

Technological Disruptiveness and the Evolution of IPOs and Sell-Outs

Donald Bowen, Laurent Frésard, and Gerard Hoberg*

January 28, 2019

ABSTRACT

We show that the recent decline in IPOs on U.S. markets is related to changes in the technological disruptiveness of startups, which we measure using textual analysis of patents from 1930 to 2010. We focus on VC-backed startups and show that those with ex-ante disruptive technologies are more likely to exit via IPO and less likely to exit via sell-out. This is consistent with IPOs being favored by firms with the potential to carve out independent market positions with strong defenses against rivals. We document an economy-wide trend of declining technological disruptiveness since World War II that accelerated since the late 1990s. This trend predicts fewer IPOs and more sell-outs, and we find that roughly 20% of the recent dearth of IPOs, and 49% of the surge in sell-outs, can be attributed to changes in firms' technological characteristics.

Key words: Initial Public Offerings (IPOs), Acquisitions, Sell-Outs, Technology,

Disruptiveness, Venture Capital

JEL classification: G32, G34, G24

*Virginia Tech, the University of Lugano, Swiss Finance Institute, and the University of Southern California respectively. Bowen can be reached at donbowen@vt.edu, Frésard can be reached at laurent.fresard@usi.ch, and Hoberg can be reached at hoberg@marshall.usc.edu. We thank Sunil Muralidhara and Lauren Phillips for excellent research assistance. We also thank Jay Ritter on sales data for IPO firms. For helpful comments, we thank Song Ma, Adrien Matray, Francesco D'Acunto, Francois Degeorge, Francois Derrien, Francesco Franzoni, Michelle Lowry, Shawn Thomas, Xuan Tian and seminar participants at American University, Arizona State University, Bocconi University, Dauphine University, the CCF conference in Bergen, the FRA conference in Las Vegas, Georgetown University, HEC Paris, Johns Hopkins University, the SFI Research days, Stockholm School of Economics, Università Cattolica, University of Virginia, Texas Christian University, Tsinghua University, CKGSB, the University of Lugano, the University of Maryland, the University of Zurich, and Virginia Tech. All errors are the authors' alone. All rights reserved by Donald Bowen, Laurent Frésard, and Gerard Hoberg.

I Introduction

Since the late 1990s, the number of initial public offerings (IPOs) in U.S. markets has sharply declined. Over the same period, the number of private firms exiting via acquisitions (i.e., trade sales) has soared. Successful startups are nowadays more likely to sell out to other (public or private) companies than seek independent public listings. Many observers in the media and policy circles worry that these trends reflect a general erosion in the ability of U.S. financial markets to spur economic growth and spread its benefits across the general public.¹ Recent studies indicate, however, that the dearth of IPOs is unlikely due to regulatory changes affecting public firms (Gao, Ritter, and Zhu (2013) and Doidge, Karolyi, and Stulz (2013)) and is partly explained by regulatory changes affecting the financing ability of private firms (Ewens and Farre-Mensa (2018)). In contrast to existing regulation-based explanations, we show that the recent shift from IPOs to trade sales is related to changes in the technological disruptiveness of U.S. startups.

Our analysis builds on the long-standing idea that exiting through an IPO or by selling out to another firm depends on which exit type enhances the growth potential of successful private firms (Bayar and Chemmanur (2011) or Gao, Ritter, and Zhu (2013)). Accordingly, we predict and find that private firms that can develop an independent market presence by creating new markets or disrupting existing ones with patented technologies are more likely to exit and scale up through a public listing.² This is optimal as such firms do not have to share the gains to achieving scale with an acquirer. In contrast, firms with less disruptive innovation, such as technology that can improve existing products, are more likely to sell out to existing firms that can more efficiently scale their technologies. Studying the exits of a large sample of U.S. startups backed by venture capitalists (VCs), we find that startups' technological disruptiveness has significantly decreased in recent

¹See for instance “The endangered public company: The big engine that couldn’t,” *The Economist* (May 19, 2012) or “US stock markets seek depth in IPO pool,” *Financial Times* (January 9, 2018).

²Consistent with this idea, Darby and Zucker (2018) show that biotechnology firms go public when they have a science base that can be successfully commercialized, Chemmanur, He, He, and Nandy (2018) report that manufacturing firms are more likely to go public than sell out when they already have a strong product market presence (i.e., market share), Poulsen and Stegemoller (2008) and Cumming and Macintosh (2003) show that firms with more growth potential favor exit through IPOs.

years, and this trend can partially explain the decline in IPOs and soaring trade sales.³

We define a technology (patent) as *ex ante* disruptive if it is newly gaining traction and has the potential to rapidly change the technological path followed by a large number of firms. Technological disruptiveness is thus distinct from the concept of innovativeness used in the literature, which is based either on patents’ use of new or unique vocabulary, or various measures based on patents’ citations (e.g., Packalen and Bhattacharya (2018)). We measure the technological disruptiveness of the 6.6 million patents filed with the U.S. Patent Office (USPTO) between 1930 and 2010 as the extent to which their text contains vocabulary that is fast-growing across all patent applications in the same year. For example, genetics-related words such as ‘peptide’, ‘clone’, or ‘recombinant’ were rapidly growing across a host of patents in 1995, and thus our measure would tag patents extensively employing such words in 1995 as disruptive.

Because our goal is to estimate predictive models of startups’ exits, we measure the technological disruptiveness of patents at the time of their application, using only text contained in all past and contemporaneous patents.⁴ We confirm empirically that disruptiveness is distinct from existing measures of innovativeness and that it identifies the breakthrough inventions of the last century that radically altered the course of technological evolution (e.g., jet engine, transistor, laser, satellite, DNA recombination, or Google’s PageRank algorithm).⁵ Furthermore, we find strong validation of our *ex ante* measure of disruptiveness, which predicts *ex post* changes in the trajectory of future inventions as indicated by high levels of *ex post* citations and estimated economic values.

³We focus on VC-backed startups because detailed data enables us to precisely link their choice and timing of exits (i.e., IPO or sell-out) to the patents they file with the U.S. Patent and Trademark Office (USPTO). This is a useful and relevant sample, as VC-backed startups account for a large fraction of the IPO and acquisition market (Ritter (2017)), produce a substantial share of innovation in the economy (Gornall and Strebulaev (2015)), and face exits due to the limited lives of VC funds.

⁴See Dahlin and Behrens (2005) for a detailed discussion of the distinction between *ex ante* and *ex post* disruptiveness. The *ex-ante* feature of our measure also eliminates look-ahead bias, reduces truncation bias (Lerner and Seru (2017)), and increases the measure’s utility to practitioners (e.g., who can use it to predict outcomes and investor returns) and regulators (e.g., who can examine the impact of policies in real time).

⁵For instance, using a list of 101 historically important patents based on the USPTO’s “Significant Historical Patents of the United States” list, we estimate that the average patent in this list of radical inventions ranks in the 77th percentile of our disruptiveness measure.

Consistent with recent evidence suggesting that new ideas are getting harder to discover and develop (Jones (2009) and Bloom, Jones, Reenen, and Webb (2017)), the average disruptiveness of U.S. patents has markedly declined since 1950, except for temporary spikes during the 1970s (i.e., computers), the 1980s (i.e., genetics), and the 1990s (i.e., the internet). As they contain pioneering and fast growing ideas that are legally protected, startups holding more disruptive patents are likely to radically change markets and capture valuable competitive advantages. We posit that these startups are more likely to go public and less likely to sell out. A comparison of aggregate technological disruptiveness versus private firm exits suggests that IPOs are indeed more prevalent during disruptive periods, whereas trade sales prevail in less disruptive times.

We confirm this interpretation at the micro-level by examining the link between the technological disruptiveness of 9,167 VC-backed startups (94,703 patents) and their exits over the 1980-2010 period. We find strong evidence that startups with more disruptive patents are significantly more likely to go public and less likely to sell out. The link between technological disruptiveness and startup exits is economically large, as a one standard deviation increase in a startup’s technological disruptiveness is associated with a 25.2% increase in its IPO rate and a 18.8% decrease in its rate of exit through trade sales. The role of disruptiveness in startups’ exits remains after controlling for startups’ age, size, financing rounds, overall financial market conditions, and other technological characteristics, such as their technological “breadth” (based on the intensity with which patent text combines vocabulary from diverse bodies of knowledge), similarity to other firms, patent citations, average word age, and originality. They also hold in specifications with fixed effects for startup cohorts, geographic locations, and technological areas.⁶

We provide additional evidence supporting the validity of our new measure of technological disruptiveness and its role in explaining startups’ exits. For example, we delineate each startup’s product market based on the textual similarity between its business description and the 10-K business descriptions of publicly-traded firms. As direct validation,

⁶Our results are also robust to changes in econometric specifications that vary the horizon over which we measure exits (ranging from the next quarter to the next five years), and to focusing on the early part of the sample to limit potential truncation bias relating to startup resolution and patent grants (Lerner and Seru (2017)).

we find that the publicly-traded firms that are most similar to a given startup discuss market disruption significantly more when that startup displays a high level of technological disruptiveness. Alternatively, we show that among startups exiting through IPOs, those with more disruptive technologies exit into less stable and more competitive markets that intuitively offer greater opportunities for disruption (Abernathy and Utterback (1978)).

Given our finding that startups' ex-ante technological disruptiveness is strongly related to their ex-post choice of exit, we next examine whether the aforementioned aggregate decline in disruptiveness in recent years can help to explain disappearing IPOs and the surge in trade sales in the time-series. Inspired by the disappearing dividends literature,⁷ we first estimate quarterly cross-sectional exit models with different sets of explanatory variables over an initial period (1980-1995), and use the fitted model to predict quarterly IPO and sell-out rates in the subsequent out-of-sample years (1996-2010) given the characteristics of firms in the latter period. We then compare the models' predictions to the actual exit rates in the out-of-sample period.

The actual probability that a VC-backed startup will experience an IPO in a given quarter is 0.33 percentage points in the out-of-sample period. When we fit the predictive model without our text-based technological characteristics, the predicted IPO rate is 0.84 percentage points per quarter. Because the actual rate is 0.51 percentage points lower than the predicted rate, we thus confirm the disappearing IPOs "anomaly" over the 1996-2010 period. Notably, when we include the text-based technological characteristics the gap declines from 0.50 to 0.42 percentage points. Our text-based variables thus explain roughly 16% of the disappearing IPOs anomaly, which rises to 20% when we train our model to look at longer predictive horizons. A similar analysis reveals that changes in technological traits account for roughly 50% of the recent rise of trade sales. These findings are robust to changes in the specifications and timing.

We further show that the role of startups' technological characteristics in the recent decline of IPOs varies significantly across IPO size and product markets. In particular,

⁷We use methodology used in Fama and French (2001) and Hoberg and Prabhala (2009) to explain disappearing dividends, and more recently used by Bates, Kahle, and Stulz (2009) to explain increasing cash holdings.

we find that technological disruptiveness is a strong determinant of the decline of “small” IPOs (defined as IPOs with pre-IPO sales below the median). This is the segment of the IPO market displaying the sharpest decline in recent years (Gao, Ritter, and Zhu (2013)). Startups’ technological traits and disruptiveness appear unrelated to the decline of large IPOs, but they explain about 40% of the decline in small IPOs. The stability of the product market is also important, as we predict that disruption will become harder in more stable markets.⁸ Following Hoberg, Phillips, and Prabhala (2014) we compute the degree of product market stability for each startup using their business description text, which is observable at the time of initial VC funding. Technological traits explain 25% of the decline in startups’ IPOs in more stable markets, compared to just 3% in fluid markets.

Our analysis adds to the recent literature examining potential explanations of the recent disappearance of IPOs, and the contemporaneous rise of trade sales.⁹ Ewens and Farre-Mensa (2018) indicate that part of the IPO decline results from the increased bargaining power of founders (relative to investors), their preference for control, and inexpensive capital in the private market. Gao, Ritter, and Zhu (2013) suggest that the decline in IPOs originates from changes in market structure that favor selling out to realize economies of scope. Doidge, Kahle, Karolyi, and Stulz (2018) argue that an increased focus on intangibles also likely plays an important role. Our paper shows that changes in firms’ technological traits (especially disruptiveness) also account for part of the decline in IPOs and the rise of trade sales in the recent period. Unlike existing papers, our analysis considers the margins of startups’ exits jointly and quantifies how much of the observed trends are attributable to changes in firms’ technological characteristics. Although more work is needed to draw clear policy recommendations, our findings suggest that the decline in IPOs might be particularly unresponsive to policy changes, although funding subsidies to disruptive entrepreneurs might be a possible exception.

⁸The intuition for this hypothesis is that the most logical inventions are discovered first. As industries mature, the best ideas become “picked over” and later inventions arise only with very high search costs, as in Abernathy and Utterback (1978).

⁹Ritter and Welch (2002) and Lowry, Michaely, and Volkova (2017) provide comprehensive surveys of the literature on IPOs.

Our findings also add to the literature studying the determinants of startup exits. While the vast majority of existing research studies IPOs or trade sales in isolation, a small number of papers examines these exit choices jointly and indicate that they depend on founders’ private benefits of control, product market presence, and firms’ growth potential (Cumming and Macintosh (2003), Bayar and Chemmanur (2011), Poulsen and Stegemoller (2008) or, Chemmanur, He, He, and Nandy (2018)). Our analysis emphasizes the pivotal role of technological characteristics in understanding startups’ exits and growth patterns.

Finally, our paper adds to a collection of recent studies using patent text to measure and characterize the nature of technology. Packalen and Bhattacharya (2018) and Balsmeier, Assaf, Chesebro, Fierro, Johnson, Johnson, Li, Luck, OReagan, Yeh, Zang, and Fleming (2018) use patent text to identify new ideas based on the first appearance of words (or sequences of words) in the patent corpus, and analyze their propagation. Closest to our paper, Kelly, Papanikolaou, Seru, and Taddy (2018) construct a measure of patent “significance” based on the textual similarity with prior and future patents to analyze technological change over the last centuries. While their measure of technological change also considers changes in textual similarity, it relies on ex-post information (i.e., the text of future patents). The intentional absence of look-ahead bias in our measure is important specifically for what we study since our goal is to explain startups’ decisions based on timely patent information, thereby limiting endogeneity or truncation concerns.¹⁰

II Text-based Technological Characteristics

In this section, we describe the patent data and explain the construction of our new measures of technological characteristics based on patent text.

A Patent Text

We use a web-crawling algorithm to gather information from Google Patents for all 6,595,226 patents that were applied for between 1930 and 2010 and granted by 2013.

¹⁰Indeed, similar to future patent citations, the textual similarity measure based on future patents could be affected by the very decisions we are attempting to analyze such as exit mechanisms (e.g., through investors interest in the developed technology upon getting VC funding, going public, or being acquired). Although measures using future information are excellent for some purposes, we believe ex-ante measurable variables are important in our context.

For each patent, we gather the publication date, application date, names of inventor(s), and initial assignee(s). We also collect the full patent text and information on the technology classification of the patents by converting the U.S. Patent Classification (USPC) into the two-digit NBER technology codes created in Hall, Jaffe, and Trajtenberg (2001). Since we are interested in measuring the technological changes pertaining to the corporate sector, we categorize each patent into groups based on four types of applicants: U.S. public firms, U.S. private firms, foreign (private or public) firms, or others (e.g., universities or foundations). For brevity, we describe this classification method in Appendix A.

[Insert Figure I about here]

The full text of each patent consists of three distinct sections: abstract, claims, and patent description. The claims section defines the scope of legal protection granted. The description section explicitly describes the characteristics of the invention/innovation. It typically includes a title, technical field, background art, specification example, and industrial applicability. The abstract contains a summary of the disclosure contained in the description and claims sections. Figure I presents an example of a typical patent textual structure (#6285999, “A method for node ranking in a linked database”, assigned to Google in 1998). We append all three sections into a unified body of text because earlier patents do not include all sections, and because the organization of patent text into the three sections may have changed over time (Kelly, Papanikolaou, Seru, and Taddy (2018) and Packalen and Bhattacharya (2018)).

Following earlier studies constructing variables from text (e.g., Hanley and Hoberg (2010) or Hoberg and Phillips (2016)), we represent the text of each patent as a numerical vector with a length equal to the number of distinct words in the union of all patent applications in a given year t . We denote this length N_t .¹¹ Following the convention in the literature, we eliminate commonly-used words (words appearing in more than 25% of all patents in a given year) and rare words (words appearing only in one patent in a

¹¹We organize patents based on their application year rather than the year of the patent grant, as this more accurately reflects the timing of innovation.

given year).¹² Each patent j applied for in year t is then represented by a vector $V_{j,t}$ (of length N_t) in which each element corresponds to the number of times patent j employs one of the unique N_t words used in year t . If patent j does not use a given word, the corresponding element of $V_{j,t}$ is set to zero. This vectorization procedure insures that all patent applications in a given year are represented by a collection of vectors that are in the same space (of dimension N_t).

Due to the large number of words used across all patents in a given year, the vectors $V_{j,t}$ are quite sparse with most elements being zero. For instance, in 1980, the number of distinct words used in an average patent is 352, and the median is 300, while there are 400,097 distinct words used across all patent applications. In 2000, the average and median are 453 and 338, and the total across all applications is 1,358,694.

B Technological Disruptiveness

We define a technology (patent) as ex ante “disruptive” if it is (1) newly gaining traction, and (2) has the potential to change the technological path followed by other firms (see for instance Dahlin and Behrens (2005)).¹³ This definition relates a patent to the activities around it in the technology space. To capture the level of disruptiveness of a given patent, we focus on the extent to which it uses vocabulary that is new and is experiencing rapid growth across all patents compared to existing knowledge. Intuitively, disruptive patents are based on new ideas that have the *potential* to radically change an industry or business strategies either by creating whole new markets or by disrupting existing ones. We thus posit that disruptive patents use novel vocabularies that are fast-growing in the universe of patents.

To quantify disruptiveness, we rely on the rate of change in the use of each word among all patent applications during the current and prior year. As our goal is to link

¹²Given the highly technical and rapidly evolving nature of text in the patent corpus, we do not implement additional filters (e.g. nouns only). While this choice might potentially introduce noise into our measurements, it maintains power.

¹³The notion that inventions could be either “incremental” (i.e., run-of-the-mill inventions) or “disruptive” has a long tradition in the literature studying technological changes. Besides “disruptive”, existing research in different fields employs various terminologies to label inventions that break with traditions, such “radical”, “discontinuous”, “generational”, or “breakthrough” (Dahlin and Behrens (2005)).

technological disruptiveness to future outcomes (i.e., startups’ exits), we construct technological disruptiveness as an ex-ante measure, using only text contained in all past and contemporaneous patents. As highlighted by Christensen (1997), disruptive inventions do not always lead to direct commercial success, so we propose to measure the technological disruptiveness of patents at the time of their application (i.e., before it has reached the market). As we note in later validation tests, disruptive patents do indeed influence the content of future inventions ex post, as measured by their citation path or economic value (see Section II.D). Our focus on ex ante measurement also alleviates any look-ahead bias and truncation in our measure, which is necessary given our objective to estimate predictive models of startup exit and to examine disappearing IPOs while avoiding endogeneity arising from the interplay of corporate decisions and future outcomes.

We define an aggregate vector Z_t (of length N_t) in each year t with elements containing the number of times a given word is used across all patent applications in year t . This vector thus represents the aggregate frequency of word usage in the patent corpus in a given year. We then compute the annual rate of change in the usage of each word (from $t - 1$ to t) by defining the (annual) vector D_t as:

$$D_t = \frac{Z_t - Z_{t-1}}{Z_t + Z_{t-1}}, \quad (1)$$

where division is element-by-element.¹⁴ The set of annual vectors D_t thus tracks the appearance, disappearance, and growth of specific technological vocabulary across all patents over time. Elements of D_t are positive if the usage of the corresponding words increases from year $t - 1$ to t , and negative if it decreases (e.g., words becoming obsolete).

[Insert Table I about here]

As an illustration, Table I displays the ten words experiencing the largest increases and decreases in usage across all patent applications in specific years. For instance, in 1995, we detect an acceleration of terms related to genetics, such as “polypeptides”, “clones”,

¹⁴To ensure Z_t and Z_{t-1} are in the same space (i.e., the union of N_t and N_{t-1}), we modify Z_{t-1} by adding zero elements for words that newly appeared in year t (as they were not originally in the $t - 1$ space). Analogously, we modify Z_t by adding zero elements for words that appeared in year $t - 1$ but not year t .

“recombinant” and “nucleic”, following rapid progress in genome sequencing. In that year, use of terms such as “cassette,” “ultrasonic,” and “tape” are sharply decreasing. In 2005, the most rapidly growing words are related to the internet and include terms such as “broadband”, “click”, “configurable”, or “telecommunications”.

To obtain the level of technological disruptiveness for a given patent j , we take the frequency-weighted average of the vector D_t based on the words that patent j uses as follows:

$$\text{Tech Disruptiveness}_{j,t} = \frac{V_{j,t}}{V_{j,t} \cdot \mathbf{1}} \cdot D_t \times 100, \quad (2)$$

where the operator “ \cdot ” denotes the scalar product between two vectors, and “ $\mathbf{1}$ ” is a unit vector of dimension N_t . Intuitively, patents using words whose usage surges across all patent applications (i.e., have positive entries in the vector D_t) have higher levels of disruptiveness. This is the case for patents that either employ words that appear in the patent space for the first time, or that use words whose usage experiences fast growth across all patents. Hence, a patent using words such as “polypeptides”, “clones”, and “recombinant” is classified as disruptive if its application year is 1995, but not in 2005. Symmetrically, patents relying on words whose usage decreases across all patents (i.e., using obsolete vocabulary such as “cassette,” or “tape” in 1995) have lower (and possibly negative) levels of disruptiveness.

C Technological Breadth and Similarities

We also develop text-based measures capturing the technological breadth of each patent, as well as their similarities with other patents of economically linked firms. We posit that these characteristics are also relevant in predicting startup exits. To measure the technological breadth of a patent, we first identify words that are strongly associated with a specific technological field using the six broad technological fields (f) defined by the first digit of the NBER technical classification.¹⁵ Specifically, we count how often a given word (in N_t) is used by patents classified into each field in each year, and keep the two fields with the highest usage of the given word. We define a word as “specialized”

¹⁵ “Chemicals”, “Computer and Communication”, “Drugs and Medicine”, “Electricity”, “Mechanics”, and “Others”.

(and associated with a field f) if its use in its most popular field is more than 150% that of its second most popular field. Each word is thus classified into one of the six fields of specialization or it is deemed an “unspecialized” word. For instance, words such as “bluetooth” and “wifi” are in the “Computer and Communication” field, and “acid” and “solvent” are in the “Chemicals” field. Second, we define as $w_{j,t,f}$ the fraction of patent j ’s specialized words that are classified into each field f . By construction, each $w_{j,t,f}$ lies in the $[0,1]$ interval, and they sum to one for each patent j . We then define technological breadth as:

$$\text{Tech Breadth}_{j,t} = 1 - \sum_{f=1}^6 w_{j,t,f}^2. \quad (3)$$

This measure is one minus the technological concentration of the patent’s vocabulary. Patents have higher technological breadth if they amalgamate vocabularies from different specialized technological fields and for which a wide range of knowledge is needed to develop and understand the invention.¹⁶ In contrast, patents with lower breadth use vocabulary that primarily concentrates on one specialized technological field.

Next, we define three measures of technological similarity by directly comparing the vocabulary of a given patent to that of patents assigned to three specific groups: lead innovators, private U.S. firms, and foreign firms. To do so, we rely on the concept of cosine similarity (see Sebastiani (2002)), which is defined as the scalar product between each patent j ’s normalized word distribution vector $V_{i,t}$ and a normalized word vector aggregating the vocabulary specific to a given group of patents.¹⁷ To capture the similarity of a given patent j with patents of “Lead Innovators” (henceforth LI), we define LIs annually as the ten U.S. public firms with the most patent applications. This set varies over time as the importance of sectors and firms changes. LIs include Microsoft and Intel in 2005, IBM and Motorola in 1995, General Electric and Dow Chemical in 1985, and General Electric, Bell Telephone, and General Motors in 1935. For each set of LIs in year t , we first identify the set of patents applied for by the LIs over the past three years (i.e., from year $t - 2$ to t). The aggregate LI word vector in year t corresponds to the

¹⁶An illustrative example of a high technological breadth patent is the satellite (patent #2835548 from 1957), which required both mechanical and electronic technologies among others.

¹⁷The result is bounded in $[0,1]$ and values close to one indicate closer textual similarity.

equally-weighted average of the resulting normalized patent vectors. We then compute the similarity of any given patent to those of the LIs as:

$$\text{LI Similarity}_{j,t} = \frac{V_{j,t}}{\|V_{j,t}\|} \cdot \frac{V_{LI,t}}{\|V_{LI,t}\|}. \quad (4)$$

Because the word vector $V_{LI,t}$ aggregates word usage across patents of lead innovators in the last three years, patents exhibiting higher levels of LI similarity contain technologies that are textually close to those of lead innovators. In contrast, patents with low levels of LI similarity use text that is unrelated to that used in the patents of lead innovators, and thus are more distant technologies.

We use similar methods to compute the similarity between the text in each patent j and the overall text of patents assigned to private U.S. firms or to foreign firms. Specifically, we form the aggregate private firm (foreign firm) word vectors $V_{P,t}$ ($V_{F,t}$) as the equally-weighted average of the normalized vectors V of patent applications by private (or foreign) firms in year t .¹⁸ We then define the similarity between each patent j and the contemporaneous patent applications of all private U.S. firms as:

$$\text{Private Similarity}_{j,t} = \frac{V_{j,t}}{\|V_{j,t}\|} \cdot \frac{V_{P,t}}{\|V_{P,t}\|}. \quad (5)$$

Analogously, the similarity between patent j and those of foreign firms is:

$$\text{Foreign Similarity}_{j,t} = \frac{V_{j,t}}{\|V_{j,t}\|} \cdot \frac{V_{F,t}}{\|V_{F,t}\|}. \quad (6)$$

These measures are high for patents whose vocabulary is technologically close to that of patents assigned to private U.S. firms or to foreign firms, respectively. As proximity to private and foreign firms indicates more contested markets for innovation, these variables allow us to examine competitive effects in exit choices.

¹⁸Because these groups contain very large numbers of patents, we aggregate over just the single year t . We also note that when a patent j belongs to a private U.S. firm or a foreign firm, we exclude it from the set of patents used to compute $V_{P,t}$ and $V_{F,t}$, respectively.

D Descriptive Statistics, Comparisons, and Examples

Table II presents descriptive statistics for our new text-based technological characteristics as well as existing patent variables from the literature for the full sample of patents (1930-2010). All variables are defined in Appendix B. We first focus on our central new variable: patents' technological disruptiveness. Across all patent applications between 1930 and 2010, we note that the empirical distribution of technological disruptiveness is highly skewed. The first row of Panel A indicates that the average disruptiveness of patents is 1.64, the median is 1.27, and the 75th percentile is 2.34. The observed asymmetry indicates that while the vast majority of patents contain incremental inventions, a smaller set of patents appear to be highly disruptive.

[Insert Table II about here]

Our ex ante measure of disruptiveness is designed to capture the potential of a patent to alter the technological path followed by other firms and to create high-value economic opportunities. To assess whether this ex ante potential is realized ex post, we examine the association between our ex ante measure of patent disruptiveness, the ex post citation path, and the patent's economic value. To do so, we collect accumulated citations for each patent as of early 2014 from Google Patents.¹⁹ For the subset of patents assigned to publicly-listed firms, we also use the implied economic value as measured using stock returns around patent grant announcements (see Kogan, Papanikolaou, Seru, and Stoffman (2016), henceforth KPSS).

Panel B of Table II reveals significant and positive correlations, indicating that patent applications exhibiting higher levels of ex ante technological disruptiveness attract significantly more citations ex-post, and have larger economic values when granted. These positive relationships remain highly significant in regressions including fixed effects for the application year, grant year, and technology categories, as well as clustering standard errors by grant years. While the technological disruptiveness of a patent significantly predicts its future citations, it is constructed solely based on prior information. Disrup-

¹⁹We eliminate patents granted in early 2014 to avoid having a partial year of data.

tiveness therefore provides a more timely characterization of patents' disruption potential than any measure based on citation counts.²⁰

Although a patent's technological disruptiveness is significantly related to future citations and economic value, the reported correlations range between 0.04 and 0.06, suggesting that not all patents with an ex ante potential to disrupt materially influence the path of future inventions or are valued highly upon being granted (consistent with Christensen (1997) and Christensen and Rosenbloom (1995)). To better understand this difference between these metrics, we study in detail the 25 patent applications occurring between 1980 and 2010 that have attracted the largest number of citations in the first five years following their grant. We display in Panel A of Table III the percentile of each such patent in the (cohort-adjusted) distribution of technological disruptiveness. Although citations are clear indications that patents are widely used ex-post, technological disruptiveness is specifically designed to capture new and fast-growing ideas ex-ante. This distinction is evident in Panel A, where patents displaying high levels of disruptiveness tend to build-out novel ideas (e.g., the use of HTML in the mid-nineties), rather than refinements or synthesis of existing ideas (e.g., the iPhone or semiconductor advances in the mid-2000s). Indeed, the average percentile of technological disruptiveness for these 25 patents is 57, as the most disruptive patents refer to more nascent technologies.

[Insert Table III about here]

Panel B displays similar figures for a collection of twelve unambiguous breakthrough patents, as identified by the USPTO's "Significant Historical Patents of the United States".²¹ The average percentile of technological disruptiveness of these patents is very high (84th percentile). The most disruptive patents in this set are "Complex computer" in 1944 (#2668661) and DNA modifications in 1980 (#4399216), both of which virtually created new industries. Other key inventions, such as the satellite, laser, and PageRank,

²⁰Additionally, our measure can be computed for all patents, whereas economic value (which may not necessarily coincide with scientific value) can only be computed for public firm patents (which represent only 28% of all patents).

²¹Listed patents applied before 1960 come from a list of historical patents at <http://www.uspat.com/historical/>. More recent patents are noted for the revenue they generated.

use vocabulary that is new and rapidly growing across many patents around the time of their application. Interestingly, some of these breakthrough inventions are barely cited. For instance, the patents related to the invention of the “television” (#1773980) and the “helicopter” (#1848389) are in the lowest percentile of the cohort-adjusted distribution of citations. Yet, our new measure classifies these patents as highly disruptive.²² In the Appendix (Table A1), we consider a more comprehensive list of 101 important patents between 1930 and 2010 identified by Kelly, Papanikolaou, Seru, and Taddy (2018), which encompasses indisputably important and radical inventions (see their Table A.6). We again consider the percentiles of the distributions of technological disruptiveness and citations, and report that the average patent in this set of breakthrough inventions is in the 77th percentile of disruptiveness, but only in the 64th percentile of citations.²³

Tables II and III also provide statistics for the other text-based measures. Unlike patents’ technological disruptiveness, patent breadth is evenly centered around its average value of 0.42, indicating less skew in specializations across patents. Remarkably, Table III indicates that patents’ technological breadth and their originality (as defined by Trajtenberg, Henderson, and Jaffe (1997)) are largely unrelated.²⁴ For instance, software-related patents (e.g., #6964374, #7630986, or #7356679) display low levels of technological breadth, as they rely on a relatively narrow vocabularies. Yet, they rank high in terms of their originality. In contrast, breakthrough inventions display above-median levels of technological breadth. For example, the invention of the satellite, computer, and the jet engine combine broad scientific language.

²²Kelly, Papanikolaou, Seru, and Taddy (2018) similarly note that some patents classified as significant based on their measure attract few citations, and provide illustrative examples, such as patent #174465 issued to Graham Bell for the telephone in 1876 having received only 10 citations until March 2018.

²³For comparison, Kelly, Papanikolaou, Seru, and Taddy (2018) report that the average patent in this same set is in the 73th percentile of the distribution of their patent significance measure, and the average percentile for the KPSS measure for this set of patents is 64.

²⁴The construction of our breadth measure is somewhat similar to that of originality, but we measure concentration based on the assignment of *words* to technology areas rather than citations. This has several advantages. First, our measure of breadth is well defined even for patents with zero or one backward citations. Second, technology links revealed by vocabularies are not influenced by strategic avoidance of citations and do not rely on patent examiners having a complete knowledge of the patent space. Third, our measure of breadth can pick up a reliance on a technology area even if no specific citation to that area is given. For example, a patent might use “textbook” information about chemistry to describe a portion of the invention without needing to cite a “Chemical” patent.

Finally, we observe some variation in similarity across patents, but the overall levels are low, which is not surprising given the large range and diversity in the vocabulary used across all patents. Panel C of Table II further indicates that the text-based measures capture distinct dimensions of the technological nature of patents. Patents with higher levels of disruptiveness tend to display lower breadth. They are also located nearer to patents of lead innovators than to patents assigned to private or foreign firms. We also note that patents nearer to lead innovators tend have lower technological breadth, while those most similar to private and foreign firms display higher breadth.

III The Decline in Technological Disruptiveness

In this section, we document the aggregate time-series properties of our text-based technological characteristics between 1930 and 2010. We then contrast these technological changes with aggregate changes in IPO and acquisition activities occurring between 1980 and 2010.

A Technological Changes in the Last Century

To track the evolution of technological disruptiveness over time, we compute the aggregate stock of technological disruptiveness.²⁵ Figure II displays the time-series from 1930 and 2010, smoothed using a four-quarters moving average. Several interesting patterns emerge, suggesting important changes in the technological characteristics of U.S. innovation. Although there is considerable fluctuation, the evolution closely maps the history of U.S. innovation. We easily detect periods of sharply increasing disruptiveness. The first peak occurs around 1950 with an average level of disruptiveness that is almost double the level of 1930. The period around 1950 is often considered a time of radical innovation in manufacturing technologies, featuring the invention of the television, transistor, jet engine, nylon, and xerography. A second peak occurs in the mid-seventies, corresponding to innovation related to the computer. The last two peaks of technological disruption ap-

²⁵To compute the aggregate stock of any patent variable, we first compute the sum of *Tech Disruptiveness* for patents applied for in a given quarter. We then compute a rolling depreciated sum of the prior 20 quarters, using a 5% quarterly rate of depreciation. Finally, we normalize the rolling stock by the number of patents applied for in the 20 prior quarters.

pear in the late eighties and mid-nineties, reflecting waves of inventions related to genetics (e.g., methods of recombination) and the mass adoption of the Internet.

[Insert Figure II about here]

Despite these periodic surges in disruption, the 1930-2010 period is characterized by a protracted and steady long-term decline in the disruptiveness of U.S. patents. Between 1950 and 2010, the average level of technological disruptiveness has significantly decreased, with levels in 2010 being roughly one quarter that of 1950. Importantly, this decline is not due to changes in the composition of patents (e.g., shifts across technology classes) as we continue to observe a similar trend after we account for broad technology and location fixed effects. Rather, the decline in technological disruptiveness indicates a widespread deceleration in vocabulary usage growth rates among U.S. patents. This trend echoes recent research highlighting the increasing difficulty to generate new innovative ideas (e.g., Jones (2009) and Bloom, Jones, Reenen, and Webb (2017)).

Figure II also reveals important changes in the technological breadth of U.S. patents. The overall level of breadth steadily increases between 1930 to 1970. Beginning in the mid-seventies, there is a twenty-year period of growth in overall patent breadth which reaches a peak in the mid-nineties that was 20% above the 1970 level. In the most recent years, however, there is a large decline in the breadth of U.S. patents, dropping by about 25% between the mid-nineties and 2010. We also find an inverse U-shaped pattern in patent similarities over the last century. All three measures steadily increase until the eighties, as the text in the average U.S. patent during this period became increasingly similar to patents assigned to private U.S. firms, foreign firms, and lead innovators. Beginning in the eighties, however, these trends reversed, leading to marked declines in the similarity measures. The recent period is thus characterized by patents becoming both more specialized (i.e., lower technological breadth) and more distinct across firms.

B Technology, IPO, and Acquisitions (1980-2010)

We next examine the aggregate time-series properties of IPOs and acquisitions. We obtain data on IPOs from Jay Ritter's website, and exclude non-operating companies, as well as

IPOs with an offer price lower than \$5 per share, unit offers, small best effort offers, bank and savings and loans IPOs, natural resource limited partnerships, companies not listed in CRSP within 6 month of their IPO, and foreign firms' IPOs. Data on acquisitions are from the Thomson Reuters SDC Platinum Database, and include all domestic completed acquisitions (of private or public firms) coded as a merger, acquisition of majority interest, or acquisition of assets giving the acquirer a majority stake.

[Insert Figure III about here]

Panel A of Figure III plots the number of IPOs for each quarter between 1980 and 2010. The patterns are similar to those reported by Gao, Ritter, and Zhu (2013), Doidge, Karolyi, and Stulz (2017), and Ewens and Farre-Mensa (2018). Remarkably, the evolution of IPO activity rather closely maps the aggregate dynamics of technological disruptiveness during this thirty-year period. The number of IPOs drops around 1990, coinciding with a decline in disruptiveness that follows the earlier surge in genetic science in the mid-1980s. There were more IPOs as the nineties progressed, when technological disruptiveness experienced a very large increase. The decline in IPO intensity then began in the early 2000s, when the average technological disruptiveness of U.S. patents also started to plummet. In the aggregate, the intensity of new public listings is substantially higher at times where the average technological disruptiveness of U.S. patent applications is elevated.

Panel B of Figure III plots the evolution of the number of acquisitions, both in total and separately for private firms. The number of acquisitions has increased since 1980, with a strong acceleration in the mid-nineties. We note subsequent declines in acquisitions in the aftermath of the tech bubble and the financial crisis. Yet, the number of acquisitions remains significantly higher since the mid-nineties when compared to the 1980-1995 period, suggesting a relationship between the surge in aggregate acquisitions and the decline in technological disruptiveness of U.S. patents. Although the aggregate pattern for trade sales is less striking than that for IPOs, the patterns suggest that acquisitions tend to be high when overall technological disruptiveness is lower.

IV Technological Disruptiveness and Startups' Exits

To better understand the interplay between technological changes and exits, we explore the cross-sectional relationship between our text-based technological characteristics and the decision of private firms to exit by going public or through trade sales. Ideally, we would study the exit strategy of *all* private firms that are plausible candidates for IPOs or acquisitions. Because data limitations preclude this, we focus on a large sample of venture-backed private firms, for which we observe both their technological characteristics and their exit choices.

A Sample of VC-backed Startups

We obtain data on VC-backed U.S. firms from Thomson Reuters's VentureXpert (Kaplan, Stromberg, and Sensoy (2002)), which contains detailed information about private startups including the dates of financing rounds and their ultimate exit (e.g., IPO, acquisition, or failure). We focus on the period 1980-2010 and restrict our attention to VC-backed companies (henceforth startups) that are granted at least one patent during the sample period. To assign patents to VC-backed startups, we follow Bernstein, Giroud, and Townsend (2016) and develop a fuzzy matching algorithm that matches the names of firms in VentureXpert to patent assignees obtained from Google Patents (see Appendix C for details). The result is an unbalanced panel of startups-quarter observations.²⁶ A startup enters our sample in the quarter it is founded (based on founding dates in VentureXpert) and exits the sample when its final outcome (IPO, acquisition, or failure) is observed based on the "resolve date" variable in VentureXpert. Startups still active as of November 2017 remain unresolved.²⁷ We exclude startups if their founding date is missing or if it is later than the resolve date. The sample begins in 1980 to guarantee reliable data on outcomes and ends in 2010 because this is when the link to publicly traded firms created by Kogan, Papanikolaou, Seru, and Stoffman (2016) ends. Our final sample con-

²⁶Lerner and Seru (2017) note that bias can occur in matching patent assignments to startups because patents can be assigned to subsidiaries with different names than their parent corporations. However, this issue is limited in our sample as VC-backed startups are small and are unlikely to have complex corporate structures.

²⁷Ewens and Farre-Mensa (2018) note that unresolved firms can result from stale data collection. Thus, we code firms as failed if it has been seven years since their last funding round.

tains 347,929 startup-quarter observations, corresponding to 9,167 unique startups and 94,703 patent applications.

We obtain the technological characteristics for each startup-quarter by aggregating each patent-level variable (text-based and others) using their depreciated sums over the past 20 quarters using a quarterly depreciation rate of 5%. For example, the technological disruptiveness of startup i in quarter q corresponds to the depreciated sum of the disruptiveness of all its patent applications in the past five years, normalized by the number of patents i applied for over that period.²⁸ We define the exit variables (IPO or sell-out) as binary variables equal to one if startup i experiences a given exit in quarter q . The construction of all variables is explained in detail in Appendix B.

[Insert Table IV about here]

As we found for the overall statistics reported in Table II, we find in Table IV a substantial asymmetry in the distribution of technological disruptiveness among VC-backed startups despite the aggregation of their patents. The other variables are overall in line with their aggregate counterparts, indicating that the technological characteristics of VC-backed startups are roughly representative of those in the economy at large. Table IV further indicates that the quarterly IPO rate (i.e. the number of IPOs in a quarter divided by the number of active startups in that quarter) is 0.38 percentage points, and the quarterly sell-out rate is 0.54 percentage points.²⁹

[Insert Figure IV and Figure V about here]

Figure IV plots the time-series of technological characteristics for all unresolved VC-backed startups in each quarter. For the sake of comparison, we compute the aggregate

²⁸Because *Foreign Similarity* and *LI Similarity* are non-trivially correlated (60% and 45%) with *Private Similarity*, in regressions, we orthogonalize *Foreign Similarity* and *LI Similarity* by subtracting *Private Similarity*.

²⁹We report additional information about the sample firms in Appendix D in Table A2. Relative to the founding date, IPOs and acquisitions play out over time. Of these, IPOs occur fastest on average. The average firm applies for its first patent after 4.42 years, and receives its first round of VC funding 5.29 years after its founding. All of these numbers are mechanically shorter when measured relative to the first patent instead of the founding year.

stock of each variable for the set of patents granted to VC-backed startups as in Figure II. The trends closely map those of the aggregate dynamics presented in Figure II, indicating that the technological changes occurring among startups is mirroring economy-wide changes. In Figure V, we compare the evolution of IPO and sell-out rates for VC-backed startups to the aggregate patterns. We scale the quarterly number of IPOs and acquisitions by lagged real GDP to obtain aggregate exit rates. The upper figure shows that the evolution of IPO rates for VC-backed startups between 1980 and 2010 closely follows aggregate IPO rates, especially the decreased IPO rates after 2000. The lower figure also shows agreement between the sell-out rates observed for VC-backed startups and the aggregate trend. In particular, we observe growth in acquisition activity in the latter part of the sample.

Although our sample does not include all firms that have the potential to go public or get acquired, VC-backed startups nevertheless represents a useful laboratory to study the interplay between technological changes and the evolution of IPOs and acquisitions. First, these firms account for a large share of the IPO market (Ritter (2017)) and the production of innovation (Gornall and Strebulaev (2015)). Second, their IPO and acquisition rates from the last thirty years appear comparable to the economy-wide patterns, as shown in the figures above. Third, they are (almost) automatically facing exit choices due to the limited lives of most venture capital funds.

B Startups’ Disruptiveness: Validation Tests

Our analysis rests on the ability of our measure of patents’ disruptiveness to identify (in a timely manner) startups that can create new markets or disrupt existing ones. To corroborate this interpretation, we examine whether the startups that we classify as technologically disruptive (based on their patents) do in fact operate in product markets that display higher risk of disruption.

First, for each startup, we identify public firms offering similar products and services, following an approach similar to Hoberg and Phillips (2016). We obtain product descriptions of startups from VentureXpert as reported in the year of their first round of funding. We then compute the (cosine) similarity between the text in each startup’s product de-

scription (in year t) and that of all public firms, obtained from the product description section of their 10-K report (also in year t). Since 10-Ks became available in electronic format in 1997, we focus on 5,417 startups whose first funding round occurred after 1996. To identify the public firms operating in a given startup’s market (in year t), we retain the 25 public firms that have the most similar product descriptions relative to the startup. This is done by measuring the cosine similarity between the startup’s business description and public firm 10-K business descriptions and taking the 25 public firms with the highest cosine similarity.

Next, for each startup, we measure whether its 25 public “peers” are actually at risk of disruption. We do so by computing the fraction of paragraphs in each public firm’s 10-K that mentions words related to technology-based disruption, considering three measures. First, we search for paragraphs that contain words having the root “technol” and also a word having the root “change” to measure whether public peers are discussing exposures to technological change, a direct form of market disruption. Second, we identify paragraphs having the words with the root “technol” and also the word root “compet” to identify firms that are explicitly discussing competition on the margin of technological expertise. Third, we consider the more strict set of paragraphs containing at least one word with the roots “disrupt”, “technol”, and “compet”. In order to be counted as a hit, a paragraph must contain all three. Table V reveals a positive association between the technological disruptiveness of startups (using our text-based patent measure) and these ex-post mentions of market disruption by related public firms. This strong positive relationship holds across specifications that include either year fixed effects or a more complete set including year, technology, location, age, and cohort fixed effects.

[Insert Table V about here]

In a separate set of validation tests, we examine the product market characteristics of 848 startups in our sample that go public after 1997. Because these startups gain public status, we are able to link their ex-ante technological disruptiveness to the product market attributes that are only measurable when firms are publicly traded. For each IPO firm, we thus consider three characteristics of their post-IPO market (measured in the year of

their initial listing): product market concentration (HHI), the total similarity to publicly traded peers (TSimm), and product market fluidity. These variables are available since 1997 and obtained from Hoberg and Phillips (2016) and Hoberg, Phillips, and Prabhala (2014). Following the life cycle theory of Abernathy and Utterback (1978) we predict that disruptable markets are those for which superior new technologies are more likely to be discovered by competing early-stage firms. For validation, we thus assess whether IPO firms displaying higher levels of technological disruptiveness exit into more competitive, less differentiated, and more fluid product markets.

[Insert Table VI about here]

Based on these direct firm-specific measures, Table VI confirms this prediction. Newly-public firms with more technologically disruptive patents indeed exit into markets that are more fluid, contested, and thus disruptable (see Christensen (1997)). In contrast, IPO firms with less disruptive technology operate in more stable markets with less competition and more product differentiation. Put together, the results in Table V and Table VI confirm that our text-based measure of technological disruptiveness is consistently associated with product markets facing greater risk of disruption.

C Disruptiveness and Startups' Exits

We now examine the relationship between startups' technological disruptiveness and their propensity to exit via IPO or sell-out.³⁰ Our central hypothesis is that firms with more disruptive technologies will be more likely to seek a public listing and less likely to sell out. Our baseline specification relies on the competing risks regression approach of Fine and Gray (1999) that explicitly models the "risk" of choosing a particular exit in quarter q given that the firm is still unresolved at that time.³¹ Startups enter the sample (i.e., become at risk of exiting) when they are founded. Their exit is modeled using competing hazards to reflect multiple potential exit strategies that are mutually exclusive. This ap-

³⁰As mentioned in Section IV.A, VentureXpert frequently fails to code firms as failed when they stop operations. Given this data limitation, we do not directly report tests relating to failure outcomes.

³¹The use of a competing risk model is relatively rare in finance. One recent exception is Avdjiev, Bogdanova, Bolton, Jiang, and Kartasheva (2017), who examine the determinants of convertible capital choice by banks.

proach allows us to estimate the relationships between startups' text-based technological disruptiveness and the full set of potential exits.

To ensure that we are not capturing the effects of other technological characteristics that may correlate with startups' technological disruptiveness, we include in the specification the log number of patent applications in the past five years, a dummy variable indicating startups with no applications in the past five years, startups' technological breadth and similarities (with lead innovators, private and foreign firms), as well as the originality of startups' patents and their patents' citations (both aggregated for each startup-quarter as is the case for our text-based variables). Following the literature on IPOs and acquisitions, we also control for overall market activity using market relative valuation and stock returns as well as an identifier for the last quarter of the year (Lowry (2003) and Pastor and Veronesi (2005)). We cluster the standard errors by startup to account for any potential within-startup dependencies over time.

[Insert Table VII about here]

The first two columns of Table VII confirm the key role played by technological disruptiveness in startups' exit choices. In the first column, we observe a strong positive link between startups' technological disruptiveness and their likelihood of exiting through an IPO in the next quarter. The point estimate is 0.252 with a t -statistic of 13.09. This indicates that a one standard deviation increase in startups' technological disruptiveness is associated with a 25.2% increase in the quarterly rate of IPOs.³² On the other hand, column (2) reveals that the odds of exiting via a trade sale are negatively related to technological disruptiveness. Indeed, the estimated coefficient is negative (-0.188) and statistically significant with a t -statistic of -7.55. Hence a one standard deviation increase in disruptiveness indicates a 18.8% decrease in the sell-out rate. In addition to being statistically significant, the baseline estimates reveal economically large relationships.

Table VII shows that the startups' other technological characteristics are also impor-

³²As explained in Fine and Gray (1999), regression coefficients from a sub-distribution hazard model denote the magnitude of the relative change in the sub-distribution hazard function associated with a one-unit change in the given covariate. Therefore, estimated coefficients reflect the relative change in the instantaneous rate of the occurrence of the event in those subjects who are event-free.

tant determinants of startups' exits. We observe for instance that firms' technological breadth is positively related to IPO incidence, and is negatively related to trade sales. Intuitively, high breadth technologies are difficult for other firms to emulate and are less redeployable toward other uses. As such, firms with more technological breadth appear less amenable to acquisitions and are more amenable to exit via IPOs. Table VII also indicates that firms whose patents are more similar to those of other private firms are significantly less likely to exit through sell-outs (t -statistic of -8.62) and are marginally more likely to go public (t -statistic of 1.70). These results are in line with the negative link between product market similarity and the likelihood of being a target documented in Hoberg and Phillips (2010) for public firms. When a firm has more peers, it is more easily replaced and hence the rate of acquisition is lower for these firms. In contrast, firms holding patents that are more similar to that of lead innovators are significantly more likely to go public (t -statistic of 4.58). We also find that future citations (originality) are positively (negatively) associated with exit via both IPO and trade sale. Also, startups are more likely to exit via IPO after periods of strong overall stock market performance, consistent with growth options or market timing stories. Periods characterized by higher overall valuations have more IPOs and sell-outs. Importantly, the inclusion of these control variables insures that the role technological disruptiveness in startups' exist is robust.

In the last two columns of Table VII, we report estimates from linear probability models where the dependent variables are indicators for whether a given exit occurs in a given quarter. Although this approach ignores the potential dependence across exits (i.e., competing risks), linear models allow us to include a wider array of fixed effects. We include year, state, technology, age, and cohort fixed effects to estimate the link between exits and technological disruptiveness among startups of the same age, those receiving first funding at the same time, those operating in the same year and state, and those innovating in the same technological fields.³³ We find that our conclusions are largely unaffected, indicating that the associations between the technological disruptiveness of startups and

³³Technology fixed effects are based on the most common NBER technology category used in a firm's patents (see Lerner and Seru (2017)).

their exits is unique. We also estimate (but do not report for brevity) separate logistic and multinomial logistic models for each exit type that include year, state, and technology fixed effects. These tests produce similar results.

[Insert Table VIII about here]

In Table VIII, we use the same specifications as reported in Table VII, but we further control for startups' financing, as previous research reports that the amount of VC funding (a proxy for startups' implied valuation) predicts startups' exits (Cumming and Macintosh (2003)). These tests are important because our interpretation of the link between technological disruptiveness and startup exits could be due to VCs providing more funding to startups developing more disruptive innovation.³⁴ Table VIII indicates that this is not the case. To account for the possible role of funding, we include startups' cumulative VC funding (from founding to quarter $q - 1$) and a binary variable identifying whether startups received funding in the last five years. Across all specifications, we confirm that the financing variables are strong determinants of startups' exit, especially sell-outs. However, our main result for technological disruptiveness is fully robust, indicating that our findings cannot be explained by financing.

Table IX explores the dynamic links between startups' technological disruptiveness and exits by increasing the measurement window for identifying startup exits from one quarter to five years using increments of one year. We focus on linear specifications that include the full set of fixed effects as described above, and only report coefficients for the technology variables for brevity. Panel A indicates that the positive associations between startups' technological disruptiveness and IPO incidence remains strong at all horizons. In contrast, Panel B reveals that the negative relation between technological disruptiveness and the propensity to sell out is only present at short horizons and then fades after two years. We further report in Appendix D that our conclusions are stable throughout our sample, as the results hold across sub-periods.³⁵

³⁴We confirm this intuition in Table A3 of the Appendix.

³⁵Therefore, it is unlikely that our results are affected by truncation biases coming from the "unresolved" status of recent startups or patent applications not yet granted as pointed out by Lerner and Seru (2017).

[Insert Table IX about here]

Finally, we also consider whether technological traits are related to the propensity of a startup to remain private for longer periods. This analysis is motivated by the evidence in Gao, Ritter, and Zhu (2013) and Ewens and Farre-Mensa (2018) that, in recent years, many startups remain private longer. Panel C of Table IX presents results from regressions of startups' odds of remaining private at different horizons on their current technological characteristics. The results show that startups with high technological disruptiveness exit more quickly than less disruptive startups. We also find that startups with higher overall technological similarity to other startups are more likely to stay private longer, and those that are more similar to foreign firms exit more quickly. Startups with high levels of private firm similarity are in more competitive technology spaces, and the results suggest that such startups require additional time to establish stable market positions for IPOs or distinguish themselves from alternative acquisition targets. Overall, these findings complement those of Gao, Ritter, and Zhu (2013) and Ewens and Farre-Mensa (2018).

V Disappearing IPOs and Surging Sell-Outs

Our results thus far indicate that, between 1980 and 2010, the technological disruptiveness of startups is systemically related to their propensity to exit via IPO and sell-out. Since aggregate technological disruptiveness has declined since the late nineties (see Figure IV), we examine whether the recent aggregate shift from IPOs to trade sales is related to changes in startups' technological characteristics.

A Prediction Errors for Startups' Exit

We use methods from the disappearing dividends literature (e.g., Fama and French (2001) or Hoberg and Prabhala (2009)) to assess the extent to which our technological characteristics relates to the disappearing IPOs and surging sell-out anomalies. We thus compare actual IPO rates to those predicted by models of IPO and Sell-out exits. To assess the impact of our new variables, we estimate predicted rates with and without our text-based technological variables.

We proceed in two steps. First, as in Fama and French (2001), we estimate two linear probability models (a “Base” model and a “Text” model) using quarter-by-quarter Fama and MacBeth (1973) regressions where the dependent variable is the incidence of IPO exits in each quarter during the initial period 1980-1995 (the “pre-period”). In the base model, the independent variables are the (log) of startups’ age and their patent stock. We augment this in the Text model by including the text-based startup’s technology traits. Second, we compute predicted values of IPO incidence for each startup-quarter in the 1996-2010 period (the “post-period”) by applying the average coefficients estimated in the pre-period for both models to the actual values of the independent variables in the post-period. We then average the predicted IPO rates across all startups in each quarter and compare them to the actual average IPO rate within the post-period. Since the coefficients associated with each characteristic are fixed at their pre-period value, variation in predicted IPO rates in the post-period is only due to changing startups’ characteristics. We repeat these steps for sell-outs to compare actual and predicted sell-out rates.

[Insert Table X about here]

Table X presents the results. Panel A indicates that using only the control variables (i.e., excluding the text-based technology traits) to estimate the quarter-by-quarter IPO regressions in the pre-period yields an average predicted quarterly IPO rate of 0.84 percentage points in the post-period. This predicted incidence is substantially higher than the actual IPO rate, which is of 0.33 percentage points per quarter in the post-period, implying a prediction error of 0.51. The predicted IPO rate is thus 2.5 times higher than the actual rate. This confirms that IPO rates in the post-period are “abnormally” low and we thus replicate the disappearing IPOs anomaly. When we include the text-based technological characteristics in the model, the average predicted IPO quarterly rate in the post-period declines to 0.75 percentage points, which is still higher than the actual incidence rate as the prediction error is 0.42. Although a significant portion remains unexplained, our estimates imply that changes in our text-based technological traits account for 16% of the disappearing IPOs anomaly. The rest of Panel A indicates that our estimates are not sensitive to the definition of the pre- and post-periods or to the fore-

casting horizons (i.e., the lags between the dependent and independent variables). Across all specifications, changes in technological characteristics account for 19% of the recent abnormally low IPO rates.

Panel B reports parallel analysis for sell-out rates. A benchmark linear model that excludes our technology variables estimated in the pre-period yields an average predicted sell-out incidence of 0.60 percentage points per quarter in the post-period. Compared to the actual rate of 0.86 per quarter, the base model’s prediction is 42% lower than the actual rate, suggesting that the prevalence of sell-outs in recent years is “abnormally” high. When we include our text-based technology variables in the model, the prediction gap narrows significantly, as we obtain a predicted sell-out rate of 0.75 percentage points per quarter. When we alter specifications in the remainder of Panel B, we observe that changes in startups’ technological characteristics explain between 26% and 71% of the surging trade sales. We conclude that roughly 50% (the average across specifications) of the surge in trade sales is accounted for by changes in startups’ technological characteristics.

B Small and Large IPOs

Existing research indicates that the recent dearth of IPOs is particularly pronounced for smaller-company IPOs (see Gao, Ritter, and Zhu (2013) and Doidge, Kahle, Karolyi, and Stulz (2018)). To the extent that IPO activity truly relates to the technological characteristics of startups, recent changes in these characteristics should account for a significant portion of the sharp decline in small IPOs.

To test this conjecture, we perform an analysis similar to that reported in Table X separately for small and large IPO exits. Data on pre-IPO sales is provided by Gao, Ritter, and Zhu (2013) and inflation adjusted to 2009 dollars. We define an IPO as “small” if its pre-IPO sales are below the median in our sample (\$25 million), and as “large” if its pre-IPO sales exceeds that amount. We then estimate the probability that a given startup exits through a small (large) IPO in a given quarter in the pre-period quarter-by-quarter. As before, we estimate the model with and without our text-based technological variables, and compare the predicted IPO rate in the out-of-sample period to the actual rate in order to assess the disappearing IPOs anomaly.

[Insert Table XI about here]

Table XI displays the results. Panel A indicates that, across six specifications which vary the definition of the pre- and post-periods or the forecasting horizon, changes in startups’ technological traits account for roughly 37% of the disappearing small IPO anomaly in the recent period. In sharp contrast, Panel B reveals that adding startups’ technological characteristics in the regression models (estimated in the pre-period) does not bring the average predicted rate of large IPOs closer to its actual value in the post-period. Across six different specifications, changes in technological traits appears unrelated to the variation in large IPO observed in the post-period. Findings in Table XI indicate that changes in startups’ technological characteristics is a particularity important factor related to the decline in small-size IPOs.

C The Role of Product Market Stability

We also explore the role of product markets in the recent shift from IPOs to trade sales. Our central hypothesis is that the recent dearth of IPOs partly stems from the aggregate decline in startups’ technological disruptiveness, and should thus vary across industries with varying product maturity. Following the product life cycle theory in Abernathy and Utterback (1978) we posit that maturing markets in transition or those reaching maturity (e.g., markets that have reached a dominant product design) are likely to experience the most extreme decline in IPO rates. In these markets, breakthrough inventions obtain only with very high search costs (Jones (2009) and Bloom, Jones, Reenen, and Webb (2017)) as the best ideas are already “picked over”, so that the decline of technological disruptiveness should be a stronger predictor of lower IPO activity. In contrast, we posit that very young “fluid” markets should feature more IPOs that might even grow during the earliest years as competing new disruptive ideas are commercialized (a process requiring some initial “time to build”).

To test this idea, we follow Hoberg, Phillips, and Prabhala (2014) and compute the degree of product market fluidity in each startup’s product market from 1980 to 2010, using the business description text that is available at the time of the first funding round

in Thomson Reuters’s VentureXpert. To do so, we first compute the aggregate change in product description vocabulary used by startups as the year-over-year change in the frequency of usage across all business descriptions. This quantity is computed separately for each word and the result is stored in an aggregate vector containing the set of word frequency changes for all words (this procedure is similar to that in Equation (1)). Second, for a given startup, we compute the frequency-weighted average of the aggregate change vector where the weights are the frequency of words used by the startup in its own business description (this calculation is similar to that in Equation (2)). The resulting variable is a product fluidity measure similar to the one used in Hoberg, Phillips, and Prabhala (2014), but defined over all VC-backed startups receiving their first money between 1980 and 2010.

[Insert Table XII about here]

To assess whether changes in startups’ technological traits account for the recent decline in IPOs differently in stable and unstable markets, we divide our startup-quarter observations into above and below median fluidity sub-samples, based on median breakpoints chosen separately for each cohort of startups (based on the year of the first funding round) and repeat the prediction procedure discussed above across each sub-sample. Panel A confirms that startups operating in stable markets are less likely to exit via IPO relative to startups in fluid markets, with IPO rates of 0.30% and 0.37% per quarter in the post-period (1996-2010). Moreover, changes in startups’ technological characteristics explain roughly 25% of the dearth of IPOs in stable markets. This figure is tightly estimated across different specifications. In contrast, changes in startups’ technological attributes account for just 3% of the dearth of IPOs in less stable markets. This figure ranges between -7% and 15% across different specifications. Although market stability appears critical for understanding the evolution of IPOs, Panel B indicates that stability has little effect in moderating our ability to explain the surging trade sales anomaly, as the average improvement is 50% and 49% in stable and fluid markets, respectively.

VI Conclusions

We develop new measures of technological disruptiveness and other characteristics using textual analysis of 6,595,226 U.S. patents that were applied for between 1930 and 2010. We document that these characteristics are highly influential in predicting which startups will exit via IPO or sell-out. These results are economically large and remain important after controlling for a host of known explanatory variables.

We find that startups with more disruptive patents are more likely to exit via IPO. Understanding the economics of disruptive patents is more intuitive when juxtaposed against patents that primarily refine or extend existing technologies. These technologies are also valuable, but they have less potential to establish independent markets for their owners. These findings suggest that IPOs are favored by firms with technologies that have the potential to establish independent market positions with strong defenses against rivals. Sell-outs are favored by firms whose technologies have less potential to create independent markets, and are more redeployable. Such technologies are more valuable in the hands of existing firms with strong independent market positions.

We also show that technological traits change dramatically over our long sample period. For example, we document an economy-wide decline in technological disruptiveness that began after World War II. Because technological characteristics can explain exit choices at the micro level, we further propose that widespread changes in disruptiveness might explain aggregate trends such as the recent decline in IPOs and the surge in sell-outs. By comparing actual IPO and sell-out intensities to predicted levels that condition on our technological characteristics, we estimate that roughly 20% of the decline in IPOs is attributable to technological changes. Analogously, roughly 50% of the surge in sell-outs can be explained by these same variables.

References

- Abernathy, William J., and James M. Utterback, 1978, Patterns of industrial innovation, *Technology Review* 80, 41–47.
- Avdjiev, Stefan, Bilyana Bogdanova, Patrick Bolton, Wei Jiang, and Anastasia Kartasheva, 2017, CoCo issuance and bank fragility, available at <https://www.nber.org/papers/w23999>.
- Balsmeier, Benjamin, Mohamad Assaf, Tyler Chesebro, Gabe Fierro, Kevin Johnson, Scott Johnson, Guan Cheng Li, Sonja Luck, Doug OReagan, Bill Yeh, Guangzheng Zang, and Lee Fleming, 2018, Machine learning and natural language processing on the patent corpus: Data, tools, and new measures, *Journal of Economics and Management Strategy* pp. 535–553.
- Bates, Thomas, Kathleen Kahle, and Rene Stulz, 2009, Why do U.S. firms hold so much more cash than they used to?, *Journal of Finance* pp. 1985–2021.
- Bayar, Onur, and Thomas Chemmanur, 2011, IPOs versus acquisitions and the valuation premium puzzle: A theory of exit choice by entrepreneurs and venture capitalists, *Journal of Financial and Quantitative Analysis* pp. 1755–1793.
- Bernstein, Shai, Xavier Giroud, and Richard Townsend, 2016, The impact of venture capital monitoring, *Journal of Finance* pp. 1591–1622.
- Bloom, Nicholas, Charles Jones, John Van Reenen, and Michael Webb, 2017, Are ideas getting harder to find?, Working Paper.
- Chemmanur, Thomas, Jie He, Shan He, and Debrashi Nandy, 2018, Product market characteristics and the choice between IPOs and acquisitions, *Journal of Financial and Quantitative Analysis* pp. 681–721.
- Christensen, Clayton, 1997, *The innovator’s dilemma: When new technologies cause great firms to fail* (Harvard Business School Press: Boston).
- , and Richard Rosenbloom, 1995, Explaining the attacker’s advantage: Technological paradigms, organizational dynamics, and the value network, *Research Policy* pp. 233–257.
- Cumming, Douglas, and Jeffrey Macintosh, 2003, Venture-capital exits in Canada and the United States, *The University of Toronto Law Journal* pp. 101–199.
- Dahlin, Kristina, and Dean Behrens, 2005, When is an invention really radical? defining and measuring technological radicalness, *Research Policy* pp. 717–737.
- Darby, Michael, and Lynne Zucker, 2018, Going public when you can in biotechnology, NBER Working Paper.
- Doidge, Craig, Kathleen Kahle, Andrew Karolyi, and Rene Stulz, 2018, Eclipse of the public corporation or eclipse of the public markets?, Working Paper.
- Doidge, Craig, Andrew Karolyi, and Rene Stulz, 2013, The U.S. left behind? Financial globalization and the rise of IPOs outside the U.S., *Journal of Financial Economics* pp. 546–573.
- , 2017, The U.S. listing gap, *Journal of Financial Economics* pp. 464–487.
- Ewens, Michael, and Joan Farre-Mensa, 2018, The deregulation of the private equity markets and the decline in IPOs, Working Paper.
- Fama, Eugene, and Kenneth French, 2001, Disappearing dividends: Changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 3–44.
- Fama, Eugene, and J. MacBeth, 1973, Risk, return and equilibrium: Empirical tests, *Journal of Political Economy* 71, 607–636.
- Fine, Jason, and Robert Gray, 1999, A proportional hazards model for the subdistribution of a competing risk, *Journal of the American Statistical Association* pp. 496–509.

- Gao, Xiaohui, Jay Ritter, and Zhongyan Zhu, 2013, Where have all the IPOs gone?, *Journal of Financial and Quantitative Analysis* 48, 1663–1692.
- Gornall, Will, and Ilya Strebulaev, 2015, The economic impact of venture capital: Evidence from public companies, Working Paper.
- Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg, 2001, The NBER patent citation data file: Lessons, insights and methodological tools, No. w8498 National Bureau of Economic Research.
- Hanley, Kathleen, and Gerard Hoberg, 2010, The information content of IPO prospectuses, *Review of Financial Studies* 23, 2821–2864.
- Hoberg, Gerard, and Gordon Phillips, 2010, Product market synergies in mergers and acquisitions: A text based analysis, *Review of Financial Studies* 23, 3773–3811.
- , 2016, Text-based network industry classifications and endogenous product differentiation, *Journal of Political Economy* 124, 1423–1465.
- , and Nagpurnanand Prabhala, 2014, Product market threats, payout, and financial flexibility, *Journal of Finance* pp. 293–324.
- Hoberg, Gerard, and Nagpurnanand Prabhala, 2009, Disappearing dividends, catering, and risk, *Review of Financial Studies* 22, 79–116.
- Jones, Benjamin, 2009, The burden of knowledge and the death of the renaissance man: Is innovation getting harder?, *Review of Economic Studies* pp. 283–317.
- Kaplan, Steven, Per Stromberg, and Berk Sensoy, 2002, How well do venture capital databases reflect actual investments?, Working Paper.
- Kaplan, Steven N, and Luigi Zingales, 1997, Do investment-cash flow sensitivities provide useful measures of financing constraints?, *Quarterly Journal of Economics* 112, 169–215.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2018, Measuring technological innovation over the long run, Working Paper.
- Kogan, Leonid, Dimtris Papanikolaou, Amit Seru, and Noah Stoffman, 2016, Technological innovation, resource allocation and growth, *Quarterly Journal of Economics* forthcoming.
- Lerner, Josh, and Amit Seru, 2017, The use and misuse of patent data: Issues for corporate finance and beyond, Working Paper.
- Lowry, Michelle, 2003, Why does IPO volume fluctuate so much?, *Journal of Financial Economics* 67, 3–40.
- , Roni Michaely, and Ekaterina Volkova, 2017, Initial public offerings: A synthesis of the literature and directions for future research, Working Paper, Cornell University.
- Packalen, Mikko, and Jay Bhattacharya, 2018, New ideas in invention, NBER Working Paper 202922.
- Pastor, Lubos, and Pietro Veronesi, 2005, Rational IPO waves, *Journal of Finance* pp. 1713–1757.
- Poulsen, Annette, and Mike Stegemoller, 2008, Moving from private to public ownership: Selling out to public firms versus initial public offerings, *Financial Management* pp. 81–101.
- Ritter, Jay, 2017, Initial public offerings: Updated statistics, University of Florida.
- , and Ivo Welch, 2002, A review of IPO activity, pricing, and allocations, *Journal of Finance* 57, 1795–1828.
- Sebastiani, Fabrizio, 2002, Machine learning in automated text categorization, *ACMCS* 34, 1–47.
- Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe, 1997, University versus corporate patents: A window on the basicness of invention, *Economics of Innovation and new technology* 5, 19–50.

Figure I: Example of a Google Patent page


This figure shows the structure of a Google Patent page. The depicted patent is 6,285,999, commonly known as PageRank. Available at <https://patents.google.com/patent/US6285999>.

Method for node ranking in a linked database

Abstract

A method assigns importance ranks to nodes in a linked database, such as any database of documents containing citations, the world wide web or any other hypermedia database. The rank assigned to a document is calculated from the ranks of documents citing it. In addition, the rank of a document is calculated from a constant representing the probability that a browser through the database will randomly jump to the document. The method is particularly useful in enhancing the performance of search engine results for hypermedia databases, such as the world wide web, whose documents have a large variation in quality.

Images (4)



Classifications

- [G06F17/30864](#) Retrieval from the Internet, e.g. browsers by querying, e.g. search engines or meta-search engines, crawling techniques, push systems
 - [G06F17/30728](#) Information retrieval; Database structures therefor; File system structures therefor of unstructured textual data based on associated metadata or manual classification, e.g. bibliographic data using citations
 - [Y10S707/99935](#) Query augmenting and refining, e.g. inexact access
 - [Y10S707/99937](#) Sorting
- [Hide more classifications](#)

Description

CROSS-REFERENCES TO RELATED APPLICATIONS

This application claims priority from U.S. provisional patent application Ser. No. 60/035,205 filed Jan. 10, 1997, which is incorporated herein by reference.

STATEMENT REGARDING GOVERNMENT SUPPORT

This invention was supported in part by the National Science Foundation grant number IRI-9411306-4. The Government has certain rights in the invention.

FIELD OF THE INVENTION

This invention relates generally to techniques for analyzing linked databases. More particularly, it relates to methods for assigning ranks to nodes in a linked database, such as any database of documents containing citations, the world wide web or any other hypermedia database.

BACKGROUND OF THE INVENTION

Due to the developments in computer technology and its increase in

US6285999B1

US Grant

[Download PDF](#) [Find Prior Art](#)

Inventor: [Lawrence Page](#)

Current Assignee: [Leland Stanford Junior University](#), [Google LLC](#)

Original Assignee: [Leland Stanford Junior University](#)

Priority date: 1997-01-10

Family: US (10)

Date	App/Pub Number	Status
1998-01-09	US09004827	Expired - Lifetime
2001-09-04	US6285999B1	Grant
Show 8 more applications		
2012	US13616965	Expired - Lifetime

Info: [Patent citations \(28\)](#), [Non-patent citations \(20\)](#), [Cited by \(812\)](#), [Legal events](#), [Similar documents](#), [Priority and Related Applications](#)

External links: [USPTO](#), [USPTO Assignment](#), [Espacenet](#), [Global Dossier](#), [Discuss](#)

Claims (29)

What is claimed is:

1. A computer implemented method of scoring a plurality of linked documents, comprising:

obtaining a plurality of documents, at least some of the documents being linked documents, at least some of the documents being linking documents, and at least some of the documents being both linked documents and linking documents, each of the linked documents being pointed to by a link in one or more of the linking documents;

assigning a score to each of the linked documents based on scores of the one or more linking documents and

processing the linked documents according to their scores.

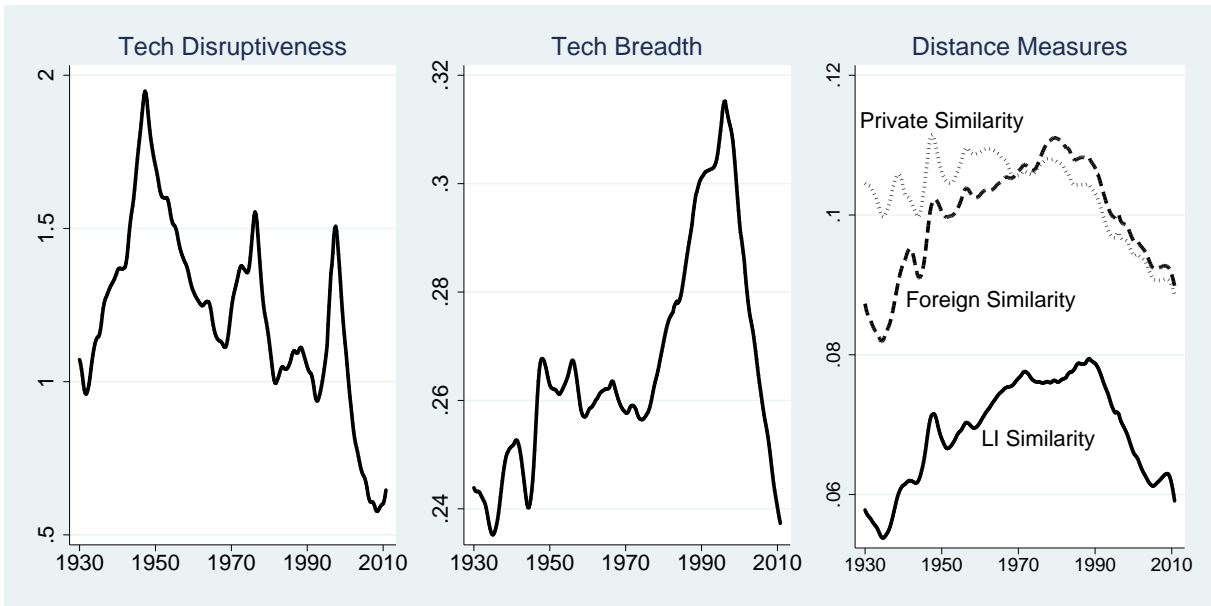
2. The method of claim 1, wherein the assigning includes:

identifying a weighting factor for each of the linking documents, the weighting factor being dependent on the number of links to the one or more linking documents, and

Figure II: Trends in Aggregate Technology Variables

This figure reports characteristics of the aggregate patent corpus from 1930 to 2010. The variables are defined at the patent level in Section II. To compute the aggregate stocks, we first compute the sum of each of the patent-level characteristics for patents applied for in a given quarter. We then compute a rolling depreciated sum of the prior 20 quarters, using a 5% quarterly rate of depreciation. Finally, we normalize the rolling stock by the number of patents applied for in the 20 prior quarters. The underlying patent level measures are winsorized at 1/99% level annually. The series presented are four quarter moving averages to smooth out seasonality.

Panel A: 1930-2010



Panel B: 1980-2010





Figure III: Trends in Agregate IPOs and Acquisitions

This figure reports the evolution of the quarterly number of IPOs in Panel A and acquisitions in Panel B. We obtain data on IPOs from Jay Ritter's website, and exclude non-operating companies, as well as IPOs with an offer price lower than \$5 per share, unit offers, small best effort offers, bank and savings and loans IPOs, natural resource limited partnerships, companies not listed in CRSP within 6 month of their IPO, and foreign firms' IPOs. Data on acquisitions are from the Thomson Reuters SDC Platinum Database, and include all domestic completed acquisitions (of private or public firms) coded as a merger, acquisition of majority interest, or acquisition of assets giving the acquirer a majority stake. In Panel B we reports acquisition volume for all targets (dotted line) and private targets (solid line). The series presented are four quarter moving averages.

Panel A: IPO volume



Panel B: Acquisition volume




Figure IV: Trends in Technology Variables for VC-backed Startups

This figure reports characteristics of the aggregate corpus of patents held by VC-backed startups from 1980 to 2010. The variables are defined at the patent level in Section II. To compute the series below, we first compute the sum of each of the patent-level characteristics for patents applied for in a given quarter by VC-backed startups. We then compute a rolling depreciated sum of the prior 20 quarters, using a 5% quarterly rate of depreciation. Finally, we normalize the rolling stock by the number of patents applied for in the 20 prior quarters by VC-backed startups. The underlying patent level measures are winsorized at 1/99% level annually. The series presented are four quarter moving averages to smooth out seasonality.




Figure V: Trends in IPOs and Sell-Outs of VC-backed Startups

This figure reports the propensity of firms to exit via IPO and acquisition over time. We obtain data on IPOs from Jay Ritter's website, and exclude non-operating companies, as well as IPOs with an offer price lower than \$5 per share, unit offers, small best effort offers, bank and savings and loans IPOs, natural resource limited partnerships, companies not listed in CRSP within 6 month of their IPO, and foreign firms' IPOs. Data on acquisitions are from the Thomson Reuters SDC Platinum Database, and include all domestic completed acquisitions (of private or public firms) coded as a merger, acquisition of majority interest, or acquisition of assets giving the acquirer a majority stake. Aggregate trends on IPOs and acquisitions of private targets and are reported in dashed lines as a fraction of lagged real GDP (left axis). Real GDP is in units of \$100m. Trends in the VC-backed startups sample are reported in solid lines as a percentage of startups that exit in the sample during the year (right axis). All series are reported as four quarter moving averages.

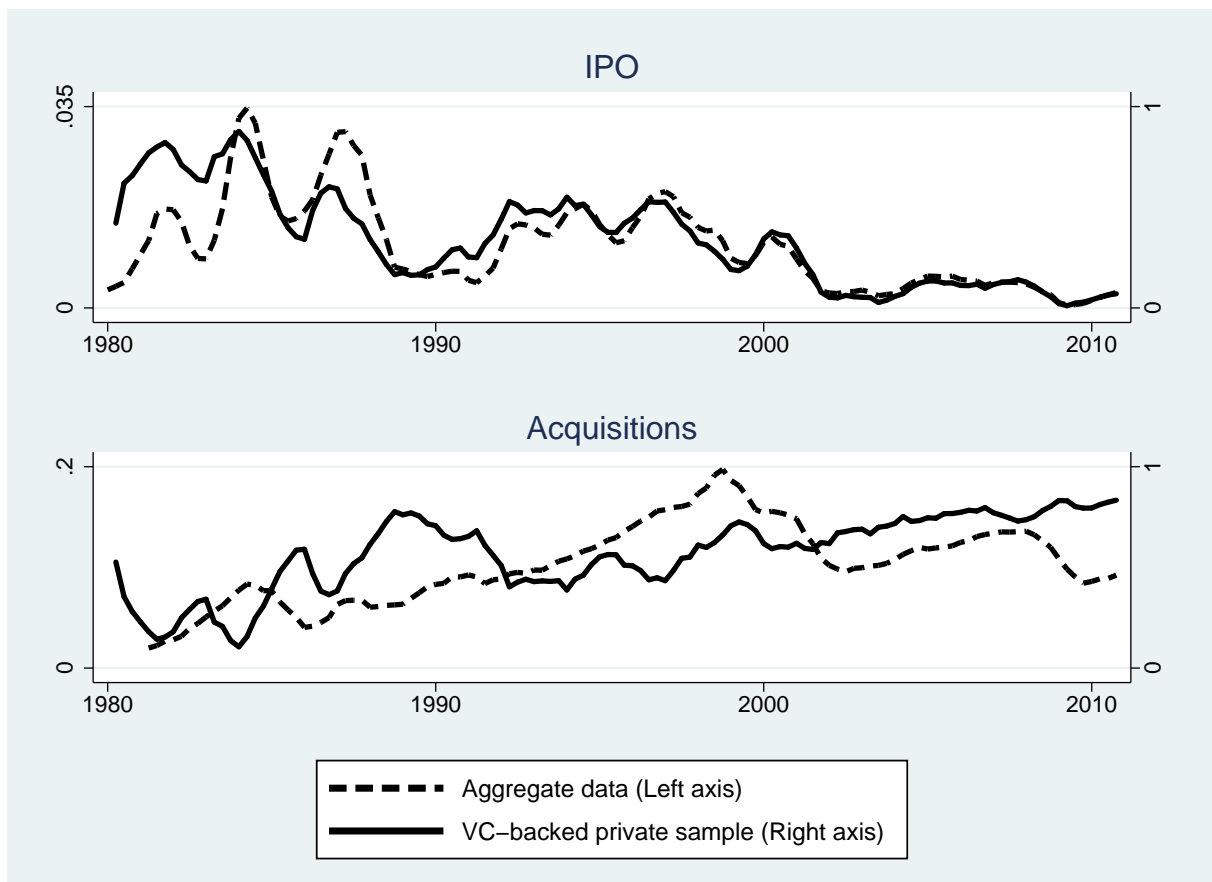


Table I: Changes in Patent Word Usage: Examples

This table reports, for five illustrative years between 1930 and 2010, how innovation has changed based on the text within patents. Panel A lists the ten words that have the largest year-over-year increase in use across all patents. Panel B lists the ten words that have the largest year-over-year decrease in use across all patents.

Panel A: Words with largest increase in use

1935	1975	1985	1995	2005
cent	bolts	laser	polypeptides	broadband
leaves	effort	japanese	deletion	intervening
axes	lithium	wavelength	clones	candidates
packing	user	publication	polypeptide	click
column	describes	blood	peptides	configurable
lead	exemplary	infrared	recombinant	luminance
coupled	entitled	polymer	cdna	abstract
notch	typically	mount	nucleic	acquiring
copper	phantom	optical	transcription	telecommunications
chain	exploded	comparative	plasmid	gamma

Panel B: Words with largest decline in use

1935	1975	1985	1995	2005
chambers	assistant	sulfuric	cassette	vegetable
crank	inventor	collection	ultrasonic	acyl
boiling	inventors	crude	machining	spiral
agent	firm	stock	abutment	gram
seats	priority	dioxide	tape	wedge
yield	john	evident	sand	gelatin
reducing	foreign	hydrocarbon	packing	crude
engine	sept	shut	bottle	oven
bell	june	circuitry	slidable	maybe
film	corporation	oxides	insofar	drilling

Table II: Summary Statistics: Patent-level Sample

This table presents descriptive statistics for patents applications between 1930 and 2010. The new technology variables *Tech Disruptiveness*, *Tech Breadth*, *Private Similarity*, *LI Similarity*, and *Foreign Similarity* are defined in Section II. Remaining variables are defined in Appendix B. Panel A reports summary statistics, where P25 and P75 denote the 25th and 75th percentiles. Panel B reports correlations between *Tech Disruptiveness* and preexisting measures of patents. Panel C reports correlations between the new technology variables. All variables are winsorized at the 1/99% level annually.

Panel A: Summary statistics

	N	Mean	SD	P25	Median	P75
Tech Disruptiveness	6,594,248	1.64	1.81	0.51	1.27	2.34
Tech Breadth	6,594,143	0.42	0.22	0.24	0.47	0.60
Private Similarity	6,594,248	0.15	0.05	0.12	0.15	0.18
LI Similarity	6,594,248	0.11	0.05	0.06	0.09	0.13
Foreign Similarity	6,594,248	0.15	0.06	0.11	0.14	0.19
KPSS Value	1,781,386	9.75	23.69	0.73	3.25	9.16
# of Cites	6,595,226	1.53	2.49	0.00	1.00	2.00
Originality	5,335,987	0.40	0.33	0.00	0.46	0.67

Panel B: Correlation with measures of patent value and quality

	(1)	(2)	(3)
(1) Tech Disruptiveness	1.00		
(2) Log(1+KPSS Value)	0.06***	1.00	
(3) Log(1+Cites)	0.04***	0.10***	1.00

Panel C: Correlation among text-based technology variables

	(1)	(2)	(3)	(4)	(5)
(1) Tech Disruptiveness	1.00				
(2) Tech Breadth	-0.24***	1.00			
(3) Private Similarity	-0.13***	0.28***	1.00		
(4) LI Similarity	0.06***	-0.09***	0.45***	1.00	
(5) Foreign Similarity	-0.08***	0.25***	0.61***	0.74***	1.00

Table III: Technological Disruptiveness: Examples of Important Patents

In Panel A, we define important patents as the 25 patent applications occurring between 1980 and 2010 that have attracted the largest number of citations in the first five years following their grant. In Panel B, we consider of twelve unambiguous breakthrough patents, as identified by the USPTO’s “Significant Historical Patents of the United States”. For each of these patents, we report the percentile of *Tech Disruptiveness* (“Disrpt”), *Cites* (received within five years of the patent grant), and *KPSS Value* (“KPSS”), *Tech Breadth* (“Brdth”) and *Originality* (“Orig”). The column labeled “Diff” reports the difference between the percentile of *Tech Breadth* and *Originality*. Red cells indicate higher values within a set of measures and blue cells indicate lower values within a set of measures.

Patent	Year	Dsrpt	Cites	KPSS	Brdth	Orig	Diff	Note
Panel A: Top 25 patents from 1980-2010 by citations								
7,674,650	2006	17	100	—	40	55	-15	Semiconductor/transistor
7,732,819	2008	14	100	—	42	55	-13	Semiconductor/transistor
7,501,293	2003	83	100	—	42	95	-53	Semiconductor/transistor
7,468,304	2006	13	100	18	26	1	25	Semiconductor/transistor
7,663,607	2004	67	100	96	29	80	-51	Multipoint touchscreen (Apple)
5,572,643	1995	99	100	—	15	72	-57	Early HTML use
7,462,862	2005	5	100	76	53	42	11	Semiconductor/transistor
7,453,065	2005	38	100	15	39	69	-30	Image sensor
7,453,087	2006	31	100	15	39	55	-16	Semiconductor/transistor
6,964,374	1998	65	100	77	12	88	-76	Storing and accessing metadata
7,630,986	2000	94	100	—	15	88	-73	Secure data interchange
7,632,985	2006	87	100	91	34	55	-21	Soybean biotech
7,411,209	2007	8	100	21	40	—	—	Semiconductor/transistor
7,181,438	2000	93	100	—	15	79	-64	User database system
7,402,506	2005	6	100	40	61	1	60	Semiconductor/transistor
7,356,679	2004	76	100	94	7	95	-88	Disk image capture
5,530,852	1994	100	100	93	16	1	15	Early HTML use
5,742,905	1994	99	100	—	12	52	-40	Personal communications
5,774,660	1996	98	100	—	20	86	-66	Early server management
7,479,949	2008	19	100	94	11	43	-32	Key iPhone patent
7,385,224	2005	10	100	14	25	81	-56	Semiconductor/transistor
5,608,786	1995	99	100	—	15	51	-36	Electronic communications
5,862,325	1996	99	100	—	9	84	-75	Electronic communications
5,708,780	1995	99	100	65	19	60	-41	Early server management
7,389,268	2002	7	100	—	12	23	-11	Tools for electronic trading
Panel B: Breakthrough patents								
1,773,980	1927	86	1	—	61	—	—	TV
1,848,389	1929	70	1	—	27	—	—	Helicopter
2,404,334	1941	73	95	—	77	—	—	Jet Engine
2,524,035	1948	82	100	71	46	87	-41	Transistor
2,569,347	1948	83	100	55	49	67	-18	Junction Transistor
2,668,661	1944	100	78	59	82	73	9	Modern digital computer
2,835,548	1957	81	68	—	99	—	—	Satellite
2,929,922	1958	90	99	87	60	—	—	Laser
4,237,224	1979	95	100	—	67	—	—	Cohen/Boyer patent
4,399,216	1980	99	99	—	73	1	72	“Axel” patent
4,681,893	1986	67	100	62	34	43	-9	Lipitor patent
6,285,999	1998	87	100	—	15	84	-69	PageRank (Google)

Table IV: Summary Statistics: Startup-Quarter Sample

This table presents summary statistics for the quarterly sample (1980-2010) of venture-backed startups. Startups are in the sample from their founding date until the quarter of their final outcome. Note that some startups remain private at the end of the sample period. The sample is further detailed in Section IV.A. The technology variables *Tech Disruptiveness*, *Tech Breadth*, *Private Similarity*, *LI Similarity*, and *Foreign Similarity* are stock variables defined in Section IV.A. Remaining variables are defined in Appendix B. The underlying patent level measures are winsorized at the 1/99% level annually.

	N	Mean	SD	P25	Median	P75
Tech Disruptiveness	347,929	0.66	1.14	0.00	0.00	0.98
Tech Breadth	347,929	0.13	0.18	0.00	0.01	0.27
Private Similarity	347,929	0.06	0.06	0.00	0.06	0.11
LI Similarity	347,929	0.04	0.05	0.00	0.03	0.07
Foreign Similarity	347,929	0.05	0.06	0.00	0.04	0.09
Log(1+Firm Age)	347,929	3.07	1.15	2.40	3.18	3.76
No PatApps[q-1,q-20]	347,929	0.47	0.50	0.00	0.00	1.00
Log(1+PatApps[q-1,q-20])	347,929	0.79	0.97	0.00	0.69	1.39
Log(MTB) (q-2)	347,929	0.15	0.08	0.11	0.15	0.19
MKT Return [q-2,q-1]	347,929	0.01	0.13	-0.08	0.02	0.09
Q4	347,929	0.25	0.43	0.00	0.00	0.00
Originality	347,929	0.16	0.20	0.00	0.00	0.31
Log(1+Cites)	347,929	0.54	0.70	0.00	0.00	1.02
IPO rate (x100)	347,929	0.42	6.43	0.00	0.00	0.00
Sell-Out rate (x100)	347,929	0.73	8.50	0.00	0.00	0.00

Table V: Validity Tests: Public Peer Discussion

This table presents validity tests based on textual analysis of the 10-Ks of public peers of VC-backed startups. The sample is a cross-section of startups measured during the year in which they receive their first round of funding. As discussed in Section IV.B, we link each startup to the 25 public firms whose product descriptions—reported by VenturXpert as of the first VC round—are most similar. In columns (1)-(3), the dependent variable is the average fraction of paragraphs in public peers 10-Ks that mention words with roots “technol” and “change”. In columns (4)-(6), the dependent variable is based on paragraphs with a root of “technol” and “compet”. In columns (7)-(9), the dependent variable is based on paragraphs with a root of “technol” and “compet”, together with either “disrupt”, “change”, or “obsoles”. Thus, these variables measure the intensity with which the public peers of a given startup discuss technology-based market disruption, as discussed in Section IV.B. All variables are defined in Appendix B and the underlying patent level measures are winsorized at the 1/99% level annually. Adjusted R² is reported as a percentage. Standard errors are heteroskedastic robust and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Discussion of:	“Tech” and “Change”			“Tech” and “Competition”			“Tech”, “Competition”, “Disrupt”		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Tech Disruptiveness	0.016*** (4.75)	0.012*** (3.89)	0.007** (2.24)	0.039*** (4.49)	0.031*** (3.64)	0.025*** (2.82)	0.010*** (4.48)	0.007*** (3.42)	0.005** (2.15)
Constant	0.370*** (112.53)			1.401*** (151.99)			0.251*** (110.54)		
Observations	5,417	5,417	5,313	5,417	5,417	5,313	5,417	5,417	5,313
R2 (%)	0.4	6.6	30.2	0.3	4.9	21.0	0.4	8.1	25.3
Year FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Technology FE	No	No	Yes	No	No	Yes	No	No	Yes
Location FE	No	No	Yes	No	No	Yes	No	No	Yes
Firm Age FE	No	No	Yes	No	No	Yes	No	No	Yes
Firm Cohort FE	No	No	Yes	No	No	Yes	No	No	Yes

Table VI: Validity Tests: Post-IPO competition

This table presents validity tests based on a sub-sample of VC-backed startups that go public after 1997 where we are able to merge in both public firm identifiers (GVKEY) and obtain data on the product space of the firm (on startups' IPO year). The dependent variables *HHI* and *TSimm*, from Hoberg and Phillips (2016), are text-based measures of industry concentration and total similarity among a firm's public peers, respectively. *Product Mkt Fluidity* is from Hoberg, Phillips, and Prabhala (2014). All variables are defined in Appendix B and the underlying patent level measures are winsorized at the 1/99% level annually. We include year fixed effects for the year of IPO. Adjusted R² is reported as a percentage. Standard errors are heteroskedastic robust and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	HHI (1)	Log(TSimm) (2)	Product Mkt Fluidity (3)
Tech Disruptiveness	-0.012** (-2.18)	0.073** (2.10)	0.245** (2.16)
Observations	523	523	524
R2 (%)	1.7	1.7	2.5
Year FE	Yes	Yes	Yes

Table VII: The Determinants of Startups' Exits - Baseline

This table presents cross-sectional tests relating startups' ex-ante technological traits to their exit. The outcomes we consider are IPO and sell-out (acquisition). The sample is a quarterly panel of VC-backed startups from 1980-2010 and is described in Section IV.A. Columns (1)-(2) use a competing risk hazard model and columns (3)-(4) use an OLS linear probability model. To facilitate interpretation, coefficients for OLS estimates report the incremental % change in a given outcome. Independent variables are lagged one quarter unless explicitly noted and all controls are standardized for convenience, except for the $Q4$ and $No\ PatApps[q-1,q-20]$ dummy variables. All variables are defined in Appendix B. *LI Similarity* and *Foreign Similarity* are orthogonalized relative to *Private Similarity*. The underlying patent level measures are winsorized at the 1/99% level annually. Technology fixed effects are based on the most common NBER-technology category across a firm's patents. Location fixed effects are based on the state reported in VentureXpert. Adjusted R^2 is reported as a percentage. Standard errors are clustered by startup and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Competing Risk Hazard		OLS	
	IPO (1)	Sell-Out (2)	IPO (3)	Sell-Out (4)
Tech Disruptiveness	0.252*** (13.09)	-0.188*** (-7.55)	0.082*** (4.04)	-0.067*** (-3.42)
Tech Breadth	0.510*** (10.02)	-0.210*** (-5.50)	0.091*** (3.28)	-0.140*** (-4.39)
Private Similarity	0.127* (1.70)	-0.479*** (-8.62)	0.037 (1.06)	-0.433*** (-9.47)
LI Similarity	0.248*** (4.58)	-0.005 (-0.12)	0.056* (1.80)	-0.059 (-1.64)
Foreign Similarity	-0.056 (-1.42)	-0.003 (-0.11)	-0.017 (-0.93)	0.059** (2.39)
No PatApps[q-1,q-20]	1.648*** (10.75)	-2.172*** (-22.17)	0.350*** (5.11)	-1.657*** (-15.09)
Log(1+PatApps[q-1,q-20])	0.327*** (9.65)	-0.056** (-2.11)	0.175*** (7.25)	-0.087*** (-2.97)
Log(MTB) (q-2)	0.133*** (5.15)	0.162*** (8.94)	0.151*** (4.20)	0.044 (0.82)
MKT Return [q-2,q-1]	0.341*** (11.92)	0.004 (0.16)	0.046*** (3.55)	0.036* (1.79)
Q4	-0.059 (-0.77)	0.114** (2.05)	0.137*** (2.86)	0.370*** (5.75)
Originality	-0.125*** (-3.11)	-0.175*** (-6.01)	-0.028 (-1.44)	-0.178*** (-7.34)
Log(1+Cites)	0.171*** (3.95)	0.118*** (3.96)	0.075*** (3.64)	0.152*** (5.38)
Year FE	No	No	Yes	Yes
Technology FE	No	No	Yes	Yes
Location FE	No	No	Yes	Yes
Firm Age FE	No	No	Yes	Yes
Firm Cohort FE	No	No	Yes	Yes
Observations	346,490	345,403	342,146	342,146
R2 (%)	N/A	N/A	0.5	0.6

Table VIII: Determinants of Startups' Exits - Financing

This table presents cross-sectional tests relating startups' ex-ante technological traits to their exit. Each of the models repeats the corresponding model from Table VII, but adds endogenous financing controls. $\log(\text{CumVCFunding})$ is the log of cumulative VC funding the firm receives between its founding and $q - 1$. $\text{No Funding}[q-1, q-20]$ is a control equal to one if the firm has not received funding in the prior 20 quarters. For brevity, we only report the new financing controls and *Tech Disruptiveness*. To facilitate interpretation, coefficients for OLS estimates report the incremental % change in a given outcome, and *Tech Disruptiveness* is standardized. Independent variables are lagged one quarter unless explicitly noted. The underlying patent level measures are winsorized at the 1/99% level annually. Adjusted R^2 is reported as a percentage. Standard errors are clustered by startup and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

	Competing Risk Hazard		OLS	
	IPO (1)	Sell-Out (2)	IPO (3)	Sell-Out (4)
Tech Disruptiveness	0.263*** (13.66)	-0.170*** (-6.82)	0.075*** (3.70)	-0.077*** (-4.00)
$\log(\text{CumVCFunding})$	0.074*** (4.56)	0.138*** (10.18)	0.059*** (8.83)	0.119*** (15.01)
No Funding[q-1, q-20]	-0.927*** (-5.35)	-2.508*** (-9.79)	-0.126** (-2.22)	0.130** (1.99)
Controls	Yes	Yes	Yes	Yes
Year FE	No	No	Yes	Yes
Technology FE	No	No	Yes	Yes
Location FE	No	No	Yes	Yes
Firm Age FE	No	No	Yes	Yes
Firm Cohort FE	No	No	Yes	Yes
Observations	346,490	345,403	342,146	342,146
R2 (%)	N/A	N/A	0.7	0.8

Table IX: The Determinants of Startups' Exits - Dynamic Responses

This table presents dynamic cross-sectional tests relating startups' ex-ante technological traits to their exit over several horizons. In Panel A, column 1 repeats the OLS model examining IPO exits from column 3 in Table VII. Columns 2-6 subsequently replace the one-period ahead IPO exit indicator with longer horizons. We repeat this analysis for sell-outs in Panel B. Panel C examines whether a firm is still private (i.e. no IPO, or sell-out). In all models, the sample, independent variables, and coefficient interpretation are the same as the OLS models in Table VII. Independent variables are standardized for convenience. *LI Similarity* and *Foreign Similarity* are orthogonalized relative to *Private Similarity*. For brevity, the control variables and fixed effects are omitted. Standard errors are clustered by startup and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Exit within next:	Qtr (1)	Year (2)	2 Years (3)	3 Years (4)	4 Years (5)	5 Years (6)
Panel A: Exit by IPO						
Tech Disruptiveness	0.082*** (4.04)	0.299*** (3.88)	0.420*** (3.14)	0.424** (2.49)	0.472** (2.34)	0.466** (2.07)
Tech Breadth	0.091*** (3.28)	0.362*** (3.54)	0.714*** (3.85)	0.973*** (3.90)	1.167*** (3.81)	1.229*** (3.52)
Private Similarity	0.037 (1.06)	0.120 (0.94)	0.215 (0.95)	0.261 (0.87)	0.492 (1.34)	0.623 (1.52)
LI Similarity	0.056* (1.80)	0.182 (1.57)	0.287 (1.40)	0.265 (0.95)	0.313 (0.91)	0.270 (0.69)
Foreign Similarity	-0.017 (-0.93)	-0.045 (-0.64)	-0.029 (-0.22)	0.007 (0.04)	-0.045 (-0.21)	-0.089 (-0.36)
Panel B: Exit by Sell-Out						
Tech Disruptiveness	-0.067*** (-3.42)	-0.191** (-2.46)	-0.145 (-0.92)	-0.055 (-0.25)	0.011 (0.04)	0.089 (0.28)
Tech Breadth	-0.140*** (-4.39)	-0.600*** (-4.85)	-1.086*** (-4.60)	-1.501*** (-4.43)	-1.899*** (-4.43)	-2.056*** (-4.10)
Private Similarity	-0.433*** (-9.47)	-1.468*** (-8.23)	-2.123*** (-6.24)	-2.052*** (-4.28)	-1.485** (-2.53)	-0.628 (-0.93)
LI Similarity	-0.059 (-1.64)	-0.164 (-1.14)	-0.012 (-0.04)	0.303 (0.75)	0.652 (1.29)	1.297** (2.20)
Foreign Similarity	0.059** (2.39)	0.158 (1.60)	0.059 (0.31)	-0.219 (-0.79)	-0.460 (-1.30)	-0.924** (-2.24)
Panel C: Still Private						
Tech Disruptiveness	-0.056* (-1.94)	-0.267** (-2.41)	-0.564*** (-2.77)	-0.745*** (-2.76)	-0.859*** (-2.70)	-0.880** (-2.54)
Tech Breadth	0.044 (0.99)	0.147 (0.88)	0.175 (0.57)	0.223 (0.53)	0.402 (0.79)	0.516 (0.89)
Private Similarity	0.359*** (5.98)	1.301*** (5.82)	2.049*** (5.01)	2.235*** (4.04)	1.778*** (2.69)	1.305* (1.77)
LI Similarity	0.045 (0.90)	0.137 (0.73)	0.092 (0.26)	-0.038 (-0.08)	-0.311 (-0.54)	-0.842 (-1.30)
Foreign Similarity	-0.088*** (-2.65)	-0.279** (-2.20)	-0.409* (-1.73)	-0.390 (-1.19)	-0.392 (-0.98)	-0.161 (-0.36)

Table X: Explaining Aggregate IPO and Sell-Out Rates

This table presents the out-of-sample performance of predictive models of startups' exit using variables standard in the IPO and acquisition literature (the "Base" model) and a model which augments the "Base" model with the new text-based technological variables (the "Text" model). Panel A examines IPO exits and Panel B examines sell-outs. In a given test (column 1), we estimate a Fama and MacBeth (1973) regression quarter-by-quarter where the dependent variable is a dummy indicating an IPO exit (Panel A) or indicating a sell-out exit (Panel B) based on the horizon listed in column 2 (ranging from one quarter to three years) and using the ex-ante measurable independent variables in Table VII. This model is fitted using the early part of our sample, which begins in 1980 and ends before the out-of-sample period (noted in column 3). These fitted Fama-MacBeth coefficients from the early period are then used in the out-of-sample post period (listed in column 3) to predict the average IPO rate and sell-out rate. These predicted exit rates are then compared to the actual rates to compute the fraction of the disappearing IPOs or surging sell-outs anomaly that is explained by either the "Base" model or the "Text" model as noted in columns (5) to (8). Column 9 reports the percentage of each anomaly that cannot be explained by the base model that is explained by the Text model. All probabilities in columns (4)-(8) are reported as percentage points.

Test (1)	Pred- iction Horizon	Post Period	True Exit Rate	Predicted Exit Rate		Error		Text Impr
	(2)	(3)	(4)	Base (5)	Text (6)	Base (7)	Text (8)	(9)
Panel A: IPO Exits								
1	1Q	[1996,2010]	0.33	0.84	0.75	-0.50	-0.42	16%
2	1Q	[1998,2010]	0.27	0.85	0.76	-0.58	-0.49	16%
3	1Q	[2000,2010]	0.22	0.85	0.75	-0.63	-0.53	16%
4	1Y	[1996,2010]	1.27	3.34	2.89	-2.07	-1.63	21%
5	2Y	[1996,2010]	2.45	6.50	5.59	-4.05	-3.15	22%
6	3Y	[1996,2010]	3.56	9.46	8.17	-5.90	-4.60	22%
Panel B: Sell-Out Exits								
7	1Q	[1996,2010]	0.86	0.60	0.75	0.25	0.11	57%
8	1Q	[1998,2010]	0.91	0.63	0.80	0.27	0.10	62%
9	1Q	[2000,2010]	0.95	0.68	0.87	0.27	0.08	71%
10	1Y	[1996,2010]	3.49	2.52	2.96	0.98	0.54	45%
11	2Y	[1996,2010]	7.19	5.30	5.94	1.88	1.25	34%
12	3Y	[1996,2010]	10.96	8.21	8.93	2.75	2.03	26%

Table XI: Explaining Aggregate IPO Rates (Small vs Big IPOs)

This table presents the out-of-sample performance of predictive models of startups' exit using variables standard in the IPO and acquisition literature (the "Base" model) and a model which augments the "Base" model with the new text-based technological variables (the "Text" model). Panel A examines small IPO exits and Panel B big IPO exits. We define an IPO as "small" if its pre-IPO sales are below the median in our sample (\$25 million), and as "large" if its pre-IPO sales exceeds that amount. In a given test (column 1), we estimate a Fama and MacBeth (1973) regression quarter-by-quarter where the dependent variable is a dummy indicating an IPO exit (Panel A) or indicating a sell-out exit (Panel B) based on the horizon listed in column 2 (ranging from one quarter to three years) and using the ex-ante measurable independent variables in Table VII. This model is fitted using the early part of our sample, which begins in 1980 and ends before the out-of-sample period (noted in column 3). These fitted Fama-MacBeth coefficients from the early period are then used in the out-of-sample post period (listed in column 3) to predict the average IPO rate and sell-out rate. These predicted exit rates are then compared to the actual rates to compute the fraction of the disappearing IPOs or surging sell-outs anomaly that is explained by either the "Base" model or the "Text" model as noted in columns (5) to (8). Column 9 reports the percentage of each anomaly that cannot be explained by the base model that is explained by the Text model. All probabilities in columns (4)-(8) are reported as percentage points.

Test (1)	Pred- iction Horizon (2)	Post Period (3)	True Exit Rate (4)	Predicted Exit Rate		Error		Text Impr (9)
				Base (5)	Text (6)	Base (7)	Text (8)	
Panel A: Small IPO Exits								
7	1Q	[1996,2010]	0.15	0.35	0.28	-0.21	-0.13	36%
8	1Q	[1998,2010]	0.12	0.38	0.29	-0.26	-0.18	32%
9	1Q	[2000,2010]	0.08	0.38	0.30	-0.30	-0.21	28%
10	1Y	[1996,2010]	0.55	1.45	1.02	-0.90	-0.47	48%
11	2Y	[1996,2010]	1.06	2.86	2.10	-1.80	-1.04	42%
12	3Y	[1996,2010]	1.53	4.17	3.22	-2.65	-1.69	36%
Panel B: Big IPO Exits								
1	1Q	[1996,2010]	0.13	0.35	0.38	-0.22	-0.25	-13%
2	1Q	[1998,2010]	0.11	0.35	0.38	-0.24	-0.27	-11%
3	1Q	[2000,2010]	0.10	0.34	0.36	-0.24	-0.27	-9%
4	1Y	[1996,2010]	0.51	1.36	1.41	-0.86	-0.91	-6%
5	2Y	[1996,2010]	0.99	2.64	2.54	-1.65	-1.55	6%
6	3Y	[1996,2010]	1.46	3.86	3.59	-2.39	-2.12	11%

Table XII: Explaining Aggregate IPO and Sell-Out Rates (Stable vs Fluid Markets)

This table presents the out-of-sample performance of predictive models of startups’ exit with variables standard in the IPO and acquisition literature (the “Base” model) and a model which augments the “Base” model with the new text-based technological variables (the “Text” model). Panel A examines IPO exits and Panel B examines sell-outs. The procedure is analogous to that described in Table X, except each test is repeated for two sub-samples: *Stable Markets* and *Fluid Markets*, which are defined in Section V.C. We omit the model-implied out-of-sample probabilities to conserve space.

Test (1)	Pred. Horizon (2)	Post Period (3)	Stable Markets				Fluid Markets			
			True Rate (4)	Base Error (5)	Text Error (6)	Text Impr (7)	True Rate (8)	Base Error (9)	Text Error (10)	Text Impr (11)
Panel A: IPO Exits										
1	1Q	[1996,2010]	0.30	-0.41	-0.31	26%	0.37	-0.62	-0.67	-7%
2	1Q	[1998,2010]	0.26	-0.45	-0.33	27%	0.29	-0.75	-0.79	-5%
3	1Q	[2000,2010]	0.24	-0.44	-0.32	29%	0.20	-0.86	-0.89	-4%
4	1Y	[1996,2010]	1.14	-1.69	-1.29	24%	1.42	-2.57	-2.29	11%
5	2Y	[1996,2010]	2.20	-3.32	-2.60	22%	2.74	-5.02	-4.27	15%
6	3Y	[1996,2010]	3.23	-4.75	-3.68	22%	3.96	-7.38	-6.51	12%
Panel B: Sell-Out Exits										
7	1Q	[1996,2010]	0.85	0.16	0.04	73%	0.86	0.37	0.21	44%
8	1Q	[1998,2010]	0.89	0.17	0.04	75%	0.93	0.40	0.19	52%
9	1Q	[2000,2010]	0.91	0.16	0.01	93%	0.98	0.40	0.17	57%
10	1Y	[1996,2010]	3.43	0.59	0.34	43%	3.55	1.46	0.79	46%
11	2Y	[1996,2010]	6.97	1.05	0.89	15%	7.39	2.90	1.56	46%
12	3Y	[1996,2010]	10.52	1.44	1.43	1%	11.38	4.32	2.32	46%

A Defining the Entity Type of Patents' Assignees

To classify if a patent is granted to (A) a private, domestic U.S. firm, (B) an international firm, or (C) a U.S. public firm, we use the following procedure. First, we find all patents assigned to public firms. We obtain the GVKEY for assignees from the NBER patent dataset, and augment this with Kogan, Papanikolaou, Seru, and Stoffman (2016). We use all assignee links for the entire 1900-2013 period. Also note that Kogan, Papanikolaou, Seru, and Stoffman (2016) contains PERMNO identifiers, which we convert to GVKEY using a link table from WRDS. When the headquarters country from CRSP-Compustat is available, we mark these firms as either international firms or U.S. public firms. Next, we output the top 3,000 remaining assignees and manually classify the entity type. After these steps, 3,126,605 patents are classified as either U.S. public firms or foreign firms.

Second, we use information from the NBER classification of assignees and manual categorization to remove patents assigned to governmental entities, research think tanks, or universities.

Third, we directly identify patents assigned to foreign firms when the last word in the assignee name is an unambiguous foreign legal identifier, such as “GMBH”, “PLC”, and “Aktiengesellschaft”. We also identify patents granted to foreign firms when the assignee is a firm (e.g. “CORP”) and USPTO data indicates that the assignee is not domestic. This step identifies 898,797 patents granted to foreign firms.

Fourth, we classify entities as U.S. private domestic firms when the assignee is a firm (e.g. “CORP”) and USPTO data indicates the assignee is domestic. Previous steps affirmatively prevent us from calling a corporation a private domestic firm if the corporate is a public firm, a think tank, or international corporation.

In total, we classify the entity type of 78% of all patents granted from 1900-2013. Moreover, during our main analysis period (1980-2010), we are able to classify the assignee entity type for 92% of patent applications. Of the 4,161,306 applied for in the main analysis period, 12% are private U.S. firms, 27% are public U.S. firms, 41% are foreign firms, 8% are unclassified, and 11% are “other”.

B Variable Definitions

Patent level variables

Tech Disruptiveness	See Equation 2 and Section II.B.
Tech Breadth	See Equation 3 and Section II.C.
LI Similarity	See Equation 4 and Section II.C.
Private Similarity	See Equation 5 and Section II.C.
Foreign Similarity	See Equation 6 and Section II.C.
KPSS Value	From Kogan, Papanikolaou, Seru, and Stoffman (2016).
# of Cites	Number of citations received in the first five years after publication by the USPTO. Citations up to December 31, 2013.
Originality	The originality of a focal patent is defined as 1 minus the HHI of the technology fields of the patents cited by the focal patent (Trajtenberg, Henderson, and Jaffe (1997)). We use the adjustment given in Hall, Jaffe, and Trajtenberg (2001) to reduce bias for patents that contain few backward citations. We convert U.S. Patent Classifications for to the NBER technology codes so that <i>Tech Breadth</i> and <i>Originality</i> are based on the same granularity of technology classifications.

Startup-quarter variables

Tech Disruptiveness	The depreciated sum of patent-level <i>Tech Disruptiveness</i> for patents the firm applied for over the prior 20 quarters. Quarterly depreciation is 5%. We normalize the depreciated sum by the number of patents the startup applied for. See Section IV.A for more.
Tech Breadth	Converted to startup-quarter like <i>Tech Disruptiveness</i> .
Private Similarity	Converted to startup-quarter like <i>Tech Disruptiveness</i> .
LI Similarity	Converted to startup-quarter like <i>Tech Disruptiveness</i> .
Foreign Similarity	Converted to startup-quarter like <i>Tech Disruptiveness</i> .
Log(1+Cites)	Log of the stock of citations. Citations for a startup-quarter is the sum of the # of Cites (patent-level variable defined above) for patents the startup applies for in the quarter. Note that this is forward-looking. The stock is computed using a quarterly depreciation of 5%.
Originality	Converted to startup-quarter like <i>Tech Disruptiveness</i> .
No PatApps[q-1,q-20]	Dummy variable equal to one if the startup has not applied for a patent (which was eventually granted) during the last 20 quarters.
Log(1+PatApps[q-1,q-20])	The log of 1 plus the number of patent applications (which were eventually granted) made by the startup in the last 20 quarters.
Log(1+startup Age)	startup age is defined as the number of years since its founding date according to VentureXpert.
IPO	One if the startup goes public in the quarter, zero before.
Sell-out	One if the startup is acquired in the quarter, zero before.

Quarterly variables

Log(MTB) (q-2)	Aggregate market-to-book is computed quarterly using all firms in the CRSP-Compustat database. We sum each subcomponent of MTB across all firms, then compute $MTB = (at - ceq + mve - txdb)/at$ as defined in Kaplan and Zingales (1997).
MKT Return [q-2,q-1]	We compute quarterly market returns from Ken French's daily factor file using geometric compounding.
Q4	Equal to one if $t - 1$ is the fourth quarter (and t is the first quarter) of the year.

C Matching patents to VentureXpert

We download all data on firms receiving venture capital funding starting in 1970 and ending in 2013 from VentureXpert using SDC Platinum. In addition to the dates of venture financing, we also download data indicating each portfolio company’s founding date, its final resolution (as IPO, acquisition, or unresolved) and date of resolution, the company’s name and the number of financing rounds it received.

Merging VentureXpert with the patent level data requires a link between firms in the patent database (the initial assignees) and firms in the VentureXpert database. We develop a fuzzy matching algorithm—outlined below—to match firms in both databases using their names. The algorithm matches 532,660 patents granted between 1966 and 2013 to 19,324 VC-backed firms.³⁶ 96.6% of the patent matches and 90.7% of the VC-backed firms are matched via exact matches on the raw firm name in both datasets or on a cleaned version of the firm name.

The matching procedure begins by standardizing assignee names in the patent dataset and in VentureXpert, using a name standardization routine from Nada Wasi.³⁷ This standardizes common company suffixes and prefixes and produces stem names. We also modify this program to exclude all information after a company suffix, as this is typically address information erroneously stored in the name field by the USPTO. After standardizing the names, we use the following steps to match firms in the two datasets:

1. We compare all *original string* names in each dataset, adjusted only to replace all uppercase characters. If a single VC-backed firm is an exact match where the patent application is after the firm’s founding date, we accept the match. This step matches 59,026 patents to VC-backed firms, or 11% of the accepted matches.
2. For the remaining patents, we compare all *cleaned string* names in each dataset. If a single VC-backed firm is an exact match where the patent application is after the firm’s founding date, we accept the match. This step matches 455,456 patents to VC-backed firms, or 86% of accepted matches.
3. For the remaining patents, we select matches using a fuzzy matching technique, with rules based on random sampling and validation checks in a hold out sample. This step matches 18,178 patents to VC-backed firms, or 3% of accepted matches. The steps are as follows:
 - (a) We compute string comparison scores by comparing all *cleaned string* names in each dataset using several different string comparison functions. We do this three separate times, requiring that (1) the first three characters are exact matches, (2) the first five characters are exact matches, and (3) the first seven characters are exact matches. We then output a random sample of patents for an RA to examine.
 - (b) The highest performing rule was a bi-gram match function with the restriction that the first seven characters were equivalent in both the patent assignee and company name. For each remaining patent, we keep as candidate matches any pair with equivalent name stems and the highest bi-gram match above 75%.
 - (c) A random subset of suggested matches, in addition all borderline suggested matches, were reviewed by hand.

³⁶Firms can receive patents before VC funding.

³⁷ <http://www-personal.umich.edu/~nwasi/programs.html>

As a result of this matching process, our patent level database contains U.S. private firms that both (A) have patents and (B) have received VC funding. Aside from imperfections in the matching process, which could be material, this database is the universe of such firms.³⁸ For each such firm, we have data indicating its final outcome and text-based data indicating the details of the firm’s patents, and when they were applied for and granted. This data allows us to examine both (A) potential drivers of VC funding among firms that have patents but have not yet received funding, and (B) final resolutions of private status as IPOs or acquisitions. Cross-sectional and time series examination of both form the basis of our hypothesis testing.

³⁸Lerner and Seru (2017) note that using string matching to identify firms suffers from a limitation when private firms have patents issued to legal entities with different names, such as subsidiaries or shell companies meant to obfuscate the owner. This limitation can not be avoided, but is reduced for our sample of interest. VC-backed private firms are typically small and thus are unlikely to have distinctly named subsidiaries for research). Moreover, obfuscation is most often used by *non-practicing entities*, often called patent trolls, which are unlikely to be a material number of firms in our 19,324 firm sample.

D Additional Tables

1. Table A1 presents percentiles of various statistics for a sample of important patents.
2. Table A2 presents information on the timing of key life events for startups in the main analysis sample.
3. Table A3 presents subsample tests of the main OLS models on the determinants of startups' exit from Table VII. The subsamples are based on the date of the observation.
4. Table A4 presents tests of regression of startups' financing on their technological characteristics.

Table A1: Percentiles of various statistics for a sample of important patents

The patents below are the 1930-2010 subset of key important patents listed in Kelly, Papanikolaou, Seru, and Taddy (2018) (henceforth, KPST) over which the textual measures in this paper are defined. For each patent, we report the percentile (across all patents) of each statistic. The percentiles for the KPST measures are taken directly from their Table A.6. *Qual* is their main innovation measure, which compares the text in a patent to text in all patents over the previous and next five years. *Cite* is citations for the first five years after a patent is granted, as reported in KPST. *Tech Disruptiveness* (“Tech Dsrpt”), *Tech Breadth* (“Tech Brdth”), Private Similarity, LI Similarity, and *Foreign Similarity* are defined in Section II and labeled as Text. *KPSS* is an estimate of the economic value of a patent based on stock market returns for the firm granted the patent from Kogan, Papanikolaou, Seru, and Stoffman (2016).

Patent	Year	KPST		Text					KPSS
		Qual	Cite	Tech Dsrpt	Tech Brdth	Priv Simm	LI Simm	Foreign Simm	KPSS
Panel A: Average of percentiles in Panel B									
		0.73	0.64	0.77	0.52	0.54	0.64	0.57	0.64
Panel B: Percentiles of various measures for breakthrough patents									
1929453	1933	0.79	0.93	0.96	0.23	0.85	0.59	0.86	
1941066	1933	0.38	0.00	0.68	0.67	0.26	0.95	0.78	
1948384	1934	0.27	0.00	0.85	0.72	0.28	0.67	0.56	
1949446	1934	0.06	0.31	0.68	0.28	0.86	0.43	0.36	
1980972	1934	0.76	0.00	0.99	0.15	0.88	0.71	1.00	
2021907	1935	0.38	0.00	0.66	0.96	0.35	0.74	0.53	
2059884	1936	0.20	0.50	0.95	0.42	0.49	0.61	0.84	0.26
2071250	1937	0.63	0.50	0.98	0.27	0.97	0.85	1.00	0.71
2087683	1937	0.68	0.00	0.89	0.59	0.43	0.76	0.85	
2153729	1939	0.81	0.00	0.99	0.12	0.74	0.63	1.00	
2188396	1940	0.97	0.00	0.83	0.15	0.86	0.56	0.76	0.15
2206634	1940	0.82	0.31	0.94	0.50	0.73	0.60	0.95	
2230654	1941	0.43	0.89	0.94	0.24	0.87	0.70	0.49	
2258841	1941	0.02	0.77	0.60	0.84	0.78	0.47	0.82	
2292387	1942	0.45	0.31	0.70	0.90	0.66	0.79	0.61	
2297691	1942	0.06	0.71	0.33	0.81	0.67	0.64	0.74	
2329074	1943	0.05	0.99	0.93	0.19	0.99	0.94	0.98	
2390636	1945	0.34	0.97	0.61	0.50	0.14	0.17	0.28	
2404334	1946	0.13	0.94	0.74	0.75	0.86	0.76	0.96	
2436265	1948	0.65	0.81	0.70	0.39	0.83	0.76	0.68	
2451804	1948	0.65	0.50	0.93	0.19	0.95	0.77	0.50	
2495429	1950	0.22	0.87	0.92	0.84	0.54	0.73	0.36	
2524035	1950	0.60	1.00	0.83	0.44	0.74	0.92	0.91	0.71
2543181	1951	0.44	0.99	0.85	0.32	0.94	0.83	0.73	
2569347	1951	0.45	1.00	0.84	0.47	0.84	0.98	0.96	0.55
2642679	1953	0.36	0.50	0.39	0.65	0.98	0.66	0.92	
2668661	1954	0.95	0.31	1.00	0.80	0.01	0.01	0.01	0.59
2682050	1954	0.63	0.00	0.91	0.99	0.40	0.67	0.49	
2682235	1954	0.48	0.77	0.93	0.97	0.16	0.27	0.19	
2691028	1954	0.88	0.00	0.91	0.04	0.98	0.92	0.95	

Continued on next page

Patent	Year	KPST		Text					KPSS
		Qual	Cite	Tech Dsrpt	Tech Brdth	Priv Simm	LI Simm	Foreign Simm	KPSS
2699054	1955	0.92	0.98	0.37	0.22	0.98	0.97	0.98	
2708656	1955	0.98	1.00	1.00	0.89	0.01	0.01	0.01	
2708722	1955	0.70	0.97	0.94	0.42	0.13	0.74	0.44	
2717437	1955	0.44	0.62	0.73	0.45	0.06	0.10	0.14	
2724711	1955	0.74	0.31	0.99	0.06	0.58	0.74	0.79	
2752339	1956	0.84	0.62	0.85	0.04	0.96	0.94	0.97	
2756226	1956	0.70	0.71	0.90	0.18	0.48	0.64	0.88	
2797183	1957	0.86	0.31	1.00	0.24	0.88	0.73	0.87	
2816721	1957	0.71	0.77	0.70	0.45	0.54	0.53	0.74	
2817025	1957	0.70	0.96	0.46	0.45	0.94	1.00	0.79	
2835548	1958	0.81	0.92	0.82	0.98	0.76	0.45	0.59	
2866012	1958	0.77	0.97	0.37	0.52	0.98	0.97	0.89	
2879439	1959	0.72	0.96	0.74	0.76	0.59	0.83	0.58	
2929922	1960	0.82	1.00	0.90	0.57	0.54	0.69	0.61	0.87
2937186	1960	0.83	0.31	0.98	0.11	0.99	0.97	1.00	
2947611	1960	0.71	0.00	0.55	0.54	0.96	0.96	0.95	0.56
2956114	1960	0.68	0.62	0.56	0.53	0.99	0.98	0.91	0.64
2981877	1961	0.96	1.00	0.84	0.41	0.71	0.83	0.52	
3057356	1962	0.88	0.93	0.88	0.51	0.67	0.95	0.63	
3093346	1963	0.89	0.87	0.92	0.79	0.53	0.57	0.56	
3097366	1963	0.48	0.62	0.42	0.98	0.90	0.67	0.79	
3118022	1964	0.70	0.50	0.51	0.68	0.79	0.81	0.63	0.86
3156523	1964	0.82	0.00	0.35	0.24	0.90	0.84	0.98	
3174267	1965	0.62	0.71	0.59	0.92	0.58	0.18	0.21	0.44
3220816	1965	0.83	0.31	0.14	0.41	0.07	0.24	0.13	
3287323	1966	0.70	0.00	0.40	0.12	0.65	0.98	0.99	0.56
3478216	1969	0.71	0.31	0.52	0.35	0.47	0.72	0.54	
3574791	1971	0.66	0.84	0.38	0.18	0.84	1.00	0.99	0.91
3663762	1972	0.59	0.87	0.55	0.65	0.21	0.52	0.16	0.82
3789832	1974	0.42	0.71	0.80	0.87	0.64	0.89	0.71	
3858232	1974	0.39	0.95	0.56	0.75	0.67	0.95	0.81	0.98
3906166	1975	0.38	0.81	0.88	0.89	0.58	0.79	0.52	0.46
4136359	1979	0.79	0.62	0.62	0.70	0.73	0.99	0.88	
4229761	1980	0.59	0.00	0.64	0.82	0.05	0.33	0.21	
4237224	1980	1.00	1.00	0.96	0.65	0.10	0.33	0.26	
4371752	1983	0.75	0.93	0.95	0.31	0.79	0.99	0.87	
4399216	1983	0.99	0.97	0.99	0.71	0.08	0.33	0.16	
4464652	1984	0.70	0.98	0.45	0.89	0.96	0.88	0.85	0.87
4468464	1984	1.00	0.50	1.00	0.61	0.08	0.32	0.25	
4590598	1986	0.70	0.31	0.69	0.83	0.85	0.95	0.87	0.91
4634665	1987	0.99	0.62	0.95	0.59	0.09	0.37	0.14	
4683195	1987	0.97	1.00	0.91	0.79	0.39	0.70	0.61	0.46
4736866	1988	1.00	0.81	1.00	0.52	0.04	0.23	0.06	
4744360	1988	0.94	0.81	0.77	0.74	0.60	0.71	0.63	
4799258	1989	0.96	0.95	0.81	0.34	0.24	0.62	0.21	
4816397	1989	0.97	0.81	0.99	0.66	0.26	0.55	0.34	
4816567	1989	0.99	0.77	0.93	0.71	0.36	0.67	0.46	0.77
4838644	1989	0.94	0.81	0.79	0.99	0.83	0.95	0.89	

Continued on next page

Patent	Year	KPST		Text					KPSS
		Qual	Cite	Tech Dsrpt	Tech Brdth	Priv Simm	LI Simm	Foreign Simm	KPSS
4889818	1989	0.98	0.99	0.92	0.78	0.53	0.82	0.69	0.45
4965188	1990	0.97	0.99	0.93	0.82	0.44	0.79	0.65	0.40
5061620	1991	0.99	1.00	0.91	0.40	0.10	0.25	0.12	
5071161	1991	0.81	0.96	0.44	0.98	0.69	0.57	0.44	
5108388	1992	0.97	0.00	0.90	0.92	0.46	0.65	0.43	0.30
5149636	1992	0.99	0.31	0.62	0.60	0.05	0.31	0.10	
5179017	1993	1.00	0.96	0.92	0.37	0.14	0.41	0.26	
5184830	1993	0.98	0.98	0.65	0.39	0.82	0.90	0.89	
5194299	1993	0.87	0.00	0.51	0.35	0.66	0.58	0.62	0.79
5225539	1993	1.00	0.99	0.97	0.62	0.09	0.34	0.19	
5272628	1993	0.99	0.92	0.99	0.27	0.13	0.51	0.15	0.99
5747282	1998	0.98	0.71	1.00	0.19	0.32	0.30	0.28	0.43
5770429	1998	0.91	0.84	0.99	0.17	0.29	0.27	0.27	
5837492	1998	0.95	0.00	1.00	0.31	0.32	0.37	0.37	0.32
5939598	1999	1.00	0.31	0.99	0.36	0.18	0.37	0.27	0.60
5960411	1999	1.00	1.00	0.84	0.10	0.12	0.47	0.08	1.00
6230409	2001	0.86	0.62	0.40	0.98	0.75	0.20	0.34	
6285999	2001	0.98	1.00	0.87	0.15	0.08	0.42	0.10	
6331415	2001	0.98	0.00	0.61	0.68	0.41	0.71	0.53	0.94
6455275	2002	0.97	0.31	0.99	0.18	0.38	0.39	0.42	
6574628	2003	0.99	0.95	0.84	0.11	0.22	0.62	0.19	
6955484	2005	0.59	0.84	0.03	0.65	0.55	0.18	0.29	
6985922	2006	0.81	0.95	0.75	0.11	0.16	0.55	0.09	

Table A2: Years between keys events for ventured-backed Startups

This table presents information of key events for startups in the main analysis sample described in Table IV and Section IV.A. A startup's first patent is based on the earliest application date for (eventually) granted patents. Information on VC funding, timing, and exits are from VentureXpert, and patenting information is from Google Patents.

Panel A: Events after the startup's founding

Event	N (startups)	Years between the startup's founding and event				
		Mean	SD	P25	Median	P75
First patent	9,167	4.42	10.76	0.75	2.25	5.75
VC funding	9,167	5.29	10.63	0.50	1.75	5.50
IPO	1,677	9.41	9.89	4.50	7.00	11.25
Acquisition	3,377	11.23	10.50	6.00	8.50	12.75

Panel B: Events after the startup's first patent

Event	N (startup)	Years between the startup's first patent and event				
		Mean	SD	P25	Median	P75
VC funding	9,167	0.87	7.78	-2.00	-0.25	2.50
IPO	1,677	3.10	8.45	-0.50	3.00	6.75
Acquisition	3,377	7.46	7.00	3.75	6.25	10.00

Table A3: Subsample analysis of startups' exit: Time

This table repeats the OLS cross-sectional tests in columns (4)-(5) from Table VII on two subsamples. The tests relate a startup's ex-ante technological traits and its ultimate outcome. We split the sample based on the observation date. Even numbered columns include observations before January 1, 1996 and odd numbered columns include observations on or after January 1, 1996. In all models, the definition of independent variables and interpretation of coefficients are the same as the OLS models in Table VII. Independent variables are lagged one quarter and standardized for convenience. Note that we standardize variables *within* the subsample of the test. *LI Similarity* and *Foreign Similarity* are orthogonalized relative to *Private Similarity*. For brevity, the control variables are omitted. All variables are winsorized at the 1/99% level annually. Technology fixed effects are based on the most common NBER-technology category across a startup's patents. Location fixed effects are based on the state reported in VentureXpert. Adjusted R² is reported as a percentage. Standard errors are clustered by startup and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Exit Type:	IPO		Acquisition	
	Before (1)	After (2)	Before (3)	After (4)
Observation before/after 1995:				
Tech Disruptiveness	0.127*** (2.73)	0.089*** (4.41)	-0.186*** (-6.32)	-0.025 (-0.98)
Tech Breadth	0.109 (1.31)	0.012 (0.47)	-0.286*** (-3.91)	-0.179*** (-4.68)
Private Similarity	-0.017 (-0.16)	-0.017 (-0.56)	-0.164* (-1.76)	-0.466*** (-9.18)
LI Similarity	-0.001 (-0.01)	-0.019 (-0.64)	0.017 (0.24)	-0.062 (-1.34)
Foreign Similarity	0.062 (0.87)	-0.008 (-0.45)	-0.036 (-0.66)	0.077*** (2.69)
Year FE	Yes	Yes	Yes	Yes
Technology FE	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes
Firm Age FE	Yes	Yes	Yes	Yes
Firm Cohort FE	Yes	Yes	Yes	Yes
Observations	112,643	229,499	112,643	229,499
Firms	3,951	7,483	3,951	7,483
R2 (%)	0.3	0.5	0.4	0.6

Table A4: The Determinants of Startups' VC Funding

This table presents OLS cross-sectional tests relating a firm's ex-ante technological traits and its VC financing. The outcomes we consider are the log of cumulative VC funding (*Cum.Funds*) the firm receives between its founding and quarter q , and *New Round*, a binary variable that equals one if a firm receives a new round of VC financing in quarter q . In all models, the sample, independent variables, and coefficient interpretation are the same as the OLS models in Table VII. Independent variables are standardized for convenience and lagged one quarter. *LI Similarity* and *Foreign Similarity* are orthogonalized relative to *Private Similarity*. For brevity, control variables are omitted. Adjusted R^2 is reported as a percentage. Standard errors are clustered by firm and are reported in parentheses. The symbols ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively.

Ind. Variable:	Cum.Funds	New round	Cum.Funds		New Round	
			pre-95	post-95	pre-95	post-95
Sample:	Whole	Whole				
	(1)	(2)	(3)	(4)	(5)	(6)
Tech Disruptiveness	0.162*** (5.02)	0.670*** (6.82)	0.919*** (5.33)	0.466*** (4.08)	0.139*** (2.59)	0.162*** (4.44)
Tech Breadth	-0.323*** (-5.91)	0.017 (0.11)	0.309 (0.87)	-0.299 (-1.59)	-0.164 (-1.50)	-0.441*** (-6.65)
Private Similarity	0.060 (0.92)	1.247*** (6.21)	0.866** (2.00)	1.265*** (5.39)	0.037 (0.28)	0.018 (0.24)
LI Similarity	0.207*** (3.52)	0.477*** (2.76)	0.415 (1.15)	0.267 (1.27)	0.128 (1.12)	0.124* (1.70)
Foreign Similarity	-0.054 (-1.30)	-0.025 (-0.21)	0.380 (1.34)	-0.198 (-1.42)	0.128 (1.38)	-0.088* (-1.85)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Technology FE	Yes	Yes	Yes	Yes	Yes	Yes
Location FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Age FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Cohort FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	342,146	347,918	114,364	233,550	112,643	229,499
Firms	9,145	9,167	3,972	7,543	3,951	7,483
R2 (%)	32.4	2.2	2.6	1.6	24.7	27.6