Beyond Patent Ownership: Learning About Technological Usefulness^{*}

Jesus Gorrin^{\dagger} Rory Mullen^{\ddagger}

First version: June 7, 2025 This version: June 15, 2025

Abstract

Technology is central to economics, but current datasets lack the scale, scope, span, and specificity economists need. We apply natural language processing and positive-unlabeled machine learning to descriptions of U.S. patents and public firms over nearly three decades to create a firm-level technology dataset that is unmatched in its combination of these qualities. For the first time, we reveal the core technologies of non-patenting firms. We estimate the value these firms derive from technological innovation, and argue that stock market participants inefficiently process their technological information, enabling profitable trading strategies.

JEL: O32, G12, C55, O34, C81

Keywords: technology, patents, firm-level, positive-unlabeled learning, textual analysis, technology spillovers, technology momentum, non-patenting firms

^{*}Working paper: comments welcome. Results are preliminary and subject to revision. We thank Vashisht Bhatt for excellent research assistance, Junxi Liu for contributions to an earlier version of the paper circulated under the title "Firm Technology Usage," and the Gillmore Centre for Financial Technology at the University of Warwick for financial support.

[†]Warwick Business School, University of Warwick, Jesus.Gorrin@wbs.ac.uk

[‡]Warwick Business School, University of Warwick, Rory.Mullen@wbs.ac.uk

1 Introduction

Technology is central to economics, driving everything from growth and business cycles (Solow, 1960; Kydland and Prescott, 1982) to labor markets and financial markets (Autor, Levy, and Murnane, 2003; Kogan, Papanikolaou, Seru, and Stoffman, 2017). Yet the majority of firms—those that do not patent—remain technologically veiled, obscuring how technology is used throughout much of the economy. Current technology datasets lack the scale, scope, span, and specificity that economists need, and often focus on the small minority of firms that own intellectual property rather than the broader universe of firms that use technology. While many technology datasets exist, each has important limitations.

In this paper, we construct a novel dataset linking firms to technologies using positive and unlabeled machine learning. Our firm-level technology dataset offers an unprecedented combination of scale, scope, span, and specificity. In scale, it covers all U.S. public firms and a sample of 50,000 U.S. utility patents per year, amounting to hundreds of millions of firm-patent pairs. In scope, it covers thousands of distinct technology categories curated by patent examiners with subject-specific expertise, representing all major areas of economically important innovation. In span, it covers nearly three decades of innovation, including the rise of internet and mobile computing, advances in biotechnology and medical technology, and the growth of renewable energy and electric vehicles. In specificity, it characterizes the usefulness of each individual patent to each firm as a continuous probability, allowing us to characterize each firm's technological associations in granular detail and by degrees. This contrasts with the available binary indicators—such as patent ownership or innovation survey responses—which frequently cover few firms or few technologies, lack specificity, or convey no variation in the strength of association.

Existing datasets typically rely on patent ownership or innovation surveys. Patent ownership datasets have several limitations. First, few firms receive patent grants—fewer than 1% of firms in the U.S. Census Bureau's Business Register (Graham, Marco, and Miller, 2015) and only 15.64% of CRSP firms in an average year. Second, patent-owning firms are atypical—they are unusually large and concentrate in manufacturing (Mezzanotti

and Simcoe, 2023). Third, patent ownership does not always signal technology use—firms often file patents for reasons other than protecting actively used technology, including patent blocking, use as bargaining chips in negotiations, and prevention of suits (Cohen, Nelson, and Walsh, 2000; Moore, 2005). Fourth, patent ownership does not imply exclusive access—intellectual property rights must be asserted through costly litigation (Lemley and Shapiro, 2005). Fifth, patent owners regularly grant access to others—around 40% of patents are embodied in commercial products (Argente et al., 2023) and around 4% are reassigned annually (Graham, Marco, and Myers, 2018). These limitations make patent ownership data unsuitable for many economic applications. Innovation surveys overcome some limitations but lack the specificity of patent data. For example, survey respondents may report *that* innovations occurred but rarely describe *which* specific innovations.

To overcome these limitations, we construct a new dataset that emphasizes technological usefulness use over patent ownership. We start from a simple premise: if the language that describes a patent is similar to the language that describes a firm, the patent is probably useful for the firm—regardless of who owns it. Building on this premise, we develop a novel methodology that combines natural language processing with techniques borrowed from an area of machine learning called positive and unlabeled learning. Positive and unlabeled learning is new to the economics literature, and allows us to view patent ownership as a positive signal that a patent is useful to the firm that owns it, without viewing the patent as useless to other firms. With positive and unlabeled learning, we train a classifier on positive usefulness labels exclusively, and use the classifier to predict the usefulness of patents to all firms, regardless of ownership.

We construct our dataset in two steps. First, we measure similarities in the language used to describe patents and firms. The patent descriptions come from a corpus of patent filings with the U.S. Patent and Trademark Office. The firm descriptions come from a corpus of annual reports filed with the U.S. Securities and Exchange Commission. These sources have each been extensively studied in the economics literature, but they have not been combined with the aim of constructing a comprehensive firm-level technology dataset covering non-patenting and patenting firms for use in economic research. We represent each description numerically using traditional term frequencies and inverse document frequencies and using a modern natural language model built on the transformer architecture. We then compute similarity scores between these numerical representations for all firm-patent pairs. While the techniques we use in the first step of our methodology are now commonplace in the economics literature, we take an important second step that is novel.

In the second step, we use similarity scores to estimate probabilities that each firm finds each patent useful. We refer to these estimates as *usefulness probabilities*. Traditional supervised machine learning requires both positive and negative labels; in our setting, these labels would indicate that a firm finds a patent useful or useless, respectively. But these labels are not always available in our setting. When a patent is owned by a firm, we assume the firm is likely to find the patent useful, and assign a (possibly noisy) positive label to the pair. But for other firm patent pairs, we cannot immediately assign negative labels, because patented innovations may be useful to firms who do not own them. Lacking negative labels, we turn to positive and unlabeled machine learning techniques specifically designed to overcome this challenge.

Positive and unlabeled learning, while novel in the economics literature, is wellestablished in the machine learning literature. We adopt a classic two-stage procedure from this literature, which was proposed by Liu et al. (2002). In the first stage, we assign negative labels to a random selection of positive firm-patent pairs, and use these "spies" to estimate a probability threshold for identifying firm-patent pairs where the patent is reliably useless to the firm. We then train a second-stage classifier on positive and reliably-negative firm-patent pairs. The spy procedure has an intuitive appeal and imposes few assumptions on the labeling process.

Using this approach, we uncover surprising similarities between non-patenting and patenting firm types. After conditioning on industry and firm size, we find that within-type variation in technological associations substantially exceeds between-type variation. We also find difference: non-patenting firms associate with broader but shallower technological portfolios and exhibit higher rates of technological instability over time. The technological associations we uncover have economic consequences: technology momentum portfolios of non-patenting firms generate monthly alphas of 1.97%, significantly exceeding strategies restricted to patenting firms. The superior returns appear to be driven by slower information diffusion about non-patenting firms' technological profiles. Event study evidence confirms that technological spillovers extend beyond patent ownership—firms experience positive abnormal returns of approximately 0.9 basis points per useful patent over 30 days following patent announcements, even for patents they do not own. Together, these results demonstrate that technological associations matter for firm value, even for non-patenting firms.

Related Literature. We contribute to three strands of literature. First, we build on efforts to construct firm-level technology and innovation datasets. Early approaches used direct profiling methods. The UN Industrial Development Organization's Profiles of Manufacturing Plants documented equipment use across manufacturing plants worldwide in the late 1960s and early 1970s (United Nations, 1971), but these profiles lacked standardization, covered few plants per industry, and the series was discontinued after three editions. The U.S. Census Bureau's Survey of Manufacturing Technology in the late 1980s and early 1990s systematically surveyed over 10,000 U.S. manufacturing plants about their use and planned adoption of seventeen advanced technologies (see Dunne, 1994, for a description). However, the survey focused on a narrow set of technologies and was also discontinued after three editions.

Standardized innovation surveys emerged to provide broader, ongoing coverage of firm innovation activities. The E.U. Community Innovation Survey (launched in 1992) and U.S. National Science Foundation Business R&D and Innovation Survey (launched in 2008) now measure firm-level innovation under a common framework. These surveys represent a significant advance in scale and span over earlier profiling efforts. Yet their indicators lack scope and specificity, relying on firms' self-reported yes/no responses and coarse categorizations ("product" versus "process") that limit the utility of the data.

In parallel, a tradition developed around patent-based measures of firm technology. A

major NBER initiative launched in the 1980s created the first large-scale patent ownership dataset, with early contributions collected in Griliches (1987), a retrospective in Hall, Jaffe, and Trajtenberg (2001), and an update in Arora, Belenzon, and Sheer (2021). Patent datasets offer advantages in scale, scope, span, and specificity compared to survey approaches. However, patent data exclude the majority of firms—those that do not patent—creating significant coverage gaps. We contribute to this literature by providing highly-detailed technological profiles of non-patenting firms using positive-unlabeled machine learning. Unlike Bloom, Hassan, Kalyani, Lerner, and Tahoun (2021), who study 29 disruptive technologies, we study a comprehensive set of legacy and disruptive technologies.

Second, we contribute to a literature that uses the datasets described above to study how technology and innovation relate to firm productivity and performance. The literature is too large to survey here (Lerner and Seru, 2021, identify over 80 papers in top economics journals between 2005 and 2020 using patent data), but seminal contributions include: Griliches (1990), who surveys patents as economic indicators; Lerner (1994), who examines patent scope and firm value; Hall, Jaffe, and Trajtenberg (2005), who show patent citations better predict firm value than counts; Bloom, Schankerman, and Van Reenen (2013), who identify productivity spillovers from patented innovations; Kogan et al. (2017), who measure how innovation drives firm growth and aggregate productivity; and Akcigit and Kerr (2018), who examine how different innovation types affect firm productivity. We contribute by comparing, for the first time, the technological profiles of non-patenting and patenting firms. We find that the technology use of non-patenting firms rivals that of patenting firms, after conditioning on industry and firm size, with both groups exposed to similar shocks and spillovers.

Third, we contribute to the growing literature on natural language processing and machine learning methods in economics and finance. Gentzkow, Kelly, and Taddy (2019) and Ash and Hansen (2023) survey recent work in economics, and Loughran and McDonald (2020b) and Kelly and Xiu (2023) survey recent work in finance. Papers studying firms' regulatory filings include Hoberg and Phillips (2016) on dynamic product-market industries,

Hoberg and Phillips (2018) on industry momentum, Lopez-Lira (2019) on risk factors for asset pricing, and Loughran and McDonald (2020a) and Cohen, Malloy, and Nguyen (2020) on firm complexity. Papers studying patent descriptions include Myers and Lanahan (2022) on R&D spillovers, Lerner et al. (2021) on patented financial technologies, Kogan et al. (2021) on technology and labor productivity, and Kakhbod et al. (2024) on innovation spillovers. To our knowledge, no prior work has linked regulatory filings with patent data to produce a comprehensive firm-level technology dataset covering non-patenting and patenting firms for use in economic research.

The paper is organized as follows. Section 2 describes our data sources. Section 3 describes our methodology for classifying patents as useful to firms. Section 4 compares the technological profiles of non-patenting and patenting firms. Section 5 studies a technology momentum investment strategy. Section 6 estimates the effect of new patent grants on stock market performance. Section 7 concludes.

2 Data Sources

Our study combines traditional financial data with textual data on public firms and patent grants in the United States. The financial data are CRSP daily and monthly stock files and Fama-French return factors (Fama, 2023). The textual data are business descriptions extracted from annual reports from the Electronic Data Gathering, Analysis, and Retrieval database (EDGAR) of the U.S. Securities and Exchange Commission (SEC) and patent descriptions from the PatentsView database of the U.S. Patent and Trademark Office (USPTO).

Figure 1 plots annual counts of business and patent descriptions. These descriptions offer broad coverage of technology users and technological innovations over an extended period. The SEC data describes the products, operations, and intellectual property of firms, while the USPTO data describes the novelty, function, and intended use of technology. These complementary perspectives motivate our basic hypothesis: that a patent is likely to be useful to a firm when the patent and firm are described in similar

Figure 1: Annual Business and Patent Description Counts



Notes. The figure shows annual counts of SEC business descriptions and USPTO patent descriptions. For business descriptions, the solid line plots the total number of available annual reports, while the dashed line plots the number of business descriptions we were able to extract. For patent descriptions, the solid line plots the total number of patent grants, while the dashed line plots the number of utility patent grants. We sample 50,000 utility patents randomly from each yearly total.

language—even if the firm does not own the patent.

We provide details about business descriptions from SEC filings in Section 2.1 below, and about patent descriptions from the USPTO filings in Section 2.2. We describe the Cooperative Patent Classification (CPC) system used to categorize patents in Section 2.3, and we describe our approach to identifying patent owners in Section 2.4.

2.1 Business Descriptions

The Securities Exchange Act of 1934 mandates that certain firms register their securities and file annual reports with the SEC (U.S. Congress, 1934). The mandate applies to firms with securities registered under Section 12 of the Act, which includes firms listed on a national securities exchange as well as unlisted firms that own assets exceeding \$10 million in value and issue equity securities held by more than 2,000 persons or more than 500 persons who are not accredited investors.¹

The SEC requires annual reports to be filed on Forms 10-K and 20-F. Form 10-K filings must include non-financial information as outlined in Regulation S-K (U.S. SEC, 2013). Form 10-K is used by domestic U.S. issuers, while Form 20-F is used by foreign

¹The act is amended from time to time, with updated thresholds and other changes. Exemptions or modified requirements apply for certain types of securities and issuers, such as securities issued by banks, savings and loan associations, and religious, educational, or charitable organizations.

private issuers with shares listed on U.S. national securities exchanges and follows its own form-specific disclosure requirements. Item 101 of Regulation S-K requires 10-K filing firms to disclose their dominant business segments and markets served, competitive conditions, principal products and services, material contracts and customer dependencies, government contracts, material government regulation and compliance costs, distribution methods, sources and availability of raw materials, research and development activities, and patents, trademarks, and other intellectual property.² Comparable disclosures are required in Form 20-F.³ In prior research, Hoberg and Phillips (2016) have primarily focused on product market information contained in business descriptions, but Regulation S-K requires disclosures to contain operational and technological information that goes beyond product markets.

We use the SEC's annual index files to identify relevant filings; these files record an identifier, the filing date, form type, and company name for each filing.⁴ We then scrape the full text of relevant filings and extract business descriptions from these. Over the period 1997 to 2023, we achieve an average annual extraction rate of 82%. We obtain business descriptions for a total of 29,807 unique firms or an average of 8,048 unique firms per year over the sample period.

2.2 Patent Descriptions

The Patent Act of 1952 governs patenting in the United States, authorizing the USPTO to issue patents for inventions that are new, useful, and non-obvious (U.S. Congress, 1952). Usefulness is broadly defined under U.S. patent law: the invention must provide a specific, substantial, and credible utility, including a practical application in industry or research, solving a real-world problem, or performing a useful function (USPTO, 2013). Patent

 $^{^{2}}$ For smaller reporting firms, the disclosure requirements are simplified. They must still report their principal products, markets, and competitive conditions, government regulations, environmental compliance, key suppliers, and material customer dependencies, though the level of detail required is reduced (U.S. SEC, 2013, Item 101(h)).

³Form 10-K business descriptions appear in Item 1. Form 20-F business descriptions appear in Item 4, which provides an extensive overview with information on the company's operations, products, markets, raw materials, important dependencies, competitive position, regulatory context, organization structure, and property, plant, and equipment.

⁴The SEC maintains a complete list of all filings on its full-index web page (link).

applications must include a description that establishes the usefulness of the innovation, that is detailed enough to enable any person skilled in the relevant field to make and use the invention. As such, patent descriptions must provide enough detail for reviewers—and later, for researchers like ourselves—to infer whom the invention will benefit and the context in which it will be applied.

We obtain digital records of patent filings made available on the USPTO's PatentsView platform, covering all filings from 1976 to present. The records include detailed patent descriptions and metadata including patent title, assignee, application date and grant date, and Cooperative Patent Classification (CPC) codes. We focus on utility patents granted between 1997 and 2023, totaling 6,226,101 filings over the sample period. For our main analysis, we randomly sample 50,000 patents per year, representing over 20% of the yearly average of 230,596 utility patents in the PatentsView dataset.

2.3 Cooperative Patent Classification

The USPTO has a congressional mandate to maintain a classification system for patents, to facilitate prior art searches and examination (U.S. Congress, 1836; U.S. Congress, 2011). The Cooperative Patent Classification (CPC) system currently serves this purpose; it builds on the International Patent Classification (IPC) system, adopting identical rules and principles but extending the number of classification codes to more than 260,000 (Simmons, 2014; EPO and USPTO, 2017; Lobo and Strumsky, 2019).⁵ Patents are frequently reclassified, and PatentsView maintains both current and historical CPC classifications. We use the current classification to ensure consistency over our sample period, in line with previous studies (Strumsky, Lobo, and Van der Leeuw, 2012; Lobo and Strumsky, 2019).

Following IPC principles, the CPC classifies patents according to either intrinsic function or particular application (WIPO, 2024, paragraphs 81-87), and classifications

⁵The CPC emerged from a collaboration with the European Patent Office (EPO) and was formally adopted by the USPTO in 2015, replacing a legacy system, the United States Patent Classification (USPC), that had been in use since 1836 (Simmons, 2014). The CPC is revised multiple times per year to keep pace with technological change, and Notices of Change are published regularly to the USPTO website (link).

CPC	Number of	Average Annual Patents per Category								
Level	Categories	Median	Median Min		IQR	Skewness	Kurtosis			
Section	8	5,360	356	14,017	$5,\!630$	0.48	-1.03			
Class	121	99	1	$6,\!576$	282	4.43	21.35			
Subclass	577	21	1	$3,\!991$	62	9.15	109.50			
Group	3,887	3	1	1,075	7	11.03	195.72			

Table 1: Descriptive Statistics for Levels of the CPC

Notes. The table describes four levels of the Cooperative Patent Classification (CPC) system: Section, Class, Subclass, and Group. For each level, we report the number of categories, and the minimum, median, and maximum number of patents per category, as well as the interquartile range (IQR), skewness, and kurtosis, as averages of yearly statistics over the sample period, 1997 to 2023. The skewness and kurtosis measures indicate distributional asymmetry and tail heaviness, respectively, particularly at finer levels of classification.

often, but not always, cross industry boundaries. For example, patents for mixing and agitation are grouped together, regardless of whether they are used for washing clothes, mixing paint, or churning butter (Simmons, 2014). Similarly, valves characterized by aspects of their construction are grouped together, independent of the fluid they control; however, a valve specially for use in a heart, as a particular application, would be classified separately (WIPO, 2024, paragraph 85).⁶

The CPC organizes patents into a hierarchy with four main levels: Sections, Classes, Subclasses, and Groups, progressing from broad to narrow categories. Table 1 characterizes these levels. Our sample of 50,000 patents per year includes patents from 8 Sections, 121 Classes, 577 Subclasses, and 3,887 Groups.⁷ Lower levels of the CPC are highly skewed with a few large and many smaller categories. For example, the skewness coefficient rises from 0.48 for Sections to 11.03 for Groups, and the ratio of maximum to median annual patent count rises from 2.62 for Sections to 324.86 for Groups. Similar skew also arises in other patent classification systems and may result from stochastic growth and

⁶The example from Simmons (2014) refers to the legacy USPC system, but carries over to the CPC, where patents for mixing and agitation are classified under Subclass B10F (link). Likewise, the example of values from the WIPO's IPC manual carries over to the CPC, where, for instance, values with pivoted discs or flaps are classified under F16K 1/18 (link) and heart values are classified under A61F 2/24 (link).

⁷The CPC system distinguishes between Main Groups (denoted by classification codes ending in /00) and Subgroups (with additional digits after the slash). Throughout this paper, "Group" refers exclusively to Main Groups, which number 3,887 in our sample. While the complete CPC system contains over 260,000 categories when including all Subgroups, we focus on Main Groups as they provide sufficient granularity for our analysis.

category-splitting dynamics of the system over time (Lafond and Kim, 2019).

2.4 Patent Ownership

We identify patent owners by matching disambiguated assignee names from USPTO patent grants with company names from SEC annual reports using natural language processing. Our approach essentially follows the methodology pioneered by Bound, Cummins, Griliches, Hall, Jaffe, et al. (1982), while adopting recent refinements introduced by Kogan et al. (2017) and Arora, Belenzon, and Sheer (2021). While we adopt a similar name-matching approach to establish initial patent-to-owner links, our work departs from this literature by treating these links as a first step that is followed by a machine learning classification step, rather than an end goal.

We pre-process patent assignee names and company names by standardizing cases, removing non-standard characters and punctuation, and stripping any suffixes that do not aid in matching. We then use exact and fuzzy matching techniques to identify potential matches between the pre-processed names, considering common substrings and other similarity metrics. Initial matches are supplemented with a set of manually curated matches for the most active patenting firms, and the results are iteratively refined to ensure reliability. Only high-confidence matches are retained, with ambiguous cases flagged for manual review.

We identify 712.48 unique patenting firms in an average year and 19,237 in total over the period 1997 to 2023, matching 9,653.00 out of 50,000 sampled patents in an average year and 260,631 patents in total. Only 15.64% of CRSP firms receive patent grants in an average year, and while these patenting firms are often large, they account for just 46.29% of CRSP market capitalization—less than half.⁸ Even among patenting firms, concentration is high: while the average firm receives 13.55 patents annually, the yearly maximum reaches 1,078.37 patents. Figure 2 shows that the low share of patenting firms has remained remarkably stable from 1997 to 2023. These patterns highlight the

⁸The numbers rise to 25.07% of CRSP firms and 54.08% of CRSP market capitalization when firms with patents granted in five-year rolling windows are included. For comparison, Lee et al. (2019) use firm-patent matches from Kogan et al. (2017) and report an annual average of 956 patent owners accounting for 52.56% of CRSP market capitalization over the period 1963 to 2012, similar to what we find.



Figure 2: Prevalence of Patent Owners in CRSP

Notes. The figures show the number and market capitalization of patent-owning firms relative to the number and market capitalization of SEC filing firms in CRSP monthly data from 1997 to 2023. The left figure shows the number of patent owners divided by the number of CRSP SEC filing firms each year, while the right figure shows the total market capitalization of patent owners divided by the total market capitalization of CRSP SEC filing firms each year. Solid lines plot values for patent owners defined as firms with at least one patent granted in a given year, while dashed lines plot values for patent owners defined as firms with at least one patent granted in a backward-looking five-year rolling window.

importance of expanding technology research beyond patent ownership to include the majority of firms that use technology but do not own patents.

Table 2 compares non-patenting to patenting firms by industry group and size class. Industry groups are defined by SIC codes: Resource (0100–1799 and 4900–4999), Manufacturing (2000–3999), Service (4000–4899, 5000–5999, and 7000–8999), and Finance (6000–6399 and 6411). The table excludes real estate, holding companies, public administration, and firms with missing SIC codes. Size classes are based on market capitalization across all industries within each year: Large Cap (top 10%), Mid Cap (next 20%), Small Cap (bottom 70%), and Private (firms with SEC filings but no CRSP data). The table reports firm counts, market capitalization, and industry shares for each category.

Patenting activity varies substantially across industries. Patenting firms account for 74.50% of manufacturing market capitalization, 45.35% of services, and only 18.53% of finance. Over one in three manufacturing firms patent, compared to one in ten service firms and fewer than one in fifty finance firms. Large firms dominate patenting across all industries, accounting for 48.90% of total market capitalization while representing only 2.42% of all firms. These patterns show that focusing only on patenting firms excludes most firms, particularly in service-oriented industries. As we show later, however, non-patenting and patenting firms have surprisingly similar technological profiles within industry and size groups.

Industry	Size	Firm	Count	Firm	ı Size	Industry Share		
Group	Class	NP	Р	NP	Р	NP	Р	
Finance	Large	61.67	8.74	29.20	54.42	61.80	18.12	
	Mid	125.26	2.42	3.05	3.57	13.74	0.38	
	Small	523.96	3.40	0.32	0.57	5.93	0.08	
	Private	414.22	2.44					
	All	1,125.11	16.48	3.03	35.19	81.47	18.53	
Service	Large	87.41	40.37	24.35	59.62	37.38	41.96	
	Mid	237.52	43.93	3.07	3.28	12.35	2.64	
	Small	799.89	98.37	0.37	0.42	4.92	0.75	
	Private	828.33	38.70				—	
	All	1,953.15	221.37	2.92	16.47	54.65	45.35	
Manufacture	Large	49.48	121.52	29.79	46.06	17.67	66.22	
	Mid	155.44	169.41	3.05	3.21	5.36	6.09	
	Small	736.52	471.56	0.29	0.38	2.47	2.19	
	Private	653.26	126.11					
	All	1,594.70	888.59	2.30	8.06	25.50	74.50	
Resource	Large	31.33	7.19	15.19	24.60	50.28	18.33	
	Mid	68.52	7.62	3.13	2.99	22.37	2.39	
	Small	162.78	13.74	0.36	0.39	6.13	0.59	
	Private	415.93	10.41					
	All	678.56	38.67	2.72	6.60	78.78	21.22	
All	Large	229.89	177.81	25.51	49.54	32.38	48.90	
	Mid	586.74	222.81	3.08	3.23	9.65	3.88	
	Small	$2,\!223.15$	586.81	0.33	0.38	3.89	1.30	
	Private	3,121.19	193.07					
	All	6,160.96	1,180.52	2.71	10.10	45.92	54.08	

Table 2: Comparison of Non-Patenting and Patenting Firms by Industry and Size

Notes. The table compares non-patenting firms (NP) and patenting firms (P) by industry group and size class. Firm count is the number of firms of each type. Firm size is the average market capitalization (in millions) for firms of each type. Industry share is the total market capitalization for firms of each type expressed as a percentage of the total market capitalization of each industry group. These statistics are computed by industry group and size class. Industry groups are defined by SIC four-digit codes: Resource includes SIC 0100–1799 and 4900–4999; Manufacture includes SIC 2000–3999; Service includes SIC 4000–4899, 5000–5999, and 7000–8999; and Finance includes SIC 6000–6399 and 6411. We exclude firms with SIC codes 6400–6410, 6412–6499, 6500–6599, 6700–6799, 9000–9999, and firms with missing SIC codes from the table. Size classes are determined by market capitalization across all industries within each year: Large Cap (top 10%), Mid Cap (next 20%), Small Cap (bottom 70%), and Private (firms with SEC filings that do not appear in CRSP). We define patenting firms as firms with patents granted in five-year rolling windows.

3 Estimating Usefulness Probabilities

Using textual similarity scores between business and patent descriptions, we train a binary classifier to estimate the probability that a given patent would, in principle, be useful to a given firm—even if the firm does not own the patent. We refer to the estimates as *usefulness probabilities*. For training, we use a two-stage positive and unlabeled machine learning technique. Despite the growing importance of machine learning in economics and finance, we believe our paper is the first in these fields to use positive and unlabeled learning.

We interpret patent ownership as a (possibly noisy) signal of usefulness. Under this assumption, we assign positive usefulness labels to firm-patent pairs where the firm owns the patent. All remaining firm-patent pairs, which constitute the vast majority of pairs, are left unlabeled. We then identify "reliably negative" pairs from the unlabeled set and assign them negative labels in the first stage. In the second stage, we use logistic regression to train a binary classifier on the positive and reliably negative labeled data. We evaluate our model on a test set of firm-patent pairs that were not used in training, using a performance metric appropriate for positive and unlabeled learning. This approach allows us to bring formal statistical tools to the patent usefulness classification problem, revealing for the first time the technological profiles of non-patenting firms.

3.1 Measuring Document Similarity

The core features used by our classifier to predict usefulness are textual similarity scores between business descriptions and patent descriptions. To obtain these features, we vectorize the text of descriptions using three methods: term frequency (TF), term frequency-inverse document frequency (TF-IDF), and Sentence-BERT (SBERT). TF and TF-IDF are traditional methods that focus on word frequencies within and across documents. In contrast, SBERT is a modern pre-trained transformer model that produces embeddings (numerical vectors) capturing the contextual and semantic meanings of text.

For each of the three document vectorization methods, we compute the cosine similarity

between all possible pairs of business descriptions and patent descriptions. We use the resulting similarity scores as the core features in our positive and unlabeled learning classifier. We estimate the classifier separately for each year, considering only contemporaneous pairs of business and patent descriptions. We discuss document vectorization in more detail in Appendix A.1, and cosine similarity scoring in Appendix A.2. In addition to the similarity scores, we include indicators for SIC Divisions and CPC Sections, as well as textual characteristics as control features in some specifications of our classifier.

Including document-level textual characteristics as control features is important because cosine similarity can be sensitive to superficial textual properties such as document length, writing style, or the presence of non-technological language (Brown and Tucker, 2011). Without accounting for these factors during the classification step, we risk conflating true technological relevance with stylistic or structural artifacts of the documents.

The three textual characteristics we examine are word count (the number of words after removing standard English stop words), corpus overlap (the number of unique words from one corpus that appear more than once in the other corpus), and lexical diversity (measured using the MTLD metric of McCarthy and Jarvis 2010). Figure 3 shows the cross-sectional distributions of these characteristics over time. Unlike traditional firm-level characteristics like market capitalization and industry assignment, which differ markedly between non-patenting and patenting firms as Table 2 shows, these textual characteristics show more moderate differences.

Panel 3a shows textual characteristics of business descriptions for non-patenting firms (left) and patenting firms (right). Averaging across years, we find cross-sectional median word counts of 5,700.7 for non-patenting firms and 7,895.2 for patenting firms—a difference of 38.5%. Corpus overlap medians are 1,309.1 versus 1,632.3—a difference of 24.7%. Lexical diversity medians are 77.0 versus 85.9—a difference of 11.6%. The plots show broadly similar interquartile ranges and 90-10 percentile ranges for the cross-sectional distributions of all three metrics for both firm types, and interquartile ranges within firm type are nearly twice the difference in medians across firm types for all three metrics. In other words, differences within type are much larger than differences across types.



Figure 3: Business and Patent Description Characteristics

(a) Business Descriptions

Notes. The plots in Panels 3a and 3b show the cross-sectional distribution of characteristics for business and patent descriptions, respectively, from 1997 to 2023. Black lines shows the median and shaded regions show the 25-75, 90-10, and 92.5-97.5 quantile ranges of values across descriptions each year. In Panel 3a, the left column shows the characteristics of business descriptions for non-patenting firms, while the right column shows the same characteristics for patenting firms. In Panel 3b, the left column shows the characteristics of patent descriptions for patents with no identified firm owner (these patents may be owned by universities, government, or firms that do not file annual reports with the SEC, including some foreign firms), while the right column shows the same characteristics for patents with identified owners. The plotted characteristics are word count, corpus overlap, and lexical diversity. We define word count as the number of words (including repeated occurrences) in a given description after removing standard English stop words, corpus overlap as the number of unique words in a given description from one corpus that appear more than once anywhere in the second corpus, and lexical diversity as the MTLD metric of McCarthy and Jarvis (2010). Panel 3b shows textual characteristics of patent descriptions for patents without identified owners (left) and with identified owners (right) in our sample of firms.⁹ Averaging across years, we find median word counts of 2,146.4 for patents with identified owners and 2,672.4 for those without—a 24.5% difference. Corpus overlap medians are 424.1 versus 539.1—a 27.1%. Lexical diversity medians are 35.3 versus 35.1—a -0.6% difference. Additionally, as with business descriptions, the cross-sectional distributions are broadly similar for patents with and without identified owners. Furthermore, interquartile ranges within patent type are nearly twice the difference in medians across patent types for all three metrics, indicating again that differences within type are much larger than differences across types.

Business descriptions are substantially longer, contain greater corpus overlap, and exhibit higher lexical diversity than patent descriptions. Both document types show gradual increases in median word count and corpus overlap over time. From 1997 to 2023, business description word counts grew 386.4% for non-patenting firms and 337.4% for patenting firms, while patent word counts grew 155.6% and 156.2% for patents without and with identified owners, respectively. Lexical diversity trends diverged: business descriptions grew by 19.8% (non-patenting) and 18.0% (patenting), while patent descriptions shrunk by -27.1% (no identified owner) and -26.9% (identified owner), respectively.

Importantly, these patterns hold consistently across both non-patenting and patenting firm types, indicating that textual characteristics are not driven primarily by firm type. This similarity matters for our empirical approach: our training data consists of patenting firms paired with patents that have identified owners—though we emphasize that most pairs are patenting firms that do not own the specific paired patent. If non-patenting firms or patents without identified owners had substantially different textual characteristics than the pairs we use for training, our classifier might perform poorly when applied to the broader universe. The evidence in Figure 3 suggests this is not a major concern—the textual features we use for classification are comparable across all firm and patent types,

⁹Patents without identified owners in our sample of firms may be owned by private individuals, governments, or firms that do not file annual reports with the SEC, which would include some small U.S. firms and foreign firms.

both in terms of cross-sectional distributions and time trends.

3.2 Positive and Unlabeled Learning

The machine learning literature has developed several approaches to positive and unlabeled learning, each suited to different assumptions about how positive instances are labeled. Bekker and Davis (2020) review these approaches and summarize the assumptions that underlie them. When all positive instances from a population are equally likely to be labeled positive, irrespective of the characteristics of each instance, the positive instances are said to be selected completely at random (SCAR). If the probability of a positive instance being labeled depends on observable characteristics, the instances are said to be selected at random (SAR). If the labeling mechanism depends on the probability of the instance truly being positive, even controlling for observable characteristics, the instances are said to be selected not at random (SNAR).

Many approaches to positive-unlabeled learning rely on the SCAR assumption, but our setting does not satisfy this assumption. Our positive labels arise from patent grants, and differences between patenting and non-patenting firms are well established. While these differences are less pronounced in our text-based features, they persist—suggesting SAR rather than SCAR conditions. Moreover, the labeling mechanism itself likely depends on the true positive status: firms that receive patent grants are more likely to find those patents genuinely useful, a characteristic of SNAR settings. Given these challenges, the SCAR assumption cannot be justified in our setting.

Instead, we employ a methodology that requires minimal assumptions about the labeling mechanism. Our approach belongs to a class of two-step methodologies that rely on assumptions of separability and smoothness (Bekker and Davis, 2020). Under separability, the classes are assumed to be separable such that a classifier exists that can map positive and negative instances to opposite sides of a decision threshold. Under the smoothness assumption, instances with similar features have similar label probabilities. Under these assumptions, a first-stage classifier can identify reliable negative instances based on their distance from labeled positive instances in feature space. A second-stage





SBERT Cosine Similarity

Notes. The scatter plot shows a class-balanced random sample of 250 positive and 250 unlabeled firmpatent pairs, highlighting the separation between these groups. Firm-patent pairs where the firm owns the patent and our model predicts usefulness are marked as transparent black circles with solid black centers, while those predicted useless are marked as transparent black circles only. Pairs where the firm does not own the patent and our model predicts usefulness are marked as transparent red circles with solid red centers, while those predicted useless are marked as transparent red circles only. The horizontal axis represents SBERT cosine similarity scores, and the vertical axis represents TF-IDF cosine similarity scores. Kernel density plots above and to the right show the class-conditional distributions of each feature for the random sample. While the scatter plot demonstrates the predictive results of the simple two-feature Model 2a in Table 3 to illustrate class separation in two dimensions, our preferred Model 3c in Table 3, which incorporates additional features, achieves substantially better performance.

classifier can then be trained on the labeled positives and the reliable negatives that were identified in the first stage.

Figure 4 illustrates the class separation achievable using just two document similarity measures. The scatter plot shows that labeled positive and unlabeled firm-patent pairs exhibit substantial, though not perfect, separation in the feature space defined by SBERT and TF-IDF cosine similarities. The marginal distributions overlap but show distinct modes for each class–and importantly, since the unlabeled class contains both positive and negative instances, the true separation between positive and negative classes would be even greater than what we observe here. This separation suffices for a two-stage methodology, as Bekker and Davis (2020) note that perfect separation is not required. Rather, there should exist regions where positive density significantly exceeds negative density, enabling the first stage to identify reliable negative examples with high confidence. The clear separation in the marginal distributions, particularly for SBERT, suggests our feature space meets this requirement. Additionally, the clustering of positive examples in the scatter plot, rather than random dispersion among unlabeled examples, supports the smoothness assumption that nearby points share similar class labels. While these two similarity measures alone achieve reasonable separation, our preferred model incorporates additional features and controls for improved performance.

Given this evidence for separability and smoothness, we implement a version of the two-stage spy methodology developed by Liu et al. (2002). In the first of the two stages, we create spies by removing the labels from a subset of positively labeled firm-patent pairs. We assign negative labels to the spies and to all unlabeled observations and train a first-stage classifier on this manipulated data. The first-stage classifier predicts probabilities that we use to identify reliable negatives. Specifically, we mark observations with first-stage probabilities below those of the spies as reliably negative.¹⁰ This first stage filters out unlabeled firm-patent pairs that are likely positive. In the second stage, we train a classifier on the positive and reliably negative pairs.

3.3 Model Training

Our positive and unlabeled data is inherently class imbalanced: there are far fewer firmpatent pairs with positive labels than without labels. The imbalance arises because each patent has only one owner, and we rely on ownership for our positive labeling. This class imbalance can significantly affect our classifier's performance, producing a bias towards

 $^{^{10}}$ The original spy methodology of Liu et al. (2002) uses the minimum positive label probability of the spies to establish the threshold. In our application, for robustness against possible noise in our positive labels, we take the tenth percentile rather than the minimum probability from the set of spies.

predicting the majority class, which in our case are the reliably-negative firm-patent pairs. Because far fewer than one percent of our firm-patent pairs have positive labels, a classifier that never predicts positive labels would achieve near-perfect accuracy. But such a classifier would be worthless.

The machine learning literature has developed a range of sampling and ensemble methods for dealing with class imbalance (Galar et al., 2011). One such method, bootstrap resampling, is particularly well-suited to positive and unlabeled learning, especially in large datasets with few labeled observations (Mordelet and Vert, 2014). The method of bootstrap aggregation ("bagging"), as adapted to the positive-unlabeled setting by Mordelet and Vert, entails repeatedly under-sampling the unlabeled observations to produce a number of re-balanced training samples for an ensemble of classifiers. A simple average of the classifiers in the ensemble can then be used for prediction.

We adopt Mordelet and Vert's bootstrap aggregation procedure, using logistic regression as our core classification model. We prefer logistic regression over alternative models because it provides coefficients that indicate how each feature affects the model's decision, is efficient to train on large datasets, and is familiar to most economists. We form an ensemble of 10 logistic regression classifiers and class-balanced subsets of labeled and unlabeled observations in the bootstrapped training samples.

Table 3 reports log odds ratios and z statistics for three models, where each model is estimated without document controls, with document controls, and with document controls, SIC Division indicators, and CPC Section indicators. Document controls include the word count, corpus overlap, and lexical diversity of each document in a given firmpatent pair. Model 1, which uses only SBERT similarity scores, shows that semantic similarity is a strong predictor of patent usefulness, with odds ratios ranging from 18.36 to 53.83 depending on the specification. Model 2 adds TF-IDF similarity scores, which contribute additional predictive power while moderating the effect of SBERT similarity. In our preferred specification, Model 3c, which incorporates TF, TF-IDF, and SBERT similarity scores, both SBERT and TF-IDF maintain strong positive associations with patent usefulness (odds ratios of 20.17 and 6.56 respectively), while TF shows a weak

Features	Average Odds Ratios										
		Model 1	-		Model 2			Model 3			
	(a)	(b)	(c)	(a)	(b)	(c)	(a)	(b)	(c)		
SBERT	18.36	44.50	53.83	9.14	18.58	22.70	7.14	19.78	20.17		
	(15.20)	(34.38)	(33.94)	(19.06)	(26.31)	(28.89)	(23.54)	(29.27)	(32.07)		
TF-IDF				2.05	3.37	3.46	11.44	6.59	6.56		
				(17.18)	(19.80)	(18.36)	(22.48)	(16.51)	(17.19)		
TF							0.15	0.43	0.42		
							(-22.45)	(-8.70)	(-9.05)		
F_{1c}	0.48	0.61	0.61	0.51	0.65	0.65	0.56	0.65	0.67		
$\operatorname{Precision}_{c}$	0.54	0.59	0.60	0.55	0.62	0.62	0.57	0.62	0.63		
Recall	0.73	0.80	0.80	0.74	0.80	0.81	0.76	0.80	0.81		
Doc Ctrls	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes		
Sic Ctrls	No	No	Yes	No	No	Yes	No	No	Yes		
Cpc Ctrls	No	No	Yes	No	No	Yes	No	No	Yes		
Penalty	L2	L2	L2	L2	L2	L2	L2	L2	L2		
Obs/Reg	15489	15480	15470	15472	15485	15465	15486	15477	15515		
$\operatorname{Reg}/\operatorname{Ens}$	10	10	10	10	10	10	10	10	10		
RRS	2	2	2	2	2	2	2	2	2		

Table 3: Logistic Regression on Positive and Unlabeled Data

Notes. The table shows alpha estimates for low and high decile technology momentum portfolios upper panel of the table shows odds ratio estimates from three logistic regression models trained on positive and unlabeled data, with z-statistics reported in parentheses under each estimate. Model 1 uses SBERT, Model 2 uses SBERT and TFIDF, and Model 3 uses SBERT, TFIDF, and TF similarity scores as features. We estimate each model without controls (columns a), with document controls (columns b), and with document, SIC, and CPC controls. All models are estimated with L2 ridge penalties applied.

negative association. The addition of document, industry, and patent category controls improves model performance across all specifications, with model evaluations rising substantially from the baseline Model 1a to our preferred Model 3c.

3.4 Model Evaluation

Evaluating classifiers trained on positive and unlabeled data presents unique challenges that limit the informativeness of traditional performance metrics. In standard supervised learning, classifier performance is typically assessed using metrics derived from the confusion matrix—a matrix containing true positives (TP), false positives (FP), true negatives (TN), and false negatives (FN). Two fundamental metrics are precision and recall. For the positive class, precision is defined as TP/(TP + FP), and measures the fraction of positive predictions that are correct. Again for the positive class, recall is defined as TP/(TP + FN), and measures the fraction of actual positive cases that are correctly identified. Both metrics also have a symmetric definition for the negative class. The F_1 score, given by 2 × (precision × recall)/(precision + recall), provides a single summary measure that balances these two objectives.

These traditional metrics must be interpreted with caution in the positive-unlabeled learning context. While recall can be reliably estimated using only positive examples, precision requires false positives—positive predictions on truly negative instances. In positive and unlabeled learning, we fundamentally cannot observe false positives since unlabeled examples may be either positive or negative. If traditional classifier performance metrics are applied to non-traditional classifiers, positively-predicted unlabeled instances are treated as false positives (the standard approach) and understate precision.¹¹

This understatement is particularly severe in settings with significant class imbalance. Consider a classifier that correctly identifies a labeled positive instance but also predicts positive for three out of 1,000 unlabeled instances. Even if all three of these "false positives" are actually correct predictions of unlabeled positive cases, treating them as errors results in a precision of 25%. Thus, even strong classifier performance can appear poor when evaluated using traditional precision with positive and unlabeled data in a class-imbalanced setting.

To partially address the problem of class imbalance in our model evaluation, we follow Siblini et al. (2020) and "calibrate" our precision and F_1 scores to a reference class ratio of 1, making them more interpretable. This calibration does help with class imbalance, but does not correct for mistaken "false positives." We therefore rely primarily on a robust metric that, while less interpretable, is suitable for determining the optimal decision threshold for our classifier—that is, the threshold value above which predicted probabilities are classified as positive. Rather than choosing the threshold to maximize the calibrated F_{1c} score, which remains problematic in our setting, we maximize a modified

¹¹Under the SCAR assumption, analytical corrections can be made (Elkan and Noto, 2008), but we do not make the SCAR assumption in our setting, as explained above.

	Average Classifier Performance								
	$\operatorname{Precision}_{c}$	Recall	F_{1c}	Support					
Unlabeled Class	0.98 (0.00)	0.90 (0.01)	0.94 (0.01)	1,422,892 (33)					
Labeled Class	(0.00) (0.28) (0.02)	(0.01) (0.72) (0.03)	(0.01) (0.39) (0.02)	(33)					
Weighted Average	0.98 (0.00)	0.90 (0.01)	0.94 (0.01)	1,424,812 (0)					
Macro Average	0.63 (0.01)	0.81 (0.01)	0.67 (0.01)	1,424,812 (0)					

Notes. The tables shows traditional performance metrics for our classifier trained on positive and unlabeled data. In the upper panel of the table, we report average annual precision, recall, and F_1 performance metrics computed from individual classifiers trained on yearly patent and business description data. In the lower panel, we report weighted and unweighted (macro) averages across classes of the average annual class-specific performance metrics. To account for the substantial class imbalance in our positive and unlabeled data, we report calibrated precision and F_1 scores, indicated by the subscript $_c$ in the first and third columns.

version of the performance measure proposed by Lee and Liu (2003),

Modified Lee-Liu Score:
$$\lambda_{\gamma} = \frac{r^{\gamma}}{Pr(\hat{y}=1)}$$
, (1)

where r is recall, $Pr(\hat{y} = 1)$ is the fraction of instances classified as positive, and the parameter γ can be adjusted to place greater emphasis on recall. The threshold λ_{γ} can be reliably estimated from positive and unlabeled data, making it suitable for positiveunlabeled learning. We set $\gamma = 3$ to reflect our preference for higher recall.¹²

We present traditional evaluation metrics (with the caveats noted above) in Table 4. These metrics are computed using repeated random sub-samples (RRS) of our training data. The RRS procedure involves repeatedly re-partitioning the data into training and test sets, training a new model on each training partition, and evaluating it on the corresponding test partition. This approach provides more robust performance estimates than a single train-test split. Within each year, we compute performance metrics and

¹²Lee and Liu (2003) show that their original metric with $\gamma = 2$ is proportional to the product of precision and recall $(r \times p)$. Since the F_1 score is the harmonic mean of precision and recall, their metric captures similar information while being computable in the positive-unlabeled setting. Our choice of $\gamma = 3$ places more emphasis on recall, which is conceptually similar to using an F_β score with $\beta > 1$.

their standard errors across RRS iterations, then average both the metrics and standard errors across sub-samples and years. After the model evaluation stage, we use all available data to train a model for final predictions.

For the unlabeled class, our classifier achieves a calibrated precision of 0.98 and a recall of 0.90, yielding a calibrated F_{1c} score of 0.94. These results reflect the extreme class imbalance in our data—with such a low base rate of positive cases, even a naive classifier that predicts "unlabeled" for all observations would achieve near-perfect performance here.

The more challenging metrics are those for the labeled class, where our classifier achieves a calibrated precision of 0.28 and recall of 0.72. The low precision understates true performance—without true negative examples, many "false positives" are likely correct predictions of technological usefulness not captured in our labeled set. The higher recall indicates that our classifier identifies a substantial majority of labeled instances where a patent is owned by a particular firm and is therefore likely to be useful to the firm. Recall may also understate true performance, if our positive labels are noisy and contain instances of useless patents owned by firms.

The support column in Table 4 shows the scale of our evaluation, with metrics computed over more than 1.4 million firm-patent pairs per year. The standard errors, shown in parentheses, are computed across RRS iterations within each year and then averaged across years. These errors indicate that our performance estimates are stable across different random sub-samples. The substantial difference between weighted and macro averages (0.94 versus 0.67 for F_{1c}) shows how class imbalance affects the traditional performance metrics even after calibration.

The ultimate test of our estimates lies in their ability to predict or explain economic phenomena. Before turning to economic applications, however, we first examine the usefulness probabilities estimated by the model, and use the estimates to compare the technological associations of non-patenting and patenting firms. Despite the large number and economic importance of non-patenting firms, their technological profiles have remained largely unexplored by researchers due to data limitations. Our estimates remove these limitations, revealing surprising similarities between the technological associations of non-patenting and patenting firms.

3.5 Model Predictions

We use Model 3c from Table 3 to predict firm-patent associations, and plot the unconditional results of these predictions for all years in Figure 5. The figure displays the distribution of predicted usefulness probabilities for non-patenting firms (left) and patenting firms (right). The upper panels show probability density functions (PDFs) with annual distributions plotted as thin gray lines and their pointwise average as a thick black line. The lower panels show cumulative distribution functions (CDFs) with the same line conventions. Thin dashed lines mark the 50th and 90th percentiles, while the thin solid vertical line indicates the classification threshold that maximizes the modified Lee-Liu score λ_{γ} , which equals 0.82 when averaged across years.

Without conditioning on industry or firm size, we find that non-patenting firms have lower probabilities of associating with patents overall. The median predicted usefulness probability for non-patenting firms is 0.06, compared to 0.1 for patenting firms. Similarly, the 90th percentile values are 0.55 and 0.76 respectively, indicating that non-patenting firms have lower predicted associations than their patenting counterparts even at the upper end of the distribution. However, these raw differences primarily reflect the distinct industry and size compositions of the two firm types rather than fundamental differences in technological associations, as we emphasize in the following section.



Figure 5: Annual Distributions of Predicted Usefulness Probabilities

Notes. The figure shows the distribution of predicted usefulness probabilities for non-patenting firms (left) and patenting firms (right). Distributions are estimated from 1% samples of firm-patent usefulness probabilities drawn from each firm's set of usefulness probabilities each year, using kernel density functions with reflecting barriers at 0 and 1. In the upper panels, annual PDFs are plotted as thin gray lines, and the pointwise average of annual PDFs is plotted as a thick black line. In the lower panels, annual CDFs are plotted as thin gray lines, and the pointwise average of annual S0th and 90th quantile values are indicated by thin dashed lines. The average of the annual classification threshold that maximizes the modified Lee-Liu score is indicated by a solid vertical line.

4 Characterizing Technological Associations

Having estimated usefulness probabilities for all firm-patent pairs, we now examine how technological associations vary between non-patenting and patenting firm types. We examine these associations at two levels. First, we examine aggregated patterns by studying the intensity of technological associations between industries and technology categories by firm type. Second, we examine firm-level patterns by studying the technological portfolios of individual firms. At both levels, we find surprising similarities between non-patenting and patenting firms. Cross-sectional distributions of several measures of technological association are similar across firm types, particularly in their central tendencies. Indeed, we find that within-type differences are much larger than between-type differences, after conditioning on industry and firm size.

4.1 Aggregated Technological Associations

We use Sankey diagrams to illustrate the intensity of technological associations between industries and technology categories. We define the aggregated intensity of technological association between SIC Division s and CPC Section c for firm type $\tau \in \{NP, P\}$ as:

$$I_{sc}^{(\tau)} = \frac{A_{sc}^{(\tau)}}{F_s^{(\tau)} \times P_c} \tag{2}$$

where $A_{sc}^{(\tau)}$ is the number of firm-patent associations between Division s and Section c for firm type τ , $F_s^{(\tau)}$ is the number of firms of type τ in Division s, and P_c is the number of patents in Section c. This intensity measure normalizes raw association counts by the total number of possible firm-patent pairs in each industry-technology combination, providing a measure that is adjusted for compositional differences in firm and patent distributions across industries and technology categories. The Sankey diagrams in Figures 6 and 7 show the aggregated intensity-based associations for non-patenting and patenting firms, respectively. We provide Sankey diagrams based on raw association counts in Appendix B.1.

We draw two conclusions from the intensity-based Sankey diagrams. First, each SIC Division shows substantial associations with multiple CPC Sections, suggesting that technologies frequently cross industry boundaries. Second, the patterns for non-patenting firms in Figure 6 closely resemble those for patenting firms in Figure 7. While Sankey diagrams based on raw association counts in Appendix B.1 show greater differences, the intensity-based view reveals strong similarities after adjusting for compositional effects. That said, some differences between the intensity-based diagrams for non-patenting and patenting firms are apparent upon close inspection.

To facilitate the direct comparison of aggregated technological associations between non-patenting and patenting firms, we construct a heatmap representation in Figure 8 that quantifies differences between the Sankey diagrams in Figures 6 and 7 for each industry-technology combination. The heatmap displays a relative intensity metric R_{sc} ,



Figure 6: CPC-SIC Sankey Diagram: Intensity-Based, Non-Patenting Firms

Notes. The figure shows a Sankey diagram of associations between CPC Sections (left) and SIC Divisions (right). Flows represent intensities of association, defined as the number of the number of associated firm-patent pairs for a given CPC-SIC combination, *relative* to the number of possible firm-patent pairs for that combination, aggregated over all sample years. CPC Sections: Human Necessities (A), Performing Operations; Transporting (B), Chemistry; Metallurgy (C), Textiles; Paper (D), Fixed Constructions (E), Mechanical Engineering; Lighting; Heating; Weapons; Blasting (F), Physics (G), and Electricity (H). SIC Divisions: Agriculture, Forestry, and Fishing (0100–0999), Mining (1000–1499), Construction (1500–1799), Manufacturing (2000–3999), Transportation, Communications, Electric, Gas and Sanitary Service (4000–4999), Wholesale Trade (5000–5199), Retail Trade (5200–5999), Finance and Insurance (6000–6799, excl 6500–6599 and 6700–6799), and Services (7000–8999). Some CPC and SIC names have been shortened for the figure.



Figure 7: CPC-SIC Sankey Diagram: Intensity-Based, Patenting Firms

Notes. The figure shows a Sankey diagram of associations between CPC Sections (left) and SIC Divisions (right). Flows represent intensities of association, defined as the number of the number of associated firm-patent pairs for a given CPC-SIC combination, *relative* to the number of possible firm-patent pairs for that combination, aggregated over all sample years. CPC Sections: Human Necessities (A), Performing Operations; Transporting (B), Chemistry; Metallurgy (C), Textiles; Paper (D), Fixed Constructions (E), Mechanical Engineering; Lighting; Heating; Weapons; Blasting (F), Physics (G), and Electricity (H). SIC Divisions: Agriculture, Forestry, and Fishing (0100–0999), Mining (1000–1499), Construction (1500–1799), Manufacturing (2000–3999), Transportation, Communications, Electric, Gas and Sanitary Service (4000–4999), Wholesale Trade (5000–5199), Retail Trade (5200–5999), Finance and Insurance (6000–6799, excl 6500–6599 and 6700–6799), and Services (7000–8999). Some CPC and SIC names have been shortened for the figure.





Notes. The figure shows the relative intensity metric defined in equation (3) for technological associations between CPC Sections and SIC Divisions. Red cells indicate stronger associations for non-patenting firms, gray cells indicate stronger associations for patenting firms, and white cells indicate equal associations. CPC Sections: Human Necessities (A), Performing Operations; Transporting (B), Chemistry; Metallurgy (C), Textiles; Paper (D), Fixed Constructions (E), Mechanical Engineering; Lighting; Heating; Weapons; Blasting (F), Physics (G), and Electricity (H). SIC Divisions: A (Agriculture, Forestry, and Fishing: 0100–0999), B (Mining: 1000–1499), C (Construction: 1500–1799), D (Manufacturing: 2000–3999), E (Transportation, Communications, Electric, Gas and Sanitary Services: 4000–4999), F (Wholesale Trade: 5200–5999), H (Finance, Insurance, and Real Estate: 6000–6799, excl 6500–6599 and 6700–6799), I (Services: 7000–8999).

which we define as

$$R_{sc} = \frac{I_{sc}^{(N)} - I_{sc}^{(P)}}{I_{sc}^{(N)} + I_{sc}^{(P)}},$$
(3)

where superscripts N and P denote non-patenting and patenting firms, respectively. This bounded measure ranges from -1 to +1, with positive values (red shading) indicating stronger technological associations for non-patenting firms and negative values (grey shading) indicating stronger associations for patenting firms.

The heatmap reveals several systematic differences between the intensity of technological associations for non-patenting and patenting firms. We caution, however, that some of the strongest differences arise in industry-technology combinations with overall low intensities. For example, the services industry shows strongly positive relative intensity values in fixed constructions and mechanical engineering, indicating that non-patenting service firms associate more intensively with these technologies than patenting service firms. However, the service industry has low-intensity associations with these technologies overall. Conversely, chemical and metallurgical technologies show consistently negative values across most industries, indicating that patenting firms associate more strongly with these technologies than non-patenting firms; this is especially true in agriculture, which shows a high-intensity association with chemistry.

Overall, a nuanced picture emerges from the aggregated results, where technologies frequently cross industry boundaries, and industry-technology associations on an intensity basis are surprisingly similar for non-patenting and patenting firms, with some important differences between firm types in isolated cases. Next, we turn to firm-level evidence.

4.2 Firm-Level Technological Associations

We consider three metrics that characterize firm-level technological associations. To assess technological breadth and depth, we count each firm's associated patents and CPC categories. To assess technological instability over time, we compute add and drop rates for each firm's associated CPC categories. To assess technological generality, we compute the average number of industries associated with each firm's associated patents and CPC categories. We compute metrics at the patent level and at the CPC Group, Subclass, Class levels. This approach allows us to assess firms' associations with increasingly broad technology categories. For each metric at each level, we find that differences within non-patenting and patenting firm types exceed differences between firm types.

Categorical Associations. Our positive and unlabeled learning framework produces usefulness probabilities that associate individual firms with individual patents. We use a binomial testing framework to convert these firm-patent associations into probabilistic firm-category, patent-industry, and category-industry associations, which we require in order to compute technological breadth and depth, instability, and generality metrics at each level of the CPC system.

To associate firms with CPC categories, we compare the count of a firm's patent associations within the category, relative to the count of all patents in the category, with the count of all firms' patent associations within the category, relative to the count of all firms multiplied by the count of all patents in the category. The binomial probability of A_{ic} associations between firm *i* and patents in CPC category *c* is given by

$$\pi_{ic} = \begin{pmatrix} P_c \\ A_{ic} \end{pmatrix} \pi_{0c}^{A_{ic}} (1 - \pi_{0c})^{P_c - A_{ic}}, \quad \text{with} \quad \pi_{0c} \equiv \sum_i \frac{A_{ic}}{F \times P_c}, \tag{4}$$

where F is the number of firms, P_c is the number of patents in category c, and π_{0c} is the baseline probability of a firm-patent association in category c. Under the null hypothesis, $\pi_{ic} = \pi_{0c}$ and firm i is no more likely to associate with patents in category c than the average firm. We test the alternative hypothesis that $\pi_{ic} > \pi_{0c}$, and associate firms with categories when we fail to reject the alternative hypothesis at the 5% significance level.

This approach provides a statistical framework for identifying when a firm's relationship with a technology category is unlikely to have arisen by chance. It adjusts for CPC category size, requiring firms to have more patent associations in larger categories. And, because it uses category-specific baseline probabilities, it accounts for systematic differences in the number of associations across categories while maintaining a consistent threshold for statistical significance.

We adopt a similar procedure to associate patents with four-digit SIC industries. The binomial probability of A_{js} associations between patent j and firms in SIC industry s is given by

$$\pi_{js} = \begin{pmatrix} F_s \\ A_{js} \end{pmatrix} \pi_{0s}^{A_{js}} (1 - \pi_{0s})^{F_s - A_{js}}, \quad \text{with} \quad \pi_{0s} \equiv \sum_j \frac{A_{js}}{F_s \times P}, \tag{5}$$

where F_s is the number of firms in industry s, P is the number of patents, and π_{0s} is the baseline probability of a firm-patent association in industry s. Under the null hypothesis, $\pi_{js} = \pi_{0s}$ and patent j is no more likely to associate with firms in SIC industry s than the average patent. We test the alternative hypothesis that $\pi_{js} > \pi_{0s}$, and associate patents with industries when we fail to reject the alternative hypothesis at the 5% significance level. We use the same procedure to associate CPC categories with four-digit SIC industries, replacing patent j with category c and number of patents P with number of categories Cin the binomial probability calculation. Non-Patenting and Patenting Firm Comparisons. With the firm-level category and industry associations in place, we can compute metrics for technological breadth and depth, instability, and generality at the patent level and at the level of CPC Classes, Subclasses, and Groups. We present the results in Figure 9, which plots the cross-sectional distribution of each metric, separately for non-patenting and patenting firms, conditional on broadly-defined industry group and size class. The industry groups (finance, service, resource, and manufacture) and size classes (private, small, medium, and large) are those used throughout the paper, and defined in Section 2.

Within industry group and size class, distributions for non-patenting and patenting firms are compared on a common scale, but the scales do differ across industry groups and size classes. We deliberately omit numerical values from the figure, to avoid distracting from the main insight: conditional on industry group and size class, the differences within non-patenting and patenting firm types exceed differences between firm types for each metric at each CPC level. We report extensive tabular results with statistical tests for differences in means and medians between firm types in Appendix B.

Panel 9a plots cross-sectional distributions of technological breadth and depth, measured at the firm level as the count of patents and CPC categories with which firms associate. We tabulate these results in Appendix B.2. While we find large and statistically significant differences in counts between industry groups, we find smaller differences between size classes within industry group, and yet smaller differences between nonpatenting and patenting firms within industry group and size class. While differences between median non-patenting and patenting firms remain statistically significant after conditioning on industry and size, these differences are much smaller than the interquartile range of firm-level values found within the non-patenting and patenting firm types.

We interpret counts at the patent level as measures of technological depth, and do find consistently higher median counts for patenting firms at this level, in particular outside of finance. As we consider increasingly broad CPC categories, we interpret the counts as measures of technological breadth, and find that median counts for non-patenting firms frequently exceed median counts for patenting firms. Overall, the results suggest deeper

Figure 9: Firm-Level Technological Associations Within Industry Group and Size Class(a) Technological Breadth and Depth

Class				
Subclass				
Group	Ţ ₽ ₽ ₽ ₽₽ ₽₽	$\begin{bmatrix} T \top & T \top & T \\ \bullet \bullet & \bullet & \bullet \\ \bullet \bullet & \bullet & \bullet & \bullet \\ \bullet \bullet & \bullet &$	$\begin{array}{c} \mathbf{T}_{T} & \mathbf{T}_{T}^{T} & \mathbf{T}_{T}^{T} & \mathbf{T}_{T}^{T} \\ \mathbf{T}_{T}^{T} & \mathbf{T}_{T}^{T} & \mathbf{T}_{T}^{T} & \mathbf{T}_{T}^{T} \end{array}$	$\begin{array}{cccccccccccccccccccccccccccccccccccc$
Patent		$\begin{array}{cccccccccccccccccccccccccccccccccccc$		
(b) Technolo	gical Instability			
Class	T TT TT •• •• ••	$\begin{array}{cccccccccccccccccccccccccccccccccccc$	I TITIT I⊒ II II II I⊒ II II II	I I I II II II II II II II II
Subclass		$\begin{array}{cccccccccccccccccccccccccccccccccccc$	T TT TT TT TT TT TT	
Group				
(c) Technolog	gical Generality			
Class				$\begin{array}{cccccccccccccccccccccccccccccccccccc$
Subclass	$\begin{array}{c} \mathbf{T} \\ $			TI II II II II II II II
Group				$\begin{array}{cccccccccccccccccccccccccccccccccccc$
Patent	$\underline{\underline{\mathbf{H}}} = \underline{\underline{\mathbf{H}}} = \underline{\underline{\mathbf{H}}} = \underline{\underline{\mathbf{H}}} = \underline{\underline{\mathbf{H}}}$			$\begin{array}{cccccccccccccccccccccccccccccccccccc$
	P S M L	P S M L	P S M L	P S M L Manufacture
	гшансе	Dervice	nesource	manufacture

Notes. The figure shows cross-sectional distributions of three metrics characterizing firms' technological associations within industry group and size class. Panel 9a shows technological breadth and depth measured by counts of patent and CPC category associations. Panel 9b shows technological instability measured by CPC category churn rates (the average of add and drop rates). Panel 9c shows technological generality measured by cross-industry usage (the average number of industries associated with each firm's associated patents or CPC categories). Box plots show medians, interquartile ranges, and the 10th and 90th percentiles, separately for non-patenting firms (dark) and patenting firms (light). Distributions for non-patenting firms are compared on a common scale within industry group and size class, but scales do differ across industry groups and size classes. Industry groups and size classes are defined in Section 2. We report extensive tabular results with statistical tests for differences in means and medians between firm types in Appendix B.

technological associations for patenting firms, and broader technological associations for non-patenting firms outside of finance.

Panel 9b plots cross-sectional distributions of technological instability, measured at the firm level as the rate of churn in firms' associations with CPC categories over time. We tabulate these results in Appendix B.3. For each firm-year, we first calculate two measures, the add rate and the drop rate of CPC categories, defined as the percentage of CPC categories added to or dropped from a firm's associations in year t + 1 relative to the average number of the firm's associated categories across years t and t + 1. We then define the churn rate as the average of add and drop rates. As with the breadth and depth metrics, we find that differences in the median measure of technological instability between non-patenting and patenting firm types, conditional on industry group and size class, are substantially smaller than interquartile ranges of firm-level values found within either firm type.

That said, we do find statistically significant differences in median instability measures between non-patenting and patenting firm types, conditional on industry group and size class, with non-patenting firms generally showing higher degrees of technological instability. Differences are generally more pronounced at the level of narrowly-defined CPC Groups, and less pronounced for broadly-defined CPC Classes. We note, however, that exceptions can be found; see, for example, private financial firms or large resource firms. Overall, our results suggest that non-patenting firms have more flexible technological portfolios, potentially adapting more quickly to changing technological opportunities. This interpretation would align with our finding that non-patenting firms maintain broader but shallower technological portfolios.

Panel 9c plots cross-sectional distributions of technological generality, measured at the firm level as the average number of four-digit SIC industries associated with each firm's associated patents or CPC categories. We tabulate these results in Appendix B.4. For example, a firm with associated patents A and B would have a generality value of two at the patent level if patent A were associated with one industry and patent B were associated with three industries (2 = (1+3)/2). The generality metric quantifies the extent to which

firms' associated technologies are in general use across multiple industries.¹³ Once again, differences in technological generality are larger within firm type than between firm types, conditional on industry group and firm size, though statistically significant differences across types do exist.

For CPC Classes and Subclasses, which are broadly-defined technology categories, we find that median patenting firms tend to associate with more general technologies than median non-patenting firms, particularly in services. At the CPC Group level, we find no clear pattern across industry groups and size classes. At the patent level, the pattern reserves, and we find that median non-patenting firms associate with more general technologies than median patenting firms. This result suggests that patenting firms associate with more specific patents, but technology categories with broader cross-industry appeal.

Taken together, these patterns in firm-level technological associations suggest that patenting and non-patenting firms differ in systematic ways: patenting firms exhibit deeper, more stable, and more specific technological focus than non-patenting firms, with some exceptions in particular industry groups and size classes. However, while these differences are statistically significant, they are small relative to the within-type variation across all three metrics, after conditioning on industry group and size class. Ultimately, we find modest differences and surprising similarities in the technological associations of non-patenting and patenting firms.

5 Technological Momentum

We now explore a first application of our patent usefulness probabilities to asset pricing, constructing technological momentum portfolios that include, for the first time, a large set of publicly-traded non-patenting firms whose technological profiles have previously been inaccessible to researchers. Our approach extends recent work by Lee, Sun, Wang, and

¹³This measure of technological generality relies on the assumption that four-digit SIC industries are defined with equal granularity in all parts of the economy, which may not hold in practice. However, by comparing non-patenting firms with patenting firms within broad industry groups and size classes, we mitigate problems arising from violations of this assumption.

Zhang (2019) and Bekkerman, Fich, and Khimich (2023) documenting the profitability of technological momentum strategies applied to patenting firms. These authors argue that technological momentum works because markets are slow to process technological information, especially for technologically intensive firms with limited investor attention. Non-patenting firms use technology as intensively as patenting firms but are typically smaller, with less analyst coverage and more opaque technological profiles. We therefore expect the technological momentum strategy to be particularly effective when extended to these firms.

Our methodology differs from prior approaches in two key ways. First, we expand coverage to include a majority of publicly-traded firms that do not patent and were previously excluded. Second, we measure technological similarity based on firms' exposure to useful technology rather than ownership of intellectual property. This distinction matters because firms often patent for strategic reasons unrelated to their core technological activities (Cohen, Nelson, and Walsh, 2000). Overall, these differences allow us to capture broader technological relationships for a wider set of firms.

5.1 Methodology

Technological momentum strategies exploit predictable patterns in how technology-related information affects stock prices. The strategy identifies firms with similar technological profiles, then takes long positions in firms whose technological peers recently performed well and short positions in firms whose peers performed poorly. Prior research demonstrates that technological peer performance predicts future returns, generating significant alpha.

In implementations of this strategy by Lee et al. (2019) and Bekkerman, Fich, and Khimich (2023), technological similarity is measured using the patent portfolios of patenting firms. Lee et al. (2019) computes the distribution of patents owned by patenting firms across USPTO classes, and correlates these distributions for pairs of patenting firms to assess firm-to-firm technological similarity, following Jaffe (1986). Bekkerman, Fich, and Khimich (2023) measure patent-to-patent textual similarity as the cosine similarity between TF-IDF vectors for each patent, and then average these over the patents owned by patenting firms to assess firm-to-firm technological similarity. Both methods consider only patenting firms, excluding the majority of publicly-traded firms from the analysis.

To extend the technological momentum strategy to this excluded majority, we use our estimated usefulness probabilities to identify technological peers. These probabilities are available for all U.S. public firms and an annual sample of 50,000 patents, covering the period from 1997 to 2023. As noted in Section 2.4, only 15.64% of CRSP firms receive patent grants in an average year, accounting for 46.29% of CRSP market capitalization (or 25.07% of CRSP firms accounting for 54.08% of CRSP market capitalization, when counting patents in five-year rolling windows), so the expansion to non-patenting firms significantly improves coverage.

Because we estimate similarity scores between each firm and each patent each year, we can assess firm-to-firm technological similarity by computing the cosine similarity between vectors of patent-level usefulness probabilities directly. With 50,000 entries, each vector constitutes a highly granular technological profile of each firm. The usefulness probabilities in most of these entries lie below the classification threshold that maximizes the modified Lee-Liu score from the positive and unlabeled learning step.¹⁴ We set usefulness probabilities below the threshold to zero before the similarity calculation.

For focal firm i and month t, we calculate the weighted average return on a portfolio of the focal firm's top 100 technologically similar peers,

$$\text{TECHRET}_{it} = \frac{\sum_{j \neq i} w_{ijt} \times \text{RET}_{jt}}{\sum_{j \neq i} w_{ijt}}$$
(6)

where RET_{jt} is the return of firm j in month t and w_{ijt} is the lagged technological cosine similarity between firms i and j. We lag technological cosine similarity weights to avoid look-ahead bias. Specifically, the weight w_{ijt} used in month t is the firm-to-firm technological similarity computed from patent grants and annual reports in year $y_t - 1$, where y_t denotes the year in which month t occurs. Since annual reports are filed once per year and our similarity measures are calculated annually, all months within a given calendar year use the same set of lagged similarity weights from the previous year's annual

¹⁴The threshold varies from year to year, averaging 0.82 across all years, 1997 to 2023.



Figure 10: Technology Momentum Decile Transition Matrices

Notes. The figure shows monthly transition probabilities for firms sorted into deciles based on the prior month's stock market performance of technologically-similar peers, for non-patenting firms (left) and patenting firms (right). The transition probabilities are estimated using monthly returns data over the period 1998 to 2023. The portfolios are formed on the universe of CRSP firms with available annual reports, after a sample selection rule is applied. Technological similarity is measured using the cosine similarity between vectors of each firm's patent usefulness probabilities. Each firm's peer group is defined as the 100 most technologically-similar firms. The transition probabilities reflect the likelihood of a firm moving between decile portfolios from one month to the next, indicating the persistence of performance for decile portfolios formed on technological momentum.

reports and patent grants. For example, portfolio returns calculated in any month of 2020 use similarity weights computed from 2019 annual reports and patent grants.

We then sort firms into deciles based on lagged TECHRET values from the previous month, and form portfolios that go long the top decile (firms whose technological peers performed well in the previous month) and short the bottom decile (firms whose technological peers performed poorly). These portfolios are rebalanced monthly to maintain either equal or value weights. To be included in a decile portfolio, firms must have filed an annual report in the previous year and, to ensure tradability, must be above the 10th percentile of CRSP firms' one-month lagged prices, market values, and trading volumes.

Figure 10 shows the transition matrices for technological momentum deciles for nonpatenting firms (left) and patenting firms (right). Each cell (m, n) represents the probability of a firm moving from decile m in the current month to decile n in the next month. The diagonal elements indicate performance persistence—firms remaining in the same decile



Figure 11: Monthly Alpha of Technology Momentum Decile Portfolios

Notes. The figure shows monthly alpha estimates in percentage points for equal-weighted and valueweighted decile portfolios formed on the recent performance of technological peers, where L represents the lowest decile and H the highest. We estimate alpha using a four-factor model that includes market, size, value, and cross-sectional momentum factors. Lines extending from each bar show 95% confidence intervals. The portfolios are formed on the universe of CRSP firms, after a sample selection rule is applied. Portfolios are rebalanced monthly over the period January 1998 to December 2023 based on technology momentum decile rankings.

from one month to the next. The transition matrices show that non-patenting firms are substantially more likely than patenting firms to remain in the highest and lowest deciles from one month to the next. This result suggests that technological information diffuses more slowly for non-patenting firms, creating more persistent return predictability. As we demonstrate next, this improved predictability translates into improved performance for the technological momentum strategy.

5.2 Empirical Results

Figure 11 presents the monthly alphas for technological momentum portfolio excess returns across deciles, reported separately for non-patenting and patenting firms and for equal and value-weighted portfolios. The alphas are estimated from a four-factor model that controls for exposure to market, size, value, and cross-sectional momentum factors.¹⁵ For

¹⁵These four factors, commonly denoted MKT, SMB, HML, and MOM, were downloaded from the data library on Ken French's website (link) on 10 March 2024.

equal-weighted portfolios, the results reveal a striking monotonic pattern across deciles, with higher technological momentum deciles consistently earning higher risk-adjusted returns. This pattern is particularly pronounced for non-patenting firms, where the spread between high and low deciles is substantially larger than for patenting firms.

Table 5 reports the factor loadings for the technological momentum strategy. The high-minus-low portfolio loads negatively on the market factor and positively on the size factor for both non-patenting and patenting firms. Importantly, these factor exposures do not explain away the substantial alphas generated by the strategy, suggesting that the strategy represents a distinct anomaly not captured by standard risk factors, consistent with the findings of both Lee et al. (2019) and Bekkerman, Fich, and Khimich (2023). Table B43 in Appendix B.5 demonstrates that these results are robust to alternative asset pricing models, with monthly alphas for non-patenting firms ranging from 1.79% to 2.48% depending on the specification. The technological momentum results remains economically and statistically significant across all model specifications.

	Non-Patenting Firms							Patenting Firms				
	Decile	Alpha	MKT	SMB	HML	MOM	All	pha	MKT	SMB	HML	MOM
Equal- Weighted	High	$0.89 \\ (3.75)$	1.03 (18.43)	0.97 (12.62)	-0.04 (-0.59)	-0.09 (-1.85)	0 (2	.74 .24)	0.91 (11.67)	1.39 (12.95)	-0.55 (-5.68)	-0.11 (-1.52)
	Low	-1.24 (-4.66)	1.16 (18.65)	0.83 (9.61)	-0.10 (-1.26)	-0.27 (-4.86)	$-0 \\ (-1)$.28 .04)	1.23 (19.27)	0.70 (7.96)	-0.19 (-2.33)	-0.22 (-3.80)
	High–Low	$1.97 \\ (4.74)$	-0.13 (-1.36)	0.15 (1.09)	$0.06 \\ (0.45)$	0.18 (2.02)	0 (1	.87 .69)	-0.32 (-2.66)	$0.69 \\ (4.14)$	-0.37 (-2.44)	0.11 (1.01)
T 1	High	0.51 (2.22)	0.97 (17.85)	0.07 (0.89)	0.08 (1.17)	0.08 (1.58)	0 (0	.24 .92)	0.90 (14.44)	0.17 (1.95)	-0.30 (-3.89)	0.01 (0.17)
Value- Weighted	Low	-0.38 (-1.58)	1.12 (19.95)	0.22 (2.79)	-0.14 (-2.02)	-0.14 (-2.78)	0 (1	.27 .06)	1.13 (19.24)	-0.03 (-0.31)	-0.17 (-2.34)	-0.13 (-2.45)
	High–Low	0.74 (1.88)	-0.15 (-1.63)	-0.15 (-1.17)	0.22 (1.90)	0.21 (2.60)	-0 (-0	.18 .42)	-0.23 (-2.30)	0.19 (1.39)	-0.13 (-1.04)	0.14 (1.51)

Table 5: Technology Momentum Four-Factor Model Loadings for Non-Patenting and Patenting Firms

Notes. The table shows factor loadings from a four-factor model estimated using monthly excess returns from 1998 to 2023 for equal-weighted and value-weighted technology momentum portfolios, for non-patenting and patenting firms. Each sub-table shows factor loadings for high-decile and low-decile portfolios, as well as for a high-minus-low portfolio. T-statistics are reported in parentheses under each alpha estimate.

Our equal-weighted long-short portfolio generates a monthly alpha of 1.97% (*t*-statistic = 4.74) after controlling for the market, size, value, and momentum factors. This performance exceeds the 1.17% monthly alpha reported by Lee et al. (2019), though their alpha is estimated over a longer period 1963–2012. Bekkerman, Fich, and Khimich (2023) report a 1.29% monthly alpha for their text-based approach over the period 1977-2016. However, as Bekkerman, Fich, and Khimich (2023) document, the performance of technological momentum strategies has fallen over time, not risen, so the historical monthly alpha likely overstates performance in recent years. In that sense, our shorter and more recent sample period (1998–2023) works against finding strong results, reducing both statistical power and alpha. Despite this challenging sample period, the alpha we find for non-patenting firms remains statistically significant and exceeds that of previous studies, showing the effectiveness of our approach in identifying technological relationships that extend beyond traditional patent-based measures.

Figure 12 illustrates the cumulative performance of the high-decile technological momentum portfolio from 1998 to 2023, compared against standard technology benchmarks, on both natural and logarithmic scales. The shaded areas indicate NBER-dated recessions. For non-patenting firms, both equal and value-weighted high-decile portfolios substantially outperform the benchmarks over this period, with the equal-weighted portfolio performing particularly well. In contrast, a value-weighted technological momentum strategy for patenting firms performs no better than the technology benchmarks.

5.3 Economic Mechanism

What explains the superior performance of our strategy compared to previous approaches? We propose that the key factor is our ability to identify technological relationships for firms that do not directly own patents. These firms are substantially exposed to technology shocks through their use of technology in production, their dependence on complementary innovations, or their positions in broader technology networks, but their exposures have remained opaque to market participants.

Consistent with Lee et al. (2019) and Bekkerman, Fich, and Khimich (2023), investors



Figure 12: Cumulative Growth of High-Decile Technology Momentum Portfolio(a) Natural Scale

Notes. The figure shows the cumulative growth of a \$1 investment in equal-weighted and value-weighted high-decile technology momentum portfolios of stocks for non-patenting and patenting firms from 1998 to 2023. For comparison, the figure also shows the cumulative growth of four benchmark indices: NYSE Technology, S&P North American Technology Sector, FTSE All-World Technology, and Russell 2000 Technology. Panel 12a plots this growth on the natural scale, while Panel 12b plots it on the logarithmic scale. Technology momentum decile portfolios are formed on the universe of CRSP firms, after dropping the bottom 10% of firms by price, market capitalization, and trading volume. Portfolios are rebalanced monthly based on the recent performance of each firm's technological peers. Shading indicates U.S. recession dates.

appear to systematically underreact to technological information. This underreaction is more pronounced for non-patenting firms with less transparent technological profiles. The stronger persistence in decile membership for non-patenting firms, shown in Figure 10, supports this interpretation. Our results extend previous research by demonstrating that technological momentum effects are not limited to patenting firms but affect the broader universe of publicly-traded companies.

6 Technological Spillovers

The size and direction of technological spillovers between firms remains an unresolved question in the economics of innovation. When innovations emerge, do they mainly benefit the broader economy, or do they mainly benefit innovators at the expense of competitors?

Theoretical arguments for positive spillovers are compelling. Jaffe (1986) shows that R&D spillovers flow through technological proximity, with firms in related spaces benefiting from each other's innovations. Romer (1989) emphasizes that knowledge is nonrivalrous—multiple firms can use the same ideas without diminishing their value. Jones and Williams (1998) shows how knowledge accumulation creates positive feedback loops. These mechanisms suggest innovation generates social returns beyond what innovators capture.

Yet persuasive arguments also exist for negative spillovers. Schumpeter's creative destruction theory posits that innovation displaces existing technologies and the firms using them. Aghion and Howitt (1992) formalize how innovation makes existing products obsolete, imposing losses on incumbents while producing economic growth. Arora, Ceccagnoli, and Cohen (2008) document how patenting firms extract rents from competitors, while Bloom, Schankerman, and Van Reenen (2013) show that product market rivalry can cause innovations to harm technologically-related firms. These competitive dynamics generate what Kogan et al. (2017) term "substantial growth, reallocation, and creative destruction."

The stock market approach to measuring these spillover effects builds on pioneering

work by Pakes (1985), who documented empirical relationships between patents, R&D expenditures, and stock market valuations. Building on this foundation, Austin (1993) developed event-study methods to estimate individual patent values in biotechnology, finding that product-linked patents generate significantly higher returns and that rival firms experience negative returns from competitor patents. Kogan et al. (2017) extended this approach by developing a comprehensive measure of patent values spanning 1926-2010, using filtered stock price reactions around patent grant dates. More recently, Chen, Wu, and Yang (2019) applied similar techniques to FinTech innovations, using machine learning to classify and value patents in this emerging sector.

However, these approaches share a fundamental limitation: they capture only the private returns to patent owners. Even Kogan et al. (2017)'s findings of creative destruction where competitor innovation reduces firm growth—rely on comparing patenting firms within industries. Their aggregate innovation index, while useful for studying economywide effects, cannot distinguish between spillovers to non-patenting firms and private benefits captured by innovators. This limitation is consequential given that the majority of public firms do not patent, leaving the net effect of technological spillovers empirically unresolved.

We revisit the question of technological spillovers using our novel usefulness probabilities to measure firms' exposure to patents. This approach overcomes a limitation of prior work by examining stock market reactions across all firms affected by a patent grant, not just the patent owner. Importantly, we separately estimate effects on non-patenting firms who comprise the majority of public firms but whose distinct responses have not been isolated in prior work—from effects on patenting firms. This distinction matters because non-patenting firms may experience different spillover patterns than their patenting counterparts, yet their aggregate response determines whether innovation generates net social benefits or merely redistributes profits among competitors.

6.1 Methodology

To analyze how innovation affects technologically-related firms, we use an event study methodology. Following Austin (1993), Kogan et al. (2017), and Chen, Wu, and Yang (2019), we examine stock returns around patent event dates (grant dates as opposed to application dates, in our case) to measure patents' impact on firm value.

Our setting presents unique challenges. Traditional event studies examine how patenting firms react to their own patent grants, which are relatively infrequent. In contrast, we measure how all firms react to all patent grants based on our estimates of technological usefulness. This creates two complications. First, with thousands of patents granted weekly, nearly every firm is "treated," at some level, on patent grant days. Second, firms are exposed to patents with varying degrees of technological relevance, as captured by our continuous usefulness probability measures.

We address these challenges using a two-stage approach. In the first stage, we calculate cumulative abnormal returns (CARs) for every firm following every patent grant date. We estimate normal returns using the Fama-French five-factor model over days -250 to -31before each patent grant, excluding the 30 days immediately prior to avoid information leakage. Since patents are always granted on Tuesdays, we include day-of-week fixed effects to separate patent impacts from weekly trading patterns, and to control for the generalized effect of patent grants.

In the second stage, we then compute the number of expected useful patents for each firm on each event date:

$$\operatorname{SumProb}_{it} = \sum_{j \in \mathcal{P}_t} \hat{p}_{ij} \cdot \mathbf{1}[\hat{p}_{ij} > \lambda_{\gamma}], \qquad (7)$$

where \mathcal{P}_t denotes the set of patents granted on date t, \hat{p}_{ij} is our estimated usefulness probability for patent j to firm i, and λ_{γ} is the modified Lee-Liu threshold from our positive-unlabeled learning classifier. The threshold λ_{γ} serves two purposes: it filters out noise from low-relevance patents and ensures our results aren't mechanically driven by the total volume of patents granted on any given day. We then regress the first-stage CARs





Notes. The figure shows coefficient estimates and 95% confidence intervals (measured in basis points) from regressions of cumulative abnormal returns on firm-level patent usefulness measures at different horizons following patent grant announcement dates. Abnormal returns are calculated using the Fama-French five-factor model. The dependent variable is the firm-level cumulative abnormal return from day 1 to day t, where t ranges from 1 to 30 days after the USPTO patent grant announcement date (typically occurring on Tuesdays). The independent variable is the number of patents predicted to be useful (at a 90% usefulness probability threshold) on each announcement date for each firm. The black dots show the estimated coefficients at each horizon, while the gray shaded area represents 95% confidence intervals. The results indicate that firms' stock prices respond positively and significantly to the announcement of patents predicted to be useful to them, with the effect growing to approximately 3 basis points per useful patent over 30 days, suggesting markets gradually incorporate information about technological spillovers to non-patent-owning firms.

on the expected number of useful patents. Our regression specification is given by

$$CAR_{i,t+h} = \beta_h Sum Prob_{it} + \varepsilon_{i,t+h}, \qquad (8)$$

where $\text{CAR}_{i,t+h}$ is the cumulative abnormal return for firm *i* over horizon *h* following patents granted on date *t*, spanning horizons from 1 to 30 days. We cluster standard errors at the patent grant date level to address correlation from multiple firms responding to the same patent grants (see Kolari and Pynnönen, 2010, for a related discussion).

6.2 Empirical Results

Figure 13 shows coefficient estimates β_h at each horizon from 1 to 30 days following patent grant dates. The left panel presents results for non-patenting firms, the right panel for patenting firms. Shaded areas indicate 90% and 95% confidence intervals.

Three key findings emerge. First, cumulative abnormal returns increase significantly

and persistently following the grant of technologically-relevant patents. The effect builds gradually over the 30-day window, suggesting markets require time to incorporate information about technological spillovers. Second, both non-patenting and patenting firms experience positive abnormal returns when exposed to useful patents. Non-patenting firms gain 0.92 basis points per useful patent over 30 days (p < 0.10), while patenting firms gain 0.98 basis points (p < 0.05). These results provide direct evidence of positive technological spillovers—firms benefit from innovations even when they lack ownership rights. Third, the magnitude of spillovers is similar across firm types. The 0.06 basis point difference, while suggesting patenting firms may be slightly better positioned to exploit innovations, is economically small. This similarity challenges the view that only innovative firms can effectively absorb new technologies. These findings contribute new evidence to a long-standing empirical debate. Using our technological usefulness probabilities, we find that positive spillovers dominate: innovations generate benefits that extend beyond patent owners to the broader economy.

7 Conclusion

We develop a novel measure of firm-level technological usefulness by applying positive and unlabeled machine learning to patent and business descriptions. Our methodology addresses a fundamental limitation in the economics literature: existing datasets focus on the small minority of firms that own patents, while the technological profiles of most firms remain unobserved. By measuring the usefulness of patents to all firms rather than just their owners, we create a technology dataset that covers all U.S. public firms and all economically important technology categories.

Our analysis reveals three key findings. First, technological associations between non-patenting and patenting firms are remarkably similar after controlling for industry and size. Non-patenting firms maintain broader but shallower technological portfolios than patenting firms, with higher rates of technological instability, and greater technological generality. Second, we document substantial return predictability from technological momentum strategies, with monthly alphas of 1.97% for non-patenting firms—significantly exceeding the performance of similar strategies applied only to patenting firms. Third, event study evidence shows that firms experience positive abnormal returns following announcements of patents we identify as useful to them, even when they do not own the patents. These results suggest that technological spillovers extend far beyond patent ownership.

Our approach has several limitations that future research should address. Our method assumes that firms' technological activities are reflected in their public disclosures, though firms may use technologies they do not mention or strategically withhold information. Additionally, while our usefulness probabilities strongly predict patent ownership patterns, they remain estimates that may not fully capture the complexity of how firms use technology. Furthermore, our textual analysis depends on the quality and specificity of language in patent and business descriptions, though modern language models help mitigate concerns about ambiguity or technical jargon. Finally, our focus on U.S. firms and patents may limit generalizability to other institutional or international contexts.

Despite these limitations, our work makes important contributions. Methodologically, we introduce positive and unlabeled learning to economics, demonstrating how machine learning can address missing data problems that have constrained empirical research. Substantively, we provide the first comprehensive view of technology usage across all public firms, revealing that non-patenting firms—which account for over half of market capitalization—have rich technological profiles that matter for asset prices and firm performance. By moving beyond patent ownership to measure technological usefulness, we open new avenues for research on innovation, productivity, and technological change across the entire economy.

References

- Aghion, P. and P. Howitt (1992). "A Model of Growth Through Creative Destruction." *Econometrica* 60.2.
- Akcigit, U. and W. R. Kerr (2018). "Growth through heterogeneous innovations." Journal of Political Economy 126.4, pp. 1374–1443.
- Argente, D. et al. (2023). "Patents to products: Product innovation and firm dynamics."

- Arora, A., S. Belenzon, and L. Sheer (2021). "Matching patents to Compustat firms, 1980–2015: Dynamic reassignment, name changes, and ownership structures." *Research Policy* 50.5, p. 104217.
- Arora, A., M. Ceccagnoli, and W. M. Cohen (2008). "R&D and the patent premium." International Journal of Industrial Organization 26.5, pp. 1153–1179.
- Ash, E. and S. Hansen (2023). "Text algorithms in economics." Annual Review of Economics 15.
- Austin, D. H. (1993). "An event-study approach to measuring innovative output: The case of biotechnology." *The American economic review* 83.2, pp. 253–258.
- Autor, D. H., F. Levy, and R. J. Murnane (2003). "The skill content of recent technological change: An empirical exploration." *The Quarterly journal of economics* 118.4, pp. 1279– 1333.
- Bekker, J. and J. Davis (2020). "Learning from positive and unlabeled data: A survey." *Machine Learning* 109, pp. 719–760.
- Bekkerman, R., E. M. Fich, and N. V. Khimich (2023). "The effect of innovation similarity on asset prices: Evidence from patents' big data." *The Review of Asset Pricing Studies* 13.1, pp. 99–145.
- Bloom, N., M. Schankerman, and J. Van Reenen (2013). "Identifying technology spillovers and product market rivalry." *Econometrica* 81.4, pp. 1347–1393.
- Bloom, N. et al. (2021). *The diffusion of disruptive technologies*. Tech. rep. National Bureau of Economic Research.
- Bound, J. et al. (1982). "Who does R&D and who patents?"
- Brown, S. V. and J. W. Tucker (2011). "Large-sample evidence on firms' year-over-year MD&A modifications." *Journal of Accounting Research* 49.2, pp. 309–346.
- Chen, M. A., Q. Wu, and B. Yang (2019). "How valuable is FinTech innovation?" *The Review of Financial Studies* 32.5, pp. 2062–2106.
- Cohen, L., C. Malloy, and Q. Nguyen (2020). "Lazy prices." The Journal of Finance 75.3, pp. 1371–1415.
- Cohen, W. M., R. Nelson, and J. P. Walsh (2000). Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not).
- Dunne, T. (1994). "Plant Age and Technology use in U.S. Manufacturing Industries." *The RAND Journal of Economics* 25.3, p. 488.
- Elkan, C. and K. Noto (2008). "Learning classifiers from only positive and unlabeled data." *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 213–220.
- European Patent Office (EPO) and United States Patent and Trademark Office (USPTO) (2017). Guide to the CPC (Cooperative Patent Classification). Version 1.0.
- Fama, E. (2023). "Production of US Rm-Rf, SMB, and HML in the Fama-French Data Library." *Chicago Both Paper* 22-23.
- Galar, M. et al. (2011). "A review on ensembles for the class imbalance problem: bagging-, boosting-, and hybrid-based approaches." *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 42.4, pp. 463–484.
- Gentzkow, M., B. Kelly, and M. Taddy (2019). "Text as Data." Journal of Economic Literature 57.3, pp. 535–574.
- Graham, S. J., A. C. Marco, and R. Miller (2015). "The USPTO patent examination research dataset: A window on the process of patent examination." *Georgia Tech Scheller College of Business Research Paper No. WP* 43.

- Graham, S. J., A. C. Marco, and A. F. Myers (2018). "Patent transactions in the marketplace: Lessons from the USPTO patent assignment dataset." Journal of Economics & Management Strategy 27.3, pp. 343–371.
- Griliches, Z. (1987). *R&D*, *Patents and Productivity*. Conference report. University of Chicago Press.
- (1990). "Patent Statistics as Economic Indicators: A Survey." Journal of Economic Literature 28.4.
- Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001). The NBER patent citation data file: Lessons, insights and methodological tools.
- Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). "Market Value and Patent Citations." *The RAND Journal of Economics* 36.1, pp. 16–38.
- Hoberg, G. and G. Phillips (2016). "Text-based network industries and endogenous product differentiation." *Journal of Political Economy* 124.5, pp. 1423–1465.
- (2018). "Text-based industry momentum." Journal of Financial and Quantitative Analysis 53.6, pp. 2355–2388.
- Jaffe, A. (1986). "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value." *American Economic Review* 76.5, pp. 631–664.
- Jones, C. I. and J. C. Williams (1998). "Measuring the Social Return to R&D." *The Quarterly Journal of Economics* 113.4, pp. 1119–1135.
- Kakhbod, A. et al. (2024). "Measuring Creative Destruction." Available at SSRN 5008685.
- Kelly, B. T. and D. Xiu (2023). "Financial machine learning." Available at SSRN.
- Kogan, L. et al. (2017). "Technological innovation, resource allocation, and growth." *The Quarterly Journal of Economics* 132.2, pp. 665–712.
- Kogan, L. et al. (2021). Technology-skill complementarity and labor displacement: Evidence from linking two centuries of patents with occupations. Tech. rep. National Bureau of Economic Research.
- Kolari, J. W. and S. Pynnönen (2010). "Event Study Testing with Cross-sectional Correlation of Abnormal Returns." The Review of Financial Studies 23.11, pp. 3996– 4025.
- Kydland, F. E. and E. C. Prescott (1982). "Time to Build and Aggregate Fluctuations." *Econometrica* 50.6, p. 1345.
- Lafond, F. and D. Kim (2019). "Long-run dynamics of the US patent classification system." *Journal of Evolutionary Economics* 29.2, pp. 631–664.
- Lee, C. M. et al. (2019). "Technological links and predictable returns." Journal of Financial Economics 132.3, pp. 76–96.
- Lee, W. S. and B. Liu (2003). "Learning with positive and unlabeled examples using weighted logistic regression." *ICML*. Vol. 3. 2003, pp. 448–455.
- Lemley, M. A. and C. Shapiro (2005). "Probabilistic patents." Journal of Economic Perspectives 19.2, pp. 75–98.
- Lerner, J. and A. Seru (2021). "The Use and Misuse of Patent Data: Issues for Finance and Beyond." *The Review of Financial Studies* 35.6, pp. 2667–2704.
- Lerner, J. et al. (2021). Financial innovation in the 21st century: Evidence from us patents. Tech. rep. National Bureau of Economic Research.
- Lerner, J. (1994). "The importance of patent scope: an empirical analysis." *The RAND Journal of Economics*, pp. 319–333.
- Liu, B. et al. (2002). "Partially supervised classification of text documents." *ICML*. Vol. 2. 485. Sydney, NSW, pp. 387–394.

- Lobo, J. and D. Strumsky (2019). "Sources of inventive novelty: two patent classification schemas, same story." *Scientometrics* 120.1, pp. 19–37.
- Lopez-Lira, A. (2019). "Risk Factors That Matter: Textual Analysis of Risk Disclosures for the Cross-Section of Returns." *mimeo*.
- Loughran, T. and B. McDonald (2020a). "Measuring firm complexity." Journal of Financial and Quantitative Analysis, pp. 1–55.
- (2020b). "Textual analysis in finance." Annual Review of Financial Economics 12, pp. 357–375.
- McCarthy, P. M. and S. Jarvis (2010). "MTLD, vocd-D, and HD-D: A validation study of sophisticated approaches to lexical diversity assessment." *Behavior research methods* 42.2, pp. 381–392.
- Mezzanotti, F. and T. Simcoe (2023). Innovation and appropriability: Revisiting the role of intellectual property. Tech. rep. National Bureau of Economic Research.
- Moore, K. A. (2005). "Worthless patents." *Berkeley Technology Law Journal* 20.4, pp. 1521–1552.
- Mordelet, F. and J.-P. Vert (2014). "A bagging SVM to learn from positive and unlabeled examples." *Pattern Recognition Letters* 37, pp. 201–209.
- Myers, K. R. and L. Lanahan (2022). "Estimating spillovers from publicly funded R&D: Evidence from the US Department of Energy." American Economic Review 112.7, pp. 2393–2423.
- Pakes, A. (1985). "On patents, R & D, and the stock market rate of return." Journal of political economy 93.2, pp. 390–409.
- Romer, P. (1989). "Endogenous Technological Change."
- Siblini, W. et al. (2020). "Master your metrics with calibration." International Symposium on Intelligent Data Analysis. Springer, pp. 457–469.
- Simmons, H. J. (2014). "Categorizing the useful arts: Part, present, and future development of patent classification in the united states." *Law Libr. J.* 106, p. 563.
- Solow, R. M. (1960). "Investment and Technical Progress." Mathematical methods in the social sciences 1, pp. 48–93.
- Strumsky, D., J. Lobo, and S. Van der Leeuw (2012). "Using patent technology codes to study technological change." *Economics of Innovation and New technology* 21.3, pp. 267–286.
- U.S. Congress (1836). Act of July 4, 1836, ch. 357, 5 Stat. 117. https://www.loc.gov/ item/uscode18360704/.
- (1934). Securities Exchange Act of 1934, 15 U.S.C. §§ 78l, 78m. https://www.govinfo.gov/content/pkg/USCODE-2011-title15/pdf/USCODE-2011-title15-chap2B.pdf.
- (1952). 35 U.S.C. §§ 101-103 (Patent Act of 1952). https://www.govinfo.gov/ content/pkg/USCODE-2012-title35/pdf/USCODE-2012-title35-partII-chap10sec101.pdf.
- (2011). 35 U.S.C. § 8 (Leahy-Smith America Invents Act of 2011). https://www.law. cornell.edu/uscode/text/35/8.
- U.S. Securities and Exchange Commission (2013). *17 CFR Part 229 (Regulation S-K)*. https://www.govinfo.gov/content/pkg/CFR-2013-title17-vol2/pdf/CFR-2013-title17-vol2-sec229-101.pdf.
- United Nations (1971). *Profiles of Manufacturing Establishments*. Vol. I, II, and III. Industrial Planning and Programming Series 4. New York: United Nations.

United States Patent and Trademark Office (USPTO) (2013). MPEP § 2107: Guidelines for Examination of Applications for Compliance with the Utility Requirement.
World Intellectual Property Organization (2024). Guide to the International Patent Classification (2024 Edition). World Intellectual Property Organization. Geneva, Switzerland.