J Evol Econ (2016) 26:271–290
DOI 10.1007/s00191-015-0439-z

REGULAR ARTICLE

Learning from successes and failures in pharmaceutical R&D

Jing-Yuan Chiou[1] · Laura Magazzini[2] · Fabio Pammolli[3] · Massimo Riccaboni[3,4]

Published online: 6 January 2016

Abstract In this paper, we build a cumulative innovation model to understand the role of both success and failure in the learning dynamics that characterize pharmaceutical R&D. We test the prediction of our model by means of a unique dataset that combines patent information with R&D projects, thus distinguishing patents related to successfully marketed products from those covering candidate drugs that failed in clinical trials. Results confirm model predictions showing that patents associated with successfully completed projects receive more citations than those associated with failed projects. However, we also show that failed projects can be in turn cited more often than patents lacking clinical or preclinical information. We further explore the 'black box' of innovation, providing evidence that both successes and failures contribute to R&D investment decisions and knowledge dynamics in science-driven sectors.

Keywords R&D competition · Pharmaceutical industry · Knowledge flows · Technological trajectories

JEL Classification D23 · D83 · O34

✉ Massimo Riccaboni
massimo.riccaboni@imtlucca.it

[1] Department of Economics, National Taipei University, Taipei, Taiwan

[2] Department of Economics, University of Verona, Verona, Italy

[3] IMT Institute for Advanced Studies, Lucca, Italy

[4] Department of Managerial Economics, Strategy and Innovation, K.U. Leuven, Leuven, Belgium

1 Introduction

Innovation is a cumulative process, where researchers learn from past results and adjust their behavior based on previous experience (Herriot et al. 1985; Cyert and March 1992). Even though inventors tend to privilege learning from successes, also driven by the strong emphasis on positive outcomes in scientific journals and the business press (Denrell 2003), theoretical and empirical contributions highlight the essential role of failures in learning (Levitt and March 1993; Levinthal and March 1993; Teerlak and Gong 2008).

Against this background, in this paper we build a cumulative innovation model where both successful and failed research project may provide useful information for future research. Differently from the standard single-research-line model, such as Green and Scotchmer (1995) and O'Donoghue (1998), we consider a scientific space consisting of multiple research lines (Aghion et al. 2008), or technological paths (Dosi 1982). Several research lines are possible candidates to reach successfully a final invention (i.e. to develop and launch a drug in the market); that is, different research lines compete with each other for success, and it may well happen that no research line can lead to the final invention. In this case, a piece of "positive" information that raises the successful likelihood of a line of research will at the same time reduce that of the other possible paths. Similarly, "negative" information, e.g., when a research project on some line of research fails, although reducing the successful probability of the experimented research line, will simultaneously raise the prospects of other lines of research. In a simple way, this captures the idea that both success and failure convey useful information for subsequent innovation.

Learning by the experience of others is a fundamental aspect of most cognitive and choice activities (Levitt and March 1993). It is at the very heart of phenomena such as diffusion processes (Stoneman 1981; Chatterjee and Eliashberg 1990; Arthur and Lane 1993), trust and reputation building (Cabral 2005). Recently, more attention has been paid to the role of negative and positive feedback in the evolution of technologies and markets. We contribute to this literature by considering a case where multiple research lines compete for success, and both good and bad news contribute to the evolution of knowledge about the best technological solution available.

We argue that, in high-uncertainty technological domains, such as the biopharmaceutical industry, firms build their product development strategies both on success and failure, analogously to the economic value of findings and nonfindings spanning from basic research (David et al. 1992). However, since open science is strongly biased toward the publication of positive results, patents and clinical trials play a fundamental role in pharmaceuticals to disclose R&D failures information (Magazzini et al. 2009). Empirical analysis in the pharmaceutical domain reveals that, indeed, technological competencies are accumulated building both on successfully developed compounds as well as on failures, i.e. compounds that do not pass through all the stages involved in drug development, due, for example, to lack of effectiveness or toxicological effects (Magazzini et al. 2012).

Recently, dedicated incentive schemes such as a market for R&D failures (Shalem and Trajtenberg 2009) have been implemented to enhance the transmission of the information regarding failures among competitors. Following the lead of the

UK's Medical Research Council (MRC), the National Institute of Health (NIH) launched in 2012 an initiative to bring academic scientists and eight of the largest pharmaceutical companies together to find new uses for abandoned compounds (Allison 2012). On the one side, it is not surprising that marketed products play an important role in guiding subsequent research efforts of both the innovating firm and its rivals. On the other side, failures substantially spur innovative efforts, since abandoned drugs can be redirected, repurposed, repositioned and reprofiled (Ashburn and Thor 2004).

The pharmaceutical industry represents an ideal testbed for our model, given that: (i) the industry is a textbook example of a "science-based" sector (Pavitt 1984): strong linkages exist between drug development and the scientific advances in the "Open Science", leading firms to dissect and to analyze an increasing number of techniques, research lines and exploration strategies (Orsenigo et al. 2001); (ii) patents play an important role as means for protecting innovation – in exchange of the full disclosure of the characteristics of the innovation (Mansfield 1986; Cohen et al. 2000; Arundel and Kabla 1998; Pammolli and Riccaboni 2007); (iii) the pharmaceutical R&D process is characterized by a large presence of knowledge spillovers, as rival research results are positively correlated with firm productivity (Henderson and Cockburn 1994; 1996); (iv) the innovation process is characterized by strong uncertainty and high failure rates in drug discovery and development (Pammolli and Riccaboni 2004; Munos 2010; Pammolli et al. 2011).

Our model is tested using data from a comprehensive dataset about the innovative activity of pharmaceutical and biotechnology firms, including R&D project level data, patents, and citations, allowing us to explore the nature of technological advances and of the underlying learning regime, shaping the industrial patterns of innovative activity (Nelson and Winter 1982; Dosi 1982, 1988; Winter 1994; Malerba and Orsenigo 1993; Breschi et al. 2000). We exploit the information provided by patents and patent citations as a proxy for research efforts and to identify the different research lines in order to characterize the dynamic nature of the innovation process (Trajtenberg 1994; Lanjouw and Schankerman 1999; Harhoff et al. 1999; Jaffe et al. 2000; Trajtenberg et al. 1997; Jaffe and Trajtenberg 2002).

The paper is organized as follows. Section 2 describes the model and its main assumptions. This also sets the background of the subsequent empirical analysis, developed in Section 3, and the results are reported in Section 4. Section 5 concludes.

2 Theory

We consider scientific research as a cumulative process that consists of a series of experiments to discover the right or successful approach for product development (e.g., drugs). The technology space contains $I > 1$ paradigms. Within each paradigm, there are $J > 3$ candidate approaches or lines of research. Different research lines in a paradigm compete for success, but not different paradigms. For simplicity, we consider the case where at most one line of research is the "right" approach, that is, within one paradigm, there is at most one research line leading to success (e.g., a marketable drug); but whether a paradigm contains a successful line of research is not

affected by that of other paradigms. The true state of nature concerning a paradigm is described by which line of research, if any, is the successful one.

A research project is an experiment on a line of research. An experiment generates a piece of evidence concerning whether the experimented line of research can eventually succeed. An experiment on line j generates one of three outcomes $\tau^j \in \{s, f, n\}$, where a result $\tau^j = s$ means "success," while $\tau^j = f$ and n represent "failure" and "no result," respectively. If research line j will succeed, then the experiment returns with a good sign $\tau^j = s$ with probability β_1; and if research line j will fail, the outcome of the experiment will be bad, $\tau^j = f$, with probability β_0, that is, β_1 and β_0 are the probability that, conditional on the state, the experimental result coincides with the true state of nature.

With probability γ, the experiment generates no useful result: the occurrence of the signal $\tau^j = n$ does not depend on the true state. With the remaining probability $1 - \beta_1 - \gamma$ (probability $1 - \beta_0 - \gamma$), the experiment delivers the wrong result when the true state is that research line j will succeed (will fail, respectively). We impose the following assumption so that a result $\tau_j \in \{s, f\}$ remains informative.

Assumption 1 (Informative experiments) $\beta_1 + \beta_0 > 1 - \gamma > 0$.

A lower γ raises the likelihood that the experiment delivers some outcome (success or failure), and a higher β_0 or β_1 ensures that this outcome provides more information: it is more aligned with the true state of nature. These parameters capture the "quality" of experiment and determine the informativeness of an experimental outcome.[1] Note that, to apply Bayesian updating, firms need to know the value of the triplet $(\beta_1, \beta_0, \gamma)$, namely, the quality of experiment. In our empirical context, this information may be disclosed via clinical trials, patent documents or related scientific publication. In addition, firms may differ in their ability to conduct scientific experiments. There will be no conceptual difficulty in introducing firm heterogeneity regarding the probability triplet. Since it would not change the fundamental insight of the model, we do not pursue this complication here, but leave it for future treatment.

Given this information structure, we use Bayesian updating to calculate the successful probability in the research process. After t experiments conducted in the same paradigm, denote the profile of each research line's success probability as $\left\{\hat{\alpha}_t^j\right\}_{j=1,2,\ldots,J}$, with $0 \leq \sum_{j=1}^{J} \hat{\alpha}_t^j \leq 1$ due to the mutually exclusive success of each line of research. Let the $t+1$th experiment be conducted on research line j. The

[1]This type of research is "applied" research in the sense that the experiment is conducted only on one line of research within one paradigm, despite its informational externality, to be shown below. On the other hand, "basic" or "fundamental" research can be thought of as experiments delivering direct results about several lines of research, within or across paradigms, according to how "fundamental" the research could be. This distinction between basic and applied research stresses not the timing of invention, but the contribution to the knowledge accumulation process. Another type of research is research tools, which can be modeled as an invention that increases the precision of applied research, i.e., a better research tool increases β_1 and β_0.

realization of the experiment outcome $\tau_{t+1}^{j} \in \{s, f, n\}$ causes new assessment of the success probability according to Bayesian updating. Let $k \in \{1, 2, \ldots, J\}$ and $k \neq j$,

- $\tau_{t+1}^{j} = s$: it occurs with probability $\hat{\alpha}_t^j \beta_1 + \left(1 - \hat{\alpha}_t^j\right)(1 - \beta_0 - \gamma)$, and

$$\hat{\alpha}_{t+1}^{j} = \frac{\hat{\alpha}_t^j \beta_1}{\hat{\alpha}_t^j \beta_1 + (1 - \hat{\alpha}_t^j)(1 - \beta_0 - \gamma)} > \hat{\alpha}_t^j, \quad \hat{\alpha}_{t+1}^{k} = \frac{\hat{\alpha}_t^k (1 - \beta_0 - \gamma)}{\hat{\alpha}_t^j \beta_1 + (1 - \hat{\alpha}_t^j)(1 - \beta_0 - \gamma)} < \hat{\alpha}_t^k; \quad (1)$$

- $\tau_{t+1}^{j} = f$: it occurs with probability $\hat{\alpha}_t^j (1 - \beta_1 - \gamma) + \left(1 - \hat{\alpha}_t^j\right)\beta_0$, and

$$\hat{\alpha}_{t+1}^{j} = \frac{\hat{\alpha}_t^j (1 - \beta_1 - \gamma)}{\hat{\alpha}_t^j (1 - \beta_1 - \gamma) + \left(1 - \hat{\alpha}_t^j\right)\beta_0} < \hat{\alpha}_t^j, \quad \hat{\alpha}_{t+1}^{k} = \frac{\hat{\alpha}_t^k \beta_0}{\hat{\alpha}_t^j (1 - \beta_1 - \gamma) + \left(1 - \hat{\alpha}_t^j\right)\beta_0} > \hat{\alpha}_t^k; \quad (2)$$

and

- $\tau_{t+1}^{j} = n$: it occurs with probability $\hat{\alpha}_t^j \gamma + \left(1 - \hat{\alpha}_t^j\right)\gamma$, and

$$\hat{\alpha}_{t+1}^{j} = \frac{\hat{\alpha}_t^j \gamma}{\hat{\alpha}_t^j \gamma + (1 - \hat{\alpha}_t^j)\gamma} = \hat{\alpha}_t^j, \quad \hat{\alpha}_{t+1}^{k} = \frac{\hat{\alpha}_t^k \gamma}{\hat{\alpha}_t^j \gamma + (1 - \hat{\alpha}_t^j)\gamma} = \hat{\alpha}_t^k. \quad (3)$$

By mutual exclusivity, a research line's positive outcome "crowds out" the prospects of other research lines in the same paradigm. More interestingly, a failed experiment reduces the success probability of that line of research (and the whole paradigm), but at the same time *increases* the probability that the successful route may hide in other research lines. A failed experiment not only indicates that the experimented line of research is less likely to be the right approach, but also suggests that the future endeavor, if we continue the search in this paradigm, probably should look elsewhere. This piece of information could also be valuable to guide future research efforts. In other words, both successful and failed experiments are informative.

Now denote $\hat{\alpha}_t^* \equiv \max\left\{\hat{\alpha}_t^j\right\}$ as the highest successful probability in the paradigm after t experiments, and let the corresponding line of research be j^*. We now consider a simple R&D decision process, in which we assume that a firm's incentives to start a research project in this paradigm is increasing in the probability $\hat{\alpha}_t^*$. Post entry, the firm, with the belief profile $\left\{\hat{\alpha}_t^j\right\}$, conducts experiment on the research line j^*. An uninformative result $\left(\tau_{t+1}^{j^*} = n\right)$ does not change any beliefs and so $\left\{\hat{\alpha}_{t+1}^j\right\} = \left\{\hat{\alpha}_t^j\right\}$. But a positive sign $\left(\tau_{t+1}^{j^*} = s\right)$ raises the highest successful probability. A positive experimental result maintains the status of j^* as the most promising line of research, and raises its successful probability, $\hat{\alpha}_{t+1}^{j^*} > \hat{\alpha}_t^{j^*}$. Therefore, a successful experiment will increase entry incentives, whereas this is not true for one experiment with no informative outcome. If the experiment outcome is negative, $\left(\tau_{t+1}^{j^*} = f\right)$, then the successful probability of line of research j^* becomes lower. But the information brought by a failed experiment may sufficiently boost the successful probability of

other research lines, such that the highest successful probability after incorporating the new information is larger than $\hat{\alpha}_t^{j^*}$. That is, there may exist $k \neq j^*$ such that

$$\hat{\alpha}_{t+1}^{k} = \frac{\hat{\alpha}_t^k \beta_0}{\hat{\alpha}_t^{j^*}(1-\beta_1-\gamma) + \left(1-\hat{\alpha}_t^{j^*}\right)\beta_0} > \hat{\alpha}_t^{j^*}. \tag{4}$$

Since, in our empirical analysis, a failed experiment ($\tau^j = f$) is a project that is stopped due to the emergence of toxicological effects or lack of effectiveness, we further impose the following assumption so that a failed experiment indicates that the experimented line of research cannot succeed.

Assumption 2 (Failure means failure) $\beta_1 = 1 - \gamma$.

Given this assumption, $1 - \beta_1 - \gamma = 0$; that is, if the true state is that the line of research j will succeed, then the experimental outcome cannot wrongly return a failure result $\tau^j = f$. Therefore, a failed experiment shows unambiguously that the line of research can never succeed. Together with the assumption that $\beta_0 > 1$, this assumption also ensures that Assumption 1 will hold. Under this assumption, condition (4) becomes

$$\frac{\hat{\alpha}_t^k \beta_0}{\left(1-\hat{\alpha}_t^{j*}\right)\beta_0} = \frac{\hat{\alpha}_t^k}{1-\hat{\alpha}_t^{j*}} > \hat{\alpha}_t^{j*}. \tag{5}$$

When this condition holds, a failed outcome will also attract more entry than one with an uninformative outcome. In general, this condition requires that the difference in successful probability between the most promising and the second most promising line of research is not too large. For instance, if at the initial state, i.e., before any experiments are run, the prior belief is characterized by uniform distribution, $\hat{\alpha}_0^j$ is the same for all j, then the condition holds. Or, if the number of candidate research lines J is large enough, and if past experiments are all failures, then there is always some "untested" line of research k. The success probability of research line k is the same as $\hat{\alpha}_t^{j^*}$; the condition also holds.

Second, we compare positive and negative experimental outcomes. Fixing $\hat{\alpha}_t^{j^*}$, compare the highest successful probability after an experiment is conducted on line of research j^*. For $k \neq j^*$, under Assumption 2,

$$\frac{\hat{\alpha}_t^{j^*}\beta_1}{\hat{\alpha}_t^{j^*}\beta_1 + \left(1-\hat{\alpha}_t^{j^*}\right)(1-\beta_0-\gamma)} > \frac{\hat{\alpha}_t^k \beta_0}{\hat{\alpha}_t^{j^*}(1-\beta_1-\gamma) + \left(1-\hat{\alpha}_t^{j^*}\right)\beta_0} \tag{6}$$

$$\Leftrightarrow \quad \frac{\hat{\alpha}_t^{j*}(1-\gamma)}{\hat{\alpha}_t^{j*}(1-\gamma) + \left(1-\hat{\alpha}_t^{j*}\right)(1-\beta_0-\gamma)} > \frac{\hat{\alpha}_t^k}{1-\hat{\alpha}_t^{j*}} \tag{7}$$

$$\Leftrightarrow \quad \beta_0 \hat{\alpha}_t^k > (1-\gamma)\left[\frac{\hat{\alpha}_t^k}{1-\hat{\alpha}_t^{j*}} - \hat{\alpha}_t^{j*}\right]. \tag{8}$$

When condition (8) holds, a positive result must raise the highest success probability by a larger amount than a negative result, and so must attract more research efforts. This condition requires that β_0 be sufficiently large. Since Assumption 2 ensures that an experimental outcome of failure perfectly indicates the true state of nature, a high β_0 boils down to the requirement that a successful experiment also gives a sufficiently accurate, though still imperfect, assessment of the prospect.

To summarize, we will test the following proposition:

Proposition 1 *(Information spillover) Suppose that Assumption 2 holds. Under Condition* (8)*, a successful outcome increases entry incentives by a larger amount than a failed experiment or an experiment the result of which is inconclusive. When condition* (5) *holds, an experiment corresponding to a failed outcome also increases entry incentives more than one corresponding to a project that is inconclusive.*

Clearly, not all experiments lead to a patented invention and results are usually kept secret. However, in our empirical analysis we will focus on downstream applied research in the pharmaceutical industry. Therefore, we analyze preclinical and clinical experiments on patented drugs to measure the impact of failed and successful clinical trials on subsequent R&D efforts. According to the outcome of the associated R&D project, we classify patented drugs into three categories. A patent is called: (1) a success, when it is matched to a R&D project that leads to a drug being marketed; (2) a failure, when the associated R&D project is terminated during clinical trials (due to, for instance, toxicological effects or lack of effectiveness); and (3) inconclusive (or no information), when we (and rival companies) have no information about effectiveness in humans (i.e. no clinical trial has been conducted). We emphasize that, by a successful or failed patent, we do not mean that the patent itself is a success or failure. A patent is granted for its technological or scientific contribution, and thus represents some success in advancing knowledge. These modifiers refer to the subsequent experimental outcome of R&D projects associated with the patent. Since all new drug candidates are patented well before entering into trials and all companies must successfully pass through clinical trials to launch innovative drugs, the pharmaceutical industry provides an ideal setting to investigate the value of failed R&D attempts. As a proxy of the value of an experiment for follow-on research, we use the number of forward citations of the related patent. This implies that, when one starts a new experiment, the corresponding patent is likely to cite previous patents in the same paradigm in proportion to their contribution to resolve uncertainty about the probability of success of different lines of research.

3 Empirical strategy

The model presented in the previous section is tested in the context of the worldwide pharmaceutical industry. The empirical analysis builds on a comprehensive dataset on the innovative activity undertaken within the pharmaceutical industry, including R&D project level data, patents, and citations maintained by IMT Lucca, Italy. The database contains information about all pharmaceutical and biotechnology

patents granted by the USPTO since 1965, including backward and forward citations.[2] Firm data at the level of specific R&D projects worldwide in the last 30 years are also available. The database tracks the development history of more than 22,000 R&D projects, starting from patent application to the latest stage of drug development through preclinical, clinical development and commercialization. In the case of aborted projects, the database reports the time when the firm announces that the research around the compound has been terminated. By exploiting the information about the patents protecting the compound, the project data have been linked to patent data (number of forward citations up to May 2004, the application and grant date, and the name of the assignee(s)). For projects listing a patent granted by a patent office other than the USPTO, we considered the US patent in the same family as the one listed in the database. In case no US patent is identified, the project is not considered in the analysis. This choice has been driven by the fact our data sources only provide citation data for the US patents. Moreover, different patent examination procedures characterize the US and European patent offices, leading to large differences in the average number of citations per patent (Breschi and Lissoni 2004; Michel and Bettels 2001). Focusing only on US citations avoids the emergence of spurious results driven by different institutional settings. We further excluded old molecules and/or natural products, which do not have any associated patent. A total of 16,342 new molecules have been considered. Patent history is available for 49 per cent of the projects included in the database (7,971 projects).[3]

The analysis takes into account all the patents associated with R&D projects of successful candidate drugs (s), i.e. a new product launched on the market, or a failure (f), i.e. the project was stopped due to the emergence of toxicological effects or to lack of effectiveness. Many projects are still ongoing and, for some abandoned projects, the information about the reason of discontinuation is not available. The combination of the two selection criteria, i.e. at least one US patent covering the candidate drug and known outcome (either success or failure), largely restricts the set of projects we consider in our analysis. The final database encompasses information on about 2,000 R&D projects with informative outcomes and their associated patents. Henceforth, we refer to successful (failed) patents as the patent associated with successful (failed) R&D projects.

Put differently, each patent corresponds to an experiment and clinical trials are needed to distinguish failures (f) from successes (s). If no preclinical or clinical development is known after the patent is granted, the experiment is labeled as uninformative (n).[4] Next, we use forward citations to link research efforts within a research

[2]US patents are selected in the database on the basis of the International Patent Classification (IPC) and US classification. Pharmaceutical patents are defined as those in IPC classes A61K and A01N (Lanjouw and Cockburn 2001) and we further include patents in US classes 424, 435, 514, and 800.

[3]The matching of the different sets of data proved to be a formidable, large-scale task, that tied up a great deal of our research efforts for a long time, providing us a unique dataset that monitors R&D activities of pharmaceutical and biotechnology firms from patenting to commercialization (if any) of the protected compound.

[4]What is relevant about uninformative patented compounds is that no information about their therapeutic (lack of) effectiveness has been made available. In a sense, they can also be considered as early failures to be compared to *informative* later failures and successes.

line or technological path. We build on the idea, well-established in the literature, that if patent A cites patent B, it means that patent A contains a piece of knowledge patent B is building upon.

In order to test Proposition 1, citation patterns of our sample of patents is compared with patents the protected compound of which has not entered preclinical or clinical development. For each (successful or failed) patent in the original sample, a patent has been randomly matched from the set of bio-pharmaceutical patents with the same application year, publication year, and IPC class, but no information about preclinical or clinical development (n, no result).[5] By comparing the citation pattern of failed and successful patents against the sample of uninformative patents, we will be able to ascertain the level of knowledge utilization and diffusion and the related R&D competition dynamics. The three groups resemble the three research outcomes considered in the model (successful: s; failed: f; no information on preclinical or clinical development: n).

Since the focus is on the pharmaceutical industry, only citations from subsequent patents in the pharmaceutical domain have been taken into account. We disregard self-citations and focus on the patterns of citations made by other companies, proven to be a good proxy for knowledge spillovers (Jaffe et al. 2000).[6]

First, we employ the double-exponential function to model the citation lag distribution for successful and failed patents, against biopharmaceutical patents with no information about development. The model provides a flexible framework for studying the process underlying the generation of citations, where an exponential process by which knowledge diffuses is combined with a second exponential process by which knowledge becomes obsolete (Jaffe and Trajtenberg 1996; Caballero and Jaffe 1993).

The analysis allows us to characterize knowledge dynamics of successful and failed experiments (as identified through patents and patent citations) with respect to the uninformative patents issued within the bio-pharmaceutical domain (Proposition 1).

Next, we run a regression where the dependent variable is the number of citations received by the patents adjusted on the basis of the estimated citation lag distribution, in order to reflect life-time citations (Hall et al. 2005). Variables on the right hand side measure the importance of scientific references and the level of knowledge accumulated within the relevant domain of research, controlling for characteristics of the

[5]Besides grant year and application year, by matching patents on the basis of IPC class, we are able to control for the technological field. We have built three different "matched" samples in order to check the robustness of our results. Estimated coefficients across the three samples do not change substantially.

[6]On the contrary, self-citations are considered to be indicators of the cumulative nature of the technology and a measure of the extent to which innovators are able to reap the benefits of their own research (Hall et al. 2001). A known source of noise in citation studies comes from the fact that citations in the final patent document are not only those declared by the inventors, but also added by the examiner. Recent literature shows that analysis based on pooled sets of citations may suffer from bias (Alcácer and Gittelman 2006). In the analysis, we will make *relative* comparison across the citations to different groups of patents (patents with outcome equal to f, s, n.) As long as the number of citations added by the examiner is unrelated to the outcome of the associated R&D project (which is unknown when the patent is granted), our relative comparison is unaffected by the examiner-citation issue.

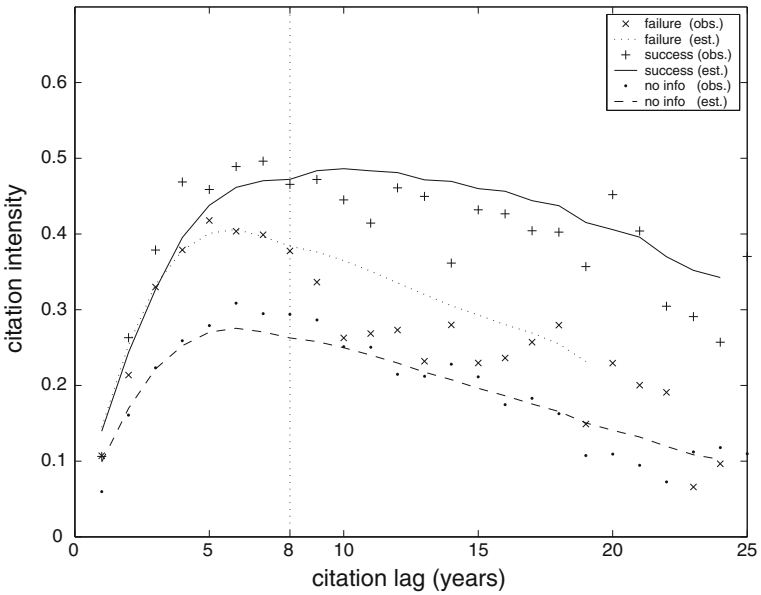

Fig. 1 Observed and estimated citation lag distribution

technological classes (defined on the basis of the International Patent Classification, henceforth IPC), assignees and the patent-innovation itself.

The two sets of results will provide a comprehensive view of the dynamics underlying R&D competition in the biopharmaceutical domain, allowing us to show the dynamics associated with successful and failed experiments in drug development.

4 Results

Figure 1 compares the observed and estimated citation lag distribution functions for successful (s) and failed (f) patents, taking as a benchmark the citation lag distribution function of patents with no information (n) about preclinical or clinical development. The x-axis reports the citation lag, i.e. the difference between the citing and the cited patent grant year. It represents the time elapsed from the grant date. The vertical line drawn eight years after patent grant corresponds to the average length from patent application to termination of the project.[7] The y-axis depicts the (average) observed and estimated citation intensities, i.e. the likelihood that any patent will be cited by the patents granted x years apart (Jaffe and Trajtenberg 1996).

[7]This is actually a few months longer for marketed compounds, being equal to 7.8 years for failed R&D projects and to 8.3 years for marketed R&D projects. The value is consistent with previous studies analyzing the average duration of the drug development process (Abrantes-Metz et al. 2004).

The observed citation lag distribution is computed as the ratio between the number of citations received by patents granted in year t from patents granted in year T and the theoretical number of potential citations ($T - t$ is the citation lag):

$$p(t, T, \tau) = \frac{C(t, T, \tau)}{N(t, \tau)\, N(T)}$$

where t indicates the grant year of the cited patent, T is the grant year of the citing patent, and τ represents the outcome of the associated R&D project ($\tau \in \{s, f, n\}$). The potential number of citations is given by the number of citations that would have been observed if all patents granted in year T would cite all patents granted in year t with outcome τ, that is equal to the product of the number $N(T)$ of patents granted in the citing year and the number $N(t, \tau)$ of patents granted in the cited year with a known outcome τ.

In order to estimate the theoretical citation lag distribution, we consider the following specification of the double-exponential function, modeling the likelihood that a patent granted in year T will cite a patent granted in year t, with $T > t$ (Jaffe and Trajtenberg 1996):

$$p(t, T, \tau) = \delta_0(t, T, \tau) \exp[-\delta_1(\tau)(T - t)](1 - \exp[-\delta_2(T - t)]).$$

where δ_0 is linked to the overall likelihood of receiving a citation, whereas δ_1 and δ_2 are indicators of, respectively, the rate of obsolescence of knowledge (i.e., the rate at which new knowledge replaces the existing knowledge) and the rate of diffusion of the knowledge related to the invention protected by the patent. We claim that the grant year of the citing and the cited patent only affect the average citation intensity δ_0, while the outcome of the project affect both the average citation intensity δ_0 and the rate of obsolescence δ_1.[8] Due to identification problems, the rate of diffusion (δ_2) is considered constant over time and across the three sets of patents.

According to our model, we expect that successful patents receive on average a higher number of citations with respect to failed patents. The model specification, however, allows us also to dig further into the dynamics of knowledge utilization and diffusion and to analyze the speed and extent to which existing knowledge is "picked up" in the case of failures and successes, as a proxy for the diffusion and utilization of the associated research line. Results, reported in Table 1, are obtained by nonlinear least squares estimation, weighting each observation by $[N(t, \tau)\, N(T)]^{1/2}$. The lines depicted in Fig. 1 are obtained by taking the average,[9] for each lag, of the fitted values from Model 2.

Consistent with our model and previous literature showing that the number of citations received by a patent is positively associated with its technological impact,[10] citations turn out to be related to the outcome of the project. The observed and estimated distributions indicate that, on average, failed patents receive a number

[8] As in previous empirical literature dealing with this model, convergence problem forbids the estimation of the model where all the cited-year effects are considered. The problem is solved by introducing the cited-year effects defined on the basis of 5-year time periods.

[9] Both in the case of observed and estimated citation lag distributions, weighted averages are considered, where the weights are the same as the ones used in the estimation process.

[10] See, e.g., Jaffe and Trajtenberg (2002).

Table 1 Double-exponential function, results of estimation, dependent variable: citation intensity

Coefficients	Model 1		Model 2	
δ_0^f	1.264	*(.182)**	1.329	*(.111)**
δ_0^s	1.309	*(.158)**	1.511	*(.108)**
δ_1	0.107	*(.008)**	0.084	*(.013)**
δ_1^f	0.827	*(.135)**	1.049	*(.108)**
δ_1^s	0.554	*(.087)**	0.604	*(.077)**
δ_2	0.114	*(.014)**	0.248	*(.042)**
Cited year effects	no		yes	
Citing year effects	no		yes	
R-squared	.609		.641	

* statistically significant at 5 % level
Standard errors in parenthesis. In square brackets: t-stat. for H0: parameter = 1 (if relevant)

of citations lower than the number of citations received by patents associated with successful projects (compare the estimated values of the δ_0^f and δ_0^s coefficients). Furthermore, consistent with Proposition 1, both sets receive a higher number of citations than the average patent in the bio-pharmaceutical domain with no information about clinical or preclinical development ($\tau = n$). The estimate of the δ_0 coefficients associated with failed and successful patents is higher than one, indicating that patents associated with preclinical or clinical development, irrespective of their outcome, are more likely to receive a citation than bio-pharmaceutical patents with no informative results (taken as the reference category). The analysis shows that, even though the compound associated with the patent will never reach the market (e.g., due to the emergence of toxicological problems or lack of effectiveness), the results of the research project are subsequently exploited by firms other than the original innovator.

It is worth noticing that, within the first five years from patent grant, no significant difference is detected between failed and successful patents, whereas starting from year five the two series start to diverge. The analysis of the estimates of δ_1 reveals an important difference between successful and failed compounds in terms of knowledge obsolescence. In this respect, failed and the uninformative patents exhibit very similar dynamics (the δ_1 associated with failed patents is very close to one, pointing to no differences between failed patents and the uninformative patents). On the contrary, the knowledge embedded in patents protecting marketable compounds becomes obsolete less quickly than the other patents in the biopharmaceutical domain (the δ_1 associated with successful patents is lower than one). Indeed, the citation intensity of marketed compounds is rather stable after commercialization, whereas the citation intensity of failed patents decreases substantially.[11]

[11] Also note that the larger departures between the estimated and observed citation lag distribution in the case of failed patents is registered right after the average time when the project is stopped. This might point to the fact that the termination of the research around a compound/mechanism of action is a major signal for rival firms. Nonetheless failed patents regain interest after a few years from the time of discontinuation and their citation intensity is still higher than the citation intensity of uninformative patents, also many years after discontinuation. On this issue, we asked a pharmacologist to inspect extensively the patents citing failed projects in search of a reason for the citation, finding no instance of "negative" citations. Rather citations refer to pharmacological action or the structure of the compound (i.e., to the research line of the original innovator).

Table 2 Description of the variables

Variable	Char.	Description	Mean	σ
science	(pt)	Science Index*	0.34	0.35
cite-sf	(pt)	Number of backward citations with known outcome (either failure or success)	0.25	0.59
selfc	(pt)	Share of self-citations of the patent*	0.14	0.29
orig	(pt)	Index of originality of the patent*	0.43	0.37
timeb	(pt)	Average time lag*	5.56	4.50
importb	(pt)	Importance of cited patents*	122.4	640.7
nimp	(ipc)	Number of firms operating in the same IPC class	110.21	126.32
conc	(ipc)	Concentration of the IPC Class (Herfindhal index)	0.10	0.21
coree	(asg)	Share of firm patent within the same technology class (IPC)	0.14	0.23
dbf	(asg)	Dummy equal to 1 if the originating firm is a dedicated biotechnology company	0.19	0.39
pro	(asg)	Dummy equal to 1 if the originator is a public research organization	0.09	0.29

* defined as in Trajtenberg et al. (1997)
Char.: characteristic of (pt) patent; (ipc) IPC class; (asg) assignee

The maximal citation frequency for failed patents is earlier in time than the maximal citation frequency of successful patents.[12]

The disclosure of the information about the compound under study in patents and the advances in science set the ground for a "race" in reaching the market, where competitors start exploring the new research arena, pursuing parallel research trajectories even though the outcome is still highly uncertain. Competition on the R&D side in the pharmaceutical industry is substantial and firms entering the new research arena build both on failures and successes (Fig. 1).

We now turn to a regression model that looks at failures and successes separately and at the drivers of their citation intensity. We build a sort of production function, where the characteristics of the experiments (patents) and of the field of research are taken as determinants of the technological impact (as proxied by the number of forward citations) of the invention.

[12] Our estimates are coherent with the estimates of the Drugs and Medical sector presented by Hall et al. (2001), with the exception of the estimated δ_2, which is lower. This might be explained by the fact that we only consider citations by institutions other than the original assignee, which can require a longer time span with respect to self-citations. Moreover, an interesting pattern emerges in their results when comparing Drugs and Medical to other sectors. The citation lag distribution for this sector is flatter, whereas the citation lag distribution functions for the sectors of Computers and Communications, Electrical and Electronics, Chemical, and Mechanical have higher peaks earlier in time. Knowledge in the Drugs and Medical sector diffuses less rapidly and takes a longer time to become obsolete. Important information about the protected compounds in terms of toxicological effects and effectiveness are revealed over time, leading to a lengthier process of citation within this industry.

Table 3 Correlation table

	ncit	science	cite-sf	selfc	orig	timeb	importb	nimp	conc	coree	dbf
ncit											
science	.1378										
cite-sf	.0375	.1429									
selfc	−.0741	.0322	−.0830								
orig	−.0686	.3544	−.0682	.2539							
timeb	.0287	.3957	.1743	.0191	.3154						
importb	.3485	.1342	.1424	−.0337	−.0894	.1106					
nimp	−.0318	.1353	.0469	−.0317	−.0077	.0380	.0945				
conc	−.0025	−.0290	−.0208	−.0019	−.0088	.0088	−.0154	−.3436			
coree	.0707	.0409	−.0060	−.1554	−.0098	.0191	.0491	.0541	−.0401		
dbf	.2037	.2474	.0593	−.0820	−.0487	.0863	.1784	.0984	−.0165	.2180	
pro	.0412	.1170	−.0220	−.0749	−.0330	−.0509	.0010	.0517	−.0026	.0172	−.1571

The dependent variable is the number of citation received during the life time of the patent (Hall et al. 2005). The observed citation frequency has been adjusted using the estimated coefficient of the citation lag distribution function on the basis of estimated coefficients in Table 1.

The independent variables are listed in Table 2 and aim at capturing the characteristics of the cited patent (pt), the IPC class (ipc) and the patenting firm (asg). Table 3 shows the correlation among the variables.

Patent characteristics are measured using the indicators developed by Trajtenberg et al. (1997) on the basis of backward citations. The importance of scientific sources with respect to technological ones within the patent is captured by $science$, which is the ratio between the non-patent references and the total number of references (previous patents or previous scientific literature) listed in the patent. The closer to one, the larger the scientific underpinnings of the research, relying more heavily on the scientific literature rather than on previous patents. The variable has been considered as a proxy for the basicness of the innovation and a positive coefficient is expected, as more basic inventions are expected to have larger technological impact. This is especially true in the pharmaceutical industry, the archetype example of a "science-based" sector (Pavitt 1984). The index of originality of the patent ($orig$) is also considered, which measures the breadth of its technological roots,[13] whereas the importance of the previous patents cited by the patent under investigation is measured by $importb$, which takes into account the number of backward citations in the

[13]The index is computed as an Herfindahl index of diversification, considering the share of backward citation in each IPC class. The closer *orig* is to one, the broader are the technological roots of the underlying research, i.e. they span many different IPC classes. The index is zero when all backward citations contained in the patent are classified within the same IPC class.

patents and the number of citations they receive.[14] Besides these variables, the history of the research line is taken into account by counting the number of citations with known outcomes (a proxy for the level of information that characterizes the patent's research trajectory; $cite - sf$).

We also consider the share of self-citations in the patents ($selfc$) that measures the extent to which benefits from research antecedents are appropriated by the firm and help in understanding whether the patent belongs to a research line strongly rooted within the company. The expected sign for this variable is negative: it is considered a measure of the appropriability of the line of research, and so a lower level of spillovers to rivals is expected.

Finally, $timeb$ measures the time distance between the citing and the cited patents. The higher $timeb$, the older the sources the patent builds upon.[15]

As for the characteristics of the technological class of the patent, we consider the number of firms active in the IPC class, and the Herfindahl index of concentration computed at the technology class level on the basis of patent counts.

We also take into account the share of firm patents within the same technology class (IPC), and two dummy variables indicating whether the patentee is a dedicated biotechnology company (DBF) or a public research organization (PRO). The largest share of patents in our sample are assigned to pharmaceutical companies: 19 per cent of patents are assigned to DBF, and nine per cent to PRO. Previous literature has shown that, in the pharmaceutical sector, smaller biotechnology firms are more likely to average a higher citation rate (Hall et al. 2000). PROs are expected to run more basic research, and so a higher rate of citations is expected (see also Magazzini et al. 2012).

Finally, cited year dummies are included in all the specifications.

Results of the estimation of a Poisson regression model are reported in Table 4, where we distinguish between successes and failures. The Poisson modeling has been preferred to simple regression of a log-linearized equation, as recent research shows that, under heteroskedasticity, OLS estimation of log-linearized models lead to biased estimates of the true elasticities (Santos Silva and Tenreyro 2006). Furthermore, we record 20.77 % patents receiving zero citations.

Even though the magnitude of the coefficients differ between successes and failures, statistical significance of the regressors does not change. The one exception is the coefficient associated with the number of citations with known outcomes ($cite - sf$) that is negative for failures and seems to exhibit no effect for successes. The lower the number of patents (previous research) with informative outcomes, the higher the number of citations to failed patents. Put differently, failures receive a higher number of citations (i.e. are the basis for subsequent innovations) in nascent fields as compared to fields where more knowledge has been accumulated, leading

[14] The higher the value of $importb$, the higher the number of backward citations contained in the patent and the citations they receive.

[15] As compared with the descriptive statistics reported in Trajtenberg et al. (1997), no difference emerges with respect to the value of $selfc$. On the contrary, the average value of $timeb$ in our sample is lower, indicating younger sources for our sample of patents, whereas the values of $orig$, $science$, and $importb$ are higher. One important difference with the sample in Trajtenberg et al. (1997) is that we only consider pharmaceutical patents, and citations are counted only within the pharmaceutical technological classes.

Table 4 Regression results (Poisson model; pseudo-ML estimation)

Variable	Failure		Success	
science	.9689	*(.1729)**	.4841	*(.1744)**
cite-sf	−.2704	*(.0895)**	.0174	*(.0794)*
conc	−.2452	*(.3553)*	.0997	*(.2850)*
nimp	−.0009	*(.0005)*	.0005	*(.0005)*
selfc	−.2626	*(.1863)*	−.1983	*(.1909)*
orig	−.8878	*(.1368)**	−.5642	*(.1534)**
timeb	−.0030	*(.0106)*	−.0253	*(.0130)*
importb	.0002	*(.4E-4)**	.0003	*(.6E-4)**
coree	.1886	*(.2292)*	−.0587	*(.1630)*
dbf	.7057	*(.1085)**	.6450	*(.1289)**
pro	.6216	*(.1914)**	.3617	*(.1775)**
constant	2.672	*(.2422)**	3.247	*(.2853)**
Obs.	1,554		725	
Log lik.	−21123.2		−11679.7	

Dependent variable: number of adjusted forward citations
(Robust) standard errors in parenthesis
* statistically significant at 5 % level
Cited year (application) included in all regressions

us to claim that condition (5) is actually satisfied within the pharmaceutical domain. By contrast, citations to success do not seem to be driven by the level of knowledge accumulated in a given research domain.

As expected, patents with a predominance of scientific sources over technological ones contribute more heavily to subsequent research.

The estimated coefficient of *orig* shows that patents with sparse technological roots receive a lower number of citations by other firms. These are likely patents within narrow fields of application, being relevant only to the firms and institutions working within the same technological domain. Patents building on an important (in the sense of highly cited) knowledge base are more often subsequently cited.

As far as the characteristics of the patent assignee are concerned, patents by DBF and PRO receive on average a higher number of citations. Within the pharmaceutical domain, these are characterized as running more basic research, and therefore are expected to have higher technological impact.

Finally, the characteristics of the IPC class of the patent do not seem to exert a significant effect on the number of subsequent citations received by the patent.

5 Concluding discussion

Innovation is a cumulative process in which both successful inventions and failures play a fundamental role. As compared to success stories and positive feedback, the role of failures has been under-investigated in the theoretical and empirical literature on innovation and learning. In the same vein, little has been done from the societal

point of view to reap the fruit of negative outcomes in science and technology. Our contribution to fill this gap is twofold.

First, we develop a cumulative innovation model by analyzing R&D competition and learning mechanisms along different lines of research that posits a role for successes as well as (under suitable conditions) failures as inputs to the innovative activities of the innovation process. According to our model, failures are particularly informative under conditions of high uncertainty among multiple plausible research lines and strong competition, or selection. Learning by the negative experience of others presupposes full disclosure of both failed and successful trials. This matter deserves much more attention than in the past, both from positive and a normative perspectives. In this paper, we intentionally kept our model as simple as possible, to highlight the contribution of failures to learning and innovation. In future work, we plan to generalize the model to the case of milder selection regimes (multiple successful lines of research in a paradigm) and heterogeneity of beliefs about the prospects of different research lines. Another important limitation to the model we introduced is that failure means failure under Assumption 2. The role of false negatives should be further investigated to find the most efficient solutions to identify and successfully recover abandoned lines of research.

Second, the model has been tested in the context of the worldwide pharmaceutical industry. The pharmaceutical industry is a suitable testbed for the model in its current form, since drug development is characterized by high cost, delayed feedback and strong uncertainty. Moreover, information about patented inventions as well as failures in clinical trials are publicly available. By relying on a comprehensive dataset of innovative activities in the pharmaceutical industry, we have provided evidence of the value of failures in the pharmaceutical innovation process. The analysis suggests the existence of a social value associated with discontinued projects and with the disclosure of the associated information. If the FDA rejects an application, the letter to the manufacturer stating the reasons is confidential. Unlike the FDA, the European Medicines Agency (EMA) has recently begun to post reasons for non-approval on its website. However, in most of cases, evidence required to judge accurately drug's effectiveness and safety is missing from the public domain. Much of the information collected in unsuccessful drug development trials is inaccessible to the research community (Hakala et al. 2015). Non-publication deprives research systems of crucial evidence that should inform the planning of future clinical trials.

Within this scenario, the discussion about patent scope becomes crucial, where research is highly cumulative in nature and firms enjoy knowledge spillovers spanning from internal and, to some extent, from external R&D projects, pointing to a trade-off that cannot be easily resolved. This poses problems for the optimal design of patent law. On the one side, it is necessary to reward fully early innovators for the technological foundation they provide to later innovators. In this respect, dedicated incentive schemes such as a 'market for R&D failures' (Shalem and Trajtenberg 2009) and 'inside-out' open innovation practices (Chesbrough and Chen 2015) have been put forward to enhance the transmission of the information regarding failures

among competitors. This is important, since many initially rejected drugs are eventually approved (Sacks et al. 2014). On the other side, follow-on innovators should be rewarded adequately for their improvements on successful products (Magazzini et al. 2015).[16]

Our results in this paper provide both a theoretical justification and strong empirical evidence in favor of recent initiative to disclose fully the information on clinical trials and to find new uses for abandoned compounds (Doshi et al. 2013; Mello et al. 2013). About 90 % of the investment companies make per successful drug development program has gone to failed compounds. Unfortunately, most of the time the information about failed clinical trials is kept secret: trials of unapproved drugs are much less likely to be published than trials of licensed drugs, limiting the scientific knowledge that supports further drug development (Hakala et al. 2015). Moreover, information voluntarily disclosed by pharmaceutical companies in scientific publications tends to be incomplete as compared with unpublished clinical trial reports (Wieseler et al. 2013). Therefore, we strongly support the view that the FDA should follow the EMA's lead and make non-approval decisions public. Commercial confidentiality should not prevent the scientific community from accessing crucial information about the viability of alternative research lines in clinical trials.

References

Abrantes-Metz RM, Adams CP, Metz A (2004) Pharmaceutical Development Phases: A Duration Analysis, Bureau of Economics, US FTC

Aghion P, Dewatripont M, Stein J (2008) Academic Freedom, Private-sector Focus, and the Process of Innovation. RAND J Econ 39(3):617–635

Alcácer J, Gittelman M (2006) Patent Citations as a Measure of Knowledge Flows: The Influence of Examiner Citations. Rev Econ Stat 88(4):774–779

Allison M (2012) NCATS Launches Drug Repurposing Program. Nat Biotechnol 30(7):571–572

Arora A, Gambardella A (1994) The Changing Technology of Technological Change: General and Abstract Knowledge and the Division of Innovative Labour. Res Policy 23(5):523–532

Arthur WB, Lane DA (1993) Information Contagion. Struct Chang Econ Dyn 4(1):81–104

Arundel A, Kabla I (1998) What Percentage of Innovations are Patented? Empirical Estimates for European Firms. Res Policy 27(2):127–141

Ashburn TT, Thor KB (2004) Drug Repositioning: Identifying and Developing New Uses for Existing Drugs. Nat Rev Drug Discov 3(8):673–683

Breschi S, Malerba F, Orsenigo L (2000) Technological Regimes and Schumpeterian Patterns of Innovation. Econ J 110(463):388–410

Breschi S, Lissoni F (2004) Knowledge Network From Patent Data: Methodological Issues and Target Research. CESPRI, Milan. Working Paper

Caballero RJ, Jaffe AB (1993) How High are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth, NBER Working Paper

Cabral LM (2005) The Economics of Trust and Reputation: A Primer. New York University and CEPR

Chatterjee RA, Eliashberg J (1990) The Innovation Diffusion Process in a Heterogeneous Population: A Micromodeling Approach. Manag Sci 36(9):1057–1079

Chesbrough H, Chen EL (2015) Using Inside-Out Open Innovation to Recover Abandoned Pharmaceutical Compounds. J Innov Manag 3(2):21–32

[16] See Scotchmer (1991) for a detailed discussion of the optimal patent scheme in the case of cumulative knowledge.

Cohen WM, Nelson RR, Walsh J (2000) Protecting Their Intellectual Assets: Appropriability Conditions and Why US Manufacturing Firms Patent (or Not), NBER Working Paper
Cyert RM, March JG (1992) A Behavioral Theory of the Firm, 2nd edn. Blackwell, Cambridge
David PA, Mowery D, Steinmueller WE (1992) Analysing The Economic Payoffs From Basic Research. Economics of Innovation and New Technology 2(1):73–90
Denrell J (2003) Vicarious Learning, Undersampling of Failure, and the Myths of Management. Organ Sci 14:227–243
Doshi P, Dickersin K, Healy D, Vedula SS, Jefferson T (2013) Restoring Invisible and Abandoned Trials: A Call for People to Publish the Findings. Br Med J 346:227–243
Dosi G (1982) Technological Paradigms and Technological Trajectories. Res Policy 11(3):147–162
Dosi G (1988) Sources, Procedures, and Microeconomic Effects of Innovation. J Econ Lit 26(3):1120–1171
Green J, Scotchmer S (1995) On the Division of Profit in Sequential Innovation. RAND J Econ 26:20–33
Hakala A, Kimmelman J, Carlisle B, Freeman G, Fergusson D (2015) Accessibility of Trial Reports for Drugs Stalling in Development: A Systematic Assessment of Registered Trials. Br Med J 350: h1116
Hall BH, Jaffe AB, Trajtenberg M (2000) Market Value and Patent Citations: A First Look, NBER Working Paper
Hall BH, Jaffe AB, Trajtenberg M (2005) Market Value and Patent Citations. RAND J Econ 36(1):16–38
Hall BH, Jaffe AB, Trajtenberg M (2001) The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools, NBER Working Paper
Harhoff D, Narin F, Scherer FM, Vopel K (1999) Citation Frequency and the Value of Patented Inventions. Rev Econ Stat 81(3):511–515
Henderson R, Cockburn I (1994) Racing to Invest? The Dynamics of Competition in Ethical Drug Discovery. J Econ Manag Strateg 3(3)
Henderson R, Cockburn I (1996) Scale, Scope, and Spillovers: The Determinants of Research Productivity in Drug Discovery. RAND J Econ 27(1):32–59
Herriot SR, Levinthal D, March JG (1985) Learning from Experience in Organizations. Am Econ Rev 75:298–302
Jaffe AB, Trajtenberg M (1996) Flows of Knowledge from Universities and Federal Labs: Modeling the Flowof Patent Citations Over Time and Across Institutional and Geographic Boundaries, NBER Working Paper
Jaffe AB, Trajtenberg M, Fogarty MS (2000) The Meaning of Patent Citations: Report on the NBER/Case-Western Reserve Survey of Patentees, NBER Working Paper
Jaffe AB, Trajtenberg M (2002) Patents, Citations and Innovations: a Window on the Knowledge Economy. MIT Press
Lanjouw J, Schankerman M (1999) The Quality of Ideas: Measuring Innovation with Multiple Indicators, NBER Working Paper
Lanjouw JO, Cockburn IM (2001) New Pills for Poor People? Empirical Evidence after GATT. World Dev 29(2):265–289
Levinthal DA, March JG (1993) The Myopia of Learning. Strat Manag J 14:95–112
Levitt B, March JG (1993) Organizational Learning. Annu Rev Sociol 14:319–340
Magazzini L, Pammolli F, Riccaboni M, Rossi MA (2009) Patent Disclosure and R&D Competition in Pharmaceuticals. Economics of Innovation and New Technology 18(5):467–486
Magazzini L, Pammolli F, Riccaboni M (2012) Learning from Failures or Failing to Learn? Lessons from Pharmaceutical R&D. Eur Manag Rev 9:45–58
Magazzini L, Pammolli F, Riccaboni M (2015) Real Options and Incremental Search in Pharmaceutical R&D Project Portfolio Management. Creativity and Innovation Management. doi:10.1111/caim.12119
Malerba F, Orsenigo L (1993) Technological Regimes and Firm Behavior. Ind Corp Chang 2(1):45–71
Mansfield E (1986) Patents and Innovation: An Empirical Study. Manag Sci 32(2):173–181
Mello MM, Francer JK, Wilenzick M, Teden P, Bierer BE, Barnes M (2013) Preparing for responsible sharing of clinical trial data. N Engl J Med 369(17):1651–1658
Michel J, Bettels B (2001) Patent Citation Analysis. A Closer Look at the Basic Input Data from Patent Search Reports. Scientometrics 51(1):185–201
Munos B (2010) Lessons From 60 Years of Pharmaceutical Innovation. Nat Rev Drug Discov 8:959–968
Nelson RR, Winter SG (1982) An Evolutionary Theory of Economic Change. Belknap Press
O'Donoghue T (1998) A Patentability Requirement for Sequential Innovation. RAND J Econ 29:654–79

Orsenigo L, Pammolli F, Riccaboni M (2001) Technological Change and Network Dynamics: Lessons from the Pharmaceutical Industry. Res Policy 30(3):485–508
Pammolli F, Magazzini L, Riccaboni M (2011) The Productivity Crisis in Pharmaceutical R&D. Nat Rev Drug Discov 10:428–438
Pammolli F, Riccaboni M (2004) Market structure and drug innovation. Health Aff 23:48–50
Pammolli F, Riccaboni M (2007) Innovation and Industrial Leadership: Lessons from Pharmaceuticals, Center for Transatlantic Relations. John Hopkins University
Pavitt K (1984) Sectoral Patterns of Technical Change: Towards a Taxonomy and a Theory. Res Policy 13(6):343–373
Powell WW, Koput KW, Smith-Doerr L (1996) Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology. Adm Sci Q 41(1):116–145
Sacks LV, Shamsuddin HH, Yasinskaya YI, Bouri K, Lanthier ML, Sherman RE (2014) Scientific and Regulatory Reasons for Delay and Denial of FDA Approval of Initial Applications for New Drugs, 2000–2012. JAMA 311(4):378–384
Santos Silva JMC, Tenreyro S (2006) The Log of Gravity. Rev Econ Stat 88(4):641–658
Scotchmer S (1991) Standing on the Shoulders of Giants: Cumulative Research and the Patent Law. J Econ Perspect 5(1):29–41
Shalem R, Trajtenberg M (2009) The market for R&D failures, Working paper. Available from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1547268
Stoneman P (1981) Intra-Firm Diffusion, Bayesian Learning and Profitability. Econ J 91:375–388
Teerlak A, Gong Y (2008) Vicarious Learning and Inferential Accuracy in Adoption Processes. Acad Manag Rev 33:846–868
Trajtenberg M (1994) A Penny for Your Quotes: Patent Citations and the Value of Innovations. RAND J Econ 21(1):172–187
Trajtenberg M, Henderson R, Jaffe A (1997) University vs. Corporate Patents: A Window on the Basicness of Innovations. Economics of Innovation and New Technology 5(1):19–50
Wieseler B, Wolfram N, McGauran N, Kerekes MF, Vervolgyi V, Kohlepp P, Kamphuis M, Grouven U (2013) Completeness of Reporting of Patient-Relevant Clinical Trial Outcomes: Comparison of Unpublished Clinical Study Reports with Publicly Available Data. PLoS Med 10(10):e1001526
Winter SG (1994) Schumpeterian Competition in Alternative Technological Regimes. J Econ Behav Organ 5:287–320