Alice Guerra[1] – Francesco Parisi[2] – Andres Sawicki[3]

Rewarding Failure

ABSTRACT: Research is a messy enterprise, fraught with dead ends and missed opportunities. When research is successful, intellectual property law offers market-driven rewards. When research is unsuccessful, however, the investment is treated as a waste—an unlucky gamble in the horse race of innovation. But failed research investments are not valueless. Instead, they generate useful information that increases the odds of success on subsequent research efforts. From a social perspective, then, firms should disclose information about research failures so that other firms can avoid duplicative, fruitless lines of inquiry. Unlike with successful research investments, however, firms have no easy opportunity to capture the value of information generated by their unsuccessful research. Indeed, the law offers no coherent approach for harvesting the social value of the information generated by failed research efforts. As a result, discoveries of what "does not work" are not optimally disclosed. In this paper, we develop a simple sequential-stage model to identify the shortcomings of this one-sided reward system. We explore possible legal and institutional solutions to the problem, including contractual arrangements, trade secrets, research subsidies, and patent-like "propertization" mechanisms. Our results shed light on the potential and limits of these mechanisms for aligning private and social incentives for research.

Keywords: Value of failure, trade secrets, negative patents, research subsidies

JEL Codes: C72, D72

---

[1] Associate Professor of Economics, University of Bologna, Department of Economics. E-mail: alice.guerra3@unibo.it.

[2] Oppenheimer Wolff & Donnelly Professor of Law, University of Minnesota, School of Law and Professor of Economics, University of Bologna, Department of Sociology and Business Law. Visiting Professor, University of Miami, School of Law (2023-24). E-mail: parisi@umn.edu.

[3] Professor of Law & Director, Business of Innovation, Law, & Technology Concentration, University of Miami, School of Law. E-mail: a.sawicki@umiami.edu.

Acknowledgments: The authors would like to thank Keith Hylton, Louis Kaplow, Saul Levmore, and the other participants to the 2024 Annual Meeting of the American Law & Economics Association, as well as Bernard Chao, Joseph Fishman, Lisa Ouellette, and the other participants of the 2024 Intellectual Property Scholars Conference for their valuable comments and suggestions. We are also grateful to Erik Birkeneder, Gen. Vincenzo Camporini, and Arturo Parisi for suggestions and examples drawn from their respective areas of expertise. Samuel Makikalli at the University of Minnesota Law School provided valuable research assistance.

Table of Contents

"*It's fine to celebrate success but it is more important to heed the lessons of failure.*"

Bill Gates

1. Introduction

Research can be thought of as the pursuit of knowledge or truth about something. The modern economy rewards successful research in many ways. One of the most important rewards is intellectual property. Through patents and copyrights, the legal system grants protection for economic utilization of valuable information. By granting discoverers of valuable information exclusive rights in the exploitation of such information, the legal system allows parties to internalize the benefit of their research, thus creating efficient incentives for conducting research.

Not all research, however, leads to successful discoveries. Sometimes, research shows that a seemingly promising opportunity is not viable. Even though researchers do not reach an exploitable discovery in these scenarios, they nonetheless generate valuable information for other researchers, helping them avoid unproductive paths and focus on more promising alternatives. Contrary to the case for successful research, the legal system does not have effective instruments for firms to internalize these informational externalities. In this paper, we develop a model to identify the shortcomings of the existing, one-sided reward system and investigate potential solutions to the problem.

This paper is structured as follows. In Section 2, we carry out a brief review of the literature and develop a simple sequential-stage model to consider the conditions under which negative information and research failures generate a social benefit that is not captured by the private incentives of a researcher. This allows us to identify the conditions under which we are likely to observe a discrepancy between private and social incentives for research. We start by distinguishing two kinds of informational value generated by unsuccessful research: (1) the value of avoiding the same dead ends; and (2) the value of learning. Information regarding dead ends benefits other firms by reducing the total number of paths they might pursue. This information increases the expected value of subsequent research by allowing firms to focus on fewer paths of research, increasing the likelihood that they may reach a successful result in the future. We will refer to this informational benefit as the "value of failure." Information from past research can also benefit future research in other ways. In particular, findings from unsuccessful research can percolate to other lines of inquiry, and possibly contribute to the general knowledge of an industry or scientific field. We will refer to these broader informational benefits as the "value of learning." In the subsequent analysis, we develop a model that focuses on the vastly underexplored value of failure problems, showing how the value of failure creates a discrepancy between private and social incentives for research. The identified wedge in research incentives adds to the learning externalities discussed in the existing literature. The distinctive contribution of our model lies in capturing, in a simplified manner, the value of failure for innovation. The value of failure yields two important distortions. The first follows from the fact that the information generated by one firm's failure increases the expected value of subsequent research in the same field. This creates strategic incentives for researchers to delay investment so that they may take advantage of the negative information generated by prior research. This leads to socially sub-optimal *timing* in research investments. The second arises because firms do not fully internalize the value of the information they create. This leads to socially sub-optimal *levels of investment* in research.[4]

In Section 3, we consider alternative policy instruments and evaluate them analytically to assess their effectiveness as possible solutions to the value of failure problem. We start by considering firms' private incentives when the information generated by research failures is neither protected nor rewarded. These results are used as a benchmark to evaluate the effect of alternative legal regimes analytically. We consider the merits and limitations of contracts, trade secrets, patents, research subsidies, and patent-like propertization solutions to the information externality problem that stems from failed research.

In Section 4, we explore the respective merits and practical limits of alternative protective and rewarding instruments to incentivize timely and efficient investments in and disclosure of research. Under the intellectual property protection for negative information that is currently available, we observe a tension between the prospective and retrospective effects of alternative protection and reward mechanisms. For example, through trade secrets, the discoverer of negative information obtains "protection" of his information, thereby preventing possible competitors from free riding. This mitigates the incentive to delay investment in research by eliminating the unwanted positive externality that would otherwise favor

[4] Although the timing problem may be more significant in conjunction with information externalities arising from the finding of dead-ends (value of failure), suboptimal investment problems may also be significant in the presence of information externalities that benefit other lines of research (value of learning).

competitors. However, the protection obtained through trade secrets does not provide other firms with access to the failed research information; this results in duplicative research investments within the industry. We contemplate how alternative systems, such as tailored subsidies for public disclosure or propertization mechanisms, might foster the spread of negative information through an industry. We consider how a hypothetical market for negative intellectual property could function and explore alternative instruments to yield an optimal internalization of positive and negative research externalities in equilibrium. Our analysis points to the implementation problems that would need to be addressed to boost the effectiveness of alternative policy instruments.

Section 5 concludes with some final considerations and ideas for future research.

2. The Value of Failure

2.1 Literature Review

Innovation is a learning process where researchers rely on past results to inform their investment decisions. The vast theoretical literature on research and development ("R&D") and technological innovation as well as recent empirical research on spillovers in R&D processes have explored how firms build on their own and each other's failures and successes. A brief review of this literature will aid our understanding of the academic framing of the innovation process.

Very early in this literature, Tarde (1903) identified the initial steps of the innovation-decision process as "first knowledge" and "forming an attitude." Several studies have dealt with the first point, seeking to understand how firms learn from prior positive or negative experiences. Moreover, prior studies highlight how both positive *and* negative experiences hold important lessons for innovators. To elaborate on this point, it seems perfectly reasonable that, to learn the secrets of success, it would be efficient to study successful people and organizations. Thus, for example, authors of popular business books describe successful companies that have focused on one key product and argue that such focus caused the companies' success. However, organizational behavior research carried out by Denrell (2003) suggests that studying successes without also looking at failures tends to create a misleading, if not entirely wrong, picture of what it takes to succeed. Indeed, Denrell (2003) highlights Kodak and Xerox as examples of mono-focused companies that performed poorly over time.

While scholars generally acknowledge the potential learning benefits of failures (Lant and Montgomery 1987; Sitktin 1992; Marcus and Nichols 1999), the concept of failure has received limited attention in the law and economics literature, probably due to the shared wisdom that firms tend to privilege learning from success and are incapable of finding insights from failures (Denrell 2003; Baumard and Starbuck 2005). In focusing excessively on successes, scholars overlook the important fact that failing companies do many of the same things as companies that succeed.

Drawing on these principles, the literature on research and learning is vast and methodologically diverse. Several contributions have pointed out that research organizations can learn by vicariously observing other players in the same field, including their failures (Levitt and March 1988; Chuang and Baum 2003; Baum and Dahlin 2007; Terlaak and Gong 2008; Francis and Zheng 2010). More recent papers have refined these insights (Haunschild and Sullivan 2002; Lapré and Tsikriktsis 2006; Baum and Dahlin 2007).

From a broader perspective, the outcome of research cannot be predicted. Indeed, one key function of the patent system is to encourage the pursuit of highly uncertain research (Merges 1992). Moreover, most research might be classified as a failure in the sense that the immediate results are not useful or exploitable. A problem arises in that organizations' and researchers' failures are not always highly visible. The knowledge and publication of failures might help organizations to indirectly gain learning benefits without directly experiencing the reputational and operational costs related to failure. Thus, Seymore (2012) proposes the creation of a "null patent" to facilitate dissemination of failures through the Patent Office.

Our project closely aligns with Levmore (2013), who argues that the growing significance of knowledge in the modern economy necessitates a reevaluation of the traditional emphasis in patent law on new inventions and technological innovations. Several other papers have studied the optimal scope and efficiency of patents and trade secrets when research generates knowledge with potential positive spillovers on other researchers (Denicolò and Franzoni, 2003 and 2012; Franzoni, 2016; and Franzoni and Kaushik, 2016).

The strand of literature on open innovation is also relevant to our research. Over the past three decades, open-access models have emerged as stable models for provisioning critical segments of information. The success of these practices has led prominent scholars to question the need for intellectual property instruments to foster optimal investments in research—these scholars point to the expanding frontiers of open source innovation as the emerging example of knowledge flow using a mix of pecuniary and non-pecuniary incentives (see, e.g., Benkler 2002, 2012, 2017). Open innovation scholarship criticizes the widespread reliance on intellectual property instruments for rewarding innovation, and their criticism would, to some extent, apply to our focus on instruments that reward research failures. We do not plan to engage in this debate. We hope, however, that our analysis may lead open innovation scholars to investigate the extent to which the incentives for open innovation considered in their models also account for the benefits of research failures and incentives to openly disclose negative information.[5]

2.2 A Simple Sequential-Stage Model: Basic Setup

In this section, we build a sequential-stage patent race model to analyze how the information spillovers of research failures impact decisions regarding both timing and levels of research investments. Our model describes the innovation-decision process in a competitive market, where firms choose to pursue one out of several research paths.

The distinctive contribution of our model lies in capturing, in a simplified manner, the value of failure for innovation. In the following, we present the basic framework and assumptions of the model. In Section 2.3, we derive the socially optimal timing and levels of research investments. In Section 3, we use these results to show that, in competitive scenarios, the unsuccessful research of early-movers offers later entrants an opportunity to benefit from past research, which leads to two possible distortions: (i) firms might delay entry into research, and (ii) research investments may be suboptimal. We then consider the effect of alternative legal regimes in mitigating these distortions and compare their effects analytically.

Consider a sequential-move game in which N firms are involved in an innovation race—e.g., for the development of a new vaccine. The firm that develops the vaccine first will have patent and FDA exclusivity

[5] For a recent review of this literature, see Bertello, De Bernardi, and Ricciardi (2023). For a methodological presentation of this research field, see Chesbrough and Bogers (2014); Holgersson and Granstrand (2021); and The Economist Group (2022).

over it, enabling it to receive a financial return from its discovery. A firm enters the race as a first mover. Those engaging in research in later periods can observe the outcomes of earlier research.

Each Firm $i \in [1, N]$ invests e_i in research. Not all research is successful in the sense of producing an exploitable invention; research fails when it reveals that something will not work. Let us refer to the two possible conclusive research outcomes as "positive" or "negative."

(i) *Positive Result*. A positive research result allows the firm to capture the applicative value of its discovery (e.g., a pharmaceutical firm obtains FDA approval after its clinical trials demonstrate that a vaccine candidate is safe and effective). We shall broadly refer to the applicative value of positive results as the "value of success," V.

(ii) *Negative Result*. A negative research result provides evidence that the chosen research path is a dead end (e.g., a pharmaceutical firm conducts clinical trials that demonstrate that a vaccine candidate is toxic or ineffective). This negative result has informational value— other firms can avoid duplicative research efforts because they now know that they should not allocate resources this research path (e.g., other pharmaceutical firms will not conduct clinical trials on the same vaccine candidate). We shall broadly refer to the value of negative results as the "value of failure," F. For the purpose of our analytical presentation, the value F does not include the value of the information that might benefit other lines of research and/or contribute to the general knowledge of an industry (i.e., the "value of learning" generated by the firm's research).

In addition, research may be inconclusive in the sense that the results leave unclear whether a research path can produce an exploitable invention (e.g., a pharmaceutical firm conducts clinical trials, but the sample size turns out to be too small to produce statistically significant results). The probability of reaching a conclusive result, $p(e_i)$, is an increasing function of a firm's own research investment—i.e., $p' > 0$; with probability $1 - p(e_i)$, Firm i's research is inconclusive. For our analysis, we assume that when a firm is unable to reach any positive or negative result, no applicative or informational value is generated, and the expenditure e is largely a sunk waste of resources.

For ease of presentation and without loss of generality, let us consider a case where (1) a successful result can be reached following $M \geq N > 1$ possible lines of research; (2) only one of the M lines of research can lead to a successful discovery;[6] and (3) firms do not reenter the patent race after failure.[7] A real-world example of this scenario would be a field in which start-up firms are founded with specialized expertise and attempt to introduce a novel product to market; if a start-up's effort fails, it runs out of funding and exits.

The first mover's research investment can lead to either a positive result (hereinafter, "success") with probability $p(e_1)\frac{1}{M}$ or a negative result (hereinafter, "failure") with probability $p(e_1)\frac{M-1}{M}$. In the former case, the first mover captures the value of his success, *V.* In the latter case, the negative result only reveals that the already-pursued research path led to a dead end, and that it is not worth pursuing again in the future.

When considering a firm's decision to invest in research, it is important to consider two countervailing effects of a firm's timing:

(i) *Lost opportunity effect*. Firms that delay their investment face the risk that earlier movers may reach a successful result before them, thus losing the opportunity to make the sought-after discovery. This effect may induce firms to expedite their investment.

(ii) *Informational benefit effect*. The information generated by past research accrues over time during the patent race. Firms that delay research can obtain more information and thereby increase their odds of success by avoiding known dead-ends. This effect may induce firms to delay their investment.

In the following, we will proceed to illustrate these two effects analytically.

[6] The same qualitative results can be reached assuming that a larger subset, *m*, of the *M* lines of research can lead to success, where $M > m > 1$.

[7] Our qualitative results would hold even in the more general case where firms have a positive probability to reenter the race, but the wedge between private and social optima will be narrower in that case.

2.2.1 Lost Opportunity Effect

As a first mover, Firm 1 engages in research before all other firms, and will thus always have the full opportunity to pursue research, $\theta_1 = 1$. Subsequent firms' opportunity to engage in research hinges upon the probability that no other firm reached a successful result before them. Firm 2 will thus have a reduced opportunity to pursue research, given by the probability that Firm 1 reaches a successful result:

$$\theta_2 = 1 - \frac{1}{M} = \frac{M-1}{M} \qquad 2.1$$

Firm 3's opportunity to engage in research will be further reduced by the probability that Firm 2 reaches a successful result:

$$\theta_3 = \left(1 - \frac{1}{M}\right)\left(1 - \frac{1}{M-1}\right) = \frac{M-2}{M} \qquad 2.2$$

And so on for all subsequent firms:

$$\theta_{N-1} = \left(1 - \frac{1}{M}\right)\left(1 - \frac{1}{M-1}\right) \dots \left(1 - \frac{1}{M-N+3}\right) = \frac{M-N+2}{M} \qquad 2.3$$

$$\theta_N = \left(1 - \frac{1}{M}\right)\left(1 - \frac{1}{M-1}\right) \dots \left(1 - \frac{1}{M-N+2}\right) = \frac{M-N+1}{M} \qquad 2.4$$

Hence, for a generic Firm $i \in [1, N]$:

$$\theta_i = \frac{M-i+1}{M} \qquad 2.5$$

Opportunities to engage in research thus decrease as the race progresses: [8]

$$\theta_1 > \theta_2 > \cdots > \theta_{N-1} > \theta_N \qquad 2.6$$

The "lost opportunity effect" may therefore induce firms to expedite their investment.

---

[8] Also note that, as the race progresses, waiting imposes incrementally larger losses of opportunity: $(\theta_1 - \theta_2) < (\theta_2 - \theta_3) < \cdots < (\theta_{N-2} - \theta_{N-1}) > (\theta_{N-1} - \theta_N)$.

2.2.2 Informational Benefit Effect

If publicly known, the failures of the early entrants in a patent race create an informational benefit to later researchers. Firms that delay research can increase their odds of success. To illustrate, consider the case of two competing firms. Firm 1's research generates a negative result. Firm 2 benefits from the failure of Firm 1 by discarding the research path that led to that failure, thereby increasing its own probability of success:

$$\frac{1}{M} < \frac{1}{M-1} \qquad \forall M > 1 \qquad 2.7$$

Firm 2's odds of success are greater than Firm 1's odds of success because Firm 2 is choosing from one of $M-1$ paths rather than from one of M paths; this difference between choosing from one of $M-1$ paths rather than from one of M paths can be understood as the informational benefit that Firm 1 confers upon Firm 2. More generally, Firm $i+1$ benefits from the failure of Firm i because it has one fewer path to choose from.[9] Let us accordingly denote as f_i the informational benefit created by Firm i's failure that is received by all the following firms, which in our case can be quantified as follows:[10]

$$f_1 = \frac{1}{M-1} - \frac{1}{M} \qquad 2.8$$

$$f_2 = \frac{1}{M-2} - \frac{1}{M-1} \qquad 2.9$$

$$\ldots$$

[9] The reader might note that we are treating the situation as one of Knightian uncertainty. Future work may consider the more complex case in which firms' estimate of the probability of success of any given path are "updated" with the information generated by prior failures. Through this extension, the value of failure would be augmented, increasing the benefits of waiting in undertaking research.

[10] The failure of the Nth firm creates no informational benefit for other firms because, in the future, no other firms will subsequently invest in that line of research.

$$f_{N-1} = \frac{1}{M-N+1} - \frac{1}{M-N+2} \quad 2.10$$

$$f_N = 0 \quad 2.11$$

The informational benefits generated by the failures of early movers percolate down to all subsequent firms, while the failures of later movers do not have backward effects on earlier movers (e.g., the informational benefit f_1 carries over to all remaining $N-1$ firms, whereas the informational benefit f_{N-1} only benefits the last entrant, Firm *N*). A firm's expectations regarding the outcome of its research should therefore be influenced by both (1) the probability of reaching a conclusive result as a function of its own effort and (2) the information received from the failures of prior firms.

$$V_2^f(e_2) = p(e_2)\left[\frac{1}{M} + f_1\right]V = p(e_2)\left[\frac{1}{M-1}\right]V \quad 2.12$$

$$V_3^f(e_3) = p(e_3)\left[\frac{1}{M} + f_1 + f_2\right]V = p(e_3)\left[\frac{1}{M-2}\right]V \quad 2.13$$

...

$$V_{N-1}^f(e_{N-1}) = p(e_{N-1})\left[\frac{1}{M} + f_1 + f_2 + f_3 \dots + f_{N-2}\right]V = p(e_{N-1})\left[\frac{1}{M-N+2}\right]V \quad 2.14$$

$$V_N^f(e_N) = p(e_N)\left[\frac{1}{M} + f_1 + f_2 + f_3 \dots + f_{N-1}\right]V = p(e_N)\left[\frac{1}{M-N+1}\right]V \quad 2.15$$

As shown in Eq. 2.12 – 2.15, a generic Firm $i \in [1, N]$ receives an informational benefit from the failures of previous firms, i.e., $\sum_1^{i-1} f_i$, and at the same time creates a flow of informational benefit for all subsequent firms, from $i+1$ to N. To capture this latter effect, let us define F_i as the *cumulative* informational benefit generated by Firm i to all subsequent firms' chances of reaching a successful discovery:

$$F_i = (N - i)f_i, \qquad \forall i \in [1, N] \tag{2.16}$$

From the above, it should be noted that the *marginal* informational benefits and resulting increments in success probabilities increase as the patent race progresses, i.e.,

$$f_1 < f_2 < \cdots < f_{N-1} \tag{2.17}$$

Since early failures benefit all subsequent firms in the race, the *cumulative* informational benefit created by early movers is higher than that of late movers:

$$1 > F_1 > F_2 > \cdots > F_{N-1} > 0 \tag{2.18}$$

In our setting with sequential entries, the nature of the race intimates that firms that receive the most information are those that will generate less informational benefit to others, and vice-versa. At the beginning of the race, early-movers receive little information (Firm 1 receives no information; indeed, $V_1^f(e_1)$ is omitted from Eq. 2.12 – 2.15 because f_0 does not exist); but generate a large flow of informational benefit for others (Firm 1's negative results benefit all subsequent firms in the race, as shown in Eq. 2.18). At the end-tail of the race, late-movers receive large amounts of information (Firm *N* receives information from all prior firms, as shown in Eq. 2.15) but generate a small flow of informational benefit on others (indeed, at the very end of the race, Firm *N* does not benefit any subsequent firm in the race, as shown in Eq. 2.18).

The "informational benefit effect" may thus induce firms to delay their investments in research.

2.3 Socially Optimal Incentives

With a certain probability, any given firm will fail to reach its desired research objective. Disappointing as this may be for the failing firms, reaching a negative result in research often serves as a steppingstone to others' successes, thereby contributing to the overall progress of research.

Socially optimal research investments are those that would foster the progress of research, maximizing the *aggregate* expected returns for all participating firms. In such an ideal world, firms would not only focus on the expected value of their success, *V*, but would also consider the informational value generated by their failures, *F*, net of research costs:

$$R_1^S\,(e_1) = p(e_1)\left[\frac{1}{M} + \frac{M-1}{M}F_1\right]V - e_1 \qquad 2.19$$

$$R_2^S\,(e_2) = p(e_2)\,\theta_2\left[\frac{1}{M} + f_1 + \frac{M-2}{M-1}F_2\right]V - e_2 \qquad 2.20$$

$$\cdots$$

$$R_{N-1}^S\,(e_{N-1}) = p(e_{N-1})\;\theta_{N-1}\left[\frac{1}{M} + \sum_{i=1}^{N-2} f_i + \frac{M-N+1}{M-N+2}F_{N-1}\right]V - e_{N-1} \qquad 2.21$$

$$R_N^S\,(e_N) = p(e_N)\;\theta_N\left[\frac{1}{M} + \sum_{i=1}^{N-1} f_i\right]V - e_N = p(e_N)\frac{1}{M}V - e_i \qquad 2.22$$

2.3.1. Socially Optimal Timing

In a social optimum, firms will enter a research race if the expected return from research is positive. Firms will have no incentive to strategically delay their entry. This can be verified by comparing the firms' respective payoffs and observing that the social return of the first-moving firm is higher than that of the second-moving firm, and so on for all subsequent firms, i.e., $R_1^{**} > R_2^{**} > \cdots > R_N^{**}$. The first-moving firms generate larger social value because (i) they have the first opportunity to reach a successful result (i.e., no lost opportunity effect); and (ii), in case of a negative result, they generate the greatest flow of informational benefit onto the industry, which percolates down to all the subsequent firms (i.e., greatest informational benefit effect).

2.3.2. Socially Optimal Investment

In our sequential race, a generic Firm $i \in [2, N]$ gets to enter the race and invests in research only if no other firm previously reached a successful result. That is, Firm $i \in [2, N]$ enters the race only if it has not lost the opportunity to make successful research; formally, given $\theta_i = 1$. At that time, firms choose their investment knowing that past firms failed. Formally, Firm *i*'s social maximization problem becomes:

$$\max_{e_i} R_i^S (e_i) \mid (\theta_i = 1) \;\; i \in [1, N] \tag{2.23}$$

Given that knowledge and conditional on that probability $\theta_i = 1$, when firms enter the research race, social optimality requires them to invest in research knowing that past firms have failed. Additionally, entering firms must include the value of their possible failure for subsequent firms in their optimization problem.

The firms' socially optimal levels of effort would thus be given by the following equations:

$$e_1^{**}: p'(e_1) \left[\frac{1}{M} + \frac{M-1}{M} F_1 \right] V = 1 \tag{2.24}$$

$$e_2^{**}: p'(e_2) \left[\frac{1}{M} + f_1 + \frac{M-2}{M-1} F_2 \right] V = 1 \tag{2.25}$$

....

$$e_{N-1}^{**} = p'(e_{N-1}) \left[\frac{1}{M} + \sum_{i=1}^{N-2} f_i + \frac{M-N+1}{M-N+2} F_{N-1} \right] V = 1 \tag{2.26}$$

$$e_N^{**} = p'(e_N) \left[\frac{1}{M} + \sum_{i=1}^{N-1} f_i \right] V = 1 \tag{2.27}$$

Later firms have a reduced opportunity to enter the race, but, when those opportunities arise, they have increased probabilities of success. As a result, the socially optimal investments increase as the patent race progresses, $e_N^{**} > \cdots > e_2^{**} > e_1^{**}$.

In the following sections, we will consider the interdependent role of "protection" and "reward" instruments on firms' research incentives. We will separately analyze the firms' (i) timing incentives and (ii) investment incentives for research under these alternative instruments, comparing them to the benchmark case presented in Section 3.1 where negative information is neither protected nor rewarded.

3. Incentivizing Success and Failure: Protection vs. Rewards

Negative results in research often serve as a steppingstone to others' successes, contributing to the overall progress of research. However, firms would largely ignore the informational benefits that their negative research results may confer on others, at least in the absence of legal or institutional mechanisms that encourage them to take those benefits into account.

This is most clear in cases of failing firms that do not plan to reenter the patent race, because the only way a firm can reap the benefits of its own negative information is by reengaging in the patent race. The lessons drawn from past failures only benefit failing firms that plan to remain active in that field of research in the future. Even in such contexts, the firms' cost-benefit calculations would ignore the fact that their failed research generates valuable information to other firms. Each firm's investment would be solely incentivized by the hope to capture the value V and would not account for the informational value of F to the industry. From an economic perspective, then, a public good problem arises because third parties can benefit from negative information without contributing to the cost of generating it—freeriding would take place and research would advance at both *suboptimal timing* and *investment* levels.

Intellectual property law addresses this issue in the context of positive results by granting property rights to the party that invested in successful research. Intellectual property rights, such as patents and copyrights, simultaneously “protect” and “reward” successful discoveries. They “protect” discoveries by allowing firms who carried out research to exclude third parties from using their information. They “reward” discoveries because the owner can obtain a financial reward for successful research by selling the invention at supra-competitive prices or by licensing or selling the patent or copyright to others. This intellectual property law protection and reward system enables those investing in research to internalize the positive externalities generated by their findings.

However, these mechanisms fall short in addressing the public good problem associated with research failures. In this context, the “public good” is the information that can enhance the chances of success in future research endeavors. When such “negative” information is not protected, firms’ incentives for research deviate from the social optimum because firms will be less willing to move first (“suboptimal *timing* problem”). Similarly, when the information from failed research efforts is not rewarded, firms’ expected returns fail to fully internalize the benefits created by the research, resulting in a departure from the socially optimal level of research investment (“suboptimal *investment* problem”).

To address these divergences between private and social incentives, we consider alternative protection and reward mechanisms for research failures. In the following, we explore four possible categories of instruments, as outlined in Table 1, that provide protection or reward to firms that failed to achieve successful discoveries.

<table>
<tr><td colspan="2"></td><td colspan="2">Reward</td></tr>
<tr><td colspan="2"></td><td>No
($\phi = 0$)</td><td>Yes
($\phi = 1$)</td></tr>
<tr><td rowspan="2">Protection</td><td>No
($\omega = 0$)</td><td>1. None (Benchmark)</td><td>3. Research Subsidies</td></tr>
<tr><td>Yes
($\omega = 1$)</td><td>2. Trade Secrets</td><td>4. Negative Patents</td></tr>
</table>

Table 1: *Protecting and Rewarding Failures*

Let us briefly outline our key findings. The "protection" of information granted through trade secrets allows firms to withhold information, thereby preventing competitors from exploiting firms' failed research. But trade secrets do not provide financial rewards for failed research, thereby preventing firms from capturing the financial value of their failed research. Our findings show that protecting negative information without a corresponding financial return mitigates the suboptimal timing problem because early firms and later entrants initiate their research with the same starting set of information, giving no advantage to first- or last-movers. Nevertheless, from a social point of view, trade secrets hinder information diffusion, potentially leading to recurrent research failures among competing firms. This prevents firms from optimally targeting their research efforts, hence perpetuating the suboptimal investment problem.

The shortcoming of trade secrets resides in the fact that giving firms a right to exclude others to negative information is not, on its own, sufficient for firms to internalize the value of their failures. Property over the firms' negative information entails multiple rights in the classic bundle of sticks sense: trade secrets grant protection and the right to exclude others, but in this context, they do not confer an effective right to transfer the information, which is needed in order to obtain rewards when the information is not itself exploitable (as in the case of negative information).

To achieve socially optimal research incentives, it is crucial to introduce instruments that reward firms for failures. This entails providing financial compensation for negative information, which allows firms to internalize the value of both successes and failures. Our findings show that rewarding negative information through research subsidies effectively resolves both the suboptimal timing and investment problems. Alternatively, a comparable outcome can be also attained by establishing tradable property rights over negative information. This only mitigates rather than entirely resolves the timing and investment problems, but it may be a more practically implementable solution compared to research subsidies.

This section is structured as follows: Section 3.1 outlines the benchmark case of failures without protection or rewards. Sections 3.2, 3.3, and 3.4 respectively explore the impact of protection (trade secrets), rewards (research subsidies), and a combination of both (negative patents) on addressing failures. Section 3.5 compares the effects of protective versus rewarding remedies on research outcomes. Next, Section 4 discusses implementation difficulties of the proposed instruments, along with market and legal solutions.

3.1 The Benchmark Case: Failures Without Protection or Rewards

Consider a sequential-move game in which *N* competitors engage in research and make their research investment decisions after observing the outcomes of earlier research. A firm enters a patent race as a first mover. Its research investment can lead to either a positive result (i.e., "success") with probability $p(e_1)\frac{1}{M}$, a negative result (i.e., "failure") with probability $p(e_1)\frac{M-1}{M}$, or no results with probability $1-p(e_1)$. In the absence of legal or institutional mechanisms that protect or reward the informational value of the firm's failure, the first mover's expected return from research is only driven by the expected value of its success, net of research costs.

As a first mover, Firm 1 has the first opportunity to pursue a successful result, and thus faces no lost opportunity effect from prior firms' attempts. For the same reason, Firm 1 cannot receive any informational benefit from the failure of its predecessors. For Firm 1, the expected return from the research endeavor is thus given by:

$$R_1(e_1) = p(e_1)\frac{1}{M}V - e_1 \qquad 3.01$$

If later entrants make investment decisions after observing the research outcomes of the earlier movers, the first firm's failure creates valuable information. If Firm 1 reached no results, no information is generated. Firm 2 will thus find itself in the same position as Firm 1, facing the same expected returns and investment incentives of Firm 1. If Firm 1 reached some results, Firm 2 enter the race considering Firm 1's research outcomes. With probability $\frac{M-1}{M}$, Firm 2 will consider Firm 1's negative results and capture the informational benefit of Firm 1's failure. However, with probability $\frac{1}{M}$, Firm 1 could achieve a successful result and Firm 2 would face a loss of opportunity. Firm 2's expected return would thus become:

$$R_2(e_2) = p(e_2)\,\theta_2\left[\frac{1}{M} + f_1\right]V - e_2 = p(e_2)\frac{1}{M}V - e_2 \qquad 3.02$$

As can be seen from Eq. 3.02, the two effects—informational benefit and loss of opportunity—offset one another, giving Firm 2 the same expected returns of Firm 1. Likewise, if Firm 1 or Firm 2 failed, Firm 3 will get the informational benefit from its predecessors but will also face the risk that one of them may succeed, foreclosing the opportunity to enter the race. Firm 3's research prospects will thus generate an expected return identical to that of Firm 1:

$$R_3(e_3) = p(e_3)\,\theta_3\left[\frac{1}{M} + f_1 + f_2\right]V - e_3 = p\,(e_3)\frac{1}{M}V - e_3 \qquad 3.03$$

Hence for the N^{th} firm:

$$R_N(e_N) = p(e_N)\theta_N\left[\frac{1}{M} + \sum_{i=1}^{N-1} f_i\right]V - e_N = p(e_N)\frac{1}{M}V - e_N \qquad 3.04$$

It should be noted that the full offsetting of informational benefits and the lost opportunity throughout the research race is driven by the fact that our model assumes all paths to have an equal likelihood of success. By relaxing that assumption and assuming that firms have reasonable estimates of whether a given path will be the successful one, then the lost opportunity effect would be larger than the informational benefit effect. On the other hand, the ability to produce reasonable estimates likely increases over time as firms can benefit from the information generated by prior failures, improving the accuracy of their estimates of whether a given path will be successful; this, in turn, could yield informational benefits that outweigh the loss of opportunity effect. We leave the analysis of these extensions to future work.

3.1.1 Timing Incentives

Comparing the firms' respective payoffs as described in Eq. 3.01 - 3.04, we observe that the private returns are identical for all firms, i.e., $R_1 = R_2 = \cdots = R_N$, with the last mover's expected returns coinciding with the socially optimal ones ($R_N^S = R_N$). Late movers can capture the informational benefit created by earlier firms, but their delayed entry will cause loss of opportunities. Both informational benefits and lost opportunities increase as firms delay their research, and in this benchmark case the two effects will have fully offsetting effects.

These results reveal that, in a regime where the negative information from failures is neither protected nor rewarded, there will be no advantage for firms to commence a research race early, yielding to a pace of research that falls below the socially optimal level.[11]

[11] Due to the positive information externalities that arise from failures in research, as the number of research paths and competing firms grows larger, it is possible that no firm would choose to be the first to enter a new field of research. If the number of research avenues and the number of competitors tends toward infinity, the expected profit for the 1st competitor would tend toward zero and the expected profit for the Nth competitor would tend toward the value of innovation: $\lim_{M,N\to\infty} R_1 = \lim_{M,N\to\infty} \frac{1}{M}V - e_N = 0$ and $\lim_{M,N\to\infty} R_N = \lim_{M,N\to\infty} \frac{1}{M-N+1}V - e_N = V$. This would lead to a stall

3.1.2 Investment Incentives

As shown in Section 3.1.1, when making their timing decisions, firms consider the alternative expected returns of early-movers and late-movers in the race, choosing the timing that maximizes their expected payoff. This decision is carried out based on their respective probabilities of success and failure, but with no knowledge of how those probabilities will play out in their future race. On the contrary, when making their investment decision, firms have full knowledge of how past events materialized. Firms get to invest in research only if no firm reached a successful result before them. At that time, firms will thus choose their investment based on the knowledge of past events. Similarly to Eq. 2.23, a generic Firm *i* will maximize the expected private return given that it has not lost the opportunity to make a successful discovery (i.e., that no other prior firms have reached the success); formally:

$$\max_{e_i} R_i\ (e_i) \mid (\theta_i = 1)\ \ i\ \in [1, N] \qquad 3.05$$

Hence, the privately optimal investments in research are given by the following equations:

$$e_1^*: p'(e_1)\frac{1}{M}V = 1 \qquad 3.06$$

$$e_2^*: p'(e_2)\frac{1}{M-1}V = 1 \qquad 3.07$$

….

$$e_{N-1}^*: p'(e_{N-1})\ \frac{1}{M-N+2}V = 1 \qquad 3.08$$

$$e_N^*: p'(e_N)\ \frac{1}{M-N+1}V = 1 \qquad 3.09$$

in the research process, even in the hypothetical case where the total expected benefits of research exceed the total expected costs.

Just as we observed in the socially optimal efforts in Section 2.2, in this benchmark case the information externalities created by early movers in the research race increase the probability of success for later entrants by reducing the number of research paths from which later entrants choose. As a result, unlike what we have seen when considering expected payoffs, the investment incentives of firms that have an opportunity to enter the race are not constant, but grow stronger as the patent race progresses, $e_N^* > \cdots > e_2^* > e_1^*$.

Comparing these results to the socially optimal values reveals that omitting the *cumulative* value of failure from firms' private objective functions creates a wedge between the privately and socially optimal incentives to invest in research. Specifically, in the absence of protection or reward mechanisms that apply to the information generated by a firm's failure to subsequent firms, firm research efforts will fall below the socially optimal levels, $e_i^{**} > e_i^*$ for all Firms $i \in [1, N)$, with only the last-mover converging toward optimality, $e_N^{**} = e_N^*$. This discrepancy stems from the fact that firms' socially optimal efforts are determined by the sum of the expected *marginal and cumulative* values of success, *V,* and failure, *F*, while privately optimal efforts are only driven by the corresponding expected *marginal* values.

In the following Sections 3.2 through 3.4, we will carry out a formal recasting of these results when protection or reward mechanisms are introduced.

3.2 Protecting Failure: Trade Secrets

Let's proceed to consider a state of the world, $\omega = 1$, corresponding to Case (2) in Table 1, where the information generated by a firm's negative results is fully protected. Protective instruments such as those afforded by trade secrets prevent firms that enter a patent race from benefitting from the negative information of prior firms. Granting protection to the information generated by their failures does not change the earlier firms' research outlook. Like in the benchmark case discussed in Section 3.1, firms' returns remain tied to their successful research result, net of research costs. So, although firms can prevent others from taking advantage of the negative information, they are unable to capture the informational value that their failure could confer on subsequent firms. Meanwhile, the follower firms also cannot take advantage of the negative information to increase the likelihood of success, so they face a similar return outlook as the first mover. Firms' expected returns under this protective regime are:

$$R_1^{\omega}(e_1) = p(e_1)\frac{1}{M}V - e_1 \tag{3.10}$$

$$R_2^{\omega}(e_2) = p(e_2)\theta_2\frac{1}{M}V - e_2 \tag{3.11}$$

$$\cdots$$

$$R_N^{\omega}(e_N) = p(e_N)\theta_N\frac{1}{M}V - e_N \tag{3.12}$$

3.2.1 Timing Incentives

By nullifying the informational benefits for the subsequent movers, the protection regime eliminates the informational benefit of waiting, hence extinguishing the last-mover advantage identified in the previous regime. Waiting reduces the opportunities for late movers: as the race progresses, it is increasingly likely that one of the firms achieves a successful result, eliminating the opportunity to enter for subsequent firms. In a protective regime, waiting only leaves late movers with the lost opportunity effect, thus leading to a first-mover advantage, $R_1^{\omega} > R_2^{\omega} > \cdots > R_N^{\omega}$.[12]

[12] This result may explain the empirical finding showing that patents are less important in industries where firms can effectively rely on a combination of first-mover effects and trade secrecy. See Graham et al. (2009).

3.2.2 Investment Incentives

Similarly to Eq. 3.05, a generic Firm i will maximize the expected private return given that it has not yet lost the opportunity to make a successful discovery, formally:

$$\max_{e_i} R_i^{\omega}\ (e_i) \mid (\theta_i = 1)\ \ i\ \in [1, N] \qquad 3.13$$

In a protective regime, every firm carries out research without access to the information acquired by prior firms' research and failed investments. As a result, it is straightforward to see that all competitors will exert the same research efforts, $e_N^{\omega *} = \cdots = e_2^{\omega *} = e_1^{\omega *} = e_1^{*}$. So, in a protective regime, every firm's effort level—including that of the first firm—will fall short of socially optimal levels, $0 < (e_1^{**} - e_1^{\omega *}) < (e_2^{**} - e_2^{\omega *}) < \cdots < (e_{N-1}^{**} - e_{N-1}^{\omega *}) < (e_N^{**} - e_N^{\omega *})$. Unlike in the benchmark case (Section 3.1), the last-mover's effort does not converge toward optimality, $e_N^{**} < e_N^{\omega *}$.

3.3 Rewarding Failure: Research Subsidies

Information externalities, just like any other externality can be tackled with tax/subsidy systems. In the following we shall consider a state of the world $\phi = 1$ corresponding to Case (3) in Table 1, where firms' negative research results are rewarded with subsidies. For example, under this regime, a firm can voluntarily disclose the details of its failed research attempts, publishing the negative results and making them publicly available to the industry, in exchange for a subsidy. The subsidy would be likely paid by the pool of likely beneficiaries of the information generated by the failing firm. As discussed below, the amount of the subsidy would need to be gauged to the external benefit created the failing firm's information, and this information may not be easy to ascertain by a central subsidizing agency.[13]

In this scenario, we continue to assume that failing firms do not reenter the patent race, thus they will all strictly prefer to disclose their negative results to obtain the subsidy. All subsequent firms can take advantage of the information made publicly available through this instrument. As will be shown in 3.3.2, the amount of the subsidy must be gauged along the patent race in order to bring about socially optimal levels of research effort.

When engaging in research, each firm's expected return consists of the expected value of its positive result (i.e., "success") plus the expected reward, $S_i > 0, \forall i \in [1, N-1]$, obtainable from the publication of its negative result (i.e., "failure"), net of research costs.

$$R_1^{\phi}(e_1) = p(e_1)\left[\frac{1}{M}V + \frac{M-1}{M}S_1\right] - e_1 \qquad 3.14$$

$$R_2^{\phi}(e_2) = p(e_2)\,\theta_2\left[\left(\frac{1}{M}+f_1\right)V + \frac{M-2}{M-1}S_2\right] - e_2 \qquad 3.15$$

...

$$R_{N-1}^{\phi}(e_{N-1}) = p(e_{N-1})\;\theta_{N-1}\left[\left(\frac{1}{M} + \sum_{i=1}^{N-2} f_i\right)V + \frac{M-N+1}{M-N+2}S_{N-1}\right] - e_{N-1} \qquad 3.16$$

$$R_N^{\phi}(e_N) = p(e_N)\,\theta_N\left[\frac{1}{M} + \sum_{i=1}^{N-1} f_i\right]V - e_N = p(e_N)\frac{1}{M}V - e_N \qquad 3.17$$

[13] In most situations, the value of a firm's failure only reaches other firms that are active in the same industry or line of research. In these cases, an industrial subsidy paid through revenue raised within the specific industry would be the appropriate instrument, to avoid inter-industrial effects. Conversely, when the benefits of the information extend to broader networks and to society at large (such as for the more general case of learning), governmental subsidies may be a viable instrument.

3.3.1 Timing Incentives

By comparing Eq. 3.14 through 3.17, we can observe that the private return of the first-moving firm is higher than that of the second-moving firm, and so on for all subsequent firms, i.e., $R_1^{\phi}(e_1) > R_2^{\phi}(e_2) > \cdots > R_N^{\phi}(e_N)$. This follows the same sequence observed for the socially optimal timing (Section 2.2.1). If $S_i = F_i V < V, \forall i \in [1, N-1]$, the subsidy regime will create optimal timing incentives for firms. Firms will engage in research as soon as an opportunity arises for them. As a first mover, the first firm can capture the first opportunity to succeed (and obtain V) or fail (and obtain the largest available subsidy, S_i), and so on for all subsequent firms. Under a properly gauged subsidy regime, private and social timing incentives will be aligned for all N firms: no firm will choose to delay its entry into the race, when an opportunity arises: $\left(R_1^S - R_1^{\phi}\right) = \left(R_2^S - R_2^{\phi}\right) = \cdots = \left(R_{N-1}^S - R_{N-1}^{\phi}\right) = \left(R_N^S - R_N^{\phi}\right) = 0$.[14]

3.3.2 Investment Incentives

As discussed above, when making their investment decision, firms have full knowledge of how past events materialized. Their opportunity to invest in research arises only when all previous firms have failed. Formally, firms' private optimization function becomes:

---

[14] As it is often the case with centrally administered taxes and subsidy mechanisms, the implementation of a subsidy system is not without limitations. Ascertaining the appropriate size of the subsidy may be difficult for a central entity, and distortions of information by the failing firms may arise. If the industry or the government cannot (or do not) verify negative results, a firm can assert that its research attempts failed and collect the subsidy without having actually engaged in the research or (more likely) while having invested inadequately in the research path. That is, at times it might be difficult for a central authority administering the subsidies to distinguish a negative result that identified a dead end from an unsuccessful or aborted research venture, that generated no new information. If so, then firms might devote an inadequate amount of effort to research, collecting subsidy based on false claims to have reached a negative result. The problem would be worse in the earlier stages, when the subsidy is large, than the later stages, when it is small.

$$\max_{e_i} R_i^{\phi}\ (e_i) \mid (\theta_i = 1)\ \ i\ \in [1, N] \tag{3.18}$$

The firms' privately optimal efforts in research under a research subsidy regime will thus be based on the knowledge of those past events, and their respective effort levels are given by the following equations:

$$e_1^{\phi *}: p'\ (e_1)\left[\frac{1}{M}V\ + \frac{M-1}{M}S_1\right] = 1 \tag{3.19}$$

$$e_2^{\phi *}:\ p'(e_2)\left[\left(\frac{1}{M}+f_1\right)V + \frac{M-2}{M-1}S_2\right] = 1 \tag{3.20}$$

$$\dots$$

$$e_{N-1}^{\phi *}\ : p'(e_{N-1})\left[\left(\frac{1}{M} + \sum_{i=1}^{N-2} f_i\right)V + \frac{M-N+1}{M-N+2}S_{N-1}\right] = 1 \tag{3.21}$$

$$e_N^{\phi *}: p'(e_N)\left[\frac{1}{M} + \sum_{i=1}^{N-1} f_i\right]V = 1 \tag{3.22}$$

Note that, as before, the value of each firm's negative information to other firms changes as the patent race progresses. So, for an optimal alignment of incentives, the government or industry social planner would have to carefully gauge and tailor the amount of subsidy to award to each firm for the expected cumulative value of its respective negative information in the amount of $S_i = F_i V < V, \forall i \in [1, N-1]$,[15] with F_i as expressed as in Eq. 2.16 (i.e.: $F_i = (N-i)f_i,\ \forall i \in [1, N]$).

Given the decreasing cumulative marginal value of firms' failures observed in Eq. 2.18 (i.e., $1 > F_1 > F_2 > \cdots > F_{N-1} > 0$), this would require reducing the amount of the subsidy as the patent race progresses, with higher subsidies given to the firms that pioneered newer research.

In an ideal world where subsidies could be determined efficiently, each firm's incentives would align with the social optimum. Given the public availability of the research information of earlier movers in the

[15] The failing firm's informational benefit increases the probability of achieving a discovery in the future and, by definition, it is lower than the actual value of the discovery.

research race, the probability of success for later entrants would increase. As a result, under this regime, the privately optimal research efforts would grow stronger as the patent race progresses, $e_N^{\phi *} > \cdots > e_2^{\phi *} > e_1^{\phi *}$, with each firm's privately optimal effort levels coinciding with the socially optimal levels: $(e_1^{**} = e_1^{\phi *}) < (e_2^{**} = e_2^{\phi *}) < \cdots < (e_{N-1}^{**} = e_{N-1}^{\phi *}) < (e_N^{**} = e_N^{\phi *})$.

In sum, the use of research subsidies that reward firms for the expected value of their failure can solve the suboptimal timing problem observed in the benchmark case (Section 3.1) and the suboptimal investment problem observed in the protection regime (Sections 3.2).

3.4 Protecting and Rewarding Failure: Negative Patents

Next, we will consider a state of the world corresponding to Case (4) in Table 1, where negative research results are fully protected by an enforceable right to exclude, $\omega = 1$, and this right to exclude is transferable to others so that firms have an opportunity to capture some of the financial value of the information generated from their failures, $\phi = 1$. In this scenario, the right to exclude is what "protects" the information and the right to transfer the information is what "rewards" the firm. Such a regime would grant property-type protection to a firm's negative information, much like the existing patent regime currently provides property-type protection to positive information: firms would be able to freely sell (and resell) the protected property rights in negative information to others. We shall metaphorically refer to this regime as the "negative patents" regime. Under this regime, negative information is not made publicly available—only the firm that owns the information has access to it.

In this scenario, firms do not reenter the patent race. Given the opportunity to sell the information to others, the failing firm's reservation price would therefore be zero because it cannot otherwise benefit from the negative information. Meanwhile, an entering firm would instead have a positive valuation and willingness to pay for that information. The immediate benefit that a firm would obtain from Firm *i*'s information is $f_i V$. The buying firm also knows that, if it also fails, it will be able to resell acquired information to other firms. Its valuation and willingness to pay for Firm i's information would thus rise to $F_i V$. In this setting, both firms would strictly prefer to trade the information, so the negative patent will be transferred at a price ranging between 0 and $F_i V$.

When engaging in research under this negative patent regime, firms' expected returns would include the costs of acquiring information from earlier-moving firms and the revenue obtainable through the sale of negative information (both acquired and produced) to subsequent firms. Under our symmetric information trading environment, we can follow the standard convention and assume that the parties will share the gains from trade in proportion to their relative bargaining power.[16] Let α and $1-\alpha$ respectively denote the bargaining power of the seller (i.e., the failing firm, selling the negative patent) and the buyer (i.e., the following firm, purchasing the negative patent), with $\alpha \in [0,1]$.[17]

[16] For our analysis, it is sufficient to observe that failing firms can internalize a larger or smaller share of the informational value of their negative patent depending on their bargaining power relative to that of prospective buyers. For analytical convenience and ease of interpreting the results, we assume α to be an exogenous parameter, denoting the bargaining power of sellers (i.e., failing firms). This assumption can be relaxed to allow for idiosyncratic differences of α across firms without affecting our qualitative results. In the existing literature, starting from the seminal paper by Dunlop and Higgins (1942), patience in negotiations and competition are identified as the two main exogenous determinants of bargaining power. In our setting, failing firms are the sole sellers of their negative patent, whereas multiple firms may face time pressure and compete with one another to acquire the negative patent, giving the selling firm an advantage in capturing the informational value of its failure. Things may play out differently if only one prospective buyer approached the selling firm to acquire the negative patent. In this case, the new firm cannot effectively engage in research without the information from the failing firm, and the failing firm has no use for the negative information if it is not sold to the new firm. Therefore, both firms have zero outside options and a Nash (1950) bargaining split of the value of the negative patent would likely occur. See Aghion and Holden (2011) and Holden and Malani (2014). See also the relevant real-world discussion of competition as a restraint to bargaining power in Choi and Triantis (2012).

[17] Under the alternative bargaining protocol used in the literature, the seller and buyer make a demand/offer with certain probability, exogenously given. Under this alternative protocol, the probabilities α and 1-α can be understood as proxies for the bargaining power of the seller and buyer, respectively. Assume the failing firm (seller) makes a take-it-or-leave-it (TIOLI) demand with probability $\alpha \in (0,1)$; the new entrant firm (buyer) makes a TIOLI offer with the remaining probability, $(1-\alpha)$. When the failing firm gets to make a TIOLI offer, its demand will approach $F_i V$. On the other hand, if the buying firm gets to make the TIOLI offer, the offering price will approach the reservation value of the failing firm, 0.

Let us consider Firm 2 acquiring a negative patent that gives it access to and exclusive ownership of the negative information generated by Firm 1's research failure. This negative patent will therefore help Firm 2 pursue its research objectives with a direct informational benefit of f_1. In case of its own failure, Firm 2 will also be able to resell the negative patent to the subsequent firm, capturing the downstream informational value $(N-2)f_1$.[18] In case of failure, Firm 2 will additionally have the negative information generated by its own failure, F_2. The total information owned by Firm 2 after its own failure will thus be $(F_1 - f_1) + F_2$.

The same reasoning applies to the value of negative patents for subsequent entrants. After its own failure, Firm 3 will own the bundle of information generated by Firms 1 and 2, with a total informational value of $(F_1 - 2f_1) + (F_2 - f_2) + F_3$, and so on. The information acquired by Firm *N* has no resale value, since no other firm will enter the race and be willing to pay for it.

Let us denote the value of the negative patent that a generic Firm *i* can transfer to Firm *i* +1 as:

$$\pi_i = \{F_i + \sum_{j=1}^{i-1}[F_j - (i-j)f_j]\}V \qquad 3.23$$

Recall that with probability $1 - p(e_i)$, Firm *i* will reach no result. In this case, Firm *i*'s research adds no new information, $F_i = 0$. Firm *i* will exit the race, reselling the previously acquired information at price $\alpha\,\pi_i$, in this case equal to $\alpha\,\pi_{i-1}$ for $\forall i \in [1, N-1]$.

Additionally, for Firm 1, $\pi_1 = F_1 V$, because the patent π_1 conveys only the information derived from Firm 1's failure (no prior failures occurred before Firm 1). Under our trading environment, Firm *i* will thus be able to capture a share of the value of its negative patent, obtaining a price $\alpha\,\pi_i$, for $\forall i \in [1, N-1]$.

Under a negative patent regime, the firms' expected returns will thus be: [19]

---

[18] The sum of the direct and patent resale values of the information will thus equal $(N-1)f_1$. For a generic Firm *i*, the sum of the direct and resale patent values of the information will thus be equal to $F_i = (N-i)f_i,\ \forall i \in [1, N]$, as defined in Eq. 2.16.

[19] Firm 2's expected returns in Eq. 3.25 directly derives from these steps: $R_2^{\rho} = \theta_2 \left\{p(e_2)\left[\left(\frac{1}{M}+f_1\right)V + \frac{M-2}{M-1}\alpha\,\pi_2\right] + (1 - p(e_2))[\alpha\,\pi_1] - \alpha\,\pi_1 - e_2\right\} = \theta_2 \left\{p(e_2)\left[\left(\frac{1}{M}+f_1\right)V + \frac{M-2}{M-1}\alpha\,\pi_2\right] - p(e_2)\,\alpha\,\pi_1 - e_2\right\} =$

$$R_1^\rho(e_1) = p(e_1)\left[\frac{1}{M} + \frac{M-1}{M}\alpha\, F_1\right]V - e_1 \quad 3.24$$

$$R_2^\rho(e_2) = \theta_2\, \left\{p(e_2)\left[\left(\frac{1}{M}+f_1\right) + \alpha\, f_1\right] V - e_2\right\} \quad 3.25$$

...

$$R_{N-1}^\rho(e_{N-1}) = \theta_{N-1}\left\{p(e_{N-1})\left[\left(\frac{1}{M} + \sum_{i=1}^{N-2} f_i\right) + \alpha\, f_{N-2}\right]V - e_{N-1}\right\} \quad 3.26$$

$$R_N^\rho(e_N) = \theta_N\left\{p(e_N)\left(\frac{1}{M} + \sum_{i=1}^{N-1} f_i\right)V - \alpha\, \pi_{N-1} - e_N\right\} \quad 3.27$$

where π_{N-1} in Eq. 3.27 is defined as in Eq. 3.23.

3.4.1 Timing Incentives

By comparing Eq. 3.24 through 3.27, we can observe that, when the sellers have zero bargaining power, $\alpha = 0$, the private return of early movers is lower than that of last movers, i.e., $R_1^\rho(e_1) < R_2^\rho(e_2) < \cdots < R_N^\rho(e_N)$, since the first mover produces negative information for its followers without obtaining any revenue, and without any information from prior firms.

The opposite result occurs when the sellers have full bargaining power, $\alpha = 1$, for which $R_1^\rho(e_1) > R_2^\rho(e_2) \lesseqgtr \cdots \lesseqgtr R_{N-1}^\rho(e_{N-1}) > R_N^\rho(e_N)$. This occurs because the first-moving firm's expected returns include the highest opportunity to reach successful research (i.e., $1 = \theta_1 > \theta_2 > \cdots > \theta_N$) and the highest cumulative informational benefits (i.e., $\frac{M-1}{M}F_1$), and not the costs of acquiring information from past firms. This results in the first mover's expected returns being greater than those of the other competitors.

Conversely, at the end of the race, the last mover's expected returns only include the costs of acquiring information from past firms (i.e., π_{N-1}), but not the revenue obtainable through the sale of their negative

$\theta_2\, \left\{p(e_2)\left[\left(\frac{1}{M}+f_1\right)V + \alpha\left(\frac{M-2}{M-1}\,\pi_2 - \pi_1\right)\right] - e_2\right\}$. By substituting π_i as expressed in Eq. 3.23, $R_2^\rho = \theta_2\, \left\{p(e_2)\left[\left(\frac{1}{M}+f_1\right)V + \alpha\left(\frac{V}{M(M-1)}\right)\right] - e_2\right\}$, where $\frac{V}{M(M-1)} = f_1 V$, hence Eq. 3.25. Similar steps apply to compute the expected returns of the remaining $N-1$ firms.

information to subsequent firms. Taking also into account that the last mover has the lowest opportunity to reach successful research (i.e., $\theta_1 > \theta_2 > \cdots > \theta_N$), these "lowering" factors outweigh the "positive" informational benefits that the last mover can derive from previous failures (i.e., $\sum_{i=1}^{N-1} f_i$). This results in the last mover's expected returns being lower than those of the other competitors.

As for what occurs between these two extremes, the ranking between the other competitors' expected returns cannot be clearly stated. Indeed, as the race progresses, two countervailing factors affect the expected revenues: the *decreasing* opportunity to reach successful research vs. the *increasing* informational benefits from previous competitors' failures. For example, by comparing Firm $N-1$'s vs Firm 2's expected returns, we have that $\theta_{N-1} < \theta_2$, but $\sum_{i=1}^{N-2} f_i > 2f_1$. This results in $R_2^{\rho}(e_2) \lesseqgtr \cdots \lesseqgtr R_{N-1}^{\rho}(e_{N-1})$.

3.4.2 Investment Incentives

Similarly to Eq. 2.23, a generic Firm *i* will maximize the expected private return given that it has not yet lost the opportunity to make a successful discovery, formally:

$$\max_{e_i} R_i^{\rho}(e_i) \mid (\theta_i = 1) \;\; i \in [1, N] \qquad 3.28$$

The firms' privately optimal efforts in research under a negative patent regime will thus be based on the knowledge of those past events, and is given by the following equations:

$$e_1^{\rho *}: p'(e_1)\left[\frac{1}{M} + \frac{M-1}{M}\alpha F_1\right]V = 1 \qquad 3.29$$

$$e_2^{\rho *}: \; p'(e_2)\left[\left(\frac{1}{M} + f_1\right) + \alpha f_1\right]V = 1 \qquad 3.30$$

…

$$e_{N-1}^{\rho*} : p'(e_{N-1}) \left[\left(\frac{1}{M} + \sum_{i=1}^{N-2} f_i \right) + \alpha\, f_{N-2} \right] V = 1 \qquad 3.31$$

$$e_N^{\rho*} : p'(e_N) \left(\frac{1}{M} + \sum_{i=1}^{N-1} f_i \right) V = 1 \qquad 3.32$$

In the limiting case where the sellers have full bargaining power, $\alpha = 1$, the first- and last movers' research investments will approach the socially optimal level $e_i^{\rho*} \rightarrow e_i^{**}$, with, $e_N^{\rho*} > e_1^{\rho*}$, while the other competitors' research investments will be below the socially-optimal level. Considering that $\sum_{i=1}^{N-1} f_i > \sum_{i=1}^{N-2} f_i + f_{N-2} > \cdots > 2f_1$, it follows that $e_N^{\rho*} > \cdots > e_2^{\rho*} > e_1^{\rho*}$.

Where the sellers have zero bargaining power, $\alpha = 0$, and buyers can exploit the entire value of the seller's information, the firms' incentives will revert to those observed in the absence of protection or reward (Section 3.1), $e_i^{\rho*} \rightarrow e_i^*$. Hence $e_N^{\rho*} > \cdots > e_2^{\rho*} > e_1^{\rho*}$, where only the last mover makes the socially optimal effort, $e_N^{\rho*} = e_N^{**} = e_N^*$.

3.5 Effects of Protective vs Rewarding Remedies on Research

In the previous sections, we developed a simple model to show the conditions under which negative information from research failures generates a misalignment between the private and social incentives to research. Failure is an important and vastly underexplored component of the value of research. Information about past failures allows other firms to avoid running into the same dead ends, increasing the odds that future researchers will succeed. Unless there are instruments that allow firms to internalize the value of their failures, private and social incentives for research will diverge.

Building on the results derived in Sections 3.1 through 3.4, we will now compare the effect of protective and rewarding remedies in fostering an optimal pace of research. Specifically, in Section 3.5.1, we look at the first mover's returns from research under the four regimes under consideration. Expected returns are important to predicting which regimes will yield a socially optimal return to capital and foster timely participation of private firms in the research industry. In Section 3.5.2, we look at the resulting aggregate level of investments in research to see which regime will induce the industry to make socially optimal levels of investments.

3.5.1 Effects on Returns from New Research

Our analysis revealed that socially optimal returns from research investments can be more easily achieved through rewarding instruments, such as research subsidies and negative patents, than through protective instruments, such as trade secrets.

The critical feature of both types of rewarding instruments is the financial reward that a failing firm may obtain in exchange for the disclosure of negative information. Research subsidies and negative patents differ in that they rely on regulatory mechanisms and markets, respectively, for generating the financial incentives to disclose negative information. With research subsidies, the failing firm obtains the financial reward by complying with regulatory guidelines and disclosing information to the entire industry. With negative patents, the failing firm obtains the financial reward through the market by selling the negative information to another firm.

Rewarding instruments produce two important effects: (i) retrospectively, they incentivize firms to disclose the negative information that they acquired, making that information available to other firms in the industry; (ii) prospectively, they allow firms that invest in research to benefit not only from successful discoveries but also from the knowledge generated by their failures. The effects of subsidies and negative patents diverge when we consider the firms' expected returns under the two regimes. Unlike the financial reward obtainable with research subsidies, the financial reward obtainable under the negative patent regime (for all $\alpha > 0$) comes from subsequent firms paying the patent price. These payments impose a cost on all firms $\in (1, N]$, which reduces their expected returns below the level observed under the subsidy regime.

With protective instruments, such as trade secrets, failing firms obtain "protection" of their information, thereby preventing subsequent firms from getting a free informational advantage from earlier-movers' failures. However, the protection obtained through trade secrets does not provide other firms with access to the failed research information, which creates a risk of duplicative research investments within the industry. Trade secret protection also does not allow firms to internalize the value of the knowledge generated by their failures.

By comparing the first mover's expected returns vis-à-vis the social optimum, we observe that:

$$\left(\underbrace{R_1^S}_{Social} = \underbrace{R_1^{\phi}}_{Subsidy} = \underbrace{R_1^{\rho}(\alpha = 1)}_{\substack{Negative\\Patent\\(\alpha=1)}}\right) > \underbrace{R_1^{\rho}(\alpha \in (0,1))}_{\substack{Negative\\Patent\\(\alpha\in(0,1))}} > \left(\underbrace{R_1}_{Benchmark} = \underbrace{R_1^{\omega}}_{\substack{Trade\\Secret}} = \underbrace{R_1^{\rho}(\alpha = 0)}_{\substack{Negative\\Patent\\(\alpha=0)}}\right) \quad (3.33)$$

Regimes that yield first mover's expected returns that approach the returns observed in a social optimum—i.e., research subsidies and negative patents when the sellers have full bargaining power—will foster optimal timing with respect to when firms initiate new research. More specifically:

(i) A research subsidy regime, with $S_i = F_i V$, brings expected returns from research to a socially optimal level for both first movers and for the aggregate.

(ii) For the first mover, a negative patent regime could also achieve the same result in the limiting case where sellers have full bargaining power $\alpha = 1$. In this scenario, the first failing firm can price the negative patent at, $\pi_1 = F_1 V$, capturing the full informational value of its failure. Also in this case, the first mover's returns will approach social optimality and incentivize optimal investments in research. However, the aggregate expected returns under this negative patent regime with $\alpha = 1$ will be at a lower level, given the cost of paying for the informational benefit from prior firms that is imposed on subsequent firms.

(iii) In the remaining cases (i.e., benchmark case, and negative patents with $\alpha = 0$ or $\alpha \in (0,1)$), firms—including the first mover—do not capture the full benefit or any of the informational benefit created by their failures. In these regimes, both the first mover's returns and aggregate returns to investment are lower than the social optimum.

It should be noted that the observed wedge between the expected revenues in the subsidies and negative patents regimes comes from the source of the reward for the failing firms. Subsidies are paid by the industry or the government to failing firms. But the subsidies are not a free lunch and have to be funded somehow. If they are funded by the government out of general revenues, then the cost is dispersed throughout society (this spreading of the cost might be sensible for research that has widely-dispersed benefits); if they are funded by the industry, then the cost falls on a smaller group of firms that are active in that line of business (this concentration of the cost might be sensible for research that has more narrowly-felt benefits). If the subsidies are funded from an industry-specific tax (or similar mechanism), then its impact on the expected revenues from research would be similar to that of the negative patent regime.

3.5.2 Effects on Aggregate Investments in Research

Our analysis revealed that adopting rewarding instruments, such as research subsidies and negative patents, can more easily generate socially optimal incentives for research than can protective instruments, such as trade secrets.

By comparing the aggregate investments ($Ae = \sum_{i=1}^{N} e_i$) that firms will make under the alternative regimes we examined vis-à-vis the social optimum , we observe that:

$$\left(\underbrace{Ae^S}_{Social} = \underbrace{Ae^\phi}_{Subsidy}\right) > \underbrace{Ae^\rho(\alpha = 1)}_{\substack{Negative\\Patent\\(\alpha=1)}} > \underbrace{Ae^\rho(\alpha \in (0,1))}_{\substack{Negative\\Patent\\(\alpha\in(0,1))}} > \left(\underbrace{Ae}_{Benchmark} = \underbrace{Ae^\rho(\alpha = 0)}_{\substack{Negative\\Patent\\(\alpha=0)}}\right) > \underbrace{Ae^\omega}_{\substack{Trade\\Secret}} \quad (3.34)$$

Equation 3.34 shows the effects of protective vs rewarding regimes in incentivizing research investments. Specifically:

(i) A research subsidy regime, $S_i = F_i V$, yields aggregate investments in research that equal those observed in a social optimum. These investments will foster a socially optimal pace of research.

(ii) The remaining cases will exhibit lower aggregate investments in research. Trade secrets will result in the lowest level of investment. Then, the benchmark case and the negative patents case where $\alpha = 0$ will lead to a comparatively higher level of investment. Finally, a negative patents regime where $\alpha \in (0,1)$, negative patents with $\alpha = 1$ will yield the highest levels of investment, though such levels will still be lower than the research subsidy regime. In these cases, failing firms do not capture the full, or any of the informational benefit created by their failed research.

In our setting where firms do not reenter the race after a failure, rewarding instruments can yield socially optimal outcomes, both in terms of timing and investments in research. Conversely, protective instruments like trade secrets do not allow failing firms to capture the value of their failures. Rather, trade secrets merely eliminate the positive information externality that favors follower firms while providing no financial return for the negative results generated by firm failures. Beyond this, trade secrets deprive follower firms of access to failed research information, thereby eliminating another potential benefit. This results in duplicative research investments within the industry, rendering the lowest levels of research returns and investments in research compared to alternative solutions. This suggests that alternative, reward-based instruments can result in better research conditions and outcomes.

4. Protecting and Rewarding Failures: Market and Legal Solutions

The lack of sharing of negative information exists in other research settings beyond the private R&D context. In the academic sphere, it commonly manifests as "publication bias" (Stanley, 2005; Franco and Simonovits, 2014). Publication bias constitutes the selective and exclusive publication of results that reject a null hypothesis. This practice causes the bulk of published results to be misleading because the excluded negative results might otherwise cast doubt on published positive results (Dubben and Beck-Bornholdt, 2008; Stanley, 2008). The academic sphere seeks to address this challenge by establishing channels for sharing failures. For instance, several academic journals have arisen that exclusively publish negative results.[20] Meanwhile, others are establishing registration protocols that allow researchers to make exclusive submissions to journals before running an experiment, with the journal committing to publish their results—positive or negative—once the experiment proposal is accepted.[21]

Similar challenges are faced by defense departments with respect to information generated by failed military operations and errors in technology. For example, when air vessels are pushed to the limits of their design boundaries, this natural experimentation occasionally yields negative results. But reputational norms discourage officers and their commanders from publicizing these errors and failures. Yet the information created could generate a valuable data set for future pilots and technology developers. In recent years, the reluctance to disclose military failures and technology errors has been addressed by military forces and

[20] Several journals specialize in publishing negative or null results across various disciplines, reflecting a growing commitment to transparency and openness in scientific research. Noteworthy examples include the Journal of Negative Results (focusing on ecology and evolutionary biology); Journal of Negative Results in Biomedicine; Journal of Pharmaceutical Negative Results; Journal of Interesting Negative Results (specializing in natural language processing and machine learning). Moreover, the Journal of Trial and Error (JOTE) stands out by not only accepting negative results but also publishing rejected grant applications. For more details, refer to their website: https://journal.trialanderror.org/pub/sciencefails/release/4.

[21] This submission method is commonly known as "Registered Reports." For additional information, please refer to the webpage https://www.cos.io/initiatives/registered-reports. As of February 10, 2024, the website reports that "over 300 journals use the Registered Reports publishing format either as a regular submission option or as part of a single special issue. Other journals offer some features of the format." (Center for Open Science, n.d.). In law and economics research areas, journals accepting Registered Reports include the *Journal of Development Economics*, *Oxford Open Economics*, and *Law and Human Behavior*. Several other journals are currently considering the adoption of this submission protocol.

counterbalanced by the creation of an obligation to record all failed operations and errors in a register, shared with all allied forces.[22]

However, the emerging "failure-sharing" norms of academia and the military do not transfer perfectly to private sector R&D. Academic researchers, defense departments and private companies may enjoy similar benefits when learning from others' failures—namely, they avoid inefficient repetitions of similar errors, facilitating their pursuit of positive results. But academics and defense departments, on the one hand, differ from private firms on the other because private research investments are driven by expected financial returns, rather than academic rewards or governmental objectives. Academics may derive reputational benefit from having an additional publication in their name, even if their research leads to a negative result such as failing to refute a null hypothesis. Academia in general may also be more amenable to forming collaborations on the basis for failures—publishing negative results may help academics find others with like-minded research interests and other complementary skills, thereby creating new opportunities for academics to work together. These reputational benefits and opportunities for cooperation may offset the potential advantages of publishing first in the academic context; nor is there a governmentally granted monopoly in the future use of positive scientific findings and ideas. Military officers can be required to follow a failure-sharing norm by their command-and-control structure. The market does not offer a meaningful reputational reward or a hierarchical enforcement system for the sharing of negative information by a given firm.

Conversely, as we demonstrated in Section 3.1, when the value of research failures is not internalized, a divergence will occur between privately optimal and socially optimal equilibria. Firms can benefit from accessing the negative information produced by earlier research failures insofar as they can avoid allocating their own scarce resources to dead ends. When firms do not compensate the failing firm for its negative information, an externality arises that equals the aggregate value of the informational benefits conferred on all competing firms. Accordingly, the market equilibrium will not maximize the sum of the surplus for (1) the firm that produced the negative information and (2) the beneficiaries of the information externality.

[22] We would like to thank former Italian Secretary of Defense, Arturo Parisi, and joint Chief of Staff of the Italian Armed Forces, Gen. Vincenzo Camporini for pointing out this information and sharing personal narratives about the existence of this problem at the national and EU levels, which was addressed during their terms in office (2006-2008 and 2008-2011, respectively), with the creation of a duty to record all erred flight operations, and the establishment of national database, shared with all allied national forces, and the gradual acceptance of a "failure-disclosure norm" within the armed forces.

Positive externalities from research failures thus lead markets to engage in research at a slower pace or with a lower level of investment than is socially desirable.

We discussed several possible regimes for dealing with this problem in Sections 3.2 through 3.5. However, the regimes examined above are of little use if they cannot be implemented in the real world. The following discussion accordingly turns to existing and potential solutions to the information externality problem created by research failures. The economic literature has identified several ways to correct externality-based problems in general. These include Pigouvian taxes (for negative externalities), subsidies (for positive externalities), private solutions (e.g., contracts, joint ventures, and mergers), and other regulatory and mixed solutions. These solutions have been extensively applied to externality-based problems associated with research that produces positive information.

Existing literature has not, however, adequately explored the unique policy challenges posed by negative information. In the following, we consider the benefits of and limits to existing and potential legal, regulatory, and private solutions in the context of negative information. We aim to make two contributions here. First, we aim to show precisely how and why existing mechanisms do not adequately align private and social equilibria. Second, we aim to show that, although potential solutions are not unachievable theoretical ideals, serious implementation problems must be overcome before they can be applied as real policy instruments.

In Section 4.1, we begin by considering the case of a market in which the legal system has no special rules applicable to the negative information associated with research failures. In this scenario, negative information would be treated much as if it were run-of-the-mill confidential information; more particularly, we set aside the direct use of trade secret law to protect negative information. Firms could, in principle, contract over negative information. Absent an agreement, however, no firm could be liable for using negative information. In Section 4.2, we consider how the availability of trade secret protection for negative information influences our analysis. In doing so, we identify several shortcomings of the existing trade secret regime as applied to negative information. We also explain why there are inherent limits to the use of trade secrecy as a solution to the problems we identified. In Section 4.3, we turn to the patent system. Although the results of research failures cannot themselves be directly patented, we explain how several existing doctrines may influence a firm's decision on whether to disclose negative information. We also explain how the influence of the patent system depends on the particular research context to which it is applied. In Section 4.4, we explore the possibility of combining several mechanisms into a market institution that resembles our hypothetical system of "negative intellectual property." We consider the way in which a market for such negative intellectual property could serve as an instrument for internalizing positive and negative research externalities and identify the intrinsic limitations of such hypothetical solutions. Finally, in Section 4.5, we conclude by examining the use of potential combinations of IP and non-IP innovation incentives and allocation mechanisms, like grants and tax subsidies, to influence firms' decisions to produce and disclose negative information.

4.1 Contracting over Negative Information

We begin here by evaluating a legal system that offers no special protection for negative information. That is, we assume first that negative information does not qualify for protection under trade secret or patent law. In this scenario, a firm contemplating the possibility that its research investment might yield negative information is not entirely without recourse—instead, it could try to limit spillovers by relying on contracts (1) with its own employees; and (2) with other firms.[23] A wide array of such arrangements have been extensively analyzed in the context of positive information (Merges, 1999; Arora and Merges, 2004). However, as we will explain, such contracts would entail significant transaction costs because of the unique practical challenges posed by negative information. Our analysis reveals that contracts standing alone should be understood as something like an intermediate case between the "benchmark" scenario of Section 3.1 and the "protect" scenario of Section 3.2; that is, firms can use contracts to modestly reduce the ability of competitors to observe the results of their research, and they cannot realistically use contracts to obtain

[23] Of course, if the research produced positive information, then the firm could rely on trade secret or patent law to internalize a substantial fraction of the value of that information.

significant rewards for negative information. Thus, if the only way for firms to obtain protection or rewards for negative information were to rely on contracts, firms would underinvest in research.

Begin with the contracts that a firm might use to limit spillovers by its employees.[24] Firms regularly demand that their employees sign confidentiality agreements as a condition of employment (Hrdy and Seaman, 2023). Such agreements purport to restrict employees' ability to use or disclose information acquired during their employment. And they apply whether *vel non* such information can be protected as a trade secret or form the basis for a patent application (Hrdy and Seaman, 2023). As a general matter, then, they appear to limit rivals' ability to free-ride on negative information; firms might therefore try using them to protect the results of their research efforts.

In the context of negative information, however, enforcement costs will severely limit the efficacy of such contracts. To see why, suppose there are three potential treatments for Alzheimer's: X, Y, and Z. A research firm conducts a well-powered trial that reveals that candidate treatment X produces toxic side-effects, such that X is not suitable for use in humans. An employee who has signed a confidentiality agreement then leaves the research firm and begins to work at a competing firm.

Even without disclosing this information, the employee could "use" it in a variety of ways that would likely evade enforcement by the research firm. For example, if the employee tried to persuade her new coworkers to research candidate treatment Y, her arguments would be motivated by her knowledge of the negative information associated with candidate treatment X, and this is so even if she never disclosed the

[24] The discussion that follows assumes, contra our model, that the firm that conducts the initial research does not exit upon failure. If that firm did exit upon failure, ex-employees would likely obtain jobs at competing firms, where they would be able to use and disclose negative information with relative impunity because the firm that conducted the research would no longer be around to enforce whatever contracts it had made prior to exit. Thus, although our aim is to show why contractual mechanisms are ineffective solutions to the problem of negative information, we note that the problem is likely even more severe in real-world environments that more closely resemble our "exit upon failure" model (e.g., industries where startup firms attempt to introduce a single, innovative product and wind down if their initial R&D efforts fail) because firms in those scenarios cannot even rely on the (ineffective) contractual mechanisms we describe.

negative information itself.[25] Crucially, however, it would be exceptionally difficult for the research firm to observe such uses and still more difficult for the firm to discover evidence it could rely on as the evidentiary basis for a lawsuit—there is rarely an overt act or affirmative statement that the research firm might discover to indicate that its former employee is using her knowledge of the negative information associated with candidate treatment X to motivate her advocacy for candidate treatment Y.[26] Accordingly, observability and verifiability will often pose insurmountable barriers to the use of confidentiality agreements to limit spillovers of negative information.[27]

Because confidentiality agreements cannot adequately protect negative information, a research firm may also try to use noncompete agreements with its employees. Indeed, because the research firm can readily observe whether a former employee is working for a competing firm, a noncompete agreement imposes much less significant enforcement costs than does a confidentiality agreement. That said, noncompete agreements entail significant social costs because, *inter alia*, they reduce labor market mobility. The legal

---

[25] In the analogous context of trade secrecy, "use" has been construed to cover "relying on the trade secret to assist or accelerate research or development." (Oakwood Laboratories LLC v. Thanoo, 999 F.3d 892, 909, 3d Cir. 2021). If the employee successfully persuades her new firm to allocate resources toward candidate treatment Y, that would "assist or accelerate research or development" insofar as it avoids the dead-end of candidate treatment X.

[26] At least absent disclosure. Of course, if the former employee were to tell her new coworkers that they should pursue candidate treatment Y because she learned in her prior job that candidate treatment X is a dead end, then there would be the sort of evidence that could form the basis for a successful lawsuit by the research firm. Still, without such disclosure, it is hard to imagine that the research firm could prove "use" of negative information outside of very unusual cases (e.g., the research firm's former employee persuades her new employer to suddenly stop researching candidate treatment X by strenuously arguing in favor of candidate treatment Y, in a scenario in which X and Y are the only plausible candidates). For an exception, consider Anthony Levandowski's departure from the self-driving car startup Waymo to form his own firm, Otto, which was quickly acquired by Uber. Waymo alleged that Levandowski downloaded large amounts of confidential information from his work devices immediately before his departure, and then used that information to benefit Otto and Uber. Negative information likely comprised a significant fraction of the allegedly misappropriated information, and Waymo was apparently able to detect its misuse, although perhaps in part because it was combined with positive information. The case ultimately settled for a quarter-billion dollars in 2018 (Marshall, 2018).

[27] This scenario is distinguishable from scenarios involving positive information, where the research firm could plausibly detect the competing firm's sales of a product or process that incorporates information about what does work. Observability accordingly does not necessarily pose an insurmountable barrier to the use of confidentiality agreements to limit spillovers of positive information. We do not claim that confidentiality agreements are necessarily sufficient to solve externality-based problems associated with positive information. We only claim that, even if such agreements are sometimes adequate in the context of positive information, they are much less (if ever) adequate in the context of negative information.

system accordingly requires that the duration and scope of a noncompete agreement be "reasonable." This standard typically limits the effective duration of a noncompete to two years or less.

To be sure, this temporal limit would not necessarily render noncompete agreements completely ineffective. This is because the value of negative information eventually drops to zero—indeed, once the successful path has been identified, information regarding the failure of alternative paths becomes worthless.[28] If a noncompete were in place until the successful research path has been identified, it would prevent competing firms from acquiring negative information by hiring the research firm's employees, at least for the period during which the negative information has value. Thus, so long as the successful path is identified within two years, noncompete agreements might effectively protect negative information, even if they wouldn't provide effective rewards.

But in many cases, it will take more than two years to identify the successful path. In those cases, negative information will retain value longer than the maximum extent of a noncompete agreement.[29] As a result, noncompete agreements can only partially protect negative information.[30] Moreover, even a noncompete agreement that effectively prevents an employee of the research firm from moving to a competitor cannot effectively prevent that employee from simply disclosing the negative information to employees of competing firms, as the preceding discussion regarding confidentiality agreements suggested.

Although firms might try to protect negative information by contracting with their own employees, in order to obtain rewards, they must contract with other firms too. Such efforts, however, will likely be frustrated by search costs, bargaining costs, and (again) enforcement costs. First, potential buyers and sellers

---

[28] In keeping with the model we develop here, this assumes that there is only one successful research path. If there are multiple successful research paths, then the value of negative information would still eventually drop to zero. Instead of occurring upon the discovery of a successful research path, however, it would occur once all paths have been explored. We discuss this possibility, *infra*, in our analysis of patent law's enablement requirement.
[29] This assumes, of course, that the employee departs, and the clock on the noncompete agreement's duration starts, immediately upon discovery of the negative information.
[30] And, again, longer noncompetes likely impose social costs that would outweigh any potential benefits as instruments for solving externalities associated with negative information.

of a particular piece of negative information will have a hard time finding each other. A company that is pursuing a particular research path is unlikely to publicly disclose that it has not yet found a positive solution to a given question. Thus, a prospective buyer (i.e., a company engaged in the same research area) would likely struggle to identify prospective sellers of negative information. Similarly, companies that possess valuable negative information might not be able to assess, based on publicly available information, what research paths other companies are actively pursuing. Therefore, a prospective seller would likely have trouble identifying prospective buyers (i.e., companies that are considering but have not yet pursued the research path to which the negative information relates).

Second, even if potential buyers and sellers of negative information find each other, bargaining costs will make it difficult for them to agree on a price because they will face Arrow's information paradox, which suggests that bargaining for information through partial disclosure leads to suboptimal valuation and trading of that information (Arrow, 1962). As Arrow explained the paradox, an information seller will demand payment before disclosure in order to avoid the consequences of the nonexcludable character of information. An information buyer, however, will demand the opposite sequencing—disclosure before payment—because that is the only way for her to value the information before deciding whether the seller's price is appropriate. The seller, however, will refuse because if she discloses too much of the information, the buyer effectively acquires it without paying and the seller has no recourse to demand compensation. As a result, most treatments of the problem conclude that information sharing is only feasible in a system that legally protects information (e.g., through patents); information generated during stages of the research process before legal protection arises cannot easily be traded (Arrow, 1962; Bar-Gill and Parchomovsky, 2009).

Still, some work indicates that staged disclosures or relational contracts enable information exchange even without effective patent or trade secret protection (Burstein, 2012). Consider a biotech firm that has conducted some research on candidate treatment X. The biotech firm wants to contract with a large pharmaceutical firm that could do the additional research required to obtain FDA approval. This is the classic scenario described by Arrow: as soon as the biotech firm discloses the chemical structure of candidate treatment X, the pharmaceutical firm could simply use that information without paying for it. The biotech firm must therefore find alternatives to revealing the structure of candidate treatment X.

For one possibility, the biotech firm might disclose ancillary information regarding candidate treatment X without disclosing information regarding its structure (assuming both that the structure is where the exploitable value lies and that information regarding the structure is wholly nonexcludable). In some cases, the large pharmaceutical firm can rely on that ancillary information to assess the value of information regarding its structure. For example, information regarding the therapeutic area, biological targets, and preliminary testing results may enable firms to negotiate for disclosure of the structure of candidate treatment X without creating an intolerable risk that the large pharmaceutical firm will appropriate the most valuable information (i.e., that regarding its structure). In such instances, staged disclosures supported by non-disclosure agreements, term sheets, and other preliminary contracts may facilitate sales of positive information (Burstein, 2012).

In other cases, the biotech firm and the large pharmaceutical firm envision a long-term relationship involving the bilateral transfer of information. That is, in addition to the chemical structure of candidate treatment X, the firms may expect to exchange a range of information about how to research, manufacture, distribute, or market candidate treatment X. Some information will originate with the biotech firm; other information will originate with the pharmaceutical firm. Importantly, at least some of this information will be tacit, rendering it partially excludable; in such instances, the firms might enter into contracts that establish governance mechanisms to facilitate ongoing information exchanges (Gilson et al., 2009). Such contracts do more to dictate how the firms would resolve disputes regarding exchanges of information than they dictate precisely what the firms must disclose and how much they would pay for such disclosures.

But these solutions are unworkable in the negative information scenarios we have described. In particular, staged disclosure will not work because a large pharmaceutical firm would not continue to negotiate over the right to develop a compound whose ancillary information was not promising. It is the ultimate right to bring a treatment to market that motivates the transaction in the staged disclosure scenarios. However, if preliminary testing results demonstrate that candidate treatment X does not work, the ultimate right to bring that treatment to market is worthless and there is nothing to motivate the transaction.

In addition, the long-term relationships envisioned by Gilson et al. (2009) rely on complementarity between tacit information (which is at least partially excludable) and codified information (which is nonexcludable). In the context of negative information, as we define it here, there is no tacit information—the only possible outcomes are that the research path failed, that it succeeded, or that it generated no information. In these settings, relational contracts seem unlikely to work.

To be sure, there may be scenarios in which negative information about one research path is relevant to other research paths. And in such scenarios, it's possible that tacit information is required to traverse the gap between the failed research path and other research paths. For example, research on the degree to which candidate treatment X binds to a biological target may help a scientist develop intuitions about the viability of other candidate treatments, even if they have not codified this information. If so, then the solutions described by Burstein (2012) and Gilson et al. (2009) may be viable. More generally, we have made here the simplifying assumption that negative information regarding one research path does not influence firms' assessments of the relative appeal of other research paths or their ability to achieve conclusive results when exploring other research paths—negative information is valuable in our model only because it reveals that one potential research path will in fact not work. That said, real-world situations will often depart from this assumption, which could make these other solutions more viable. Future work may profitably explore these possibilities.

Finally, as in the case of a research firm trying to protect its negative information from use or disclosure by its employees, a research firm that tries to sell negative information to other firms will confront the enforcement problem that drives Arrow's information paradox. Once a potential buyer learns of a failure, it can use the negative information by simply avoiding the disclosed dead-end. This reality presents a unique evidentiary challenge for proving misappropriation of negative information. Again, the use of positive information might be observed in a competitor's sales of products or processes. Conversely, the use of negative information materializes through the absence of action—that is, by avoiding a "blind alley." Proving that a company used negative information in order to not do something will often be exceptionally challenging—courts and juries may require a "smoking gun," such as a company memo stating that the company will abandon a particular research path that it might otherwise have pursued because of information it learned about another firm's failure. Uncertainty and inefficiency associated with such enforcement mechanisms would likely dissuade firms from engaging in direct, contract-based sales of negative information without some form of special legal rights to negative information.

4.2 Protective Instruments under Current Trade Secret Law

Turn now to trade secrets. As we'll see, both legal and practical limits render trade secret ineffective at aligning private incentives with social welfare in this context; as was the case with respect to contractual mechanisms, trade secret law offers only modest protection of negative information and cannot be used to obtain significant rewards for it.

Since the widespread adoption of the Uniform Trade Secrets Act by the states and the enactment of the Defend Trade Secrets Act by the federal government, trade secret law now indisputably encompasses negative information.[31] Because trade secrecy (nominally) protects negative information from misappropriation, trade secrets (nominally) provide research firms at least one of the features of

[31] Comment 5 to Section 1 of the Uniform Trade Secrets Act notes that "[t]he definition includes information that has commercial value from a negative viewpoint, for example the results of lengthy and expensive research which proves that a certain process will not work could be of great value to a competitor." (Uniform Trade Secrets Act, 2005). The UTSA has been adopted by forty-seven jurisdictions.

propertization that permits them to internalize the value of their failures: exclusivity. Rivals cannot benefit from negative information produced by the research firm insofar as a research firm has a right of action against rivals who misappropriate negative information.

To be sure, trade secret rights of action may not be effective because of the same sorts of enforcement costs that limit the efficacy of contracts. Observing a rivals' "use" of negative information will be challenging, and obtaining verifiable evidence of such "use" even more so. Still, compared to a scenario in which rivals could freely obtain and use negative information produced by a research firm, trade secret law reduces the externality associated with the production of negative information. To the extent that it does so, it reduces the informational benefit of waiting, inducing earlier entry than would otherwise be the case.

However, although trade secret law does not theoretically distinguish between positive and negative information, in practice, courts tend to enforce trade secret interests in negative information primarily in the context of departing employees. When they have done so, they have invoked the "inevitable disclosure" theory, a trade secret analogue to or substitute for contractual noncompete agreements. The inevitable disclosure theory provides that a former employee, in their new position, will inevitably use trade secrets belonging to their previous employer; injunctions are, from this perspective, appropriate to prevent the inevitable misappropriation of the previous employer's trade secrets.[32] Inevitable disclosure thus appears at

[32] See Avery Dennison Corp. v. Finkle, No. CV010757706, 2002 WL 241284, at *2 (Conn. Super. Ct. 2002) (issuing a permanent injunction on the inevitable disclosure theory of negative knowledge); International Business Machine, Corp. v. Seagate Technology, Inc., No. Civ. 3-91-630, 1991 WL 757821, at *2 (D. Minn. 1991) (issuing a preliminary injunction restraining an employee from starting a new job relying partially on the knowledge of failures); Novell, Inc. v. Timpanogos Research Group, Inc., 1998 WL 177721 (D. Utah 1998) (issuing an eighteen-month injunction restraining an employee from employment in the relevant competitive industry on theory it is inevitable the employee will use negative knowledge). Other cases have suggested recognition of a proprietary right to negative trade secrets but have not gone so far as to award damages solely on that basis. See, e.g., Courtesy Temporary Service, Inc. v. Camacho, 222 Cal. App. 3d 1278, 1287–88 (Cal. Ct. App. 1990) (stating that a customary list necessarily contains negative information about which customers not to call, and such negative research is a protectable trade secret); see also Graves (2007). But see Metallurgical Industries v. Fourtek, Inc., 790 F.2d 1195, 1202–03 (5th Cir. 1986) (finding the distinction between negative and positive knowledge "unavailing" and declining to recognize the existence of distinct, negative trade secrets).

first glance to give a research firm a mechanism to prevent rivals from accessing negative information by hiring the research firm's employees, just as (unlimited) noncompete agreements would.

But allowing research firms to obtain unlimited injunctions in this context would lead to perverse and conceptually unworkable effects (Graves, 2007), similar to those that have led legislatures to limit the scope of noncompete agreements. When an employee leaves one firm for another, the employee possesses both positive and negative information about efforts made by their previous employer. The previous employer could conceivably classify such negative information as proprietary trade secrets (Graves, 2007). If the former employer succeeds, however, the ex-employee would be unable to make full use of their expertise to benefit their new employer because people tend to develop skills and expertise by learning from mistakes; negative information derived from the prior employer's failures forms part of the foundation for the employee's general knowledge, skills, and experience. Indeed, extreme cases might force an employee to knowingly let their new employer repeat the former employer's failed research pursuits in order to avoid liability for misappropriating their former employer's negative information (Graves 2007). Rather than leading to socially optimal information sharing and research strategies, overreaching trade secret protections may simply enable anticompetitive and inefficient conduct.

To thread this needle between protecting a firm's negative information and allowing a former employee to make use of their general knowledge, skills, and experience, courts limit the application of the inevitable disclosure theory. Most pertinently for our purposes, court typically only issue injunctions on this basis for short durations—i.e., no more than a few months. Thus, again as noted in the context of noncompete agreements, the efficacy of the inevitable disclosure theory as a mechanism to protect negative information depends on whether firms will exhaust all research paths during the period of time in which an inevitable disclosure injunction might be in effect. For this reason, the real trade secret regime cannot provide the kind of complete protection described by the idealized trade secret regime in our model.

Moreover, trade secrets are unlikely to provide the other feature of propertization that research firms require in order to internalize the value of their failures: alienability. That is, trade secrecy will not be an effective tool for firms to obtain rewards for the negative information they produce. In addition to the transaction costs we have described in the context of purely contractual mechanisms, trade secret law denies protection to information that is generally known within an industry.[33] This puts a ceiling on the number of buyers with which a research firm can engage.

Indeed, in the later periods of our sequential-stage model, the sale of a trade secret to a rival might render the negative information entirely ineligible for protection. For a simple example, suppose there are ten firms in an industry. Firm 1 determines that candidate treatment X is toxic, and then sells this negative information to Firm 2. So far, trade secret law would continue to protect the information. Firm 2 might then sell this negative information about candidate treatment X to Firm 3, Firm 3 would eventually sell it to Firm 4, and so on. Once five or more firms know that candidate treatment X is toxic (and likely even before then), a court would conclude that the information is "generally known" in the industry and therefore ineligible for trade secret protection at all. Similarly, negative information might become publicly known through mandated government disclosures (e.g., failed FDA approval of a drug); in those cases, again, trade secret protection would terminate. The upshot is that firms will be unable to use trade secret law to internalize the full downstream value of the negative information they generate in earlier periods. As a result, even if trade secret law mitigated the timing problem firms face in the early periods of a race, it would not mitigate the timing problem in later periods.

Finally, the weakness of trade secrecy as a tool for rewarding firms that produce negative information means that, if it were the only mechanism available, firms would not account for the full social value of negative information that their research might produce. As a result, even if trade secrecy might somewhat

[33] UTSA §1(4)(i) (defining a "trade secret" as information that "derives independent economic value, actual or potential, from not being generally known . . .").

mitigate the timing problem, we do not expect that trade secrecy would do much to mitigate the suboptimal investment problem we describe.

4.3 The Treatment of Negative Information in the Patent System

The patent system offers research firms robust protection and rewards for creating and disclosing positive information. Patents are available to any firm that "invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof" (35 U.S.C. § 101). A firm that obtains a patent can then enjoin and obtain damages from anyone who "makes, uses, sells, or offers to sell" the patented invention without authorization (35 U.S.C. § 271). In addition, because they have "the attributes of personal property," patents can readily be sold or licensed (35 U.S.C. § 261). And because protection does not end when information becomes widely known, there is no inherent limit to the number of potential buyers of patented positive information. Firms that generate positive information regarding a "new and useful process, machine, manufacture, or composition of matter," accordingly often use patents to obtain both protection and rewards.

Patents cannot, however, be so readily used to internalize the value of negative information. Begin with scenarios like the one we model—that is, suppose that patents issue only at the end of a race, when a firm discovers positive information that has relatively immediate exploitable value. In these scenarios, the utility doctrine would prevent firms from using patents to protect and obtain rewards for the negative information they produce along the way. The utility doctrine requires that "the invention be operable to achieve useful results" (*In re Swartz*, 232 F.3d 862, Fed. Cir. 2000). Negative information cannot satisfy this requirement because it reveals what isn't operable, rather than what is.

However, patents can sometimes be obtained while the race is still ongoing. In such scenarios, firms might be able to combine patents with other mechanisms in order to protect and obtain rewards for the negative information they generate before any firm discovers the corresponding positive information.[34]

Suppose, for example, that a startup firm synthesizes a new compound but does not yet have convincing evidence demonstrating that it will be safe and effective in humans. Although the utility doctrine prevents a patent from issuing when the applicant has no plausible use for a new compound—that is, when the patent can be understood as a mere "hunting license" (*Brenner v. Manson*, 383 U.S. 519, 1966)—the utility doctrine is not so strict as to require the sorts of safety and efficacy evidence that the FDA demands. Instead, all that is required to satisfy the utility doctrine is evidence that, in the view of a skilled artisan, reasonably supports the applicant's assertions regarding the ultimate therapeutic uses (*In re Brana*, 51 F.3d 1560, 1995). Thus, so long as the startup firm can point to, say, *in vitro* results supporting the compound's efficacy as a treatment for some condition, it should be able to patent the compound. This patent on the compound can be understood as the start of the race.

Now, with this patent in hand, the firm can conduct further research to determine whether the compound will in fact be an effective treatment for a particular condition. Suppose that those trials produce negative information—that is, they reveal that the compound is toxic or ineffective. Although the startup no longer has enough capital to continue researching other alternatives, because it has patented the compound, it is not without recourse. Instead, it can sell the patent to another firm to test whether the compound is an effective treatment for another condition (as in the exit scenario we describe).

Importantly, the negative information the startup generated about the compound's use for a particular condition can be incorporated into the buyer's purchase of the patent on the compound. First, the (now

[34] Our focus is on the incentives that firms confront when deciding whether to invest in research that could generate negative information. For an interesting proposal to create a new mechanism to capture and disseminate whatever negative information happens to be produced, see Seymore (2012).

failing) startup can reveal the structure of the patented compound to the buyer. Because the patent on the compound prevents all others from using the compound for any purpose without authorization, this disclosure does not create a risk of appropriation by the buyer. The seller can then offer to include the negative information as "know how" that will accompany the sale of the patent. Such an offer would likely be structured as a staged disclosure of the sort described by Burstein (2012), in which the seller provides some information about the failed research path that still stops short of providing the buyer all the information it needs in order to avoid the dead-end. Because there is still the possibility of discovering positive information, though, there is still enough here to motivate the transaction, with the seller only completing disclosure of the failed research path after the buyer has agreed to purchase it. In addition, the seller would likely have significant bargaining power because of the patent on the compound. As a result, when the utility doctrine permits the issuance of a patent early in a race, we might have a scenario that resembles the "negative patent" one described in Section 3.4.[35]

Returning to more common scenarios in which the utility doctrine prevents firms from obtaining patents until the end of the race, patent law may nonetheless influence firms' decisions whether to disclose negative information. We have assumed that the first firm to generate positive information obtains a patent that enables them to appropriate the full social value of the invention, V. But patents are not guaranteed, even for the discovery of positive information. In particular, patents issue only for inventions that would have been "nonobvious" to a "person having ordinary skill in the art" (*KSR International Co. v. Teleflex, Inc.*, 550 U.S. 398, 2007).

[35] In addition, patents can be issued on the basis of credible predictions about how an invention will work, even though the applicant has not actually verified those predictions (Freilich 2019). In a considerable number of cases, those predictions turn out to be incorrect—patents therefore sometimes claim inventions that turn out not to work as described. And as Seymore (2018) and Kapczynski and Syed (2013) have explained, firms lack adequate incentives to produce and disclose negative information that is generated after the filing of a patent application.

Importantly, disclosures of negative information in the earlier stages of the race can influence the likelihood that the eventual winner is able to overcome this nonobviousness hurdle. Courts assess several factors, known as "secondary considerations," to determine whether an invention is nonobvious. Two of those have particular relevance in the context of negative information: (1) "failure of others"; and (2) "obvious to try." As we'll see, these two factors push firms in opposite directions.

The "failure of others" factor provides that an invention is more likely to be nonobvious when an inventor achieves a goal that others in the field have tried, but failed, to accomplish (*Graham v. John Deere*, 383 U.S. 1, 18, 1966). The intuition underlying this factor is that, if it were indeed obvious for a person of skill in the art to produce the invention, then prior efforts to do so should have succeeded. Thus, reports that others before the inventor had tried, but failed, to produce the invention "serv[e] as a simulated laboratory test of the obviousness of the solution to a skilled artisan" (*Symbol Technologies., Inc. v. Opticon, Inc.*, 935 F.2d 1569, 1578–79, Fed. Cir. 1991).

Plainly, this factor encourages firms to conceal negative information. This is because a failing firm's disclosure of its negative information would reveal that an individual (putatively) of ordinary skill in the art had tried, but failed, to achieve the invention. As a result, such a publication would incrementally increase the odds that the ultimately successful firm will be able to patent her invention. To avoid that result, the failing firm should conceal its failure.

Past a certain point, however, the "obvious to try" test pushes in the opposite direction. That test provides that an invention is likely to be obvious when there are "a finite number of identified, predictable solutions" and a person of ordinary skill has "a good reason to pursue the known options" (*KSR International Co. v. Teleflex, Inc.*, 550 U.S. 398, 2007). Nevertheless, an invention is unlikely to be obvious when an "inventor would have had to try all possibilities in a field unreduced by direction of the prior art," such that she would have had to "vary all parameters or try each of numerous possible choices until one possibly arrived at a successful result" (*Bayer Schering Pharma AG v. Barr Laboratories., Inc.*, 575 F.3d 1341, Fed. Cir. 2009).

Thus, in order for an invention to be "obvious to try," the "number of options" available must have been "finite, and, in the context of the art, small or easily traversed." (*Ortho-McNeil Pharmaceutical, Inc. v. Mylan Laboratories*, 520 F.3d 1358, Fed. Cir. 2008).[36]

For a research firm that discovered negative information, the implications are as follows. Suppose that, at the outset, there are many research paths available. Disclosure of negative information about one of those paths would likely be deemed evidence of "failure of others." That said, subsequent disclosures will continue to reduce the "number of options" remaining; eventually, the number of remaining paths will be "small or easily traversed," such that it is "obvious to try" them. At that point, the ultimate success will be obvious and therefore unpatentable. As a result, these factors might suppress the disclosure of negative information early in a race but encourage it later in the race.[37]

For similar reasons, the "failure of others" and "obvious to try" factors will have opposite impacts on the lost opportunity effect. Recall that the lost opportunity effect encourages entry because waiting means that a firm might find the successful research path and obtain the corresponding patent before others even have a chance to participate in the race; in order to have the chance of winning, firms enter earlier. But the successful firm is not guaranteed a patent, and the "failure of others" factor increases the likelihood that a patent will issue as the number of prior failures increases. As a result, if later opportunities actually

---

[36] An interesting example is Wm. Wrigley Jr. Co. v. Cadbury Adams USA LLC, 683 F.3d 1356 (Fed. Cir. 2012). The patent in that case claimed a chewing gum comprising, *inter alia*, a combination of menthol and a substance named "WS-23." (Wm. Wrigley Jr. Co. v. Cadbury Adams USA LLC, 683 F.3d 1356, 1362, Fed. Cir. 2012). The court held that the combination was invalid as obvious in part because of a prior art reference that reported the results of testing 1200 compounds to determine which might have a "synergistic effect" with menthol (Wm. Wrigley Jr. Co. v. Cadbury Adams USA LLC, 683 F.3d 1356, 1362–63, Fed. Cir. 2012). One of the successful compounds reported there was WS-23 (Wm. Wrigley Jr. Co. v. Cadbury Adams USA LLC, 683 F.3d 1356, 1362–63, Fed. Cir. 2012). Still, the mere testing of 1200 compounds may have been sufficient to render the remaining handful "obvious to try" even if that reference had not reported the successful combination of menthol and WS-23 (Wm. Wrigley Jr. Co. v. Cadbury Adams USA LLC, 683 F.3d 1356, 1362–63, Fed. Cir. 2012). The point is that reducing the number of research paths by disclosing one firm's failures can render an ultimate success by another firm unpatentable.

[37] To be sure, the failing firm would only actually disclose if the gain to it from reducing its competitors' odds of obtaining a patent is greater than the gain to it from forcing its competitors to devote resources to the exploration of a path that the failing firm knows is a blind alley.

materialize, they will be more valuable than earlier ones because they will carry better odds that the invention will overcome the obviousness hurdle. The “failure of others” factor might therefore delay entry.

On the other hand, the “obvious to try” factor might somewhat amplify the lost opportunity effect and accelerate entry in the later stages of the race. As the number of remaining paths dwindle, it becomes increasingly likely that any resulting success is unpatentable because it is obvious to try the ultimately successful one. The upshot is that the likelihood of a patent issuing upon the production of positive information will drop rapidly towards the end of the race. Later opportunities would accordingly be less valuable than earlier ones. In combination with the “failure of others” factor, this suggests that there is a “sweet spot” for entry—neither too early, nor too late.

Finally, although our model describes a failing firm as exiting the market and a field in which there is only one successful option, it’s worth considering here a situation in which a failing firm can continue researching after failure and there is more than one successful option. If a firm initially fails but ultimately succeeds, disclosure of the negative information it generated along the way would be socially desirable insofar as we would not want to waste resources pursuing those dead-ends in a futile effort to design around the patent.

The enablement doctrine may have this desirable effect. That doctrine demands that the patent’s specification “enable any person skilled in the art to make and use the invention” (*Amgen v. Sanofi*, 598 U.S. 594, 2023). To be sure, negative information tells a skilled artisan *what not to do*; standing alone, negative information does not tell a skilled artisan *what she must do* in order to make the invention and use it herself.

However, as seen most recently in *Amgen v. Sanofi* (598 U.S. 594, 2023), enablement may require that applicants disclose some negative information alongside the traditional disclosure of positive information. In that case, the Supreme Court evaluated a patent claiming a class of antibodies. The claims were drawn in functional, rather than structural, terms—that is, they defined the class of antibodies in terms of their ability to bind to a naturally occurring protein, rather than in terms of their amino acid sequences. Amgen’s patent did, however, describe the amino acid sequence for 26 antibodies that performed the required functions; in our terms, Amgen provided positive information regarding 26 antibodies. However, there were literally millions of unexplored research paths—that is, there were millions of other antibodies whose animo acid sequences made them plausible candidates for the functionally-defined class, but that had not yet been evaluated for compliance with the functional claim limitations. Moreover, there was nothing in the patent specification or the state of the art that indicated which antibodies actually performed the required function of binding to the naturally occurring protein. As a result, a skilled artisan was left "to create a wide range of candidate antibodies and then screen each to see which happen to bind to PCSK9 in the right place and block it from binding to LDL receptors"—a laborious task of trial-and-error.

From our perspective, this is (in part) a problem of negative information. The problem with Amgen's patent was that it left too many research paths unexplored in a setting in which there was reason to think that most paths would result in failure. Had Amgen found and disclosed sufficient negative information—that is, if Amgen had explored enough of the potential research paths to know which ones would succeed and which would fail—it could have satisfied the enablement requirement by, for example, "identifying] a quality common to every functional embodiment." (*Amgen v. Sanofi*, 598 U.S. 594, 2023). In other words, Amgen could have narrowed the scope of the claims by disclosing the negative information consisting of categories of candidate antibodies that could not achieve the functional limitations required by the claims; had it done so, it may well have been able to satisfy the enablement requirement.[38] The upshot is that a robust application of *Amgen*'s interpretation of the enablement doctrine may demand disclosure of negative information alongside positive information, at least by firms that ultimately discover a successful path.

[38] See Amgen v. Sanofi, 598 U.S. 594 (2023) (explaining that "the more a party claims, the broader the monopoly it demands, the more it must enable"). The Court also acknowledged that Amgen's specification adequately enabled the 26 antibodies it identified by amino acid sequence. The problem was that the claims swept "much broader than those 26 antibodies" because it included a large swathe of unexplored amino acid sequences with no indications of which ones would work and which ones would fail.

4.4 Toward a System of "Negative Intellectual Property"?

Although real-world evidence of negative information transfers is limited, the sale of negative trade secrets may be possible provided that (1) the information has adequate legal protection; and (2) the knowledge in question can be accurately defined. Because the Uniform Trade Secrets Act has been enacted in 49 states—and the Defend Trade Secrets Act, which mimics the relevant portions of the UTSA, has been enacted by the federal government—it is now clear that trade secret law encompasses information relating to failed research paths.[39] Still, as explained above, the ordinary trade secret regime is inadequate to the task of facilitating transfers of negative information. Instead, a more structured market may have to be designed and implemented. We sketch out here some crucial details of that market.

First, transferable information must be defined sufficiently precisely so as to allow the buyer and seller to accurately assess its scope and value ex ante. As one possibility, a company could use publicly available information to create discrete categories of information. Ideally, this would provide transacting companies with an adequate foundation for determining the scope and value of potential information. If the value of information cannot be determined ex ante, companies might preemptively agree to an instrument for deferred valuation—for instance, the purchaser firm could agree to a compulsory license or to pay royalties (based on a predetermined time limit) for products developed from the transferred information. Although companies would still have to estimate the prospective value of negative information probabilistically, such mechanisms could help to reduce uncertainty.

[39] The Conference comment to the UTSA states that its "definition includes information that has commercial value from a negative viewpoint, for example the results of lengthy and expensive research which proves that a certain process will *not* work could be of great value." UTSA Comment § 1 (1985). *See also* Pace (1995).

Second, to address search and bargaining costs, a market for negative information could benefit from third-party intermediaries who facilitate exchange. A third-party intermediary could help avoid the bargaining costs associated with Arrow's information disclosure paradox—through mutual disclosure of the contents of negative information and a buyer's interests, the third-party consultant could determine the value of the information to the buyer. With adequate confidentiality agreements in place, the information would not be disclosed to the buyer until they agreed to purchase the information. Thus, the third-party intermediary helps to eliminate the danger of a prospective buyer making use of the negative information before paying the seller. For such a system to succeed, the third-party intermediaries would need to establish reputations for fair, accurate, and discrete valuations.

A possible solution to search costs could consist of private intermediaries operating a negative trade secret posting database that is analogous to real estate listings. First, potential sellers could post circumstantial information relating to the negative information's value, including the cost of research, time spent, and generalized goal of the research. When a buyer identifies a potential match, the third-party consultant could step in to more conclusively evaluate whether the information matches the buyer's interests and to determine the value of the information. Ultimately, this could help companies avoid needing to guess what others are actively researching.

4.5 Combining IP and Non-IP Mechanisms

Finally, our analysis concluded that subsidies are most likely to lead to optimal timing of and investment in research that can produce negative information. Those subsidies can take various forms, including prizes, grants, and tax credits; moreover, depending on the form, they could be provided by government or private actors. And more generally, innovation policy levers like patents, trade secrets, prizes, grants, and tax credits perform two distinct functions: they operate as both (1) innovation incentives; and (2) allocation mechanisms (Hemel and Ouellette, 2019).[40] Depending on the circumstances, the optimal policy mix may comprise "matching" the innovation incentive of an IP policy lever (i.e., patent or trade secret) with the allocation mechanism of a non-IP policy lever (i.e., prizes, grants, or tax credits), or vice versa. Policy makers might also consider "mixing" IP and non-IP policy levers on the same side of the innovation-incentive/allocation-mechanism divide.

The case for "mixing" in this context seems strong because both government-set and market-set rewards face serious obstacles as innovation incentives for the production of negative information. We suspect that it will be difficult for the government to ascertain the social value of any given piece of negative information. Precisely because negative information does not directly translate into a good or service, it is challenging to know just how much it is worth. Negative information has value primarily because it helps firms minimize opportunity costs. As a result, the most relevant considerations for valuing negative information relate to which other paths the firm might have taken and how efficiently the firm's investments translate to conclusive results (that is, how likely the firm is to reach either a positive or negative result for a given level of effort). This is information that is uniquely in the hand of private actors. Market-set rewards, like patents and trade secrets, thus seem preferable to government-set ones.

On the other hand, we have described the serious obstacles that research firms face when attempting to transact over negative information. These obstacles suggest that market signals will be poor proxies for the social value of negative information. This would counsel in favor of using grants, in-house research at federal agencies, and government-set prizes as innovation incentives (Hemel and Ouellette, 2019). In

[40] Thus, for example, patents provide (1) an innovation incentive in the form of "an ex post reward" that uses "a market-generated estimate of social value" and (2) an allocation mechanism of monopoly pricing by virtue of the patentee's right to exclude (Hemel and Ouellette, 2019). Conversely, grants provide (1) an innovation incentive in the form of an ex ante reward that uses a government-set estimate of social value and (2) an allocation mechanism of open-access. This is the case for grants standing alone insofar as they do not themselves restrict use of the information generated by using the grant money. That said, in many instances, grants are combined with other innovation policy tools such that the actual allocation mechanism is not open-access. For example, the Bayh-Dole Act of 1980 (94 Stat. 3015) permits universities to obtain patents on the results of federally-funded research.

addition, because effort influences the likelihood of achieving results, ex post rewards like prizes would be preferred over ex ante rewards like grants (Hemel and Ouellette, 2019).[41] That returns to our previous point: because the value of negative information depends on the respective values of the other paths a firm might explore instead of the dead end, governments will have a hard time setting appropriate prize amounts. Because both government-set and market-set rewards are noisy signals of social value and are unlikely to be correlated with each other, there is therefore good reason to prefer mixing of innovation incentives (Hemel and Ouellette, 2019).

Meanwhile, tax credits (unless fully refundable) require that the overall firm be profitable. In our models, the research firm cannot be profitable because it exits upon failure. Indeed, more generally, the production of negative information is unlikely to lead directly to profitability—research failures, by definition, do not themselves generate viable commercial products. And while one could imagine a fully refundable tax credit that reimburses failing firms for all or a substantial portion of the costs of their failure, it seems unlikely that such a policy would be politically viable. Government payouts in these contexts may resemble the sorts of politically unpopular bailouts that followed the 2008 financial crisis, only without the purported advantage of preventing widespread economic harm that accounted for the political viability of those measures. Nevertheless, future work that explores in depth the possibility that research firms will remain in the market after failure or scenarios in which negative information is complementary with positive information ought to consider the use of tax credits as innovation incentives.

[41] We note in passing that three conditions that tend to favor ex ante rewards—capital constraints, risk-aversion, and present-bias—do not seem to be significant in these contexts, or at least not any more significant with respect to negative information than positive information.

With respect to allocation mechanisms, none of the considerations that counsel in favor of IP-based allocation appear to be present for negative information (Hemel and Ouellette, 2019).[42] It thus seems likely that open-access is the preferred allocation mechanism for negative information.

In sum, optimal innovation policy likely entails mixing innovation incentives—some combination of IP-like exclusivity from trade secrets and ex post rewards like prizes—matched with a non-IP allocation mechanism—open-access. The challenge is to develop a mechanism that enables the free distribution of negative information protected under a trade secret regime while providing adequate compensation to the research firm that produced the negative information. Patent buyout schemes like those proposed by Kremer (1998) may serve as useful models. Kremer's particular design is not promising here insofar as the auction would likely require disclosure of the negative information, raising the specter of free-riding by participants who lose the auction because of the enforcement challenges we identify. Still, it seems desirable to attempt to leverage private information to determine the price at which the government might buy-out the negative-information trade-secret holder, and thereby place the negative information in the public domain. We leave such efforts for future work.

5. Conclusions

In this paper, we examined the distortions created by the lack of rewards obtainable for negative research results. Whereas academic researchers may be encouraged to share negative information by the opportunity

[42] Unlike luxury goods, there's no general reason to think that a user-pays system for negative information is desirable from a distributive perspective. Nor does consumption of negative information generate negative externalities, as in the case of overconsumption of antibiotics, such that restricting consumption via monopoly pricing would be attractive. And, finally, negative information cannot itself be commercialized; the sorts of collective action problems that undergird the "commercialization" rationale for patents thus appear to be absent (Kitch, 1977; Sichelman, 2010). While regulation or Pigouvian taxation are more efficient ways to control access, monopoly pricing may be more viable for political economy reasons (Hemel and Ouellette, 2019).

to publish research in dedicated journals, private firms face a starker incentive to avoid giving rivals a competitive edge through disclosure of negative information. However, firms face multiple, severe challenges to using contracts to suppress the flow of negative information; accordingly, trade secrecy's protection for negative know-how looms large in this context. To foster a socially optimal level of research, markets and legal remedies should provide incentives to pull firms out of the trade secret regime, encouraging disclosure of information that firms might otherwise keep secret. The firms' inability to capture the full social value of negative information weakens firms' incentives to invest in high-expected-return projects that also entail multiple likely dead ends in research. In these situations, the reward structure of the patent regime and particular patent doctrines (including utility and enablement) are unable to serve as effective mechanisms to generate optimal levels of research investments in a deconcentrated market structure. The asymmetric treatment of successes and failures generates suboptimal levels of research investments in several common situations. Extensions of our model should examine the robustness of our results by relaxing some of the simplifying assumptions of our model. Here are a few points worthy of investigation:

1. Our model assumes that all unexplored paths are equally likely and that researchers have no ability to reduce uncertainty on the success probability of any given research path. This simplifying assumption overlooks that fact that acquiring information about someone else's failure may also improve the firms' ability to estimate the likelihood that any given remaining research path will succeed (e.g., information that one compound doesn't bind to a target should allow firms to update their beliefs as to the likelihood that structurally similar compounds will bind to that target).

2. In our model, we focused on the value of failure, setting aside the interrelated value of learning. As a result of our modeling focus, our function $p(e)$ is not sensitive to information about prior research failures. This simplifying assumption could be relaxed to consider situations where prior research failures also generate learning (e.g., skills or techniques that increase new firms' ability to perform their own research).

A firm's ability to reach a conclusive result (positive or negative) would require less effort as the stock of learning increases: in nascent fields, it might be relatively more difficult to achieve successful outcomes than in mature ones. Our policy results would need to be recalibrated in light of this extension, accounting for the fact that failures reveal something about the underlying reality that improves firms' ability to ultimately succeed.

3. Further analysis is also warranted before drawing any specific normative conclusion on the desirability of specific policy instruments. For example, our model assumes that research proceeds sequentially, such that each new entering firm can acquire information from prior failing firms, but this sequential process cannot operate when firms engage in parallel research. In these situations, a subsidy regime would outperform a negative patent regime, since it would render the negative information available to all parallel competitors, while negative patents would transfer the negative information only to one or a few firms, excluding all other parallel researchers from it.

4. Our analysis focused on the effects of information externalities from failed research on the timing of and level of investment in productive research. Further explorations of this topic should consider ways in which these effects interact with other strategic incentives that may counterbalance one or both of these effects. Specifically, our results should be evaluated in conjunction with the results of recent literature on patent races, which suggest that the incentives created by the current patent system may lead to an excessive rate of research and innovation. The delayed research that our paper identifies may mitigate or offset the excessive rate of research identified in the patent race literature. The understanding of which of these two effects (i.e., the "race effect" or the "procrastination effect") will dominate in real life situations would truly benefit from additional theoretical and empirical research.

References

[1] Aghion, Philippe, and Richard Holden. 2011. "Incomplete Contracts and the Theory of the Firm: What Have We Learned over the Past 25 Years?" 25 *Journal of Economic Perspectives* 181–97.

[2] Arora, Ashish and Robert Merges. 2004. "Specialized Supply Firms, Property Rights, and Firm Boundaries," 13 *Industrial and Corporate Change* 451-475.

[3] Arrow, Kenneth. 1962. Economic Welfare and the Allocation of Resources for Invention, in R.R. Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton: Princeton University Press.

[4] Bar-Gill, Oren, and Gideon Parchomovsky. 2009. "Law and the Boundaries of Technology Intensive Firms," 157 *University of Pennsylvania Law Rev.* 1649–89.

[5] Baum, Joel A.C., and Kristina B. Dahlin. 2007. "Aspiration Performance and Railroads' Patterns of Learning from Train Wrecks and Crashes," 18 *Organization Science* 368–85.

[6] Baumard, Philippe, and William H. Starbuck. 2005. "Learning from Failures: Why it May Not Happen," 38 *Long Range Planning* 281–98.

[7] Benkler, Yochai. 2002. "Coase's Penguin, or Linux and the Nature of the Firm," 112 *Yale Law Journal* 369–446.

[8] –––. 2012. "Open Wireless vs. Licensed Spectrum: Evidence from Market Adoption," 26 *Harvard Journal of Law and Technology* 69–163.

[9] – – –. 2017. Open Access and Information Commons, in F. Parisi, ed., *Oxford Handbook of Law & Economics*. Oxford: Oxford University Press.

[10] Bertello, Alberto, Paola De Bernardi, and Francesca Ricciardi. 2023. "Open Innovation: Status Quo and Quo Vadis - An Analysis of a Research Field," *Review of Management Science*. https://doi-org.ezproxy.unibo.it/10.1007/s11846-023-00655-8.

[11] Burstein, Michael J. 2012. "Exchanging Information Without Intellectual Property," 91 *Texas Law Review*, 227–82.

[12] Chesbrough, Henry, and Marcel Bogers. 2014. Explicating Open Innovation: Clarifying an Emerging Paradigm for Understanding Innovation, in H. Chesbrough, W. Vanhaverbeke, and J. West, eds., *New Frontiers in Open Innovation*. Oxford: Oxford University Press, 3-28.

[13] Choi, Albert, and George Triantis. 2012. "The Effect of Bargaining Power on Contract Design, 98 *Virginia Law Review* 1665–1743.

[14] Chuang, You-Ta and Joel A.C. Baum. 2003. "It's All in the Name: Failure-Induced Learning by Multiunit Chains," 48 *Administrative Science Quarterly* 33–59.

[15] Denicolo, Vincenzo and Luigi A. Franzoni. 2003. "The Contract Theory of Patents," 23 *International Review of Law & Economics* 365-380

[16] Denicolò, Vincenzo and Luigi A. Franzoni. 2012. "Weak Intellectual Property Rights, Research Spillovers, and the Incentive to Innovate," 14 *American Law and Economics Review* 111-140.

[17] Denrell, Jerker. 2003. "Vicarious Learning, Undersampling of Failure, and the Myths of Management," 14 *Organization Science* 227–43.

[18] Dubben, Hans-Herman, and Hans-Peter Beck-Bornholdt. 2008. "Systematic Review of Publication Bias in Studies on Publication Bias," 331 *British Medical Journal*, 433–35.

[19] Dunlop, John T., and Benjamin Higgins. 1942. "'Bargaining Power' and Market Structures," 50 *Journal of Political Economy* 1–26.

[20] Francis, John, and CongCong Zheng. 2010. "Learning Vicariously from Failure: The Case of Major League Soccer and the Collapse of the North American Soccer League," 35 *Group & Organization Management* 542–71.

[21] Freilich, Janet. 2019. "Prophetic Patents," 53 *U.C. Davis Law Review* 663–731.

[22] Franco, Annie, Neil Malhotra, and Gabor Simonovits. 2014. "Publication Bias in the Social Sciences: Unlocking the File Drawer," 345 *Science* 1502–05.

[23] Franzoni, Luigi A. 2016. "Trade Secrets Law." In: *Encyclopedia of Law and Economics*. New York : Springer.

[24] Franzoni, Luigi A. and Kaushik, Arun K. 2016. "The optimal scope of trade secrets law." 45 *International Review of Law and Economics* 45-53.

[25] Graves, Charles T. 2007. "The Law of Negative Knowledge: A Critique," 15 *Texas Intellectual Property Law Journal* 387–416.

[26] Gilson, Ronald J., Charles F. Sabel, and Robert E. Scott. 2009. "Contracting for Innovation: Vertical Disintegration and Interfirm Collaboration," 109 *Columbia Law Review* 431–502.

[27] Graham, Stuart J.H. and Merges, Robert P. and Samuelson, Pamela and Sichelman, Ted M. 2009. "High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey," 24 *Berkeley Technology Law Journal* 4: 255-327.

[28] Haunschild, Pamela R., and Bilian Ni Sullivan. 2002. "Learning from Complexity: Effects of Prior Accidents and Incidents on Airlines' Learning," 47 *Administrative Science Quarterly*, 609–43.

[29] Hemel, Daniel J., and Lisa L. Ouellette. 2019. "Innovation Policy Pluralism," 128 *Yale Law Journal* 544–614.

[30] Holden, Richard T., and Anup Malani. 2014. "Renegotiation Design by Contract," 81 *University of Chicago Law Review* 151–78.

[31] Holgersson, Marcus, and Ove Granstrand. 2021. "Value Capture in Open Innovation Markets: The Tole of Patent Rights for Innovation Appropriation," 25 *European Journal of Innovation Management* 320–39.

[32] Hrdy, Camilla A., and Christopher B. Seaman. 2024. "Beyond Trade Secrecy: Confidentiality Agreements That Act Like Noncompetes," 133 *Yale Law Journal* 669–754.

[33] Kapczynski, Amy and Talha Syed. 2013. "The Continuum of Excludability and the Limits of Patents," 122 *Yale Law Journal* 1900–1963.

[34] Kitch, Edmund W. 1997. "The Nature and Function of the Patent System," 20 *Journal of Law and Economics*, 265–90.

[35] Lant, Theresa K., and David B. Montgomery. 1987. "Learning From Strategic Success and Failure," 15 *Journal of Business Research* 503–18.

[36] Lapré, Michael A., and Nikos Tsikriktsis. 2006. "Organizational Learning Curves for Customer Dissatisfaction: Heterogeneity Across Airlines," 52 *Management Science* 352–66.

[37] Levitt, Barbara, and James G. March. 1988. "Organizational Learning," 14 *Annual Review of Sociology* 319–38.

[38] Levmore, Saul. 2013. "The Impending iPrize Revolution in Intellectual Property Law," 93 *Boston University Law Review* 139-62.

[39] Marcus, Alfred A., and Mary L. Nichols. 1999. "On the Edge: Heeding the Warnings of Unusual Events," 10 *Organization Science* 482–99.

[40] Marshall, Aarian. 2018. "Uber and Waymo Abruptly Settle for $245 Million," *WIRED*. Available at https://www.wired.com/story/uber-waymo-lawsuit-settlement/.

[41] Merges, Robert P. 1999. "The Law and Economics of Employee Inventions," 13 *Harvard Journal of Law & Technology*.

[42] Merges, Robert P. 1993. "Uncertainty and the Standard of Patentability," 7 *(Berkeley) High Technology Law Journal* 1-70.

[43] Nash, John F. (1950). "The Bargaining Problem," 18 *Econometrica* 155–62.

[44] Pace, Christopher R.J. 1995. "Case for a Federal Trade Secrets Act," 8 *Harvard Journal of Law & Technology* 427–70.

[45] Seymore, Sean B. 2019. "Patenting Around Failure," 166 *University of Pennsylvania Law Review* 1139–1187.

[46] Seymore, Sean B. 2012. "The Null Patent," 53 *William & Mary Law Review* 2041–2105.

[47] Sichelman, Ted. 2010. "Commercializing Patents," 62 *Stanford Law Review* 341–411.

[48] Sitkin, Sim B. 1992. "Learning Through Failure: The Strategy of Small Losses," 14 *Research in Organizational Behavior* 231–66.

[49] Stanley, Tom D. 2005. "Beyond Publication Bias," 19 *Journal of Economic Surveys* 309–45.

[50] –––. 2005. “Meta-Regression Methods for Detecting and Estimating Empirical Effects in the Presence of Publication Selection,” 70 *Oxford Bulletin of Economics and Statistics* 103–27.

[51] Tarde, Gabriel (1903). *The Laws of Imitation*, E.C. Parsons, trans. New York: Henry Holt. (Original work published 1880).

[52] Terlaak, Ann, and Yan Gong. 2008. “Vicarious Learning and Inferential Accuracy in Adoption Processes,” 33 *Academy of Management Review* 846–68.

[53] The Economist Group. 2022. “The Open Innovation barometer. Assessing the Collaborative Spirit of a World in Flux.” Available at https://impact-economist-com.ezproxy.unibo.it/projects/open-innovation/Open%20Innovation%20Briefing%20Paper.pdf (last accessed 08/02/2023)