as design trading mechanisms and policy implications that would account for obstacles such as moral hazard, verification, over-investment and private gains.

1 Introduction

Failures are probably the most commonly widespread by-product of R&D. Thomas Edison, for example, had tested over three thousand filaments before he came up with his version of a practical light bulb. Moreover, the idea that the accumulated knowledge consisting of past failures is what advances the future progress of the R&D endeavor, has been long standing and widely agreed upon, as conveyed by the above Edison quote. Likewise, it has also been said that "although it is obviously not patentable, even knowledge about research failures can be useful to **others**, since it may suggest novel lines of approaching a problem and at least permits avoidance of the same mistakes."[1].

Unfortunately, the unique informational structure of R&D suggests that firms are normally exposed to much less information regarding the R&D failures of their competitors, than to that of their successes. The existence of this phenomenon stems from the fact that knowledge of successes tends to disseminate, either by means of reverse engineering of new products or because of the issuance of patents[2]. Added to all that is the natural human tendency to be more expressive of one's successes than failures, which exacerbates this phenomenon.

Granted that knowledge of failures does not disseminate efficiently, it seems that it would have been quite natural to try to constructing mechanisms that would enable the existence and then the use of such flow of knowledge. Surprisingly, however, the idea of sharing or trading failures has not, as yet, been followed through - with the distinct exception of Haller and Pavlopoulos (2002) whose contribution will be discussed in more details later on. The literature shying away from this topic is in spite of the fact that the academic world does engages, from time to time, in an informal ongoing debate regarding the importance of publicizing failed research attempts.

Our paper is related to those strands of the literature that deal with R&D cooperation, such as joint ventures and open science, yet it is most closely linked to the strand that deals with licensing interim knowledge of R&D. The following paragraphs, then, will try to place our contribution within this latter part of the literature. It should be said upfront that although

---

[1]Stevenson 1980, pp.9-10. Was quoted in Yosha 2003.

[2]As a matter of fact, one of the declared objectives of the patent system is the disclosure of such successes.

this literature has studied various forms of cooperation, to the best of our knowledge, our paper is the first to analyze the potential of contracts between rivals that disclose information of failures of one agent, for the purpose of helping the R&D endeavor of his competitor, as well as discuss possible caveats that might hinder the implementation of such contracts.

There are several features that distinguish our paper from previous literature. Firstly, papers that have dealt with licensing interim knowledge, which constitute a relatively small portion of the licensing literature, have either considered some interim stages that need to be reached or very implicitly discussed some general know-how of R&D, or basic research knowledge, which could potentially be licensed. One of our contribution is to explicitly state that accumulated failures constitute the most prevalent form of interim knowledge; hence, the analysis should focus on their specific caveats as well as on their latent advantages. Since our approach explicitly examines failures, it requires particular modeling choices, which will subsequently be reflected in the results. Secondly, whereas the focus of most previous papers is on the licensing strategies and the bargaining power of the licensor versus the licensee, our paper focuses on characterizing the type of industry conditions, which makes a market for R&D failures more likely to thrive, as well as the kind of contracts that may be useful in overcoming some of the intrinsic impediments that are expected to be found in such a market.

The prominent papers dealing with licensing interim knowledge are by Bhattacharya, Glazer, and Sappington (1992), d'Aspremont, Bhattacharya, and Gerard-Varet (2000), Bhattacharya and Guriev (2006) and Spiegel (2007). In their research, d'Aspremont et al. (2000) have modeled a patent race between two firms, in which the success rate is Poisson distributed; yet, one firm exhibits a different hazard rate due to its superior interim knowledge[3]. Licensing the technology reduces the average time till the discovery is made, thus creating a positive surplus to be shared. Their paper then focuses on bargaining mechanisms for licensing that knowledge, which have some appealing properties. The other three papers assume that a few firms try to develop a cost reducing innovation, which cannot be patented. The probability of success for each firm depends on the interim knowledge it possesses. After the race is over, *Bertrand* competition takes place between the firms in the product market. Based on the nature of such product market competition, their joint value is maximized when

---

[3]Bhattacharya and Ritter (1983) have considered a similar setup for studying a model of signalling with partial disclosure in financial markets.

exactly one firm wins the competition, a fact that ultimately affects the licensing decisions of the firms. The interim knowledge in these papers has an effect solely on the chance of success[4]. Bhattacharya et al. (1992) assume that the interim knowledge is produced within a research joint venture of N firms. The researchers characterize efficient mechanisms which ensure that the interim knowledge would be shared efficiently, and at the same time induce the firms to maintain the right levels of R&D intensity in the first stage. Bhattacharya and Guriev (2006) assume that the interim knowledge is held by an outside research lab, which could license it to two firms, and that knowledge tends to partially leak. Their paper focuses on the tradeoffs between patenting and licensing the knowledge or alternatively maintaining it as a trade secret. Finally, Spiegel (2007) uses this setup with three firms, in order to study how the existence of a third firm affects the outside options of the other two, and consequently the outcome of the bargaining game between them.

Bhattacharya et al. (1992), d'Aspremont et al. (2000) as well as Spiegel (2007) - except in one of the extensions he considers - assume that the interim knowledge of the firms can be ordered in a Blackwell sense, i.e. whenever any two firms have different interim knowledge, then the knowledge of one is a subset of the knowledge of the other. As will be shown, this assumption is not necessarily plausible when the interim knowledge being considered is a stock of failures.

A distinct exception in this literature is Haller and Pavlopoulos (2002), who have indeed recognized the potential commercial value of knowledge of past failures, that, if properly exploited, can serve in preventing future failures. Their research focuses on showing how an outside research lab, that has accumulated some failures in the past, could extract the highest rent from current participants of a patent race. Though this question, per se, is an interesting and innovative one, the fact that they have focused their attention on such a setup, trivially implies that the research lab would want to sell its knowledge. Furthermore, although they do analyze, in one of the extensions, whether the lab would choose to get in the race or not, they do not explicitly analyze whether a market for R&D failures could possibly serve rival firms, nor do they expose the factors that could potentially hinder the emergence of such a market. Moreover, as their analysis begins only after the failures have already been accumulated, they

---

[4]Bhattacharya and Chiesa (1995) have considered a similar setup, with only two firms, in order to compare and contrast two alternative financing arrangements, namely the loan market vs. bilateral bank-borrower ties.

ignore the question of whether there is a potential distorting effect of such a transaction on the research effort of the lab.

Our paper is divided into three main parts. In the first part we theoretically explore the merits in establishing a market of R&D failures. The economic value of the knowledge of R&D failures, as shown in this section, is a positive result that, in turn, gives rise to the possibility that a market for that knowledge could be constructed. Having establishing that, we then confront the question of why such markets have hardly yet invaded our reality. In the second part of the paper we examine some of the theoretical obstacles that might impede the emergence of such a market, and also attempt to tackle some of these obstacles, perhaps as a first step in an effort to introduce a market for R&D failures into the economic reality. Finally, the third part suggests a few alternative ways of disseminating knowledge of R&D failures, by means other than direct trading of that knowledge, as well as illuminates the role governmental policy could play in that context.

2   The Economic Rationale for a Market for R&D Failures

There are at least two potential social merits in a market for R&D failures. The first is the ability to avoid repeating the same mistakes others have already made, including the ability to avoid certain catastrophes or other costs which are associated with failures. The second is the ability to update one's beliefs on the technological space, based on known R&D failures. These potential benefits may both reduce the average cost of R&D, as well as shorten the average time required in attaining success. Since both the effective cost of R&D and its potential gains are higher, the existence of a market for R&D failures may also ultimately imply that more R&D would be conducted in equilibrium.

This section explores a framework of a patent race, so as to illustrate those aspects related to the merit of avoiding repeating mistakes, and to analyze the question of whether trading in R&D failures may be possible in this context. Though the second merit will be touched upon in the subsequent section, more rigorous analysis of this and other merits, as well as their exploration in non-trivial economic contexts, other than patent races, is left for future work.

As mentioned in the introduction, it is quite reasonable to assume that knowledge of past failures allows for better progress in R&D, and yet the patent race literature, by and large, starting from Loury (1979), has taken quite a different approach. In order to obtain closed-form solutions in such dynamic setups, this literature has almost exclusively adopted distributions with memorylessness properties, such as Poisson, in describing the time when success should occur. Since distributions of this kind are equivalent to sampling with replacement, using them implicitly assumes away the value of failures; thus, the literature has strayed away from analyzing the possible economic value of knowledge of R&D failures. In recent years, however, numerical models have started to become more acceptable in the Industrial Organization literature, and so some change in approach can be seen, as in Doraszelski (2003) who has examined the idea of knowledge accumulation.

An exception in this line of research is the paper by Fershtman and Rubinstein (1997). They have studied a simple game-theoretic model of searching for a prize, hidden in one of many boxes - a framework constructed as an analogy to a patent race. Their simple setup also assumes that there is neither per-period nor per-sampling costs of R&D. They have found three intrinsic inefficiencies in patent races; one of which has to do with the fact that the different players might frequently search the same empty box. This property is attributed to the fact that none of the players attains any information on the past failures of the others. The researchers, however, do not attempt to offer a mechanism that would solve the problem. Haller and Pavlopoulos (2002) have used a similar framework, yet with per-sampling cost of R&D, for analyzing optimal licensing contracts of failures in a patent race by an outside player. We use a very similar economic framework for the purpose of checking whether participants in a patent race would be inclined to trade the knowledge of their accumulated R&D failures.

2.1 The Baseline Model - A Simple Patent Race Model with Failures Trading

Assume that two pharmaceutical firms are involved in a patent race for the development of a new drug, and that only the first firm to develop the drug will be able to patent it, and then receive a lump sum of V. If both firms successfully finish at the same time, each will have an

equal chance for patenting the drug. The R&D process involves testing the effect that each of s possible chemical substances has on the disease, where $s > 1$. For simplicity assume, for now, that it is common knowledge that exactly one of these substances will be effective in treating the disease, and that each of the substances has an equal chance of being the effective one. In every period, each firm can test exactly one substance at the cost of c. Let n_i denote the number of substances that firm $i \in \{1, 2\}$ has tested until time t, which is the time when the analysis begins. Without loss of generality, we assume that at the given time t, the stock of failures of firm 1 is at least as large as that of firm 2; namely that

$$n_1 \geq n_2. \tag{1}$$

For simplicity, assume that both firms are risk natural. The time preference of the firms between any two consecutive periods will be represented by a discount factor, $\delta \in (0, 1]$.

We assume that the knowledge of its failed attempts is private information to each of the firms. We also assume that the sampling is done randomly, when the set is restricted to substances which have not yet been tested by the firm. Appendix 1 shows that sampling randomly by both firms is an equilibrium behavior - though, possibly, not the only equilibrium. It also shows that in equilibrium no firm will choose a pure sequence, because of the ability of its rival to preempt its attempts. This implies that in any equilibrium there will always be some inefficiency resulting from the possibility that a firm might test a substance which has already been tested by its rival.

At each stage of the patent race, the expected profits of the two firms, denoted by π_1 and π_2, are both functions of the size of the stock of failures held by each one of them:

$$\pi_1\left(n_1, n_2\right) = \frac{V}{s-n_1}\left(\frac{s-n_2-1}{s-n_2} + \frac{1}{s-n_2}\frac{1}{2}\right) - c + \frac{s-n_1-1}{s-n_1}\frac{s-n_2-1}{s-n_2}\delta\pi_1\left(n_1+1, n_2+1\right), \tag{2}$$

$$\pi_2\left(n_1, n_2\right) = \frac{V}{s-n_2}\left(\frac{s-n_1-1}{s-n_1} + \frac{1}{s-n_1}\frac{1}{2}\right) - c + \frac{s-n_1-1}{s-n_1}\frac{s-n_2-1}{s-n_2}\delta\pi_2\left(n_1+1, n_2+1\right). \tag{3}$$

In each of the payoff functions, the first term represents the probability that the firm would be successful in the current period, times the expected payoff from success, given the probability of a simultaneous discovery. The second term is the per-period cost of R&D. The third term is the probability that no success will be made in the current period by either firm, multiplied by the discounted value of the expected payoffs in the rest of the race.

The recursive form of the payoff functions can be transformed into a normal form, as shown in Appendix 2. For the case in which $\delta = 1$, the payoff functions are as follows:

$$\pi_1 \left(n_1, n_2 \right) = V \frac{s + n_1 - 2n_2}{2(s - n_2)} - c \left[1 + \frac{(s - n_1 - 1)(2s + n_1 - 3n_2 - 1)}{6(s - n_2)} \right], \tag{4}$$

$$\pi_2 \left(n_1, n_2 \right) = V \frac{s - n_1}{2(s - n_2)} - c \left[1 + \frac{(s - n_1 - 1)(2s + n_1 - 3n_2 - 1)}{6(s - n_2)} \right]. \tag{5}$$

In each of the payoff functions, the first term is the value of the patent, times the overall probability of being the winning firm. The second term is the per-period R&D cost, times the expected duration of the race.

We assume that V is large enough, relative to c, such that $\pi_1, \pi_2 > 0 \;\; \forall n_1, n_2$. A sufficient condition for that is:

$$V > 2sc. \tag{6}$$

The payoff function of each firm is an increasing function of the size of its stock of failures, and a decreasing function of the size of the stock of failures of its rival. Disclosure of research failure to a rival firm might, obviously, result in an increase in the probability that the rival firm will win the race. Hence, the firm that performs the disclosure incurs an ex-ante loss. There could be a potential economic added-value of disclosure, if, and only if, this loss is more than being offset by the ex-ante increase in the profit of the rival firm. The existence of a positive economic added-value implies that there are potential gains from trade. It also implies that there is a non-degenerate interval of market prices which corresponds to such potential transaction between the two firms - one which can improve the ex-ante profit of both. We shall next show that such transaction is actually always feasible - except in one trivial case - which implies that the incompleteness of markets for research failures is a source of economic inefficiencies.

The disclosure process is modeled as follows. Suppose, for the simplicity of the analysis, that both firms have hired a third party, to whom all the failures are disclosed. This third party could then identify those failures that have been encountered by one firm, but not by its rival. In this baseline model only those failures could potentially have an economic value. Let k denote the number of chemical substances, which are disclosed in such a transaction by one of the firms, but have not yet been tested by the other firm. Let Π denote the joint expected profits of the two firms, i.e. $\Pi \equiv \pi_1 + \pi_2$. The disclosure value,

when firm 2 is doing the disclosure, is $\Pi\left(n_1 + k, n_2\right) - \Pi\left(n_1, n_2\right)$. When $\delta = 1$, this term is simply:

$$\Pi\left(n_1 + k, n_2\right) - \Pi\left(n_1, n_2\right) = \frac{ck\left(k + s + 2n_1 - 3n_2\right)}{3\left(s - n_2\right)}. \tag{7}$$

However, when firm 1 is the one to disclose its failures, it is sometimes possible that following the disclosure, firm 2 will take the lead. Explicitly, condition (1) might not hold, following the disclosure. When this is not the case, and condition (1) does hold following the disclosure, then the value of disclosure will be $\Pi\left(n_1, n_2 + k\right) - \Pi\left(n_1, n_2\right)$. When $\delta = 1$, this term is simply:

$$\Pi\left(n_1, n_2 + k\right) - \Pi\left(n_1, n_2\right) = \frac{ck\left[\left(s - n_1\right)^2 - 1\right]}{3\left(s - n_2\right)\left(s - k - n_2\right)}. \tag{8}$$

When one firm discloses its failures to its rival, the expected duration of the race may be shorter, due to the ability of the rival to avoid testing failing substances. This ability, in turn, may induce gains from trade, which are both due to the possible reduction in total expected costs of R&D, and to an earlier discovery of the right substance. The following proposition summarizes the cases in which trades can take place.

Proposition 1 (The feasibility of a market for R&D failures): *In a two-player patent race, there is a non-degenerate interval of prices for which the two firms would like to trade information of failed research. The only exception is when $s = n_1 + 1$, in which case firm 1 has only one substance left to test. The gains from such information trading is monotonically increasing with the amount of disclosed failures.*

Proof. See Appendix 3.

The important implication of Proposition 1 is that, unless firm 1 is already in the last stage of the race, there is always a range of prices, for which failures can be traded, and gains from trade will be positive. This proposition also implies that full disclosure, of all the failures accumulated by both firms, is likely to happen. Hence, these results constitute a strong indication supporting the feasibility of a market for R&D failures. Note that other forms of cooperation in patent races are not always that robust. Silipo (2005), for instance, shows that research joint ventures are not likely to be formed when the firms are in different positions in the race, nor when cooperation following the discovery is impossible or weak.

When $\delta = 1$, the value of disclosure is not a function of the value of the invention, V, but rather that of the cost of development, c. Hence, the entire value of disclosure is created from the cost reduction of the R&D process. Yet, it does not imply that the upper bound of the price of a single failure is c. In fact, from the standpoint of the firm purchasing the knowledge of a failure, it may be much more valuable than c, since the disclosure also affects the probability that the firm will win the race. For instance, in the case in which the disclosure is done by firm 2, the following terms exhibit the maximum price that firm 1, the buyer, would be willing to pay, versus the minimum price that firm 2, the seller, would charge:

$$\pi_1 \left(n_1 + k, n_2\right) - \pi_1 \left(n_1, n_2\right) = \frac{k \left[c \left(k + s + 2n_1 - 3n_2\right) + 3V\right]}{6 \left(s - n_2\right)}, \tag{9}$$

$$\pi_2 \left(n_1 + k, n_2\right) - \pi_2 \left(n_1, n_2\right) = \frac{k \left[c \left(k + s + 2n_1 - 3n_2\right) - 3V\right]}{6 \left(s - n_2\right)}. \tag{10}$$

If the cost of development was zero, for the case in which $\delta = 1$, the R&D race would actually be a zero-sum game. In that case there would not be any means for creating an economic value through the disclosure of failures in this baseline model. Indeed, the only exception to full disclosure, in our model, is the case where the condition $s > n_1 + 1$ does not hold, which implies that firm 1 is known to be only one period away from discovery. This, of course, implies that the disclosure of any failed result can not possibly shorten the time till the discovery is made. Hence, there is no potential cost reduction by means of disclosure, but rather just a possible different split of the value V between the two firms.

By contrast, when $0 < \delta < 1$, the value of disclosure is always a positive function of the term $[(1 - \delta) V + 2\delta c]$. This implies that the value of disclosure is comprised of both the reduction in the research costs, and the increase in the discounted value of discovery (i.e. V), caused by the shortened average time until the discovery is made. As δ decreases, the weight of V, relative to c, increases in the disclosure value. Of course, if we considered total welfare, shortening the average time until the discovery is made, would also imply a positive welfare effect on the potential users of that drug.

It is worth discussing the practicality of the assumption that only relevant chemical substances are disclosed, and in that context, the need for a third party to mediate the transaction. Note that both firms are indifferent to the disclosure of an irrelevant substance, that is, one which was originally included both in n_1 and in n_2. Still, the value of the disclosure is

ex-ante a stochastic variable, since the number of relevant substances is unknown to each of the firms prior to disclosure. However, as both firms agree ex-ante on the distribution of k, it follows that they can agree on a price, even when k is ex-ante unknown. This implies that the role of the third party, in making sure that no irrelevant substances are disclosed is, in fact, redundant in such transactions. At this point, however, a third party is still needed in order to verify that every failure disclosed, does indeed represent a substance which has been tested - an issue that will be discussed in the next section.

The rest of the paper refers mostly to the case in which $\delta = 1$, as it is a much simpler case to analyze. The qualitative results of this paper apply to the more general case as well. However, it should be noted that as we abstract away from the element of time in the model, important related issues, such as the implications on sequential innovations, are also put aside.

3   Potential Obstacles in Implementing a Market for R&D Failures

Having thus far established some of the possible merits of markets for R&D failures, we next turn to probe the question of why such markets are almost nonexistent in reality. Their scarcity suggests that in reality there are probably some obstacles that hinder attempts to trade R&D failures. Actually, trading any kind of interim knowledge should prove to be challenging, as foreseen by Kenneth Arrow in his 1962 seminal work[5]: "To appropriate information used as a basis for further research is much more difficult than to appropriate its use in producing commodities; and the value of information for use in developing further information is much more conjectural than its use in production and therefore much more likely to be underestimated, so that if a price is charged for the information, the demand is even more likely to be sub-optimal.".

Moreover, knowledge of R&D failures poses its own challenges, stemming from the specific characteristics of failures. One such trait, which differentiates failures from successes, is that the former can be relatively easily and inexpensively generated. A second

---

[5]Pg. 14

challenge is due to the fact that, unlike most commodities, when knowledge of failures is bought, the purpose of the purchaser is not to directly use the failures, but rather to be able to avoid them. This detachment of the buyer from the product may induce the seller to try and deceive the buyer by lowering the quality of the information sold. Finally, another difficulty results from the lack of structural Intellectual Property Rights (IPRs) over R&D failures, the kind of protection provided by patents to R&D successes.

In this section we shall revisit the aforementioned model, of trading in knowledge of R&D failures in a simple patent race, where the existence of a market for R&D failures has been found to be very plausible. We shall relax some of the assumptions and otherwise alter this simple model in order to surface and expose possible difficulties which are likely to occur in reality, when trying to trade knowledge of R&D failures. In addition, we shall explore the kind of contracts that could serve the trading parties in overcoming some of the difficulties they are likely to encounter.

3.1   Moral Hazard: Contract-Based Approach

Whenever one of the firms in our patent race is paid to disclose its failures, it has an incentive to deceive its rival. Thus, the firm may be tempted to report as failures, substances it has never actually tested. The possibility of such a hidden action, by the seller, might, consequently, lead to a potential impediment in the performance of the market. Hence, the challenge is to design contracts ensuring that the sold information, regarding past failures, is indeed accurate.

Suppose that in our baseline model one firm inaccurately reports a set of allegedly failing substances to its competitor, which in fact includes the successful substance. Since it is common knowledge that one success in fact exists, it follows that one of the firms will eventually find and then patent this substance, thus unraveling the lie of the untruthful seller. Thus, there is a one-to-one link, in the baseline model, between the hidden action of a deceptive seller, i.e. his potential lie, and a verifiable and contractible element, i.e. that the right substance will ultimately be patented and be unraveled to all. This link implies that the hidden action in the baseline model is, in fact, hidden only for a limited time, making this problem simple to overcome, contractually, as will later be explained.

Since the moral hazard problem of untruthful disclosure does not fully present itself in the baseline model, in this section we would like to deviate from our baseline model and discuss more generally the way in which such moral-hazard problems may affect the market for R&D failures. We would like to account for extensions such as one in which the probability that a success would eventually be found might be less than one, due, for instance, to budget constraints. The importance of this type of extension is in its implication that there might be a positive probability that a fraud would never be detected. Consider, for example, a case in which a hundred substances have not yet been tested, while each of the firms has enough resources to check only two more substances. In this case, if one of the firms tries to deceive the other by selling knowledge of a substance it has never tested, both firms might then proceed to test other substances and the fraud would never be detected. Under such circumstances the buyer would be forced to engage in active auditing in order to ensure truthful revelation by the seller.

Another type of extension, to which this section is applicable, involves a hidden effort-level in the research endeavor of the seller. For instance, suppose that, like in the baseline model, the per-testing cost, c, reflects the testing cost when the firms exert the maximum effort. However, suppose that, contrary to our baseline model, the firms could also choose a lower effort level, such that the cost would be lower, but so would the probability that if the successful substance were tested, it would be identified as such. This type of extension exhibits the aforementioned feature of failures - the relatively easy and inexpensive way they can be generated. In such a model even if the seller could prove that he had actually performed the tests and failed, the quality of the information sold would still need to be contractually addressed.

The following analysis focuses on comparing the performance of a market for R&D failures, in which deceiving the rival firm is possible, to a similar, counterfactual, market, in which deception is impossible. We shall denote the outcome in the market with no deception "the first-best outcome". Consider, then, two competing firms, such that firm j would want to sell knowledge of failures to firm i. Let π_i and π_j denote the expected payoffs of firm i and firm j prior to the transaction, and π_i^* and π_j^* denote their payoffs after the transaction, given that the seller, j, is truthful. For such a transaction to be feasible in the counterfactual world

with no possibility of deception, the net disclosure value has to be positive, meaning that:

$$\pi_i^* + \pi_j^* - \pi_i - \pi_j > 0. \tag{11}$$

By contrast, in a world in which it is possible for the seller to deceive the buyer, let π_i^{**} and π_j^{**} denote the payoffs to the firms, should firm i acts as if firm j is truthful, yet firm j deceives firm i by selling some fake test results from tests it has never performed. Moral-hazard implies that:

$$\pi_j^{**} > \pi_j^*, \pi_i^{**} < \pi_i^*. \tag{12}$$

Let $P_S \in (0, 1]$ denote the probability that if the seller is deceitful, a success will, in fact, exist amongst the sold, so called, failures. We assume, as we did in the baseline model, that with probability 1, none of the failures sold by a truthful seller could potentially be a success. When the seller was deceitful, let P_N denote the probability that a fraud would eventually be detected, i.e. the probability of detection, given that a success, in fact, exists amongst the sold results, and assuming that the buying firm believes that the selling firm is truthful. Let P_U denote the probability that a fraud would be detected if the buying firm pursued the research strategy it had intended to pursue before the contract was signed. We assume that:

$$P_U > 0,$$

which implies that prior to the transaction firm i had a positive probability of attaining, on its own, the test results revealed to it by the transaction. Note that in the baseline model, for instance, $P_N = P_U = 1$.

The following proposition summarizes the outcome when $P_N > 0$, and a contingent contract is signed:

Proposition 2 (No auditing): *If deception is possible by the seller, and $P_N > 0$, then the first-best outcome is achievable by imposing a finite state-contingent fine.*

Proof. See Appendix 4

Proposition 2 demonstrates that when $P_N > 0$, the fine can be made contingent upon detecting a fraud, without actively participating in auditing, as there is a strictly positive probability that the fraud will naturally be exposed.

Before proceeding to the case in which $P_N = 0$, we would like to comment on another extension, which offers an alternative, and perhaps a more realistic, assumption, that even when the seller has been truthful there could still be a certain positive probability that one of the disclosed failures would, in fact, be a potential success. This might be the result of a setting in which, when the right substance is tested, there is a positive probability that the test will go wrong and inaccurately indicate that the substance is a failure. Now, even if the seller is truthful, there is still a positive probability $\hat{P}_S \in (0, 1)$ that one of the disclosed failures is in fact a potential success. This implies that there is a strictly positive probability that he will still have to pay the fine - what is sometimes referred to as a type-1 error. The following two constraints - IR_j and IR_i, respectively - ensure that both firms will agree to participate in the transaction, given that the seller is truthful:

$$\pi_j^* - \pi_j + p - P_N \hat{P}_S F \geq 0 \tag{13}$$

$$\pi_i^* - \pi_i - p + P_N \hat{P}_S F \geq 0, \tag{14}$$

which simply implies that, in such a model, the upfront payment should also compensate the truthful seller for the probability that he might have to pay the fine. The following constraint (IC_j) ensures that firm j, the seller, will not wish to deceive the buyer:

$$\pi_j^* - \pi_j + p - P_N \hat{P}_S F \geq \pi_j^{**} - \pi_j + p - P_N P_S F, \tag{15}$$

which is also not difficult to satisfy, as long as $\hat{P}_S < P_S$, meaning that substances which have been tested and failed are more likely to actually be failures than those which have not been tested yet. Thus, since the fine is not an exogenous social cost, but rather a side payment between the two firms, and since the firms are risk-neutral, the fine has no net social cost. Its effect washes out, as the price of the transaction, p, compensates the truthful seller for the probability that the disclosed test result would turn out to be a success.

Examining the incentives of the buyer, reveals a more substantial difference from the model in which $\hat{P}_S = 0$. By distorting his own research plan, it might become possible for the buyer to affect the probability that the fine would be paid by the seller. Any such deviation by

the buyer will yield some $\pi_i^D \leq \pi_i^*$, and increase the probability of discovering a success, if it exists amongst the disclosed substances, to $P_D \geq P_N$. For any feasible pair $\{\pi_i^D, P_D\}$ the IC constraint for firm i implies that for any such deviation, the following condition must hold:

$$\pi_i^* - \pi_i - p + P_N \hat{P}_S F \geq \pi_i^D - \pi_i - p + P_D \hat{P}_S F, \tag{16}$$

which is equivalent to:

$$F \leq \frac{\pi_i^* - \pi_i^D}{(P_D - P_N) \hat{P}_S}. \tag{17}$$

This condition will always be satisfied if, like in our baseline model, success is guaranteed to occur, that is if $P_N = 1$. In other cases, this upper bound constraint on the size of the fine, F, might ultimately hinder the implementation of the first-best outcome. The mechanism, which will be described next, of "auditing via undoing", will be able to deal with this problem, since the fine will only be potentially imposed when the auditing procedure is being exercised, in which case the buyer will be left unaware of the specific failures involved in the transaction.

Let us examine now the case in which $P_N = 0$, where without some form of active auditing, a deceitful seller will never be identified as such. In our model the natural way for the buying firm to verify that the substances disclosed to it are indeed failures, is by testing them. We shall consider a specific form of auditing in our model, such that requires the buying firm to sometimes revert to the research plan it had prior to the transaction. Evidently, in such an undoing procedure, there is an ex-ante strictly positive probability that the seller would be caught, since $P_U > 0$. Let P_A denote the probability that such an auditing procedure would be implemented[6]. The following proposition summarizes the outcome of implementing such a contract:

Proposition 3 (Auditing via undoing): *If deception was possible by the seller, and $P_N = 0$, then an outcome infinitely close to the first-best would be attainable by imposing an infinitely large state-contingent fine. A similar implementation with a finite fine would result in an*

---

[6]Note that we have assumed that firm i could commit to the auditing process. This could be implemented by keeping the sold knowledge, sealed in an envelope at the hand of a third party. This third party would then perform a lottery in order to determine whether auditing should be implemented. If the auditing has to be implemented, then the sealed envelope will be disclosed to the buyer only after the race is over, in order to find out if the fine has to be imposed.

outcome different from the first-best, though it would not distort the transactions made in the market for R&D failures.

Proof. See Appendix 5.

Proposition 3 presents an approach not too different from the "high fine-low probability" result by Becker (1968)[7], since in order to incentivize the seller there is a need to balance between the fine and the auditing probability, while only the latter bears a real social cost. It is, however, quite unique that in our model even with the imposition of a finite fine, the transactions are left unchanged. This feature results from the fact that the auditing costs in our model are not exogenous costs, but instead they merely undo the transaction. This implies that the net utility resulting from the transaction does actually drop, but not so much as to make the transaction unprofitable to the firms.

3.2 Moral Hazard: Reputation-Based Market

The theory of repeated games has paved the way for the emergence of some alternative theoretical solutions, for the establishment of long-term mutually beneficial trades. We mention this alternative approach, as it is not an uncommon practice in many markets in reality in addressing various moral-hazard problems. In fact, in recent years, increasing attention has been given to these theoretical ideas, as e-commerce has developed and some of these mechanisms have been implemented in trading platforms, such as eBay. However, as mentioned above, knowledge of R&D failures, as a commodity, has its own special characteristics. For example, unlike the reputation system used by eBay, in the case of selling knowledge of R&D failures, the buyer is often left unaware of the actual quality of the object of purchase.

Reputation-based mechanisms are especially important in cases where direct mechanisms cannot be used, such as when contracts cannot be legally enforced. However, a model

---

[7]Becker's model of crime fighting shows that since fines are costless transfers, between the convicted offender and the government, and detection has a net cost to society, the government should set the fine to be equal to the offender's entire wealth, and complement it with the appropriate probability, in order to achieve optimal deterrence. That way the government still provides for optimal deterrence, but saves resources on law enforcement.

of reputation-based market for R&D failures, might technically be very challenging, as each phase of our patent race is not identical to the next, making it difficult to be modelled as a repeated game. In the reminder of this section we will make some simplifications in order to overcome this problem, and demonstrate a formal model of a reputation-based market.

We assume that two firms are engaged in an infinite sequence of R&D projects, all of which have identical *ex-ante* properties. For simplicity, assume that the projects are engaged in sequence, one after the other, and that there is never overlapping between projects. Let each such R&D race be denoted as a "stage game". The time span between the starting time of any two consecutive projects is assumed to be constant, and in particular independent of the time when each project terminates, as well as of the actions of the two firms. Let $\hat{\delta}$ denote the discount factor, appropriate for discounting this time span. For technical simplicity, assume that all the payments and costs related to each project, are actually made at the point in time when the next project begins.

The following definitions will be used for denoting the possible payoffs for firm $i \in \{1, 2\}$ from each R&D project. Let π_i denote the expected payoffs for the firm, given that each firm does not disclose any R&D failures to its rival. Let π_i^* denote the net expected payoffs for the firm, given that it is common knowledge that the firms do not try to deceive one another, and that R&D failures are optimally traded[8]. Let π_i^a denote the expected payoff that the firm can attain, when it chooses to cheat its rival by picking action, a, out of the set of its possible cheating actions, and given that the rival expects firm i not to cheat. Obviously, it is assumed that $\forall a, \pi_i^a > \pi_i^* > \pi_i$, which is implied by the moral-hazard problem.

Following Green and Porter (1984) and Rotemberg and Saloner (1986), we will look for an equilibrium with trigger strategies, that enforces cooperation between the two firms. The game alternates between two possible "phases". In the "cooperative phase" the knowledge of all the R&D failures is traded by the two firms, as if there was no moral-hazard problem. Once a fraud is detected, the game then switches for a fixed number of projects to the "punishment phase", in which no trade takes place. Each firm faces a stationary two-state (cooperation and punishment) Markov dynamic programming problem, in which it is an optimal policy not to trade knowledge of R&D failures in punishment periods, and to truthfully

---

[8] π_i^* would be the expected payoff for firm i, if there were no moral-hazard problem.

trade R&D failures in cooperation periods, as if there was no moral-hazard problem.

It is important to state, that we assume that the contract-based approach cannot be implemented. Hence, the moral hazard problem implies that the Nash equilibrium of each stage game, i.e. each R&D race, is such with no disclosure at all. Thus, one possible Nash equilibrium of the entire game would be such with no disclosure, in any of the races, with an expected payoff of π_i for firm $i \in \{1,2\}$, in each race. This ensures that if a defection from cooperation actually occurs, the non-defecting firm will agree to mete out the proposed punishment. This conventional way of implementing the punishment ensures sequential rationality; hence, we also restrict attention to strategies of this kind.

Let P_F^a denote the probability that a fraud of type a will be detected, either naturally, or by means of auditing, such as the one which was specified in the previous section. Moreover, we assume that the probability of an honest mistake equals zero.

Proposition 4 (Reputation-based market): *If the probability of an honest mistake equals zero, and the firms are patient enough, such that the following condition holds, $\forall a$:*

$$\hat{\delta} \geq \frac{\pi_i^a - \pi_i^*}{P_F^a \left(\pi_i^* - \pi_i\right) + \left(\pi_i^a - \pi_i^*\right)}, \tag{18}$$

then the reputation-based market is as efficient as the best contract-based implementation.

Proof. See Appendix 6.

The implication of condition (18) is that if the firms are patient enough, such that the condition is satisfied, $\forall a$, then the first best is achievable. That is to say, that in equilibrium the trades would be as if there was no moral-hazard problem; there would be no attempts of deception, and the punishment phase would never take place.

It should be noted that condition (18) demonstrates that the most profitable cheating action is not necessarily the one associated with the maximal expected immediate profits, i.e. π_i^a, but rather it also critically depends on the chance that this specific action will result in a punishment, i.e. P_F^a. Therefore, condition (18) requires that the firms would be patient enough, such that the threat of punishment would offset any pair (π_i^a, P_F^a).

By contrast, whenever the probability of an honest mistake is not zero, the reputation-based mechanism, vis-a-vis the contract-based mechanism, will have some disadvantages, in terms of welfare. Similar to Green and Porter (1984), an infinite punishment phase, in this case, will not be the natural equilibrium to focus on[9]. If the firms are patient enough, a finite-period punishment phase might be enough to deter cheating altogether. However, as in the classical Green and Porter (1984) model, even if cheating is never done in equilibrium, punishment phases will sometimes appear in equilibrium, since the firms will be unable to distinguish between an honest mistake and a deviation from cooperation. Since in such a model, the punishment to the seller is formed by means of creating foregone trades, then each time a punishment is implemented it results in welfare loss. This will not occur using a contract-based mechanism, in which the fine is simply a side payment between the two firms.

Finally, it should be noted that if it is the case that both firms are active in more than one project at a time, then the punishment can also be carried out by limiting the trades of R&D failures in all projects. This kind of mechanism, though similar to the one that has just been described, might be more efficient in cases in which the discount factor is not large enough.

3.3   Endogenous Research Capacity and Over-Investment

One of the underlying assumptions in the baseline model is that the intensity of the research efforts of the firms is exogenous, and equal to one substance per period. Suppose, instead, that the firms could affect their research capacity, namely the number of substances tested per period. It is well known that in a typical patent race there might be an inherent inefficiency associated with overinvestment. Considering, then, multi-stage models, such as that of Grossman and Shapiro (1987)[10], it becomes clear that it is typically the case that the over-investment problem is exacerbated when the two rivals are running "neck to neck", i.e. when their probabilities of success is close to each other. Thus, one can speculate that perhaps

---

[9]Even if did, there would be a strictly positive probability that, at some point, an honest mistake would result in a punishment phase, hence eliminating any further trades.

[10]Similar properties appear in the patent race model of Fudenberg, Gilbert, Stiglitz and Tirole (1983), in which firms in the same position compete fiercely, dissipating the rent from innovation.

in our model trading of failures may have an effect on that issue. This section explores the question of overinvestment in our model, in connection to the existence of a market for R&D.

Though it is not straightforward, it turns out that the easiest way to account for the R&D intensity of firm 1 and firm 2, would be to change n_1 and n_2 respectively. For example, suppose that $s = 100$, $n_1 = 50$ and that, as in the baseline model, both firms have a search intensity of one substance per period. If firm 1 wanted to change its research capacity such that it would then be able to test two substances per period, it would seem as if firm 1 had 25 pairs of remaining substances, and was testing a pair at each period. In terms of the payoff functions of the two firms, which appear in equations (2) and (3), it would seem as if firm 1 has increased n_1 to being equal to 75. More generally, if firm $i \in \{1, 2\}$ wanted to change its research capacity by some factor $k_i \geq 1$, the term n_i in the payoff functions of the two firms would have to be replaced with the term $[n_i + (s - n_i) \frac{k_i - 1}{k_i}]$. A second adjustment which needs to be made is to multiply the per-period cost, c, by the R&D intensity of the firm, k_i, in order to account for the fact that c is a per-testing cost, and not a per-period cost. For instance, in the given example, each pair tested by firm 1 should cost $2c$ to test. Having accounted for all that, the payoff function, not including the costs of establishing the research capacity, and for $\delta = 1$, is now:

$$\pi_1\left(n_1, n_2, k_1, k_2\right) = V \frac{2k_1(s-n_2)-k_2(s-n_1)}{2k_1(s-n_2)} - ck_1 \frac{(s+k_1-n_1)[k_1(3s+k_2-3n_2)-k_2(s-n_1)]}{6k_1^2(s-n_2)}, \qquad (19)$$

$$\pi_2\left(n_1, n_2, k_1, k_2\right) = V \frac{k_2(s-n_1)}{2k_1(s-n_2)} - ck_2 \frac{(s+k_1-n_1)[k_1(3s+k_2-3n_2)-k_2(s-n_1)]}{6k_1^2(s-n_2)}. \qquad (20)$$

Furthermore, we shall now replace the assumption which appears in equation (1) with the more elaborated assumption stating that firm 1 has more chances of winning the race, that is:

$$\frac{s-n_2}{k_2} \geq \frac{s-n_1}{k_1}. \qquad (21)$$

Assume, for example, that $s = 100$, condition (21) allows for the case where $n_1 = 60$ and $n_2 = 70$, if the research capacity of firm 1 was two substances per period ($k_1 = 2$) and that of firm 2 was one per period ($k_2 = 1$). Consequently, firm 1 has a higher chance of winning the race, since it is as if it only had 20 more pairs to sample, which is equivalent to $n_1 = 80$, with one substance tested at each period.

It should be noted, that any intensity strictly larger than one, implies some inherent inefficiency for the firm. For instance, suppose that the intensity was equal to three, and the

right substance was one of the three substances currently being tested by the firm. Because the firm was testing all three substances simultaneously, it had to incur the testing cost, c, for all three. Alternatively, had the intensity of the firm been one, once the right substances have been tested, the firm would immediately stop its research; hence potentially enabling it to avoid the testing of one or two of the other failing substances.

We assume that in order for firm $i \in \{1, 2\}$ to build a research capacity k_i, it is required to invest $f_i(k_i)$, which is assumed to be continuously differentiable, where $f(1) = 0$ and $f(k_i) > 0$. It is assumed that the research capacity could be built up at any point, so that if the original research capacity is $\bar{k}_i$, and the firm wants to increase it to $\hat{k}_i > \bar{k}_i$, it will have to pay $f(\hat{k}_i) - f(\bar{k}_i) > 0$. However, this investment is assumed to be irreversible, in the sense that the firm cannot reduce its research capacity, once built up.

Suppose that the race starts when $n_1 = n_2 = 0$. It is obvious, then, that in a model with no discounting, the optimal intensity, from the collective standpoint of the firms, would be that one of the firms, say firm 2, would remain idle, while the other would invest in a minimum research capacity of one substance per-period. The payoffs would then be:

$$\pi_1^c = V - c\frac{s+1}{2}, \tag{22}$$

$$\pi_2^c = 0. \tag{23}$$

However, when the firms compete against each other, each has an incentive to overinvest so as to increase its chances of winning the race. For instance, in the above case, if firm 2 actively participated in the race, with the minimum research capacity of one, the payoffs would be:

$$\pi_1^p = \pi_2^p = \tfrac{1}{2}V - c\tfrac{(s+1)(2s+1)}{6s}, \tag{24}$$

while the total profits of both firms would be:

$$\pi_1^p + \pi_2^p = V - c\tfrac{(s+1)(2s+1)}{3s}. \tag{25}$$

Note that since it is assumed that $V > 2sc$, then $\pi_2^p > 0$, therefore firm 2 is definitely better off participating in the race. This, per se, is an indication of overinvestment. Moreover, once both firms participate in the race, private gains from increasing capacities even further will

21

always be greater than the collective gains:

$$\frac{\partial \pi_1}{\partial k_1} - \frac{\partial (\pi_1 + \pi_2)}{\partial k_1} = \frac{k_2 (s - n_1) [2ck_2 (s - n_1) + 3k_1 (V - cs + cn_2)]}{6k_1^3 (s - n_2)}, \quad (26)$$

$$\frac{\partial \pi_2}{\partial k_2} - \frac{\partial (\pi_1 + \pi_2)}{\partial k_2} = \frac{ck_1^2 + (s - n_1) (3V - cs + cn_1)}{6k_1 (s - n_2)}. \quad (27)$$

Since we have assumed that $V > 2cs$, then for both firms the difference between their private and collective gains is obviously positive. This difference between the private gains from research capacity increase, and the collective gain, might be translated into further overinvestment, depending on the shape of the cost structure of that investment, $f(k_i)$.

Having established that overinvestment is a feature of our patent race, we next proceed to analyze the way in which incentives to invest in research capacity could change, depending on the stock of failures possessed by each firm. The dynamic problem, which involves both uncertainty and irreversibility of investment, is left beyond the scope of this paper. We focus on the myopic marginal willingness of each firm to invest in research capacity, relative to its position in the race, which is denoted by γ_i:

$$\gamma_1 \equiv \frac{\partial \pi_1}{\partial k_1} = \frac{k_2 (s - n_1) (3V - cs + cn_1) - ck_1^2 (3s + k_2 - 3n_2)}{6k_1^3 (s - n_2)}, \quad (28)$$

$$\gamma_2 \equiv \frac{\partial \pi_2}{\partial k_2} = \frac{2ck_2 (s - n_1)^2 - ck_1^2 (3s + 2k_2 - 3n_2) + 3k_1 (s - n_1) (V - cs + cn_2)}{6k_1^2 (s - n_2)}. \quad (29)$$

Proposition 5 (Overinvestment): *Unless firm 1 is currently at the last stage of the race, then the myopic willingness of both firms to invest in research capacity is decreasing in the difference between the size of their stocks of failures, $n_1 - n_2$. Consequently, overinvestment will be exacerbated following a transaction which marginally advances firm 2*, yet keeps it the laggard. By contrast, *overinvestment will be ameliorated following a transaction which marginally increases firm 1's lead by increasing its stock of failures.*

Proof. See Appendix 7.

The following figure illustrates Proposition 5:

Figures 1
The Effect of the Size of the Stock of Failures of Each Firm on Incentives to Overinvest
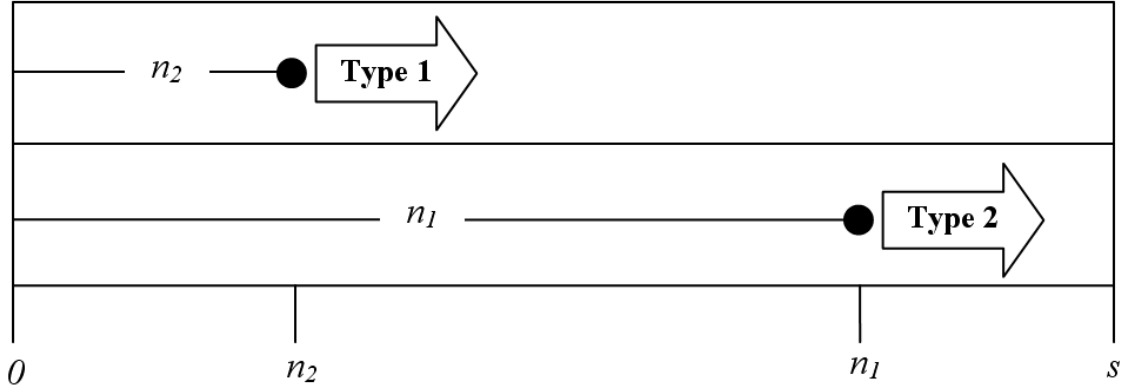
Type 1 transactions - exacerbate overinvestment.
Type 2 transactions - ameliorate overinvestment.

All in all we have found that, except for the case in which firm 1 is certain to make the discovery in the next period, a transaction which drives the laggard closer to the leader may increase incentives to overinvestment, whereas a transaction which drives the leader further away from the laggard produces the opposite effect. Though these results seem to resemble those of Grossman and Shapiro (1987), there is a significant difference, since in their model the position of each firm in the race has been measured in terms of milestones that must first be covered before moving on, whereas in our model the position of each firm is measured by its stock of failures, which is a by-product of the R&D process, not directly necessary for making the discovery. Therefore, it follows that in a transaction in which firm 1 discloses failures to firm 2, the inherent positive disclosure value may be offset by the overinvestment deadweight loss.

One simple way of overcoming the overinvestment problem is for both firms to mutually commit not to invest in a research capacity larger than one. In certain contexts this kind of a commitment naturally arises, since the cost of increasing research capacity, $f(k_i)$, is extremely high, as has been implicitly assumed in the baseline model. In other contexts, however, creating a credible commitment not to over-invest may prove to be very hard.

A second way, of contractually overcoming the possibility of overinvestment, is by

introducing side payments which are contingent upon winning the race. For instance, suppose that each firm agreed to pay a monetary sum of βV to its rival, where $\beta \in (0, 1)$, should it win the patent race. The contract might also include a non-contingent side-payment, $\alpha \in (-\infty, \infty)$, to ensure that both firms would want to sign it. This arrangement obviously dilutes the incentives of each firm to win the race, and hence to further invest in research capacity. In its extreme form, such a contract may induce one of the firms not to operate at all, and is thus not so different from an M&A arrangement. A less extreme form would simply dilute the incentives enough, so that the post-transaction willingness of the firms to invest in research capacity would equal their pre-transaction willingness to pay.

3.4 Exit

The market for R&D failures, as portrayed in this paper, is meaningful whenever more than one firm simultaneously tries to achieve a common goal by concentrating their efforts in similar directions. This context is similar to the inherent problem of patent races raised by Dasgupta and Maskin (1987), in which efficiency requires less correlation in the direction chosen by different firms. Indeed, if firms can coordinate their efforts, so as not to follow the exact same goal, their joint profit may increase. Hence, in a patent race with more than one participant, exit, and perhaps entry deterrence, may be welfare increasing. These claims would be explored in this section.

Before analyzing the effects that the market for R&D failures may have on exit, it is important to first relate to entry. As shown in Proposition 1, trading knowledge of R&D failures may, ultimately, increase the profits of both firms participating in the patent race. However, in order to discourage a potential entrant from entering the race, by reducing the perceived profits after entry, an incumbent might want to convince the potential entrant that such trades would not take place. It is true, of course, that such a threat would be non-credible, as Proposition 1 ensures that once the entrant has entered the race it would be profitable for both parties to trade their knowledge of R&D failures. The potential implication of this incentive structure is that it may induce firms to make such trades non-public. Moreover, firms active in several patent races, either simultaneously or over time, might have additional complicated reputational effects, thus ultimately turning entry deterrence into a hurdle on the

market for R&D failures.

In order to explore the effect of exit on profits, let us now drop the assumption that $V > 2sc$, which ensures that both firms would wish to stay in the race until the discovery is made, and concentrate first on a model without trading failures. Once in the race, the leader of the race, firm 1, would always want to stay in the race, since its expected per-period profits increase as the race progresses:

$$\left[V \frac{s-n_2-1-0.5}{(s-n_1-1)(s-n_2-1)} - c\right] - \left[V \frac{s-n_2-0.5}{(s-n_1)(s-n_2)} - c\right] = V \frac{(s-n_2-2)(s-n_2)+0.5(n_1-n_2+1)}{(s-n_1-1)(s-n_2-1)(s-n_1)(s-n_2)} > 0. \quad (30)$$

However, this is not the case for firm 2, as there are situations in which the expected per-period profits might decrease, leading to a point in time when the firm would choose to abandon the race[11]:

$$\left[V \frac{s-n_1-1-0.5}{(s-n_1-1)(s-n_2-1)} - c\right] - \left[V \frac{s-n_1-0.5}{(s-n_1)(s-n_2)} - c\right] = V \frac{(s-n_1-2)(s-n_1)+0.5(n_2-n_1+1)}{(s-n_1-1)(s-n_2-1)(s-n_1)(s-n_2)}. \quad (31)$$

Note that if the per-period profits of firm 2 turn negative, they will stay negative throughout the race, consequently firm 2 will choose to abandon the race at the first time its per-period profits are no longer positive.

It is also true that in this model the departure of the laggard increases the joint profits of the two firms, as it prevents the future wasteful duplication of the R&D endeavor, which is due to the possibility of the two firms testing the same failing substance. In other words, efficiency in this model calls for a single firm to be active in the patent race. The following Lemma formalizes this statement.

Lemma 6: *The joint profits of the two firms are a decreasing function of the stage in which firm 2 abandons the race.*

Proof. See Appendix 8.

---

[11]Consider the example in which $V = 100,000$, $c = 90$, $s = 100$, $n_1 = 95$ and $n_2 = 1$. The current-period expected profits of firm 2 is 9.1. However, if a discovery is not made by either firm at the current period, the next period expected profits of firm 2 will drop to -7.1, and remain negative throughout the race.

It follows from Lemma 6, that when the disclosure of failures facilitates exit, then the gains from such a disclosure will be at least as large as those of the baseline model. By contrast, when a transaction impedes exit it may, as a result, reduce the net disclosure value, compared with the baseline model, and in some cases might even turn the net disclosure value to be negative. The following two propositions summarize these insights.

Proposition 7 (Trades which are not exit-discouraging): *Consider a patent race, in which, prior to the transaction, one of the firms, i, has not planned on abandoning the race until a discovery is made. Then, (i) there is a non-degenerate interval of prices, that supports a transaction in which information of failed research is sold to firm i by its rival, as long as $s > n_i + 1$; and (ii) the value of such a transaction is monotonically increasing in the amount of disclosed failures.*

Proof. See Appendix 9.

One of the implications of Proposition 7 is that firm 2 would always be willing to sell its failures to firm 1, since, as shown above, the leader of the race, firm 1, would stay in the race until a discovery has been made. Moreover, when the number of failures accumulated by firm 2 is relatively close to that of firm 1, it is more likely that firm 2 would also be reluctant to abandon the race before all substances have been tested. Thus, if this is the case, then a transaction in which firm 1 sells its failures to firm 2 may also be possible, according to Proposition 7.

It should be noted that the gains from such transaction are comprised of both a reduction in the expected costs of the buyer, and in some cases also the chance that the seller would abandon the race earlier, due to the disclosure. An interesting implication of this analysis is that disclosure of R&D failures might induce exit, and in some cases even an immediate exit. This implies that in reality, transactions of failures disclosure might be disguised as buyouts or mergers, since what is observed is a transaction after which one of the firms immediately exits the market.

In some contexts the laggard, firm 2, would choose to abandon the race at a certain stage before all substances have been tested, as its costs of testing an additional substance are

higher than the expected gains from that test. In this case, selling knowledge of R&D failures to firm 2 may improve its R&D efficiency, thus discouraging it from exiting. The tradeoff which then arises is between the efficiency gains from failure disclosure, and the potential deadweight loss from postponing exit. When the potential of the latter force is relatively small, the market for R&D failures should still be active, without any further modifications. The following proposition summarizes whether these kinds of transactions could take place, granted that all the non-exit-discouraging trades have already been conducted:

Proposition 8 (Exit-discouraging trades): *Consider a patent race in which firm 2 has already disclosed all of its failures to firm 1, and is planning to abandon the race if and when it accumulates $\bar{n}_2$ failures. Then, the net gains from full disclosure by firm 1 are (i) negative if $\bar{n}_2 = n_2$, (ii) non-negative if $\bar{n}_2 = n_2 + s - n_1$, and (iii) increasing with $\bar{n}_2$, for $\bar{n}_2$ between n_2 and $n_2 + s - n_1$.*

Proof. See Appendix 10.

Proposition 8 reveals the difficulties in disseminating knowledge of R&D failures, which stem from the fact that it may discourage exit. However, as in the possibility of over-investment, there is a contractual way of decoupling the failures dissemination from the exit decision, by introducing side payments which are contingent upon winning the race. In this context too, it is possible to dilute the profits of the buyer so as to induce him to exit no later than at the stage in which he initially intended to exit.

The following example will illustrate the claims made in Proposition 8. Assume that the values of the technological parameters were $s = 1,000$, $n_1 = 995$ and $n_2 = 0$. It then follows that the race would go on for, at most, five more stages. If the economic parameters were such that $V = 10,000$ and $c = 8.9$, then $\bar{n}_2 = 1$, hence firm 2 would exit after one failure. If, however, it was the case that $V = 10,000$ and $c = 5.5$, then $\bar{n}_2 = 4$, hence firm 2 would exit after accumulating four failures. If firm 1 was to disclose all of its failures to firm 2, then they would both stay in the race until a discovery is made. On the one hand, full disclosure would make the R&D endeavor of firm 2 much more efficient, as it would be able to search only among the 5 remaining substances, rather than among all the possible

$1,000$. On the other hand, it would induce firm 2 to prolong its participation in the race, which entails an inherent inefficiency, as can been seen in Lemma 6. In the case in which $c = 5.5$, disclosure implies that firm 2 would stay in the race for, at most, one more stage, relative to the case of no disclosure, since even before the disclosure it had planned on testing four substances before exiting; hence, the direct efficiency gains from disclosure overcome the inefficiencies associated with exit-discouragement. However, in the case in which $c = 8.9$, disclosure implies that firm 2 would not abandon the race before all five substances are tested, instead of abandoning it after failing in one test; thus, the exit-discouragement inefficiencies would be much greater, in this case, surpassing the direct efficiency gains of disclosure.

Let us now consider the possibility of partial disclosure. Since $\bar{n}_2$ is not a continuously differentiable function of the number of disclosed substances, k, then the net disclosure value itself is a non-continuously differentiable function of k. That being the case, there may be some parameter values for which partial disclosure is possible, while full disclosure is not. As in the example discussed in the previous paragraph, if $k \leq 15$ for the case in which $c = 8.9$, the disclosure will not affect $\bar{n}_2$, hence the exit-discouragement effect will not exist[12]. The following figure illustrates, for the two cases discussed in the example, how the net disclosure value changes, as a function of the number of failures disclosed by firm 1.

---

[12]In the case in which $c = 5.5$, the exit-discouragement effect does not exist $\forall k \leq 86$.

Figure 2

Exit-Discouragement and the Disclosure Value
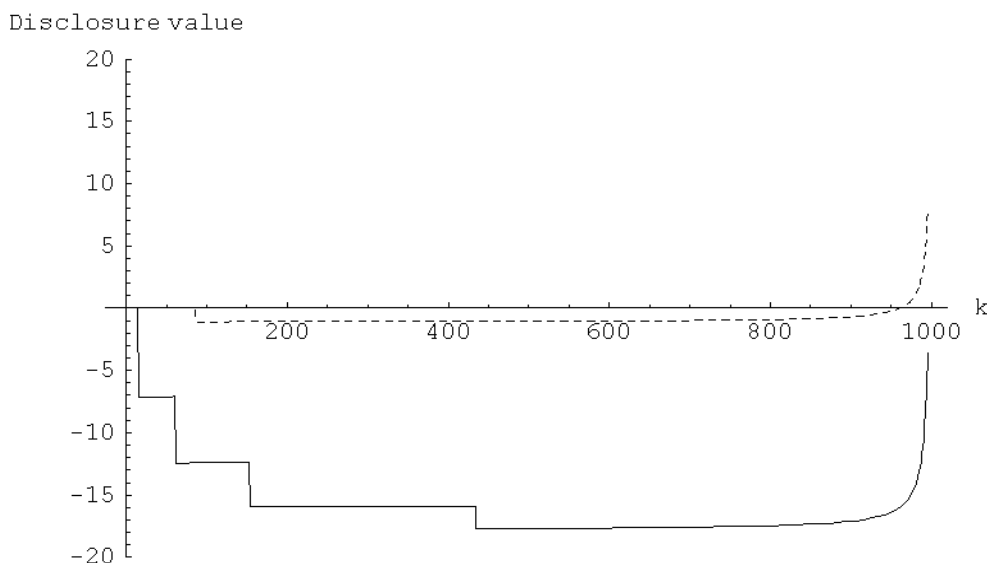
The net disclosure value, for both firms, is shown as a function of the number of disclosed substances, k. In the lower line $c = 8.9$, while in the upper, dashed, line $c = 5.5$.

3.5 Entry and Exit - with a Positive Probability of No Success

In this section we consider an extension of the baseline model that allows for a positive probability that, in fact, no success exists. It is of particular interest to apply this extension in the context studying the effect that the market for R&D failures has on entry and exit. In such a model the firms develop a certain level of expectations, regarding the probability of whether success exists within the set of yet untested substances. In addition to the effects demonstrated in the previous section, whenever failures are disclosed in this kind of model, each firm may also become either more pessimistic or more optimistic regarding this probability, which may consequently alter its incentives to exit or reenter the market.

Let $p \in (0, 1)$ denote the prior probability that one success in fact exists within the set of s substances being explored. Using Bayes' law, the firms can update their beliefs on that probability, given the knowledge of their failed experiments. Let $\hat{P}(n_1, n_2)$ denote the probability that a success exists among the substances being tested, given that n_1 and n_2

substances have already been tested and failed, by firm 1 and 2 respectively:

$$\hat{P}_{n_1,n_2} \equiv \frac{p\left(\frac{s+1}{s}\right)^2 \left(\prod_{i=0}^{n_1} \frac{s-i}{s-i+1}\right)\left(\prod_{j=0}^{n_2} \frac{s-j}{s-j+1}\right)}{(1-p)+p\left(\frac{s+1}{s}\right)^2 \left(\prod_{i=0}^{n_1} \frac{s-i}{s-i+1}\right)\left(\prod_{j=0}^{n_2} \frac{s-j}{s-j+1}\right)} = \frac{p\left(s-n_1\right)\left(s-n_2\right)}{s^2+pn_1n_2-ps(n_1+n_2)}. \quad (32)$$

Now in each period the probability of discovery by each of the firms equals $\hat{P}\left(n_1,n_2\right)$ times the probability of finding the success, given that one exists. If both firms do not exit before all the substances are tested, their expected payoff will be:

$$\pi_1\left(n_1,n_2\right) = \hat{P}_{n_1,n_2}\frac{s-n_2-0.5}{(s-n_1)(s-n_2)}V - c + \left(1 - \hat{P}_{n_1,n_2}\frac{2s-n_1-n_2-1}{(s-n_1)(s-n_2)}\right)\pi_1\left(n_1+1,n_2+1\right), \quad (33)$$

$$\pi_2\left(n_1,n_2\right) = \hat{P}_{n_1,n_2}\frac{s-n_1-0.5}{(s-n_1)(s-n_2)}V - c + \left(1 - \hat{P}_{n_1,n_2}\frac{2s-n_1-n_2-1}{(s-n_1)(s-n_2)}\right)\pi_2\left(n_1+1,n_2+1\right), \quad (34)$$

Where $\pi_1\left(s,n_2\right) = \pi_2\left(s,n_2\right) = 0, \forall n_2$.

As the race advances, the expected per-period profits may decline. This is due to the fact that the firms ascribe smaller probability to the state in which a success in fact exists within the explored set. Had the firms been aware of the union of failures both have accumulated, which is of size $n \in [n_1, \min(s, n_1 + n_2)]$, then the probability ascribed to success existing within the set would have been:

$$\hat{P}_n \equiv \frac{p\left(\frac{s+1}{s}\right)\left(\prod_{i=0}^{n} \frac{s-i}{s-i+1}\right)}{(1-p)+p\left(\frac{s+1}{s}\right)\left(\prod_{i=0}^{n} \frac{s-i}{s-i+1}\right)} = \frac{sp-np}{s-np}. \quad (35)$$

This probability may be either higher or lower than $\hat{P}_{n_1,n_2}$. This fact implies that full disclosure of the failures by both firms, for example, may either make each of the firms more optimistic regarding the chance that a success exists, or more pessimistic. The following proposition summarizes the effect that the disclosure of failures may have on the entry and exit decisions, of the two firms.

Proposition 9 (Positive probability of no success): *In a model with a strictly positive probability of no success, in some cases disclosure of failed experiments may induce the seller, and possibly also the buyer, to exit the market, while in other cases such disclosure may extend or*

renew the race. The latter effect implies that even if all the disclosed substances have been previously tested by buyer, disclosure may increase the expected profits of the firms.

Proof. Suppose that $p = 0.6$, $s = 3$, $n_1 = n_2 = 2$. If following full disclosure $n_1 = n_2 = 3$, then the two firms will immediately exit, as they realize that the entire set of substances has already been covered.

Suppose, by contrast, that for these parameters $c = 65$, and $V = 420$. It is now the case that $\hat{P}_{n_1,n_2} = \frac{1}{7}$. The current-period expected profits of each of the firms would be -35. If only one firm abandoned the race, then its rival would face current-period expected profits of -5, and so, without disclosure of failures, both firms should abandon their research endeavor, and the race should be over. Suppose that, in fact, both firms have tested the same substances. If this information was disclosed they could both update their beliefs regarding whether a success in fact exists, and then $\hat{P}_n = \frac{1}{3}$ [13]. Both firms would then opt to return to the race, as each of their current-period expected profits would be 5. If only one of them returned to the race, its current-period expected profits would be 75. ∎

In many technological races, as in the case of the atomic bomb project, undertaken by the superpowers in the mid 20th century, an "encouragement effect" occurs once one of the participants of the race reaches a **success**, since the other participants then learn that success is possible. The importance of this effect is often considerable, as uncertainty is one of the fundamental features of R&D. The most unique feature of Proposition 9 seems to lie in the possibility of an "encouragement effect", that may be caused by disclosure of **failures**. Note that in all previous sections a firm has only considered knowledge of failed tests, made by its rival, to be valuable granted that the firm has not yet conducted these tests on its own. In the current model, however, the very knowledge that the rival has failed directions identical to the firm's own failures, may have a positive, and sometimes even considerable, economic value. This value is created as a result of the fact that such disclosure affects the extent to which the firm is optimistic regarding the existence of a success. An increase in optimism, in that case,

---

[13] The reason why $\hat{P}_{n_1,n_2} < \hat{P}_n$ is that before the disclosure of the failures each of the firms had ascribed a probability of about 44% to the event that all three substances had been tested by at least one the firms, while after the disclosure that probability is known to be zero. Hence, after the disclosure the firms become much more optimistic regarding the chance that a success exists.

may induce both firms not to abandon the race prematurely, before all substances have been tested. In a way, this effect demonstrates the famous saying, ascribed to Winston Churchill, that sometimes "success consists of going from failure to failure without loss of enthusiasm".

3.6 Strategic Manipulation

It seems that in reality one of the main reasons making firms reluctant to even partly cooperate with their rivals, is their apprehension of the ability of their rival to infer additional information from any disclosure, ultimately turning such cooperation unprofitable for the seller. Hence, for instance, it is fairly easy to extend our baseline model, to account for the possibility that different groups of substances have different costs and probabilities of success, by introducing a menu of prices which corresponds to this diversity. Yet, a more challenging problem arises when the two firms are differently informed regarding the success rates; hence, the mere announcement of a menu of prices by one of the firms might convey information to its rival. Moreover, in such a model there is a concern that a firm, anticipating the possible existence of a future market for R&D failures, may distort its own research effort, in order to send its rival down the wrong path. This kind of behavior has been demonstrated by Chatterjee and Evans (2004), in a dynamic patent race, in which the firms observe each other's research choices.

The following example has features which resemble the setting used by Chatterjee and Evans (2004), in order to demonstrate their disinformation effect. Suppose, in our baseline model, that before the race begins, firm 1 with probability $\theta \in (0, 1)$ receives a signal, that accurately informs it whether the successful substance is within the first half of the group of substances or within the last half. Firm 2 gets no signal and is not informed about whether or not firm 1 has observed the signal.

In case firm 1 does not try to misinform firm 2, then receiving the signal is exactly equivalent to accumulating $\frac{s}{2}$ failures. As has been shown in the baseline model, there is a market price for which firm 2 would like to buy this knowledge and firm 1 would like to sell it. Moreover, if this signal is not verifiable, the potential moral-hazard problem can be contractually dealt with, in the same manner that has been demonstrated above.

Suppose instead, that there was a potential equilibrium with disinformation. First note that firm 1 could not sell false information to its rival, since such a possibility could be treated by using a contract which addresses moral-hazard problems. Therefore, the only possibility in our model that disinformation might present itself, is if when being informed, firm 1 decides not to sell its entire information to its rival, and firm 2 does not know for sure whether firm 1 is informed or not. Suppose then, that such disinformation has prevailed in equilibrium, and that it has induced a probability of discovery P_2 to firm 2, and an average cost of cD. If instead, the informed firm decided to sell its information to its rival, the probability of discovery by firm 2 would rise to $\bar{P}_2 > P_2$, and the average cost would drop to $cd < cD$. Firm 2 would then be willing to pay any price lower than $(\bar{P}_2 - P_2)V - c(d - D)$ for that information, while firm 1 would be willing to sell for any price higher than $(\bar{P}_2 - P_2)V + c(d - D)$. This implies that for any equilibrium with disinformation there is a market price for that information which induces the informed firm to deviate from its disinformation strategy, and sell its information to its rival.

What is the reason, then, that our model does not exhibit strategic manipulation, similar to that of Chatterjee and Evans (2004)? The answer lies in two of the main differences between their model and ours. Firstly, in their model the research direction of each firm is observable, and secondly there are no side-payments. This combination implies that distorting one's own research effort might also impede the rival's chances of success. By contrast, in our model the market mechanism allows to fully internalize the effect of the information on the joint value of the two firms, and agree on a price that would induce them to distribute their privately held information. Indeed, one of the key features promoting efficiency in our model is that both firms agree on the value of the information held by the informed firm. This would prove to be a reasonable assumption, for instance, in a case in which this information if verifiable, at the time it is sold. If, however, this information is not verifiable, then once again it raises the possibility that the informed firm might try to deceive its rival. The problem then boils down to the one which is introduced in the moral-hazard section, and the contractual solutions offered in that section could be implemented in the same manner.

Notwithstanding the effects demonstrated in this section, strategic manipulation, in a context of disclosure between rivals, remains, to our view, an interesting and important open question for future research.

3.7 Additional issues

Finally, in this section we would like to discuss some additional issues which are of specific interest in the context of the market for R&D failures.

3.7.1 Information Leakage and Resale

Knowledge, unlike most other commodities, can, generally speaking, be resold at a very low, or even zero, cost. For the market for R&D failures, the possibility of information resale may pose an even greater problem than that of licensing positive interim knowledge. For one thing, since there are no structural IPR for failures, there might be a problem in enforcing exclusive contracts. Then, even if an exclusive agreement is contracted once the failures are disclosed, it becomes particularly hard to prove infringement, on account of the fact that the knowledge of failures is important in the sense it conveys to the buyer what he should **not** do. However, firms are **not** doing a lot of things, hence it is difficult to prove that they avoided doing a particular thing as a result of gaining access to the stock of failures of their competitors.

Perhaps more importantly, due to this inability to prove infringement, a black market for R&D failures might emerge. For example, suppose that a formal market for R&D failures existed. An employee of one of the firms could then be tempted to sell the knowledge of failures, created by his employer, to his competitor. Unless caught in the act of selling this information, it would probably be impossible for the employer to know for sure that such an information leakage has ever occurred.

Unfortunately, unlike the issues we have previously discussed, it seems that this problem cannot be contractually solved. Nevertheless, on the bright side, it also hints that through labor mobility we may see **some** failures dissimilate naturally, even without a market for R&D failures. Yet, it does not rule out the chance that the mere possibility of a black market could result in an equilibrium in which neither the market for R&D failures will exist, nor a black market.

3.7.2   Infinite Set of Likely Research Possibilities

Admittedly, it has puzzled us that, generally speaking, research in Economics has not as yet actually opened the door for publicizing failures. Perhaps one of the plausible explanations for this phenomenon is that there is a large, or even infinite, number of research directions, as well as a strictly positive cost of processing research results of others. In such a framework, even if one accumulates a lot of failures, their contribution in attaining success may be very small, and the strictly positive costs could, thus, turn it unprofitable. Only if one could infer from the results on possibilities which have not yet directly been tested, such knowledge would become valuable. Probably the most prominent examples of such results, which have helped avoiding dead-ends, are impossibility theorems such as Arrow (1951), though we feel extremely uncomfortable labeling this contribution "a failure". The model of d'Aspremont, Bhattacharya, and Gerard-Varet (2000), which was mentioned in the introduction, is, in a way, consistent with this kind of framework. The reason for that is that the Poisson distribution they have assumed implies that with or without that knowledge the firms sample with replacement, hence, such impossibility theorems maintain an infinite sampling set, while increasing the probability of success.

Another possible explanation is the multiplicity of research objectives, implying that without knowledge of the exact objective of a competitor, one may not be able to trade failures. Moreover, potential competitors may be reluctant to disclose their exact research objective, in order to restrict the number of participants in the race.

4   Indirect Dissemination of Knowledge of R&D Failures

Thus far, we have shown, that in many cases there is a market price for which knowledge of R&D failures can be traded, while at the same time there are also many complications which need to be taken into account in order to successfully execute such a transaction. It seems, though, that in reality such transactions are hardly a common practice, and that formal markets for R&D failures are quite scarce. It is the purpose of this section to point out some of the indirect mechanisms which may help disseminate knowledge of R&D failures, and

serve as alternatives for the market for R&D failures.

The first such indirect mechanism relates to policy intervention. R&D involves externalities, such as knowledge spillovers and other kinds, which account for the fact that public returns to R&D are higher than the private ones. Two prominent means of bridging this gap and inducing more innovation are the patent system and R&D subsidies. The former approach rewards an inventor, who reaches a certain level of R&D success, with a temporary monopoly status, in return to making its discovery public. The latter approach rewards the inventor with a monetary compensation at an early stage, before a success has been made. In some countries, unlike patents, subsidies are not bundled with a disclosure requirement[14]. If policymakers believed that there are inefficiencies involved in the way that knowledge of failures diffuses, i.e. the market for R&D failures does not function as it should, they could intervene and directly promote that knowledge diffusion. One way of doing so is by bundling R&D subsidies with a disclosure requirement, perhaps with some lag. Even very coarse data, such as the disclosure of the form, submitted by firms applying for the subsidy, can serve other firms in understanding which attempts have been tried and failed. In that case, the subsidy can also serve as means of compensating its recipient for this partial disclosure to the market, thus bypassing the relatively complicated bilateral contracts which have been discussed in this paper.

Ventures Capitalists (VCs) may help disseminate knowledge of R&D failures in a similar manner. As the VCs often invest in quite a few R&D-intensive firms within a similar field, they can direct their firms in accordance with their overall stock of failures. Like in the case of the subsidy, in this context too, the firms may be compensated for this disclosure, both by the direct investments of the VCs, as well as by the disclosure of the other firms.

Mergers and Acquisitions (M&As) could serve firms as means for acquiring knowledge, including that of R&D failures, in a way which might also align their interests, and reduce incentives for deception. Furthermore, as has been established in the section dealing with entry and exit, our theoretical model also predicts that in many cases when a firm sells its failures to its rival, it may induce it to abandon the race. In this case, the implication of

---

[14]In Israel, for example, the "law for promoting industrial R&D" has a secrecy close, which specifically states that any material handed in, by a firm requesting a governmental R&D subsidy, must not be disclosed to any third party.

that theoretical prediction is that in reality we may see deals, which are, in essence, failures procurement, disguised as M&As.

Finally, another indirect way of acquiring failures from competitors is by labor mobility. Indeed, we believe that a lot of the experience possessed by R&D personnel is its stock of failures. Furthermore, many of the popular managerial books often mention that failures are equal to experience, and that one should learn from his mistakes in order to evolve professionally. However, it remains an unresolved and interesting open empirical question, whether firms recruiting new employees value the stock of failures these employees have accumulated from their past experience. In that context it is also unclear whether an employee would be willing to openly admit to his past failures, as they may send a signal of incompetence to the potential employer. These questions are of a particular interest in the context of technological personnel, since failures are a natural by-product of R&D.

5   Summary and Discussion

Knowledge of past R&D failures constitutes an important part of the stock of knowledge necessary for promoting and advancing R&D. The importance of such knowledge stems from the fact that failures are a natural by-product of R&D activity, hence their widespread prevalence in every R&D-intensive environment. One of the main contributions of this study is our claim that the information structure of R&D is characterized by the fact that information regarding failures is transmitted less efficiently amongst competitors, than that of success. Therefore, it is particularly interesting to explore the questions whether knowledge of R&D failures is likely to be traded between competitors, and why is it that in reality markets for failures are so scarce.

We have shown that in a patent race environment, direct competitors are likely to engage in the trading of R&D failures. However, a closer scrutiny of the specific characterizations of R&D failures as commodity reveals possible complexities involved in contracting such a transaction. As opposed to many other commodities, failures can be relatively easily and inexpensively generated. In addition, as the purpose of the purchase is not to directly use the failures, but rather to learn how to avoid repeating them, it may induce the seller to

try deceiving the buyer, by lowering the quality of the information sold. There may also be all sorts of externalities involved with such a transaction. Finally, lack of structural IPR over R&D failures, as opposed to the protection that patents provide for R&D successes, is another hurdle which needs to be addressed. This paper both analyzes these complexities and offers mechanisms that may help in coping with some of them. It seems that the free market might prove to be an efficient device, which by means of contingent and non-contingent side-payments between the firms, could help in overcoming many of the obstacles and reconcile incentives of competitors.

In more complicated environments there is a sense in which it is more difficult for transactions, in which the seller is the leader of the race, to take place. Nevertheless, this restriction, does not provide explanation as to why the trading of R&D failures is so scarce in reality, as it always seems to be profitable for the laggard to disclose its failures. A clue to this puzzle may come from the fact that, in some cases, such a transaction may encourage the laggard to quit the race. Failure-disclosing transactions may, therefore, be disguised in reality as M&A deals. In other cases, the disclosure of failures could create an "encouragement effect", similar to the effect that the disclosure of successes might have, by inducing entry.

Along with M&A, labor mobility and VCs may also play an important role in disseminating this important knowledge. Yet, owing to shortcomings and limitations found in each of these methods, they can probably serve only as an imperfect substitute to the market for R&D failures.

We have also shown that policy intervention, and in particular the introduction of transparency into existing policy, may be helpful in the dissemination of this important knowledge.

As a final remark, we would also like to point to the fact that our study also has some indirect implications regarding the publication academic research. For instance, we believe that editors of the academic journals should encourage authors to courageously publish their failing attempts too, if they suspect other researchers can make use of it. In fact, we were happy to learn of the existence of the Journal of Negative Results in Biomedicine[15], which seems to implement the approach that encourages the distribution of such unique knowledge.

---

[15]See www.jnrbm.com .

We feel that the exploration of the knowledge of negative results, such as failures, remains a fertile research ground. In future research we intend to further explore the possibility of sharing knowledge of R&D failures in non-trivial economic contexts, other than within patent races. We are quite certain that many of the features of contracting such knowledge, as presented and demonstrated in this paper, will carry over into other contexts, while some new features will surely present themselves. The study of the alternative means of disseminating R&D failures, as in the policy intervention we have mentioned, and the labor mobility issues, seems to us as additional fascinating topics for future research.

6 Appendix

6.1 The Decision Regarding the Testing Sequence

Suppose that a market for R&D failures does not exist. In the baseline model we have assumed that each of the firms tests its remaining substances randomly; the question, however, is whether this is an equilibrium strategy. If, for instance, one of the firms chose the sequence $\{1, 2, ..., s\}$, while the other chose the sequence $\{s, s - 1, ..., 1\}$, then there would not be any failing substance that would be tested by both firms, and the R&D process would be efficient. For the simplicity of the exposition we assume in this appendix that both firms start with no stock of failures. This section first shows that when both firms choose a uniformly random sequence, it is an equilibrium. In addition, this section proves that there is no equilibrium in which either firm chooses a sequence as a pure action. Finally, it shows that though it is possible that there are equilibria other than the uniformly random one, calculating the best response of each firm is an NP-hard problem, with no closed-form solution, and so in reality any other equilibrium is probably esoteric.

Although the patent race is a dynamic game by nature, since the firms do not learn throughout the game about the failed tests of their rivals, it can be analyzed as a static game. A pure strategy can be written as the sequence of tests that the firm plans. It turns out, however, that many of the mixed strategies are equivalent for both firms, in terms of payoffs, and that the only thing that matters for the payoffs is the probability that each firm ascribes to testing

each substance in each of the stages. We can therefore, write any mixed strategy of firm $i \in \{1, 2\}$ as a Doubly-stochastic matrix, D^i, of size $s \times s$, in which an element $P_i^{jk} \in [0, 1]$ denotes the probability that firm i will test substance j at stage k. A pure strategy is a matrix such that all the elements are $P_i^{jk} \in \{0, 1\}$. The expected profits of the firms, as a function of the strategies chosen by both firms are:

$$\pi_1 = \frac{1}{s} \sum_{i=1}^{s} \left\{ V \left[\sum_{l=1}^{s} \left(P_1^{li} \sum_{k=i}^{s} P_2^{lk} \right) \right] - V \left[\sum_{l=1}^{s} \frac{P_1^{li} P_2^{li}}{2} \right] - c \left[\sum_{l=1}^{s} \left(\sum_{k=i}^{s} P_1^{lk} \right) \left(\sum_{k=i}^{s} P_2^{lk} \right) \right] \right\}, \quad (36)$$

$$\pi_2 = \frac{1}{s} \sum_{i=1}^{s} \left\{ V \left[\sum_{l=1}^{s} \left(P_2^{li} \sum_{k=i}^{s} P_1^{lk} \right) \right] - V \left[\sum_{l=1}^{s} \frac{P_2^{li} P_1^{li}}{2} \right] - c \left[\sum_{l=1}^{s} \left(\sum_{k=i}^{s} P_2^{lk} \right) \left(\sum_{k=i}^{s} P_1^{lk} \right) \right] \right\}. \quad (37)$$

We assume now that for any feasible strategies by the two firms, the expected payoffs are positive.

Proposition A1: *Both firms playing a uniformly random strategy constitute a Nash equilibrium.*

Proof. Recall that for any $i \in \{1, 2\}$, $\sum_{j=1}^{s} P_i^{jk} = 1 \quad \forall k$ and that $\sum_{k=1}^{s} P_i^{jk} = 1 \quad \forall j$. Suppose that one of the firms chooses a uniformly random strategy, i.e. for that firm $P_i^{jk} = \frac{1}{s} \quad \forall j, k$. It then follows that the profit functions are the same as in equations (4) and (5), regardless of the strategy chosen by the other firm. This implies, by definition, that having both firms choose a uniformly random strategy is a Nash equilibrium in the game. ∎

Let us assume from now on that $s > 3$. Suppose that firm 1 chooses a pure strategy, such as the sequence $\{1, 2, 3, ..., s\}$, which corresponds to $P_i^{jj} = 1 \quad \forall j$. Then we might expect that firm 2 would choose a preemption sequence, namely $\{2, 3, 4, ..., s, 1\}$. This kind of a strategy allows firm 2 to win the patent race with probability $\frac{s-1}{s}$, since it preempts firm 1 in all the substances but the first.

Proposition A2: *In equilibrium no firm would choose a pure sequence as a strategy.*

Proof. Suppose that one of the firms chose a pure strategy. Without loss of generality, let it be firm 1. If firm 2 then used the "preemption" strategy, their payoffs would be:

$$\pi_1^P = \frac{1}{s} V - \frac{s^2 - s + 2}{2s} c, \quad (38)$$

40

$$\pi_2^P = \frac{s-1}{s}V - \frac{s^2-s+2}{2s}c. \tag{39}$$

While if any of the firm deviated to the uniformly random strategy their payoffs would be:

$$\pi_1^{UR} = \frac{1}{2}V - \frac{2s^2+3s+1}{6s}c, \tag{40}$$

$$\pi_2^{UR} = \frac{1}{2}V - \frac{2s^2+3s+1}{6s}c. \tag{41}$$

Any best response by firm 2 to the pure strategy played by firm 1 would have to guarantee firm 2 a payment of at least π_2^P. However, firm 1 could always guarantee for itself a payment of π_1^{UR}, since this payment is attainable by deviating to the uniformly random strategy, regardless of the strategy chosen by firm 2. Therefore if in equilibrium firm 1 chose a pure strategy, then the joint payoffs of the firm would have to be at least $\pi_1^{UR} + \pi_2^P$. However, the largest joint payoff by the two firms is achievable if they cooperate and avoid any duplication in their research process. In that case the discovery will be made on average after $\frac{s}{4}$ periods. Therefore, the joint profits of the firm will have to satisfy this upper bound:

$$\pi_1^{UR} + \pi_2^P \leq V - 2c\frac{s}{4}, \tag{42}$$

which is equivalent to the following condition:

$$\frac{V}{c} \leq \frac{2s^2+7}{3s-6}. \tag{43}$$

However, since we have assumed that for any feasible strategies by the two firms, the expected payoffs are positive, we require that $\pi_1^P \geq 0$, which is equivalent to:

$$\frac{V}{c} \geq \frac{s^2-s+2}{2}. \tag{44}$$

Combined with the previous condition this implies that:

$$s(s-3)(3s-4) \leq 26. \tag{45}$$

When $s = 4$ the left hand side equals 32, and since for any $s \geq 4$ it is obviously a monotonically increasing function of s, this condition never holds. This implies that a pure strategy by any of the firms can never be an equilibrium behavior. ∎

Finally, we would like to show that calculating the best response of any of the firms is an NP-hard problem, with no closed form analytical solution. For that purpose, let us

concentrate on the best response of firm 1. Its payoff function can be rewritten as:

$$\pi_1 = \frac{1}{s} \sum_{i=1}^{s} \sum_{j=1}^{s} P_1^{ij} \left[V \left(\sum_{k=j}^{s} P_2^{ik} \right) - \frac{1}{2} V \left(P_2^{ij} \right) - c \sum_{l=1}^{j} \left(\sum_{k=l}^{s} P_2^{ik} \right) \right]. \tag{46}$$

Hence, given the strategy chosen by firm 2, the objective function is in fact linear. We can, therefore, write the problem of finding the best response as the following Linear Programming (LP) problem:

$$\max_{P_1^{ij}} \frac{1}{s} \sum_{i=1}^{s} \sum_{j=1}^{s} P_1^{ij} \left[V \left(\sum_{k=j}^{s} P_2^{ik} \right) - \frac{1}{2} V \left(P_2^{ij} \right) - c \sum_{l=1}^{j} \left(\sum_{k=l}^{s} P_2^{ik} \right) \right], \tag{47}$$

s.t.

$$\sum_{j=1}^{s} P_1^{ij} = 1 \qquad \forall i, \tag{48}$$

$$\sum_{i=1}^{s} P_1^{ij} = 1 \qquad \forall j \tag{49}$$

and

$$P_1^{ij} \geq 0 \qquad \forall i, j. \tag{50}$$

This is a variation of the well known "Quadratic Assignment Problem", which is known to be NP-hard.

6.2   Proof of Proposition 1 - The Feasibility of a Market for R&D Failures

The easier case is when $\delta = 1$. In this case the payoff function are:

$$\pi_1 \left(n_1, n_2 \right) = V \frac{s+n_1-2n_2}{2(s-n_2)} - c \left[1 + \frac{(s-n_1-1)(2s+n_1-3n_2-1)}{6(s-n_2)} \right], \tag{51}$$

$$\pi_2 \left(n_1, n_2 \right) = V \frac{s-n_1}{2(s-n_2)} - c \left[1 + \frac{(s-n_1-1)(2s+n_1-3n_2-1)}{6(s-n_2)} \right]. \tag{52}$$

Consider first transactions involving disclosure by firm 2. When $k = 0$ then, by definition, the disclosure value in equation (7) equals zero. The marginal contribution of a disclosed failure is:

$$\frac{\partial \Pi \left(n_1, n_2 \right)}{\partial n_1} = \frac{c \left[2 \left(n_1 - n_2 \right) + \left(s - n_2 \right) \right]}{3 \left(s - n_2 \right)} > 0. \tag{53}$$

Since $s > n_1 \geq n_2$ and $c > 0$, the disclosure value is always positive, and monotonically increasing in the number of failures disclosed. The length of the interval of prices equals the net disclosure value, that is shown in equation (7), which is strictly positive since $n_1 \geq n_2$, $k > 0$ and $s > n_2$. Hence, we have established that any disclosure by firm 2 has a positive net value.

Consider now transactions involving disclosure by firm 1. For the cases in which following the disclosure condition (1) still holds, the disclosure value is represented in equation (8). When $k = 0$ then, by definition, the disclosure value in equation (8) equals zero. As long as firm 1 does not lose its lead the marginal contribution of a disclosed failure is:

$$\frac{\partial \Pi\left(n_1, n_2\right)}{\partial n_2} = \frac{c\left[\left(s - n_1\right)^2 - 1\right]}{3\left(s - n_2\right)^2}, \tag{54}$$

which is positive for every $s > n_1 + 1$. This proves that the added-value of disclosure is monotonically increasing in the number of failures disclosed, as long as firm 1 has more accumulated failures than its rival. Thus, as the disclosure value is always positive, the length of the interval of prices, which is given in equation (8), must also be strictly positive.

Since disclosure by firm 2 has always a positive net value, and is monotonically increasing, then it follows that even in those cases in which firm 1 loses its lead, subsequent to the disclosure, the disclosure value is still monotonically increasing in the number of failures disclosed.

Now for the more complex case in which $0 > \delta > 0$. The recursive form of the payoff function, which appears in equations (2) and (3), can be transformed into the following normal form:

$$\pi_1\left(n_1, n_2\right) = V \frac{\left(1 - \delta^{s-n_1}\right)(1+\delta) - (1-\delta)\delta^{s-n_1}\left(n_1 - n_2\right)}{2(1-\delta)^2(s-n_1)(s-n_2)} + V \frac{1}{(1-\delta)(s-n_1)} - \tag{55}$$

$$-c \frac{(1+\delta)\left(\delta - \delta^{s-n_1+1}\right) + (1-\delta)\delta^{s-n_1+1}\left(n_1 - n_2\right)}{(1-\delta)^3(s-n_1)(s-n_2)} + c \frac{\delta(2s-n_1-n_2)}{(1-\delta)^2(s-n_1)(s-n_2)} - c \frac{1}{1-\delta},$$

and

$$\pi_2\left(n_1, n_2\right) = \pi_1\left(n_1, n_2\right) - V \frac{\left(n_1 - n_2\right)\left(1 - \delta^{s-n_1}\right)}{(1-\delta)(s-n_1)(s-n_2)}. \tag{56}$$

Note that under the assumption of the proposition, $s - n_1 \geq 2$ and $0 > \delta > 0$. We start by analyzing the net disclosure value of failures by firm 1, which keeps firm 1 the leader of

the race (hence after the disclosure, $n_1 \geq n_2$). The disclosure value is defined as in Equation (8). The first derivative with respect to n_2 is,

$$\frac{\partial \Pi\left(n_1, n_2\right)}{\partial n_2} = \left[(1-\delta) V + 2\delta c\right] \frac{\left(1+\delta^{s-n_1}\right)(s-n_1)(1-\delta) - \left(1-\delta^{s-n_1}\right)(1+\delta)}{(s-n_2)^2(s-n_1)(1-\delta)^3}. \tag{57}$$

Under the assumptions of the model it is easy to see that the derivative is positive if and only if,

$$\Gamma \equiv \left(1 + \delta^{s-n_1}\right)\left(s - n_1\right)\left(1 - \delta\right) - \left(1 - \delta^{s-n_1}\right)\left(1 + \delta\right) \geq 0. \tag{58}$$

First note that for $\delta \to 1$, Γ approaches 0, and when $\delta \to 0$, Γ approaches $s - n_1 + 1 > 0$. Furthermore, Γ is a continuous and differentiable function of δ. We next show that whenever $\Gamma = 0$, then $\forall \delta \in (0,1)$, the derivative of Γ with respect to δ is negative, and so Γ is in fact positive $\forall \delta \in (0,1)$[16]. The partial derivative of Γ with respect to δ is:

$$\frac{\partial \Gamma}{\partial \delta} = (s - n_1 + 1) \left[(s - n_1)\delta^{s-n_1-1}(1-\delta) - \left(1 - \delta^{s-n_1}\right)\right]. \tag{59}$$

At $\Gamma = 0$, we have:

$$\left.\frac{\partial \Gamma}{\partial \delta}\right|_{\Gamma=0} = -(s - n_1 + 1)(s - n_1)\frac{1-\delta}{1+\delta}\left(1 - \delta^{s-n_1}\right) < 0, \tag{60}$$

which proves that Γ is always positive, and so the disclosure value is an increasing function of the number of failures disclosed by firm 1.

Now let us consider the case in which firm 2 discloses failures to firm 1. The disclosure value is defined as in equation 7. The first derivative with respect to n_1 is

$$\frac{\partial \Pi\left(n_1, n_2\right)}{\partial n_1} = \left[(1-\delta) V + 2\delta c\right] \left[\frac{\left[(s-n_2+1)(1-\delta)-2\right]\left(1-\delta^{s-n_1}\right)}{(s-n_1)^2(s-n_2)(1-\delta)^3} - \right. \tag{61}$$
$$\left. - \frac{\log[\delta]\delta^{s-n_1}(s-n_1)\left[1+\delta-(n_1-n_2)(1-\delta)\right]}{(s-n_1)^2(s-n_2)(1-\delta)^3}\right].$$

Under the assumptions of the model it is easy to see that the derivative is positive if and only if

$$\zeta \equiv \left[(s - n_2 + 1)(1 - \delta) - 2\right]\left(1 - \delta^{s-n_1}\right) - \tag{62}$$
$$- \log[\delta]\delta^{s-n_1}(s - n_1)\left[1 + \delta - (n_1 - n_2)(1 - \delta)\right] \geq 0.$$

---

[16]This approach is not too different from that used for proving Lemma 2 in: Milgrom, Paul R. and Robert J. Weber, 1982, "A Theory of Auctions and Competitive Bidding", *Econometrica* 50, pp. 1089-1122.

This proof, unfortunately, requires many intermediate steps, however what we would eventually like to show is that $\zeta \geq 0$, when $n_2 \to n_1$, and then show that ζ is a decreasing function of n_2. Let us define $\psi \equiv \lim_{n_2 \to n_1} \zeta$. We therefore have:

$$\psi = \left[(s - n_1 + 1)(1 - \delta) - 2 \right] \left(1 - \delta^{s-n_1} \right) - \log \left[\delta \right] \delta^{s-n_1} (s - n_1)(1 + \delta). \qquad (63)$$

Note that or $\delta \to 1$, ψ approaches 0, and when $\delta \to 0$, ψ approaches $s - n_1 + 1 > 0$. Furthermore, ψ is a continuous and differentiable function of δ. We next show that whenever $\psi = 0$, then for any δ, the derivative of ψ with respect to δ is negative, and so ψ is in fact positive $\forall \delta \in (0, 1)$. The partial derivative of ψ with respect to δ is:

$$\frac{\partial \psi}{\partial \delta} = -(s - n_1 + 1) - \delta^{s-n_1-1} \{ (s - n_1)^2 (1 - \delta) - (s - n_1 + 1)\delta + \qquad (64)$$
$$+ (s - n_1) \left[(s - n_1)(1 + \delta) + \delta \right] \log \left[\delta \right] \}.$$

If we plug in the condition that $\psi = 0$ (namely substituting the term $(s - n_1) \log \left[\delta \right]$), we will get:

$$\frac{\partial \psi}{\partial \delta} \bigg|_{\psi=0} = \frac{(s - n_1) \left[\left(1 + \delta^2 \right) \left(1 - \delta^{s-n_1} \right) - (s - n_1) \left(1 - \delta^2 \right) \right]}{\delta (1 + \delta)}. \qquad (65)$$

This term is always negative, since it equals 0 when $\delta \to 1$, and its derivative with respect to δ is always positive, as it equals:

$$\frac{\partial^2 \psi}{\partial \delta^2} \bigg|_{\psi=0} = \frac{s - n_1}{\delta^2 (1 + \delta)^2} \{ \left[s - n_1 - 1 + \delta(s - n_1 - 2) + \delta^2 \right] \left(1 - \delta^{s-n_1} \right) + \qquad (66)$$
$$+ (s - n_1)(1 + \delta) \left(\delta - \delta^{s-n_1+2} \right) \},$$

which proves that ψ is always positive.

Having established that $\zeta \geq 0$, when $n_2 \to n_1$, what is left to be proven is that ζ is a decreasing function of n_2. The derivative of ζ with respect to n_2 is:

$$\lambda \equiv \frac{\partial \zeta}{\partial n_2} = (1 - \delta) \left[\delta^{s-n_1} - \log \left[\delta \right] \delta^{s-n_1} (s - n_1) - 1 \right]. \qquad (67)$$

First note that for $\delta = 1$, $\lambda = 0$, and when $\delta = 0$, $\lambda = -1$. λ is also a continuous and differentiable function of δ. Therefore, what we would like to prove is that whenever δ is strictly between 0 and 1, then if $\lambda = 0$, its derivative with respect to δ is either always positive or always negative, and so λ itself would never be strictly positive within that interval. The derivative of λ with respect to δ is:

$$\frac{\partial \lambda}{\partial \delta} = 1 - \delta^{s-n_1} - \log \left[\delta \right] (s - n_1) \delta^{s-n_1-1} \left[(s - n_1)(1 - \delta) - \delta \right]. \qquad (68)$$

For $\delta \in (0,1)$, $\lambda = 0$ if and only if the following holds:

$$1 - \delta^{s-n_1} = -\log[\delta]\,\delta^{s-n_1}\,(s-n_1), \tag{69}$$

and so if we plug that into the derivative of λ with respect to δ we will get:

$$\left.\frac{\partial\lambda}{\partial\delta}\right|_{\lambda=0} = -\log[\delta]\,(s-n_1)^2\,\delta^{s-n_1-1}\,(1-\delta) > 0. \tag{70}$$

Hence we have established that under our assumptions that δ is strictly between 0 and 1 and that $s - n_1 > 1$, we have $\lambda \leq 0$ and so ζ is a decreasing function of n_2. Since we have proved that $\zeta \geq 0$ for $n_2 \to n_1$, we get that ζ is always positive under our assumptions. ■

6.3 Proof of Proposition 2

The contract includes an upfront payment by the buyer, p, and a fine, F, which will be paid by the seller should one of the sold test results turns out to be a success. The following two constraints (IR_j and IR_i, respectively) ensure that both firms will agree to participate in the transaction, given that the seller is truthful:

$$\pi_j^* - \pi_j + p \geq 0, \tag{71}$$

$$\pi_i^* - \pi_i - p \geq 0. \tag{72}$$

Since we have assumed that $\pi_i^* - \pi_i > \pi_j - \pi_j^*$, then $\exists\, p$, which satisfies the two IR conditions, such that:

$$\pi_i^* - \pi_i \geq p \geq \pi_j - \pi_j^*. \tag{73}$$

The following constraint (IC_j) ensures that firm j, the seller, will not try to deceive the buyer:

$$\pi_j^* - \pi_j + p \geq \pi_j^{**} - \pi_j + p - P_S P_N F, \tag{74}$$

which is equivalent to:

$$F \geq \frac{\pi_j^{**} - \pi_j^*}{P_S P_N}. \tag{75}$$

Hence, if the fine is large enough, the seller will be truthful. Since this is true $\forall \pi_j^{**}, P_S \in (0,1)$, then $\exists\, F$ which discourages any deviation from the truthful strategy.

Since this contract exactly replicates the transactions of the counterfactual world, then in equilibrium both the buyer and the seller will follow the research strategy as they would follow in a world with no deception. Therefore, the contract implements the first-best outcome. ∎

6.4 Proof of Proposition 3

The contract includes an upfront payment by the buyer, p, and a fine, F, which will be paid by the seller should one of the sold test results turns out to be a success. The following two constraints (IR_j and IR_i, respectively) ensure that both firms will be willing to participate in the transaction, given that the seller is truthful:

$$(1 - P_A)\left(\pi_j^* - \pi_j\right) + p \geq 0, \tag{76}$$

$$(1 - P_A)\left(\pi_i^* - \pi_i\right) - p \geq 0. \tag{77}$$

Since we have assumed that $\pi_i^* - \pi_i > \pi_j - \pi_j^*$ and $P_A \in (0, 1)$, then $\exists\, p$, which satisfies the two IR conditions, such that:

$$\pi_i^* - \pi_i \geq \frac{p}{1 - P_A} \geq \pi_j - \pi_j^*. \tag{78}$$

The following constraint (IC_j) ensures that firm j, the seller, would not wish to deceive the buyer:

$$(1 - P_A)\left(\pi_j^* - \pi_j\right) + p \geq (1 - P_A)\left(\pi_j^{**} - \pi_j\right) + p - P_A P_S P_U F, \tag{79}$$

which is equivalent to:

$$F \geq \frac{(1 - P_A)\left(\pi_j^{**} - \pi_j^*\right)}{P_S P_A P_U}. \tag{80}$$

Hence, for any $P_A \in (0, 1)$, there exist a fine large enough, to guarantee that the seller will be truthful. Since this is true $\forall \pi_j^{**}, P_S \in (0, 1), P_A \in (0, 1)$, then $\exists\, F$ which will discourage any deviation from the truthful strategy. Taken together with the fact that there is always a price, p, which can satisfy the IR conditions, it implies that $\forall P_A \in (0, 1)$, the transactions in the market for R&D failures are left undistorted, compared with the counterfactual world.

The net disclosure value in that case is $(1 - P_A)\left(\pi_j^* - \pi_j + \pi_i^* - \pi_i\right)$, which is infinitely close to the first-best, only if P_A is infinitely small, hence the corresponding fine, F, must be infinitely large. ∎

6.5 Proof of Proposition 4

Let K denote the number of periods, which each punishment phase consists of. For the moment we take K to be exogenous. No deception implies that $\forall a$:

$$\sum_{l=0}^{\infty} \hat{\delta}^l \pi_i^* \geq \pi_i^a + P_F^a \left[\sum_{l=1}^{K} \hat{\delta}^l \pi_i + \sum_{l=K+1}^{\infty} \hat{\delta}^l \pi_i^* \right] + (1 - P_F^a) \left[\sum_{l=1}^{K} \hat{\delta}^l \pi_i^* + \sum_{l=K+1}^{\infty} \hat{\delta}^l \pi_i^* \right]. \quad (81)$$

On the left hand side of the inequality is the expected payoffs of firm i, given that the two firms never try to deceive each other and trade R&D failures optimally. On the right hand side, the first expression is its maximal expected payoffs when it unilaterally deviates, and cheats its rival in the current race. Next, with probability P_F^a, the fraud is detected, the firm is then faced with K periods of punishment, and no trade in R&D failures, and then returning to the cooperative phase forever. With probability $1 - P_F^a$ the fraud is left undetected and the cooperation phase continues. This inequality is equivalent to:

$$P_F^a \sum_{l=1}^{K} \hat{\delta}^l \left(\pi_i^* - \pi_i \right) \geq \pi_i^a - \pi_i^*, \forall a. \quad (82)$$

which simply states that the maximal expected profits from deception in the current race should not exceed the discounted loss of the punishment phase.

Since the probability of an honest mistake is zero, then the punishment phase could be infinite, because in equilibrium the punishment phase never occurs[17]. When indeed $k \to \infty$, the condition for cooperation becomes:

$$P_F^a \frac{\hat{\delta}}{1 - \hat{\delta}} \left(\pi_i^* - \pi_i \right) \geq \pi_i^a - \pi_i^*, \forall a, \quad (83)$$

which is equivalent to condition (18). ∎

6.6 Proof of Proposition 5

A transaction which marginally increases the lead of firm 1, by increasing its stock of failures, will affect γ_i in the following way:

$$\frac{\partial \gamma_1}{\partial n_1} = -\frac{k_2 \left(3V - 2cs + 2cn_1 \right)}{6k_1^2 \left(s - n_2 \right)}, \quad (84)$$

---
[17] As in Rotemberg and Saloner (1986).

$$\frac{\partial \gamma_2}{\partial n_1} = -\frac{4ck_2\left(s-n_1\right)+3k_1\left(V-cs+cn_2\right)}{6k_1^2\left(s-n_2\right)}. \tag{85}$$

Since it is assumed that $V > 2cs$, the effect is negative for both firms; hence, such a transaction decreases the myopic willingness of both firms to invest in research capacity.

A transaction which marginally advances firm 2, yet keeps it the laggard, will affect γ_i in the following way:

$$\frac{\partial \gamma_1}{\partial n_2} = \frac{k_2\left(s-n_1\right)\left(3V-2cs+2cn_1\right)+ck_2\left[\left(s-n_1\right)^2-k_1^2\right]}{6k_1^2\left(s-n_2\right)^2}, \tag{86}$$

$$\frac{\partial \gamma_2}{\partial n_2} = \frac{3Vk_1\left(s-n_1\right)+2ck_2\left[\left(s-n_1\right)^2-k_1^2\right]}{6k_1^2\left(s-n_2\right)^2}. \tag{87}$$

In this case if firm 1 is not sure to make the discovery in the next period, meaning that $s - n_1 \geq k_1$, then such a transaction will increase the myopic willingness of both firms to invest in research capacity. If, by contrast, $s - n_1 < k_1$, then the direction of this effect is inconclusive. ∎

6.7 Proof of Lemma 6

Suppose that firm 2 is known to quit the race if it accumulates $\bar{n}_2 \in [n_2+1, n_2+s-n_1-1]$ failures. The expected profits of the two firms will then be:

$$\pi_1\left(n_1,n_2,\bar{n}_2\right)=\frac{\left[\sum\limits_{i=0}^{\bar{n}_2-n_2-1}V\left(s-n_2-i-\frac{1}{2}\right)-c(s-n_1-i)(s-n_2-i)\right]+(s-\bar{n}_2)\left[\sum\limits_{l=\bar{n}_2-n_2}^{s-n_1-1}V-c(s-n_1-i)\right]}{(s-n_1)(s-n_2)}, \tag{88}$$

$$\pi_2\left(n_1,n_2,\bar{n}_2\right)=\frac{\left[\sum\limits_{i=0}^{\bar{n}_2-n_2-1}V\left(s-n_1-i-\frac{1}{2}\right)-c(s-n_1-i)(s-n_2-i)\right]}{(s-n_1)(s-n_2)}. \tag{89}$$

The joint profits of the two firms, denoted by Π, if firm 2 is to abandon the race once it tests $\bar{n}_2$ substances will be:

$$\Pi\left(n_1,n_2,\bar{n}_2\right)=V-c\Big[\tfrac{2(\bar{n}_2-n_2)(\bar{n}_2-n_2+1)(3n_1+2\bar{n}_2-5n_2+1)}{6(s-n_1)(s-n_2)}+ \tag{90}$$
$$+\tfrac{3(s-\bar{n}_2-n_1+n_2)[(3s-4n_2+\bar{n}_2+4)(\bar{n}_2-n_2)+(s-\bar{n}_2)(s-n_1+1)]}{6(s-n_1)(s-n_2)}\Big].$$

Although Equations (88) and (89) are not well defined for $\bar{n}_2 = n_2$ and for $\bar{n}_2 = n_2 + s - n_1$, Equation (90), in fact, holds for $\forall \bar{n}_2 \in [n_2, n_2 + s - n_1]$. If firm 2 is to stay

in the race until all substances are tested, the joint profits will be as in the baseline model, meaning that:

$$\Pi\left(n_1, n_2, \bar{n}_2\right)\big|_{\bar{n}_2=n_2+s-n_1} = \Pi\left(n_1, n_2\right) = V - c\frac{(s-n_1+1)(2s+n_1-3n_2+1)}{3(s-n_2)}, \qquad (91)$$

and the joint profits when firm 2 quits immediately will be, as can be expected:

$$\Pi\left(n_1, n_2, \bar{n}_2\right)\big|_{\bar{n}_2=n_2} = \pi_1\left(n_1\right) = V - c\frac{s-n_1+1}{2}. \qquad (92)$$

The change in the joint profit, as a result of a decrease in $\bar{n}_2$ is:

$$\Pi\left(n_1, n_2, \bar{n}_2 - 1\right) - \Pi\left(n_1, n_2, \bar{n}_2\right) = c\frac{(s-\bar{n}_2-n_1+n_2+1)(s-\bar{n}_2+n_1-n_2+2)}{2(s-n_1)(s-n_2)}. \qquad (93)$$

Since $\bar{n}_2 \in [n_2, n_2 + s - n_1]$ and $s > n_1 \geq n_2$, the change in the joint profits is positive. ∎

6.8   Proof of Proposition 7

Let j denote the firm which discloses its failures. As in the baseline model, let n_i and n_j denote the number of failures each firm has accumulated, and k denote the number of substances disclosed by firm j, which have not yet been tested by firm i. Let $\bar{n}_j$ denote the number of failures which firm j had planned on accumulating before abandoning the race, prior to the disclosure of firm j, and $\bar{n}'_j$ denote that number following the disclosure.

First note that $\bar{n}'_j \leq \bar{n}_j$, since, *Ceteris Paribus*, the per-period expected profits of firm j is a decreasing function of the number of failures known to firm i. Moreover, firm i would not abandon the race following the transaction, as its per-period profits have increased following the transaction:

$$\partial\left[V\frac{s-n_j-0.5}{(s-n_i)(s-n_j)} - c\right]/\partial n_i = V\frac{s-n_j-1}{(s-n_i)^2(s-n_j)} > 0, \qquad (94)$$

$$\partial\left[V\frac{s-n_i-0.5}{(s-n_i)(s-n_j)} - c\right]/\partial n_i = -V\frac{1}{2(s-n_i)^2(s-n_j)} < 0. \qquad (95)$$

Let us denote $\hat{n}_j \equiv \min(\bar{n}_j, n_j + s - n_i - k)$. As a first step, assume that firm j decided not to change the number of substances it intended to check before abandoning the race. This implies that $\bar{n}'_j = \hat{n}_j$, since if $\bar{n}_j \geq n_j + s - n_i - k$, then firm i will finish

checking all the substances before firm j will quit. Now if $n_j \geq n_i$ then firm j was the leader of the race before the disclosure, and so it did not intend to abandon it before checking all the substances. In this case Proposition 1 holds, since there is no early exiting either before or after the transaction. If, alternatively, $n_j < n_i$ and $\bar{n}'_j = \hat{n}_j$, then two cases should be considered. In the first case $\hat{n}_j = \bar{n}_j$, and so the net disclosure value will be[18]:

$$
\begin{aligned}
&\Pi\left(n_i + k, n_j, \hat{n}_j\right) - \Pi\left(n_i, n_j, \bar{n}_j\right) = \\
&= -c \sum_{l=0}^{\bar{n}_j - n_j - 1} \left[\frac{2(s-n_i-l-k)(s-n_j-l)}{(s-n_i-k)(s-n_j)} - \frac{2(s-n_i-l)(s-n_j-l)}{(s-n_i)(s-n_j)} \right] - \\
&- c \sum_{l=\bar{n}_j - n_j}^{s-n_i-k-1} \left[\frac{(s-\bar{n}_j)(s-n_i-k-l)}{(s-n_i-k)(s-n_j)} - \frac{(s-\bar{n}_j)(s-n_i-l)}{(s-n_i)(s-n_j)} \right] - c \sum_{l=s-n_i-k}^{s-n_i-1} \left[-\frac{(s-\bar{n}_j)(s-n_i-l)}{(s-n_i)(s-n_j)} \right] = \\
&= ck \frac{3(s-\bar{n}_j)(s-n_i)(s-n_i-k) + (\bar{n}_j-n_j)(\bar{n}_j-n_j-1)(3s-\bar{n}_j-2n_j+2)}{6(s-n_i)(s-n_i-k)(s-n_j)}.
\end{aligned}
\tag{96}
$$

Since $n_i + k < s > \bar{n}_j > n_j$, then $\forall k \geq 1$, this term is strictly positive.

Otherwise, in the second case $\hat{n}_j \neq \bar{n}_j$, and so the net disclosure value will be:

$$
\begin{aligned}
&\Pi\left(n_i + k, n_j, \hat{n}_j\right) - \Pi\left(n_i, n_j, \bar{n}_j\right) = \\
&= -c \sum_{l=0}^{s-n_i-k-1} \left[\frac{2(s-n_i-l-k)(s-n_j-l)}{(s-n_i-k)(s-n_j)} - \frac{2(s-n_i-l)(s-n_j-l)}{(s-n_i)(s-n_j)} \right] - \\
&- c \sum_{l=s-n_i-k}^{\bar{n}_j - n_j - 1} \left[-\frac{2(s-n_i-l)(s-n_j-l)}{(s-n_i)(s-n_j)} \right] - c \sum_{l=\bar{n}_j - n_j}^{s-n_i-1} \left[-\frac{(s-\bar{n}_j)(s-n_i-l)}{(s-n_i)(s-n_j)} \right] = \\
&= ck \frac{(\bar{n}_j-n_j)(s-\bar{n}_j+2n_i-2n_j+2)+(s-n_i)(s+\bar{n}_j+2n_i-4n_j+2k-2)}{6(s-n_i)(s-n_j)} + \\
&+ c \frac{[(\bar{n}_j-n_j)-(s-n_i-k)-1][(s-n_i)-(\bar{n}_j-n_j)](s-\bar{n}_j+2n_i-2n_j+2)}{6(s-n_i)(s-n_j)}.
\end{aligned}
\tag{97}
$$

The fact that $\hat{n}_j \neq \bar{n}_j$, implies that $\bar{n}_j - n_2 - 1 \geq s - n_1 - k$, and so because $n_i + k < s > \bar{n}_j > n_j$, then $\forall k \geq 1$, this term too is strictly positive. All in all, then, it follows that this disclosure value which does not take into account earlier quitting of firm j, is always strictly positive.

If firm j now decides to abandon the race at an earlier stage, such that $\bar{n}'_j < \hat{n}_j$, then the value of that change $\forall \hat{n}_j$ will be:

$$
\begin{aligned}
&\Pi\left(n_i + k, n_j, \bar{n}'_j\right) - \Pi\left(n_i + k, n_j, \hat{n}_j\right) = \\
&= c \frac{\left(\hat{n}_j - \bar{n}'_j\right)\left[3\left(s - \hat{n}_j - n_i - k + n_j\right)\left(s - \bar{n}'_j + n_i + k - n_j + 1\right) + 3\left(s - \bar{n}'_j\right) + \left(3n_i + 3k - 3n_j + \hat{n}_j - \bar{n}'_j\right)\left(\hat{n}_j - \bar{n}'_j\right) + 2\right]}{6(s-n_i-k)(s-n_j)}.
\end{aligned}
\tag{98}
$$

---

[18]When $\bar{n}_j = n_2 + s - n_1 - k$, though the summation is not well-defined, the result remains the same.

Lemma 6 implies that this change, too, has a positive impact on the net disclosure value[19].

∎

6.9 Proof of Proposition 8

The maximal number of substances which firm 2 will further accumulate cannot exceed the number of substances remaining for firm 1. It is therefore true that $0 \leq \bar{n}_2 - n_2 \leq s - n_1$. Since firm 1 is assumed to have already acquired all the failures which are known to firm 2, then full disclosure implies that the number of failures known to firm 2 will increase from n_2 to n_1 ("direct disclosure effect"), and that both firms will stay in the race until a discovery is made ("exit-discouragement effect").

The change in expected profits as a result of the exit-discouragement effect is:

$$\Pi\left(n_1, n_2\right) - \Pi\left(n_1, n_2, \bar{n}_2\right) = -c\frac{(s-n_1-\bar{n}_2+n_2+1)(s-n_1-\bar{n}_2+n_2)(s-\bar{n}_2+2n_1-2n_2+2)}{6(s-n_1)(s-n_2)}, \qquad (99)$$

which is strictly negative for $\forall \bar{n}_2 \in [n_2, n_2 + s - n_1)$, and is zero for $\bar{n}_2 = n_2 + s - n_1$. The change in expected profits as a result of the direct disclosure effect is:

$$\Pi\left(n_1, n_1\right) - \Pi\left(n_1, n_2\right) = c\frac{(n_1-n_2)\left[(s-n_1)^2-1\right]}{3(s-n_1)(s-n_2)}, \qquad (100)$$

which is strictly positive for $s > n_1 + 1$, and equals zero for $s = n_1 + 1$.

The net change in expected profits for both firms, which is comprised of the two effects, is hence:

$$\Pi\left(n_1, n_1\right) - \Pi\left(n_1, n_2, \bar{n}_2\right) = \Pi\left(n_1, n_2\right) - \Pi\left(n_1, n_2, \bar{n}_2\right) + \Pi\left(n_1, n_1\right) - \Pi\left(n_1, n_2\right) =$$

$$= c\frac{2(n_1-n_2)\left[(s-n_1)^2-1\right] - (s-n_1-\bar{n}_2+n_2+1)(s-n_1-\bar{n}_2+n_2)(s-\bar{n}_2+2n_1-2n_2+2)}{6(s-n_1)(s-n_2)}. \qquad (101)$$

When $\bar{n}_2 = n_2$ the net value from disclosure is:

$$\left[\Pi\left(n_1, n_1\right) - \Pi\left(n_1, n_2, \bar{n}_2\right)\right]\big|_{\bar{n}_2=n_2} = -c\frac{(s-n_1+2)(s-n_1+1)}{6(s-n_1)} < 0, \qquad (102)$$

---

[19]It is also true that since firm j decides to quit earlier, its own expected profits should consequently increase. As the expected profits of firm i are a decreasing function of the time in which firm j is intended to quit, its profits must also increase.

while when $\bar{n}_2 = n_2 + s - n_1$ it is:

$$[\Pi(n_1, n_1) - \Pi(n_1, n_2, \bar{n}_2)]|_{\bar{n}_2 = n_2 + s - n_1} = c\frac{(n_1-n_2)\left[(s-n_1)^2-1\right]}{3(s-n_1)(s-n_2)} \geq 0. \qquad (103)$$

The marginal effect of $\bar{n}_2$ on the expected profits is:

$$\frac{\partial\left[\Pi(n_1, n_1) - \Pi(n_1, n_2, \bar{n}_2)\right]}{\partial\bar{n}_2} = c\frac{3(s-\bar{n}_2+n_1-n_2+2)(s-n_1-\bar{n}_2+n_2)+3(n_1-n_2)+2}{6(s-n_1)(s-n_2)}, \qquad (104)$$

which is strictly positive since $\bar{n}_2 - n_2 \leq s - n_1$ and $s > n_1 > n_2$. ∎

7 References

Arrow, Kenneth J., 1951, *Social Choice and Individual Values.* New York: John Wiley and Sons.

Arrow, Kenneth J., 1962, "Economic Welfare and the Allocation of Resources for Invention", *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pp. 609-626.

d'Aspremont, Claude, Sudipto Bhattacharya and Louise-Andre Gerard-Varet, 2000, "Bargaining and Sharing Innovative Knowledge", *Review of Economic Studies* 67, pp. 255-271.

Becker, Gary S., 1968, "Crime and Punishment: An Economic Approach", *Journal of Political Economy* 76(2), pp. 169-217.

Bhattacharya, Sudipto and Gabriella Chiesa, 1995, "Proprietary Information, Financial Intermediation, and Research Incentives", *Journal of Financial Intermediation* 4, pp. 328-357.

Bhattacharya S., J. Glazer, and D. Sappington, 1992, "Licensing and the Sharing of Knowledge in Research Joint Ventures", *Journal of Economic Theory* 56, 43-69.

Bhattacharya S. S. Guriev, 2006, "Patents vs. Trade Secrets: Knowledge Licensing and Spillover", *Journal of the European Economic Association* 4(6), 1112-1147.

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

Chatterjee, Kalyan and Robert Evans, 2004, "Rivals' Search for Buried Treasure: Competition and Duplication in R&D", *Rand Journal of Economics* 35, pp. 160-183.

Dasgupta, Partha and Eric Maskin, 1987, "The Simple Economics of Research Portfolios", *Economic Journal* 97, pp. 581-595.

Doraszelski, Ulrich, 2003, "An R&D Race with Knowledge Accumulation", *Rand Journal of Economics* 34, pp. 20-42.

Fershtman, Chaim and Ariel Rubinstein, 1997, "A Simple Model of Equilibrium in Search Procedures", *Journal of Economic Theory* 72, pp. 432-441.

Fudenberg, Drew, Richard Gilbert, Joseph Stiglitz and Jean Tirole, 1983, "Preemption, Leapfrogging and Competition in Patent Races", *European Economic Review* 22, pp. 3-31.

Green E. and R. Porter, 1984, "Noncooperative Collusion Under Imperfect Price Information," *Econometrica*, **52**, 87-100.

Grossman, Gene M. and Carl Shapiro, 1987, "Dynamic R&D Competition", *Economic Journal* 97, pp. 372-387.

Haller, Hans and Anthony Pavlopoulos, 2002, "Technological Failure, Economic Success", *International Game Theory Review* 4, pp. 415-434.

Loury, Glenn C., 1979, "Market Structure and Innovation", *Quarterly Journal of Economics* 93, pp. 395-410.

Rotemberg J. and G. Saloner, 1986, "A Supergame-Theoretic Model of Price Wars During Booms," *American Economic Review*, **76**, 390-407.

Silipo, Damiano Bruno, 2005, "The Evolution of Cooperation in Patent Races: Theory and Experimental Evidence", *Journal of Economics* 85, pp. 1-38.

Spiegel, Yossi, 2007, "Licensing Interim R&D Knowledge", *Working Paper.*

Stevenson, R., 1980, *Corporations and Information*, Baltimore and London: The Johns Hopkins University Press, New-York: Free Press.

Yosha, Oved, 2003, "Financing Innovation: Is Transparency a Two-Edged Sword?", *Working Paper*.