

FINDING DIAMONDS IN THE ROUGH: DATA-DRIVEN OPPORTUNITIES AND PHARMACEUTICAL INNOVATION *

Matteo Tranchero
UC Berkeley-Haas

October 8, 2023

PLEASE CLICK [HERE](#) FOR THE LATEST VERSION

Abstract

Big data are increasingly used to make predictions about uncertain investments, thereby helping firms identify innovation opportunities without the need for domain knowledge. This trend has led to questions about which firms will primarily benefit from the availability of these data-driven predictions. Contrary to existing research suggesting that data-driven predictions level the playing field for firms lacking domain knowledge, I argue—using a simple theoretical framework—that these predictions actually reinforce the competitive advantage of firms with domain knowledge. In innovation contexts, where returns are skewed and not all leads can be pursued, domain knowledge helps evaluate predictions and avoid false positives. I test this idea in the context of pharmaceutical innovation, exploiting the features of genome-wide association studies (GWASs) that provide data-driven predictions about new drug targets. The results show that GWASs stimulate corporate investments, but around one-third of these resources are misallocated toward false positive predictions. Companies lacking domain knowledge react more strongly but are disproportionately likely to fall into the trap of false positives. Instead, domain knowledge helps firms make fewer investments that target only the best opportunities. Together, the results show that even if data-driven predictions hold value when searching for innovations, domain knowledge remains the crucial source of competitive advantage in the age of big data.

*E-mail: m.tranchero@berkeley.edu. Website: <https://www.matteotranchero.com/>. I thank Lee Fleming, Abhishek Nagaraj, Carolyn Stein, Toby Stuart, and the members of the Management and Organizations Group at Haas for their precious advice and encouragement throughout this project. I gratefully acknowledge the Will Mitchell Dissertation Research Grant from the Strategic Management Society that made this research possible. All errors remain my own.

“Petabytes allow us to say: ‘Correlation is enough’ ”

Chris Anderson (2008)

1 Introduction

The data revolution is transforming how companies operate (Adner et al., 2019), with firms increasingly using machine learning and artificial intelligence (AI) to analyze large amounts of data and make predictions about future events or uncertain choices (Agrawal et al., 2018; Brynjolfsson et al., 2021; Cockburn et al., 2019; Sun et al., 2021). In addition to being used to inform operational decisions, data-driven approaches are becoming common when making strategic decisions in high-risk contexts such as innovation or entrepreneurship. For instance, venture capitalists use predictive algorithms to identify successful start-ups (Bhatia and Dushnitsky, 2023), consumer goods corporations rely on social media data to forecast the revenues of prospective products (Allen and McDonald, 2022), and pharmaceutical firms try to predict the therapeutic properties of millions of potential drugs (Heaven, 2023; Lou and Wu, 2021). In both traditional and high-tech sectors, data-driven predictions are reshaping how firms find innovation opportunities.

Given the prevalence of this phenomenon, a central question is how predictive technologies affect market competition (Iansiti and Lakhani, 2020). Traditionally, scholars have argued that domain knowledge allows firms to foresee the value of alternative choices, thus enhancing performance by recognizing opportunities (Cohen and Levinthal, 1990; Gruber et al., 2008; Taylor and Greve, 2006). However, data-driven predictions can serve the same role without needing domain knowledge (Agrawal et al., 2023; Balasubramanian et al., 2022). Since obtaining predictions is now cheaper and faster than accumulating expertise, performance gains are likely to favor those lacking domain knowledge. Recent research in contexts like writing or gaming supports this idea: predictive technologies help mostly novices, hence substituting for their lack of knowledge (Choi et al., 2022; Noy and Zhang, 2023). Despite still lacking direct evidence, one would expect a similar democratization effect when extrapolating this logic to firm innovation. Data-driven predictions might erode the role of domain knowledge as a source of competitive advantage, leveling the playing field through what some have hailed as the “death of theory” (Anderson, 2008).

Even though these arguments suggest that domain knowledge might no longer be necessary in the age of big data, there are several reasons to think otherwise. The limited existing evidence shows that data-driven predictions can, on average, yield significant performance gains (Nagaraj, 2022; Kao, 2023), but the skewed distribution of innovation returns means that most predictions are worthless. Said otherwise, even data-driven predictions that are valuable on average can mask several false positive findings (Berman and Van den Bulte, 2022). When validating predictions is relatively inexpensive,

such as in operational decisions, firms can afford a “spray and pray” approach to experimentally learn which ones unlock performance improvements. However, this logic might not extend to innovative contexts, where testing a prediction is costly and dynamically forecloses other paths (Gans et al., 2019; Levinthal, 2021). Financial and organizational constraints mean that the firm cannot follow every prediction, and betting on the wrong ones can dramatically hurt performance (Eggers, 2012). These features of the innovation process imply that achieving gains from predictive technologies poses unique challenges.

Motivated by this logic, in this paper I argue that domain knowledge is a crucial complement to data-driven predictions in the context of innovation. While it is true that predictions can identify opportunities without requiring comprehension, domain knowledge becomes essential in evaluating their plausibility to rule out false positives (Allen and Choudhury, 2022; Arora and Gambardella, 1994). This suggests a two-stage process, where predictive algorithms generate a shortlist of potential opportunities, while domain knowledge is used to rank them and funnel resources to the most promising ones. To formalize this concept, I introduce a simple theoretical framework showing how the same data-driven predictions can have heterogeneous effects among firms facing equal investment costs. My framework yields a few testable implications. Data-driven predictions will increase aggregate investments, mostly by firms without domain knowledge; however, such firms will be less likely to uncover actual opportunities in the data despite pursuing a larger number of predictions. Acknowledging the consequences of false positives suggests the opposite of existing research: data-driven predictions in innovation might reinforce the advantage of incumbents with the domain knowledge to evaluate them rather than helping novices catch up.

The validity of my theoretical framework is ultimately an empirical question. However, empirically assessing how data-driven predictions and domain knowledge interact in shaping innovation performance presents several issues. First, it requires the ability to observe the risk set of potential opportunities that firms could pursue, usually coupled with the impossibility of learning the actual value of counterfactual investments that were not chosen. Second, there also needs to be a way to know which specific data predictions were available and how they aligned with any given firm’s knowledge. The last challenge presents a crucial identification issue: if the quality of data-driven predictions correlates with firms’ domain knowledge, then it becomes impossible to distinguish the prediction’s quality from the ability to correctly interpret it.

I address these challenges by focusing on the setting of pharmaceutical innovation, a field with empirical features ideally suited for this paper. The drug development process is a lengthy and costly endeavor, starting with the difficult identification of genes that can act as drug targets. The recent emergence of genomics has enabled a new approach to locate the genetic roots of human diseases,

called genome-wide association studies (GWASs). GWASs are studies that compare the genomes of individuals with and without a specific disease to identify genetic mutations that correlate with the presence of a disease. In practice, they yield predictions on the value of genes as drug targets, bypassing the need to grasp the gene's role in the disease. This allows an exploration of how pharmaceutical firms react to the arrival of data-driven predictions that inform their selection of genetic targets for diseases. I use patent applications as an indicator of early stage investments and the discovery of new drugs to trace downstream outcomes following a GWAS.

To identify the causal role of data-driven predictions, I exploit three unique features of GWASs. First, these studies scan the entire genome and do not target specific genes ex-ante, ensuring that any correlation between genes and the disease of interest is not the result of endogenous selection by the researcher (Uffelmann et al., 2021). Second, GWASs are primarily conducted by academic researchers, who make their findings publicly available to every firm. New gene-disease associations are usually unexpected by researchers and even more so by firms, who learn about them when they are published in scientific journals. Third, GWASs uncover both breakthrough opportunities and false positive genetic targets. To separate them, I exploit the fact that diseases are often subject to multiple GWASs of increasing statistical power over time (Marigorta et al., 2018). Subsequent GWASs provide an intuitive way to identify findings that are not robust to replications, revealing that they were likely false positives.

This setting allows me to test how firms leverage data-driven predictions. However, on average, the GWAS papers in my data present 14 new gene-disease associations, of which 12 are false positives. This presents a challenge as conventional measurements, such as patent-to-paper citations, do not permit linking investments to individual predictions within the same GWAS. To overcome this challenge, I adopt a novel text-based approach to extract the specific associations targeted in each patent and map them onto an empirical landscape of all possible gene-disease combinations. This approach enables me to track firms' investments over the search landscape before and after the arrival of data-driven GWAS findings while bypassing the use of patent citations. I then use firms' publications to measure the genes researched before GWASs' arrival. As confirmed in my interviews, firms with an active research program on a gene can leverage this knowledge to assess GWAS findings even if the genetic expertise was developed for another disease (Cattani, 2005). This allows me to study the impact of data-driven predictions separately for firms with and without domain knowledge to evaluate them.

I find that the average number of patent applications on a gene-disease combination more than doubles after it is reported by a GWAS. Event study specifications confirm the absence of pre-trends and the validity of the research design. However, I also find that over a third of patent applications are directed to pursuing false positive findings. Investments made based on these false positives fail

to yield downstream outcomes such as highly cited patents or new drug molecules, suggesting that misallocating resources can hurt firm performance. Next, I explore the interaction between firms' domain knowledge and the availability of data-driven predictions. The results indicate that firms lacking domain knowledge react more strongly after a GWAS is published. However, they are disproportionately likely to select false positive findings. In contrast, firms with domain knowledge exhibit a more discerning approach, making fewer investments that target the most promising opportunities, simultaneously reducing commission (type I) and omission (type II) errors.

I further investigate these findings with a firm-level research design. Even after controlling for firm and firm-by-disease fixed effects, the analyses confirm that genetic knowledge makes firms more efficient and targeted in their investments. To explore the mechanisms underlying my results, I categorize the specific nature of firms' domain knowledge (Zollo and Winter, 2002). Regression analyses show that the ability to avoid false positives does not reflect the sheer amount of past investments or publications; instead, it seems to be the purview of firms that can translate basic scientific principles to the evaluation of GWAS associations. Additional results rule out social connections with scientists, generic organizational capabilities, or different firm-level strategies as alternative mechanisms. Taken together, these findings provide intriguing evidence that the ability to assess data-driven predictions stems from understanding the mechanisms of what makes an opportunity valuable (Camuffo et al., 2023b).

This paper makes several contributions. First, it contributes to a developing research agenda on how data-driven innovation through tools like AI will shape innovation (Agrawal et al., 2023; Allen and McDonald, 2022; Bessen et al., 2022; Cockburn et al., 2019; Jin and McElheran, 2019) and its limits (Cao et al., 2023; Hoelzemann et al., 2023; Tranchero, 2023; Wu et al., 2020). By shedding light on the implications of false positive predictions, my results show the importance of going beyond the average effects of predictive technologies by directly considering firm-level heterogeneity. In particular, my novel framework highlights the possession of domain knowledge as a critical determinant of who benefits from data-driven predictions in innovation. Relative to research showing a democratizing effect of data availability (Nagaraj, 2022), accounting for false positives implies a complementarity with domain knowledge that reinforces incumbents' advantage (Kao, 2023). Greater attention to the mechanisms and boundary conditions of data-driven innovation may provide insights into why the promise of big data seems slow to materialize in the context of innovation (Brynjolfsson et al., 2021; Lou and Wu, 2021).

Moreover, the paper contributes to the literature on firm capabilities (Cattani, 2005; Dosi et al., 2001; Teece, 2007) and their interaction with predictive technologies (Balasubramanian et al., 2022). Using a new empirical approach to map firms' investments onto an empirical search landscape, I show

the continuing importance of absorptive capacity, intended as the use of domain knowledge to both recognize and exploit opportunities (Cohen and Levinthal, 1990). However, while the literature has often focused on the latter aspect, this paper suggests that assessing predictions and understanding which ones to trust poses a crucial challenge for firms (Bikard, 2018; Eggers and Kaplan, 2009; Kaplan et al., 2003). The paper also adds to the theory-based view in management (Camuffo et al., 2023b; Wuebker et al., 2023) by showing the changing role of domain knowledge: no longer originating ideas but instead separating the wheat from the chaff among the data. Last, the paper builds on research about the use of predictive tools in hiring, marketing, and other firm operations (Cowgill, 2018; Hoffman et al., 2018; Kim, 2023a). I show how this work can extend to innovation, where tool-specific skills (Yue et al., 2023) are likely insufficient to improve performance since domain knowledge is pivotal to interpreting data-driven predictions.

The paper proceeds as follows. Section 2 presents the theoretical framework. Section 3 presents the drug discovery setting, the measurement strategy, and the data. Section 4 discusses the research design. Section 5 reports the results, while Section 6 discusses mechanisms at the firm level. Section 7 concludes.

2 Theoretical Framework

Survey evidence shows that the majority of U.S. firms use predictive tools, even in traditional sectors like manufacturing (Brynjolfsson et al., 2021). Yet, little is known about their impact on innovation, which entails searching for opportunities in complex technological spaces. This section addresses this gap by theorizing the mechanisms through which data-driven predictions change innovation and which firms benefit from them.

2.1 Using Data-Driven Predictions to Find Innovation Opportunities

A firm’s ability to select the best opportunities is a critical determinant of its competitive advantage (Teece, 2007). Firms devote substantial effort in evaluating potential investments to allocate their limited resources (Criscuolo et al., 2017; Krieger and Nanda, 2022). This is especially true in innovation, where technological landscapes are too vast to experiment with all the potential combinations (Fleming, 2001). Traditionally, firms have navigated this complexity through two main strategies: accumulating domain knowledge (Arora and Gambardella, 1994; Cohen and Levinthal, 1990; Gruber et al., 2008) or sourcing ideas externally (Fabrizio, 2009; Luo et al., 2021; Piezunka and Dahlander, 2015). For example, discovering a new drug requires searching in a chemical space of $\sim 10^{60}$ molecules. Given this enormity, it is impossible to test all of them. Therefore, pharmaceutical firms have historically focused on internal research and academic collaborations to funnel their resources toward the most promising leads (Cockburn and Henderson, 1998).

However, the recent availability of big data and predictive technologies¹ has significantly altered this status quo. Firms can now use data-driven predictions to assess the value of uncertain alternatives (Agrawal et al., 2023; Lou and Wu, 2021). Data on past successful combinations can be analyzed to find patterns that help assess the viability of potential investments (Kim, 2023b), thus triaging them before committing any actual resource (Kao, 2023; Nagaraj, 2022). Consider again the setting of drug discovery. A few ago, Stokes et al. (2020) used data on past antibiotics to score the therapeutic potential of new molecules from their shape, predicting those with the highest antibacterial properties. This process yielded a new drug without needing knowledge of chemistry, bacteriology, or biology.

The power of data-driven predictions lies in their ability to facilitate a large-scale inductive method for technological search (Strasser, 2019). Algorithms help identify association patterns between variables that elude even the most careful observation, making the search process effectively “automated” (Mullainathan and Rambachan, 2023). The advantage is that data-driven predictions is not confined to technological components of known function or to relationships with established theories (Tranchoero, 2023). Instead, they offer a theory-free approach to prioritizing search, offering guidance even if the underlying mechanisms for success are unknown to the innovator. Returning to the case of antibiotics discovery, no generalized theory exists of how molecule structures map into antibacterial effects. The algorithm developed by Stokes et al. (2020) sidestepped this issue because it recognized patterns pointing at short-changed molecules whose effects on pathogens were thus far unknown.

Predictive technologies are fundamentally reshaping the type of knowledge that organizations use to identify opportunities, with data-driven predictions becoming the basis for innovation decisions instead of domain expertise. This shift may deter firms from investing in domain knowledge since data-driven predictions could be a cheaper and more fungible substitute (Balasubramanian et al., 2022; Puranam, 2019). The consequence could be what the popular press has called the death of theory (Anderson, 2008; Schmidt, 2023; Zittrain, 2019), potentially eroding the role of domain knowledge as a source of superior performance in innovation. This change has important implications for strategy and competitive dynamics (Adner et al., 2019; Iansiti and Lakhani, 2020). For example, data-driven predictions may democratize competition by helping new entrants disrupt incumbents that have traditionally relied on large knowledge bases to keep them at bay. In summary, the rise of predictive technologies raises the question of whether domain knowledge remains relevant in innovation and who will benefit from these trends.

¹I define predictive technologies as algorithms designed to use large amounts of data to make predictive statements about future or unknown events (Brynjolfsson et al., 2021). This includes tools of varying levels of sophistication, from simple correlations to deep learning and AI (Kim, 2023b). However, data-driven predictions are conceptually different from descriptive uses of data analytics (Berman and Israeli, 2022).

2.2 Do Data-Driven Predictions Substitute Domain Knowledge?

Recent literature has started exploring the performance implications of adopting predictive algorithms. While still nascent, this work consistently finds that the most sizable performance improvements from predictive algorithms are seen among novices across a variety of tasks, from customer support to creative writing (Brynjolfsson et al., 2023; Choi et al., 2022; Doshi and Hauser, 2023; Gaessler and Piezunka, 2023). It also suggests that predictive algorithms can level the playing field for those lacking domain knowledge (Noy and Zhang, 2023).² However, these studies focus on tasks where data-driven predictions can be quite accurate, such as moves in board games or resume screening (Cowgill, 2018; Hoffman et al., 2018). More work is needed to explore if the same results apply in innovation, where the precision of predictive tools is likely lower due to the complexity of the prediction task (Tranchoero, 2023) and the incomplete training data (Choudhury et al., 2020; Kim, 2023b).

A parallel strand of work examines the impact of data-driven predictions on innovation (Agrawal et al., 2023; Allen and McDonald, 2022; Cockburn et al., 2019). However, these studies seldom address the interaction between such predictions and domain knowledge, leaving unanswered questions. The findings are mixed: some studies show that predictive technologies increase incremental innovation (Wu et al., 2020), while others suggest that they lead to more breakthroughs (Tranchoero, 2023). It is thus unclear if domain knowledge is essential to identifying valuable opportunities. Similarly, the types of firms that stand to benefit from data-driven predictions are ambiguous. Some research has found that new entrants can use data to recognize opportunities ahead of experienced firms (Jin and McElheran, 2019; Nagaraj, 2022), while others have reached the opposite conclusion (Kao, 2023). These conflicting results are puzzling if one thinks of data as reducing the cost of locating opportunities, which would imply that they substitute domain knowledge and benefit firms lacking expertise.

It is equally difficult to form expectations from the extensive literature investigating the role of domain knowledge in innovation (Cattani, 2005; Cohen and Levinthal, 1990; Volberda et al., 2010). This work finds that domain knowledge is associated with the ability to recognize opportunities (Arora and Gambardella, 1994; Kaplan and Vakili, 2015) but also that it limits search and leads to competency traps (Levinthal and March, 1993). A potential implication is that data-driven predictions might counteract organizational inertia and find breakthroughs outside domains mastered by the firm. However, because these predictions are atheoretical, experts may be skeptical (Boudreau et al., 2016), and research shows that actors with more domain experience are less willing to follow the guidance of predictive tools (Allen and Choudhury, 2022; Glaeser et al., 2022). Additional research in organizational contexts should explore if domain knowledge offsets the novelty introduced by data-

²A notable exception is the recent work by Jia et al. (2023), who show that if predictive tools only substitute mundane tasks, then experts will benefit by having more time to focus on problem-solving tasks.

driven technologies or if the two can work together as complements. From the existing evidence and frameworks, it is not easy to conclude what the role of domain knowledge will be as big data and predictive tools diffuse among firms.

2.3 Data-Driven Predictions and Domain Knowledge: A Simple Conceptual Framework

The difficulty in reconciling existing results points to the need for a new conceptual framework to understand the relationship between data-driven predictions and domain knowledge. Most of the existing literature neglects that innovation, unlike other prediction tasks, has specific features. First, false positives are frequent in technological search (Berman and Van den Bulte, 2022). Components interact in complex ways, meaning that associations between variables are often spurious (Ghosh, 2021; Tranchero, 2023). Second, validating a lead is much more costly than generating it and induces commitment that can foreclose other paths (Gans et al., 2019). Technological commitments are often irreversible (Levinthal, 2021), and the consequences of investing in suboptimal projects can be severe (Eggers, 2012). Last, the returns from innovation are extremely skewed. These three features imply that firms cannot afford to invest in every data-driven prediction: when technologies require sizable investments and take years to mature, such as deep tech or drug discovery, the price of a mistake far outweighs its learning spillovers (Greve, 2003).

This discussion implies that even if data-driven predictions are, on average, valuable, their gains are concentrated on a few high-value leads hidden between many false positives. The challenge for firms is thus assessing the plausibility of data-driven predictions and deciding which ones to follow. This suggests that domain knowledge could become a *complement* in choosing between potential opportunities in the data. In other words, the role of knowledge shifts from idea generation to filtering out false positive leads. One way to visualize this argument is as a two-stage process. First, data-driven predictions provide a shortlist of technological combinations that are valuable on average but are too many to test individually. Then, domain knowledge helps by ranking the perceived value of potential opportunities, thus channeling resources to the most promising predictions.³ This implies that firms with robust knowledge bases may continue to outperform in innovation, even if that knowledge was initially developed for other purposes (Cattani, 2005; Cohen and Levinthal, 1994).

In analytical terms, assume that risk-neutral firms search in a landscape of m potential technological combinations. Investing in any of them has a constant cost of C for every firm. Combinations can be of high or low value, but their average value is negative ($E(V) < C$), reflecting the skewness in innovation returns. Then, one can envision data-driven predictions as providing a shortlist of $n < m$

³For instance, domain knowledge might help firms understand the interdependencies between components, allowing them to rule out implausible associations in the data (Fleming and Sorenson, 2004).

combinations. Since the predictions are informative on average, firms know that the predicted value of these shortlisted combinations is higher than the cost ($E(\hat{V}_n) > C$). However, each prediction \hat{V} is a potentially noisy representation of the true value V . Firms will try to rank the predictions to start investing in the most promising ones and avoid false positives (Aguiar and Waldfogel, 2018). I propose that domain knowledge can be operationalized as giving firms a better assessment of the true value of the combination, conditional on observing the same prediction. This means that firms with domain knowledge will have a subjective ranking of data-driven predictions that is closer to the true underlying value of those combinations (i.e., closer to $V_1 > V_2 > \dots > V_i > \dots > V_n$).⁴ Now assume that firms choose sequentially and they experimentally learn the true value of a combination only after the investment. To maximize returns, they will start from the combination that they ranked highest, gradually going down their ranking until the prediction \hat{V}_i such that the expected value of the remaining data-driven predictions is below the investment cost ($E(\hat{V}_{n-i}) < C$).

This simple framework yields three testable implications. First, the data-driven predictions point, on average, to valuable opportunities and will lead all firms to increase their investments. Second, the noisier the ranking of a given firm, the higher the expected value of the residual $n - i$ data-driven predictions after the first i investments, since earlier investments will miss some high-value ones. This means that firms with less precise rankings will continue investing sequentially until the prediction $j > i$ such that $E(\hat{V}_{n-i-j}) < C$. Since firms lacking domain knowledge have a noisier ranking, one expects that they will end up pursuing a higher number of predictions. Third, those firms will make both more commission (or type I) and more omission (or type II) errors, being more likely to select low-value combinations that they mistakenly ranked high at the expense of true opportunities. Compared to firms with domain knowledge, those lacking it should be less likely to find the best opportunities among data-driven predictions despite investing quantitatively more.

Overall, this simple framework offers insights into how domain knowledge interacts with data-driven predictions and what we should expect to observe if the proposed mechanism holds (see also Appendix A for a numerical example). To examine whether these implications are supported by the data, the rest of the paper tests them in the context of pharmaceutical innovation. In particular, I investigate which firms can benefit from data-driven predictions that identify potential genetic targets for drug discovery.

⁴In other words, it is as if firm k observes a noisy firm-specific data-driven prediction: $\hat{V}_k = V + \gamma + \epsilon_k$, where $\gamma \sim N(0, \eta^2)$ is the noise in the prediction itself, $\epsilon_k \sim N(0, \sigma_k^2)$ is the noise in the firms' own assessment of the prediction, and γ reflects the precision of the prediction itself. Instead, the variance of the firm-specific error (σ_k^2) captures firm k 's domain knowledge, thus being smaller for more knowledgeable firms. Firm k will determine a subjective ranking of every combination i that received a data-driven prediction, ordering its $\hat{V}_{i,k}$ from largest to smallest. The smaller the error σ_k^2 , the closer k 's ranking is to the correct ranking based on the true underlying value of the combinations (Appendix A).

3 Empirical Setting, Measurement Strategy, and Data

3.1 Drug Discovery and GWASs

Developing a new drug is a risky process that takes, on average, 10-15 years and costs \$2.6 billion (Kao, 2023). Choosing the right *genetic target* is the first crucial step: namely, firms must identify genes that can be modulated to affect the outcomes of specific diseases (Nelson et al., 2015).⁵ This task is complex because, in principle, each genetic disease could be caused by mutations in any of the over 19,000 protein-coding human genes. Common diseases are often polygenic—for instance, diabetes has been tied to over 150 DNA mutations. The result is a vast combinatorial space of tens of millions of potential gene-disease combinations, of which only a minor fraction has therapeutic value. Given the inherent complexity and risks associated with identifying the rare gene-disease pairs that can lead to a drug (Tranchoero, 2023), it is unsurprising that pharmaceutical firms often rely on university research to reduce the risks involved in target selection (Fleming and Sorenson, 2004).

Starting from the early 2000s, the completion of the Human Genome Project and the steep decrease in the cost of collecting genetic data prompted the emergence of GWASs. GWASs are case-control studies where researchers sequence human genomes to find genetic mutations that are more likely to appear among subjects with a specific condition compared to healthy subjects (Uffelmann et al., 2021; Visscher et al., 2017). Figure 1 shows the abstract of one early GWAS and a stylized depiction of how a typical study unfolds. Researchers start by collecting DNA samples from several people, some healthy and some affected by the disease of interest. They then use microarrays to sequence DNA locations called “markers” to reconstruct the genetic constitution (or genotype) of the people in the sample. Last, they test for statistically significant differences in genotypes that correlate with the presence of the disease.⁶ Intuitively, by comparing the genetic makeup of people with and without a condition, one can predict which genes might be involved in a given disease and thus serve as drug targets (see Appendix B for more details).

GWASs provide data-driven leads into gene-disease combinations that can be valuable in research and development (R&D), and anecdotal evidence points to several drugs developed thanks to GWAS findings (Visscher et al., 2017). Nevertheless, there is still no consensus on the ability of GWASs to identify targets with high therapeutic potential (Struck et al., 2018; Tam et al., 2019; Uffelmann

⁵Genes are sequences of DNA bases that encode the “instructions” to synthesize gene products (e.g., proteins) that allow the organism to function. When genes acquire mutations in their sequence, they might alter their behavior with significant consequences for health. Knowing the genetic roots of diseases is vital in designing pharmaceutical drugs because genes that cause a disease could be targeted for therapeutic purposes (see Appendix B for more details).

⁶Association tests are adjusted for multiple hypothesis testing, usually imposing a very high threshold for statistical significance. The results are often graphically represented as a “Manhattan plot” and show the p-value of multiple statistical tests between DNAs in the case and control groups. The y-axis usually reports $-\log_{10}(\text{p-value})$, and hence higher values correspond to stronger associations.

et al., 2021). One major limitation is that GWASs cannot explain the underlying mechanisms of gene-disease associations, raising the possibility of false positives with no actionable causal pathway (Goldstein, 2009; Hermosilla and Lemus, 2019). The average GWAS can find tens of potentially promising gene-disease associations, many of which could be worthless or spurious. This creates a challenge for scientists and firms interested in therapeutic applications due to the high costs involved in following a data-driven lead.⁷ As a result, questions remain regarding how many GWAS associations are valuable and whether they generate useful knowledge for drug discovery at all (Goldstein, 2009).

3.2 Measuring the Impact of GWASs with Text-Based Empirical Landscapes

For my empirical analysis, I need a precise method to measure which firms invest in the gene-disease associations uncovered by GWASs. Traditional approaches, which rely on patent-to-paper citations, are not well-suited for this context. On the one hand, a typical GWAS identifies multiple genes linked to a disease, but it is unclear which are true positives among the false positives. Thus, it is impossible to determine from a patent citation to the GWAS paper exactly which association the patent is leveraging. On the other hand, using direct citations is ineffective for tracing the unpredictable path between foundational basic research and innovation. For instance, patents might cite later research that validates a particular GWAS association but neglect to mention the original GWAS because such discovery has become common knowledge in the field.

To address these challenges, I develop a novel measurement strategy (see Appendix C for details and examples). Taking a recombinant view of innovation (Fleming, 2001), each GWAS association can be considered as a prediction of a valuable gene-disease combination. This suggests that the innovation impact of a GWAS is proxied by the change in firm patenting experienced by the gene-disease pairs “treated” compared to otherwise similar combinations not reported by the GWAS. Next, named entity recognition algorithms can be used to extract the genes and diseases from the text of each patent application, providing a way to infer which gene-disease combinations a firm invests in. Figure 2 offers an intuitive representation of how the increase in patents targeting specific gene-disease pairs captures the heterogeneous impact of individual GWAS associations. Assuming a standard parallel trends assumption, the before-and-after changes in the number of patents mentioning treated and control combinations provide an estimate of the GWAS’s effect.⁸ Notably, this approach permits the empirical study of firms’ exploration decisions in a combinatorial landscape of genes and diseases.

⁷This theme emerged prominently in my interviews with industry professionals: “I don’t think anyone who is doing good science would just trust a GWAS, even an accurate GWAS, because R&D is so expensive” (Interview, 17 October 2022).

⁸Nagaraj (2022) uses a similar idea to assess the impact of satellite images on gold discovery at the level of individual blocks of the earth. Kao (2023) adopts the same logic to study the effect of large-scale cancer maps on the number of clinical trials on 627 genes. The approach detailed in Appendix C formalizes their intuition and expands it to additional use cases, including quantifying science-to-technology spillovers without the limitations of citation data.

This approach offers two major advantages over traditional citation-based methods. First, employing text-based knowledge entities allows for distinguishing the impact of different contributions made in the same paper. This affords granularity at the gene-disease level that traditional citations lack. Second, by tracking the mention of genes and diseases directly in the patent’s text description, the approach is well suited to objectively capturing the impact of basic research findings that might not be explicitly acknowledged in downstream patent applications. Additionally, it must be noted that this measurement approach is much more general than the present setting (Ding et al., 2013). Appendix C presents conceptual details of how this measurement strategy can be generalized to other applications where document texts capture knowledge spillovers, bypassing the flaws of patent citations.

3.3 Data Sources

To empirically test how firms respond to the arrival of data-driven predictions, three key ingredients are required. First, I need data on new gene-disease associations identified by GWASs. Second, I need information on firms’ domain knowledge, investment decisions, and drug development outcomes. Third, it is essential to know the specific genes and diseases targeted by each GWAS, patent, and drug in order to link them. Below I summarize the main data (see Appendix D and E for details):

GWAS Catalog: To identify the data-driven predictions available to firms, I rely on the GWAS Catalog, a manually curated source managed by the European Bioinformatics Institute (EBI). The GWAS Catalog is a comprehensive list of GWASs published in peer-reviewed journals.⁹ Studies are eligible for being included in the GWAS Catalog if they include an array-based genome-wide scan that does not target any specific gene ex-ante. The GWAS Catalog also collects the details of the findings as appearing in the original study. My sample includes 17,965 gene-disease associations first reported by 1,259 distinct GWASs between 2005 and 2019 (Appendix Figure E.1). These associations include 404 unique diseases¹⁰ and 5,080 protein-coding genes.

SciBite/EBI Patent Data: Through a partnership with EBI, I obtained proprietary data to measure the gene-disease pairs targeted by each USPTO patent application (2001–2019). The data have been compiled from complete patent texts using TERMite (TERM identification, tagging, and extraction), a named entity recognition software developed by the Elsevier-owned startup SciBite (see Appendix D for details). A manual validation of 200 random patents finds that SciBite’s entity recognition algorithm has a precision rate of 95%–97% and a recall rate of 91%–92% and is thus highly reliable. I merge these data with information on assignees and patent characteristics taken from PatentsView. Commonly used

⁹The first GWAS was published in *Nature* in 2005.

¹⁰Diseases are identified as unique MeSH concepts (Azoulay et al., 2021). Since the MeSH taxonomy is a hierarchical tree, in my analysis I include all diseases at level 4 of the tree. The sample is further restricted to diseases that receive more than one GWAS because I exploit subsequent GWASs to code which gene-disease associations are not replicated and are thus likely to be false positives. See Section 3.4 for details.

indicators of patent quality are from the OECD Patent Quality Indicators Database (Squicciarini et al., 2013) and the data of Kogan et al. (2017). My final sample includes all pharmaceutical companies applying for at least one USPTO patent between 2001 and 2019.

PubTator Central Publication Data: I use publicly available data from PubTator Central, which provides computer-annotated genes and diseases for each paper published in PubMed. Next, I match each patenting firm in my sample with their respective publication records using the information on authors' affiliations.¹¹ Doing this allows me to code which firms had prior knowledge of the biology of a *specific human gene* before deciding whether to invest in a relevant GWAS finding. This procedure yields a very granular measure of gene-specific expertise, unlike previous studies that used corporate publications as a generic firm-level proxy for absorptive capacity (Gambardella, 1992). Additional bibliographic characteristics of firms' publications are from the National Institutes of Health's iCite database.

Cortellis Drug Data: I supplement my data with proprietary information on drug molecules collected by the Clarivate Analytics Cortellis Competitive Intelligence Database. For this study, I use the drug development records in Cortellis, which contain development information for over 70,000 drugs up to July 2020 (Krieger, 2021). Cortellis aggregates information from various sources to provide a comprehensive list of historical development milestones for each drug molecule. I use data about new molecules observed entering the earliest phases of drug development (what Cortellis records as the "discovery phase" and the "pre-clinical phase").

Open Targets Score: Open Targets is a public-private partnership that collects all the available evidence on the strength of gene-disease associations and summarizes it in a synthetic score. Experts in the field consider Open Targets as providing the most comprehensive assessment of the genetic roots of human diseases based on the available knowledge. I download the publicly available Open Targets scores and merge them with my data.¹² The score is available for 594,353 gene-disease pairs (8% of my sample), spanning 17,437 genes and 366 diseases.

By putting together these data sources, I can trace an empirical search landscape that captures all potential combinations of genes and diseases that pharmaceutical firms could select. Table 1 reports summary statistics at the gene-disease combination level, constituting the paper's primary unit of analysis. Around 21.8% of the 7,223,924 gene-disease pairs have received investments, as proxied by their appearance in at least one patent application. Yet, only 0.19% of these pairs have advanced to drug development stage, confirming the extremely skewed distribution of innovation

¹¹The coverage of authors' affiliations in PubMed is low and includes only the first author for papers published before 2013. To address this shortcoming, I obtained affiliation records for all PubMed articles from Dimensions.

¹²The data can be downloaded from <https://platform.opentargets.org/>. See Appendix E.2 for additional details.

returns. GWASs have uncovered 17,965 new associations, constituting 0.25% of this vast combinatorial landscape. The average gene-disease pair receives 0.13 patent applications per year. However, there is a considerable variation, with some genetic targets receiving over 1,000 annual patents. Around three-quarters of patents are filed by firms that have not previously published research on the targeted gene. Descriptive evidence confirms that firms are more likely to patent gene-disease combinations with higher therapeutic potential when they have been researching the gene involved (Appendix E.3).

3.4 Identifying False Positive Associations

GWASs conduct an unbiased, atheoretical search for genetic mutations associated with human diseases. This feature is both a strength and a weakness: despite very stringent statistical significance thresholds, it is well-known that many GWAS findings turn out to be statistical noise or spurious correlations (Goldstein, 2009; Storey and Tibshirani, 2003). For the purposes of this paper, I need to empirically measure which GWAS predictions proved false. The main challenge is that ex-post direct assessments of gene-disease associations' quality are only available for findings that received investments, resulting in a classic missing data problem. Specifically, researchers can only observe the actual state of the prediction when a decision maker chooses to act on it (Hoffman et al., 2018; Kleinberg et al., 2018). In my setting, this issue would make it impossible to understand if GWAS associations that received little attention were correctly avoided due to low potential or were mistakenly overlooked by firms failing to recognize their value.

To distinguish true positives from false positives, I exploit the fact that diseases are often subject to multiple GWASs over time. In particular, I categorize gene-disease associations as true positives only if they are initially reported by a GWAS and later confirmed by subsequent GWASs focusing on the same disease. This provides an intuitive way to identify likely false positives regardless of the level of investment they received.¹³ Using this approach, I find that 84.3% of the 17,965 gene-disease associations in my sample fail subsequent replications using different samples.¹⁴ This high number aligns with accounts from experts who have long warned about the risks posed by big data in genomics (MacArthur, 2012). Non-replicable findings can lead to an inefficient use of R&D resources, misdirecting firms' limited resources to targets without therapeutic potential (Freedman et al., 2015).

¹³Put otherwise, the genome-wide design of subsequent GWAS replications serves as a retest for all findings, thereby providing assessments even for associations that have not received investments from firms. Note however that this approach provides a conservative metric of which findings are false positive in a *therapeutic* sense since even replicable associations might not offer avenues for treatment (Goldstein, 2009; Hermosilla and Lemus, 2019).

¹⁴The share of replicable associations increases from 15.7% to 25.4% when I consider GWASs targeting diseases not in my sample. When including findings in intergenic regions, Marigorta et al. (2018) find that around 40% of associations are replicable. One reason behind the lower replicability for findings related to protein-coding genes is the frequent mistakes made by GWAS' authors in mapping mutations into the correct corresponding genes (Visscher et al., 2017). This type of mistake in reporting GWAS results provides an alternative empirical strategy that I leverage in Subsection 5.3.

I validate the measure of false positives through three methods. First, I compare my coding with the Open Targets score to measure the ground truth value of each gene-disease pair.¹⁵ Appendix Table F.2 shows that replicable gene-disease associations are 114%–191% more likely to rank in the top decile of the Open Targets score. Second, I investigate the characteristics of the study design that predict which GWAS discoveries will replicate in subsequent studies. The data in Appendix Table F.3 support the intuition that more robustly designed GWASs are less likely to report associations classified as false positives. Finally, I present evidence in Appendix Table F.4 to confirm that true positives are more likely to appear in articles that receive more citations in clinical trials and a lower share of citations with a negative tone, as captured by the Scite data (Nicholson et al., 2021). Together, these tests corroborate the approach used to identify false positives in GWAS associations.

4 Research Design

There are two main challenges to assessing the impact of data-driven predictions on innovation: measurement and endogeneity. The first challenge arises from the difficulty of determining which predictions a firm has observed. This becomes problematic when firms with domain knowledge obtain systematically better predictions, thus conflating the prediction’s quality with the ability to interpret it. The second challenge relates to identifying the causal effects of data-driven predictions. In practice, firms are likely to mine for predictions where they expect the highest returns (e.g., for genes known to be druggable), potentially leading to upwardly biased estimates. An ideal experiment would bypass these issues by assigning all firms identical information on gene-disease associations. The causal effect of data-driven predictions would then be evident from changes in patenting involving “treated” gene-disease pairs relative to the others. Subsequently, I could then test my theoretical framework with heterogeneity analyses by firms’ pre-existing domain knowledge.

I approximate this ideal experiment using the staggered publication of GWAS associations in scientific journals. First, GWASs are mainly conducted by academic research teams,¹⁶ and their findings are publicly accessible upon publication. As such, their findings are available to every pharmaceutical firm simultaneously and are not driven by unobservable proprietary data or expertise that they might have. Second, GWASs scan for genetic variants across the whole genome (Uffelmann et al., 2021). This method, by design, avoids the issue of endogenous sorting since it does not focus on specific genes. GWAS discoveries are therefore entirely unforeseen by the research team and even more so by firms that later read about the findings in journals. Importantly, unbeknownst to the scientists conducting the GWAS, only some of their findings will be confirmed in future publications.

¹⁵Out of the 17,965 gene-disease associations in my data, there are 16,298 pairs with an Open Targets Score.

¹⁶In my sample, 324 GWASs are co-authored by corporate scientists, and another 117 of them acknowledge funding from pharmaceutical firms. However, all my results are robust to the exclusion of those GWASs (Appendix Table F.9).

This presents a unique opportunity to study how firms react to true and false positive associations and how their responses vary based on their domain knowledge.

I use OLS to estimate the following specification at the gene-disease combination level:

$$Y_{i,j,t} = \alpha + \beta Post_t \times GWAS_{i,j} + \gamma GD_{i,j} + \delta_t \times Gene_i + \omega_t \times Disease_j + \epsilon_{i,j,t}, \quad (1)$$

where $Y_{i,j,t}$ is either a measure of investments or innovation outcomes by firms in year t involving gene i and disease j . In alternative specifications, $Y_{i,j,t}$ is coded separately for firms with domain knowledge ($k = 1$) and without it ($k = 0$). $Post_t \times GWAS_{i,j}$ equals one after the first GWAS reports a gene-disease association and equals zero otherwise. $GD_{i,j}$ are fixed effects for the combination of gene i and disease j , which account for pair-specific differentials in research potential. I control for time-varying disease-level differences in market size by including disease-year fixed effects ($\omega_t \times Disease_j$). Similarly, gene-specific time trends ($\delta_t \times Gene_i$) consider the growth of interest in specific genes. Estimates report standard errors clustered two-way by gene and disease.¹⁷ The coefficient of interest β captures the change in $Y_{i,j,t}$ for genes and diseases found correlated in early GWASs relative to those reported later (or never within my sample period).

The main threat to identification is the potential sorting of data-driven predictions on gene-disease pairs that would receive investments even without GWAS findings. However, in this study, this issue is mitigated by two key factors. First, GWASs are distinct from targeted scientific studies in that they scan every human gene. The genome-wide nature of GWASs ensures that there is no gene-level selection conditional on the choice of disease to study, as reflected by the pattern of the associations uncovered (Appendix Figures F.1 and F.2). Second, the timing of the findings is plausibly exogenous to trends in the dependent variable. Indeed, the temporal order in which gene-disease pairs are reported in a GWAS is not influenced by previous inventive activity related to those pairs (Appendix Table F.1 and Figure F.3). In what follows, I exploit this plausibly exogenous within disease timing variation between genes to assess the causal effect of GWASs on firms' early stage innovation investments. Directly testing for the absence of pre-trends in the outcome variables will further confirm the validity of my identification strategy.

5 Results

5.1 Do GWASs Lead to More Investments and Drugs in Aggregate?

I begin by examining the aggregate impact of GWAS findings on pharmaceutical innovation. Given their uncertain reliability, it is unclear whether the publication of gene-disease associations will

¹⁷The precision of the results is robust to alternate methods of clustering standard errors, such as by gene, by disease, or by gene-disease pairs.

increase firms' willingness to invest in a gene-disease combination. Columns 1 and 2 of Table 2 present the baseline results, using patent applications as an indicator for firm investment in a given gene-disease pair. The main finding is a positive impact of GWAS, amounting to an average increase of 125% in patents filed. The design of GWASs should ensure that the effects reported are not due to researchers endogenously publishing on the most promising gene-disease pairs. Figure 3 directly checks the validity of this assumption with an event study version of Equation 1. The plot confirms flat pre-trends and a persistent effect after the GWAS is published, with the estimates stabilizing around the value of the primary estimate of Table 2. Additional heterogeneity analyses in Appendix Table F.5 reveal greater increases in patenting for associations that are statistically more robust or larger in magnitude.

GWASs are known to harbor many false positive associations (MacArthur, 2012). In my sample, 84.3% of the associations involving a protein-coding gene fail to replicate in follow-up GWASs targeting the same disease. The high incidence of misleading predictions would not be problematic if firms could recognize and avoid them. Unfortunately, as shown in Columns 3 and 4 of Table 2, this is not the case. The estimates imply that around 34.8% of the increase in patent applications is directed at non-replicable gene-disease associations. Nevertheless, GWAS findings that were later confirmed, albeit a minority, mustered up to three times more investments than the average effect. This large increase is likely compounded by firms initially investing in false positives and then pivoting away after learning about their mistake.

A case study exemplifies these dynamics. Panel (a) of Figure 4 shows patenting activity for the PLA2G7-myocardial infarction pair after it was reported in a GWAS by Suchindran et al. (2010). While GlaxoSmithKline was already investigating this specific pair, the GWAS led to a spike in investments by new firms. Unbeknownst to them, the finding would later fail to translate into therapeutic advances.¹⁸ In 2014, GlaxoSmithKline announced the failure of two clinical trials for darapladib, a molecule targeting PLA2G7. This led other firms to redirect their investments toward alternative targets (Krieger, 2021). The same dynamics can be seen in an event study regression where I keep only false positive GWAS associations in the sample (Panel (b) of Figure 4). Patenting initially rises and then declines, most likely when firms experimentally discover that the finding is not robust.

The allocation of investments toward false positives hurts firm performance because only valid associations yield successful innovation outcomes. Table 3 indicates that highly cited and high-value patents increase for GWAS associations that are later replicated, but not for the others. The estimates

¹⁸To date, the precise function of PLA2G7 is still unclear. The consensus is that it may be a non-causal biomarker (i.e., an indicator) for the risk of cardiovascular disease, which could contribute to explaining the low Open Targets score (0.045) and the failure to be confirmed in GWAS replications. By comparison, the IL23R-IBD pair uncovered by Duerr et al. (2006), and discussed in Appendix B.2, has an Open Targets score of 0.468 and has been robustly replicated.

relating to drug discovery are also positive but noisier, probably because successful drug development is an infrequent outcome. In Column 4, I follow the approach of Dranove et al. (2022) and weigh each drug by its scientific novelty, captured by the number of previous drugs that adopted the same molecular-targeting design.¹⁹ The coefficient becomes more precise, suggesting that GWAS leads are especially helpful in designing drugs that leverage new scientific approaches. Overall, these results suggest that GWAS can boost pharmaceutical investments, consistent with the first prediction of my theoretical framework. However, their impact on innovation is more nuanced because they also produce several false positive leads. If these are not avoided, they can lead to a waste of resources and reduce firms' technological performance.

5.2 Data-Driven Predictions, Domain Knowledge, and Firm Performance

This subsection explores how data-driven predictions interact with firms' domain knowledge. Previous studies have used R&D expenditures (Cohen and Levinthal, 1990; Fabrizio, 2009) or the number of corporate publications (Cockburn and Henderson, 1998; Gambardella, 1992) to identify a firm's knowledge. However, these proxies do not capture the specific domains in which the firm has the knowledge to recognize opportunities. I make progress on this issue by recording the specific genes that each firm has researched in its publications. This empirical approximation is grounded in the specificities of the setting. In my interviews with industry professionals, they confirmed that firms in the bio-pharma sector tend to specialize in genetic space.²⁰ Some firms go as far as making the selection of their target genes their fundamental competitive hypothesis. For instance, the biotech company Denali Therapeutics even mentioned its strategic focus on what it calls "degenogenes" in its IPO filings (Appendix D.3). My empirical approach captures the fact that when a new gene-disease association involving one of these genes is published, Denali's scientists will have the expertise to assess it.²¹

My research design exploits between-firm variation in genetic knowledge predating GWASs, as schematically illustrated in Appendix Figure F.4 (Cattani, 2005). Columns 1 and 3 of Table 4 present results comparing patent applications filed by firms with and without previous publications on the gene involved in a GWAS finding. While both groups of firms significantly increase their investments,

¹⁹Dranove et al. (2022) refer to the term "molecular-targeting design" as the mechanism by which drugs produce a pharmacological effect (e.g., darapladib inhibits Lp-PLA₂, thus being an "Lp-PLA₂ inhibitor"). This measure captures which drugs are more novel from a scientific perspective. Other approaches in the literature include considering the molecule's chemical novelty instead of its biological mechanism (Krieger et al., 2022).

²⁰The patent portfolio of the median firm in my sample spans only 19 genes, a number that further decreases to 11 if considering only patents that will eventually be granted (Appendix D).

²¹As one of my interviewees at Denali put it, whether to pursue or not a new genetic target appearing in the literature is an assessment based on their extensive domain expertise: "We are so entrenched in the neuro field that we kind of know off-the-cuff" (Interview, 17 October 2022).

there is a stronger reaction by firms lacking gene-specific knowledge. Compared to the sample mean, the regression coefficients imply that patenting increases by one-third more among firms that had not previously researched the gene before the GWAS was published. This confirms the second empirical implication of my theoretical framework: firms without domain knowledge will invest more after observing identical data-driven predictions.²²

Next, I investigate if firms with relevant domain knowledge can correctly rank data-driven predictions and avoid those likely to be false positives (Arora and Gambardella, 1994; Cohen and Levinthal, 1994). Appendix Figure F.6 descriptively shows the reaction of firms with and without domain knowledge to newly published gene-disease associations. Strikingly, the increase in patenting based on a false positive is over three times larger for firms without past research on the gene. I validate the statistical significance of this pattern with difference-in-differences regressions (Columns 2 and 4 of Table 4). The results confirm that the increase in patenting on false positive associations is significant only for firms without prior research on the gene, supporting the idea that such firms lack the knowledge to evaluate the findings. Mapping this result back to my theoretical framework, patent applications building on false positive associations constitute errors of omission that distort resources away from pursuing valuable leads.

However, the above results leave open the possibility that domain knowledge simply increases skepticism of data-driven predictions (Boudreau et al., 2016; Denrell and Fang, 2010). Under this alternative explanation, avoiding false positives might mechanically result from lower investments made by conservative firms, potentially at the cost of curtailing the exploration of valuable opportunities. I rule out this possibility using the Open Targets data to capture the underlying therapeutic value of gene-disease pairs with a continuous score.²³ Figure 5 shows that firms with domain knowledge are better at selecting the most valuable opportunities among the data-driven predictions. This result also implies that firms without domain expertise are more likely to make omission errors despite investing quantitatively more, while firms with genetic expertise are less likely to miss out on breakthroughs despite investing quantitatively less.²⁴ Taken together, the results support the third empirical implication of my theoretical framework: domain knowledge increases the efficiency of investments, allowing firms to focus their efforts on promising data-driven opportunities.

²²Since larger firms also have broader knowledge bases, one might assume that this result reflects organizational inertia. To rule out this concern, I repeat my analysis for small and large firms separately. I find the same result in both sub-samples: GWAS generate stronger reactions from firms without domain knowledge regardless of their size (Appendix Figure F.5).

²³While available only for a subset of gene-disease pairs in my sample, this metric offers a way to rank the “ground truth” value of the combinations highlighted by GWAS associations. See Appendix E.2 for a detailed discussion.

²⁴The better allocation of early stage investments by firms with domain knowledge translates into more drug molecules entering the discovery stage (Appendix Table F.6). However, these results should be interpreted with caution since firms in this industry frequently license promising compounds for later stages of drug development (Gambardella et al., 1995).

5.3 Robustness Checks

Several robustness checks further validate these findings. First, I replicate the main results by exploiting a unique feature of GWASs. After spotting genetic mutations correlated with the disease, scientists must map the mutation to the correct gene (Visscher et al., 2017).²⁵ Especially at the beginning of the sample, the gene reported in the original study was sometimes incorrect but later reclassified with more precision by the curators of the GWAS Catalog. In my data, there are 9,273 such cases. Confirming my primary results, Table 5 shows that only firms without domain knowledge increase investments based on these wrongly reported data-driven associations. This additional test has the advantage of not depending on my method for identifying false positives in the primary analysis.

Second, if domain knowledge truly helps in assessing the value of ambiguous data-driven predictions, then its beneficial effects should be proportionally higher when the GWAS findings are harder to evaluate. For instance, evaluating an association should be more difficult when firms cannot leverage publicly available scientific information or when the GWAS reports too many new findings to validate. Indeed, firms without domain knowledge are more likely to fall into the trap of false positives when the findings involve less-studied genes (Appendix Table F.7), but they make fewer mistakes when the GWAS reports fewer associations to evaluate (Appendix Table F.8). Results are also robust to the exclusion of GWASs co-authored with industry researchers, suggesting that social connections with scientists are not a mechanism through which firms active in research can evaluate GWAS findings (Appendix Table F.9).

Third, I examine whether patenting based on false positive findings is motivated by strategic reasons, implying that doing so is not detrimental to firms' performance, even if it does not result in discovering new drugs. A few pieces of evidence help rule out this possible explanation. First, using two alternative methods to identify patents likely to be motivated by strategic concerns (Abrams et al., 2018; Righi and Simcoe, 2020), I find the firms are more likely to apply for such patents based on true positive associations (Columns 1 and 2 of Appendix Table F.10). Intuitively, this finding is consistent with the idea that there is little strategic value in targeting associations that do not have therapeutic value. Second, patents targeting false positive leads are less likely to end up in litigation or being renewed by the firm, suggesting their actual strategic usefulness is low (Columns 3 and 4 of Appendix Table F.10).

Finally, it could be that pursuing false positive predictions results in learning spillovers for the firms. This possibility would imply that the cost of selecting the wrong target is much lower for the

²⁵More specifically, they have to identify the location of molecular markers on the genome relative to the coordinates of known genes (Vaughan and Srinivasasainagendra, 2013). For earlier GWASs, this step was not fully routinized, and it was not uncommon to attribute the mutation to the wrong gene (usually a neighboring one). Wrongly mapped associations are 62.3% less likely to be replicated compared to correctly reported findings.

firm, consistent with recent work in a similar context (Frankel et al., 2023). To test this, I examine innovation outcomes for gene-disease pairs involving the same gene and a similar disease of the GWAS associations.²⁶ Appendix Table F.11 shows no evidence of spillover effects on innovation outcomes, suggesting that the beneficial effects of experimentally learning from mistakes are likely low in this context.

6 Firm-Level Mechanisms

The previous section showed that domain knowledge allows firms to recognize false positive predictions and focus investments on valuable opportunities. In this section, I investigate the mechanisms behind these results. First, although my analysis used corporate publications to find the specific genes on which firms had expertise, the results might reflect organizational capabilities that also apply to unstudied genes. Second, the ability to recognize opportunities could stem from different organizational learning mechanisms (Di Stefano et al., 2023; Zollo and Winter, 2002). In particular, domain knowledge could either decrease the cost of directly testing predictions (“online learning”) or help firms apply theoretical principles to evaluate predictions (“offline learning”) (Gavetti and Levinthal, 2000). In what follows, I provide suggestive evidence to explore these mechanisms.

6.1 Empirical Strategy

The findings in the previous section are drawn from split-sample regressions at the gene-disease level that exploit between-firm variation in gene knowledge. In this section, I exploit an alternative source of variation: the within-firm heterogeneity in the available knowledge about different human genes. Panel (a) of Figure 6 provides a graphical intuition for this research design. The basic intuition is akin to the preceding section since firms might better understand the biological mechanisms of specific genes than others due to their past research choices. However, this research design benefits from the ability to include firm fixed effects in the regression models. As a result, the estimates effectively control for all non-time-varying unobservable firm characteristics that correlate with their ability to recognize valuable opportunities.

For this analysis, I construct a dataset where each observation represents a potential firm investment in a gene-disease pair supported by GWAS evidence. This data structure allows me to investigate the investment decisions of firms that observe data-driven predictions. I assume that firms consider all GWAS findings about diseases in which they have previously invested.²⁷ Pharmaceutical firms are usually active in specific diseases (that can be thought of as markets), which gives a convenient way

²⁶The hierarchical nature of the MeSH disease taxonomy provides an easy way to find diseases that share the same etiology and biological mechanisms. More specifically, I consider a disease (i.e., a four-digit MeSH code) similar to the one targeted by the GWAS if it shares the same “parent” disease (i.e., the same three-digit MeSH code).

²⁷This choice is consistent with what emerged from my interviews: “I would say we definitely pay attention to GWAS

to assemble the risk set of GWAS associations that are likely to catch their attention (Bikard, 2018; Krieger, 2021). While the earlier regressions at the gene-disease level offer a more accurate estimate of GWASs’ aggregate impact, the within-firm design provides a tighter way to isolate the impact of domain knowledge on firm investments. However, variation in firms’ past research portfolios is not random, and so the following results should be interpreted as suggestive correlations.

I use OLS to estimate the following specification at the firm-gene-disease level:

$$Y_{f,i,j} = \alpha + \beta \text{Domain Knowledge}_{f,i,j} + \mu_{f,j} + \theta_{i,j} + \epsilon_{f,i,j}, \quad (2)$$

where $Y_{f,i,j}$ is the number of patents by firm f targeting the gene-disease combination $\langle i, j \rangle$. $\text{Domain Knowledge}_{f,i,j}$ is a dummy that equals one if the firm possesses domain knowledge about the gene involved in the gene-disease association. $\mu_{f,j}$ are firm-by-disease dummies, accounting for firm specialization in specific disease areas as well as potential differences in how firms weigh errors of omission against errors of commission (Csaszar, 2013).²⁸ $\theta_{i,j}$ is a time effect that controls for the year when the gene-disease association is published. The coefficient of interest β captures how a firm’s investment decisions change when it can leverage genetic knowledge to evaluate the GWAS finding. Estimates report standard errors clustered at the firm level.

6.2 How Does Domain Knowledge Help?

I begin by replicating the main results of the paper using within-firm variation. Table 6 presents results consistent with the gene-disease-level evidence. Column 1 shows that firms are significantly less likely to invest in a false positive when they can draw upon their knowledge of gene biology.²⁹ This result is robust to including stringent firm \times disease fixed effects that better account for firm specialization in specific disease areas (Column 2). Columns 3 and 4 symmetrically show that domain knowledge increases the likelihood of recognizing valuable data-driven opportunities. Using the Open Targets score, Panel (b) of Figure 6 visually displays how theoretical knowledge enables firms to be more efficient and targeted in their investments, consistent with my theoretical framework. In

studies. Any time there is a GWAS study that, you know, that chose an indication that we’re interested in, we definitely pay attention” (Interview, 9 December 2022). In the appendix, I experiment with alternative ways to define the appropriate risk set of GWASs evaluated by a firm, finding consistent results (Appendix Tables F.13 and F.14).

²⁸For instance, an alternative explanation of my previous findings could be that firms follow different strategies regarding false positives. Firms that prioritize minimizing the risk of missing a valuable target might be willing to incur several commission errors, and vice versa. Including firm fixed effects helps rule out that the results are driven by firms pursuing different strategies regarding the weight given to type I vs. type II errors. This is because if the relative disutility of a false negative relative to a false positive is a fixed firm-specific parameter, the firm-fixed effect will absorb it. The additional inclusion of firm-by-disease fixed effects would also consider the more likely case that this parameter depends on the specific disease (e.g., for Abbvie, missing out on a Crohn’s disease target might be worse than doing so for other diseases where it does not specialize).

²⁹Note that including firm fixed effects changes the interpretation of the estimates. In this case, the coefficient captures the decrease in patenting based on false positive associations when a firm has domain knowledge, compared to other false positives where the firm lacks such expertise.

addition to confirming earlier analyses, these findings suggest that the main results are not due to organizational-level capabilities but rather to the specific distribution of knowledge across genes.

Next, I explore what type of domain knowledge is conducive to my results. A long tradition in the management literature has emphasized how organizations can learn routines to perform even complex operations, such as selecting innovation projects (Dosi et al., 2001; March and Simon, 1958; Nelson and Winter, 1982). This line of work emphasizes the tacit and often poorly understood nature of these heuristics (Zollo and Winter, 2002). In contrast, a growing strand of research highlights the importance of developing theory-based understandings to guide action (Camuffo et al., 2023b; Ehrig and Schmidt, 2022; Gavetti and Levinthal, 2000; Wuebker et al., 2023). Recent evidence shows that experience unlocks higher performance only when accompanied by an understanding of the causal relationships between antecedents and outcomes (Di Stefano et al., 2023). In my setting, this would suggest that the ability to recognize valuable GWAS findings stems from understanding the underlying mechanisms of what makes an opportunity valuable rather than from mere practice or testing capabilities that permit cheap validation of alternatives.

I categorize firms' knowledge types to shed some light on this issue by distinguishing firms' publications that denote testing capabilities from those focused on the genetic mechanisms of diseases. Following the approach of Azoulay et al. (2021), I record which firm publications have a *clinical* focus and which ones are *translational* research that applies basic genetic research to therapeutic purposes.³⁰ I also code firms that have conducted disease-oriented research or filed patents that target the same disease but not the same gene. I then run separate regressions using each of these proxies for domain knowledge. Figure 7 shows that only translational knowledge aids in correctly avoiding false positive associations among the GWAS findings. Instead, clinical research or generic experience with the disease do not confer any advantage when assessing GWAS findings, nor do general-purpose inventive capabilities (Arora and Gambardella, 1994). This evidence suggests that my results are not driven by cheaper experimentation costs but by the ability to apply basic scientific principles to rule out spurious gene-disease associations (Di Stefano et al., 2023; Fleming and Sorenson, 2004).

Finally, one might question why firms do not wait for the uncertainty surrounding a gene-disease association to dissipate. For instance, a firm could hold back until further research validates the finding or strategically delay investments to observe competitors' actions (Krieger, 2021). Appendix Table F.15 shows suggestive evidence that waiting before investing in a GWAS finding reduces the likelihood

³⁰Intuitively, I divide domain knowledge into testing knowledge and translational knowledge. Firms with clinical experience can cheaply learn the plausibility of GWAS associations with wet lab testing ("online learning"), while those with translational experience can evaluate findings based on their theoretical understanding of genetics ("offline learning"). The data on clinical publications are from NIH's iCite database while the definition of translational research is taken from Azoulay et al. (2021).

of targeting a false positive association. However, it seems to come at the cost of lower-impact patents and lower chances of discovering a drug (albeit the results are noisier in this case). Moreover, since patenting results in claiming intellectual property space, being late might result in fewer avenues for future innovation (Eggers, 2012). This result illustrates how data-driven predictions generate a tension between carefully validating them and investing ahead of competitors. By enabling the “offline” evaluation of alternatives (Gavetti and Levinthal, 2000; Levinthal, 2021), domain knowledge becomes a crucial means to benefit from data-driven predictions.

7 Conclusion

While the field of strategic management has extensively researched how firms can exploit innovation opportunities, the rapid diffusion of big data and predictive tools is transforming how firms identify them. In particular, data-driven predictions offer new avenues for assessing the value of potential investments instead of leveraging domain knowledge. This paper provides some of the first empirical evidence of the consequences of this phenomenon. I introduce a new conceptual and empirical framework to analyze how data-driven predictions and firms’ domain knowledge interact to shape innovation performance. Moving beyond the average effects of data-driven predictions in triggering investments, I show how the distribution of false positive findings is crucial to understanding their heterogeneous impact on innovation outcomes.

My results underscore the continuing importance of domain knowledge in the innovation process but with a new role: since the abundance of data-driven predictions is not accompanied by a generalized reduction in the cost of validating them, domain knowledge becomes essential to avoid wasting resources on dead ends. My framework also highlights the boundary conditions of this reasoning. If predictive tools were perfectly accurate or validating their predictions was costless, the need to prioritize among data-driven predictions would be lower. However, this is not the case in many innovation and strategy contexts. My theory reconciles conflicting results in previous research by suggesting that when data-driven findings are straightforward to interpret, firms can rank them without relying on domain knowledge, thus benefiting new entrants (Nagaraj, 2022). Conversely, one should expect that understanding the domain is a necessary complement to benefit from data predictions in complex domains (Kao, 2023).

Further, the innovation process is often depicted as a search through landscapes of technological combinations (Fleming and Sorenson, 2004; Levinthal, 2021). Traditional research in this area frequently uses computer simulations to characterize optimal exploration strategies and their organizational implications (e.g., Csaszar 2013). In contrast, this paper joins a recent effort to try to empirically characterize organizational search (Kang, 2023; Nagaraj, 2022) and how different search strategies

map into innovation outcomes (Tranchoero, 2023). It conceptualizes and exemplifies a new measurement strategy to depict empirical landscapes using text-based methods, opening up new avenues for an empirical re-examination of the seminal contributions in the field of organizational search (Dosi et al., 2001; March and Simon, 1958; Nelson and Winter, 1982).

My research has practical implications for managers and policymakers navigating the changing landscape of AI and predictive technologies. To cut through the hype surrounding big data, it is essential to understand what predictive technologies can and cannot do (Agrawal et al., 2018). Specifically, this paper proposes that while data-driven predictions are valuable in shortlisting potential opportunities—which is no small feat—they are not without limitations. This is especially true in the context of what Camuffo et al. (2023a) call “low-frequency/high-impact” organizational decisions. In these contexts, the source of competitive advantage lies in the ability to interpret and evaluate what the data suggest. Managers should be mindful that predictive technologies are no free lunch since their benefits may be most evident to organizations that already have a solid base of domain knowledge. Similarly, governments funding large-scale data efforts, which are especially common in the bio-pharmaceutical world, should note how they might affect competition dynamics. Instead of being a panacea to foster entry and market dynamism, large-scale mapping efforts might contribute to stifling competition. Policymakers may need to take action to avoid unintended anti-competitive effects from databases that require complementary assets to be exploited.

More broadly, predictive technologies allow innovators to identify promising areas of the technological landscape, shifting the focus from theory-driven searches to inductive methods (Strasser, 2019). As a result, data-based correlations and extrapolations increasingly guide organizational decision-making instead of logical reasoning and domain knowledge. While the popular press often hails the benefits of these developments (Anderson, 2008; Schmidt, 2023), the risk of not understanding why innovations work could be the accumulation of an “intellectual debt” (Kim, 2023b; Zittrain, 2019). This risk is especially evident in fields like drug discovery, where AI is used to find drugs with unknown mechanisms of action, thus preventing the anticipation of potential side effects (Heaven, 2023). With firms hungry to find innovation opportunities while progressively shying away from investing in domain knowledge (Arora et al., 2018), more research is needed to understand the nuanced implications of putting data ahead of theoretical understanding (Tranchoero, 2023). This paper serves as a first step in this direction.

Finally, despite the contributions outlined above, a few limitations of this paper should be acknowledged. First, fully capturing the aggregate effects of GWASs on innovation is beyond its scope. Data-driven predictions in genomics might give rise to a “streetlight effect” (Hoelzemann et al., 2023), potentially diverting investments into suboptimal targets with negative consequences for social welfare.

Moreover, since I do not observe costs directly, the corresponding implications for firm profitability are also beyond the scope of this study.³¹ Second, this paper does not explore the strategic responses of firms without domain knowledge. It is conceivable that a market for expertise could emerge, where firms outsource not the generation of ideas (Piezunka and Dahlander, 2015) but expert judgment (Luo et al., 2021) to assess the potential of data-driven opportunities and avoid false positives. Finally, while my research design focuses on holding the content of data-driven predictions constant to study firm responses, another critical dimension to explore is heterogeneity in prediction quality. Identifying which firms can obtain better predictions and understanding the organizational factors that influence this are first-order questions. Further exploration of these ideas is an exciting avenue for future research.

³¹However, back-of-the-envelope analyses suggest sizable opportunity costs tied to pursuing false positive GWAS findings, and additional tests rule out learning spillovers from these false positives.

References

- ABRAMS, D. S., U. AKCIGIT, AND J. GRENNAN (2018): “Patent value and citations: Creative destruction or strategic disruption?” *NBER wp 19647*.
- ADNER, R., P. PURANAM, AND F. ZHU (2019): “What is different about digital strategy? From quantitative to qualitative change,” *Strategy Science*, 4, 253–261.
- AGRAWAL, A., J. GANS, AND A. GOLDFARB (2018): *Prediction machines: The simple economics of artificial intelligence*, Harvard Business Press.
- AGRAWAL, A. K., J. MCHALE, AND A. OETTL (2023): “Artificial intelligence and scientific discovery: A model of prioritized search,” *NBER wp 31558*.
- AGUIAR, L. AND J. WALDFOGEL (2018): “Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music,” *Journal of Political Economy*, 126, 492–524.
- ALLEN, R. AND P. CHOUDHURY (2022): “Algorithm-augmented work and domain experience: The countervailing forces of ability and aversion,” *Organization Science*, 33, 149–169.
- ALLEN, R. AND R. McDONALD (2022): “Methodological pluralism and innovation in data-driven organizational cultures,” *University of Washington and Harvard Business School*.
- ANDERSON, C. (2008): “The end of theory: The data deluge makes the scientific method obsolete,” *Wired magazine*, 16, 16–07.
- ARORA, A., S. BELENZON, AND A. PATACCONI (2018): “The decline of science in corporate R&D,” *Strategic Management Journal*, 39, 3–32.
- ARORA, A. AND A. GAMBARDELLA (1994): “Evaluating technological information and utilizing it: Scientific knowledge, technological capability, and external linkages in biotechnology,” *Journal of Economic Behavior & Organization*, 24, 91–114.
- AZOULAY, P., W. H. GREENBLATT, AND M. L. HEGGENESS (2021): “Long-term effects from early exposure to research: Evidence from the NIH “Yellow Berets”,” *Research Policy*, 50, 104332.
- BALASUBRAMANIAN, N., Y. YE, AND M. XU (2022): “Substituting human decision-making with machine learning: Implications for organizational learning,” *Academy of Management Review*, 47, 448–465.
- BERMAN, R. AND A. ISRAELI (2022): “The value of descriptive analytics: Evidence from online retailers,” *Marketing Science*, 41, 1074–1096.
- BERMAN, R. AND C. VAN DEN BULTE (2022): “False discovery in A/B testing,” *Management Science*, 68, 6762–6782.
- BESSEN, J., S. M. IMPINK, L. REICHENSBERGER, AND R. SEAMANS (2022): “The role of data for AI startup growth,” *Research Policy*, 51, 104513.
- BHATIA, A. AND G. DUSHNITSKY (2023): “The future of venture capital? Insights into data-driven VCs,” *California Management Review*.
- BIKARD, M. (2018): “Made in academia: The effect of institutional origin on inventors’ attention to science,” *Organization Science*, 29, 818–836.
- BOUDREAU, K. J., E. C. GUINAN, K. R. LAKHANI, AND C. RIEDL (2016): “Looking across and looking beyond the knowledge frontier: Intellectual distance, novelty, and resource allocation in science,” *Management Science*, 62, 2765–2783.
- BRYNJOLFSSON, E., W. JIN, AND K. MCELHERAN (2021): “The power of prediction: Predictive analytics, workplace complements, and business performance,” *Business Economics*, 56, 217–239.

- BRYNJOLFSSON, E., D. LI, AND L. R. RAYMOND (2023): “Generative AI at work,” *NBER wp 31161*.
- CAMUFFO, A., A. GAMBARELLA, AND A. PIGNATARO (2023a): “Framing strategic decisions in the digital world,” *Strategic Management Review*, 4, 127–160.
- (2023b): “Theory-driven strategic management decisions,” *CEPR DP17664*.
- CAO, R., R. M. KONING, AND R. NANDA (2023): “Biased sampling of early users and the direction of startup innovation,” *Management Science*, Forthcoming.
- CATTANI, G. (2005): “Preadaptation, firm heterogeneity, and technological performance: A study on the evolution of fiber optics, 1970–1995,” *Organization Science*, 16, 563–580.
- CHOI, S., N. KIM, J. KIM, AND H. KANG (2022): “How does artificial intelligence improve human decision-making? Evidence from the AI-powered go program,” *Mimeo*.
- CHOUHURY, P., E. STARR, AND R. AGARWAL (2020): “Machine learning and human capital complementarities: Experimental evidence on bias mitigation,” *Strategic Management Journal*, 41, 1381–1411.
- COCKBURN, I. M., R. HENDERSON, AND S. STERN (2019): “The impact of artificial intelligence on innovation,” in *The Economics of Artificial Intelligence: An Agenda*, Chicago: Chicago University Press, 115–146.
- COCKBURN, I. M. AND R. M. HENDERSON (1998): “Absorptive capacity, coauthoring behavior, and the organization of research in drug discovery,” *The Journal of Industrial Economics*, 46, 157–182.
- COHEN, W. M. AND D. A. LEVINTHAL (1990): “Absorptive capacity: A new perspective on learning and innovation,” *Administrative Science Quarterly*, 35, 128–152.
- (1994): “Fortune favors the prepared firm,” *Management Science*, 40, 227–251.
- COWGILL, B. (2018): “Bias and productivity in humans and algorithms: Theory and evidence from resume screening,” *Columbia Business School*.
- CRISCUOLO, P., L. DAHLANDER, T. GROHSJEAN, AND A. SALTER (2017): “Evaluating novelty: The role of panels in the selection of R&D projects,” *Academy of Management Journal*, 60, 433–460.
- CSASZAR, F. A. (2013): “An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation,” *Organization Science*, 24, 1083–1101.
- DENRELL, J. AND C. FANG (2010): “Predicting the next big thing: Success as a signal of poor judgment,” *Management Science*, 56, 1653–1667.
- DI STEFANO, G., F. GINO, G. P. PISANO, AND B. R. STAATS (2023): “Learning by thinking: How reflection can spur progress along the learning curve,” *Management Science*, Forthcoming.
- DING, Y., M. SONG, J. HAN, Q. YU, E. YAN, L. LIN, AND T. CHAMBERS (2013): “Entitymetrics: Measuring the impact of entities,” *PloS one*, 8, e71416.
- DOSHI, A. R. AND O. HAUSER (2023): “Generative artificial intelligence enhances creativity,” *Available at SSRN*.
- DOSI, G., R. R. NELSON, AND S. G. WINTER (2001): *The nature and dynamics of organizational capabilities*, Oxford University Press.
- DRANOVE, D., C. GARTHWAITE, AND M. HERMOSILLA (2022): “Does consumer demand pull scientifically novel drug innovation?” *The RAND Journal of Economics*, 53, 590–638.
- DUERR, R. H., K. D. TAYLOR, S. R. BRANT, J. D. RIOUX, M. S. SILVERBERG, M. J. DALY, ET AL. (2006): “A genome-wide association study identifies IL23R as an inflammatory bowel disease gene,” *Science*, 314, 1461–1463.

- EGGERS, J. P. (2012): “Falling flat: Failed technologies and investment under uncertainty,” *Administrative Science Quarterly*, 57, 47–80.
- EGGERS, J. P. AND S. KAPLAN (2009): “Cognition and renewal: Comparing CEO and organizational effects on incumbent adaptation to technical change,” *Organization Science*, 20, 461–477.
- EHRIG, T. AND J. SCHMIDT (2022): “Theory-based learning and experimentation: How strategists can systematically generate knowledge at the edge between the known and the unknown,” *Strategic Management Journal*, 43, 1287–1318.
- FABRIZIO, K. R. (2009): “Absorptive capacity and the search for innovation,” *Research Policy*, 38, 255–267.
- FLEMING, L. (2001): “Recombinant uncertainty in technological search,” *Management Science*, 47, 117–132.
- FLEMING, L. AND O. SORENSON (2004): “Science as a map in technological search,” *Strategic Management Journal*, 25, 909–928.
- FRANKEL, A. P., J. L. KRIEGER, D. LI, AND D. PAPANIKOLAOU (2023): “Evaluation and learning in R&D investment,” *NBER wp 31290*.
- FREEDMAN, L. P., I. M. COCKBURN, AND T. S. SIMCOE (2015): “The economics of reproducibility in preclinical research,” *PLoS biology*, 13, e1002165.
- GAESSLER, F. AND H. PIEZUNKA (2023): “Training with AI: Evidence from chess computers,” *Strategic Management Journal*, Forthcoming.
- GAMBARDELLA, A. (1992): “Competitive advantages from in-house scientific research: The US pharmaceutical industry in the 1980s,” *Research Policy*, 21, 391–407.
- GAMBARDELLA, A. ET AL. (1995): *Science and innovation: The US pharmaceutical industry during the 1980s*, Cambridge University Press.
- GANS, J. S., S. STERN, AND J. WU (2019): “Foundations of entrepreneurial strategy,” *Strategic Management Journal*, 40, 736–756.
- GAVETTI, G. AND D. LEVINTHAL (2000): “Looking forward and looking backward: Cognitive and experiential search,” *Administrative Science Quarterly*, 45, 113–137.
- GHOSH, S. (2021): “Experimental Approaches to Strategy and Innovation,” Ph.D. thesis, Harvard University.
- GLAESER, E. L., A. HILLIS, H. KIM, S. D. KOMINERS, AND M. LUCA (2022): “Decision authority and the returns to algorithms,” *Harvard Business School and INSEAD*.
- GOLDSTEIN, D. B. (2009): “Common genetic variation and human traits,” *New England Journal of Medicine*, 360, 1696.
- GREVE, H. R. (2003): *Organizational learning from performance feedback: A behavioral perspective on innovation and change*, Cambridge University Press.
- GRUBER, M., I. C. MACMILLAN, AND J. D. THOMPSON (2008): “Look before you leap: Market opportunity identification in emerging technology firms,” *Management Science*, 54, 1652–1665.
- HEAVEN, W. D. (2023): “AI is dreaming up drugs that no one has ever seen. Now we’ve got to see if they work,” *MIT Technology Review*.
- HERMOSILLA, M. AND J. LEMUS (2019): “Therapeutic translation of genomic science,” in *Economic Dimensions of Personalized and Precision Medicine*, Chicago University Press.
- HOELZEMANN, J., G. MANSO, A. NAGARAJ, AND M. TRANCHERO (2023): “The streetlight effect in data-driven exploration,” *UC Berkeley and University of Vienna*.

- HOFFMAN, M., L. B. KAHN, AND D. LI (2018): “Discretion in hiring,” *The Quarterly Journal of Economics*, 133, 765–800.
- IANSITI, M. AND K. R. LAKHANI (2020): *Competing in the age of AI: Strategy and leadership when algorithms and networks run the world*, Harvard Business Press.
- JIA, N., X. LUO, Z. FANG, AND C. LIAO (2023): “When and how artificial intelligence augments employee creativity,” *Academy of Management Journal*, Forthcoming.
- JIN, W. AND K. MCELHERAN (2019): “Economies before scale: Learning, survival, and performance of young plants in the age of cloud computing,” *Rotman School of Management wp 3112901*.
- KANG, S. (2023): “Objects may be closer than they appear: The impact of public scientific map on organizational search,” *UC Santa Barbara*.
- KAO, J. (2023): “Charted territory: Mapping the cancer genome and R&D decisions in the pharmaceutical industry,” *UCLA Anderson*.
- KAPLAN, S., F. MURRAY, AND R. HENDERSON (2003): “Discontinuities and senior management: Assessing the role of recognition in pharmaceutical firm response to biotechnology,” *Industrial and Corporate Change*, 12, 203–233.
- KAPLAN, S. AND K. VAKILI (2015): “The double-edged sword of recombination in breakthrough innovation,” *Strategic Management Journal*, 36, 1435–1457.
- KIM, H. (2023a): “The value of competitor information: Evidence from a field experiment,” *Management Science*, Forthcoming.
- KIM, S. (2023b): “Shortcuts to innovation: The use of analogies in knowledge production,” *Columbia Business School*.
- KLEINBERG, J., H. LAKKARAJU, J. LESKOVEC, J. LUDWIG, AND S. MULLAINATHAN (2018): “Human decisions and machine predictions,” *The Quarterly Journal of Economics*, 133, 237–293.
- KRIEGER, J., D. LI, AND D. PAPANIKOLAOU (2022): “Missing novelty in drug development,” *The Review of Financial Studies*, 35, 636–679.
- KRIEGER, J. AND R. NANDA (2022): “Are transformational ideas harder to fund? Resource allocation to R&D projects at a global pharmaceutical firm,” *Harvard Business School Working Paper 23-014*.
- KRIEGER, J. L. (2021): “Trials and terminations: Learning from competitors’ R&D failures,” *Management Science*, 67, 5525–5548.
- LEVINTHAL, D. A. (2021): “From arms to trees: Opportunity costs and path dependence and the exploration-exploitation tradeoff,” *Strategy Science*, 6, 331–337.
- LEVINTHAL, D. A. AND J. G. MARCH (1993): “The myopia of learning,” *Strategic Management Journal*, 14, 95–112.
- LOU, B. AND L. WU (2021): “AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms,” *MIS Quarterly*, 45.
- LUO, H., J. MACHER, AND M. WAHLEN (2021): “Judgment aggregation in creative production: Evidence from the movie industry,” *Management Science*, 67, 6358–6377.
- MACARTHUR, D. (2012): “Face up to false positives,” *Nature*, 487, 427–428.
- MARCH, J. G. AND H. SIMON (1958): *Organizations*, John Wiley and Sons.
- MARIGORTA, U. M., J. A. RODRÍGUEZ, G. GIBSON, AND A. NAVARRO (2018): “Replicability and prediction: Lessons and challenges from GWAS,” *Trends in Genetics*, 34, 504–517.

- MULLAINATHAN, S. AND A. RAMBACHAN (2023): “From predictive algorithms to automatic generation of anomalies,” Available at SSRN 4443738.
- NAGARAJ, A. (2022): “The private impact of public data: Landsat satellite maps increased gold discoveries and encouraged entry,” *Management Science*, 68, 564–582.
- NELSON, M. R., H. TIPNEY, J. L. PAINTER, J. SHEN, P. NICOLETTI, Y. SHEN, ET AL. (2015): “The support of human genetic evidence for approved drug indications,” *Nature Genetics*, 47, 856–860.
- NELSON, R. R. AND S. WINTER (1982): *An evolutionary theory of economic change*, Belknap Press.
- NICHOLSON, J. M., M. MORDAUNT, P. LOPEZ, A. UPPALA, ET AL. (2021): “Scite: A smart citation index that displays the context of citations and classifies their intent using deep learning,” *Quantitative Science Studies*, 2, 882–898.
- NOY, S. AND W. ZHANG (2023): “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 381, 187–192.
- PIEZUNKA, H. AND L. DAHLANDER (2015): “Distant search, narrow attention: How crowding alters organizations’ filtering of suggestions in crowdsourcing,” *Academy of Management Journal*, 58, 856–880.
- PURANAM, P. (2019): “The theorist as an endangered species?” *Journal of Marketing Behavior*, 4, 43–48.
- REAY, W. R. AND M. J. CAIRNS (2021): “Advancing the use of genome-wide association studies for drug repurposing,” *Nature Reviews Genetics*, 22, 658–671.
- REICH, D. E. AND E. S. LANDER (2001): “On the allelic spectrum of human disease,” *TRENDS in Genetics*, 17, 502–510.
- RIGHI, C. AND T. SIMCOE (2020): *Patenting inventions or inventing patents? Strategic use of continuations at the USPTO*, NBER wp 27686.
- SCHMIDT, E. (2023): “This is how AI will transform the way science gets done,” *MIT Technology Review*.
- SQUICCIARINI, M., H. DERNIS, AND C. CRISCUOLO (2013): “Measuring patent quality: Indicators of technological and economic value,” *OECD STI Working Papers*, No. 2013/03.
- STOKES, J. M., K. YANG, K. SWANSON, W. JIN, A. CUBILLOS-RUIZ, N. M. DONGHIA, C. R. MACNAIR, S. FRENCH, ET AL. (2020): “A deep learning approach to antibiotic discovery,” *Cell*, 180, 688–702.
- STOREY, J. D. AND R. TIBSHIRANI (2003): “Statistical significance for genomewide studies,” *Proceedings of the National Academy of Sciences*, 100, 9440–9445.
- STRASSER, B. J. (2019): *Collecting experiments*, University of Chicago Press.
- STRUCK, T. J., B. K. MANNAKEE, AND R. N. GUTENKUNST (2018): “The impact of genome-wide association studies on biomedical research publications,” *Human Genomics*, 12, 1–9.
- SUCHINDRAN, S., D. RIVEDAL, J. R. GUYTON, T. MILLEDGE, X. GAO, ET AL. (2010): “Genome-wide association study of Lp-PLA2 activity and mass in the Framingham Heart Study,” *PLoS genetics*, 6, e1000928.
- SUN, M., A. GOLDFARB, AND M. SHI (2021): “Data-driven decision-making in sales: Can marketing analytics enhance sales performance?” *University of Toronto*.
- TAM, V., N. PATEL, M. TURCOTTE, Y. BOSSÉ, G. PARÉ, AND D. MEYRE (2019): “Benefits and limitations of genome-wide association studies,” *Nature Reviews Genetics*, 20, 467–484.
- TAYLOR, A. AND H. R. GREVE (2006): “Superman or the fantastic four? Knowledge combination and experience in innovative teams,” *Academy of Management Journal*, 49, 723–740.

- TEECE, D. J. (2007): “Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance,” *Strategic Management Journal*, 28, 1319–1350.
- TRANCHERO, M. (2023): “Data-driven search and innovation in well-defined technological landscapes,” *UC Berkeley*.
- UFFELMANN, E., Q. Q. HUANG, N. S. MUNUNG, J. DE VRIES, Y. OKADA, A. R. MARTIN, H. C. MARTIN, T. LAPPALAINEN, AND D. POSTHUMA (2021): “Genome-wide association studies,” *Nature Reviews Methods Primers*, 1, 1–21.
- VAUGHAN, L. K. AND V. SRINIVASASAINAGENDRA (2013): “Where in the genome are we? A cautionary tale of database use in genomics research,” *Frontiers in Genetics*, 4, 38.
- VISSCHER, P. M., N. R. WRAY, Q. ZHANG, P. SKLAR, M. I. MCCARTHY, M. A. BROWN, AND J. YANG (2017): “10 years of GWAS discovery: Biology, function, and translation,” *The American Journal of Human Genetics*, 101, 5–22.
- VOLBERDA, H. W., N. J. FOSS, AND M. A. LYLES (2010): “Absorbing the concept of absorptive capacity: How to realize its potential in the organization field,” *Organization Science*, 21, 931–951.
- WU, L., L. HITT, AND B. LOU (2020): “Data analytics, innovation, and firm productivity,” *Management Science*, 66, 2017–2039.
- WUEBKER, R., T. ZENGER, AND T. FELIN (2023): “The theory-based view: Entrepreneurial microfoundations, resources, and choices,” *Strategic Management Journal*, Forthcoming.
- YUE, D., P. HAMILTON, AND I. BOJINOV (2023): “Nailing prediction: Experimental evidence on the value of tools in predictive model development,” *Harvard Business School Working Paper 23-029*.
- ZITTRAIN, J. (2019): “The hidden costs of automated thinking,” *The New Yorker*.
- ZOLLO, M. AND S. G. WINTER (2002): “Deliberate learning and the evolution of dynamic capabilities,” *Organization Science*, 13, 339–351.

8 Figures and Tables

Figure 1: Example and schema of a typical GWAS


(a) Abstract of the GWAS that identified *IL23R* as a therapeutic target for Crohn's disease

A Genome-Wide Association Study Identifies *IL23R* as an Inflammatory Bowel Disease Gene


Richard H. Duerr,^{1,2} Kent D. Taylor,^{3,4} Steven R. Brant,^{5,6} John D. Rioux,^{7,8} Mark S. Silverberg,⁹ Mark J. Daly,^{3,10} A. Hillary Steinhart,⁹ Clara Abraham,¹¹ Miguel Regueiro,³ Anne Griffiths,¹² Themistocles Dassopoulos,⁵ Alain Bitton,¹³ Huiying Yang,^{3,4} Stephan Targan,^{4,14} Lisa Wu Datta,⁵ Emily O. Kistner,¹⁵ L. Philip Schumm,¹⁵ Annette T. Lee,¹⁶ Peter K. Gregersen,¹⁶ M. Michael Bamada,² Jerome I. Rotter,^{3,4} Dan L. Nicolae,^{13,17} Judy H. Cho^{18*}

The inflammatory bowel diseases Crohn's disease and ulcerative colitis are common, chronic disorders that cause abdominal pain, diarrhea, and gastrointestinal bleeding. To identify genetic factors that might contribute to these disorders, we performed a genome-wide association study. We found a highly significant association between Crohn's disease and the *IL23R* gene on chromosome 1p31, which encodes a subunit of the receptor for the proinflammatory cytokine interleukin-23. An uncommon coding variant (rs11209026, c.1142G>A, p.Arg381Gln) confers strong protection against Crohn's disease, and additional noncoding *IL23R* variants are independently associated. Replication studies confirmed *IL23R* associations in independent cohorts of patients with Crohn's disease or ulcerative colitis. These results and previous studies on the proinflammatory role of IL-23 prioritize this signaling pathway as a therapeutic target in inflammatory bowel disease.

(b) Schema of a typical GWAS

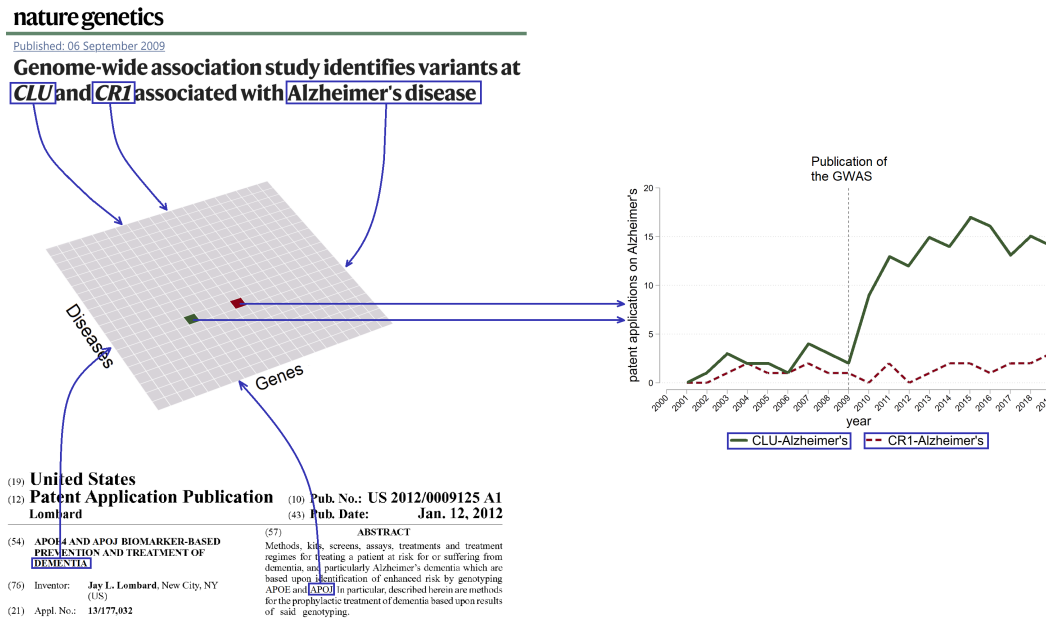


(c) Schema of Duerr et al. (2006)



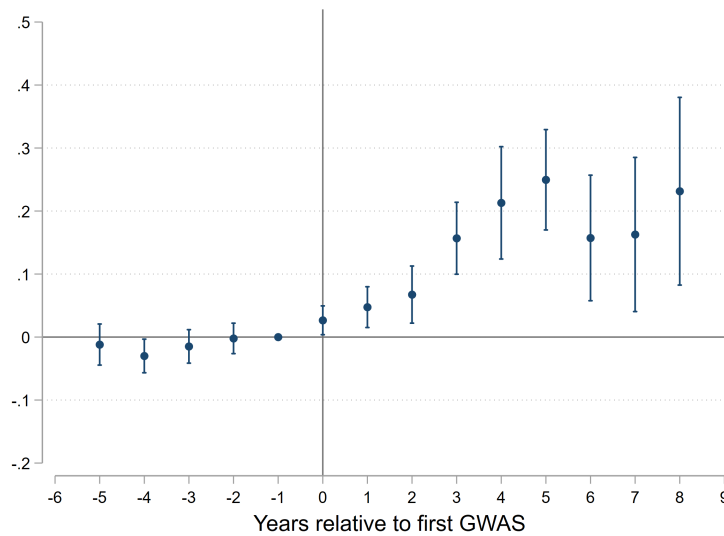
Note: Panel (a) shows the abstract of the Duerr et al. (2006) GWAS published in *Science*. The abstract highlights the potential implications of using the *IL23R* gene as a target to cure inflammatory bowel diseases. Panel (b) shows a schema of how a GWAS unfolds. First, the researchers select the disease of interest and assemble a group of cases (subjects showing the condition) and one of controls (healthy subjects). Then, the genome of people with and without the condition is genotyped in search of differences. Finally, statistical methods are used to test the association between any genetic mutation and the disease of interest. The panel at the bottom is the characteristic “Manhattan plot,” which indicates the location of the statistically significant genetic variants in the chromosome. On the Y axis, there is the strength of the finding expressed as $-\log_{10}(p\text{-value})$, hence higher values correspond to stronger associations. Panel (c) shows the same schema for the GWAS of Duerr et al. (2006). See Appendix B.2 for a detailed case study.

Figure 2: Pharmaceutical firms search for drugs in an empirical landscape of gene-disease pairs



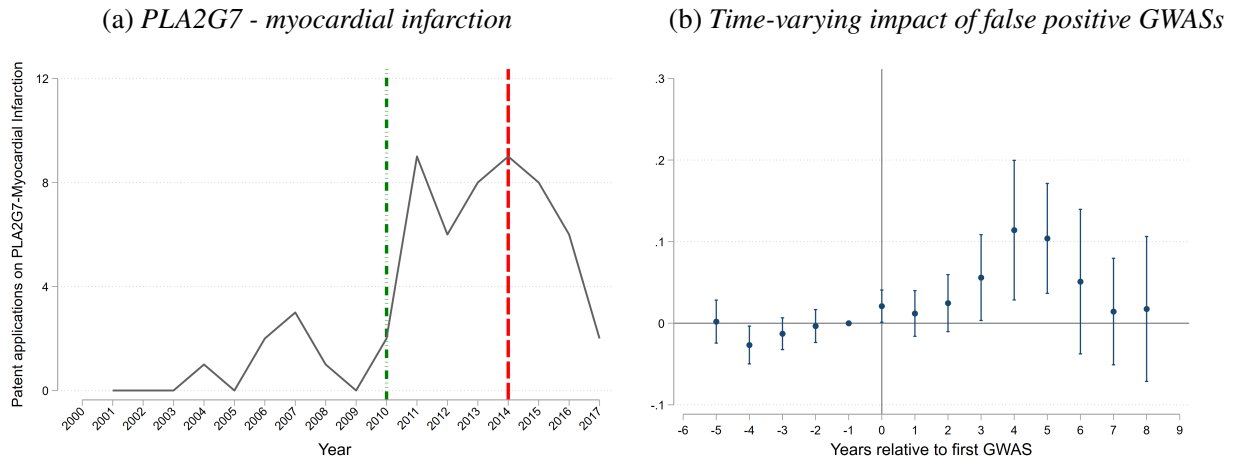
Note: The figure exemplifies the novel measurement approach adopted in this paper. Machine learning tools are used to extract genes and diseases from patent and paper texts to map them onto an empirical landscape of gene-disease combinations. The heterogeneous impact of GWASs on firm investments can be measured by tracking patenting at the gene-disease combination level, avoiding the use of patent citations. APOJ is an alternative name of the CLU gene. See the conceptual details in Appendix C.

Figure 3: Time-varying estimates of GWASs' impact on firm innovation investments



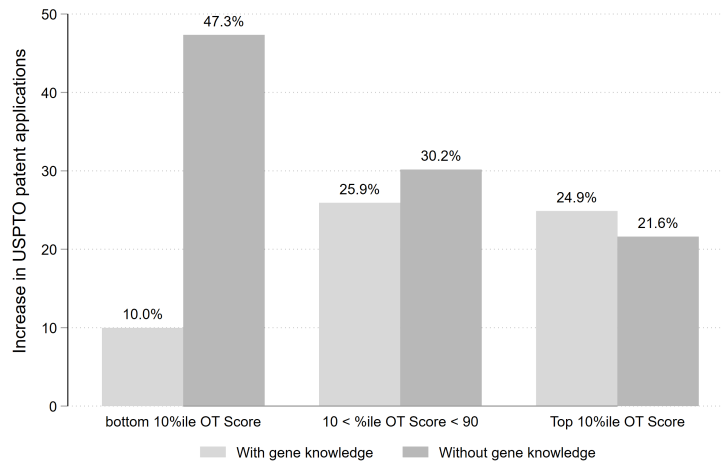
Note: The figure shows the event study coefficients estimated from the following specification: $Patent\ Applications_{i,j,t} = \alpha + \sum_z \beta_z GWAS_{i,j} \times 1(z) + \gamma GD_{i,j} + \delta_t \times Gene_i + \omega_t \times Disease_j + \epsilon_{i,j,t}$. The dependent variable is the number of USPTO patent applications targeting a specific gene-disease combination. The chart plots values of β_z for different lags z before and after the publication of the first GWAS finding about the gene-disease pair. Regressions include gene \times year and disease \times year fixed effects, as well as gene-disease combination fixed effects. Standard errors are clustered two-way at the gene and disease level.

Figure 4: Firms investing based on false positive GWAS findings pivot away over time



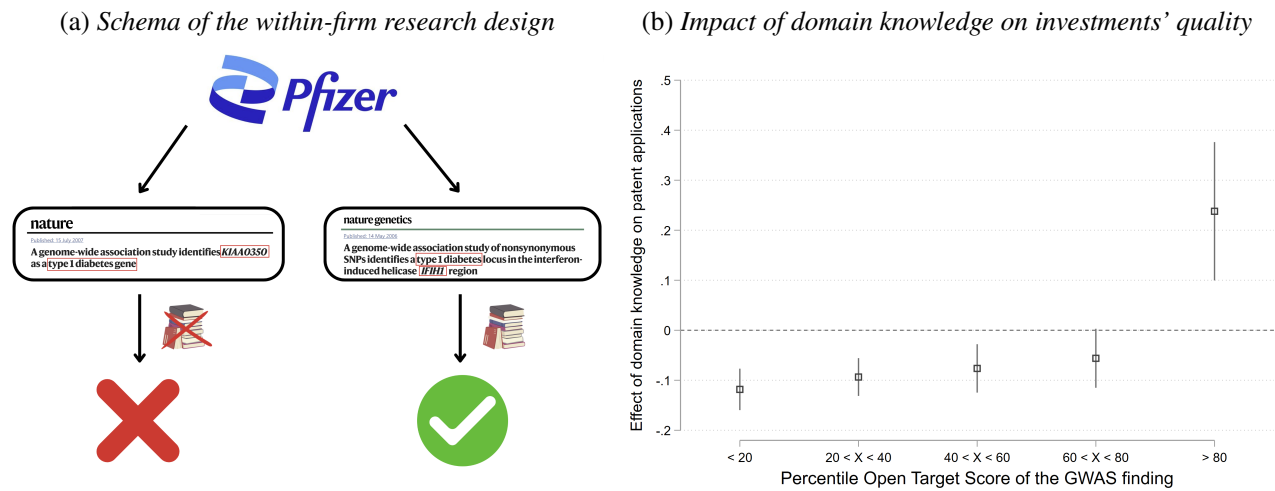
Note: Panel (a) shows the count of yearly USPTO patent applications targeting the PLA2G7 - myocardial infarction gene-disease pair. The green line in 2010 marks the publication of Suchindran et al. (2010), while the red line in 2014 marks the failure of clinical trials for Darapladib, a GlaxoSmithKline drug targeting this combination. Panel (b) shows the event study coefficients from the same specification of Figure 3. The dependent variable is the number of USPTO patent applications targeting a specific gene-disease combination. The chart plots values of β_z for different lags z before and after the publication of a GWAS finding about the gene-disease pair that is not subsequently replicated. Regressions include gene \times year and disease \times year fixed effects, as well as gene-disease combination fixed effects. Standard errors are clustered two-way at the gene and disease level.

Figure 5: Firms with domain knowledge invest proportionally more on the top tail of data-driven opportunities, while those lacking it invest mostly on gene-disease associations of low-value



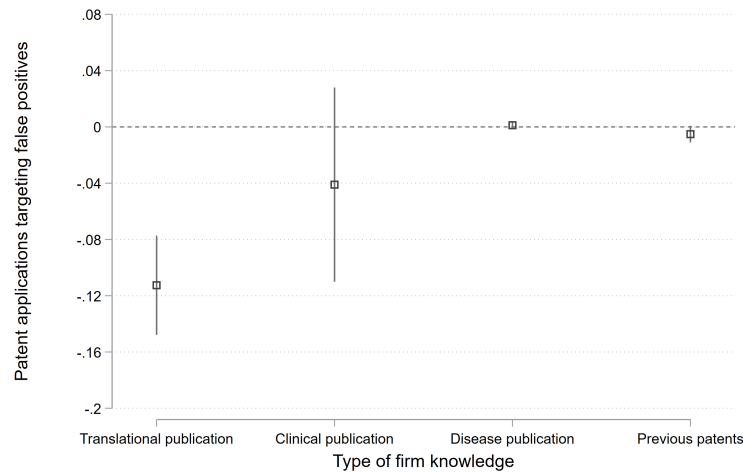
Note: This figure shows the increase in USPTO patent applications on a gene-disease pair after the publication of a GWAS, as a share of the sample average. The increase is estimated with split-sample regressions that consider gene-disease pairs of different therapeutic potential, as proxied by the Open Targets score. Details on the regressions are reported in Appendix Table F.12.

Figure 6: A within-firm research design confirms that domain knowledge improves the assessment of GWAS predictions and helps to select only the opportunities with higher Open Target scores



Note: Panel (a) exemplifies the within-firms design to isolate the role of domain knowledge in interpreting GWASs and determining performance. The research design exploits the fact that a firm might be evaluating multiple findings for the same disease, but has done genetic research only on a subset of the genes involved. Panel (b) shows coefficients from split-sample regressions evaluating the within-firm impact of genetic knowledge on firms' ability to recognize the best GWAS findings. The ground truth quality of the gene-disease association uncovered by GWAS is proxied by the percentile of the Open Targets score. The regressions include firm \times disease and year of discovery dummies.

Figure 7: The ability to avoid false positives is driven by translational knowledge that permits a scientific evaluation of GWAS findings, but not by clinical knowledge or generic disease experience.



Note: This figure shows results separately for firms with different types of domain knowledge. Each coefficient is estimated from a separate regression with the same dependent variable. *Patent applications in false positives*: count of USPTO patent applications targeting a specific gene-disease combination that is not replicated by subsequent GWASs. *Translational publication*: 0/1 = 1 if the firm has published at least one previous translational paper involving the gene, according to the definition by Azoulay et al. (2021). *Clinical publication*: 0/1 = 1 if the firm has published at least one previous clinical paper involving the gene, using the classification of NIH iCite. *Disease publication*: 0/1 = 1 if the firm has published at least one previous paper involving the disease but not the specific gene. *Previous patents*: 0/1 = 1 if the firm has applied for at least one previous patent targeting the gene involved.

Table 1: Summary statistics

	Panel A: cross-sectional descriptives					
	mean	median	st d	min	max	N
Ever a patent (0/1)	0.2182	0	0.4131	0	1	7,223,924
...by firms with gene knowledge (0/1)	0.0478	0	0.2134	0	1	7,223,924
...by firms without gene knowledge (0/1)	0.2102	0	0.4074	0	1	7,223,924
Ever a drug (0/1)	0.0019	0	0.0433	0	1	7,223,924
Ever treated by GWAS (0/1)	0.0025	0	0.0498	0	1	7,223,924
Has Open Targets Score (0/1)	0.0823	0	0.2748	0	1	7,223,924
Open Targets Score	0.0401	0.0074	0.0872	0.00004	0.8975	594,353
	Panel B: panel-level descriptives					
	mean	median	st d	min	max	N
Patent applications	0.1314	0	1.6831	0	1346	137,254,556
...by firms with gene knowledge	0.0376	0	1.0875	0	1346	137,254,556
...by firms without gene knowledge	0.0938	0	0.9243	0	1257	137,254,556
Cit-weighted patents	7.1358	0	85.5509	0	39,899	137,254,556
New drugs (total)	0.0002	0	0.0203	0	22	137,254,556
New drugs (weighted)	0.0001	0	0.0064	0	3.25	137,254,556
Treated by GWAS (0/1)	0.0005	0	0.0234	0	1	137,254,556
True positives (0/1)	0.0001	0	0.0013	0	1	137,254,556
Year	2010	2010	5.4772	2001	2019	137,254,556

Note: This table lists summary statistics at the gene-disease combination level for 7,223,924 combinations (Panel A) and at the gene-disease-year level for a balanced panel of 137,254,556 observations (Panel B). *Ever a patent*: 0/1 = 1 if the gene-disease combination appeared in at least one patent application. *Ever a patent by firms with gene knowledge*: 0/1 = 1 if the gene-disease combination appeared in at least one patent application by firms with previous publications on the gene involved. *Ever a patent by firms without gene knowledge*: 0/1 = 1 if the gene-disease combination appeared in at least one patent application by firms without previous publications on the gene involved. *Ever a drug*: 0/1 = 1 if the gene-disease combination is targeted by at least one drug in the discovery stage. *Ever treated by GWAS*: 0/1 = 1 if the gene-disease combination is ever treated by a GWAS. *Has Open Targets Score*: 0/1 = 1 if the gene-disease combination has an Open Targets Score. *Open Targets Score*: average value of the Open Target Score (for gene-disease pairs that have it). *Patent applications*: count of USPTO patent applications that target a specific gene-disease combination. *Patent applications by firms with gene knowledge*: count of USPTO patent applications that target a specific gene-disease combination filed by firms with previous publications on the gene involved. *Patent applications by firms without gene knowledge*: count of USPTO patent applications that target a specific gene-disease combination filed by firms without previous publications on the gene involved. *Cit-weighted patents*: count of USPTO patent applications that target a specific gene-disease combination weighted by citations received up to 7 years after their publication. *New drugs (total)*: number of molecules on a gene-disease combination entering the discovery stage. *New drugs (weighted)*: number of molecules on a gene-disease combination entering the discovery stage weighted by their scientific novelty (i.e., by the number of times that the same mechanism of action has been used before, following Dranove et al. 2022). *Treated by GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. *True positives*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease. *Year*: average year of observations in the panel.

Table 2: Pharmaceutical firms heavily increase their investments in gene-disease combinations that receive a GWAS, even if the association is a false positive finding

Dependent Variable:	USPTO patent applications			
Post × GWAS	0.2681*** (0.02206)	0.1637*** (0.01985)	0.1532*** (0.01779)	0.0676*** (0.01605)
...× True Positive			0.5739*** (0.08545)	0.4802*** (0.07693)
Gene-Disease FE	YES	YES	YES	YES
Disease FE	YES	NO	YES	NO
Gene FE	YES	NO	YES	NO
Year FE	YES	NO	YES	NO
Disease-Year FE	NO	YES	NO	YES
Gene-Year FE	NO	YES	NO	YES
N of Observations	137,254,556	137,254,556	137,254,556	137,254,556
N of Gene-Diseases	7,223,924	7,223,924	7,223,924	7,223,924
Mean of Dep Var:	0.1314	0.1314	0.1314	0.1314

Note: *, ***, **** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications (by all firms) filed in a given year that target a specific gene-disease combination. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. *True Positive*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease.

Table 3: Only investments based on true positive GWAS associations yield valuable patents or new drugs entering the discovery stage

Dependent Variable:	Cit-weighted patents	Patent market value	Drugs (total)	Drugs (weighted)
Post × GWAS	-0.2738 (0.51521)	0.2338 (0.14649)	-0.0004 (0.00026)	-0.0000 (0.00008)
...× True Positive	5.2464** (1.83822)	2.6963*** (0.74601)	0.0011 (0.00393)	0.0009* (0.00043)
Gene-Disease FE	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES
N of Observations	137,254,556	137,254,556	137,254,556	137,254,556 7
N of Gene-Diseases	7,223,924	7,223,924	7,223,924	7,223,924
Mean of Dep Var:	7.1358	0.6317	0.0002	0.0001

Note: *, ***, **** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. *True Positive*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease. *Cit-weighted patents*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination, weighted by the number of patent citations received up to seven years after patent publication. *Patent value*: estimated stock market value (in constant USD) of patents granted to public firms using data from Kogan et al. (2017). *Drugs (total)*: number of molecules on a gene-disease combination entering the discovery stage. *Drugs (weighted)*: number of molecules on a gene-disease combination entering the discovery stage weighted by their scientific novelty (i.e., by the number of times that the same mechanism of action has been used before, following Dranove et al. 2022).

Table 4: Firms with genetic domain knowledge are able to avoid false positive associations but invest proportionally less in gene-disease combinations that receive a GWAS

Dependent Variable:	USPTO patent applications by...			
	...firms with gene knowledge		...firms w/out gene knowledge	
Post × GWAS	0.0381*** (0.01009)	0.0138 (0.00821)	0.1256*** (0.01429)	0.0539*** (0.01052)
... × True Positive		0.1216** (0.03917)		0.3586*** (0.05837)
Gene-Disease FE	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES
N	137,254,556	137,254,556	137,254,556	137,254,556
N of Gene-Diseases	7,223,924	7,223,924	7,223,924	7,223,924
Mean of Dep Var:	0.0376	0.0376	0.0938	0.0938

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination; the count is then divided between firms with and without previous publications on the gene. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. *True Positive*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease.

Table 5: Firms with genetic domain knowledge avoid investing in GWAS associations that mistakenly report the wrong gene

Dependent Variable:	USPTO patent applications by...		
	...all firms	...firms with gene knowledge	...firms w/out gene knowledge
Post × Wrong Gene	0.0986* (0.04945)	-0.0467 (0.03994)	0.1453*** (0.02609)
Gene-Disease FE	YES	YES	YES
Disease-Year FE	YES	YES	YES
Gene-Year FE	YES	YES	YES
N	136,935,850	136,935,850	136,935,850
N of Gene-Diseases	7,207,150	7,207,150	7,207,150
Mean of Dep Var:	0.1304	0.0372	0.0931

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. The sample is slightly smaller because it excludes gene-disease pairs that receive a correctly reported GWAS association. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination; the count is then divided between firms with and without previous publications on the gene. *Post × Wrong Gene*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS that includes a *wrongly* reported protein-coding gene (i.e., that is later reclassified by the curators of the GWAS Catalog).

Table 6: Genetic domain knowledge allows recognizing both false and true positive GWAS associations in a within firm-design, ruling out that the effects are due to generic firm-level capabilities

Dependent Variable: Sample:	Sum of USPTO patent applications			
	False Positive Associations		True Positive Associations	
Has Gene Knowledge (0/1)	-0.1297*** (0.01525)	-0.1330*** (0.01441)	0.1944*** (0.05453)	0.1996*** (0.05633)
Firm FE	YES	NO	YES	NO
Disease FE	YES	NO	YES	NO
Firm-Disease FE	NO	YES	NO	YES
N of Firms	3,787	3,768	3,651	3,632
N of Observations	4,493,871	4,483,884	1,432,671	1,432,671
Mean of Dep Var:	0.0057	0.0057	0.0225	0.0225

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the firm-gene-disease level represent all the potential GWAS opportunities considered by a firm. Std. err. clustered two-ways at the firm level. *Sum of USPTO patent applications*: sum of USPTO patent applications filed in a given year that target a specific gene-disease combination. *Has Gene Knowledge (0/1)*: 0/1 = 1 if the firm has at least one publication about the specific gene before the GWAS associated it with a disease. Columns 1 and 2 include only GWAS associations that fail subsequent replications, while columns 3 and 4 include only GWAS associations that are later replicated.

Finding Diamonds in the Rough:

Data-Driven Opportunities and Pharmaceutical Innovation

Appendix

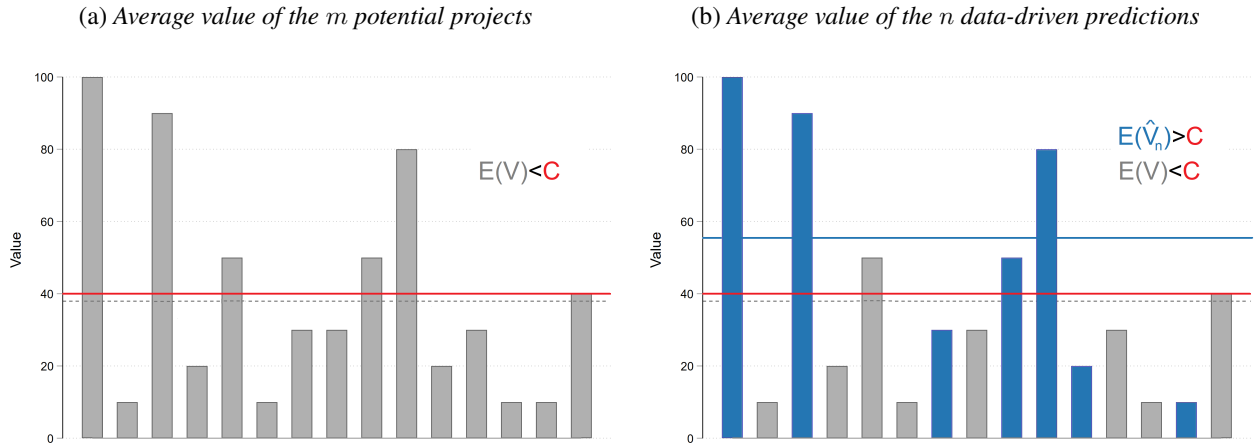
Matteo Tranchero
UC Berkeley-Haas

A	Data-Driven Predictions and Domain Knowledge: A Simple Numerical Example	2
B	Additional Details on GWASs	6
B.1	A Scientific Primer	6
B.2	Case Study: The GWAS of Duerr et al. (2006)	7
C	Measurement Approach	9
C.1	Shortcomings of Patent-to-Paper Citations	9
C.2	Using Entitymetrics to Measure Scientific Impact	9
C.3	An Application of Entitymetrics in Pharmaceutical Innovation	11
D	SciBite/EBI Patent Data	14
D.1	Sample and Validation	14
D.2	Descriptive Statistics	15
D.3	Example: Denali Therapeutics	17
E	Other Data Sources	19
E.1	The NHGRI-EBI GWAS Catalog	19
E.2	Open Targets Score	20
E.3	PubTator Central Publication Data	22
E.4	Cortellis Drug Data	25
F	Additional Figures and Tables	29

A Data-Driven Predictions and Domain Knowledge: A Simple Numerical Example

This appendix provides a numerical example to clarify the conceptual framework outlined in Section 2.3. To fix ideas, imagine that firms are risk-neutral and search for valuable innovations in a landscape of $m = 15$ potential technological combinations, each costing $C = 40$. The distribution of values of these projects is as follows: $\langle 10, 10, 10, 10, 10, 20, 25, 30, 30, 30, 40, 50, 50, 80, 90, 100 \rangle$, such that $E(V) = 39 < C = 40$. This example captures the skewness of innovation returns, with a few projects guaranteeing a high return ($V > C$) and the majority of them having little value ($V < C$). One can assume that the firms can only know the expected value of the investments $E(V) = 39$ but do not know the specific value of any given project. Therefore, under these conditions, risk-neutral firms will not invest since this will result in making an expected loss. Panel (a) of Figure A.1 presents the example graphically: the 15 potential projects are represented as grey bars, each as high as its value V . The red line represents the cost C , which lies above the average value of the opportunities $E(V)$. The opportunities are not sorted by value since firms do not know the precise V_i of any opportunity i .

Figure A.1: Data-driven predictions provide a shortlist of potential opportunities of higher expected value, thus encouraging firms to invest

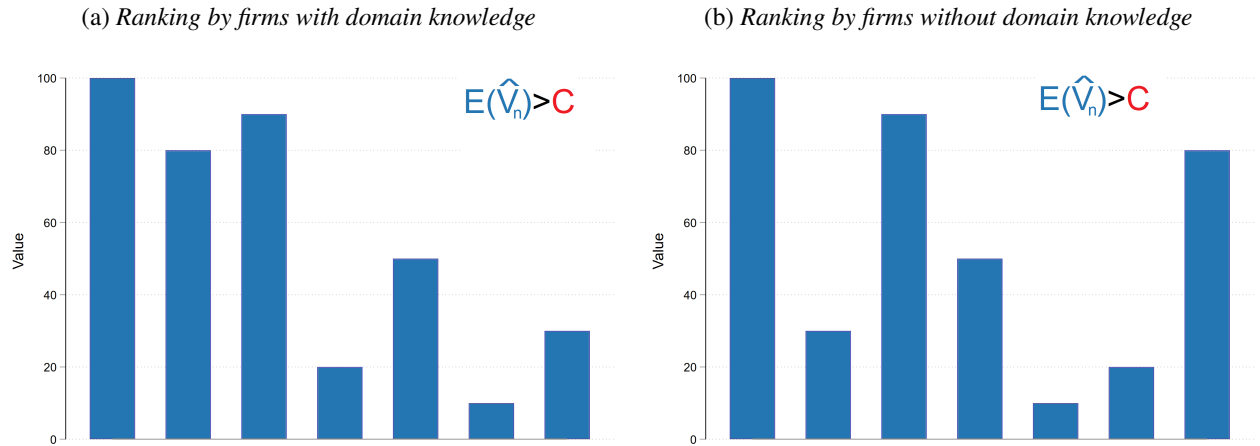


Against this backdrop, one can envision data-driven predictions \hat{V} as providing a shortlist of $n = 7 < m = 15$ combinations. Since the predictions are informative on average, firms know that the predicted value of these shortlisted combinations is higher than the cost; in this example, $E(\hat{V}_n) = 55 > C = 40$. The firms will have the incentive to start investing once they focus their attention on the subset of data-driven predictions (Panel (b) of Figure A.1). However, firms know that the predictions \hat{V}_n are noisy, i.e., imperfectly correlated with the true values V_n . The firm can learn the actual value V_i of a generic technological combination i only after investing, which captures firms' "online learning" (Gavetti and Levinthal, 2000). Since the firms try to maximize their returns, it is rational for them to *rank* data-driven predictions and start investing sequentially from those believed to be of higher value.

The key dimension of firm heterogeneity is their ability to assess the true value of data-driven

opportunities. More precisely, each firm k observes a noisy prediction \hat{V}_i and comes up with its firm-specific assessment $V_{i,k} = V_i + \gamma_i + \epsilon_k$, where $\gamma_i \sim N(0, \eta_i^2)$ is the noise in the prediction itself and $\epsilon_k \sim N(0, \sigma_k^2)$ is the noise in the firms' own assessment. Borrowing the idea of Aguiar and Waldfogel (2018), one can formalize domain knowledge as increasing the predictability of the returns from observing a given data-driven prediction, i.e., shrinking the firm-specific σ_k^2 . Therefore, domain knowledge allows the firm to come up with a ranking closer to the correct ranking based on the true underlying value of the combinations. Figure A.2 exemplifies this idea using the parameters of the example above. The figures make clear that firms with domain knowledge (panel (a)) will be able to rank the data-driven predictions better, while the other firms will order them in a way that is less reflective of the true underlying value of the combinations (panel (b)).

Figure A.2: Firms with domain knowledge are able to recognize the most promising opportunities among data-driven predictions, hence ranking them higher in their order of investments

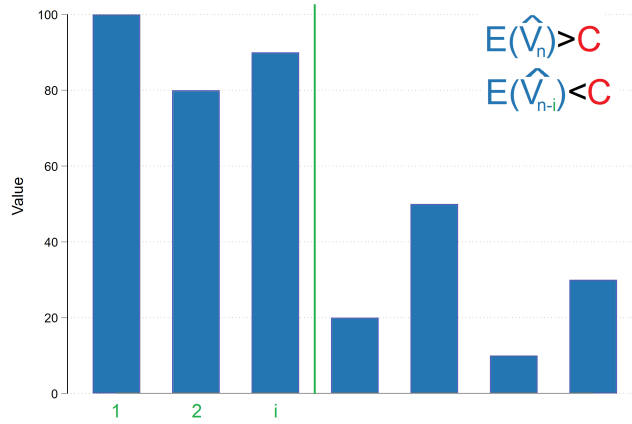


Note: Both panels show the ranking (from left to right) of the same data-driven predictions made by firms that have heterogeneous domain knowledge. Firms with domain knowledge have a better understanding of the true value V_i underlying the data-driven prediction \hat{V}_i , hence their ranking of data-driven opportunities is closer to the ground truth $V_1 > V_2 > \dots > V_i > \dots > V_n$.

Firms will then sequentially invest in each combination from those ranked highest until the marginal combination i such that $E(\hat{V}_{n-i}) < C$. Consider first the firms with domain knowledge ($k = 1$), whose ranking of data-driven predictions is depicted in Figure A.3. The firms will start investing in the first project on the left, of value $V_1 = 100$, thus making a profit $\pi = V_1 - C = 60$. At that point, the firms will know that the remaining combinations have $E(\hat{V}_{n-1}) = 47.5 > C = 40$,³² so they will continue investing. This sequential process will continue until the opportunity $i = 3$, because, at that point, the expected value of continuing to invest is below the cost ($E(\hat{V}_{n-3}) = 28.8 < C = 40$). In this example, companies with domain knowledge will make three investments of value $< 100, 80, 90 >$, hence achieving a profit of $\pi_{k=1} = 150$.

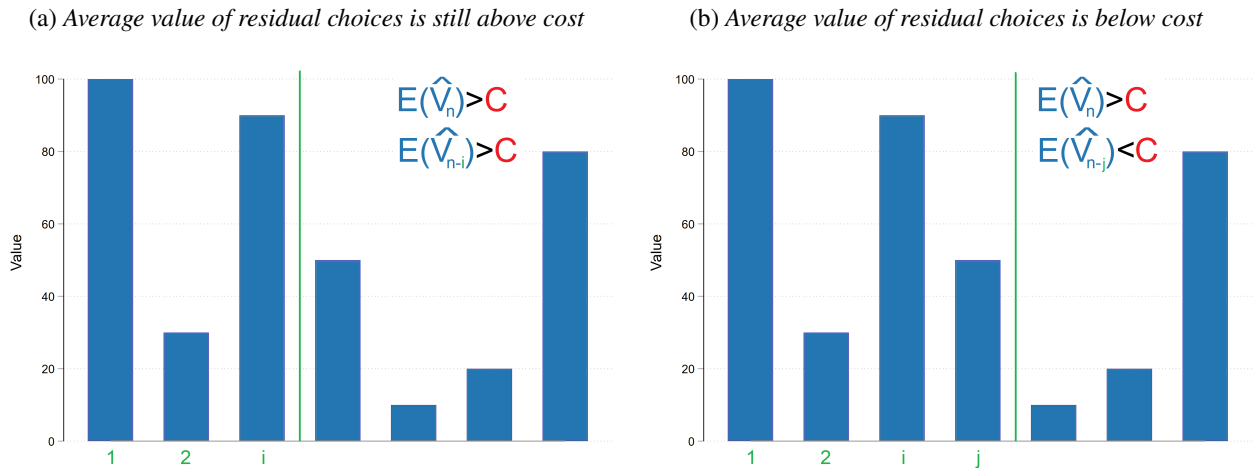
³²Since I assumed that the firms know $E(\hat{V}_n) = 55$, they will be able to compute the residual $E(\hat{V}_{n-1})$ after experimentally learning the value of their first investment.

Figure A.3: Firms will invest sequentially starting from the data-driven predictions ranked higher and until the expected returns from the remaining ones are below the cost.



Now consider the firms without domain knowledge ($k = 0$), whose ranking is exemplified in Figure A.4. Those firms will follow the same process. However, in this case, the ranking is noisier, which means that the first three investments will be, on average, of lower value (in this example, they will be $\langle 100, 30, 90 \rangle$). At that point, the expected value from continuing to invest will be $E(\hat{V}_{n-3}) = 41.3 > C = 40$, thus leading to further investments. The firms will stop at opportunity $j = 4$ because $E(\hat{V}_{n-4}) = 38.3 < C = 40$. The final result will be four investments of value $\langle 100, 30, 90, 50 \rangle$, hence achieving a total profit of $\pi_{k=0} = 110$.

Figure A.4: Firms without domain knowledge will invest in more predictions because of their lower ability to recognize opportunities

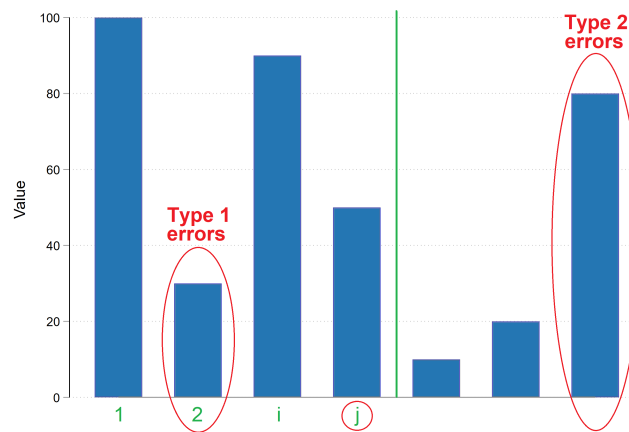


Note: Both panels show the ranking (from left to right) of the same data-driven predictions made by firms without domain knowledge. Panel (a) shows that after investing in the opportunity i , the expected value of the residual opportunities is still above cost. Panel (b) shows that the firm will invest until prediction j such that the expected value of the residual opportunities is below cost.

This simple framework offers three implications amenable to empirical testing. First, the data-driven

predictions point, on average, to valuable opportunities ($E(\hat{V}_n) = 55 > C = 40 > E(V_m) = 39$). This will lead all firms to increase their investments and, on average, introduce more innovations. Second, firms without domain knowledge will pursue more data-driven predictions because they are less capable of identifying the most promising ones. In practice, they will make more mistakes in their early investments, hence having an incentive to continue searching. Third, those firms will make both more commission (or type I) and more omission (or type II) errors, being more likely to select low-value combinations that they mistakenly ranked high at the expense of actual opportunities. Figure A.5 visually illustrates these implications. Compared to firms with domain knowledge, those without domain knowledge should be *less* likely to find the best opportunities among data-driven predictions despite investing quantitatively *more*.

Figure A.5: Firms without domain knowledge will invest in more data-driven predictions, while also making more commission and omission errors.



B Additional Details on GWASs

Genomics is the branch of biological sciences concerned with the study of genomes, i.e., the entire collection of an organism's genes. In turn, genes are sequences of DNA bases that encode the "instructions" to synthesize gene products, most notably proteins. The normal functioning of a gene can be altered by a mutation, potentially giving rise to severe health conditions. Knowing the roots of genetic diseases has practical consequences in the design of pharmaceutical drugs (Nelson et al., 2015): scientists can design drugs that bind to the malfunctioning gene (or its products) to produce a desired therapeutic effect by leveraging their understanding of gene-disease relationships.

B.1 A Scientific Primer

Diseases caused by individual gene mutations are called *Mendelian disorders*. However, Mendelian diseases are typically rare because they tend to be eliminated by evolutionary pressures due to their severity. More common are *complex* or *polygenic diseases* that are not due to a single genetic factor but rather by many genes. For complex diseases, any genetic mutation can increase the risk of presenting the condition even without being neither necessary nor sufficient, which means that it is responsible for only a tiny proportion of the heritability of the disease. Although complex disorders often cluster in families, they do not have a predictable inheritance pattern because they are determined by convoluted interactions between genes and environmental factors. This means that scientists need to search through the $\sim 19,000$ protein-coding genes to find the mutations involved in each of the thousands of polygenic diseases (Tranchoero, 2023).

Over the years, researchers have come to the conclusion that common disorders are likely influenced by genetic mutations also common in the population (Reich and Lander, 2001). Instead of looking for individual genes with strong effects on phenotypes, the field has moved toward the study of common, generic *variants* that individually have a small effect on the likelihood of having a disease (Bush and Moore, 2012). But what precisely is a variant? At the most fundamental level, two genomes differ in a specific genetic locus if they present an alternative single nucleotide (adenine, thymine, cytosine, or guanine) in that location. Such mutation in one DNA basis is called single-nucleotide polymorphism (SNP) when it appears in at least 1% of the population. One approach for associating SNPs with a disease relies on the fact that a causative variant should be found more frequently in cases than in control subjects. In practice, researchers look for statistical correlations between specific genetic variants and diseases in large samples of unrelated people.

Building over this logic, genome-wide association studies (GWASs) are hypothesis-free methods for identifying associations between genetic regions and traits (Visscher et al., 2017). Using genotyping technologies on large samples, GWASs compare genetic differences between affected and unaffected individuals. In a typical GWAS project, researchers obtain DNA from two groups of participants, patients with the disease studied and healthy individuals with comparable demographics. Then, selected SNPs on the chromosome are scanned using arrays that can genotype up to millions of SNPs for each individual. Variants significantly more likely to appear in the affected patients could be

biologically relevant for the disease and are thus involved with its etiology. Such SNPs might affect gene expression and function, especially when located within a protein-coding gene. It is essential to underscore that array-based genome-wide studies do not sequence the DNA base by base since they only determine the presence or absence of a relatively small number of SNPs (usually $< 0.1\%$ of the genome). Nevertheless, exploiting the fact that the co-occurrence of variants in proximal genetic loci is not random (a phenomenon called linkage), researchers can use reference genomes (such as the HapMap) to parsimoniously infer the characteristics of the whole genome from the much smaller number of SNPs genotyped (Bush and Moore, 2012).

GWAS findings entail a substantial risk of false positives (Boyle et al., 2017; MacArthur, 2012). Even if the association underlies a causal relationship, understanding the biological mechanisms through which it affects human health requires additional study. Moreover, most associations explain a small fraction of the variation in disease susceptibility, which means that the therapeutic benefit from intervening in them could be quite limited (Goldstein, 2009). These limitations notwithstanding, GWASs have proven extremely useful in uncovering drug targets that can assist in identifying compounds suitable for drug repurposing (Reay and Cairns, 2021; Visscher et al., 2017). By pointing out new conditions that might be addressed acting on a given target, GWASs permit to identify new uses for existing drugs that are outside the scope of the original medical indication (Andriani and Cattani, 2022; Pushpakom et al., 2019).

B.2 Case Study: The GWAS of Duerr et al. (2006)

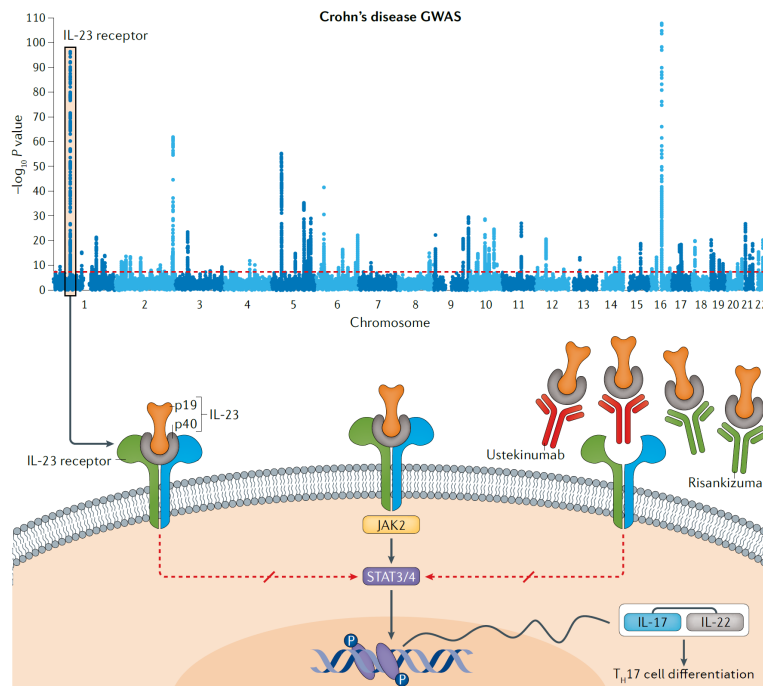
In 2015, an estimated 3 million U.S. adults (around 1.3% of the population) reported being diagnosed with chronic inflammation of the gastrointestinal tract, known as inflammatory bowel disease (IBD). The two most frequent IBD conditions are Crohn's disease and ulcerative colitis (Dahlhamer et al., 2016). Given the prevalence and severity of these diseases, researchers have been studying Crohn's and related diseases intensively. According to the DisGeNET data, IBD was in the top 1% of the diseases with the highest genetic research intensity in the pre-GWAS era (Tranchoero, 2023). A few genes, such as NOD2, had been identified by 2005, but without fully explaining the genetic risk for the disease.

In December 2006, Duerr et al. (2006) published one of the very first GWASs. The study involved 567 European ancestry patients affected by IBD and 571 healthy controls. The GWAS identified genetic mutations in the interleukin 23 receptor gene (IL23R) as significantly associated with Crohn's disease and ulcerative colitis. Before this finding, IL23R was among the least studied human genes. However, many scientists followed up on the promising lead uncovered by the Duerr et al. paper, often with the objective of elucidating the causal mechanisms (called "pathways") through which IL23R exerts its effects. It is now understood that the IL23R gene provides instructions for making a protein called the interleukin 23 receptor, an essential constituent of antibodies that identify foreign substances and defend the body against infections by promoting local inflammations. The IL-23 receptor interacts with a protein called IL-23, binding together like a lock and key (Bianchi and

Rogge, 2019). When IL-23 binds to its receptor, it triggers chemical signals to develop and activate Th17 cells, a specific type of lymphocytes that promote inflammation to fight foreign invaders such as viruses. Malfunctioning IL23R genes might misdirect such inflammatory reactions toward human tissues, giving rise to autoinflammatory diseases like Crohn's.

The work of Duerr et al. (2006) uncovered the crucial role that malfunctioning of IL23R has in inducing inflammations in the intestines. Furthermore, the discovery of IL23R's role in IBD suggested that drugs could be engineered to interfere with malfunctioning IL23R genes. This is indeed what ustekinumab, a fully human monoclonal antibody, does by blocking the p40 subunit of IL23 and preventing its binding with the IL23 receptor (Figure B.1). As of 2023, ustekinumab (Stelara, by Janssen Pharmaceuticals) is a drug approved for treating chronic inflammatory diseases in several jurisdictions, including the United States, Europe, and Australia. Several other drugs are being designed to target the IL23 receptor complex, including risankizumab (Skyrizi, by AbbVie), tildrakizumab, and guselkumab. Interestingly, all these molecules have been repositioned as Crohn's disease intervention from their initial indication for psoriasis (Reay and Cairns, 2021). This highlights how uncovering new drug targets through GWAS can enable the rapid repurposing of gene-specific molecules developed for other diseases (Andriani and Cattani, 2022).

Figure B.1: GWAS findings on Crohn's disease provided new opportunities for drug discovery



Note: The figure exemplifies how GWASs can guide drug development by finding new diseases that can be addressed with existing molecules. Once Duerr et al. (2006) located a new variant in IL23R that disrupts its function, monoclonal antibodies that target the interleukin-23 receptor pathway were repositioned as Crohn's disease interventions from their original indication for psoriasis. The image is taken from Reay and Cairns (2021).

C Measurement Approach

C.1 Shortcomings of Patent-to-Paper Citations

Measuring the technological impact of a scientific project usually involves counting the patent citations received by the publications where its output is codified (Arora et al., 2021; Azoulay et al., 2019; Fleming et al., 2019). Methodological studies showed that patent-to-paper citations are a good way to proxy knowledge flows from science to technology (Narin et al., 1997; Roach and Cohen, 2013). It is not an exaggeration to say that almost all existing empirical studies use non-patent literature citation counts to measure the innovation impact of science.³³ The recent diffusion of open access databases of patent-to-paper citations have further increased the appeal of this approach (Bryan et al., 2020; Marx and Fuegi, 2020). Nonetheless, this measurement approach suffers from two shortcomings.

First, relying on explicit references to science usually provides a downward biased measure of impact (Myers and Lanahan, 2022). Direct citations, for instance, would not capture foundational intellectual influences that become common knowledge in a field,³⁴ nor knowledge flowing through more complicated citation patterns (e.g., a patent citing a publication that cites the focal paper of interest). These examples fall under what Roach and Cohen (2013) define as “errors of omission.” Such underestimate is possibly magnified for basic research with spillovers in fields very different from the one where the idea originated (Azoulay et al., 2019) and through unpredictable channels (Cohen et al., 2002).

Second, and similarly to what happens with academic papers, non-patent literature citations can be made for various reasons (Teplitzkiy et al., 2022). Up to half of the citations are possibly devoted to irrelevant prior art (Jaffe et al., 2000), likely for strategic reasons (Kuhn et al., 2020; Lampe, 2012). This is what Roach and Cohen (2013) define as “errors of commission,” namely citations not corresponding to knowledge relevant to the invention. But even if the reference captures an actual knowledge flow, it is hard to know what part of the study the patent builds on. Academic papers often make multiple contributions; the citation might refer to any of those. Assessing the value of individual contributions made in a paper is impossible from sheer citation counts.

C.2 Using Entitymetrics to Measure Scientific Impact

In this paper, I develop a new measurement approach based on the notion of *knowledge entity* (Ding et al., 2013). I consider any individual carrier of knowledge as an entity, be it embodied into an artifact (e.g., a piece of computer hardware) or a more abstract unit of domain knowledge (e.g., biological entities, such as genes and diseases). Each publication or patent can be characterized by the knowledge entities it studies and recombines (Fleming, 2001). By taking this perspective, one can summarize written documents by compiling a list of their knowledge entities. I propose to assess the technological

³³Advances in natural language processing are enabling to sidestep citation trails by capturing text similarity between documents (Arts et al., 2018; Kaplan and Vakili, 2015; Younge and Kuhn, 2019). However, with very few exceptions (Myers and Lanahan, 2022), these methods have only been used to measure the similarity between patent specifications.

³⁴One can hardly find citations to Keynes’ *General Theory* even in the papers of so-called New-Keynesian economists.

impact of a paper by quantifying the change in activity experienced by the knowledge entities “treated” by it. The basic idea is that impactful projects generate innovation opportunities involving the entities recombined. The increase in patents including such entity combinations relative to similar control combinations is a measure of impact that has the advantage of also capturing patents that do not directly cite the focal paper (Kumar Panda and Tranchero, 2023).³⁵

More in detail, my approach combines machine learning with causal inference in three steps. First, one needs to extract knowledge entities from the relevant documents – in this case, papers and patents. This can be done using automated machine learning procedures, such as state-of-the-art Bidirectional Encoder Representations from Transformers (or BERT). BERT can be tailored to specific domains, such as the life sciences, to increase accuracy and contextual stability (Xu et al., 2020). After extracting accurate knowledge entities using BERT, one can normalize and assign a unique ID to each. The result is that text documents will be characterized by a vector of their knowledge entities. This approach is more flexible than relying on dictionary-based entity extraction methods.

Second, the union of all knowledge entities extracted from the relevant document corpus can be used to trace an empirical knowledge landscape (Rzhetsky et al., 2015). This constitutes the “ground truth,” that is, the space of entities over which researchers and firms carry out their research activities. In some instances, this landscape is known *ex ante* (e.g., the $\sim 19,000$ genes constituting the human genome), while in other cases, it can be backed out from the entity extraction task (e.g., the landscape of research topics described by Boyack et al. 2020). The advantage of studying research as happening on a landscape is that citations are not needed to track knowledge evolution (Ding et al., 2013). Instead, one can simply measure the change in follow-on work relating to the entities themselves before and after the project of interest is completed.

Finally, one can quantify the impact of a research project using a difference-in-differences framework with staggered adoption at the level of each entity combination. Said otherwise, combination-level regressions allow assessing the attention change the entity combination $\langle i, j \rangle$ received after being “treated” by the article m of interest. The basic specification is at the level of individual combination $\langle i, j \rangle$ and time t and takes the form of the following equation:

$$Y_{i,j,t} = \alpha + \beta Post_t \times Article_m + \gamma_{i,j} + \delta_{i,t} + \omega_{j,t} + \epsilon_{i,j,t} \quad (C.1)$$

Where $Y_{i,j,t}$ is any technological impact metric of interest about the treated entities (e.g., patents, clinical trials, or other follow-on outcomes), $Article_m$ equals one for the combinations treated and zero for the control ones, $Post_t$ is a binary variable that takes value one only after the project of interest is completed, $\gamma_{i,j}$ are combination fixed effects, while $\delta_{i,t}$ and $\omega_{j,t}$ are entity-specific time trends. Combination fixed effects consider the potential of different combinations, while the entity time trends avoid confounding from heterogeneous changes in attention to individual entities.

The coefficient of interest, β , captures the marginal impact of the project completion on the treated

³⁵Similarly to my approach, Iaria et al. (2018) trace the diffusion of frontier knowledge measuring the occurrence of new scientific concepts in patent texts. Suh (2023) extracts the chemical compounds mentioned in the body of patent texts to recognize innovations that rely more on technologies where the Soviet Union had a scientific lead.

entities. Under the assumption of parallel trends in inventive activities between treated and non-treated entities, this procedure provides an estimate of project impact over time that does not depend on patent-to-paper citations. In particular, it solves the measurement issues highlighted in the previous section. First, by capturing the mention of entities directly in the text description of the technological application, this approach is well suited to objectively measure the impact of basic research findings that might not be cited in downstream applications. Second, tracking the reuse of entity combinations permits distinguishing different contributions made in the same article, offering a granularity that citations do not have.

C.3 An Application of Entitymetrics in Pharmaceutical Innovation


The bio-medical sector offers an ideal setting to demonstrate the advantages of entitymetrics (Xu et al., 2020). Pharmaceutical innovation heavily depends on science (Cohen et al., 2002), and knowledge entities are well-defined and have clear meaning (e.g., genes and diseases). I leverage data from European Bioinformatics Institute to measure the gene-disease pairs targeted by each GWAS (2005-2019) and USPTO patent application (2001-2019). The entities from patent texts are extracted using TERMite, the proprietary named entity recognition software developed by SciBite.

Example 1: Undermeasurement of Basic Research Influences: The GWAS of Duerr et al. (2006) was the first to implicate the IL23R gene in the etiology of inflammatory bowel disease (IBD). This finding has proved incredibly impactful both on scientific research and pharmaceutical innovation. Duerr et al. (2006) received over 3,600 citations on Google Scholar (as of June 2022) and led to an improved theoretical understanding of how IL23R is involved in IBD (Bianchi and Rogge, 2019). On the therapeutic front, several drug molecules are now available to treat IBD by modulating the IL-23 signaling pathways, including ustekinumab and risankizumab (Reay and Cairns, 2021). Therefore, this offers a perfect case study to compare the ability of patent citations and entitymetrics to capture the considerable technological impact of this paper.

Figure C.1 shows the time series of USPTO patent applications targeting the IL23R-IBD pathway or citing the GWAS by Duerr et al. (2006), respectively. A few things are interesting to note. First, some firms were already exploring the IL23R-IBD nexus *before* the finding became known in the broader scientific community (Brusoni et al., 2001). Those firms are then the quickest to react, applying for patents that were probably already in the making and thus not citing Duerr et al. (2006), but that found a crucial validation in it (Kao, 2023). Second, it appears that just looking at patent-to-paper citations would lead to a substantial underestimate of impact. Only 46 patent applications cite Duerr et al. (2006), compared with the 88 applications that exploit its finding. Interestingly, the gap between the two measurements reduces to half among patents filed by universities and research institutions, suggesting that most undercount stems from firms' citation practices.

Direct patent-to-paper citations cannot account for citations to downstream science enabled by Duerr et al. (2006). For instance, patents might cite subsequent studies that experimentally validated the IL23R-IBD association (Ahmadpoor and Jones, 2017). Figure C.2 shows a few claims of the US

Figure C.1: Count of USPTO patent applications that directly cite Duerr et al. (2006) or that target the IL23R-IBD combinations, regardless of whether they cite Duerr et al. (2006)



Note: Data on patent citations are from Google Patents. The vertical line marks the publication of Duerr et al. (2006).

2016/0333091 A1 patent application by Boehringer Ingelheim. The patent builds on knowledge of IL23R’s role in IBD without acknowledging the GWAS by Duerr and colleagues, instead citing subsequent papers that explain the mechanisms behind the IL23R-IBD correlation (Beyer et al., 2008). The entity-based approach seems better equipped to capture the impact of basic advances triggering extensive follow-on research.

Figure C.2: Claims 28, 29, and 30 of patent application US 2016/0333091 A1

- 28) A method for treating an inflammatory disease, an autoimmune disease, a respiratory disease, a metabolic disorder or cancer comprising administering to a subject in need thereof an effective amount of an anti-IL-23p19 antibody or antigen-binding fragment or a pharmaceutical composition comprising an anti-IL-23p19 antibody or antigen-binding fragment and a pharmaceutically acceptable carrier, wherein the antibody or antigen-binding fragment thereof comprises:
 - a) a light chain variable region comprising the amino acid sequence of SEQ ID NO:19 (CDR1-L); the amino acid sequence of SEQ ID NO:20 (CDR2-L); and the amino acid sequence of SEQ ID NO:21 (CDR3-L); and
 - b) a heavy chain variable region comprising the amino acid sequence of SEQ ID NO: 63, 66, 67 or 68 (CDR1-H); the amino acid sequence of SEQ ID NO:64 (CDR2-H); and the amino acid sequence of SEQ ID NO:65 (CDR3-H).
- 29) The method according to claim 28, wherein the disease is psoriasis, inflammatory bowel disease, psoriatic arthritis, multiple sclerosis, rheumatoid arthritis, or ankylosing spondylitis.
- 30) A method for inhibiting the binding of IL-23 to the IL-23 receptor on a mammalian cell, comprising administering to the cell an anti-IL-23p19 antibody or antigen-binding fragment, whereby signaling mediated by the IL-23 receptor is inhibited.

Note: This figure shows selected claims of the patent application titled “Anti-IL-23 Antibodies” filed by Boehringer Ingelheim. The patent does not cite the study of Duerr et al. (2006) even if it clearly builds on its finding.

Example 2: Multiple Findings in the Same Paper: Scientific articles often make more than one contribution. For instance, the GWAS of Easton et al. (2007) reported four new genes as correlated with breast cancer (Panel A of Figure C.3). What was the individual impact of these four gene-disease combinations on pharmaceutical innovation? Simply counting how many patents cite this paper would not answer this question. Aveo Pharmaceuticals’ patent US 2011/0305687 A1, titled “Anti-FGFR2 antibodies,” is a good case in point. This patent cites Easton et al. (2007), and from the title alone, it

is clear that it builds on only one of its four findings. Unfortunately, this information is impossible to gather from citation patterns alone.

Figure C.3: Application to measuring the impact of multiple discoveries in the same paper

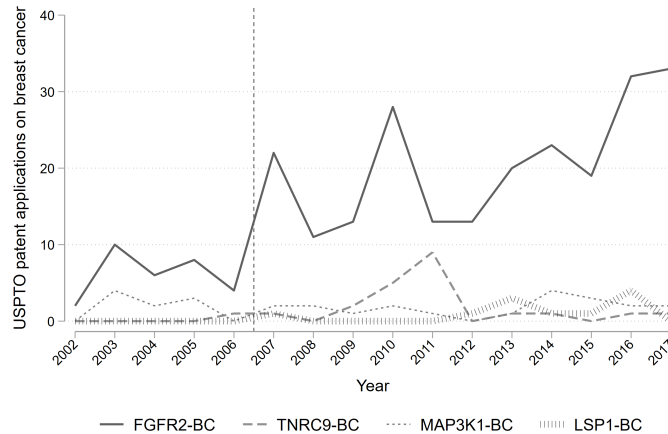
(a) *Abstract of the GWAS by Easton et al. (2007)*

nature
Genome-wide association study identifies novel breast cancer susceptibility loci
 Douglas F. Easton, Karen A. Pooley, Alison M. Dunning, ... Bruce A. J. Ponder
Nature 447, 1087–1093 (2007) | 12k Accesses | 1778 Citations | 48 Altmetric

Abstract

Breast cancer exhibits familial aggregation, consistent with variation in genetic susceptibility to the disease. Known susceptibility genes account for less than 25% of the familial risk of breast cancer, and the residual genetic variance is likely to be due to variants conferring more moderate risks. To identify further susceptibility alleles, we conducted a two-stage genome-wide association study in 4,398 breast cancer cases and 4,316 controls, followed by a third stage in which 30 single nucleotide polymorphisms (SNPs) were tested for confirmation in 21,860 cases and 22,578 controls from 22 studies. We used 227,876 SNPs that were estimated to correlate with 77% of known common SNPs in Europeans at $r^2 > 0.5$. SNPs in five novel independent loci exhibited strong and consistent evidence of association with breast cancer ($P < 10^{-7}$). Four of these contain plausible causative genes (*FGFR2*, *TNRC9*, *MAP3K1* and *LSP1*). At the second stage, 1,792 SNPs were significant at the $P < 0.05$ level compared with an estimated 1,343 that would be expected by chance, indicating that many additional common susceptibility alleles may be identifiable by this approach.

(b) *USPTO patent applications after Easton et al. (2007)*



Note: The top panel shows the abstract of the GWAS by Easton et al. (2007). The bottom panel shows USPTO patent applications that target the four findings of the GWAS by Easton et al. (2007). The vertical line marks the publication of Easton et al. (2007), showing substantial heterogeneity in the impact of individual findings.

My novel measurement strategy can bypass this limitation by extracting each recombination introduced by a given article and then relating it to patenting dynamics on the recombination itself. Figure C.3 shows that the GWAS by Easton et al. (2007) triggered substantial patent applications only on the FGFR2-breast cancer combination. Patenting on this gene-disease combination tripled, with almost no impact on the other combinations. Relying on patent-to-paper citations would completely obscure the heterogeneous impact of the GWAS by Easton et al. (2007). This is all the more important since some of these associations might be false positives. Instead, text-extracted real-world entities are especially suited to measure the impact of basic research with cumulative impact across disparate domains.

D SciBite/EBI Patent Data

This appendix describes in detail the data on the genes and diseases targeted by USPTO patent applications (2001-2019).

D.1 Sample and Validation

The primary source for the data used in the paper is a proprietary database from the European Bioinformatics Institute (EBI). The data have been compiled using TERMite (TERM identification, tagging & extraction), a named entity recognition software developed by the Elsevier-owned startup SciBite. TERMite scans and semantically annotates raw text with entities from over 50 biopharma and biomedical topics. The entities are drawn from VOCabs, a manually curated vocabulary with over 20 million synonyms specifically tuned for named entity recognition text analytics. For instance, this permits recognizing that SEPT1 is the symbol for the SEPTIN1 gene, not a date. Importantly, TERMite has built-in relevance detection, distinguishing between terms that are casual mentions and those that constitute the critical bio-entities of a document.

The data include all the protein-coding genes and diseases extracted from complete patent texts. Genes are matched to their unique NCBI IDs, while diseases are mapped into MeSH Unique IDs. Figure D.1 shows how the TERMite software works. The figure shows USPTO patent application US 2011/0301182 A1, a patent that was granted on September 30, 2014 to Boehringer Ingelheim. This patent is listed in the FDA's *Approved Drug Products with Therapeutic Equivalence Evaluations* (also known as Orange Book) as the intellectual property behind Tradjenta, the brand name for Linagliptin. Tradjenta is a medication used to treat type 2 diabetes by acting as a DPP-4 inhibitor, i.e., by altering the function of the gene DPP4 which plays a significant role in glucose metabolism. Figure D.1 shows how the entity-based approach correctly captures that this patent targets the DPP4-diabetes combination.

Figure D.1: Example of SciBite's entity recognition algorithm used to extract genes and diseases targeted by each USPTO patent application

US-20110301182-A1 / 2011-12-08

Treatment for **diabetes** in patients with inadequate glycemic control despite **metformin** therapy comprising a **DPP-IV** inhibitor

ABSTRACT

The present invention relates to the finding that certain **DPP-4 inhibitors** are particularly suitable for improving glycemic control in type 2 **diabetes** patients with inadequate glycemic control despite metformin therapy.

Note: This figure shows an example of how SciBite's entity recognition algorithm (called TERMite) extracts genes and diseases targeted by the USPTO patent application US 2011/0301182 A1.

I carry out a manual validation to assess the performance of TERMite in terms of recall and precision, following the procedure of Marx and Fuegi (2020). Recall measures how many actual bio-entities of a patent the algorithm found. This is equivalent to 1 minus the percentage of false negatives, i.e.,

entities targeted by a patent that the software failed to find. To assess this metric, I sampled 100 patent applications, obtained their full specification from Google Patents, and randomly extracted one gene and one disease they mentioned. Then I assessed whether the same entities were listed in my data for that specific patent, finding that 91% of genes and 92% of diseases were correctly captured. Symmetrically, precision is given by the share of reported bio-entities that are correct. This metric is computed as 1 minus the percentage of false positives, i.e., entities mistakenly extracted from the patent. I evaluated precision by extracting one gene and one disease for 100 patents in my data and then manually checking whether the entity in question was present in the patent specification. Overall, 95% of genes and 97% of diseases were true positives. Taken together, the F_1 score is equal to 92.96 in the case of genes and 94.43 in the case of diseases, proving the high reliability of my data.

D.2 Descriptive Statistics

My sample includes 148,232 USPTO patent applications published from 2001 to 2019 inclusive. Of these, 73,255 are eventually granted as of summer 2021. All my primary analyses rely on the entire sample of applications, but I also present descriptive statistics for the subset of granted patents for comparison. Table D.1 presents the main descriptive statistics. Each patent application targets, on average, 6.3 genes and 12.3 diseases; this number is only slightly smaller for granted patents.³⁶ However, the sample shows a large variance, with a few patents listing hundreds of genes and diseases as targets. This is probably reflective of strategic disclosure behaviors and offers an exciting avenue for future research on the conditions under which patent text can or cannot be relied upon to gauge a patent’s technological content. Finally, the average patent covers 188 gene-disease pairs, mostly due to a few outliers since the median patent is focused on a much smaller set of 13 gene-disease combinations.

Table D.1: Descriptive statistics at the patent level.

	USPTO patent applications						USPTO granted patents					
	mean	median	st d	min	max	N	mean	median	st d	min	max	N
Genes per patent	6.285	2	31.1786	1	4060	148,232	5.625	2	31.3932	1	4,060	73,255
Diseases per patent	12.257	5	22.5936	1	911	148,232	12.095	5	20.4855	1	853	73,255
Gene-disease pairs per patent	188.416	13	3,671.136	1	318,560	148,232	165.865	12	3645.6	1	289,212	73,255
Year of patent publication	2010.704	2011	5.0519	2001	2019	148,232	2011.062	2011	5.0451	2001	2019	73,255

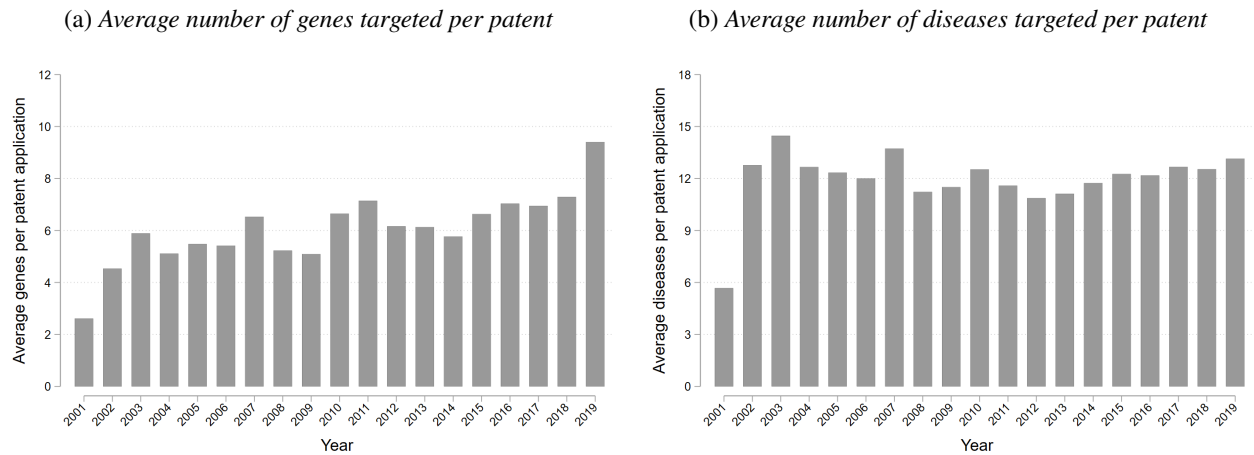
Note: This table presents descriptive statistics at the level of individual patent applications. The left part of the table presents descriptives considering all patent applications, while the right panel presents descriptives considering only applications that eventually result in a granted patent as of 2021. *Genes per patent*: count of genes targeted by the patent; *Diseases per patent*: count of diseases targeted by the patent; *Gene-disease pairs per patent*: count of gene-disease pairs targeted by the patent; *Year of patent publication*: year when the patent specification was published.

The average number of diseases each patent targets is roughly constant over my sample period (Figure D.2), while there seems to be a slightly upward trend in the number of genes targeted. Patents published in the year 2001 seem to reference an abnormally low number of bio-entities, possibly because of idiosyncrasies of that specific year (which is the first when patent applications started to be published and it corresponded to the release of the first draft of the human genome). Together,

³⁶This is consistent with evidence showing that patent examiners tend to restrict the scope of patent applications during the granting process (Marco et al., 2019).

the two graphs help rule out structural breaks in disclosure practices that could potentially bias my measurement approach.

Figure D.2: The average number of genes and diseases targeted by each USPTO patent application is relatively constant over time



Note: Panel (a) shows the average number of genes targeted by each USPTO patent application per year. Panel (b) shows the average number of diseases (MeSH Unique IDs at level 4) by each USPTO patent application per year.

Table D.2 shows the descriptive statistics at the firm level. On average, firms' patent portfolio targets around 64 genes and 41 diseases, but there is significant variation. Some firms span up to thousands of genes in their R&D efforts, while the median firm explored only 19 genes. The dispersion in the number of diseases is generally smaller. Interestingly, the table shows that focusing only on granted patents would lead to missing much of the firms' exploration in genetic space. This validates the choice of looking at patent applications to capture the earliest stages of pharmaceutical innovation and suggests that firms' competencies are much more comprehensive than what traditional research can capture (Brusoni et al., 2001). Firms are active on 1,437 gene-disease pairs on average, but the median firm explored only 162 pairs.

Table D.2: Descriptive statistics at the firm level.

	USPTO patent applications						USPTO granted patents					
	mean	median	st d	min	max	N	mean	median	st d	min	max	N
Genes per firm	64.285	19	195.922	1	5,467	4,117	37.637	11	130.255	1	4,579	3,866
Diseases per firm	41.362	27	46.173	1	338	4,117	31.530	17	40.546	1	305	3,866
Gene-disease pairs per firm	1,436.791	162	8,680.241	1	239,160	4,117	842.248	76	5,821.625	1	166,437	3,866

Note: This table presents descriptive statistics at the level of individual firms. The left part of the table presents descriptives considering all patent applications, while the right panel presents descriptives considering only applications that eventually result in a granted patent as of 2021. *Genes per firm*: count of genes targeted by a firm in its patents; *Diseases per firm*: count of diseases targeted by a firm in its patents; *Gene-disease pairs per firm*: count of gene-disease pairs targeted by a firm in its patents.

Finally, I provide some descriptive evidence of how the bio-entities recombined by a patent relate to common empirical proxies of its value. Simple OLS regressions reveal that the number of diseases is associated with patents of higher economic value (Table D.3). Innovations targeting more diseases have

higher market value and larger patent family sizes, indicating that the number of potential applications for a drug is a predictor of its economic value. Instead, the number of genetic targets of a patent seems to be associated with higher technological impact, as proxied by the number of forward patent citations received. These patents are also more likely to end up in litigation, confirming the intuition that they might cover a larger swath of the technological space and block other applications.

Table D.3: USPTO patent applications targeting multiple diseases have a higher market value, while USPTO patent applications targeting multiple genes have a higher technological impact.

Dependent Variable:	Patent family size		Litigated patent (0/1)		Patent citations		Market value	
Genes per patent	0.0004 (0.00222)		0.0004* (0.00017)		0.04659*** (0.01253)		-0.0274 (0.03072)	
Diseases per patent	0.0329** (0.01043)		0.0002 (0.00026)		0.0076 (0.02196)		0.1312* (0.06300)	
Year of application FE	YES	YES	YES	YES	YES	YES	YES	YES
N	148,226	148,226	148,226	148,226	148,226	148,226	30,019	30,019
Mean of Dep Var:	10.7969	10.7969	0.4366	0.4366	21.0813	21.0813	28.8594	28.8594

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the patent application level. Std. err. clustered at the assignee level. *Genes per patent*: count of genes targeted by a patent application; *Diseases per patent*: count of diseases targeted by a patent application; *Patent family size*: number of patent applications in the same patent family; *Litigated patent (0/1)*: 0/1 = 1 if the USPTO patent application is involved in a litigation. *Patent citations*: number of forward patent citations received by the USPTO patent application up to seven years after its publication; *Market value*: estimate of the market value (in constant USD) of the patent using data from Kogan et al. (2017); note that this measure is available only for applications that are eventually granted, hence the smaller sample size.

D.3 Example: Denali Therapeutics

In 2015, three top researchers left Genentech to start a new company: Denali Therapeutics. Aptly named after the tallest mountain in North America, Denali focuses on treating and curing neurodegenerative diseases like Alzheimer’s, amyotrophic lateral sclerosis, and Parkinson’s. Based in San Francisco, the company has raised over \$350 million in venture capital and \$250 million from its IPO in 2017. On its website, Denali lists ten compounds at different stages of clinical development as of summer 2022. Some of them are being developed together with large pharmaceutical companies, including Biogen, Sanofi, and Takeda Pharmaceuticals.

Advances in the genetics, pathology, and cell biology underlying chronic neurodegenerative disease have identified new pathways that trigger neurodegeneration. Denali Therapeutics was founded based on the idea that some such genes, dubbed “degenogenes,” could constitute a viable therapeutic avenue to tackle the most common neurodegenerative disorders. Degenogenes are genes that, when mutated, have a causative role in neurodegenerative disease. The focus of drug discovery activities on a handful of genetic targets constitutes not only the scientific foundation of this company but also its key competitive hypothesis for drug development.

In its IPO filings, Denali explicitly listed the key genetic targets that the company decided to focus

on.³⁷ This offers an opportunity to test the reliability of the SciBite data, that include ten patent applications by Denali Therapeutics as of 2019. Two patterns stand out. First, nine out of ten patents are indeed tagged with Alzheimer's and Parkinson's diseases, showing that the data accurately capture the markets targeted by Denali.³⁸ Second, 86.3% of the gene-disease combinations mentioned in its patents include the two key genes listed in the SEC files: RIPK1 and LRRK2. The first is a gene with an essential role in driving cell death and inflammation, and Denali was the first company to establish the safety of inhibiting RIPK1 kinase activity in humans with a Phase 1 clinical trial (Mifflin et al., 2020); the latter is a gene whose mutations increase the risk of developing Parkinson's disease, and Denali is leading the way in showing how it can be used for drug targeting (Kingwell, 2022). This example shows the potential of using bio-entities to capture a company's technological portfolio.

³⁷The list of genetic pathways and genes targets is at page 3 of the following link: <https://www.sec.gov/Archives/Denali>

³⁸Instead, the tenth patent generically addresses lysosomal storage disorders. As noted in the IPO filings of Denali, the lysosomal system is associated with several neurodegenerative diseases, including Parkinson's.

E Other Data Sources

The present appendix provides additional details and descriptive statistics about the other data used in this study.

E.1 The NHGRI-EBI GWAS Catalog

Information about genome-wide association studies is taken from the GWAS Catalog, a publicly available list of all published GWASs and association results (MacArthur et al., 2017). The Catalog includes all eligible GWAS studies since the first published in 2005, with details about associations with a high statistical significance ($p\text{-value} < 1.0 \times 10^{-5}$) (Marigorta et al., 2018). Compiling this resource requires manual curation of a large body of diverse and unstructured data from the literature, a task carried out by scientists at the European Bioinformatics Institute (EBI) with the support of the National Human Genome Research Institute Home (NHGRI). Catalog data are routinely used by biologists, bioinformaticians, and researchers aiming to translate scientific findings to medical applications and establish targets for novel therapies.

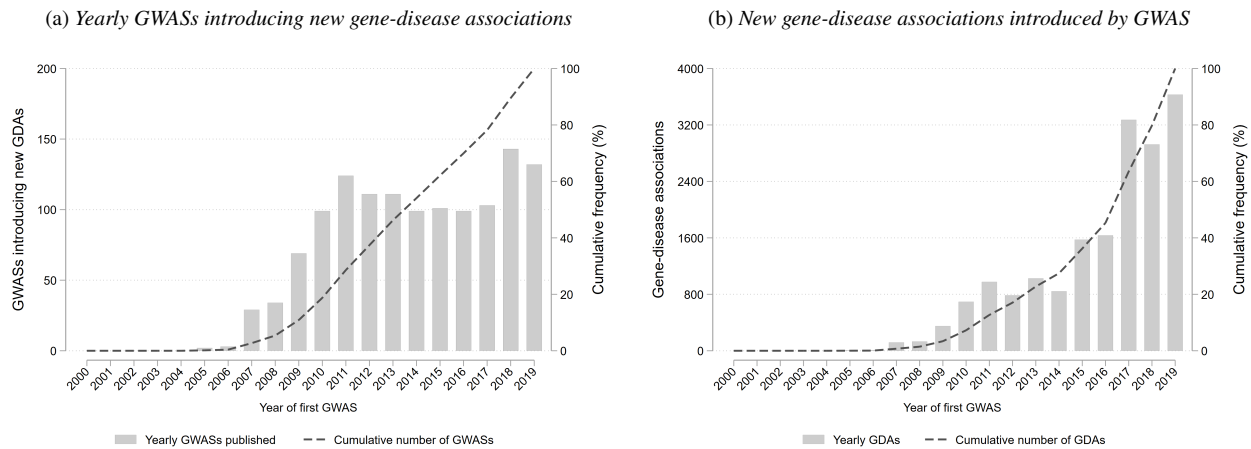
Table E.1: Descriptive statistics of GWAS papers that introduce a gene-disease association involving protein-coding genes

	mean	median	st d	min	max	N
Genes	7.848	3	25.849	1	522	1,259
Diseases	2.053	2	1.335	1	12	1,259
New gene-disease associations	14.269	4	36.276	1	522	1,259
Sample size	104,582.80	11,348	224,424.40	41	1,474,097	1,259
Replication sample (0/1)	0.639	1	0.48	0	1	1,259
Top journal (0/1)	0.314	0	0.464	0	1	1,259
High status PI (0/1)	0.261	0	0.439	0	1	1,259
Co-author industry (0/1)	0.218	0	0.413	0	1	1,259
Scientific citations	138.258	51	311.392	0	6,244	1,259
Cited by clinical trials (0/1)	0.57	1	0.495	0	1	1,259
Patent citations	1.221	0	7.459	0	226	1,259
Year	2015.823	2017	2.979	2005	2019	1,259

Note: This table presents descriptive statistics of GWASs that introduce new gene-disease associations involving protein-coding genes. *Genes*: number of protein-coding genes associated with a disease in the GWAS; *Diseases*: number of diseases studied in the GWAS; *New gene-disease associations*: number of new gene-disease associations introduced by the GWAS; *Sample size*: total number of subjects involved in the GWAS; *Replication sample (0/1)*: 0/1 = 1 for associations reported in GWASs that include also a replication analysis of their result; *Top journal (0/1)*: 0/1 = 1 for associations published in the 15 most prestigious genetics journals or the top 3 generalists scientific journals (*Science*, *Nature*, *PNAS*); *High status PI (0/1)*: 0/1 = 1 for GWASs whose last author is affiliated with one of the 20 most prestigious universities according to the QS World University Rankings for the biological sciences; *Co-author industry (0/1)*: 0/1 = 1 for GWASs with at least one industry co-author; *Scientific citations*: count of scientific citations received by the GWAS (data from NIH ICite); *Cited by clinical trials (0/1)*: 0/1 = 1 if the GWAS has received at least one citation from a clinical trial (data from NIH ICite); *Patent citations*: count of USPTO patent citations received by the GWAS (data for granted patents from Marx and Fuegi 2020); *Year*: year of publication of the GWAS.

Each entry of the Catalog includes details about the PubMed ID of the GWAS paper and the list of associated genes and diseases. Genes are identified by their NCBI IDs, while diseases are mapped into the Experimental Factor Ontology (EFO). I use the crosswalk available on the EFO website to map each disease into the corresponding MeSH Unique IDs. Note that the National Library of Medicine’s MeSH semantic keyword tree is a hierarchical tree with 13 levels of increasing specificity. For this study, I map each GWAS to the fourth level of the MeSH tree. If a more specific disease was

Figure E.1: Publication of GWASs and introduction of new gene-disease associations by year, 2000-2019



Note: Panel (a) shows the number of GWASs introducing at least one novel gene-disease association that involves a protein-coding gene, both by year and cumulatively over the sample period (2000-2019). Panel (b) shows the number of new gene-disease associations involving a protein-coding gene, both by year and cumulatively over the sample period (2000-2019).

matched (i.e., at level 5 or above), I assigned it to its parent branches up to level 4. Vice versa, if the disease matched was coarser (i.e., at level 3 or below), I assigned the finding as pertaining to all its descending level 4 branches. This procedure permits harmonizing GWASs as uncovering gene-disease associations in a landscape of NCBI IDs (unique genes) and level-4 MeSH IDs (unique diseases).

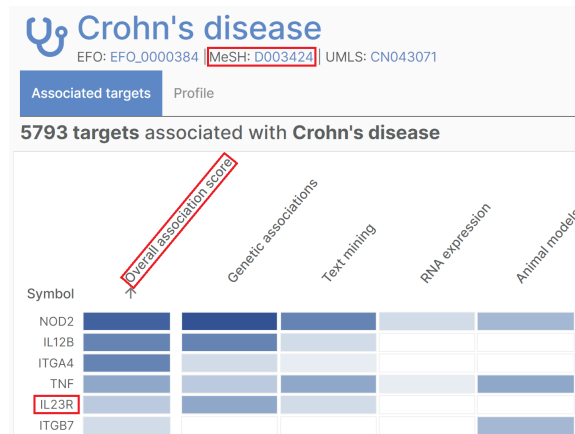
Table E.1 presents the descriptive statistics of the 1,259 GWAS papers that introduce new gene-disease associations (i.e., what constitutes the “treatment” in the primary analysis of the paper). The average GWAS targets two diseases and uncovers 14 associations. However, considerable variability exists, with a few GWASs recombining up to 522 genes. Around 64% of the GWASs include a replication sample within the same paper, which is considered a best practice to avoid reporting spurious results. Around a quarter of studies are carried out by principal investigators affiliated with high-status institutions, and little over a fifth includes an industry co-author. The average GWAS receives 138.3 scientific citations, 1.2 patent citations and has a 57% chance of being cited as the scientific background for a clinical trial. Figure E.1 shows the time series of GWAS publications and the arrival of new gene-disease associations. While the number of studies stabilized around 2010, the number of associations reported keeps growing, possibly due to increased sample sizes that allow for more statistically powered analyses.

E.2 Open Targets Score

The Open Targets Platform is a public-private partnership that aims to provide evidence for identifying and prioritizing drug targets, one of the fundamental challenges in developing new medicines (Ochoa et al., 2023). Open Targets collects all the available evidence on the strength of gene-disease

associations and summarizes it in an open-source synthetic score (Figure E.2 shows an example). In order to contextualize the relative importance of different pieces of evidence, Open Targets weights documentation on a gene-disease pair according to a scoring framework for each data source. All evidence is then mapped to the genetic target (using Ensembl gene identifiers, which I merge to NCBI IDs) and disease (using the EFO, which I merge to level-4 MeSH IDs). The Open Targets team also makes sure to minimize the presence of duplicates within the same data source. For this paper, I focus only on the Score aggregating direct evidence on the gene-disease relationship.³⁹

Figure E.2: Example of Open Targets Platform’s synthetic score for the genes associated with Crohn’s disease



Note: This figure shows an example of how the Open Targets Platform the available evidence on the strength of genetic associations with Crohn’s disease and summarizes them in an open-source synthetic score.

The Score is available for 594,353 gene-disease pairs (8% of my sample), spanning 17,437 genes and 366 diseases. However, one must take into consideration some potential limitations. Open Targets Scores partially reflect data availability about a given gene-disease pair. This means that under-studied genes or diseases are unlikely to produce high-scoring associations simply due to the lack of available evidence. Vice versa, not all pairs with available evidence can be considered legitimate genetic targets. These limitations explain why I use the replicability of a GWAS finding to find false positives and not the value of the Open Target Score, which could confound the interest in the association with its actual underlying therapeutic value. These limitations notwithstanding, this Score provides a valuable measure to rank the relative strength of genetic targets within a given disease.

Table E.2 validates that targeting gene-disease pairs with a higher Open Target Score has a technologically and economically significant effect on outcomes. Across the board, there is a strong correlation between patent applications that include gene-disease combinations with higher Open Target Scores and traditional metrics of patent value. In particular, the highest value reached by the Open Target Score of the combinations in a patent application correlates with the number of forward patent citations when including assignee fixed effects. Open Target Scores are also predictive of higher economic

³⁹The Open Targets Platform also provides scores that consider indirect evidence using the properties of the EFO disease ontology. However, this evidence is considered less robust (Ochoa et al., 2023).

Table E.2: Patent applications involving gene-disease pairs with higher Open Target Scores are of higher value

Dependent Variable:	Patent citations		Patent family size		Market value	
Max(OT Score)	4.6822 (3.4223)	7.7646* (3.10934)	3.8489*** (0.36592)	1.6012*** (0.31473)	14.7768** (4.7530)	4.0955* (2.0507)
Patent year FE	YES	YES	YES	YES	YES	YES
Firm FE	NO	YES	NO	YES	NO	YES
N	137,614	137,614	137,614	137,614	27,213	27,213
Mean of Dep Var:	26.063	26.063	11.081	11.081	30.365	30.365

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the patent application level. Std. err. clustered at the patent assignee level. *Max(OT Score)*: maximum value of the Open Target score reached by any gene-disease pair targeted in a patent application; *Patent citations*: number of forward patent citations received by the USPTO patent application up to seven years after its publication; *Patent family size*: number of patent applications in the same patent family; *Market value*: estimate of the market value (in constant USD) of the patent using data from Kogan et al. (2017); note that this measure is available only for applications that are eventually granted, hence the smaller sample size.

value of the patents, as captured by the dimension of the patent family and the monetary value of granted patents (using data from Kogan et al. 2017).

E.3 PubTator Central Publication Data

Data on the genes and diseases studied in each scientific publication are taken from PubTator Central, a web-based tool that automatically annotates biomedical concepts in PubMed abstracts and text (Wei et al., 2019). Articles are processed through concept taggers and disambiguation dictionaries to resolve annotation conflicts. The results of this process are publicly available online and include over 29 million abstracts and 3 million full-text documents. The entities extracted are matched to unique identifiers from NCBI and MeSH. The F1 score for the entity extraction pipeline is 86.70%

I obtain information on the authors' affiliations from Dimensions, a data product by Digital Science (Herzog et al., 2020). Compared to PubMed, whose recording of authors' affiliations is mainly limited to the corresponding author for papers before 2010, Dimensions has much broader coverage. I use these data to match each patenting firm in my sample to their publication portfolios using fuzzy string matching on the firm names. Then, thanks to PubTator Central, I record all the genes and diseases studied by the firms. Finally, I use NIH's iCite data to collect additional bibliometric information, including the translational focus of each article and its impact on clinical development.

Table E.3 shows the descriptive statistics of the resulting dataset. 3,346 of the 4,117 firms in my sample (81.3%) have at least one publication. Each publishing firm has on average 59 papers about 177 genes and 109 diseases. The median firm has performed basic research on 923 gene-disease pairs. A large fraction of publications falls within the scope of what Azoulay et al. (2021) define "translational research."⁴⁰ Interestingly, Figure E.3 shows that firms are more likely to carry out research involving

⁴⁰Translational articles are defined as disease-oriented studies that try to apply bench science findings to practical therapies,

Table E.3: Descriptive statistics of firms' publication portfolios


	mean	median	st d	min	max	N
Publications	59.170	9	390.944	1	10,656	3,346
Genes	176.563	52	568.922	1	12,267	3,346
Diseases	109.239	65	156.608	1	1,876	3,346
Gene-disease pairs	5,507.258	923	27,322.490	1	545,193	3,346
Biological pathway	0.381	0	2.496	0	58	3,346
Translational (cell)	12.48	2	70.712	0	1,820	3,346
Translational (animal)	11.575	1	11.575	0	1,673	3,346
Translational (human)	23.558	2	202.75	0	7,511	3,346
Year	2012.97	2014	4.672	1985	2019	3,346

This table presents descriptive statistics of firms' publication portfolios (conditional on having at least one publication). *Publications*: average number of firm publications; *Genes*: average number of genes on which the firm has published; *Diseases*: average number of diseases on which the firm has published; *Gene-disease pairs*: average number of gene-disease pairs on which the firm has published; *Biological pathway*: average number of firm publications that include a biological pathway; *Translational (cell)*: average number of firm translational publications (according to the definition of Azoulay et al. 2021) that are cellular biology-oriented; *Translational (animal)*: average number of firm translational publications (according to the definition of Azoulay et al. 2021) that involve animal organisms; *Translational (human)*: average number of firm translational publications (according to the definition of Azoulay et al. 2021) that involve human subjects; *Year*: average year of the publications.

animal model organisms than university researchers. The opposite is true when it comes to research including human subjects, possibly because of the higher costs and risks involved or because clinical trials tend to be outsourced to university hospitals (Azoulay and Fishman, 2020). Firms' publications are also disproportionately likely to be cited by clinical trials, confirming the applied nature of corporate research relative to academic science.

but without being clinical trials.

Figure E.3: Articles by firms are more likely to use animal organisms and to be relevant for clinical trials than otherwise comparable articles by university researchers.



Note: This figure shows the OLS coefficient from linear probability models that regress the type of publication on a dummy equal to 1 if the publication includes at least one author from a company in my sample. All regressions control for year and journal dummies. Standard errors are clustered at the journal level.


Table E.4: Patent applications are of higher value when they target gene-disease pairs involving genes where the firm has published before

Dependent Variable:	Patent citations		Patent family size		Market value	
Has gene knowledge (share)	-0.1496 (3.1426)	8.0016* (3.2385)	4.0945*** (0.6110)	0.9079*** (0.2152)	24.7607** (7.8265)	0.488 (1.0171)
Patent year FE	YES	YES	YES	YES	YES	YES
Firm FE	NO	YES	NO	YES	NO	YES
N	134,459	134,330	134,459	134,330	26,498	26,329
Mean of Dep Var:	26.166	26.166	11.107	11.107	30.527	30.527

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the patent application level. Std. err. clustered at the patent assignee level. *Has gene knowledge (share)*: share gene-disease pairs targeted in a patent that involve genes where the firm has published before; *Patent citations*: number of forward patent citations received by the USPTO patent application up to seven years after its publication; *Patent family size*: number of patent applications in the same patent family; *Market value*: estimate of the market value (in constant USD) of the patent using data from Kogan et al. (2017); note that this measure is available only for applications that are eventually granted, hence the smaller sample size.

Table E.4 reports suggestive evidence that having domain knowledge on a specific gene (as proxied by past publications) leads to firm patents of higher technological and economic value. The only partial exception is market value, where the effect vanishes after including firm fixed effects. However, this is likely due to how the measure of value is constructed from firms’ stock market fluctuations, which does not allow for the detection of within-firm differences in patent value (Kogan et al., 2017). Finally, Figure E.4 shows a binscatter plotting the relationship between the average Open Target Score of gene-disease pairs targeted in a USPTO patent application and the share of those pairs that include a gene that the firm has previously studied. The figure makes clear that the average therapeutic value of gene-disease combinations targeted by a patent is higher when involving genes on which the firm

Figure E.4: The average therapeutic value of gene-disease combinations targeted by a patent is higher when involving genes on which the firm has published before



Note: This figure shows a binscatter plotting the association between the average Open Target Score of gene-disease pairs targeted in a USPTO patent application and the share of those pairs that include a gene that has been previously studied by the firm. Both publication counts and Open Targets Scores are residualized by firm-disease fixed effects.

has published before, previewing the main results of this paper.

E.4 Cortellis Drug Data

I obtain drug development records from Cortellis, which contains development information for 42,896 drugs targeting at least one of the gene-disease pairs in my sample. Cortellis aggregates information from various sources to assemble a list of historical development milestones for each drug molecule. This paper’s analyses use those milestones to construct complete drug development histories for each drug whose genetic target and disease indication is recorded (Krieger, 2021). In particular, I focus on new molecules observed entering the earliest phases of drug development (what Cortellis records as the “discovery phase” and the “pre-clinical phase”). In additional analyses reported in the main text, I follow the approach of Dranove et al. (2022) and weigh each drug by its relative novelty. The basic intuition of this weighting is that the higher the number of previous drugs that adopted the same molecular-targeting design, the less scientifically novel the drug is.

Table E.5 confirms the intuition that gene-disease combinations with at least one drug molecule received a higher number of investments. Pairs with successful drug discovery activities appear in 112 patent applications and 2,176 publications, while the others only receive two patents and 47 publications on average. Interestingly, the difference is much lower when measured using the Open Target Score. This suggests that research and development inputs might disproportionately target gene-disease pairs known to be druggable but potentially miss out on other therapeutic opportunities (Oprea et al., 2018; Stoeger et al., 2018).

Table E.5: Descriptive statistics at the gene-disease pair level.

	Gene-disease pairs with drugs						Gene-disease pairs without drugs					
	mean	median	st d	min	max	N	mean	median	st d	min	max	N
Patents on GDA	112.463	25	319.223	0	13,900	13,582	2.289	0	19.631	0	8,133	7,210,342
Publications on GDA	2,175.96	425	6,192.69	0	173,306	13,582	47.250	2	424.712	0	98,941	7,210,342
Open Targets Score	0.099	0.025	0.158	0.00005	0.897	9,782	0.039	0.007	0.085	0.00004	0.874	584,571

Note: This table presents descriptive statistics at the level of gene-disease pairs. The left part of the table presents descriptives considering only pairs that have at least one drug molecule listed in the Cortellis data, while the right panel presents descriptives considering only pairs that do not have drug molecules listed in the Cortellis data as of 2020. *Patents on GDA*: average count of patents on the gene-disease pair; *Publications on GDA*: average count of publications on the gene-disease pair; *Open Targets Score*: average Open Targets Score of the gene-disease pair.

Appendix References

- ABRAMS, D. S., U. AKCIGIT, AND J. GRENNAN (2018): “Patent value and citations: Creative destruction or strategic disruption?” *NBER wp 19647*.
- AGUIAR, L. AND J. WALDFOGEL (2018): “Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music,” *Journal of Political Economy*, 126, 492–524.
- AHMADPOOR, M. AND B. F. JONES (2017): “The dual frontier: Patented inventions and prior scientific advance,” *Science*, 357, 583–587.
- ANDRIANI, P. AND G. CATTANI (2022): “Functional diversification and exaptation: The emergence of new drug uses in the pharma industry,” *Industrial and Corporate Change*, 31, 1177–1201.
- ARORA, A., S. BELENZON, AND L. SHEER (2021): “Knowledge spillovers and corporate investment in scientific research,” *American Economic Review*, 111, 871–98.
- ARTS, S., B. CASSIMAN, AND J. C. GOMEZ (2018): “Text matching to measure patent similarity,” *Strategic Management Journal*, 39, 62–84.
- AZOULAY, P. AND A. FISHMAN (2020): “The rise of for-profit experimental medicine,” *NBER wp 26892*.
- AZOULAY, P., J. S. GRAFF ZIVIN, D. LI, AND B. N. SAMPAT (2019): “Public R&D investments and private-sector patenting: Evidence from NIH funding rules,” *The Review of Economic Studies*, 86, 117–152.
- AZOULAY, P., W. H. GREENBLATT, AND M. L. HEGGENESS (2021): “Long-term effects from early exposure to research: Evidence from the NIH “Yellow Berets”,” *Research Policy*, 50, 104332.
- BEYER, B. M., R. INGRAM, L. RAMANATHAN, P. REICHERT, H. V. LE, V. MADISON, AND P. ORTH (2008): “Crystal structures of the pro-inflammatory cytokine interleukin-23 and its complex with a high-affinity neutralizing antibody,” *Journal of Molecular Biology*, 382, 942–955.
- BIANCHI, E. AND L. ROGGE (2019): “The IL-23/IL-17 pathway in human chronic inflammatory diseases—new insight from genetics and targeted therapies,” *Genes & Immunity*, 20, 415–425.
- BOYACK, K. W., C. SMITH, AND R. KLAVANS (2020): “A detailed open access model of the PubMed literature,” *Scientific Data*, 7, 408.
- BOYLE, E. A., Y. I. LI, AND J. K. PRITCHARD (2017): “An expanded view of complex traits: From polygenic to omnigenic,” *Cell*, 169, 1177–1186.
- BRUSONI, S., A. PRENCIPE, AND K. PAVITT (2001): “Knowledge specialization, organizational coupling, and the boundaries of the firm: Why do firms know more than they make?” *Administrative Science Quarterly*, 46, 597–621.
- BRYAN, K. A., Y. OZCAN, AND B. SAMPAT (2020): “In-text patent citations: A user’s guide,” *Research Policy*, 49, 103946.
- BUSH, W. S. AND J. H. MOORE (2012): “Genome-wide association studies,” *PLoS Computational Biology*, 8, e1002822.
- COHEN, W. M., R. R. NELSON, AND J. P. WALSH (2002): “Links and impacts: The influence of public research on industrial R&D,” *Management Science*, 48, 1–23.
- DAHLHAMER, J. M., E. P. ZAMMITTI, B. W. WARD, A. G. WHEATON, AND J. B. CROFT (2016): “Prevalence of inflammatory bowel disease among adults aged ≤ 18 years—United States, 2015,” *Morbidity and Mortality Weekly Report*, 65, 1166–1169.
- DING, Y., M. SONG, J. HAN, Q. YU, E. YAN, L. LIN, AND T. CHAMBERS (2013): “Entitymetrics: Measuring the impact of entities,” *PloS one*, 8, e71416.


- DRANOVE, D., C. GARTHWAITE, AND M. HERMOSILLA (2022): “Does consumer demand pull scientifically novel drug innovation?” *The RAND Journal of Economics*, 53, 590–638.
- DUERR, R. H., K. D. TAYLOR, S. R. BRANT, J. D. RIOUX, M. S. SILVERBERG, M. J. DALY, ET AL. (2006): “A genome-wide association study identifies IL23R as an inflammatory bowel disease gene,” *Science*, 314, 1461–1463.
- EASTON, D. F., K. A. POOLEY, A. M. DUNNING, P. D. PHAROAH, ET AL. (2007): “Genome-wide association study identifies novel breast cancer susceptibility loci,” *Nature*, 447, 1087–1093.
- FLEMING, L. (2001): “Recombinant uncertainty in technological search,” *Management Science*, 47, 117–132.
- FLEMING, L., H. GREENE, G. LI, M. MARX, AND D. YAO (2019): “Government-funded research increasingly fuels innovation,” *Science*, 364, 1139–1141.
- GAVETTI, G. AND D. LEVINTHAL (2000): “Looking forward and looking backward: Cognitive and experiential search,” *Administrative Science Quarterly*, 45, 113–137.
- GOLDSTEIN, D. B. (2009): “Common genetic variation and human traits,” *New England Journal of Medicine*, 360, 1696.
- HERZOG, C., D. HOOK, AND S. KONKIEL (2020): “Dimensions: Bringing down barriers between scientometricians and data,” *Quantitative Science Studies*, 1, 387–395.
- IARIA, A., C. SCHWARZ, AND F. WALDINGER (2018): “Frontier knowledge and scientific production: Evidence from the collapse of international science,” *The Quarterly Journal of Economics*, 133, 927–991.
- JAFFE, A. B., M. TRAJTENBERG, AND M. S. FOGARTY (2000): “Knowledge spillovers and patent citations: Evidence from a survey of inventors,” *American Economic Review*, 90, 215–218.
- KAO, J. (2023): “Charted territory: Mapping the cancer genome and R&D decisions in the pharmaceutical industry,” *UCLA Anderson*.
- KAPLAN, S. AND K. VAKILI (2015): “The double-edged sword of recombination in breakthrough innovation,” *Strategic Management Journal*, 36, 1435–1457.
- KINGWELL, K. (2022): “LRRK2-targeted Parkinson disease drug advances into phase III.” *Nature Reviews Drug Discovery*.
- KOGAN, L., D. PAPANIKOLAOU, A. SERU, AND N. STOFFMAN (2017): “Technological innovation, resource allocation, and growth,” *The Quarterly Journal of Economics*, 132, 665–712.
- KUHN, J., K. YOUNGE, AND A. MARCO (2020): “Patent citations reexamined,” *The RAND Journal of Economics*, 51, 109–132.
- KUMAR PANDA, B. AND M. TRANCHERO (2023): “Measuring scientific impact and knowledge spillovers using entitymetrics,” *mimeo*.
- LAMPE, R. (2012): “Strategic citation,” *Review of Economics and Statistics*, 94, 320–333.
- MACARTHUR, D. (2012): “Face up to false positives,” *Nature*, 487, 427–428.
- MACARTHUR, J., E. BOWLER, M. CEREZO, L. GIL, P. HALL, E. HASTINGS, H. JUNKINS, A. McMAHON, A. MILANO, J. MORALES, ET AL. (2017): “The new NHGRI-EBI Catalog of published genome-wide association studies (GWAS Catalog),” *Nucleic Acids Research*, 45, D896–D901.
- MARCO, A. C., J. D. SARNOFF, AND A. CHARLES (2019): “Patent claims and patent scope,” *Research Policy*, 48, 103790.
- MARIGORTA, U. M., J. A. RODRÍGUEZ, G. GIBSON, AND A. NAVARRO (2018): “Replicability and prediction: Lessons and challenges from GWAS,” *Trends in Genetics*, 34, 504–517.

- MARX, M. AND A. FUEGI (2020): “Reliance on science: Worldwide front-page patent citations to scientific articles,” *Strategic Management Journal*, 41, 1572–1594.
- MIFFLIN, L., D. OFENGEIM, AND J. YUAN (2020): “Receptor-interacting protein kinase 1 (RIPK1) as a therapeutic target,” *Nature Reviews Drug Discovery*, 19, 553–571.
- MYERS, K. R. AND L. LANAHAN (2022): “Estimating spillovers from publicly funded R&D: Evidence from the US Department of Energy,” *American Economic Review*, 112, 2393–2423.
- NARIN, F., K. S. HAMILTON, AND D. OLIVASTRO (1997): “The increasing linkage between US technology and public science,” *Research Policy*, 26, 317–330.
- NELSON, M. R., H. TIPNEY, J. L. PAINTER, J. SHEN, P. NICOLETTI, Y. SHEN, ET AL. (2015): “The support of human genetic evidence for approved drug indications,” *Nature Genetics*, 47, 856–860.
- OCHOA, D., A. HERCULES, M. CARMONA, D. SUVEGES, J. BAKER, ET AL. (2023): “The next-generation Open Targets Platform: Reimagined, redesigned, rebuilt,” *Nucleic Acids Research*, 51, D1353–D1359.
- OPREA, T. I., C. G. BOLOGA, S. BRUNAK, A. CAMPBELL, G. N. GAN, A. GAULTON, S. M. GOMEZ, R. GUHA, A. HERSEY, J. HOLMES, ET AL. (2018): “Unexplored therapeutic opportunities in the human genome,” *Nature Reviews Drug Discovery*, 17, 317–332.
- PUSHPAKOM, S., F. IORIO, P. A. EYERS, ET AL. (2019): “Drug repurposing: Progress, challenges and recommendations,” *Nature Reviews Drug Discovery*, 18, 41–58.
- REAY, W. R. AND M. J. CAIRNS (2021): “Advancing the use of genome-wide association studies for drug repurposing,” *Nature Reviews Genetics*, 22, 658–671.
- RIGHI, C. AND T. SIMCOE (2020): *Patenting inventions or inventing patents? Strategic use of continuations at the USPTO*, NBER wp 27686.
- ROACH, M. AND W. M. COHEN (2013): “Lens or prism? Patent citations as a measure of knowledge flows from public research,” *Management Science*, 59, 504–525.
- RZHETSKY, A., J. G. FOSTER, I. T. FOSTER, AND J. A. EVANS (2015): “Choosing experiments to accelerate collective discovery,” *Proceedings of the National Academy of Sciences*, 112, 14569–14574.
- STOEGER, T., M. GERLACH, R. I. MORIMOTO, AND L. A. NUNES AMARAL (2018): “Large-scale investigation of the reasons why potentially important genes are ignored,” *PLoS Biology*, 16, e2006643.
- SUH, J. (2023): “Science, startups, and novelty,” *Available at SSRN 4383392*.
- TEPLITSKIY, M., E. DUEDE, M. MENIETTI, AND K. R. LAKHANI (2022): “How status of research papers affects the way they are read and cited,” *Research Policy*, 51, 104484.
- TRANCHERO, M. (2023): “Data-driven search and innovation in well-defined technological landscapes,” *UC Berkeley*.
- VISSCHER, P. M., N. R. WRAY, Q. ZHANG, P. SKLAR, M. I. MCCARTHY, M. A. BROWN, AND J. YANG (2017): “10 years of GWAS discovery: biology, function, and translation,” *The American Journal of Human Genetics*, 101, 5–22.
- WEI, C.-H., A. ALLOT, R. LEAMAN, AND Z. LU (2019): “PubTator central: Automated concept annotation for biomedical full text articles,” *Nucleic Acids Research*, 47, W587–W593.
- XU, J., S. KIM, M. SONG, M. JEONG, D. KIM, J. KANG, J. F. ROUSSEAU, X. LI, W. XU, V. I. TORVIK, ET AL. (2020): “Building a PubMed knowledge graph,” *Scientific Data*, 7, 1–15.
- YOUNGE, K. A. AND J. M. KUHN (2019): “First movers and follow-on invention: Evidence from a vector space model of invention,” *Available at SSRN 3354530*.


F Additional Figures and Tables

Figure F.1: While there is clear evidence of time sorting on the diseases that receive a GWAS, the same is not true for genes

(a) Average past patents on the diseases by year of first GWAS

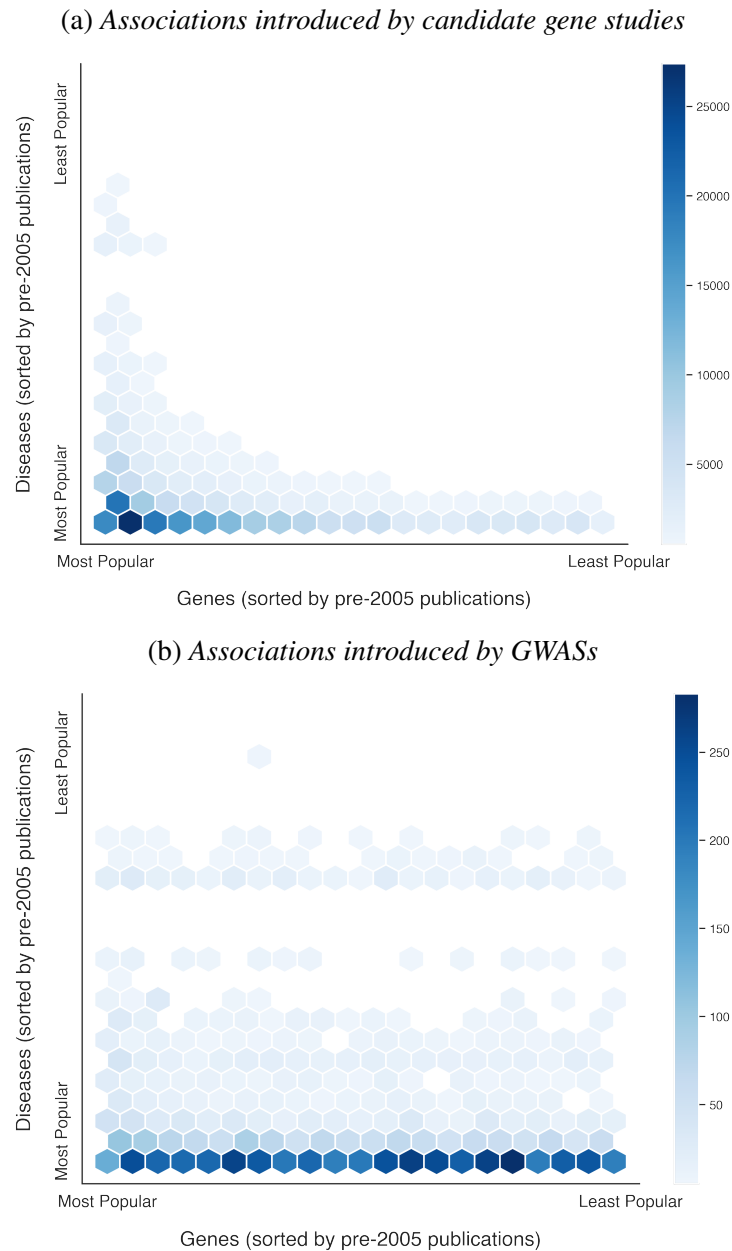


(b) Average past patents on the genes by year of first GWAS



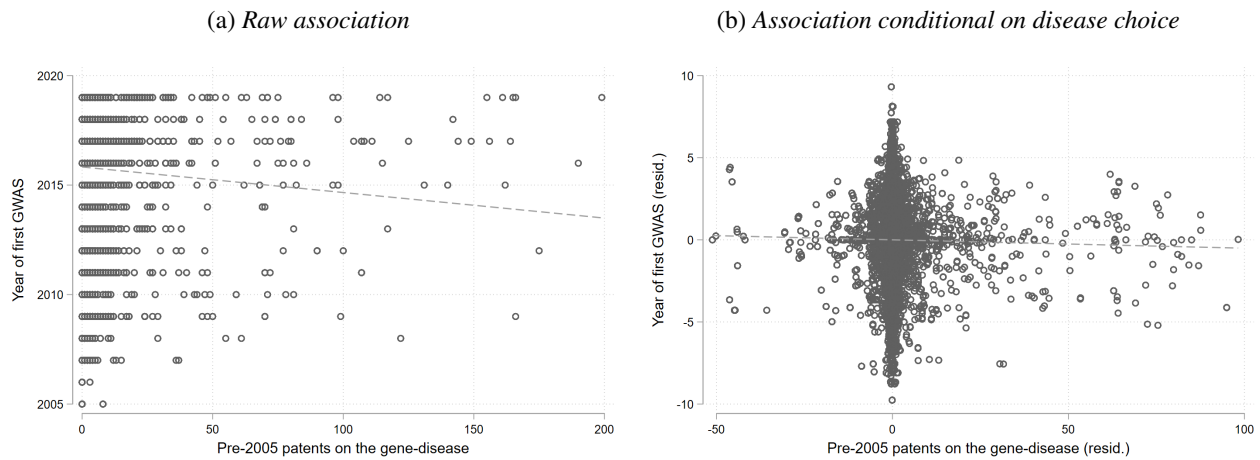
Note: Panel (a) shows the average number of patent applications received by each disease before 2005, separately by the year when the disease received its first GWAS. Panel (b) shows the average number of patent applications received by each gene before 2005, separately by the year when the gene received its first GWAS. The figure shows that diseases that received a GWAS earlier in time tend to be those with higher pre-GWAS patenting activity, confirming the prioritization of large and important diseases, while the same selection is not apparent for genes due to the research design of GWASs.

Figure F.2: The patterns of gene-disease associations reported by GWASs confirm that they do not target any gene, unlike candidate-gene studies



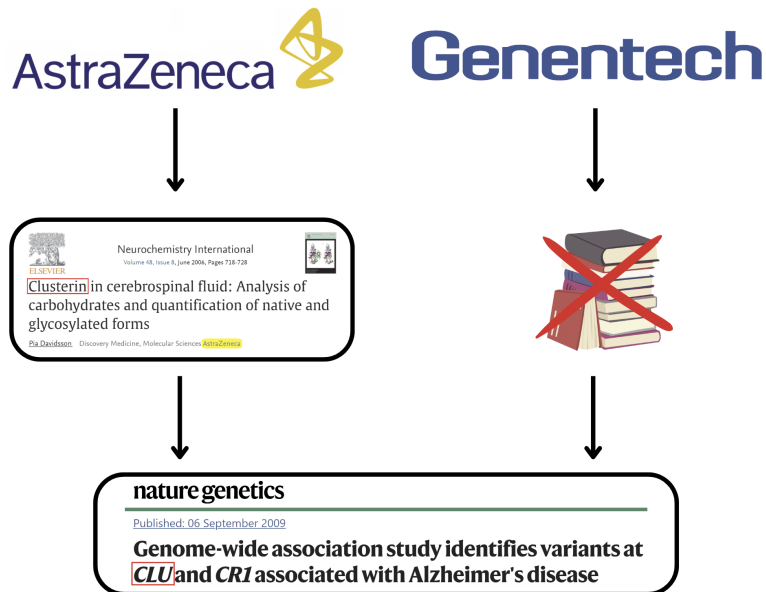
Note: Panel (a) shows a heatmap of new gene-disease associations introduced after 2005 with targeted search strategies. Panel (b) shows a heatmap of new gene-disease associations introduced after 2005 with data-driven search strategies. Both panels have 14,112 genes on the X axis, sorted from the most to the least studied in the pre-GWAS era, and 15,039 narrow disease categories on the Y axis, sorted from the most to the least studied in the pre-GWAS era. Darker areas denote a higher number of new gene-disease associations in correspondence with genes and diseases of that level of popularity in the pre-GWAS era. Figure from Tranchero (2023) using the DisGeNET database.

Figure F.3: The timing of GWAS associations is not related to past patenting on gene-diseases pairs, once controlling for disease selection



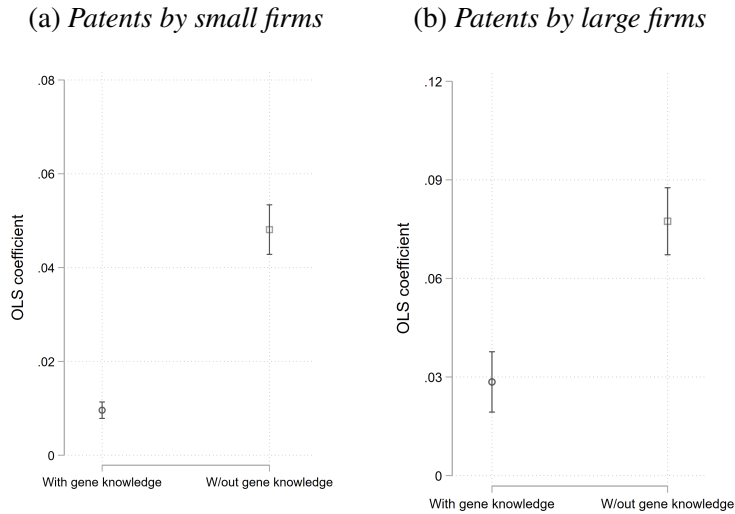
Note: These figures show the association between past patent applications of the 17,965 treated gene-disease pairs involving a protein-coding gene and the year of the first GWAS association reporting them. Panel (a) presents the raw scatterplot. Panel (b) shows the same scatterplot after residualizing for disease.

Figure F.4: Schema of the between-firms research design used for gene-disease level analyses




Note: This figure exemplifies the between-firm design used to assess the role of domain knowledge in gene-disease level regressions. Some firms evaluating the plausibility of a GWAS association can leverage their previous research on the gene involved, thus improving their assessment thanks to a deeper understanding of the gene's biology.

Figure F.5: Firms react less to the arrival of GWAS associations when they possess domain knowledge, regardless of their size



Note: The figure reports OLS coefficients capturing the main effect of GWASs on USPTO patent applications, separately by firms with and without previous publications on the gene involved. Each coefficient is estimated with the same specifications reported in Table 4. Panel (a) reports the two estimates for small firms, defined as the ones with a below-median number of patents in my samples. Panel (b) reports the two estimates for large firms, defined as the ones with an above-median number of patents in my samples.

Figure F.6: Average increase in patent applications on treated gene-disease pairs, by type of firm



Note: This figure shows the average increase in USPTO patent applications on a gene-disease pair after the publication of a GWAS (considering only treated pairs). The increase is presented separately for GWAS findings that are not replicated in subsequent studies and those that are confirmed by later GWASs. The first two columns report the patenting increase by firms with previous publications on the gene involved, while the latter two columns report the patenting increase by firms without them.

Table F.1: No evidence of association between past patenting and the timing of the first GWAS for treated gene-disease pairs

Dependent Variable:	Year of first GWAS			
Patents about gene	-0.007 [-0.65]	-0.014 [-1.87]		
Patents about gene-disease		-0.040* [-2.31]	-0.020 [-1.40]	
Disease FE	NO	YES	NO	YES
Gene FE	NO	NO	YES	YES
N of obs	17,925	17,923	16,338	16,298
Mean of Dep Var:	2015.8	2015.8	2015.8	2015.8

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease association level. Std. err. clustered two-ways at the disease and gene level. This table shows standardized beta coefficients to ease comparison and reports t-values inside square brackets. *Year of first GWAS*: the year when the gene-disease pair was first reported as significant in one GWAS paper. *Patents about gene*: sum of USPTO patent applications filed before 2005 that target a specific gene. *Patents about gene-disease*: sum of USPTO patent applications filed before 2005 that target a specific gene-disease pair.

Table F.2: GWAS findings that are subsequently replicated involve gene-disease pairs with a higher Open Targets Score

Dependent Variable:	Open Targets Score		Top 90%ile OT Score (0/1)	
Replicable association (0/1)	0.0937*** (0.00900)	0.0457*** (0.00645)	0.1909*** (0.01895)	0.1139*** (0.01442)
Gene FE	YES	YES	YES	YES
Disease FE	YES	YES	YES	YES
Year of GWAS FE	YES	YES	YES	YES
Sources Count FE	NO	YES	NO	YES
N of GDAs	16,298	8,921	16,298	8,921
Mean of Dep Var:	0.0894	0.0894	0.0998	0.0998

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease association level. Std. err. clustered at the gene level. *Open Targets Score*: value of the Open Targets Score of the gene-disease pair. *Top 90%ile OT Score*: 0/1 = 1 if the gene-disease pair is in the top decile of Open Targets Score in my sample of GWAS associations. *Replicable association*: 0/1 = 1 if the gene-disease association is replicated by subsequent GWASs about the same disease. *Sources Count FE*: number of sources aggregated by Open Targets to compute the Score of a given gene-disease pair.

Table F.3: Correlates of replicable GWAS findings: study design quality.

Dependent Variable:	Replicable association (0/1)		
Large sample (0/1)	0.3197*** (0.02069)		
Replication sample (0/1)	0.2192*** (0.01811)		
Powerful genotyping array (0/1)	0.0993*** (0.02237)		
Disease FE	YES	YES	YES
Year of GWAS FE	YES	YES	YES
N of GDAs	17,923	17,923	16,161
Mean of Dep Var:	0.1574	0.1574	0.1574

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease association level. Std. err. clustered at the gene level. *Replicable association*: 0/1 = 1 if the gene-disease association is replicated by subsequent GWASs about the same disease. *Large sample*: 0/1 = 1 for associations published in studies with a sample size larger than the median GWAS in my sample. *Replication sample*: 0/1 = 1 for associations reported in papers that include also a replication analysis of their result. *Powerful genotyping array*: 0/1 = 1 if the gene-disease association is obtained using a microarray with an above-median number of SNPs.

Table F.4: Correlates of replicable GWAS findings: citations to the article introducing them.

Dependent Variable:	Replicable association (0/1)		
Scientific citations to GWAS	172.09*** (35.417)		
Clinical citations to GWAS	0.0926*** (0.01761)		
Share of negative citations	-0.0488** (0.01525)		
Disease FE	YES	YES	YES
Year of GWAS FE	YES	YES	YES
N of GDAs	17,923	17,923	13,732
Mean of Dep Var:	0.1574	0.1574	0.1574

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease association level. Std. err. clustered at the gene level. *Replicable association*: 0/1 = 1 if the gene-disease association is replicated by subsequent GWASs about the same disease. *Scientific citations to GWAS*: number of citations from scientific papers received by the GWAS introducing the gene-disease association. *Clinical citations to GWAS*: number of citations from clinical trials received by the GWAS introducing the gene-disease association. *Share of negative citations*: share of citations with a negative tone from scientific papers received by the GWAS introducing the gene-disease association (data from Scite).

Table F.5: Firms increase their investments more when the associations are statistically stronger or the GWASs have better research designs

Dependent Variable:	USPTO patent applications				
Post × GWAS	0.0432*	0.0779***	0.0581***	0.0828***	0.0432*
	(0.02019)	(0.02019)	(0.01815)	(0.01809)	(0.01871)
...× High P-value (0/1)	0.2455***				
	(0.03989)				
...× Large effect (0/1)		0.1970***			
		(0.04246)			
...× Top journal (0/1)			0.292***		
			(0.04810)		
...× Large sample (0/1)				0.1830***	
				(0.04271)	
...× Replication sample (0/1)					0.2449***
					(0.04007)
Gene-Disease FE	YES	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES	YES
N	137,254,556	137,254,556	137,254,556	137,254,556	137,254,556
N of Gene-Diseases	7,223,924	7,223,924	7,223,924	7,223,924	7,223,924
Mean of Dep Var:	0.1314	0.1314	0.1314	0.1314	0.1314

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. *High P-value*: 0/1 = 1 for associations with a p-value larger than the median GWAS in my sample. *Large effect*: 0/1 = 1 for associations with an effect size larger than the median GWAS in my sample. *Top journal*: 0/1 = 1 for associations published in the 15 most prestigious genetics journals or the top 3 generalist scientific journals (*Science*, *Nature*, *PNAS*). *Large sample*: 0/1 = 1 for associations published in studies with a sample size larger than the median GWAS in my sample. *Replication sample*: 0/1 = 1 for associations reported in papers that include also a replication analysis of their result.

Table F.6: Firms with genetic domain knowledge introduce new drugs on true positive GWAS findings, while firms lacking domain knowledge miss these opportunities

Dependent Variable:	New drugs by...			
	...firms with gene knowledge		...firms w/out gene knowledge	
Post × GWAS	0.000113	-0.0000653	0.000221	-0.0000121
	(0.0000691)	(0.0000384)	(0.0001569)	(0.0001093)
...× True Positive		0.000892**		0.0011658
		(0.0003058)		(0.0006582)
Gene-Disease FE	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES
N	137,254,556	137,254,556	137,254,556	137,254,556
N of Gene-Diseases	7,223,924	7,223,924	7,223,924	7,223,924
Mean of Dep Var:	0.0000138	0.0000138	0.0000827	0.0000827

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *New drugs*: number of molecules on a gene-disease combination entering the discovery stage (data from Cortellis); the count is then divided between firms with and without previous publications on the gene. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. *True Positive*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease.

Table F.7: Firms without domain knowledge struggle more to recognize opportunities when the GWAS association involves less-studied genes

Dependent Variable:	USPTO patent applications by...		
	... all firms	...firms with gene knowledge	...firms w/out gene knowledge
Post \times GWAS	0.0123** (0.00462)	0.0024** (0.00092)	0.0099* (0.00431)
... \times True Positive	0.0265 (0.01458)	0.0202* (0.00876)	0.0063 (0.00891)
Gene-Disease FE	YES	YES	YES
Disease-Year FE	YES	YES	YES
Gene-Year FE	YES	YES	YES
N	66,980,776	66,980,776	66,980,776
N of Gene-Diseases	3,525,304	3,525,304	3,525,304
Mean of Dep Var:	0.02097	0.00072	0.02026

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination; the count is then divided between firms with and without previous publications on the gene. *Post \times GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS. The sample is restricted to gene-disease pairs involving genes that received a below-median number of scientific studies before 2005 (the year of the first GWAS).

Table F.8: Firms without domain knowledge invest proportionally less in false positives when the GWASs report a smaller number of associations

Dependent Variable:	USPTO patent applications by...		
	... all firms	...firms with gene knowledge	...firms w/out gene knowledge
Post \times GWAS	0.0670** (0.02019)	0.01935 (0.00982)	0.04768** (0.01385)
... \times True Positive	0.6382*** (0.10028)	0.1535** (0.05037)	0.4848*** (0.07703)
Gene-Disease FE	YES	YES	YES
Disease-Year FE	YES	YES	YES
Gene-Year FE	YES	YES	YES
N	137,084,183	137,084,183	137,084,183
N of Gene-Diseases	7,214,957	7,214,957	7,214,957
Mean of Dep Var:	0.13088	0.03741	0.09346

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination; the count is then divided between firms with and without previous publications on the gene. *Post \times GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS, excluding studies that report an above-median number of associations (i.e., larger than 55).

Table F.9: Results are robust to the exclusion of GWASs with an industry co-author or that acknowledge funding from a pharmaceutical firm

Dependent Variable:	USPTO patent applications by...					
	...all firms		...firms with gene knowledge		...firms w/out gene knowledge	
Post × GWAS	0.1263*** (0.02182)	0.0669*** (0.01856)	0.0319** (0.01146)	0.0148 (0.00979)	0.0944*** (0.01328)	0.0522*** (0.01165)
... × True Positive		0.4433*** (0.11117)		0.1278* (0.05833)		0.3155*** (0.06512)
Gene-Disease FE	YES	YES	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES	YES	YES
N	137,167,422	137,167,422	137,167,422	137,167,422	137,167,422	137,167,422
N of Gene-Diseases	7,219,338	7,219,338	7,219,338	7,219,338	7,219,338	7,219,338
Mean of Dep Var:	0.1308	0.1308	0.0373	0.0373	0.0934	0.0934

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination; the count is then divided between firms with and without previous publications on the gene. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS, excluding all GWASs with either a co-author working in a pharmaceutical company or acknowledging financial support from private firms. *True Positive*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease.

Table F.10: Investments in false positive GWAS findings seem not driven by strategic considerations, but follow patterns consistent with firms making mistakes

Dependent Variable	Strategic patent (0/1)	Continuation patent (0/1)	Litigated patent (0/1)	Expired patent (0/1)
Post × GWAS	0.0006 (0.00046)	0.0094*** (0.00139)	0.0022 (0.00119)	-0.0039*** (0.00072)
... × True Positive	0.0089*** (0.00172)	0.0306*** (0.00381)	0.0186*** (0.00340)	-0.0095*** (0.00236)
Gene-Disease FE	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES
N	137,254,556	137,254,556	137,254,556	137,254,556
N of Gene-Diseases	7,223,924	7,223,924	7,223,924	7,223,924
Mean of Dep Var:	0.0023	0.0339	0.0282	0.0100

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *Strategic patent*: 0/1=1 if the gene gene-disease combination received at least one *strategic patent* according to the definition by Abrams et al. (2018). *Continuation patent*: 0/1=1 if the gene gene-disease combination received at least one *continuation patent* according to the definition by Righi and Simcoe (2020). *Litigated patent*: 0/1=1 if the gene gene-disease combination received at least one patent that was subsequently involved in litigation. *Expired patent*: 0/1=1 if the gene gene-disease combination received at least one patent that was not subsequently renewed. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first GWAS, excluding all GWASs with either a co-author working in a pharmaceutical company or acknowledging financial support from private firms. *True Positive*: 0/1 = 1 for GWAS findings that are later replicated by another GWAS about the same disease.

Table F.11: No evidence of innovation spillovers on gene-disease pairs proximate to those treated by false positive GWAS associations

Dependent Variable:	Cit-weighted patents	Patent market value	Drugs (total)	Drugs (weighted)
Post × GWAS Spillover	0.2206 (0.21480)	-0.0361 (0.05840)	-0.000049 (0.00006)	-0.000005 (0.00002)
Gene-Disease FE	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES
N	136,913,221	136,913,221	136,913,221	136,913,221
N of Gene-Diseases	7,205,959	7,205,959	7,205,959	7,205,959
Mean of Dep Var:	7.0851	0.6229	0.0002	0.0001

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. *Post × GWAS Spillover*: 0/1 = 1 in all years after a neighboring gene-disease pair is treated by its first GWAS; neighbor pairs are defined as involving the same gene and a disease sharing the same 3-digit patent MeSH code than the ones in the GWAS. *Cit-weighted patents*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination, weighted by the number of patent citations received up to seven years after patent publication. *Patent value*: estimated stock market value (in constant USD) of patents granted to public firms using data from Kogan et al. (2017). *Drugs (total)*: number of molecules on a gene-disease combination entering the discovery stage. *Drugs (weighted)*: number of molecules on a gene-disease combination entering the discovery stage weighted by their scientific novelty (i.e., by the number of times that the same mechanism of action has been used before, following Dranove et al. 2022).

Table F.12: Firms with domain knowledge invest proportionally more on higher-quality gene-disease combinations

Dependent Variable:	USPTO patent applications by...					
	...firms with gene knowledge			...firms w/out gene knowledge		
Post × GWAS	0.0061 (0.01962)	0.0871*** (0.02619)	0.1334** (0.04119)	0.0919* (0.04219)	0.1793*** (0.03463)	0.1833** (0.05579)
Gene-Disease FE	YES	YES	YES	YES	YES	YES
Disease-Year FE	YES	YES	YES	YES	YES	YES
Gene-Year FE	YES	YES	YES	YES	YES	YES
N	588,069	9,550,920	1,091,265	588,069	9,550,920	1,091,265
N of Gene-Diseases	30,951	502,680	57,435	30,951	502,680	57,435
Mean of Dep Var:	0.0607	0.3359	0.5363	0.194	0.5944	0.8484
Sample:	<10 OT Score	10 <OT Score <90	>90 OT Score	<10 OT Score	10 <OT Score <90	>90 OT Score

Note: *, **,*** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease-year level. Std. err. clustered two-ways at the disease and gene level. The sample is limited to gene-disease pairs for which an Open Targets Score could be matched. Columns 1 and 4 include only gene-disease pairs with an Open Targets Score below the tenth percentile of the sample. Columns 2 and 5 include only gene-disease pairs with an Open Targets Score between the tenth and the ninetieth percentile of the sample. Columns 3 and 6 include only gene-disease pairs with an Open Targets Score above the ninetieth percentile of the sample. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination; the count is then divided between firms with and without previous publications on the gene. *Post × GWAS*: 0/1 = 1 in all years after a gene-disease pair is treated by its first true positive GWAS.

Table F.13: The results for false positives of Table 6 are robust to alternative definitions of the risk set of GWASs evaluated by a given firm

Dependent Variable:	USPTO patent applications				
Share of patent portfolio:	>0.005	>0.01	>0.02	>0.05	>0.1
Has Gene Knowledge (0/1)	-0.1155*** (0.01642)	-0.1082*** (0.01722)	-0.0895** (0.02690)	-0.1080*** (0.02089)	-0.1712*** (0.01818)
Firm-Disease FE	YES	YES	YES	YES	YES
N of firms	3765	3752	3666	2784	1562
N of obs	3,546,300	2,618,859	1,515,477	534,200	202,376
Mean of Dep Var:	0.0062	0.0056	0.0053	0.0043	0.0032

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the firm-gene-disease level. Std. err. clustered two-ways at the disease and gene level. Each column is estimated from a separate regression using a progressively more stringent risk set of GWASs evaluated, defined as those involving diseases constituting a certain share of a firm's patent portfolio. The sample is restricted to false-positive GWASs. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination. *Has Gene Knowledge (0/1)*: 0/1 = 1 if the firm has at least one publication about the specific gene before the GWAS associated it with a disease.

Table F.14: The results for true positives of Table 6 are robust to alternative definitions of the risk set of GWASs evaluated by a given firm

Dependent Variable:	USPTO patent applications				
Share of patent portfolio:	>0.005	>0.01	>0.02	>0.05	>0.1
Has Gene Knowledge (0/1)	0.2708*** (0.06344)	0.3336*** (0.06788)	0.4992*** (0.11716)	0.7730** (0.28434)	0.7808 (0.44003)
Firm-Disease FE	YES	YES	YES	YES	YES
N of firms	3,629	3,601	3,487	2,485	1,291
N of obs	1,159,770	867,620	514,864	192,353	75,732
Mean of Dep Var:	0.0247	0.024	0.025	0.0278	0.0032

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the firm-gene-disease level. Std. err. clustered two-ways at the disease and gene level. Each column is estimated from a separate regression using a progressively more stringent risk set of GWASs evaluated, defined as those involving diseases constituting a certain share of a firm's patent portfolio. The sample is restricted to true positive GWASs. *USPTO patent applications*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination. *Has Gene Knowledge (0/1)*: 0/1 = 1 if the firm has at least one publication about the specific gene before the GWAS associated it with a disease.

Table F.15: Descriptive evidence that firms delaying investments in a GWAS finding are less likely to choose false positives, but also less likely to obtain highly cited patents or drugs

Dependent Variable	False positive	Cit-weighted patents	Drugs (total)	Drugs (weighted)
Years Waited to First Patent	-0.0248*** (0.0028)	-1.7812*** (0.1510)	-0.0001 (0.0003)	-0.0005* (0.0002)
Firm FE	YES	YES	YES	YES
Disease FE	YES	YES	YES	YES
N of firms	225	225	225	225
N of obs	30615	30615	30615	30615
Mean of Dep Var:	0.5310	34.861	0.0081	0.0047

Note: *, **, *** denote significance at 5%, 1% and 0.1% level respectively. Observations at the gene-disease association level. Std. err. clustered two-ways at the disease and gene level. *Years Waited to First Patent*: count of years between the GWAS publication and the firm first patent application targeting the gene-disease pair. *False positive*: 0/1=1 if the gene-disease association is not subsequently replicated. *Cit-weighted patent*: count of USPTO patent applications filed in a given year that target a specific gene-disease combination, weighted by the number of patent citations received up to seven years after patent publication. *Drugs (total)*: number of molecules on a gene-disease combination entering the discovery stage. *Drugs (weighted)*: number of molecules on a gene-disease combination entering the discovery stage weighted by their scientific novelty (i.e., by the number of times that the same mechanism of action has been used before, following Dranove et al. 2022).