

R&D Alliances in the age of Artificial Intelligence

Bowen Lou

School of Business
University of Connecticut
bowen.lou@uconn.edu

Evan Rawley

School of Business
University of Connecticut
evan.rawley@uconn.edu

Abstract

Firms have long relied on research & development (R&D) alliances as an alternative to internally produced R&D. In this paper, we examine how artificial intelligence (AI) influences the frequency and productivity of R&D alliances. Our theory advances the idea that by matching discoveries to disease targets more efficiently, AI shifts the boundary of the firm toward R&D alliances. Using a rich dataset on the biopharmaceutical industry from 2010 to 2019, we show that firms with greater AI resources form 9% more R&D alliances and generate 8% additional new drugs per alliance. Moreover, we observe a key mechanism behind the results: AI is particularly useful in exploiting information held by counterparties within an alliance. The study sheds light on how artificial intelligence changes the innovation production function and shapes the economic organization of firms.

Introduction

Strategic alliances—ongoing cooperative relationships between firms—represent an important alternative to internal research and development (R&D) (e.g., Pisano, 1990). More recently, advances in artificial intelligence (AI)—the use of computer science techniques to substitute for human judgement—have promised to improve R&D efforts by directing scientists toward promising pathways for innovation (Cockburn et. al., 2018). Given the clear implications of AI for both the theory and practice, this paper studies the impact of AI on alliance formation and productivity. Building on the idea that transactions are efficiently aligned with governance choice (Williamson’s (1991), we examine how AI shifts firm boundaries by changing the relative efficiency of internal R&D versus R&D alliances at different levels of asset specificity.

Empirically, we study the biopharmaceutical industry, which is characterized by widespread R&D alliances and the rapid adoption of AI (Aggarwal and Hsu, 2009). Within the industry, it is common for large integrated pharmaceutical firms to pursue R&D alliances to access to high-potential drug candidates, or “novel chemical entities” (NCEs), identified by small, focused biotechnology firms. To help transform an NCE into a cure, pharmaceutical firms use machine learning—AI software that improves the predictive power of statistical models by continuously incorporating new data—to screen NCEs and match them to potential disease targets (e.g., Fleming, 2018). As matching improves, the path toward development becomes clearer, facilitating more targeted knowledge flows between scientists.

By directing knowledge flows more efficiently between discovery and development, AI boosts productivity. Interestingly, though, AI is most effective at matching moderately novel NCEs—drug candidates unique enough that their optimal set of targets may not be fully understood by scientists, and yet are characterized by enough structured data such that machine

learning algorithms can be trained to “understand” how they fit with potential targets (Lou and Wu, 2021).

Meanwhile, R&D alliances are particularly well-suited for governing transactions of an intermediate level of asset specificity, meaning that the assets specific to the transaction would be reduced in value by an intermediate amount in their next best application (Oxley, 1997). Thus, if novelty is a component of asset specificity, as we propose, then AI should be expected to uniquely benefit R&D alliance governance.

The main results show that firms with greater AI resources form 9% more R&D alliances and generate 8% more drugs per alliance. The baseline results are robust to standard econometric techniques for adjusting for endogeneity bias. We also demonstrate that firms with greater AI resources are better at exploiting information embedded in their alliance relationships, suggesting that improved knowledge flows are, indeed, a key mechanism behind the effects. Moreover, we provide evidence that machine learning drives the productivity result, supporting the conclusion that it is the unique analytical strength of AI, relative to general purpose IT, or other aspects of AI, that is at the heart of the AI effect on R&D alliances.

By explicating how AI influences the extensive margin (i.e., frequency) of R&D alliances, this paper contributes to the literature on contract theory. Scholars have long been curious about how characteristics of transactions and resources influence the relative efficiency of discrete structural alternatives (e.g., Pisano, 1990; Oxley, 1997; Hitt, 1999). We show that increasing a firm’s AI resources shifts the boundary of the firm by increasing the efficiency of R&D alliances. Furthermore, by documenting an AI and R&D complementarity on the extensive margin (i.e., productivity), we contribute to the literature on complementarities in organizational design (e.g., Rivkin and Siggelkow, 2003). In this view, a firm is more than a nexus of contracts, but rather a

latticework of interconnections amongst activities, resources, and contracts. We show that AI and R&D alliances represent an important organizational design complementarity. Finally, by providing evidence that AI improves knowledge flows within alliances, we unpack one mechanism for how firms use AI to gain competitive advantage; in doing so, we also contribute to the nascent literature on how AI influences firm strategy and organization (e.g., Furman and Seamans, 2019). After years of promise, AI is now having a profound effect on firms; we show that one salubrious effect is the facilitation of inter-organizational learning. In sum, the paper sheds new light on how artificial intelligence is reshaping organizations and improving the innovation production function, with implications for scholars and practitioners alike.

Theory and Hypothesis

In transaction cost economics, strategic alliances govern transactions of intermediate asset specificity, relative to market-based exchange and vertical integration (Williamson, 1991). In other words, alliances are useful when arms-length spot-market exchange fails in the sense that it does not support the development of assets that are uniquely suited to facilitate the transaction; and vertical integration fails in the sense that it would burden the transaction with incentive problems. As a result, market-based exchange predominates when asset specificity is low, and vertical integration predominates when asset specificity is high, while alliances govern transactions at mid-levels of asset specificity.

Alliances are not a panacea though. Free-riding and hold-up problems are frequent concerns within alliances, resulting in *ex post* adjustments that ameliorate, but do not eliminate, the potential for hold-up problems (Dyer 1997; Reuer, et. al., 2002).

In a typical R&D alliance between a large pharmaceutical company and a biotechnology company, the former brings a library of disease targets to the exchange, while the latter bring a

novel chemical entity (NCE; i.e., a drug candidate). Together they work to match the NCE to a disease target. It is not uncommon for the matching process to take multiple years and cost hundreds of millions of dollars (Hughes, et. al., 2011). The magnitude of the specific assets required for such an endeavor are so large that arm's length spot-market exchange is unusual in drug development. An alliance can provide more protection for specific assets (Oxley, 1997)—for example, investments in developing unique knowledge and customized proteomic, chemogenomic, and bioinformatics algorithms. But, if one side exits an alliance prematurely, or otherwise fails to perform as expected, specific investments made in the matching process would be lost. As a result, R&D alliances are limited to governing transactions with moderate levels of asset specificity (Williamson 1991).

Interestingly, in the biopharmaceutical industry, the novelty of a potential cure is correlated with asset specificity in the matching process. Matching novel drugs to disease targets, represents a development challenge that often necessitates unique scientific resources, including massive investments in data analysis (e.g., Yang, et. al., 2009). By contrast drugs that are the result of a small chemical change to an existing drug, will typically require less specific human and non-human assets. Scientists working with a well-known chemical agent can usually identify a small set of appropriate targets using their own prior knowledge without making large software investments. Therefore, novelty is a good proxy for asset specificity in drug development.

By improving predictive search, AI offers scientists a powerful analytical tool for analyzing large structured data sets (e.g., Agrawal et al. 2018). For example, machine learning has been used to infer the interplay amongst numerous biological entities such as genes, proteins, diseases and drug candidate (Cockburn et al. 2018). In that sense AI delivers a positive R&D productivity shock. But, highly novel NCEs are challenging for AI, as there is limited data for a machine

learning algorithm to exploit in matching the NCE to a potential disease target. At the other end of the novelty spectrum, drugs that are very similar to other known drugs, are relatively simple for scientists to match to drug targets without the assistance of an algorithm. Thus, AI represents a high-potential resource for improving R&D efficiency for projects with an intermediate level of chemical novelty (Lou and Wu 2021)—exactly the region where transaction cost economics predicts alliances should flourish.

To demonstrate more formally how AI should influence the frequency of R&D alliances, we begin from a Williamson’s (1991) equilibrium governance model where firms choose contracting regimes that minimize the cost of governing a particular transaction, as function of asset specificity, (see the left-hand panel of Figure 1). Next, we consider AI as a shifter of the efficiency of different contractual regimes. If we assume that pure market-based exchange is inconsequential in the biopharmaceutical industry, or that AI has no effect on market-based low-asset specificity exchange (as in the figure), it is easy to see visually how AI should influence frequency of R&D alliances.

One straightforward benchmark is the case where AI improves the efficiency of R&D alliances and vertical integration by the same proportion at each level of asset specificity—that is it acts as a productivity multiplier (see the right-hand panel of Figure 1). In that case, AI unambiguously leads to an increase in the frequency of R&D alliances.¹ As both vertical integration and alliances become increasingly more efficient per unit of asset specificity, alliances replace vertical integration at higher (but still intermediate) levels of asset specificity.

However, as noted above, we know that AI is particularly useful for facilitating knowledge flows in pursuit of matches between targets and discoveries at intermediate levels of novelty and,

¹ In the more general case where a resource shock influences all three forms of contracting proportionately, the result holds when alliances represent a “meaningful” proportion of the transactions *ex ante*.

therefore, intermediate levels of asset specificity. While non-linear of effects of AI on governance costs with respect to asset specificity could shift the equilibrium in either direction, it seems likely that the effect of AI would be to push down the efficient contracting frontier for R&D alliances even further. For example, one might expect that AI would uniquely improve R&D alliances transactions within the pre-AI range. If AI was also uniquely valuable for transactions just beyond the pre-AI border between firms and alliances, then the simple benchmark case would understate the positive effect on R&D alliances. Thus, our first hypothesis is:

(H1): Higher levels of AI resources leads to more R&D alliances.

While Hypothesis 1 implies that AI uniquely benefits R&D alliances, it could also obtain if AI simply reduced governance costs equally across the optimal contracting frontier. To wit, firms might increase both internal R&D and R&D alliances at the same time in response to investments in AI. While such behavior would be consistent with Hypothesis 1, our theory makes a stronger claim. Therefore, to further probe the idea that there is a special relationship between AI and R&D alliances, we examine whether R&D alliances are uniquely more productive as AI resources grow.

In the context of the biopharmaceutical industry, AI enables firms with deep knowledge of disease targets to use information about NCEs more efficiently. Within firms, AI improves knowledge flows between discovery and development; between firms, AI also improves knowledge flows across firm boundaries. Our theory predicts a uniquely powerful benefit of AI on cross-firm knowledge flows, because of the nature of AI and the transactions R&D alliances are used govern. R&D alliances facilitate joint investments in projects of intermediate novelty;

and AI improves the analytical power of scientists engaged in projects of intermediate novelty. Therefore, our theory predicts that AI resources and R&D alliances exhibit complementarities, in the sense of Milgrom and Roberts (1990) and Cassiman and Veugelers (2006), and summarized by equation (1):

$$(1) \text{ Innovation (R\&D alliances, AI)} > \text{Innovation (R\&D alliances)} + \text{Innovation (AI)}.$$

In words, we propose Hypothesis 2:

(H2): Firms are more innovative when they jointly increase their AI resources and R&D alliances, relative to the sum of the innovation benefits they would achieve by adopting each separately.

Hypothesis 2 implies Hypothesis 1. *Ceteris paribus*, if R&D alliances become relatively more efficient, we should observe more of them. But, Hypothesis 2 is interesting in its own right too. A long literature finds that R&D alliances are innovation enhancing in equilibrium (e.g., Powell 1998; Rosenkopf and Almeida 2003), and recent work demonstrates that AI improves R&D productivity (Lou and Wu 2021). Our theory goes one step further, advancing the idea that AI and R&D alliances are uniquely valuable when used together.

Institutional Context

The biopharmaceutical value chain consists of several phases. The first is basic research where novel drug entities are tested against certain biological targets in a laboratory setting. The next phases are stages of drug development. The first step in development is non-clinical (testing on

non-human animals). The next three phases involve human clinical trials of increasingly larger sizes. Finally, a successful drug gains FDA approval, is marketed and sold. Our study primarily focuses on the drug discovery through the non-clinical development stage, as that is where the impact of AI and R&D alliances on industry innovation are thought to be the greatest (e.g., Savage (2021)).²

The biopharmaceutical industry is an excellent laboratory to test our theory for two main reasons. First, the industry relies on a wide range of interorganizational knowledge-intensive activities to foster innovation. Over the past decades, biopharmaceutical firms have entered into systematic collaborative linkages with universities, and other research-intensive firms, which provide complementary assets necessary for innovation (Arora and Gambardella 1994). As alliances have proven to be a potential source of competitive advantage for some firms, the scope and scale of interorganizational collaborations and alliances in this sector have grown (Powell et al. 1996). As a result, successful alliances typically require extensive knowledge transference between firms (Oxley and Sampson 2004; Sampson 2007). Given the rich linkages amongst firms, markets and innovation, the biopharmaceutical industry makes a good laboratory for studying firm boundaries and R&D productivity.

Second, as the knowledge needed for discovering drug candidates is becoming more codifiable, through the process of digitization, the importance of AI has grown within the industry. As noted above, AI is most effective with large structured sources (Agrawal et al. 2018), where it can be used to find subtle linkages amongst data elements to make useful predictions about

² In the pharmaceutical industry, one of the most powerful applications of AI is to help screen potential drug candidates at the early stages of drug development (Cockburn et al. 2018). Thus, we focus primarily on the pre-clinical effect of AI (i.e., before clinical trials) where we expect the signal to noise ratio of the effect of AI to be strongest. (We do undertake some exploratory analyses on the effect of AI on later stages of drug development in Table 9).

potential innovation pathways (Wu et al. 2020). Hence, the potential drug candidates screened using AI serve as a useful starting point to improve efficiency in the drug innovation process.

In sum, we study a context where innovation and knowledge management are crucial, and strategic alliances, and AI are known to independently influence innovation. More importantly for the tests of our theory, it is also plausible that AI and strategic alliances might complement one another in the biopharmaceutical industry.

Data and measures

We study the biopharmaceutical industry from 2010 to 2019, a period that witnessed rapid growth in the application of AI, as well as the proliferation of alliance activities. A drug developed through to the end of the pre-clinical phases, represents a meaningful innovative step, as such a drug has proven scientific potential. Therefore, the dependent variable in the productivity tests is *Drugs developed*, measured as the cumulative number of drug candidates developed by a firm through the pre-clinical phase, depreciated at 15% per year, as is common practice in the innovation literature (Hall 1990; Jaffe 1986). The dependent variable in tests of the extensive margin is the number of alliances developed in a particular period (e.g., a calendar year).

Data on *R&D Alliances* are collected from the Biosci database (BiosciDB). The database tracks interfirm R&D contracts, providing the most comprehensive records of biopharma R&D alliances since the early 1980s.³ The alliance data classifies firms as “R&D” companies, firms that are primarily responsible for basic research, or “client” companies, firms that are accessing knowledge from their collaborators in the pursuit of drug development (Aggarwal and Hsu 2009; Lee et al. 2015). The dichotomy has been well defined in both academic and practical analysis of

³ Compared to some alternative alliance databases such as Securities Data Company (SDC) Platinum database on Joint Ventures and Alliances, the BiosciDB database offers wider coverage of deals in the biotechnology and pharmaceutical sector.

the R&D partnerships in the biotechnology and pharmaceutical sector (e.g., Robinson and Stuart (2007a); Robinson and Stuart (2007b)). Our analysis focuses on client firms (i.e., mostly integrated pharmaceutical firms) in their alliances with biotechnology companies. There are 2,755 R&D alliances in the data. We cumulate R&D alliances, at the firm level, to capture a firm's overall exposure to R&D alliances.

Data on drugs developed are collected from two sources that have been used extensively to study biopharmaceutical innovation (e.g., Sosa (2013) and Krieger et al. (2022)): the Informa Pharmaprojects database, and the investigational drug database from Clarivate Analytics databases. The data are very rich, but they do have one important limitation: we cannot map alliances to specific drugs. Therefore, we study changes in a firm's stock of drugs developed to capture the average complementarity effect.⁴

The data do map drugs to owners, however, because ownership of a drug can be changed through mergers and acquisitions (M&A), we use SDC Platinum and Zephyr data from the Bureau van Dijk to properly match drugs to the firms responsible for their development. After cleaning the ownership field, we have 614 unique firms that developed at least one drug through the pre-clinical phase between 2010 and 2019.

Our key explanatory variable is *Artificial Intelligence (AI) resources*, which is measured as a function of the count of a firm's AI-related patents (*AI IP*) and a count of a firm's AI-related workers (*AI skills*). *AI IP* is an intangible asset-based measure of AI resources (Webb 2019). To measure a firm's AI IP, we use global patents from a worldwide patent statistical database,

⁴ We prefer using the stock-based drug innovation outcome as a dependent variable in our estimation. The instantaneous number of drugs developed may miss "true" effects, as we don't know the precise timing of alliances and AI usage for developing specific drugs. Even though, we also examine the number of drugs developed by a firm in each year and obtain consistent estimation results.

PATSTAT, created by the European Patent Office (EPO).⁵ PATSTAT offers bibliographical data for over 100 million patents from 90 patent-issuing authorities from both leading industrialized and developing countries going back as far as the nineteenth century. It provides patent records with rich information that contains patent application year, patent technology classes, as defined by the International Patent Classification (IPC) or Cooperative Patent Classification (CPC) systems, citations, title, abstract, and legal entities (e.g., firms or any organizations) filing the patent application. Patent owners are identified using a standardization of the original name of the patentees. We use the patent filing year, as opposed to the publication year, because it more closely approximates the time when the firm produced the innovation described in the patent (Griliches et al. 1986; Hall et al. 2005).

To create a robust measure of a firm's *AI IP*, we take four steps to identify whether a patent uses or contributes to AI technologies. We start with the most obvious AI patents: USPTO class 706 for “Data Processing – Artificial Intelligence” consists of a large set of subclasses including “machine learning” and “neural network”. To link USPTO class 706 to the PATSTAT data, we apply the IPC or CPC concordance to obtain the classification code that is used for classifying global AI patents.

While class 706 explicitly captures patents at advance the knowledge frontier of AI, the narrow scope of the USPTO AI classification code may exclude patents employing AI to solve problems in an innovative way, which do not contribute to fundamental advances in AI. Thus, our

⁵ We leverage the PATSTAT data in the version 2017b to cover patent records up to 2017. As Google launched its public datasets of worldwide patents on BigQuery in 2017 (<https://cloud.google.com/blog/products/gcp/google-patents-public-datasets-connecting-public-paid-and-private-patent-data>), we further augment our patent data to cover the recent time period from 2017 to 2019 by using global patents through Google Patents Public Datasets. We adopt similar approaches to match them to other datasets used in our empirical study.

second step expands the scope of AI patents to include content-based information available in the patent documents.

Our third step is to include patents with content that fall within the well-accepted Association of Computing Machinery (ACM) Computing Classification System (CCS), which accounts for dynamic changes in AI technologies over time (WIPO 2019). The ACM/CCS methodology has been used for over 50 years to organize technological concepts and trends. Specifically, we use three major hierarchies in CCS that identify AI-related phrases, including: (i) the “artificial intelligence” hierarchy, comprised of an AI functional application, such as natural language processing, computer vision, knowledge representation and reasoning, and AI techniques used to realize those functions; (ii) a “machine learning” hierarchy, which includes numerous learning-based AI techniques; and (iii) a “life and medical sciences” hierarchy under the “applied computing” category that captures activities concerning intelligent computing for producing medicines.

Our fourth step is to expand the vocabulary list from ACM/CCS to contain phrases related to a variety of AI technologies specifically used for drug innovation in more recent years. Specifically, we include well-known off-the-shelf AI tools and systems such as PyTorch and TensorFlow (Raymond et al. 2019). Overall, we construct a comprehensive list of validated words and phrases pertaining to AI and search for these terms in both the titles and abstracts of our patents (Lou and Wu 2021) (See the Appendix for keywords used). We also distinguish machine learning from other types of AI resources (Cockburn et al. 2018; WIPO 2019). This helps us cut down our measurement of AI resources into finer elements to capture important differences regarding how AI resources affect innovation.

Based on the four steps described above, we track the development of *AI IP* over time, as the count of all AI patents, depreciating the value at 15% per year.⁶ The resulting AI stock variable aligns with the spirit of classic innovation production functions that model innovation as a function of the existing repertoire of knowledge and research resources dedicated to the production of new knowledge (Romer 1990).

Demand for personnel with AI skills represents a human capital-based measure of AI resources, as it captures the personnel-based resources firms seek to create, implement, and deploy in AI-related areas within the firm. To measure *AI skills*, we use job titles and skill requirements data from a company that collects the data from over 40,000 online job boards worldwide (Alekseeva et al. 2021). Following our systematic approach to classifying AI patents, we search for AI-related words in each job posting for each firm in our sample. Additionally, job title classifications from O*NET are also applied to identify AI-related positions, similar to the way information technology and analytics labor has been identified in prior research studies (e.g., Tambe and Hitt (2012)). Once individual AI job postings are identified, we aggregate *AI skills* to the firm level as a cumulative count.

Although *AI IP* and *AI skills* are independent measures of a firm's observable AI resources, they are almost surely jointly responsible for a firm's AI resources. For example, *AI IP* may proxy for whether the firm has: (1) tangible AI resources, such as the technological infrastructure required to develop an AI-related patent, or (2) intangible AI resources, such as the practices required to use AI for innovation (Lou and Wu 2021), or both. And, a firm's AI human capital, proxied by AI skills, are almost surely involved with the development of both tangible and intangible AI resources. Thus, we aggregate *AI IP* and *AI skills* into a single composite AI

⁶ Using depreciation rate at 15% per year is a common practice in the innovation literature (Hall 1990; Jaffe 1986). Results are directionally consistent when different depreciation rates such as 10% and 20% are used.

resources measure for our main tests, by standardizing and summing each AI-related dimension, as shown in equation (2):

$$(2) \text{ AI resources} = \text{norm}(\text{norm}(\text{AI IP}) + \text{norm}(\text{AI skills})).$$

The result is a normalized measure of a firm's AI resources that jointly considers both its asset-based and human capital-based AI resources. There are 185 firms in our sample with positive AI resources in at least one year. Because the normalization obscures some potentially interesting granular effects of the components of AI resources, we verify that the results are similar in terms of sign, significance, and approximate economic magnitude if we allow *AI IP* and *AI skills* to enter separately.

Control Variables: Three databases—Crunchbase, PitchBook and Bureau van Dijk Orbis—are used to retrieve other firm characteristics such as firm ownership status, firm age, number of employees and R&D expenses. The databases provide a wide coverage of both public and private bio-pharmaceutical firms. We account for a firm's financial ownership status over years (variable *Public Company* = 1 when the firm is publicly held, otherwise *Public Company* = 0), which has been shown to influence organizational innovation outcomes (Aggarwal and Hsu 2013). Additionally, we control for firm age, which is the difference between the observation year and the founding year of each firm in our sample. The oldest firm, Merck, is 347 years old, but the mean age is 23.81 years old, as there are many young biotechnology companies in the sample as well. We also account for each firm's size (by number of employees) and R&D expenditures, as well as the count of a firm's overall number of patents, depreciated at 15% per year.⁷

⁷ Using depreciation rate at 15% per year is a common practice in the innovation literature (Hall 1990; Jaffe 1986). Results are directionally consistent when different depreciation rates such as 10% and 20% are used.

Table 1 displays the summary statistics and bivariate correlations amongst the variables. The mean number of drugs developed is 15.33 per firm-year. The mean values of *AI IP* and *AI skills* are 1.22 and 16.12, respectively, while the mean number of alliances is 0.61. As one might expect, firms spend enormous sums on R&D—the mean value is \$3.56 billion per firm-year. In the regressions, we take the natural logarithm of a firm’s number of drugs, alliances, AI IP and skills, and all the non-categorical controls to account for their skewness, and to facilitate interpretation of the results as elasticities.

Empirical Strategy

In the ideal experiment we would randomly assign AI resources and evaluate how they altered a firm’s propensity to engage in R&D alliances. However, our analysis must instead deal with a non-random data generating process where firms choose their level of AI resources. The endogenous selection by firms into AI represents a major challenge to a causal interpretation of the results. For example, if better (worse) firms devote more resources to AI, “naïve” estimates of the effect of AI will be upward (downward) biased.

To address the most fundamental endogeneity concerns, we first undertake estimates using firm and time fixed effects that control for all time invariant heterogeneity across firms, including firm quality differences, as well as any secular time trends that influences all firms equally. Firm fixed effects also eliminate any differences in initial stocks of drugs developed, AI resources, and R&D alliances, allowing for an apples-to-apples comparison of the effect of *changes* in AI resources on *changes* in R&D alliances.

To deal with more subtle potential endogeneity effects between AI and R&D alliances, we deploy two-stage least squared (2SLS) instrumental variables analyses. Our 2SLS approach uses the total number of neighboring firms, defined as firms that cite each other’s patents, that have AI

resources, as an instrument for the focal firm's AI resources. The rationale for the instrument is that if many firms, working in the same scientific area, feel the need for AI resources to make progress, then the focal firm is partially induced to increase their AI resources, relative to firms working in scientific areas where AI is less common. If competition essentially forces a focal firm to increase their level of AI resources, a portion of the change in AI resources ceases to be a choice variable, allowing us to estimate the causal effect of AI resources using a 2SLS estimator.

To be effective and valid an instrument should be strongly correlated with the endogenous regressor (i.e., AI resources), and should solve the exclusion restriction—it should not impact a focal firm's outcomes directly. Still, unless rival firms' AI resources degrades a focal firm's innovativeness, by outpacing them in the race for specific innovations—an assumption we test directly below—it seems reasonable to assume that other firm's AI decisions would not directly influence the innovativeness of a focal firm.

Using similar firms' characteristics as instruments is a common econometric technique (e.g., Nevo, 2001; Campa and Kedia, 2002), suggesting that our 2SLS tests on the extensive margin represent a sensible solution to addressing endogeneity between AI and R&D alliances. However, in the productivity regressions, *R&D alliances* is an explanatory variable, meaning that we also need to consider how endogenous sorting into R&D alliances might bias our estimates of a complementarity between AI and R&D alliances. Unfortunately, we could not identify an instrument for R&D alliances that both met the exclusion restriction and was strong in the first stage. Therefore, we rely on a propensity score matching estimator to deal with endogenous selection into different levels of alliance formation. Propensity score estimators match firms based on their pre-treatment characteristics to ensure that one does not conflate treatment effects with observable heterogeneous pre-trends. In the absence of systematic unobservable heterogeneity that

is correlated with both the treatment (i.e., R&D alliances) and the outcome of interest, propensity score estimators can be considered “as good as” average treatment effects (Rosenbaum and Rubin 1983).⁸

Main specifications

Hypothesis 1 predicts that firms with greater AI resources should pursue more R&D alliances. Thus, we expect $\beta_l > 0$ in specification (3), defined at the firm-year level, for firm i in year t :

$$(3) \text{ Alliances}_{it} = \beta_0 + \beta_l \text{ AI resources}_{it} + X_c \beta_c + T_t + \gamma_i + \varepsilon_{it}.$$

We implement (3) with and without the firm fixed effects, γ , for periods of a single calendar year, and where *Alliances* are treated as a cumulative stock measure. A full set of year fixed effects T_t are included in all variations of (3). Subscript c indexes the vector of coefficients on time-varying firm-specific controls, including patent stock, ownership status, firm age, workforce size and R&D spending. Standard errors are clustered by firm. After first considering (3), we then implement our 2SLS estimator to address the potential endogeneity of AI resources.

Hypothesis 2 predicts an innovation complementarity between AI resources and R&D alliances, on the intensive margin—each alliance is more productive when consummated by a firm with greater AI resources. Thus, we use a fixed-effect estimator that captures the relationship between changes in R&D alliances, in conjunction with changes in AI resources, and changes in innovation, as in (4):

⁸ Other matching and weighting approaches, such as coarsened exact matching (CEM) and inverse probability weighting (IPW), yield consistent results.

$$(4) \text{ Drugs developed}_{it} = \beta_0 + \beta_1 \text{ AI resources}_{it} + \beta_2 \text{ Alliances}_{it} \\ + \beta_3 \text{ AI resources}_{it} * \text{ Alliances}_{it} + X_c \beta_c + T_t + \gamma_i + \varepsilon_{it}.$$

In (4), the coefficient β_1 represents the main effect of increasing a firm's AI, β_2 captures the main effect of increasing the number of R&D alliances, and β_3 , the coefficient on the interaction term between AI and R&D alliances, is the key coefficient of interest.

To address the potential for endogenous sorting into R&D alliances in the productivity tests we first match firms that increase their number of alliances to firms that do not, contemporaneously (i.e., in any given year), and then use a 2SLS estimator. For the matching, we use a first stage probit predicting whether a firm forms an alliance in a particular year, based on its *ex ante* characteristics. The probit corresponds to (3), except for the limited dependent variable specification ensures the predicted values are in the [0,1] interval, and the regressors are lagged by one year to address potential simultaneity bias. We then use the predicted value from the probit—the probability of being treated, or the “propensity score”—to match control and treatment groups firms, one to one, using nearest neighbor matching without replacement. The resulting matched sample allows us to compare firms that are equivalent on all observable characteristics, at any given point in time, except for their decision to increase the number of alliances in that year, or not. While propensity score matching cannot control for unobservable differences between firms, given our baseline fixed-effect estimation strategy, it would appear to provide a reasonable basis for inferring average treatment effects.

Results

We begin by visualizing a preview of our main finding in Figure 2. The graphic reveals that firms involved in above the median level of R&D alliances are more innovative, but the most innovative

firms also have high levels of AI resources. The figure provides model-free evidence consistent with the idea that R&D alliances and AI resources are complements. Whether the relationship is casual or not is the subject of much of the ensuing analysis.

In Table 2, we study the extensive innovation margin, whether AI resources influence the number of R&D alliances. For presentation purposes, all coefficients and standard errors are multiplied by 100, in Table 2 and all subsequent tables, so one can interpret the coefficient estimates as elasticities in percentages.

First, we regress the categorical variable *Alliance in year t*, which is set equal to one when an alliance is formed in a particular year *t*, and zero otherwise, on AI resources, with and without firm fixed effects. The estimation results suggest that firms with greater AI resources are more likely to enter an alliance in any given year (Column 1), even after adjusting for firm-specific heterogeneity (Column 2). Limited dependent variable models, probit and logit regressions, yield consistent results.

We also implement a fixed-effect estimator with the cumulative number of alliances as the dependent variable (Column 3). On average, a one-standard-deviation increase in a firm's AI resources is associated with a, precisely estimated, 6.68% increase in number of alliances. Although the evidence “merely” demonstrates a treatment on the treated effect on the extensive margin—firms expand the number of R&D alliances after increasing their AI resources—it still provides support for the idea that R&D alliances and AI resources are complements.

In Table 3 we show the first and second stages of the 2SLS estimators for each of the specifications Table 2. The 2SLS point estimates are substantially larger and noisier than the OLS/fixed-effect estimates, but the collective evidence points toward a causal effect of AI resources on R&D alliance formation.

Next, we study R&D productivity. To address concerns that R&D alliances are endogenous in the R&D production function, we first match firms with a new alliance to firms with no new alliances in year t , on all the observable, lagged one-year, *ex ante* characteristics of the firms available to us, using propensity score matching. Table 4 reports statistics of the observable characteristics before and after the matching. Before matching, the two groups are statistically different along every dimension. After matching, the two groups of firms have similar distributions, suggesting that selection bias is meaningfully reduced. Figure 3 shows the convergence of the propensity scores distributions graphically. We use the matched sample as the basis for our 2SLS “treatment effect” regressions.

In Table 5, we study the joint impact of AI and R&D alliances on innovation performance. The elasticity of the main effect of R&D alliances is 6.73%, the elasticity on AI resources is 6.77% and the elasticity of their interaction is 4.17%. The preliminary economic interpretation is that doubling either AI resources or R&D alliances independently, is only about 76% as effective as doubling them together.⁹

The 2SLS estimator instruments for AI resources, using the number of neighboring firms with AI resources as the instrument, within the matched R&D alliances sample. Since there are two endogenous variables remaining (i.e., after matching), *AI resources* and *AI resources x R&D alliances*, we include the main effect of the instrument and the instrument interacted with (the matched pairs of) R&D alliances. Thus, there are two first stage results to consider. The first shows that neighboring firms’ AI resources is a strong predictor of a focal firm’s AI resources.

⁹ Doubling the number of R&D alliances independently generates 6.73% more drugs, while doubling AI resources independently generates 6.77% more drugs. The sum of the two independent effects is 13.50%. Doubling R&D alliances while doubling AI resources simultaneously, generates 4.17% additional drugs. $4.17\% + 13.50\% = 17.67\%$. $13.50\% / 17.67\% = 76\%$.

Since our main concern with the validity of the exclusion restriction is that other firm's AI resources would *reduce* the focal firm's innovativeness (i.e., by allowing a rival to beat them in the race to develop a new potential drug), the first stage result appears to support the validity of the instrument. The second set of first stage results reveals that the interacted instrument is also a strong predictor of the endogenous interaction term *AI resources x R&D alliances*. While the negative coefficient on the main effect of the instrument is indicative of collinearity between the two instruments, the F-statistic on the joint significance of the instruments from the first stage is 27.85, suggesting that the instruments are jointly strong enough to generate meaningful statistical power in the second stage. The second stage results shows that the point estimate of the 2SLS estimator is 7.51%, which is almost twice as large as the estimator from our fixed-effect estimation, though noisier. The interpretation is that there is a causal complementarity between AI resources and R&D alliances.

Mechanisms, extensions, robustness

Our theory predicts a complementarity between AI resources and R&D alliances because R&D alliances provide firms with new data, while AI resources allow firms to utilize advanced software systems to improve data analysis, facilitating the creation of new knowledge. In theory, alliances serve as conduits through which information flows to a focal firm from its partner.

To understand whether knowledge flows truly undergird the innovation complementarity, we study the 2,755 alliances in our data set. Following recent work on patents and alliances, we use patent citation data to capture the changes in the citation patterns of focal firms before and after they form alliances to trace knowledge flows from its partners (Almeida et al. 2002; Oxley and Wada 2009).

To measure knowledge flows, we draw on citation patterns in an alliance firm’s patent portfolio to proxy for the changes in the relationship of its knowledge portfolio to that of its allied firm (Mowery et al. 1996). Specifically, we calculate the log difference between a focal client’s pre- and post-alliance citations to its partner R&D firm, for the two years before and after the alliance is consummated.¹⁰ The variable has a mean of 0.44, a range of zero to 678, and a standard deviation of 13.65.

We then regress knowledge flows from an R&D partner on a focal firm’s AI resources, for each alliance formed by a focal “client” pharmaceutical firm i , beginning in year t , controlling for i ’s other characteristics, as in the firm-level analysis. Specifically, the mechanism test specification is (5):

$$(5) \text{ Knowledge flows}_{it} = \beta_0 + \beta_1 \text{AI resources}_{it} + X_c \beta_c + T_t + \gamma_i + \varepsilon_{it}.$$

The test examines the extent to which knowledge can be leveraged by client firms (typically vertically integrated traditional pharmaceutical or large “platform” biotech companies) via their R&D partners (typically small biotechnology companies). We find that, on average, a one-standard-deviation increase in AI resources is associated with about an 8% increase in its knowledge inflows, as captured by the difference in post-alliance and pre-alliance patent citations to its partners (see Table 6). The 2SLS version of (5) returns a point estimate that is about twice the fixed-effect estimator, but also has a relatively large standard error. The interpretation is that AI resources cause knowledge flows to increase within R&D alliances, as predicted by the theory.

¹⁰ Using other year lags such as one-year and three-year lags produces similar results.

To see whether AI is a unique technological resource, we compare AI resources with other types of information technology (IT) resources of the firms (see Table 7).¹¹ Relative to AI resources, a firm's IT resources, and their interaction with alliances, have no significant effect on the number of drugs developed for preclinical studies. While IT is still important to support drug innovation, it has not delivered a competitive edge in recent years, possibly because most firms have already invested in traditional IT, and earlier best IT practices may have already diffused throughout the industry (Chae et al. 2014; Lou and Wu 2021).

Examining different categories of AI resources, we find that it's machine learning in particular that facilitates exploiting information held by counterparties within an alliance. Specifically, we follow Cockburn et al. (2018), who define interrelated but separate technological subfields, such as symbolic systems, and machine learning, within AI to characterize the evolution of achievements in AI. Symbolic systems include methods that can replicate the logical flow of human decision making through processing human-readable symbols. They develop applications such as expert systems that aim to emulate decision-making process of a human expert by reasoning through a certain set of heuristics and logic rules (mainly represented as if-then rules). Such systems suffer from the constraints that they are labor-intensive, rigid and difficult to scale when dealing with more complex innovation tasks and dynamic innovation environment (Harrer

¹¹ Similar to how we measure AI resources, we use IT patents (*IT IP*), and job postings that require IT skills (*IT skills*) to measure IT resources. *IT IP* is identified using Category 2 in patent classifications that include computer hardware and software, communications, computer peripherals, and information storage (Hall et al. 2001). To link the USPTO classes to the PATSTAT data, we apply the IPC or CPC concordance to obtain the classification code that is used for classifying global IT patents. We also use multiple alternative methods from Forman et al. (2016), such as incorporating electronics-related patents about electrical and semiconductor devices identified from Category 4 in Hall et al. (2001) and searching IT-related phrases on the titles and abstracts of patents. These approaches yield directionally consistent results in our estimation on the effect of IT resources. To identify *IT skills*, we use the skill requirements in the job postings as well as the job titles. For example, IT skills listed in a job posting can include software development as well as hardware support. IT-related job titles can include software engineer or systems analyst. If the job posting also contains keywords such as computer, website and telecommunication, we identify it as requiring IT skills. Similar to our construction of AI skills, we aggregate the personnel with IT skills in each firm. The IT resources of each firm can thus be calculated as the standardized sum of the standardized values of patents and skills measures related to IT: $IT\ resources = norm(norm(IT\ IP) + norm(IT\ skills))$.

et al. 2019). Recent advances in machine learning have been more useful in biopharma, as they can engage in more precise predictions of particular innovation events in the presence of data inputs from different information sources, and scale to an arbitrary level (Cockburn et al. 2018). They are suited to automatically learn hidden patterns from the data, find subtle linkages amongst data elements and make faster and better predictions. Instead of being hampered by complex interconnections, data collected from different sources can improve the accuracy of machine learning algorithms. Thus, machine learning is particularly useful for firms that exchange data within an R&D alliances.

Empirical tests suggest the distinct role of AI in leveraging knowledge flows from alliances. In Table 7, we also distinguish machine learning from other types of AI resources and estimate their effects (Cockburn et al. 2018; WIPO 2019). We find that the elasticity of the main effect of R&D alliances is 6.16%, the elasticity on machine learning resources is 4.78% and the elasticity of their interaction is 2.17%. However, the other types of AI such as expert systems have limited effect on number of drugs developed (its individual effect is significant at $p < 0.1$, but its interaction with alliances is not significant). The estimation results are directionally consistent when we put different types of AI resources together. Overall, we find that machine learning is the main type of AI resources driving the positive interaction effect of AI and alliances: 52% [$=2.17/4.17$] of the effect comes from machine learning, suggesting that the key advance in AI facilitating drug development comes from the ability to efficiently navigate a larger innovation space expanded by alliances. Machine learning is particularly well suited to reducing the cost of using the market for ideas (i.e., alliances), making alliances more useful for innovation.

Finally, in Table 8, we explore whether the AI/R&D alliance complementarity can be traced through to final drug success. Although the signal to noise ratio declines meaningfully as

we move from R&D inputs to later stage outputs, we undertake some exploratory analyses to examine later stages of drug development (i.e., phase I, phase II and phase III of clinical trials and final FDA approval). Specifically, we adopt a similar regression specification as specification (4), using a firm's stock of drugs developed in different later stages as dependent variables and controlling for the drugs developed in pre-clinical trial stage. We find that there is a positive interaction of R&D alliances and AI resources, though the elasticity of the interaction decreases across different later stages: 3.39% for drugs developed in phase I of clinical trials, 2.01% for those in phase II of clinical trials, 1.94% for those in phase III of clinical trials, and 1.55% for those in final approval stage, respectively. Although noisy, the results provide some suggestive evidence that AI leads to more drugs in the later stages of development.

Conclusion and Discussion

This research provides the conceptual and empirical basis for identifying a complementarity between a firm's AI resources and its R&D alliances in the firm's innovation production function. The main conceptual thrust is that by delivering a more capable analytical engine, AI resources reduce frictions inherent in markets for innovation (i.e., R&D alliances), allowing firms to capitalize on both the extensive margin, by making more use of the market for innovation (i.e., forming more R&D alliances), and on the intensive margin, by wringing more innovation out of each transaction (i.e., alliance).

The results show that doubling AI resources leads to 7% more R&D alliances (the extensive margin effect), and a 24% improvement in the innovation rate, compared to doubling AI resources and R&D alliances independently (the intensive margin effect). Furthermore, we show that higher levels of AI resources are associated with more knowledge sharing, precisely the mechanism at the heart of our theory.

Taken together, the evidence suggests that R&D alliances allow firms with greater AI resources to better exploit alliance partner's knowledge to improve their innovativeness. While we believe this is the first large sample study to show the complementarity between AI and R&D alliances, we are certainly not the first to evaluate the relationship amongst IT, alliances and innovation. However, compared to prior work, this study offers both more conceptual and empirical precision, and more robust measures of resources. Prior work has treated IT as the amalgam of disparate communication technologies, computer software and hardware, whereas we focus explicitly on AI as an analytical tool. Moreover, our index of AI resources includes both asset- and human capital-based resources, whereas prior work has typically considered only one or the other. We also build on exemplar prior work, by taking identification seriously and by providing evidence of the mechanism at play.

Nevertheless, this study suffers from several limitations that may indicate directions for future research. First, our empirical evidence is not based on a randomized experiment that assigns AI resources and R&D alliances at random. To deal with a non-random data generating process where firms choose both their level of AI resources and R&D alliances, we include a set of time-varying covariates and conduct fixed-effect estimations to account for observed and unobserved confounds. We use different econometric techniques to address several endogeneity concerns and we find consistent results. We also conduct some falsification tests that compare AI resources and IT resources and examine the effects of different categories of AI resources. Despite these efforts, we remain cautious in interpreting the observed effects as causal.

Second, our empirical analyses are subject to several data limitations. For example, our job posting data are not the same as the employment data. Posted job openings may not be filled, and therefore may not be fully representative of the actual human capital in firms. However, it is

encouraging that a recent study by Babina et al. (2021) finds that AI skills measured using job postings are highly correlated with those using resume data that have detailed employment history. It suggests that job posting data are a suitable tool to approximate AI skills in firms. Future research with proprietary sources of human capital data could further improve measurements of AI skills to triangulate our findings.

In addition, data limitations keep us from delving deeper into the mechanisms of the complementarity effects. We do show that AI can increase inter-organizational information flows within an alliance, but we cannot map alliances to specific NCEs or target diseases. Future research can collect more granular data to further validate the mechanism.

Finally, our study focuses on the complementarity effect of AI resources and alliances in the context of drug innovation. The findings from the study may extend beyond drug development to have broader effects on general R&D and innovation outcomes. Future research should consider the implications of AI resources and alliances on innovations in broader sectors or areas.

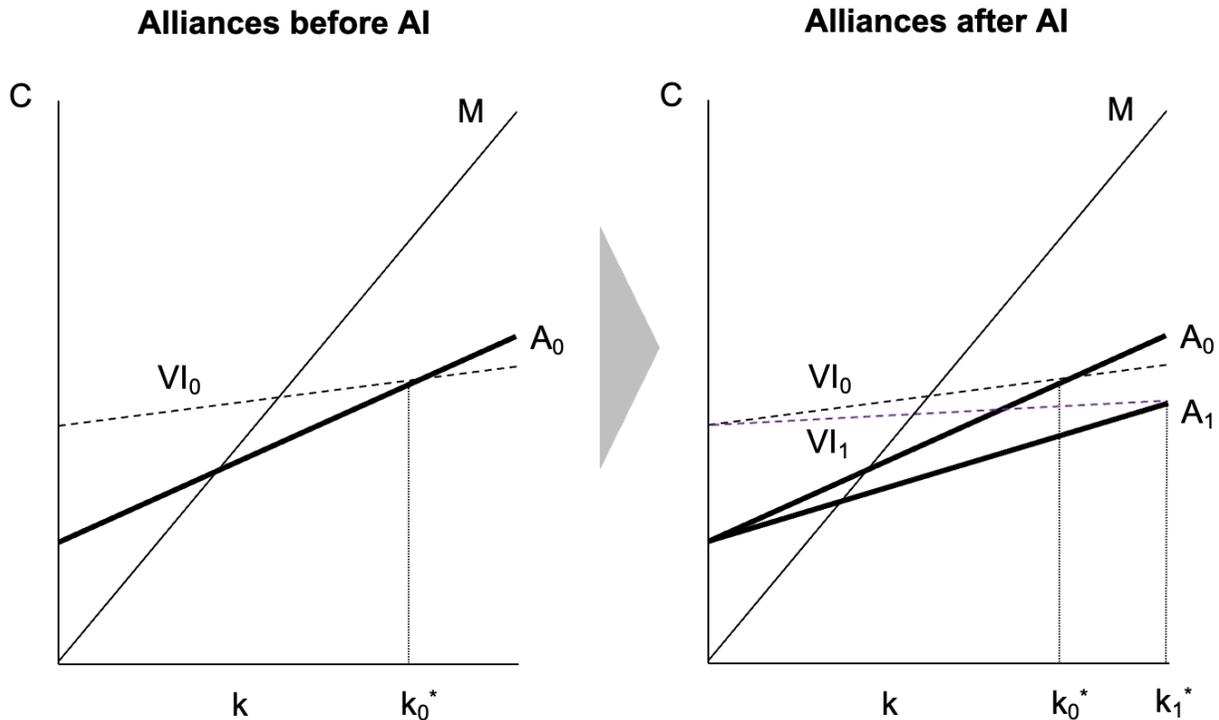
References

- Aggarwal, V. A., and Hsu, D. H. 2009. "Modes of Cooperative R&D Commercialization by Start-Ups," *Strategic management journal* (30:8), pp. 835-864.
- Aggarwal, V. A., and Hsu, D. H. 2013. "Entrepreneurial Exits and Innovation," *Management Science* (60:4), pp. 867-887.
- Agrawal, A., Gans, J., and Goldfarb, A. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Press.
- Ahuja, G. 2000. "Collaboration networks, structural holes, and innovation: A longitudinal study." *Administrative Science Quarterly*, 45: 425–455.
- Alekseeva, L., Azar, J., Gine, M., Samila, S., and Taska, B. 2021. "The Demand for Ai Skills in the Labor Market," *Labour Economics* (71), p. 102002.
- Almeida, P., Song, J., and Grant, R. M. 2002. "Are Firms Superior to Alliances and Markets? An Empirical Test of Cross-Border Knowledge Building," *Organization Science* (13:2): 147-161.
- Arora, A., Fosfuri, A., and Gambardella, A. 2001. "Markets for Technology and Their Implications for Corporate Strategy," *Industrial and corporate change* (10:2), pp. 419-451.
- Arora, A., Fosfuri, A., and Gambardella, A. 2002. "Markets for Technology in the Knowledge Economy,"
- Arora, A., and Gambardella, A. 1990. "Complementarity and External Linkages: The Strategies of the Large Firms in Biotechnology," *The journal of industrial economics*), pp. 361-379.
- Arora, A., and Gambardella, A. 1994. "Evaluating Technological Information and Utilizing It: Scientific Knowledge, Technological Capability, and External Linkages in Biotechnology," *Journal of Economic Behavior & Organization* (24:1), pp. 91-114.
- Arrow, K. 1962. "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*. Princeton University Press: 609-626.
- Babina, T., Fedyk, A., He, A. X., and Hodson, J. 2021. "Artificial Intelligence, Firm Growth, and Product Innovation," Available at SSRN: <https://ssrn.com/abstract=3651052>.
- Brynjolfsson, E., and Milgrom, P. 2013. "Complementarity in Organizations," *The handbook of organizational economics*), pp. 11-55.
- Campa, J. M., and Kedia, S. 2002. "Explaining the Diversification Discount," *The journal of finance* (57:4), pp. 1731-1762.
- Cassiman, B., and Veugelers, R. 2006. "In Search of Complementarity in Innovation Strategy: Internal R&D and External Knowledge Acquisition," *Management science* (52:1), pp. 68-82.
- Chae, H.C., Koh, C.E., and Prybutok, V.R. 2014. Information Technology Capability and Firm Performance: Contradictory Findings and Their Possible Causes. *MIS Quarterly* 38(1): 305-326.
- Cockburn, I. M., Henderson, R., and Stern, S. 2018. The Impact of Artificial Intelligence on Innovation. *National bureau of economic research* 0898-2937,.
- Cohen, W. M., and Levinthal, D. A. 1990. "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative science quarterly*): 128-152.
- Dyer, J. 1997. Effective Interfirm Collaboration: How Firms Minimize Transaction Costs and Maximize Transaction Value. *Strategic Management Journal* 18(7): 535-556.
- Fleming, N. 2018. "How Artificial Intelligence Is Changing Drug Discovery," *Nature* (557:7706), pp. S55-S55.

- Forman, C., Goldfarb, A., and Greenstein, S. 2016. "Agglomeration of Invention in the Bay Area: Not Just Ict," *American Economic Review* (106:5), pp. 146-151.
- Furman, J., and Seamans, R. 2019. "AI and the Economy," *Innovation policy and the economy* (19:1): 161-191.
- Griliches, Z., Pakes, A., and Hall, B. H. 1986. "The Value of Patents as Indicators of Inventive Activity." National Bureau of Economic Research Cambridge, Mass., USA.
- Hagedoorn, J. 2002. "Inter-Firm R&D Partnerships: An Overview of Major Trends and Patterns since 1960," *Research policy* (31:4): 477-492.
- Hall, B. H. 1990. "The Manufacturing Sector Master File: 1959-1987," *National Bureau of Economic Research*.
- Hall, B. H., Jaffe, A., and Trajtenberg, M. 2005. "Market Value and Patent Citations," *RAND Journal of economics*, pp. 16-38.
- Hall, B. H., Jaffe, A. B., and Trajtenberg, M. 2001. "The Nber Patent Citation Data File: Lessons, Insights and Methodological Tools," National Bureau of Economic Research.
- Harrer, S., Shah, P., Antony, B., and Hu, J. 2019. "Artificial Intelligence for Clinical Trial Design," *Trends in pharmacological sciences* (40:8), pp. 577-591.
- Hess, A.M., and Rothaermel, F.T. 2011. "When are Assets Complementary? Star Scientists, Strategic Alliances, and Innovation in the Pharmaceutical Industry". *Strategic Management Journal* 32(8): 895-909.
- Hitt, L. M. 1999. "Information Technology and Firm Boundaries: Evidence from Panel Data," *Information Systems Research* (10:2), pp. 134-149.
- Hughes, J.P., Rees, S., Kalindjian, S.B., and Philpott, K.L. 2011. "Principles of Early Drug Discovery," *British Journal of Pharmacology* 162(6): 1239-1249.
- Jaffe, A. B. 1986. "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits and Market Value," 0898-2937, national bureau of economic research.
- Krieger, J., Li, D., and Papanikolaou, D. 2022. "Missing Novelty in Drug Development," *The Review of Financial Studies* (35:2), pp. 636-679.
- Lee, J., Hoetker, G., and Qualls, W. 2015. "Alliance Experience and Governance Flexibility," *Organization Science* (26:5), pp. 1536-1551.
- Liu, Y., and Ravichandran, T. 2015. "Alliance Experience, It-Enabled Knowledge Integration, and Ex Ante Value Gains," *Organization Science* (26:2), pp. 511-530.
- Lou, B., and Wu, L. 2021. "Ai on Drugs: Can Artificial Intelligence Accelerate Drug Development? Evidence from a Large-Scale Examination of Bio-Pharma Firms," *MIS Quarterly* (45:3).
- Milgrom, P., and Roberts, J. 1990. "The Economics of Modern Manufacturing: Technology, Strategy, and Organization," *The American Economic Review*, pp. 511-528.
- Mowery, D. C., Oxley, J. E., and Silverman, B. S. 1996. "Strategic Alliances and Interfirm Knowledge Transfer," *Strategic management journal* (17:S2), pp. 77-91.
- Nakamura, M, Shaver J.M., and Yeung, B. 1996. "An Empirical Investigation of Joint Venture Dynamics: Evidence from U.S.-Japan joint ventures. *International Journal of Industrial Organization* 14(4): 521-541.
- Nevo, A. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica* (69:2), pp. 307-342.
- Oxley, J., and Wada, T. 2009. "Alliance Structure and the Scope of Knowledge Transfer: Evidence from Us-Japan Agreements," *Management science* (55:4), pp. 635-649.
- Oxley, J. E., and Sampson, R. C. 2004. "The Scope and Governance of International R&D Alliances," *Strategic Management Journal* (25:8-9), pp. 723-749.

- Pisano, G. 1990. The R&D Boundaries of the Firm: An Empirical Analysis. *Administrative Science Quarterly* 35(): 153-176
- Powell, W. W. 1998. "Learning from Collaboration: Knowledge and Networks in the Biotechnology and Pharmaceutical Industries," *California management review* (40:3), pp. 228-240.
- Powell, W. W., Koput, K. W., and Smith-Doerr, L. 1996. "Interorganizational Collaboration and the Locus of Innovation: Networks of Learning in Biotechnology," *Administrative science quarterly*), pp. 116-145.
- Raymond, P., Yoav, S., Erik, B., Jack, C., John, E., Barbara, G., Terah, L., James, M., Saurabh, M., and Carlos, N. J. 2019. "The Ai Index 2019 Annual Report," AI Index Steering Committee, Human-Centered AI Institute, Stanford University.
- Reuer, J.T., Zollo, M., and Singh, H. 2002. Post-Formation Dynamics in Strategic Alliances. *Strategic Management Journal* 23(2): 135-151.
- Rivkin, J., and Siggelkow, N. 2003. Balancing Search and Stability: Interdependencies among Elements of Organizational Design. *Management Science* 49(3): 290-311.
- Robinson, D. T., and Stuart, T. E. 2007a. "Financial Contracting in Biotech Strategic Alliances," *The Journal of Law and Economics* (50:3), pp. 559-596.
- Robinson, D. T., and Stuart, T. E. 2007b. "Network Effects in the Governance of Strategic Alliances," *The Journal of Law, Economics, & Organization* (23:1), pp. 242-273.
- Romer, P. M. 1990. "Endogenous Technological Change," *Journal of political Economy* (98:5, Part 2), pp. S71-S102.
- Rosenbaum, P. R., and Rubin, D. B. 1983. "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika* (70:1), pp. 41-55.
- Rosenkopf, L., and Almeida, P. 2003. "Overcoming Local Search through Alliances and Mobility," *Management science* (49:6), pp. 751-766.
- Sampson, R. C. 2007. "R&D Alliances and Firm Performance: The Impact of Technological Diversity and Alliance Organization on Innovation," *Academy of management journal* (50:2), pp. 364-386.
- Savage, N. 2021. "Tapping into the Drug Discovery Potential of Ai," *Biopharma Deal*).
- Schumpeter, J. A. 1934. "The Theory of Economic Development,")
- Sosa, M. L. 2013. "Corporate Structure, Indirect Bankruptcy Costs, and the Advantage of De Novo Firms: The Case of Gene Therapy Research," *Organization Science* (25:3), pp. 850-867.
- Tambe, P., and Hitt, L. M. 2012. "The Productivity of Information Technology Investments: New Evidence from It Labor Data," *Information Systems Research* (23:3-part-1), pp. 599-617.
- Webb, M. 2019. "The Impact of Artificial Intelligence on the Labor Market," *Available at SSRN 3482150*).
- Williamson, O. E. 1991. "Comparative Economic Organization: The Analysis of Discrete Structural Alternatives," *Administrative science quarterly*), pp. 269-296.
- WIPO. 2019. "Wipo Technology Trends 2019: Artificial Intelligence."
- Wu, L., Hitt, L., and Lou, B. 2020. "Data Analytics, Innovation, and Firm Productivity," *Management Science* (66:5), pp. 2017-2039.
- Wu, L., Lou, B., and Hitt, L. 2019. "Data Analytics Supports Decentralized Innovation," *Management Science* (65:10), pp. 4863-4877.
- Yang, Y, Adelstein, S.J., and Kassis, A.I. 2009. Target Discovery form Data Mining Approaches. *Drug Discovery Today* 14(3-4): 147-154.

Figure 1: AI and alliances: a Transaction Cost Economics approach



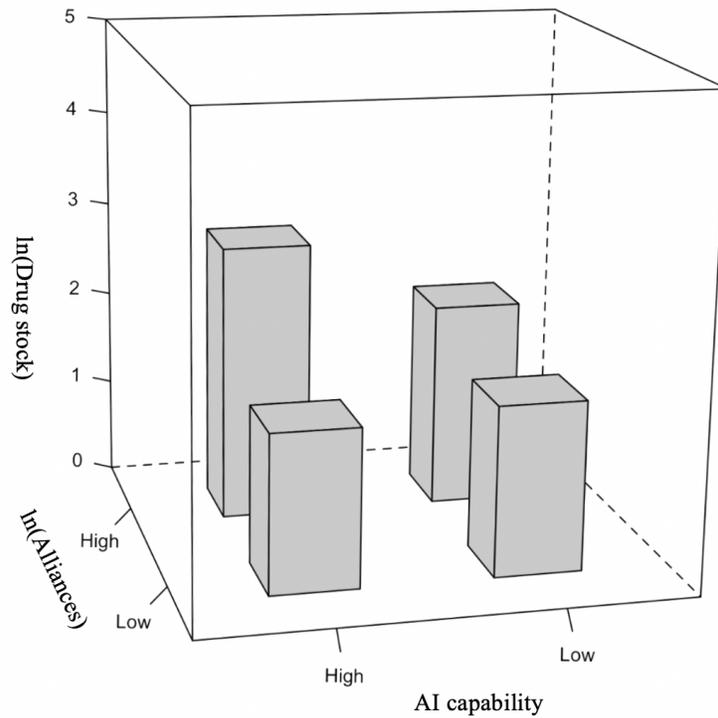
The left-hand panel of this figure is reproduced from Williamson 1991 p. 284. It shows the comparative costs of governance (C) relative to asset specificity—the cost of redeploying assets to their next best use—for three different types of transactions: market-based exchange (M), vertical-integration (VI) and alliances (A). Subscript 0 is used to denote the initial conditions of the optimal firm boundary choices before AI. We have noted where A crosses VI as point k_0^* for reference. Above k_0^* firms will choose vertical integration, below k_0^* , until A crosses M , firms will choose alliances.

The right-hand panel shows how the boundary of the firm shifts toward alliances (i.e., to point k_1^*) with the introduction of AI, assuming a proportional improvement in the comparative costs of governance for both alliances and vertical integration (i.e., the slope decreases proportionally for each).

We do not plot the point where A crosses M in either panel, because it is trivial to see that it will move downward (i.e., shifting formerly market-based transactions to alliances) with the introduction of AI, assuming that AI cannot help facilitate spot market exchange in the pharmaceutical industry.

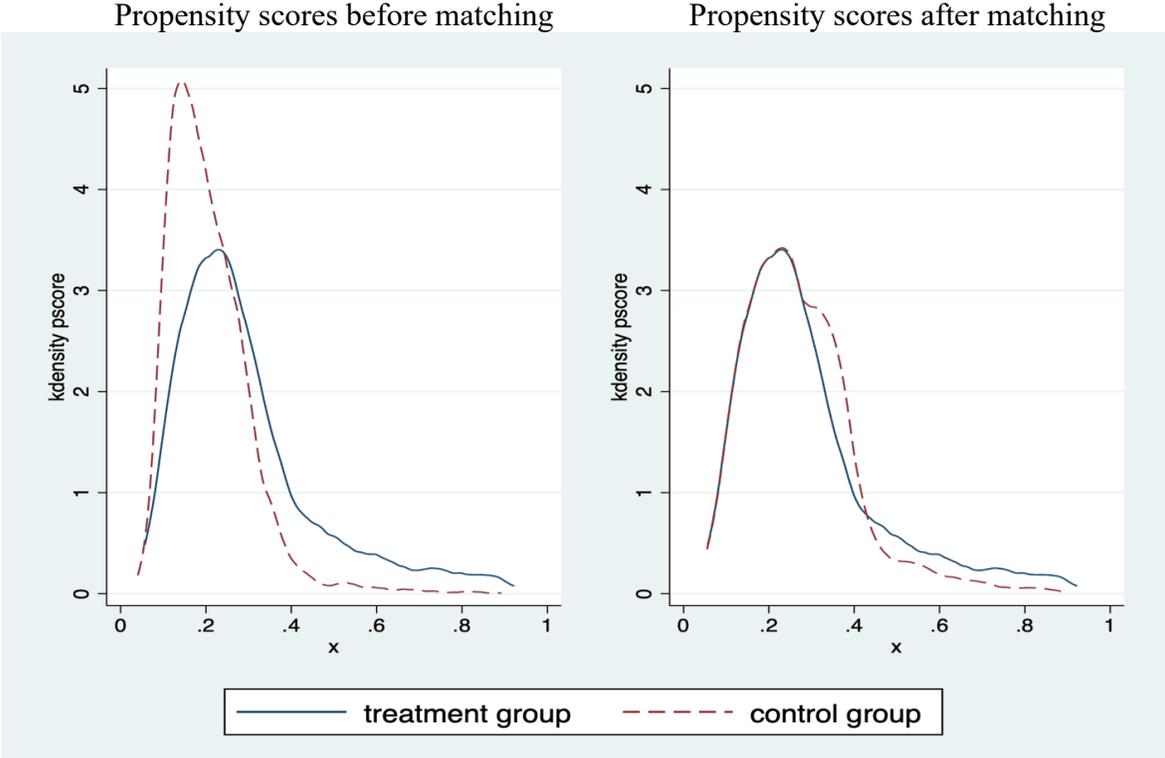
Proof: By definition (i) the slope of the alliance line is steeper than the slope of the vertical integration line: $m_{AL} > m_{VI}$ and (ii) the vertical integration intercept is greater than the alliance intercept: $b_{VI} > b_A$. The x-coordinate of $x(k_0^*) = (b_{VI} - b_A) / (m_A - m_{VI})$. Both the numerator and denominator are positive so we know $k_0^* > 0$. Now if we decrease the slope (i.e., increase efficiency) by the same proportion d , we get: $x(k_1^*) = (b_{VI} - b_A) / [(m_A - m_{VI})(1 - d)]$. So $k_1^* > k_0^*$.

Figure 2: Model-free evidence for an AI resources-R&D alliance complementarity



This figure displays the average logged number of drugs developed by type of firm (e.g., firms with high levels of AI resources and high levels of alliances vs. firms with low levels of AI resources and low levels alliances). The median level of number of alliances is zero. Thus, “Low” values of the number of alliances indicate zero alliances (at the firm level), while “High” values are greater-than-zero alliances. The dichotomy of AI resources (i.e., low vs. high) are also defined by a median split. Consistent with our theory, high AI capability is associated with improved productivity within alliance relationships.

Figure 3: Propensity score predicting an alliance before and after matching



The matched sample was created using 1:1 nearest neighbor matching on the probability of forming an alliance at time t (“the propensity score”) for firms with an alliance (the “treatment” group), compared to firms that did not have an alliance at time t (the “control” group).

Table 1. Summary statistics and correlation table (n = 4,778)

Variable	Mean	Std dev.	Min	Max	Pairwise correlations								
					1	2	3	4	5	6	7	8	
1. Drugs developed	15.33	37.35	0	561									
2. AI IP	1.22	7.31	0	142.70	0.21								
3. AI skills	16.12	143	0	4,269	0.44	0.39							
4. Alliances	0.61	2.21	0	42	0.76	0.20	0.45						
5. Patents	288.50	1,202	0	21,145	0.62	0.55	0.43	0.55					
6. Public company	0.42	0.49	0	1	0.19	0.11	0.11	0.12	0.16				
7. Firm age	23.81	28.33	0	347	0.23	0.16	0.17	0.23	0.30	0.06			
8. Employees	4,997	15,951	1	316,320	0.51	0.28	0.32	0.41	0.52	0.11	0.21		
9. R&D spend (\$B)	3.56	12.10	<0.01	289	0.10	0.05	0.04	0.08	0.07	0.08	0.03	0.08	

Notes: *AI resources* is standardized to be a mean 0, standard deviation 1, variable. For ease of interpretation, summary statistics are reported before taking logs.

Table 2. AI and R&D alliances: baseline results

<i>Dependent variable Specification</i>	(1) Alliance in year <i>t</i> OLS	(2) Alliance in year <i>t</i> FE	(3) Alliances FE
<i>AI resources</i>	6.17* (0.66)	5.91* (2.22)	6.68* (2.35)
Patents	1.71* (0.36)	-0.70 (1.27)	-0.60 (1.32)
Public company	5.94* (1.22)	0.95 (2.90)	2.12 (2.93)
Firm age	2.93* (0.81)	-4.56 (4.86)	-5.94 (4.96)
Employees	0.86* (0.25)	0.03 (0.37)	0.20 (0.37)
R&D spend	0.14 (0.22)	-0.23 (0.38)	-0.19 (0.36)
Year fixed effects	Y	Y	Y
Firm fixed effects	N	Y	Y
n	4,778	4,778	4,778
R ²	0.09	0.34	0.61

All coefficient and standard errors are multiplied by 100 for presentation purposes. All variables enter in logs, except *AI resources*, *Public company*, the fixed effects and the dependent variable in columns 1 and 2. Robust standard errors are clustered by firms in specifications with firm fixed effects. OLS regressions are reported, though the results are similar in terms of sign, significance and approximate economic magnitude using probit or logit.
* p<0.05, + p<0.1

Table 3. AI and R&D alliances: instrumental variables regressions

<i>Dependent variable</i>	(1)	(2)	(3)	(4)	(5)
<i>Specification</i>	AI resources 1 st stage	Alliance in year <i>t</i> 2 nd stage	AI resources 1 st stage	Alliance in year <i>t</i> 2 nd stage	Alliances 2 nd stage
# of neighboring firms w/AI resources	1.27* (0.12)		0.20* (0.08)		
<i>AI resources</i>		8.85* (1.60)		24.80⁺ (13.60)	32.60⁺ (18.60)
Patents	7.89* (1.77)	1.26* (0.44)	8.31* (2.36)	-2.44 (1.78)	-3.00 (2.10)
Public company	18.70* (5.13)	5.31* (1.26)	11.50* (4.30)	-1.36 (3.08)	-0.95 (3.43)
Firm age	5.74* (4.48)	2.65* (0.83)	0.51 (6.56)	-4.66 (4.58)	-6.08 (4.80)
Employees	5.03* (1.07)	0.69* (0.26)	1.11* (0.48)	-0.14 (0.38)	-0.05 (0.41)
R&D spend	2.71* (0.81)	0.05 (0.22)	0.31 (0.46)	-0.31 (0.37)	-0.29 (0.36)
Year FE	Y	Y	Y	Y	Y
Firm FE	N	N	Y	Y	Y
n	4,778	4,778	4,778	4,778	4,778
R ²	0.09	n/a	0.61	n/a	n/a

Two-stage least squares (2SLS) regression reported. All coefficient and standard errors are multiplied by 100 for presentation purposes. All variables enter in logs, except *AI resources*, *Public company*, the fixed effects and the dependent variables in specifications 1-4. Robust standard errors are clustered by firms in specifications with firm fixed effects. * p<0.05, + p<0.1

Table 4. Comparison of firm characteristics before and after matching on the propensity to form an alliance

Variable	Before Matching			After Matching		
	(1)	(2)	P-value (1) - (2)	(3)	(4)	P-value (3) - (4)
	Firms with an alliance	Firms without an alliance		Firms with an alliance	Firms without an alliance	
AI IP	0.38	0.19	<0.01	0.38	0.33	0.16
AI skill	0.71	0.19	<0.01	0.71	0.43	<0.01
Patents	4.38	3.63	<0.01	4.38	4.20	0.05
Public company	0.52	0.39	<0.01	0.52	0.52	0.85
Firm age	3.07	2.81	<0.01	3.07	3.00	0.08
Employees	6.56	5.88	<0.01	6.56	6.46	0.41
R&D spend	20.20	19.88	<0.01	20.20	20.20	0.96
Year: 2011	0.08	0.11	<0.01	0.08	0.07	0.59
Year: 2012	0.07	0.12	<0.01	0.07	0.08	0.43
Year: 2013	0.10	0.12	0.03	0.10	0.09	0.69
Year: 2014	0.15	0.11	<0.01	0.15	0.15	0.90
Year: 2015	0.16	0.11	<0.01	0.16	0.17	0.71
Year: 2016	0.16	0.11	<0.01	0.16	0.16	0.95
Year: 2017	0.13	0.10	0.05	0.13	0.13	0.89
Year: 2018	0.08	0.12	<0.01	0.08	0.08	0.55
Year: 2019	0.08	0.10	0.06	0.08	0.08	0.86
n	1,050	3,728		939	939	

The matched sample was created using 1:1 nearest neighbor matching on the probability of forming an alliance at time t (“the propensity score”) for firms with an alliance (the “treatment” group), compared to firms that did not have an alliance at time t (the “control” group). Year: 2010 is dropped because explanatory variables are all lagged by one for matching. The F-test on the joint significance of the differences in the means of the covariates is strongly statistically significant before matching, but not statistically significant at conventional levels ($p = 0.17$) after the matching.

Table 5. AI, alliances, and innovation

Dependent variable = Drugs developed

<i>Dependent variable:</i>	<i>Sample:</i>			<i>Matched sample</i>		
	Full sample			AI resources	AI resources x Alliances	Drugs developed
<i>Specification:</i>	OLS	OLS	OLS	2SLS 1 st stage-I	2SLS 1 st stage-II	2SLS 2 nd stage
	(1)	(2)	(3)	(4)	(5)	(6)
Alliances	7.62* (1.71)	7.10* (1.66)	6.73* (1.63)			
# of neighboring firms w/AI resources				0.39* (0.16)	-0.22 (0.17)	
# of neighbor. firms w/AI resources x Alliances				-0.08 (0.06)	0.38 ⁺ (0.22)	
AI resources		8.54* (2.29)	6.77* (2.24)			
<i>AI resources x Alliances</i>			4.17* (1.75)			7.51⁺ (4.17)
Patents	12.10* (2.44)	12.00* (2.41)	11.90* (2.41)	Y	Y	Y
Public company	13.20* (4.86)	12.40* (4.84)	12.30* (4.84)	Y	Y	Y
Firm age	79.30* (9.51)	79.20* (9.45)	79.60* (9.43)	Y	Y	Y
Employees	-0.49 (0.56)	-0.54 (0.56)	-0.54 (0.56)	Y	Y	Y
R&D spend	0.83 (0.54)	0.76 (0.54)	0.73 (0.54)	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y	Y
n	4,778	4,778	4,778	1,878	1,878	1,878
R ²	0.94	0.94	0.94	0.94	0.94	n/a

All coefficient and standard errors are multiplied by 100 for presentation purposes. All variables enter in logs, except *AI resources*, *Public company*, the instrument (number of neighboring firms with AI resources) and the fixed effects. Column 4-6 use the matched sample from propensity score matching. The associated F-statistic for the first stage is 27.85. Robust standard errors are clustered by firms. * p<0.05, ⁺ p<0.

Table 6. AI and knowledge flows within alliances*Dependent variable = Knowledge flows*

Specification	(3) FE	(4) 2SLS
<i>AI resources</i>	8.07* (2.94)	18.00* (8.54)
Patents	1.27 (0.97)	-0.92 (1.23)
Public company	-5.06 (9.62)	-4.86 (3.52)
Firm age	-7.47 (10.00)	0.67 (3.54)
Employees	-0.46 (0.32)	0.21 (0.28)
R&D spend	-0.32 (0.34)	-0.65 (0.46)
Year FE	Y	Y
Firm FE	Y	Y
n	2,755	2,755
R ²	0.34	n/a

All coefficient and standard errors are multiplied by 100 for presentation purposes. The sample includes all alliances in the data used in the previous tables. All RHS variables enter in logs, except *AI resources*, *Public company* and the fixed effects. Knowledge flows are measured as the natural logarithm of the difference in firms' citations to their partners before and after alliances. The instrument was strong in the first stage of the 2SLS estimator (available upon request). Robust standard errors are clustered by firms. * p<0.05

Table 7. AI vs. IT, machine learning vs. expert systems*Dependent variable = Drugs developed*

	(1)	(2)	(3)	(4)	(5)
Alliances	6.18* (1.58)	6.14* (1.51)	6.16* (1.54)	6.21* (1.57)	6.15* (1.55)
IT resources	3.82 (2.38)	-0.30 (3.77)			
IT resources x Alliances	1.95 (1.23)	-2.38 (2.64)			
AI resources		5.56+ (3.24)			
<i>AI resources x Alliances</i>		5.34* (2.66)			
Machine learning			4.78* (1.74)		3.98+ (2.15)
<i>Machine learning x Alliances</i>			2.17* (1.10)		2.30+ (1.33)
Expert systems				3.24+ (1.68)	1.56 (1.93)
Expert systems x Alliances				0.07 (0.73)	-0.23 (0.91)
Patents	10.70* (2.39)	11.10* (2.33)	11.30* (2.35)	11.50* (2.36)	11.30* (2.36)
Public company	10.90* (4.58)	10.50* (4.58)	10.90* (4.58)	11.10* (4.59)	10.90* (4.58)
Firm age	78.00* (9.32)	77.90* (9.34)	78.70* (9.41)	78.10* (9.39)	78.70* (9.41)
Employees	-0.57 (0.53)	-0.56 (0.53)	-0.55 (0.53)	-0.54 (0.53)	-0.55 (0.53)
R&D spend	0.71 (0.49)	0.73 (0.49)	0.76 (0.49)	0.77 (0.49)	0.76 (0.49)
Year FE	Y	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y	Y
n	4,778	4,778	4,778	4,778	4,778
R ²	0.94	0.94	0.94	0.94	0.94

All coefficient and standard errors are multiplied by 100 for presentation purposes. All variables enter in logs, except *AI resources*, *Public company* and the fixed effects. Robust standard errors are clustered by firms. * p<0.05, + p<0.1

Table 8. AI, alliances, and long-run innovation success

<i>Dependent variable</i>	(1) Drugs developed Clinical Trial - Phase I	(2) Drugs developed Clinical Trial - Phase II	(3) Drugs developed Clinical Trial - Phase III	(4) Drugs developed FDA Approval
Alliances	2.03 (1.35)	-1.26 (1.00)	1.50 (0.98)	1.23 (0.85)
AI resources	3.35 ⁺ (1.86)	-1.49 (1.59)	-2.33 (1.42)	-1.13 ⁺ (0.65)
<i>AI resources x Alliances</i>	<i>3.39*</i> <i>(1.68)</i>	<i>2.01*</i> <i>(1.02)</i>	<i>1.94*</i> <i>(1.05)</i>	<i>1.55*</i> <i>(0.89)</i>
Drugs developed, Pre-clinical Trial	20.3* (2.69)	3.15 (2.13)	0.71 (1.72)	-0.63 (1.00)
Patents	3.95 ⁺ (2.11)	0.21 (1.66)	1.14 (1.21)	0.69 (0.96)
Public Company	12.40* (4.16)	5.55 ⁺ (3.18)	0.80 (2.33)	0.05 (1.66)
Firm age	-4.29 (5.24)	6.23 (4.67)	-2.86 (3.18)	-5.55* (2.34)
Employees	-0.32 (0.38)	0.11 (0.32)	0.42 ⁺ (0.22)	0.27 (0.19)
R&D spend	0.24 (0.40)	0.37 (0.29)	0.22 (0.24)	0.06 (0.13)
Year FE	Y	Y	Y	Y
Firm FE	Y	Y	Y	Y
n	4,778	4,778	4,778	4,778
R ²	0.94	0.94	0.90	0.93

All coefficient and standard errors are multiplied by 100 for presentation purposes. All variables enter in logs, except *AI resources*, *Public company* and the fixed effects. Robust standard errors are clustered by firms. * p<0.05, + p<0.1

Appendix: Taxonomy used to identify AI-related patents or skills

General:	artificial intelligence, computational control, computer vision, machine intelligence, natural language processing
Expert system:	expert system, fuzzy logic, inference engine, logic program, logic system, rule-based inference, symbolic reasoning
Machine learning:	adaboost, convolutional neural network, conditional random field, decision tree, deep learning, hidden markov, latent dirichlet allocation, libsvm, keras, long short-term memory, machine learning, mxnet, pytorch, random forests, recurrent neural network, reinforcement learning, scikit-learn, semi-supervised learning, stochastic gradient descent, supervised learning, support vector machine, tensorflow, unsupervised learning, word2vec, xgboost