The Market for R&D Failures

Roy Shalem*†

June, 2007
Work in Progress

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 data that reveals their failures. This paper theoretically explores the possibility of trading the 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, and design a trading mechanism that would account for obstacles auch as moral hazard, verification, over-investment and private gains.

JEL Classification: D23, L42, O3

Keywords: R&D, Patent Races, Cooperation, Licensing, Contract Theory, Information.

*Berglas School of Economics, Tel-Aviv University; email: shalemro@post.tau.ac.il.

†I would like to thank Manuel Trajtenberg for his guidance. I would also like to thank the following people for their useful comments: Eddie Dekel, Lev Druker, David Gilo, Shane Greenstein, Joel Mokyr, Zvika Neeman, Marco Ottaviani, Mallesh Pai, Ron Siegal, Yossi Spiegel, Scott Stern, Rakesh Vohra, Yoram Weiss, Michael Whinston and participants of workshops at Tel-Aviv University and Northwestern University.

1 Introduction

"I have not failed. I've just found 10,000 ways that won't work." Thomas Edison.

Failures are probably the most commonly widespread by-product of R&D. Indeed, 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 regarding the 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 normally are exposed to much less information regarding the R&D failures of their competitors, than regarding their successes. The existence of this phenomenon is due to the fact that knowledge of successes tends to spillover, either by means of reverse engineering of new products, or because of the issuance of patents[2]. The natural human tendency of talking more about one's successes than failures exacerbates this phenomenon.

Granted that knowledge of failures does not spillover efficiently, it is quite natural to try to construct mechanisms that would enable such a 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 in the academic world there is an informal ongoing debate on the importance of publicizing failed research attempts.

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

[2] Indeed, one of the declared objectives of the patent system is the disclosure of such successes.

Our paper is related to the strands of the literature dealing with R&D cooperation, such as joint ventures and open science, but is closest to the strand dealing with licensing interim knowledge of R&D. The following paragraphs will try to place our contribution within this literature. It should be said upfront that although this literature has studied various forms of cooperation, to the best of our knowledge, this paper is the first to analyze the potential of a contracts between rivals which disclose information of failures of one agent for the purpose of helping the other's R&D endeavor, and the possible caveats of implementing them.

There are several features that distinguish our paper from previous literature. Firstly, papers which 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 can potentially be licensed. Unlike them, our contributions is to explicitly state that accumulated failures constitute he most prevalent form of interim knowledge; hence the analysis should focus on their specific caveats as well as possible advantages. Since our approach explicitly examines failures, it requires particular modeling choices, which subsequently be reflected in the results. Secondly, whereas the focus of most of these papers is on the licensing strategies and the bargaining power of the licensor versus the licensee, our paper focuses on characterizing the types of industries where a market for failures is more likely to thrive, and the kinds of contracts that would be helpful 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). d'Aspremont et al. (2000) has 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, and so there is 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 few firms try to develop a cost reducing innovation, which cannot be patented. The probability of success of each firm depends on the interim knowledge it possesses. After the race is over there *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 if exactly one firm wins the competition, which ultimately affects the licensing decisions of the firms. The interim knowledge in these papers has an affect 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 would 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 both that the interim knowledge is held by an outside research lab, which can 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 firms, and hence the outcomes 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 we will show, this assumption

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

[4]Bhattacharya and Chiesa (1995) have considered 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.

is not necessarily natural when the interim knowledge considered is a stock of failures.

A distinct exception of all this literature is Haller and Pavlopoulos (2002), who indeed recognized the potential commercial value of knowledge about past failures, that if properly exploited can serve in preventing future failures. Their research focuses on showing how an outside research lab that had accumulated some failures in the past, is able to extract the highest rent from current participates of a patent race. Though this question is in itself an interesting and innovative one, the fact that they have focused their attention on such a setup implies trivially that the research lab would like to sell its knowledge. In one of the extensions they analyze whether the lab would like to get in the race or not, yet again they do not explicitly analyze whether a market for failures between rival firms is plausible, and what can prevent such a market from existing. Finally, as their analysis begins after the failures have already been accumulated, they do not consider whether there is potential distorting effect of such a transaction on the research effort of the lab.

The paper is divided into two main parts. In the first part we theoretically explores the merits in establishing a market of R&D failures. The fact that knowledge about R&D failures has real economic value is a positive result that gives rise to the possibility that a market for that knowledge could be constructed. After establishing that, one cannot help wondering why such markets have hardly invaded our reality. The second part of the paper examines some of the theoretical obstacles that may impede the emergence for such a market. It will also attempt to tackle some of these obstacles, perhaps as a first step in an effort to introducing a market for failures into the economic reality.

2 The Merits of a Market for R&D Failures

There are at least three 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. A second is the ability to update one's beliefs regarding the technological space, based on the failures. A third possible merit is in the ability to avoid certain catastrophes or other costs which are associated with failures. All these factors may both reduce the average cost of R&D, as well as shorten the average time required in attaining a success. Since both the effective cost of R&D and its potential gains are higher, the existence of a market for failures may also ultimately imply that in equilibrium more R&D will be conducted.

This section will explore frameworks of patent races, in order to illustrate the aforementioned merit of avoiding duplications, and to analyze whether trading in R&D failures may be possible in this context. The analysis of the two other merits and exploration non-trivial economic contexts, other than patent races, are left for future versions of this paper.

2.1 A Simple Patent Race Model with Failures Trading

As mentioned in the introduction, it is quite reasonable to assume that knowledge about past failures allows for better progress in R&D. Nevertheless, the patent race literature at 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, for describing the time when success should occur. Distributions of this kind are equivalent to sampling with replacement. Hence, using them implicitly assumes away the value of failures, and so the literature has strayed away from analyzing the possible economic value of knowledge of failure. In recent years, however, numerical models have started to get more accepted in the Industrial Organization literature, and so some change in approach can be seen, as in Doraszelski (2003) who considers knowledge accumulation.

Fershtman and Rubinstein (1997) have studied a simple game-theoretic model of searching for a prize, hidden in one of many boxes. This framework was constructed as an analogy to a patent race. They have found three intrinsic inefficiencies in patent races; one has to do with the fact that the different players might frequently search the same empty box. This property is a result of the fact that none of the players has any information about the past failures of the others. The researchers, however, do not attempt to offer a mechanism that would solve the problem. They have also assumed that there is no per-period or per-sampling costs of R&D. Haller and Pavlopoulos (2002) have used a similar framework, but 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 like to trade the knowledge of their accumulated failures.

2.1.1 The baseline model

Assume that two pharmaceutical firms are involved in a patent race for the development of a new drug, and only the firm that develops it first would be able to patent it and then receive a lump sum of V. If both firms finish the development process at the same time, each will have an equal chance for patenting the drug. The development process involves testing the effect that each of s possible chemical substances has on the disease. For simplicity, assume that the firms know that exactly one of these substances will be effective for treating the disease. Each of these substances is assumed to have an equal chance of working. 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 our analysis starts from. Without loss of generality, we assume that at time t firm 1 is not falling behind firm 2 in the race, namely that

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

For simplicity assume, for now, that both firms are risk natural and time indifferent.

We assume that the knowledge of its failed attempts is private information to each of the firm. We also assume that the sampling is done randomly, when the set is restricted to substances which have not been tested yet. 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 would 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 will test a substance which has already been tested by its rival.

The expected profits of the firms at period t, is a function of the stock of failures known to each of them:

$$\begin{cases} U_1(n_1,n_2) = V\frac{1}{s-n_1}\left(\frac{s-n_2-1}{s-n_2}+\frac{1}{s-n_2}\frac{1}{2}\right) - c + \frac{s-n_1-1}{s-n_1}\frac{s-n_2-1}{s-n_2}U_1(n_1+1,n_2+1) \\ U_2(n_1,n_2) = V\frac{1}{s-n_2}\left(\frac{s-n_1-1}{s-n_1}+\frac{1}{s-n_1}\frac{1}{2}\right) - c + \frac{s-n_1-1}{s-n_1}\frac{s-n_2-1}{s-n_2}U_2(n_1+1,n_2+1) \end{cases} \tag{2}$$

After a few mathematical manipulations we can transform the recursive form of the payoff function into a normal form:

$$\begin{cases} U_1(n_1,n_2) = V\frac{s+n_1-2n_2}{2(s-n_2)} - c\left[1+\frac{(s-n_1-1)(2s+n_1-3n_2-1)}{6(s-n_2)}\right] \\ U_2(n_1,n_2) = V\frac{s-n_1}{2(s-n_2)} - c\left[1+\frac{(s-n_1-1)(2s+n_1-3n_2-1)}{6(s-n_2)}\right] \end{cases} \tag{3}$$

We assume that V is large enough, relative to c, such that $U_1, U_2 > 0 \quad \forall n_1, n_2$. A sufficient condition for that is:

$$V > 2sc. \tag{4}$$

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 would win the race. Hence, the firm which performs the disclosure incurs an ex-ante loss. There can be a potential economic added-value of disclosure if this cost 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 is a non-empty interval of market prices which can be used in a potential transaction between the two firms, one which improves the ex-ante profit of both. The next stage will be showing that such a transaction is actually always feasible. Combined, this will prove the fact that the incompleteness of markets for research failures creates economic inefficiencies.

We start by first considering disclosure of failures by firm 2. For the simplicity of the analysis, suppose that both firms hire a third party to which all the failures are disclosed. This third party can then identify which failures one firm has encountered, which its rival has not. These failures can potentially have an economic value. We begin by analyzing the case in which firm 2 sells some of its failures to firm 1. Let k denote the number of chemical substances which are disclosed in such a transaction by firm 2 and are not yet tested by firm 1. The disclosure value is then:

$$DV_2(n_1, n_2, k) = U_1(n_1 + k, n_2) - U_1(n_1, n_2) + \tag{5}$$
$$+U_2(n_1 + k, n_2) - U_2(n_1, n_2) = \frac{ck(k + s + 2n_1 - 3n_2)}{3(s - n_2)}.$$

Proposition 1: *In a two-player patent race, there is always a non-empty interval of prices, that supports a transaction in which the firm which is behind in the race sells information about its research failures, to the one which is ahead. Furthermore, the added value of such a transaction is monotonically increasing in the amount of disclosed failures, and represents the expected cost reduction of the buyer.*

Proof. When $k = 0$, then by definition the disclosure value in equation (5) equals zero. The marginal contribution of a disclosed failure is:

$$\frac{\partial[U_1(n_1, n_2) + U_2(n_1, n_2)]}{\partial n_1} = \frac{c[2(n_1 - n_2) + (s - n_2)]}{3(s - n_2)} > 0. \tag{6}$$

Hence, the disclosure value is always positive and monotonically increasing in the number of failures disclosed. The length of the interval of prices is the net disclosure value, which appears in equation (5). The length of the interval is strictly positive since $n_1 \geq n_2$ and $s > n_2$. ■

Note that the disclosure value represents solely the difference between the gain to the buyer and the loss to the seller. However, the actual price charged represents also the compensation that the buyer needs to pay the seller for the decline in the probability that the latter would win the race. The followings are the maximum price that firm 1, the buyer, would be willing to pay versus the minimum price that firm 2, the seller, would charge:

$$U_1(n_1+k, n_2) - U_1(n_1, n_2) = \frac{k\left[c(k+s) + 3V + 2cn_1 - 3cn_2\right]}{6(s-n_2)} \tag{7}$$

$$U_2(n_1+k, n_2) - U_2(n_1, n_2) = \frac{k\left[c(k+s) - 3V + 2cn_1 - 3cn_2\right]}{6(s-n_2)}. \tag{8}$$

We shall now calculate the added value of a disclosure of failures by firm 1. Note, that it is possible that following the disclosure, firm 2 will take the lead. Explicitly, condition (1) might not hold, following the disclosure. However, we know based on Proposition 1 that once $n_1 = n_2$, the disclosure value is a positive and increasing function of the number of disclosed substances by either one of the firms. Therefore we only have to check for the range in which following the disclosure firm 1 does not lose its lead. Let k now denote the number of disclosed chemical substances which are disclosed by firm 1 and were not originally tested by firm 2. In cases in which condition (1) does hold following the disclosure, the value of disclosure is

$$\begin{aligned} DV_1(n_1, n_2, k) &= U_1(n_1, n_2+k) - U_1(n_1, n_2) + \\ &+U_2(n_1, n_2+k) - U_2(n_1, n_2) = \frac{ck\left[(s-n_1)^2 - 1\right]}{3(s-n_2)(s-k-n_2)}. \end{aligned} \tag{9}$$

Proposition 2: *In a two-player patent race, as long as* $s > n_1 + 1$*, there is a non-empty interval of prices, that supports a transaction in which the firm which is ahead in the race sells information about its research failures, to the one which is behind. Furthermore, the added value of such a transaction is monotonically increasing in the amount of disclosed failures, and represents the expected cost reduction of the buyer.*

Proof. When $k = 0$, then by definition the disclosure value in equation (9) equals zero. As long as firm 1 does not lose its lead the marginal contribution of a disclosed failure is:

$$\frac{\partial[U_1(n_1, n_2) + U_2(n_1, n_2)]}{\partial n_2} = \frac{c\left[(s - n_1)^2 - 1\right]}{3(s - n_2)^2}, \tag{10}$$

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 is ahead. Taking into account the results of Proposition 1, it follows that even if firm 1 loses its lead, the disclosure value is still monotonically increasing in the number of failures disclosed. As the disclosure value is always positive, the length of the interval of prices must also be strictly positive. ■

Proposition 2 implies that there is a positive economic value from a disclosure of information by the firm which is ahead to the one which is behind. Note that when the condition $s > n_1 + 1$ does not hold, firm 1 is known to be only one period away from discovery. This, of course, means that any disclosure cannot shorten the time till the discovery is made, and thus there would be no potential cost reduction by means of disclosure, but rather just a possible different split of the value V between the two firms.

The important implication of Proposition 1 and Proposition 2 is that there is always a range of prices within which the failures can be traded. The propositions also imply that there should always be a full disclosure of all the accumulated failures by both firms. These results constitute a strong indication in the support of 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, and when cooperation following the discovery is impossible or weak.

Note that under the assumptions made here, from the collective point of view of the firms the only merit in a disclosure of a failure is in the fact that it could prevent the other firm from making the same mistake at a cost of c. 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 would win the race.

It is worth discussing the practicality of the assumption that only relevant chemical substances are disclosed, and the existence of the third party mediating the transaction. Note that both firms are indifferent to the disclosure of an irrelevant substance, namely one which is originally included both in n_1 and in n_2. Using this fact we can deduct that both Proposition 1 and Proposition 2 also hold when some irrelevant failures are disclosed. Still, the added 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. Moreover, 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, he is still needed in order to verify that every failure disclosed indeed represents a substance which has been tested - an issue that will later be addressed.

As mentioned in Proposition 1 and Proposition 2, the added value is not a function of the value of the invention, V, but rather of the cost of development, c. Hence, the entire added value is created from the cost reduction of the R&D process. Had the cost of development been zero, the R&D race would have actually been a zero-sum game. In that case there would have been no means for creating an economic added value through the disclosure of failures.

2.1.2 A model with discounting

Alternatively, assume now that the firms are not time indifferent, and instead they use a discounting factor, $\delta \in (0,1)$. The recursive form of the expected profits of the firms at period t is then:

$$\begin{cases} U_1(n_1,n_2) = \frac{V}{s-n_1}\left(\frac{s-n_2-1}{s-n_2} + \frac{1}{s-n_2}\frac{1}{2}\right) - c + \frac{s-n_1-1}{s-n_1}\frac{s-n_2-1}{s-n_2}\delta U_1(n_1+1,n_2+1) \\ U_2(n_1,n_2) = \frac{V}{s-n_2}\left(\frac{s-n_1-1}{s-n_1} + \frac{1}{s-n_1}\frac{1}{2}\right) - c + \frac{s-n_1-1}{s-n_1}\frac{s-n_2-1}{s-n_2}\delta U_2(n_1+1,n_2+1) \end{cases} \tag{11}$$

Proposition 3: *In a two-player patent race with discounting, there is a non-empty interval of prices, that supports a transaction in which information of failed research is sold either by the firm which is ahead to the one which is behind, as long as $s > n_1 + 1$, or by the firm which is behind to the one which is ahead. In either case the added value of such a transaction is monotonically increasing in the amount of disclosed failures, and represents both a reduction of the expected costs of the buyer and a drop in the average time until the discovery is made.*

Proof. See Appendix 2.

When the discounting factor is incorporated into the model, the added value from disclosure is comprised of both the reduction in the research costs and the less discounted value of discovery (i.e. V), caused by the shortened average time until the discovery is made, as it is always a positive function of the term $[(1-\delta)V + 2\delta c]$. 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.

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

Having established the possible merits of a market for R&D failures, we now turn to think 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. Indeed, trading any kind of interim knowledge should prove to be challenging, as foreseen by Kenneth Arrow in his 1962 seminal work: "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.".

Knowledge of R&D failures poses its own challenges, stemming from the specific characterization of failures. One such trait, which differentiates between failures and successes, is that former can be relatively easily and inexpensively generated. A second challenge results from the fact that when knowledge of failures is bought, as opposed to most commodities, the purpose of the purchase 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 to 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 example of failures trading in a simple patent race, in which the existence of a market for failures is found to be very plausible. We shall relax some of the assumptions and otherwise alter this simple model in order to surface possible difficulties which are likely to occur in reality, when trying to trade knowledge of R&D

failures. In addition we shall try to explore the kind of contracts which could serve the trading parties in overcoming some of these difficulties.

3.1 Moral Hazard

Each of the firms in our patent race, when being paid to disclose its failures, has an incentive to deceive its rival. It may, thus, be tempted to report as failures substances it has never actually tested. The possibility of such a hidden action by the seller might, therefore, be a potential impediment on the performance of the market. The challenge is, hence, to construct contracts which will ensure that the information sold, regarding the past failures, is indeed accurate.

Suppose in our baseline model that one firm sells inaccurate knowledge of failures to its competitor, which in fact includes the successful substance. Once one of the firms will find and patent this substance, it will unravel the lie of the untruthful seller. We therefore have in our model a one-to-one link between the hidden action of the seller, namely his potential lie, and a verifiable and contractible element, namely getting a patent over the right substance. This link implies that the hidden action in the baseline model is, in fact, hidden only for a limited time, making it very easy to overcome, contractually, as will be explained later.

Since the problem of untruthful disclosure does not fully present itself in the baseline model, in this section we would like to abstract away from our baseline model and speak more generally on how such moral-hazard may affect the market for failures. We would like to account for extensions such one in which the probability that a success will eventually be found may be less than one, for instance due to budget constraints. The importance of this extension is that it implies that there might be a positive probability that a fraud would never be detected. For instance, consider a case in which a hundred substances have not yet been tested, but each of the firms has only enough resources to check 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. In these contexts the buyer would be forced to engage in active auditing in order to ensure truthful revelation by the seller.

Another type of extension which this section is applicable to, involves a hidden effort level in the research endeavor of the seller. For instance, suppose that like in our baseline model the per-testing cost, c, reflects the testing cost when the firms exert the maximum effort. However, suppose, in oppose to our baseline model, that the firms can also choose a lower effort level, such that the cost will be lower, but so will the probability that if the successful substance is tested it will be identified as such. This extension exhibits the aforementioned feature of failures, which refers to the ability be relatively easily and inexpensively generate them. In such a model even if the seller can prove that he actually performed the tests and failed, the quality of the information sold is still something which needs to be addressed contractually.

The following analysis concentrates on comparing the performance of a market for failures in which deceiving is impossible to a similar, counterfactual, market in which deceiving the rival firm is possible. We will call the outcome in the market with no deception "the first-best outcome". Once again, we will have two competing firms, where firm j would want to sell knowledge of failures to firm i. Let U_i and U_j denote the expected payoffs of firm i and firm j prior to the transaction, and U_i^* and U_j^* denote their payoffs after the transaction, given that the seller, j, is truthful. For such a transaction to be feasible in the world with no possibility of deception, the net disclosure value has to be positive, meaning that:

$$U_i^* + U_j^* - U_i - U_j > 0. \tag{12}$$

Now in the counterfactual world, in which the seller can deceive the buyer, let U_i^{**} and U_j^{**} denote the payoffs to the firms, should firm i act as if firm j is truthful, yet firm j deceives firm i by selling some fake test results of tests it has never performed. Let $P_S \in (0, 1]$ denote the probability that if the seller was deceitful, a success would in fact exist amongst the sold, so called, failures. Moral-hazard implies that:

$$U_j^{**} > U_j^*, U_i^{**} < U_i^*. \tag{13}$$

We assume, as in the baseline model, that with probability 1, none of the failures sold by a truthful seller can potentially be a success. Suppose now that the seller is deceptive. Let P_N denote the probability that a fraud would eventually be detected, meaning 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 intended to pursue before the contract was signed. We will assume that:

$$P_U > 0,$$

which implies that prior to the transaction firm i had a positive probability of performing on its own the test results it bought. 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 4 (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. The contract includes an upfront payment by the buyer, p, and a fine, F, which would be paid by the seller should one of the test results sold would turn out to be a success. The following two constraints (IR_j and IR_i, respectively) ensure that both firms would like to participate in the transaction, given that the seller is truthful:

$$U_j^* - U_j + p \geq 0 \tag{14}$$

$$U_i^* - U_i - p \geq 0. \tag{15}$$

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

$$U_i^* - U_i \geq p \geq U_j - U_j^*. \tag{16}$$

The following constraint (IC$_j$) ensures that firm j, the seller, does not want to deceive the buyer:

$$U_j^* - U_j + p \geq U_j^{**} - U_j + p - P_S P_N F, \tag{17}$$

which is equivalent to:

$$F \geq \frac{U_j^{**} - U_j^*}{P_S P_N}. \tag{18}$$

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

Since in equilibrium both the buyer and the seller follow the research strategy they would follow in a world with no deception, the contract implements the first-best outcome. ■

Proposition 4 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 would be discovered naturally.

Before proceeding to the case in which $P_N = 0$, we would like to comment on another extension, which makes an alternative, and perhaps more realistic, assumption that even if the seller was truthful there would still be some positive probability that one of the failures sold is in fact a potential success. In our baseline model this might be the result of a setting in which when the right substance is tested, there is a positive probability that the test would go wrong and indicate the substance to be a failure. We assume, as before, that the firms sustain their R&D efforts until the success is finally found. Now even if the seller is truthful, there is still a positive probability $\hat{P}_S \in (0,1)$ that one of the failures he sells is in fact a success. This would then imply that there is a strictly positive probability that he would still have to pay the fine (this is sometimes referred to as a type-1 error). The following two constraints (IR_j and IR_i, respectively) ensure that both firms would agree to participate in the transaction, given that the seller is truthful:

$$U_j^* - U_j + p - P_N \hat{P}_S F \geq 0 \tag{19}$$

$$U_i^* - U_i - p + P_N \hat{P}_S F \geq 0, \tag{20}$$

which simply implies that now the upfront payment compensates the truthful seller also for the probability that he would have to pay the fine. The following constraint (IC_j) ensures that firm j, the seller, would not want to deceive the buyer:

$$U_j^* - U_j + p - P_N \hat{P}_S F \geq U_j^{**} - U_j + p - P_N P_S F, \tag{21}$$

which is also not difficult to satisfy, as long as $\hat{P}_S < P_S$, meaning that whatever had been tested and failed is more likely to actually be a failure than what had not been tested. 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, it 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 test result being sold would turn out to be a success.

Examining the incentives of the buyer, reveals a more substantial difference from the model in which $\hat{P}_S = 0$. It may be possible for buyer to affect the probability that the fine would be paid by the seller, by distorting his own research plan. Any such deviation by the buyer would yield $U_i^{**} \leq U_i^*$, and increase the probability of discovering a success if it exists to $P_D \geq P_N$. The IC constraint for firm i implies that for any such deviation the following condition holds:

$$U_i^* - U_i - p + P_N \hat{P}_S F \geq U_i^{**} - U_i - p + P_D \hat{P}_S F, \tag{22}$$

which is equivalent to:

$$F \leq \frac{U_i^* - U_i^{**}}{(P_D - P_N)\,\hat{P}_S}. \tag{23}$$

This condition is always satisfied if, like in our baseline model, success is guaranteed to occur, that is $P_N = 1$. In other cases, this upper bound constraint on the size of the fine, F, might ultimately hinder the implementation of the first-best outcome.

Now we shall move on to examine the case in which $P_N = 0$, where without some form of active auditing, a deceptive seller would never be identified as such. In our model a natural way for the buying firm to verify that what was disclosed to it is indeed failures, is by testing them. We shall consider a specific form of auditing in our model, which is for the buying firm to return to the research plan it had prior to the transaction. Evidently, in such a reversal, or 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[5]. The following proposition summarizes the outcome of implementing such a contract:

Proposition 5 (Auditing via reversal): *If deception is possible by the seller, and $P_N = 0$, then an outcome infinitely close to the first-best is attainable by imposing an infinitely large state-contingent fine. A similar implementing with a finite fine would result in an outcome different from the first-best, though it would not distort the transactions made in the market for failures.*

[5] Note that we assumed that firm i can commit to the auditing process. This can be implemented by keeping the knowledge sold sealed in an envelope at the hand of a third party. This third party then performs a lottery in order to determine whether an auditing will be implemented. If the auditing has to be implemented then the sealed envelope is only disclosed to the buyer after the race is over.

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

$$(1 - P_A)\left(U_j^* - U_j\right) + p \geq 0 \tag{24}$$

$$(1 - P_A)\left(U_i^* - U_i\right) - p \geq 0. \tag{25}$$

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

$$U_i^* - U_i \geq \frac{p}{1 - P_A} \geq U_j - U_j^*. \tag{26}$$

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

$$(1 - P_A)\left(U_j^* - U_j\right) + p \geq (1 - P_A)\left(U_j^{**} - U_j\right) + p - P_A P_S P_U F, \tag{27}$$

which is equivalent to:

$$F \geq \frac{(1 - P_A)\left(U_j^{**} - U_j^*\right)}{P_S P_A P_U}. \tag{28}$$

Hence, for any $P_A \in (0,1)$, there exist fines large enough, guarantying that the seller would be truthful. Since this is true $\forall U_j^{**}, P_S \in (0,1), P_A \in (0,1)$, then $\exists\ F$ which would discourage any deviation from the truthful strategy. Together with the fact that there is always a price, p, which can satisfy the IR conditions, this implies that $\forall P_A \in (0,1)$, the transactions in the market for failures are undistorted.

The net disclosure value in that case is $(1 - P_A)\left(U_j^* - U_j + U_i^* - U_i\right)$, which is infinitely close to the first-best only if P_A is infinitely small and the fine, F, is hence infinitely large. ■

Proposition 5 presents an approach not too different from the "high fine-low probability" result by Becker (1968)[6], 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 from the transaction indeed drops, but it cannot drop so much as to make the transaction unprofitable for the firms.

3.2 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), it becomes clear that it is typically the case that the overinvestment problem is exacerbated when the two rivals are running "neck to neck", i.e. when their probabilities of success is close to each other. One can therefore speculate that perhaps in our model the trading of failures may have an effect on that issue. This section explores the question of overinvestment in our model as well as the effect that the market for R&D failures has on it.

Though not straightforward, the easiest way to model the R&D intensity of firm 1 and firm 2, is by changing n_1 and n_2 respectively. For example, suppose that $s = 100$, $n_1 = 50$ and that, as before, both firms have a search intensity of one substance per period. If firm 1 wanted

[6] In Becker's model of crime fighting he 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 equal to an 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.

to change its research capacity such that it would test now two substances in every period, then it would seem as if firm 1 had remaining 25 pairs of substances, and was testing a pair in each period. In terms of the payoff functions of the two firms, which appear in equation (2), it would seem as if firm 1 has increased n_1 to being equal to 75. More generally, if firm $i \in \{1,2\}$ wants to change its research capacity by some factor $k_i \geq 1$, the term n_i in the payoff functions of the two firms will 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 above example each pair tested by firm 1 should cost $2c$ to test. After accounting for all that, the payoff function, not including the costs of establishing the research capacity, will be:

$$\begin{cases} U_1(n_1,n_2,k_1,k_2) = V\frac{2k_1(s-n_2)-k_2(s-n_1)}{2k_1(s-n_2)} - ck_1\frac{(s+k_1-n_1)[k_1(3s+k_2-3n_2)-k_2(s-n_1)]}{6k_1^2(s-n_2)} \\ U_2(n_1,n_2,k_1,k_2) = V\frac{k_2(s-n_1)}{2k_1(s-n_2)} - ck_2\frac{(s+k_1-n_1)[k_1(3s+k_2-3n_2)-k_2(s-n_1)]}{6k_1^2(s-n_2)} \end{cases} \tag{29}$$

Furthermore, we now replace assumption (1) with the more elaborated assumption stating that firm 1 is ahead in the race, that is:

$$\frac{s-n_2}{k_2} \geq \frac{s-n_1}{k_1}. \tag{30}$$

For example, assume that $s = 100$, condition (30) allows for the case where $n_1 = 60$ and $n_2 = 70$ if for instance the research capacity of firm 1 is two substances per period ($k_1 = 2$) and the research capacity of firm 2 is one per period ($k_2 = 1$), consequently firm 1 has a higher chance of winning the race, since it is as if it only has 20 more pairs to sample, which is equivalent to $n_1 = 80$, with one substance tested in 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 equals three, and the right substance is one of the three substances currently being tested by the firm. Because the firm is testing all three substances simultaneously, it has to incur the testing cost, c, for all three. Alternatively, if the intensity of the firm was one, once the right substances is tested, the firm would stop and hence is potentially able 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 $f(k_i) > 0$. This investment is assumed to be irreversible, in the sense that the firm cannot decrease its research capacity, once it was built to a certain level. However, our assumption is that research capacity can 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 $f(\hat{k}_i) - f(\bar{k}_i)$.

Suppose that the race starts when $n_1 = n_2 = 0$. It is obvious, then, that since there is no discounting in the baseline model, 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:

$$\begin{cases} U_1^c = V - c\frac{s+1}{2} \\ U_2^c = 0 \end{cases} \tag{31}$$

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

$$U_1^p = U_2^p = \frac{1}{2}V - c\frac{(s+1)(2s+1)}{6s}. \tag{32}$$

$$V - c\frac{2s^2+3s+1}{3s}. \tag{33}$$

Note that since it is assumed that $V > 2sc$, then $U_2^p > 0$, and so 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 would always be greater than the collective gains:

$$\begin{cases} \frac{\partial U_1}{\partial k_1} - \frac{\partial (U_1+U_2)}{\partial k_1} = \frac{k_2(s-n_1)[2ck_2(s-n_1)+3k_1(V-cs+cn_2)]}{6k_1^3(s-n_2)} \\ \frac{\partial U_2}{\partial k_2} - \frac{\partial (U_1+U_2)}{\partial k_2} = \frac{ck_1^2+(s-n_1)(3V-cs+cn_1)}{6k_1(s-n_2)} \end{cases} \tag{34}$$

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

After establishing that, as it is usually the case, overinvestment is a feature of our patent race, we now move to analyze the way in which incentives to invest in research capacity could change, depending on the stock of failures possesses by each firm. The dynamic problem, which involves both uncertainty and irreversibility of investment, is left beyond the scope of this paper. We will concentrate on the myopic marginal willingness of each firm to invest in research capacity, relative to its position in the race, which will be denoted by γ_i:

$$\begin{cases} \gamma_1 \equiv \frac{\partial U_1}{\partial k_1} = \frac{k_2(s-n_1)(3V-cs+cn_1)-ck_1^2(3s+k_2-3n_2)}{6k_1^3(s-n_2)} \\ \gamma_2 \equiv \frac{\partial U_2}{\partial k_2} = \frac{2ck_2(s-n_1)^2-ck_1^2(3s+2k_2-3n_2)+3k_1(s-n_1)(V-cs+cn_2)}{6k_1^2(s-n_2)} \end{cases} \tag{35}$$

Proposition 6: *A transaction which marginally increases the lead of firm* 1*, by increasing its stock of failures, would decreases the myopic willingness of both firms to invest in research capacity. In contrast, a transaction which marginally advances firm* 2*, yet keeping it the laggard would increases the myopic willingness of both firms to invest in research capacity, unless the leader is already in the midst of the last period of research.*

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

$$\begin{cases} \frac{\partial \gamma_1}{\partial n_1} = -\frac{k_2(3V-2cs+2cn_1)}{6k_1^2(s-n_2)} \\ \frac{\partial \gamma_2}{\partial n_1} = -\frac{4ck_2(s-n_1)+3k_1(V-cs+cn_2)}{6k_1^2(s-n_2)} \end{cases} \tag{36}$$

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 keeping it the laggard, would affect γ_i in the following way:

$$\begin{cases} \frac{\partial \gamma_1}{\partial n_2} = \frac{k_2(s-n_1)(3V-2cs+2cn_1)+ck_2\left[(s-n_1)^2-k_1^2\right]}{6k_1^2(s-n_2)^2} \\ \frac{\partial \gamma_2}{\partial n_2} = \frac{3Vk_1(s-n_1)+2ck_2\left[(s-n_1)^2-k_1^2\right]}{6k_1^2(s-n_2)^2} \end{cases} \tag{37}$$

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 would increase the myopic willingness of both firms to invest in research capacity. If, in contrast, $s - n_1 < k_1$, then direction of this effect is inconclusive. ■

To summarize, then, we have found that, except for the case in which firm 1 is sure 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. These results are close in spirit to those of Grossman and Shapiro (1987), though in their model the position of each firm in the race was 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 and is not directly necessary for making the discovery. It, therefore, follows that in our model a transaction in which firm 1 sells failures to firm 2, such as the ones referred to in Proposition 2, the positive inherent 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, as the cost of increasing research capacity, $f(k_i)$, is extremely high, as was implicitly assumed in the baseline model. In other contexts, however, creating a credible commitment not to overinvest 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 firm $i \in \{1, 2\}$ agrees to pay a monetary sum of $\alpha_i V$ to its rival, where $\alpha_i \in (0, 1)$, should firm i win the patent race. The contract may also include a non-contingent side-payment, $\bar{\alpha} \in (-\infty, \infty)$, which will ensure that both firms would like to sign the contract. 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 is to 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.

3.3 Strategic Manipulation

It seems that in reality one of the main reasons which makes firms reluctant to partly cooperate with their rivals is the fear that the rival would be able to infer additional information from the transaction, which will ultimately turn it to be unprofitable for the seller. For instance, it is fairly easy to extend our baseline model in order to account for the possibility that different groups of substances have different costs and probabilities of success, by introducing a menu of prices which will correspond to this diversity. However, a more challenging problem arises when the two firms are differently informed regarding the success rates, and so the mere announcement of a menu of prices by one of the firms may 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 failures, may distort its own research effort, in order to send its rival down the wrong path. Such a behavior was demonstrated in Chatterjee and Evans (2004), in a dynamic patent race, in which the firms observe each other's research choices.

The following example will have features which resemble the setting used by Chatterjee and Evans (2004) in order to demonstrate the disinformation effect. Suppose in our baseline model that before the race begins, firm 1 with probability $\theta \in (0, 1)$ gets 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 observed a signal.

In case firm 1 does not try to misinform firm 2, then whenever it receives the signal it is as if it accumulated $\frac{s}{2}$ failures. As we have shown in the baseline model, there is a market price in 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 handled in the same manner we have demonstrated above.

Suppose now that there is a potential equilibrium with disinformation. First note that firm 1 cannot sell false information to its rival, as such a possibility can be treated by using a contract which addresses moral-hazard problems. Therefore the only possibility in our model that disinformation may present itself, is if firm 1 decides not to sell its entire information to its rival, when it is informed, and firm 2 does not know for sure if firm 1 is informed or not. Suppose that such disinformation prevails in equilibrium, and that it induces 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 therefore be willing to pay any price lower than $(\bar{P}_2 - P_2)V - c(d - D)$, while firm 1 would be willing to sell for any price higher than $(\bar{P}_2 - P_2)V + c(d - D)$. This implies than 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.

Why is it, then, that our model does not exhibit strategic manipulation, similar to the one of Chatterjee and Evans (2004)? The two main differences are that in their model the research direction of each firm is observable, and there are no side-payments. This combination implies that by distorting one's own research effort it might impede the rival's chances of success. In 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 which would induce them to distribute their privately held information. One of the key features that promotes efficiency in our model is that both firms agree on the value of the information held by the informed firm. This would 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 would try to deceive its rival. The problem then boils down to the one we introduced in the moral-hazard section, and the contractual solutions we offered in that section can be implemented in the same manner.

3.4 Information Leakage and Resale

Generally speaking, knowledge, as opposed to other commodities, can be resold at a very low, or even zero, cost. More than in the case of licensing positive interim knowledge, a market for R&D failures may suffer from the possibility of information resale. Since there are no structural IPR for failures, there may be a problem to enforce exclusive contracts. Even if an exclusive agreement is contracted when the failures are sold, it would be particularly hard to prove infringement since the knowledge of failures it is important in the sense it tells you what **not** to 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, because of this inability to prove infringement, a black market for R&D failures may rise. Suppose, for example, that a formal market for failures exists. An employee from one of the firms may be tempted to sell the knowledge of failures made by his employer to its competitor. Unless he is caught in the act of selling this information, it would probably be impossible for the firm to know if such an information leakage has occurred.

Unfortunately, as opposed to the previous issues we have discussed, it seems that this issue cannot be solved contractually. Nevertheless, on the bright side it also implies that through labor mobility we may see **some** failures dissimilate naturally, even without a market for failures.

3.5 Infinite Set of Likely Research Possibilities

Finally, it puzzled us that, generally speaking, research in Economics does not allow for publicizing failures. It seems to us that one of the plausible explanations for that phenomenon is that there is a large, or even infinite, number of research directions, and a strictly positive cost of processing research results of others. In such a model, even if one accumulates a lot of failures, their contribution to achieving a success may be very small, and the costs thus may turn it unprofitable. Only if one can infer from the results on possibilities which has not been directly tested, then such a knowledge will be valuable. Probably the most prominent examples of such results, which helped avoiding dead-ends, are impossibility theorems such as Arrow (1951) or Myerson-Satterthwaite (1983), though we feel extremely uncomfortable naming these contributions "failures". The model of d'Aspremont, Bhattacharya, and Gerard-Varet (2000), which was mentioned in the introduction is, in a way, consistent with these kinds of results, since the Poisson distribution they assume implies that both with and without that knowledge the firms sample with replacement, hence such impossibility theorems, maintain an infinite sampling set, while increase the probability of success.

Another possibility is that there is a number of research objectives, and so without knowing the exact objective of his competitors, one may not be able to trade failures. Obviously potential competitors may be reluctant to disclose their exact research objective, in that case, in order to restrict the number of participants in the race.

4 Summary and Discussion

Knowledge of past R&D failures constitutes an important part of the stock of knowledge necessary for progressing R&D. The importance of such knowledge stems from the fact that failures are a natural by-product of the R&D activity, and so they are widespread 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 among competitors, than information regarding success. It is therefore particularly interesting to study whether they are likely to be traded between competitors, and why is it that in reality markets for failures are so scarce.

We show that at least in a patent race environment, direct competitors are likely to engage in the trading of R&D failures. However, a scrutiny of the specific characterizations of R&D failures as a 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 be able to avoid them, it may induce the seller to try to deceive the buyer by lowering the quality of the information sold. There may also be all sorts of externalities involved with such a transaction. This paper both analyzes these complexities and offers mechanisms which may help dealing them. It seems that the free market might prove to be an efficient device, by means of contingent and non-contingent side-payments between the firms, which allows overcoming many of the obstacles and align the incentives of competitors.

An additional potential difficulty in trading knowledge of R&D failure results from the lack of structural IPR over R&D failures, as opposed to the protection that patents provide for R&D successes. It seem to us that the lack of IPRs over R&D failures poses one of the toughest obstacles in the way of introducing market for failure into reality, and it might call for institutional changes rather than contractual solutions.

As a final remark, we also like to mention that our study also has some indirect implications on how to publish academic research. For instance, we believe that editors of the academic journals should encourage the authors to courageously also publish their failing attempts, if they suspect other researchers can make use of it. We were happy to find the existence of the Journal of Negative Results in Biomedicine[7], which seems to implement the approach which encourages distributing of such a knowledge.

We feel that the exploration of negative knowledge, such as one of failures, remains a fertile research ground. In future research it is our purpose to further explore the possible sharing of failures knowledge in non-trivial economic contexts, other than within patent races. We are sure that many of the features of contracting such a knowledge, which we demonstrated in this paper, would carry through to other contexts, while new features would surely present themselves.

5 Appendix

5.1 The Decision Regarding the Testing Sequence

Suppose that a market for failures does not exist. In the baseline model we have assumed that each of firms tests its remaining substances randomly; the question, however, is whether this is

[7] See www.jnrbm.com .

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 appendix shows, first, that both firms choosing a uniformly random sequence is an equilibrium. In addition, it will prove that there is no equilibrium in which either firm chooses a sequence as a pure action. Finally, we show that though it is possible that there are equilibria other than the uniformly random one, calculating the best response of each player 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 player $i \in \{1, 2\}$ 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 player 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 players are:

$$\begin{cases} U_1 = \frac{1}{s} \sum_{i=1}^{s} \left\{ V \left[\sum_{l=1}^{s} \left(P^1_{li} \sum_{k=i}^{s} P^2_{lk} \right) \right] - \frac{1}{2} V \left[\sum_{l=1}^{s} P^1_{li} P^2_{li} \right] - c \left[\sum_{l=1}^{s} \left(\sum_{k=i}^{s} P^1_{lk} \right) \left(\sum_{k=i}^{s} P^2_{lk} \right) \right] \right\} \\ \\ U_2 = \frac{1}{s} \sum_{i=1}^{s} \left\{ V \left[\sum_{l=1}^{s} \left(P^2_{li} \sum_{k=i}^{s} P^1_{lk} \right) \right] - \frac{1}{2} V \left[\sum_{l=1}^{s} P^2_{li} P^1_{li} \right] - c \left[\sum_{l=1}^{s} \left(\sum_{k=i}^{s} P^2_{lk} \right) \left(\sum_{k=i}^{s} P^1_{lk} \right) \right] \right\} \end{cases} \tag{38}$$

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 is a Nash equilibrium.*

Proof. Recall that for any $i \in \{1,2\}$, $\sum_{j=1}^{s} P^i_{jk} \quad \forall k$ and that $\sum_{k=1}^{s} P^i_{jk} \quad \forall j$. 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 eq. (3), 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 chooses a pure strategy. Without loss of generality, let it be firm 1. If firm 2 then uses the "preemption" strategy, their payoffs would be:

$$\begin{cases} U_1^P = \frac{1}{s}V - \frac{s^2-s+2}{2s}c \\ U_2^P = \frac{s-1}{s}V - \frac{s^2-s+2}{2s}c \end{cases} \tag{39}$$

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

$$\begin{cases} U_1^{UR} = \frac{1}{2}V - \frac{2s^2+3s+1}{6s}c \\ U_2^{UR} = \frac{1}{2}V - \frac{2s^2+3s+1}{6s}c \end{cases} \tag{40}$$

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 U_2^P. However, firm 1 can always guarantee itself a payment of U_1^{UR}, since this payment is achievable by deviating to the uniformly random strategy, regardless of the strategy chosen by firm 2. Therefore if in equilibrium firm 1 chooses a pure strategy, then the joint payoffs of the firm would have to be at least $U_1^{UR}+U_2^P$. 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 would be made on average after $\frac{s}{4}$ periods. Therefore the joint profits of the firm would have to satisfy this upper bound:

$$U_1^{UR} + U_2^P \leq V - 2c\frac{s}{4}, \tag{41}$$

which is equivalent to the following condition:

$$\frac{V}{c} \leq \frac{2s^2+7}{3s-6}. \tag{42}$$

However, since we assumed that for any feasible strategies by the two firms, the expected payoffs are positive, we require that $U_1^P \geq 0$, which is equivalent to:

$$\frac{V}{c} \geq \frac{s^2-s+2}{2}. \tag{43}$$

Combined with the previous condition this implies that:

$$s\,(s-3)\,(3s-4) \leq 26. \tag{44}$$

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:

$$U_1 = \frac{1}{s}\sum_{i=1}^{s}\sum_{j=1}^{s}P_{ij}^{1}\left[V\left(\sum_{k=j}^{s}P_{ik}^{2}\right)-\frac{1}{2}V\left(P_{ij}^{2}\right)-c\sum_{l=1}^{j}\left(\sum_{k=l}^{s}P_{ik}^{2}\right)\right] \tag{45}$$

And so, 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:

$$\begin{aligned}
&\max \tfrac{1}{s}\sum_{i=1}^{s}\sum_{j=1}^{s}P_{ij}^{1}\left[V\left(\sum_{k=j}^{s}P_{ik}^{2}\right)-\tfrac{1}{2}V\left(P_{ij}^{2}\right)-c\sum_{l=1}^{j}\left(\sum_{k=l}^{s}P_{ik}^{2}\right)\right]\\
&\textbf{s.t.}\\
&\sum_{j=1}^{s}P_{ij}^{1}=1 \qquad \forall i\\
&\sum_{i=1}^{s}P_{ij}^{1}=1 \qquad \forall j\\
&P_{ij}^{1}\geq 0 \qquad \forall i,j
\end{aligned} \tag{46}$$

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

5.2 Proof of Proposition 3 - Disclosure with a Discount Factor

A few mathematical manipulations on the recursive form expressions, which appear in equation (11), will derive the following normal form payoff functions:

$$U_1\left(n_1,n_2\right)=V\frac{\left(1-\delta^{s-n_1}\right)(1+\delta)-(1-\delta)\delta^{s-n_1}(n_1-n_2)}{2(1-\delta)^2(s-n_1)(s-n_2)}+V\frac{1}{(1-\delta)(s-n_1)}- \tag{47}$$

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

and

$$U_2\left(n_1,n_2\right)=U_1\left(n_1,n_2\right)-V\tfrac{(n_1-n_2)\left(1-\delta^{s-n_1}\right)}{(1-\delta)(s-n_1)(s-n_2)}. \tag{48}$$

Note that under the assumption of the proposition, $s-n_1\geq 2$ and $0>\delta>0$. We will 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 9. The first derivative with respect to n_2 is

$$\frac{\partial[U_1\left(n_1,n_2\right)+U_2\left(n_1,n_2\right)]}{\partial n_2}=\left[(1-\delta)V+2\delta c\right]\tfrac{\left(1+\delta^{s-n_1}\right)(s-n_1)(1-\delta)-\left(1-\delta^{s-n_1}\right)(1+\delta)}{(s-n_2)^2(s-n_1)(1-\delta)^3}. \tag{49}$$

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

$$\Gamma\equiv\left(1+\delta^{s-n_1}\right)\left(s-n_1\right)\left(1-\delta\right)-\left(1-\delta^{s-n_1}\right)\left(1+\delta\right)\geq 0. \tag{50}$$

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

$$\frac{\partial\Gamma}{\partial\delta}=\left(s-n_1+1\right)\left[\left(s-n_1\right)\delta^{s-n_1-1}\left(1-\delta\right)-\left(1-\delta^{s-n_1}\right)\right]. \tag{51}$$

If we plug in the condition that $\Gamma=0$ we get:

$$\left.\frac{\partial\Gamma}{\partial\delta}\right|_{\Gamma=0}=-\left(s-n_1+1\right)\left(s-n_1\right)\frac{1-\delta}{1+\delta}\left(1-\delta^{s-n_1}\right)<0, \tag{52}$$

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 5. The first derivative with respect to n_1 is

$$\frac{\partial[U_1(n_1,n_2)+U_2(n_1,n_2)]}{\partial n_1}= \\ =[(1-\delta)V+2\delta c]\,\frac{[(s-n_2+1)(1-\delta)-2]\left(1-\delta^{s-n_1}\right)-\log[\delta]\delta^{s-n_1}(s-n_1)[1+\delta-(n_1-n_2)(1-\delta)]}{(s-n_1)^2(s-n_2)(1-\delta)^3}. \quad (53)$$

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

$$\zeta\equiv[(s-n_2+1)(1-\delta)-2]\left(1-\delta^{s-n_1}\right)- \\ -\log[\delta]\,\delta^{s-n_1}(s-n_1)[1+\delta-(n_1-n_2)(1-\delta)]\geq 0. \quad (54)$$

This proof, unfortunately, requires many intermediate steps, however what we eventually would 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 $\psi\equiv\lim\limits_{n_2\to n_1}\zeta$. We therefore have:

$$\psi=[(s-n_1+1)(1-\delta)-2]\left(1-\delta^{s-n_1}\right)-\log[\delta]\,\delta^{s-n_1}(s-n_1)(1+\delta) \quad (55)$$

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 will 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 for any $0<\delta<1$. The partial derivative of ψ with respect to δ is:

$$\frac{\partial\psi}{\partial\delta}=-(s-n_1+1)- \\ -\delta^{s-n_1-1}\left\{(s-n_1)^2(1-\delta)-(s-n_1+1)\delta+(s-n_1)[(s-n_1)(1+\delta)+\delta]\log[\delta]\right\}. \quad (56)$$

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

$$\left.\frac{\partial \psi}{\partial \delta}\right|_{\psi=0} = \frac{(s-n_1)\left[\left(1+\delta^2\right)\left(1-\delta^{s-n_1}\right)-(s-n_1)\left(1-\delta^2\right)\right]}{\delta\,(1+\delta)}. \tag{57}$$

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

$$\begin{aligned}\left.\frac{\partial^2 \psi}{\partial \delta^2}\right|_{\psi=0} = & \frac{s-n_1}{\delta^2\,(1+\delta)^2}\{\left[s-n_1-1+\delta\,(s-n_1-2)+\delta^2\right]\left(1-\delta^{s-n_1}\right)+ \\ & +(s-n_1)\,(1+\delta)\left(\delta-\delta^{s-n_1+2}\right)\}>0,\end{aligned} \tag{58}$$

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:

$$\lambda \equiv \frac{\partial \zeta}{\partial n_2} = (1-\delta)\left[\delta^{s-n_1}-\log[\delta]\,\delta^{s-n_1}\,(s-n_1)-1\right]. \tag{59}$$

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:

$$\frac{\partial \lambda}{\partial \delta} = 1-\delta^{s-n_1}-\log[\delta]\,(s-n_1)\,\delta^{s-n_1-1}\left[(s-n_1)\,(1-\delta)-\delta\right]. \tag{60}$$

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

$$1-\delta^{s-n_1} = -\log[\delta]\,\delta^{s-n_1}\,(s-n_1)\,, \tag{61}$$

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

$$\left.\frac{\partial \lambda}{\partial \delta}\right|_{\lambda=0} = -\log[\delta]\,(s-n_1)^2\,\delta^{s-n_1-1}\,(1-\delta) > 0. \tag{62}$$

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 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.

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.

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.

Myerson, Roger B. and Mark A. Satterthwaite, 1983, “Efficient Mechanisms for Bilateral Trading”, *Journal of Economic Theory* 29, pp. 265–281.

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*.