Patent Trajectories and M&A Choices

Ali Jafarzadeh
Leeds School of Business
University of Colorado Boulder

February 1, 2026

Abstract

This paper studies how the evolution of firms' patented technologies shapes merger selection. I construct a dataset of firm-year embeddings based on firms' patent portfolios. Each firm is mapped to a time-varying centroid in a high-dimensional technology space using text from patent abstracts. I introduce a novel measure of convergence speed, which captures how quickly a potential target's innovation trajectory moves toward an incumbent. Faster convergence reflects greater future competitive pressure and stronger bargaining power. Exploiting quasi-random variation in patent outcomes generated by examiner leniency, I show causally that faster technological convergence reduces the likelihood of acquisition, conditional on distance. I also find that acquisitions are most likely at intermediate technological distances: firms that are too distant offer limited synergies, while firms that are already very close provide little incremental value and face greater regulatory scrutiny. To fix ideas and provide intuition, a simple framework explains these patterns through trade-offs between synergy, bargaining power, and regulatory costs.

1 Introduction

Mergers and acquisitions (M&A) are a primary mechanism through which corporate assets, technologies, and control rights are reallocated across firms. A large literature documents the scale of M&A activity and its mixed consequences for shareholders and operating performance (Andrade et al., 2001; Healy et al., 1992; Loughran and Vijh, 1997). More recent work shows that many deals are explicitly about repositioning firm boundaries around innovative activities and technologies rather than only around physical assets (Bena and Li, 2014; Seru, 2014). In this view, a merger is a decision about which innovation programs to internalize and which to leave outside the firm.

I take the perspective that a firm can be represented as a portfolio of inventive human capital and ideas that evolves in a technology space. This representation is consistent with models where firms grow by adding and pruning product lines and technologies and where R&D investment and patenting jointly shape firm dynamics (Klette and Kortum, 2004; Kogan et al., 2017a,b). Related work represents firms spatially using text-based measures of products or technologies and shows that firms' relative positions in this space capture competitive structure and firm boundaries (Hoberg and Phillips, 2016). Patents provide a noisy but informative proxy for this evolving portfolio. They record the timing and content of research efforts and have been shown to predict firm value, growth, and follow-on innovation in settings that exploit quasi-experimental variation in patent grants and invalidations (Farré-Mensa et al., 2019; Galasso and Schankerman, 2018; Sampat and Williams, 2019; Gaulé, 2014). In addition, assignment of applications to examiners and variation in examination practices create heterogeneity in which ideas are turned into granted patents (Frakes and Wasserman, 2014, 2020; Barber et al., 2022; Dyer et al., 2024; Tur et al., 2025). Together, these findings justify treating granted patent portfolios as an economically meaningful but imperfect measure of firms' technological positions and their evolution over time.

Existing work on innovation-related M&A typically focuses on static measures of technological or product-market relatedness at the time of the deal. Research in strategy and innovation documents that post-acquisition innovation outcomes are maximized when acquirers and targets are neither too similar nor too unrelated in technology space (Ahuja and Katila, 2001; Cloodt et al., 2006). In corporate finance, text-based measures of product similarity show that deals cluster among firms with related products and that these synergies shape merger gains and competitive effects, giving rise to an interior (inverted-U) region of merger activity (Hoberg and Phillips, 2010) in which I'll refer to as the *Golden Zone*. Related work uses similar spatial representations to study product-market threats and the evolution of competitive pressure over time (Hoberg et al., 2014), as well as the organization of production and firm scope in multi-product firms (Hoberg and Phillips, 2025). These studies establish that who buys whom depends on relatedness, but they largely treat relatedness as a static characteristic.

This paper instead focuses on the dynamic motion of firms' technologies. I treat each firm's patent portfolio as a trajectory in technology space and characterize potential targets along two primitives: their instantaneous technological distance to the incumbent and the speed at which their innovation path moves toward or away from the incumbent's core. Economically, distance summarizes how similar two firms' current capabilities are. The baseline distance is speaking to the relatedness idea. The novel measure I am proposing is the speed of convergence. The speed summarizes how quickly the candidate's portfolio is catching up to or diverging from the incumbent, capturing dynamic competitive pressure and future bargaining strength. A small, fast-moving firm that is rapidly approaching the incumbent in its own technological domain is therefore both a valuable acquisition candidate and a potential future threat if left independent. Also, if the candidate gets too close, it exists the golden zone which means low synergy in case of merger and also it will increase the risk of anti trust authority intervention and the costs associated to that.

To fix ideas and provide some intuition about the potential mechanisms, I develop a simple framework in which merger surplus, bargaining, and regulatory scrutiny depend jointly on distance and convergence speed. Following the literature on the *Golden Zone* mergers, synergy is positive only on an interior band of distances where portfolios overlap enough to generate knowledge transfer but not so much that duplication and market overlap dominate. Within this absorbable band, higher convergence speed raises the target's outside option and bargaining weight and increases expected antitrust and litigation costs if regulators worry about current and future concentration. The model implies a distance-dependent cutoff in convergence speed. Deals occur when firms are in the absorbable band and their convergence speed is below this cutoff. Candidates that converge too rapidly fall into a collision zone: they become too strong and too risky to acquire and are more likely to be left as independent competitors. Keeping in mind that they would also will result in reduced synergy mergers.

Empirically, I map firms into technology space using transformer-based embeddings of patent abstract text. I build on PatentBERT, which shows that fine-tuning a pre-trained BERT model on patent claims substantially improves patent classification relative to convolutional and word-embedding baselines (Lee and Hsiang, 2020), and on subsequent work that develops patent-specific sentence embeddings and domain-adapted transformer models for patent similarity and classification (Bekamiri et al., 2024; Althammer et al., 2021). These advances indicate that models trained directly on patent language capture technological structure more precisely than generic text models or simple bag-of-words measures. I use such embeddings to construct annual firm-level centroids in technology space and to compute both levels and changes of patent-based distances between incumbents and all potential targets. I understand the concern about the *Look Ahead Bias* about the use of LLMs, but I do not believe the mapping of text into embedding space for categorization purposes would have minimal bias, if any.

The empirical analysis delivers three main results. First, realized acquirer–target pairs sit inside an intermediate band of technological distances and exhibit steady pre-deal convergence, consistent with the absorbable region of the model and the golden-zone logic documented in prior work. Second, conditional on distance in this band, acquisition probabilities are sharply decreasing in convergence speed: incumbents are significantly less likely to acquire candidates whose portfolios move toward them too quickly. Third, when I instrument for the evolution of granted patent portfolios using examiner leniency and related patent office frictions, the negative effect of convergence speed on acquisition is amplified, suggesting that dynamic innovation motion is a causal determinant of who is acquired and who remains a competitor. In this way, the paper connects the literatures on innovation and M&A, spatial representations of firms, patent-based measurement of technologies, and examiner-based identification, and shows that the dynamics of technological motion are central for understanding merger selection.

2 Simplified Conceptual Framework

This section develops a minimal framework linking the motion of firms' patent portfolios to merger decisions. Meaning how distance, direction, and speed of innovation affect acquisition incentives. This section's goal is to not to solve a model, and I will only provide a simplified framework. The goal is to fix ideas and provide some intuition regarding the dynamics of the firms. The framework captures the geometry underlying the empirical analysis. The details of the framework are as follows:

Patent portfolios are represented as points(centroids) moving in technology space, and their motion generates an "*absorbable band*" of desirable targets and a "*collision zone*" of fast-approaching rivals.

As we will see in the next sections, the primitives of the framework are exactly the quantities observed in the data: the technological distance between an incumbent and a candidate, the direction in which the candidate's innovation path moves, and the speed of that motion.

2.1 Technology Space and Patent Portfolios

Time is continuous, $t \geq 0$. There are two firms: an incumbent I and a candidate J. Technologies live in a K–dimensional Euclidean space $\mathbb{R}^K$. Firm $f \in \{I, J\}$ holds a portfolio of patented technologies summarized by a centroid

$$c_f(t) \in \mathbb{R}^K,$$

interpreted as the "center of mass" of its innovation activity.

The relative position of the candidate with respect to the incumbent is

$$x(t) \equiv c_J(t) - c_I(t), \qquad d(t) \equiv \|x(t)\|.$$

The scalar $d(t)$ is the technological distance between the two portfolios.

For intuition, one may imagine $K = 2$ with I at the origin and J at position $x(t)$, where $d(t)$ is its radius around the incumbent's technological core.

2.2 Innovation, Direction, and Speed

Each firm chooses R&D effort over time. Let $R_f(t)$ denote the R&D effort of firm f at time t. Innovation arrives according to a Poisson process with intensity

$$\lambda_f(t) = \phi(R_f(t)), \qquad \phi'(\cdot) > 0,$$

so higher R&D increases the expected arrival rate of new ideas.

Conditional on an idea arrival at time t, firm f's centroid moves by a random vector Δv_f with mean

$$\mathbb{E}[\Delta v_f \mid \text{idea at } t] = \eta_f(t)\, u_f(t),$$

where $u_f(t)$ is a unit vector describing the *direction* of technological search and $\eta_f(t) \geq 0$ is the expected *step size* of an idea in that direction.

Aggregating many small moves over a short interval Δt, the expected motion of firm f is approximately

$$c_f(t+\Delta t) - c_f(t) \approx \lambda_f(t)\, \eta_f(t)\, u_f(t)\, \Delta t.$$

The candidate's expected motion relative to the incumbent is then

$$x(t+\Delta t) - x(t) \approx v(t)\, \Delta t, \qquad v(t) \equiv \lambda_J(t)\eta_J(t)u_J(t) - \lambda_I(t)\eta_I(t)u_I(t).$$

The relevant scalar object for merger decisions is the *radial component* of this relative velocity: the rate at which technological distance changes. Define the instantaneous convergence speed as

$$s(t) \equiv -\frac{d}{dt}d(t) = -\frac{x(t)}{\|x(t)\|} \cdot v(t).$$

A positive value $s(t) > 0$ means that the candidate is *converging* toward the incumbent, while $s(t) < 0$ means it is *diverging*. The magnitude $|s(t)|$ is the speed of this motion.

2.3 Merger Surplus, Bargaining, and Regulation

It is obvious that M&A is not a meaningful option for every pair of firms at every point in time. Most firms remain either too distant to generate transferable knowledge or too close to produce incremental value. An acquisition opportunity becomes economically relevant only when the incumbent and the candidate move into a region of technological proximity where neither extreme holds. They must be close enough for integration to create potential gains but not so close that their portfolios are essentially duplicates.

Let T denote the time at which the pair first enters this relevance region. At that moment the state is summarized by the technological distance $d \equiv d(T)$ and the convergence speed $s \equiv s(T)$. Distance captures how similar the portfolios are today. Speed captures how the gap between the portfolios is expected to evolve. I will try to show how both d and s determine the gains from merging and the division of these gains.

Merger surplus. The total surplus created by merging at state (d, s) compares the future value of the combined firm to the future values of the two firms if they remain independent:

$$S(d, s) = \Pi^M_{IJ}(d, s) - \Pi^N_I(d, s) - \Pi^N_J(d, s),$$

where $\Pi^M_{IJ}(d, s)$ is the continuation value of the merged entity and $\Pi^N_f(d, s)$ is the stand–alone continuation value of firm f if no merger occurs.

Both d and s contribute to the merger surplus. Distance d determines the current degree of technological similarity. Speed s determines the direction and intensity of future similarity. When s is positive the candidate is on a trajectory toward the incumbent's core technologies, so merging today avoids future duplication and strategic crowding. When s is negative the portfolios are drifting apart and the long–run benefits from integrating knowledge bases are smaller.

It is useful to decompose the surplus into knowledge gains and duplication losses:

$$S(d, s) = \Gamma(d, s) - \Delta(d, s).$$

- $\Gamma(d, s)$ captures knowledge–transfer and innovative gains. These gains depend on both d and s. When firms are very distant (large d) their technologies are unrelated, so the scope for learning is limited. When they are very close (small d) their portfolios overlap heavily, so additional similarity creates little new knowledge. Gains are highest at intermediate distances. Speed s matters because a positive s implies that the candidate is moving toward the incumbent's portfolio, which increases the expected overlap the merged firm can internalize. Empirically, knowledge gains peak at intermediate relatedness (Ahuja and Katila, 2001; Cloodt et al., 2006).
- $\Delta(d, s)$ captures losses from duplication, strategic crowding, and product–market overlap. These losses are largest when d is very small because the portfolios are nearly identical. They decline as d increases. Speed s matters because a positive s implies that overlap will grow in the future, raising the value of eliminating duplication through merger. When s is negative the portfolios are diverging, so the future duplication cost is lower.

Motivated by these properties, I assume that $S(d, s)$ is positive only on an interval of distances (d_H, d_L) with $0 < d_H < d_L$ and is maximized at an interior combination of (d, s). I refer to (d_H, d_L) as the absorbable band. Within this band, the firms are related enough to generate knowledge gains and the speed of convergence determines the magnitude of future duplication that merger can internalize.

Outside options and bargaining power. Let $\beta(d, s) \in [0, 1]$ denote the share of merger surplus $S(d, s)$ that the candidate receives when the firms bargain at state (d, s), so the incumbent receives $1 - \beta(d, s)$.

If we consider $\Pi_f^N(d, s)$ (the continuation value of firm f if no merger occurs) as the relevant outside option in bargaining, then it naturally depends on both primitives. A higher convergence speed influences these outside options in two ways. First, a fast–moving candidate often reflects stronger innovative capability or a more effective research team. In many innovation– driven industries patent portfolios proxy the direction and productivity of a firm's human capital, so a small firm that is rapidly entering the incumbent's technological space can credibly grow into a high–quality competitor on its own. This raises the candidate's no–merger value. Second, the same rapid movement increases the competitive pressure faced by the incumbent in the future, which lowers its no–merger value. These forces imply

$$\frac{\partial \Pi_J^N(d, s)}{\partial s} > 0 \qquad \text{and} \qquad \frac{\partial \Pi_I^N(d, s)}{\partial s} < 0.$$

Under Nash bargaining, the firm with the stronger outside option receives the larger share of the surplus. Because faster convergence strengthens the candidate's fallback position and weakens the incumbent's fallback position, the candidate's bargaining weight increases in s. I capture this pattern with the reduced–form specification

$$\beta(d, s) = \beta_0 + \beta_s s, \qquad \beta_s > 0.$$

Regulatory and integration costs. Apart from all the internal considerations of both firms, acquiring the candidate exposes the incumbent to regulatory scrutiny. Let $R(d, s)$ denote the expected regulatory and litigation cost of completing the merger when the state is (d, s). This cost depends on both primitives. First, regulators are more concerned when the two firms are already close in technology space, because high similarity can signal a loss of competition or a consolidation of innovation capability. Let d_{crit} denote the distance below which regulators view portfolios as potentially substitutable. The term $r_1(d_{\text{crit}} - d)$ captures this static concern, where $r_1 > 0$ measures how sharply scrutiny rises as d approaches the critical similarity level.

Second, convergence speed s also matters. A positive s signals a future close rival. Under the nascent–competition view, eliminating a fast–approaching competitor is more likely to be interpreted as anticompetitive. The term $r_2 s$ captures this dynamic concern, where $r_2 \geq 0$ measures how much regulators penalize the removal of a rapidly converging rival. Combining these components,

$$R(d, s) = r_0 + r_1(d_{\text{crit}} - d) + r_2 \max\{s, 0\},$$

where r_0 is a baseline regulatory cost that applies to any merger regardless of proximity. In addition, merging requires a fixed integration cost $K \geq 0$, which represents managerial distraction, cultural mismatch, and coordination costs that do not depend on (d, s).

The acquisition condition. Given the state (d, s), the incumbent's net gain from acquiring the candidate is

$$V_I(d, s) = (1 - \beta(d, s))S(d, s) - R(d, s) - K.$$

To reiterate the variables: $S(d, s)$ is the total synergy from merging at the current state (d,s). The factor $1-\beta(d, s)$ is the share of that synergy the incumbent receives after bargaining, which depends on how strong the candidate would be if it remained independent. The term $R(d, s)$ captures the expected regulatory and litigation burden of the merger, which rises when the firms are currently similar or becoming similar quickly. Finally, K is the fixed cost of integration.

The incumbent proceeds with the acquisition if and only if $V_I(d, s) \geq 0$. This inequality formalizes the requirement that synergy must be large enough, the incumbent's share must be high enough, and regulatory pressure must be low enough to justify the deal.

Speed cutoff. At any given distance d, an increase in convergence speed reduces the incumbent's net gain from acquiring the candidate through two channels. First, faster motion raises the candidate's standalone continuation value and lowers the incumbent's continuation value if no merger occurs. These outside-option effects increase the candidate's bargaining weight, so the incumbent retains a smaller fraction of the surplus: $\partial\beta(d, s)/\partial s > 0$. Second, regulators interpret rapid convergence as evidence of nascent competitive pressure, which increases the expected antitrust cost of acquisition: $\partial R(d, s)/\partial s > 0$.

Both channels imply that the incumbent's value function is strictly decreasing in s whenever

$S(d) > 0$:

$$\frac{\partial V_I(d, s)}{\partial s} = -\beta_s S(d) - r_2 < 0.$$

Because the value of acquiring falls monotonically with s, there exists a unique cutoff $\bar{s}(d)$ such that the incumbent acquires for speeds below this threshold and does not acquire above it. The cutoff is the point at which the static gains from integration exactly offset the dynamic costs generated by future motion, bargaining redistribution, and regulatory exposure.

2.4 Acquisition Characterization

Once the state (d, s) is realized inside the relevance region, the incumbent faces a well-defined choice: acquire the candidate or let the two firms continue on their independent innovation paths. The acquisition decision depends on how static technological synergy, dynamic competitive pressure, bargaining redistribution, and regulatory scrutiny jointly determine the incumbent's net gain.

From the acquisition condition to a speed cutoff. The incumbent acquires the candidate if and only if its net gain from merging is non-negative. Using the value function introduced earlier, the condition is

$$V_I(d, s) \geq 0.$$

At a fixed distance d, the incumbent compares how much static synergy exists today with how costly the merger becomes as the candidate moves faster toward its technological core. Distance determines *how beneficial* integration is; speed determines *how urgent* and *how expensive* it is to neutralize the candidate's trajectory.

Substituting the incumbent's value function,

$$V_I(d, s) = (1 - \beta(d, s))S(d) - R(d, s) - K,$$

and recalling $\beta(d, s) = \beta_0 + \beta_s s$ and $R(d, s) = r_0 + r_1(d_{\text{crit}} - d) + r_2 s$ gives:

$$(1 - \beta_0 - \beta_s s)S(d) \geq r_0 + r_1(d_{\text{crit}} - d) + r_2 s + K.$$

This expression captures a fundamental tradeoff:

Synergy $S(d)$ is the only force *raising* the merger's value, while bargaining-power drift (β_s) and regulatory drift (r_2) are the forces that make fast convergence costly.

Solving for the highest speed the incumbent tolerates,

$$s \leq \bar{s}(d) = \frac{(1 - \beta_0)S(d) - r_0 - r_1(d_{\text{crit}} - d) - K}{\beta_s S(d) + r_2}. \tag{1}$$

The numerator contains the static terms of the merger. The component $(1 - \beta_0)S(d)$ is the static share of synergy the incumbent keeps, and r_0, $r_1(d_{\text{crit}} - d)$, and K are static regulatory and integration costs. The numerator is therefore the static net gain from merging at distance d.

The denominator $\beta_s S(d) + r_2$ contains the dynamic terms. The part $\beta_s S(d)$ reflects how faster convergence raises the candidate's bargaining power, and r_2 captures the dynamic regulatory penalty. The denominator is therefore the dynamic cost of speed.

The cutoff $\bar{s}(d)$ is the ratio of static net gain to dynamic cost.

Absorbable region and the collision zone. For states where $S(d) > 0$, the speed threshold $\bar{s}(d)$ partitions feasible targets into those the incumbent chooses to acquire and those it does not. Define

$$\mathcal{A} = \{(d, s) : d \in (d_L, d_H),\ s \leq \bar{s}(d)\}, \qquad \mathcal{C} = \{(d, s) : d \in (d_L, d_H),\ s > \bar{s}(d)\}.$$

The acquisition region $\mathcal{A}$ contains firms that are neither too distant (so synergy is still positive) nor approaching too quickly (so dynamic bargaining and regulatory costs do not outweigh synergy).

The collision zone $\mathcal{C}$ contains firms whose technological trajectory is too aggressive: despite having positive synergy, they approach rapidly enough that buying them is too expensive, either because (i) they would extract a large share of the surplus, or (ii) regulators interpret the acquisition as eliminating a nascent competitor.

Thus, two things must hold simultaneously: (i) static proximity must be in the absorbable band. (ii) convergence cannot exceed $\bar{s}(d)$.

Realized acquisitions must satisfy both.

Quadratic synergy and why it matters. To derive explicit comparative statics, I assume the synergy function takes the quadratic form:

$$S(d) = \alpha(d - d_L)(d_H - d), \qquad \alpha > 0.$$

This functional form is consistent with the model setup where:
- when d is close to d_L, the firms are too similar; combining them adds little,
- when d is close to d_H, the firms are too distant; little can be recombined,
- in the interior, both portfolios contain distinct and complementary knowledge.

With this specification, we obtain:

$$\bar{s}(d) = \frac{(1 - \beta_0)\alpha(d - d_L)(d_H - d) - r_0 - r_1(d_{\text{crit}} - d) - K}{\beta_s \alpha(d - d_L)(d_H - d) + r_2}.$$

In this form, the numerator is "how much is gained by merging today," and the denominator is "how costly is the candidate's motion." This ratio determines the maximum aggressiveness the incumbent can tolerate at distance d.

Comparative statics in bargaining sensitivity β_s. Holding d fixed,

$$\frac{\partial \bar{s}(d)}{\partial \beta_s} = -\frac{S(d)\big[(1-\beta_0)S(d) - r_0 - r_1(d_{\text{crit}} - d) - K\big]}{(\beta_s S(d) + r_2)^2}.$$

The sign is negative whenever the merger is attractive for some speeds. Economically, β_s measures how quickly the candidate's bargaining power rises with speed. If bargaining reacts sharply to faster convergence, then even modest speeds make the candidate too expensive. Thus, a larger β_s shrinks the absorbable region and expands the collision zone.

Comparative statics in regulatory pressure r_1. For fixed distance and positive synergy,

$$\frac{\partial \bar{s}(d)}{\partial r_1} = -\frac{d_{\text{crit}} - d}{\beta_s S(d) + r_2}.$$

Regulators who penalize similarity strongly (high r_1) make acquisitions less tolerable, especially when $d < d_{\text{crit}}$. A higher r_1 steepens the dynamic regulatory penalty and reduces the incumbent's willingness to acquire fast-moving rivals.

Comparative statics in synergy $S(d)$. Differentiating:

$$\frac{\partial \bar{s}(d)}{\partial S(d)} = \frac{\beta_s[r_0 + r_1(d_{\text{crit}} - d) + K]}{(\beta_s S(d) + r_2)^2} > 0.$$

Here, higher synergy expands the absorbable region because it increases the static gain from integrating the portfolios while leaving the dynamic bargaining and regulatory drift essentially unchanged. This makes the incumbent more willing to tolerate faster motion.

Distance profile and why the boundary collapses. As $d \to d_L$ or $d \to d_H$, we have $S(d) \to 0$, so the cutoff becomes

$$\bar{s}(d) \to -\infty.$$

This means that when the portfolios are either too similar or too distant, the incumbent will never acquire, regardless of how slowly the candidate approaches. This captures the intuition that mergers between almost identical portfolios add little value and raise regulatory flags, while mergers between unrelated portfolios generate no synergy.

At intermediate distances, synergy is maximized and the denominator $\beta_s S(d) + r_2$ is stable. Thus, $\bar{s}(d)$ is highest in the interior. This gives rise to a donut-shaped acquisition region in (d, s)-space: the incumbent acquires targets that are neither too distant nor too similar and are converging slowly enough that bargaining and regulatory drift remain manageable.

Summary of predictions. The model generates three testable predictions. (i) Realized acquisitions cluster at intermediate distances in (d_L, d_H). (ii) Conditional on distance, acquisition probability declines in convergence speed and drops sharply above $\bar{s}(d)$. (iii) Greater negotiating sensitivity (higher β_s) or regulatory pressure (higher r_1) shrinks the absorbable region, while stronger technological synergy expands it.

2.5 Link to Empirical Specifications

The primitives of the framework map directly into the empirical measures constructed from patent embeddings. The distance $d(t)$ corresponds to the Euclidean (or cosine) distance between firm–year centroids. The convergence speed $s(t)$ corresponds to the change in distance over time divided by the horizon length. The absorbable band (d_L, d_H) corresponds to the region around the incumbent in technology space, and the collision zone is the subset of this band where observed $s(t)$ exceeds the implied threshold $\bar{s}(d)$.

The acquisition rule based on $V_I(d, s)$ motivates the empirical design in which I construct acquirer–year risk sets of feasible public targets and estimate how acquisition probabilities vary with distance and convergence speed. The examiner–leniency exposure instrument provides exogenous variation in the evolution of granted patent portfolios, and therefore in d and s, allowing me to test whether the avoidance of fast–approaching rivals is causal.

Before moving to the empirical specifications and results I quickly introduce the data sources I have used for this project.

3 Data

This section describes the construction of the patent, merger, and examiner-level datasets and the procedures used to link them through firm identifiers.

The analysis begins with the extended patent–firm database of Kogan et al. (2017b), covering all U.S. patents obtained by public firms from 1836–2023. I augment this dataset with Cooperative Patent Classification (CPC) codes, permanant company number PERMCO supplied by the authors. I merge these patents with USPTO PatentView abstracts (1941–2023) and generate 1,024-dimensional PatentBERT(Lee and Hsiang, 2020) embeddings for each abstract, L2-normalized to lie on the unit sphere. After merging with CRSP–Compustat through the CCM link, the final sample includes 2.6 million firm-linked patents for 7,926 firms during 1961–2023. These embeddings form the basis for dynamic measures of technological proximity and innovation direction. Figure 1 summarizes this distribution.

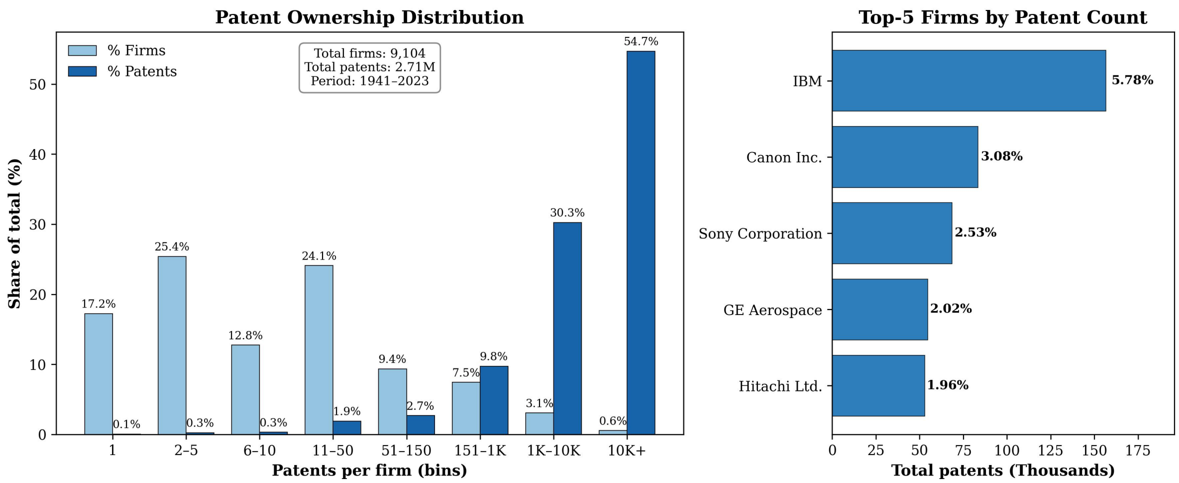

Figure 1: Distribution of historical patent counts for acquirers and targets in the matched public–public sample. The distribution is highly skewed: a small number of firms hold the vast majority of patents, while many firms possess only a limited number of patents.

The global Refinitiv SDC M&A database reports over two million announced transactions worldwide. I restrict the sample to completed deals, yielding 374,547 unique transactions. Because SDC does not provide identifiers consistent with CRSP or Compustat, acquirer and target CUSIPs are matched to CRSP's monthly security file using six-digit CUSIP roots.

For each CRSP firm, I construct trading intervals between the first and last observed month of activity and define a firm as public if its CUSIP was active on the deal completion date. This procedure identifies 68,621 public acquirers (18.3%), 5,507 public targets (1.5%), and 3,115 transactions in which both sides are public (0.83%).

Merging these CRSP-linked firms with the KPSS patent database through `PERMCO` shows that 58.1% of acquirers and 36.7% of targets hold at least one patent. The intersection yields 831 public–public transactions where both parties are patent-active. Among these, 455 targets and 144 acquirers hold fewer than ten patents, corresponding to 503 deals in which at least one side exhibits limited patenting activity. This subset forms the baseline for analyzing technological proximity and post-merger innovation outcomes.

To obtain exogenous variation in measured innovative activity, I incorporate the USPTO PatEx dataset (1970–2022). PatEx contains application–level information for more than 13 million filings, including examiner IDs, art-unit assignments, application years, and grant/denial outcomes. I used this data to construct a firm-year level leniency exposure.

4 Empirical Design

4.1 Construction of Technological Measures

Each firm–year (f, t) is represented by a 1,024-dimensional centroid of the embeddings of the abstracts of its patents (PatentBERT embeddings):

$$\boldsymbol{c}_{f,t} = \frac{1}{N_{f,t}} \sum_{p \in \mathcal{P}_{f,t}} \boldsymbol{v}_p.$$

To characterize different horizons of innovation, I construct two additional centroids:

$$\boldsymbol{c}_{f,t}^{cum} = \frac{1}{N_{f,t}^{cum}} \sum_{\tau \leq t} \sum_{p \in \mathcal{P}_{f,\tau}} \boldsymbol{v}_p, \qquad \boldsymbol{c}_{f,t}^{(5)} = \frac{1}{N_{f,t}^{(5)}} \sum_{\tau = t-4}^{t} \sum_{p \in \mathcal{P}_{f,\tau}} \boldsymbol{v}_p.$$

The cumulative centroid $\boldsymbol{c}_{f,t}^{cum}$ captures the firm's long-run technological base, while the five-year centroid $\boldsymbol{c}_{f,t}^{(5)}$ highlights its recent innovation direction. Together, these centroids trace continuous firm trajectories in patent space for all distance and convergence calculations.

Technological Distance, Convergence, and Threat: For each acquirer–candidate pair (a, j) and relative year t, I compute cosine and Euclidean distances:

$$d_{aj,t}^{cos} = 1 - \frac{\boldsymbol{c}_{a,t} \cdot \boldsymbol{c}_{j,t}}{\|\boldsymbol{c}_{a,t}\| \, \|\boldsymbol{c}_{j,t}\|}, \qquad d_{aj,t}^{euc} = \|\boldsymbol{c}_{a,t} - \boldsymbol{c}_{j,t}\|_2.$$

Smaller values of d indicate greater similarity in innovation portfolios.

Convergence speed is defined over one- and five-year horizons ($+S$ denotes convergence):

$$S_{aj,t}^{(1)} = d_{aj,t-1} - d_{aj,t}, \qquad S_{aj,t}^{(5)} = \frac{d_{aj,t-5} - d_{aj,t}}{5}.$$

The technological threat index scales convergence by current proximity:

$$T_{aj,t}^{(h)} = \frac{S_{aj,t}^{(h)}}{d_{aj,t} + \kappa}, \qquad h \in \{1, 5\}, \ \kappa = 1.$$

Pre–Deal Motion in Technology Space: To summarize the raw dynamics of technological proximity prior to mergers, I track the yearly cosine and Euclidean distances between each realized acquirer–target pair during the ten years preceding the deal. For each pair, I compute the percent change in distance relative to its value ten years before the transaction.

Figure 2 shows the average trajectory. Both distance measures indicate systematic pre–deal convergence: on average, acquirers and targets move closer in technology space as the deal approaches.

This pattern provides descriptive support for the idea that realized mergers occur within the intermediate region where portfolios gradually become more related.

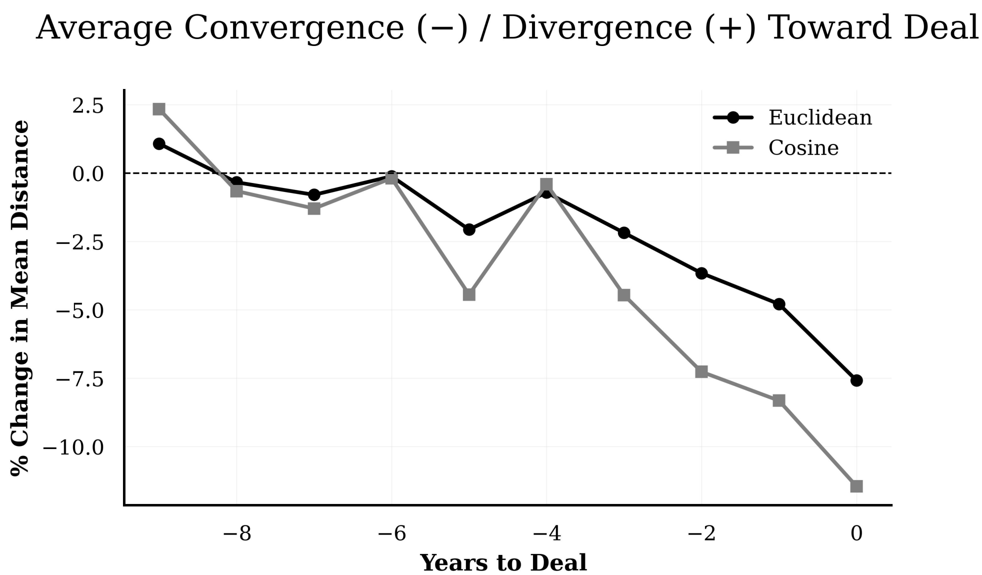

Figure 2: Percent change in cosine and Euclidean distance between acquirer and target over the ten years before the deal. Distances are normalized to zero at $t = -10$.

4.2 Constructing the Risk Set of Potential Targets

A central empirical challenge is that observing only realized M&A pairs does not reveal which alternative targets were actually feasible in a given year. To estimate target–choice behavior, it is necessary to define the set of potential target firms that an acquirer could plausibly have purchased instead of the observed target.

For each completed public–public deal $(i \to j)$ in year t, I treat i as the acquirer and j as the realized target. I first define the *focal target profile* using j's most recent characteristics measured over the window $[t-5, t]$.

The candidate pool consists of all other (Limiting to 25 candidates per target) patenting firms $k \neq j$ that:

(i) operate in the same 2-digit SIC industry during $[t-5, t]$, (ii) are active patenters during the same window, and (iii) satisfy the economic calipers:

$$|\ln A_k - \ln A_j| \leq 1, \; |\xi_k - \xi_j| \leq 5, \; |CPC_k - CPC_j| \leq 10.$$

Within this set, all firm characteristics are standardized, and I select the $K = 25$ nearest neighbors in Euclidean distance:

$$d_{jk} = \|\tilde{\mathbf{x}}_{j,t} - \tilde{\mathbf{x}}_{k,t}\|_2, \qquad k \in \arg\min d_{jk}.$$

These K firms represent the set of feasible targets available to acquirer i in year t. The realized target j is then compared to its feasible alternatives on technological distance, convergence speed, threat, and financial covariates. This construction allows estimation of target–choice models that respect realistic substitution patterns and acquirer-specific opportunity sets.

Table 1 reports balance between realized targets and their matched candidate firms across standard financial and innovation variables. The lack of systematic differences supports the empirical validity of the constructed risk sets.

Table 1: Balance diagnostics for target–candidate matching

Variable	**Target mean**	**Candidate–Target mean**	**Difference**	***t*-stat**	***p*-value**
$\ln AT$	5.907	5.136	-0.771	-0.485	0.630
ROA	-0.035	-0.028	0.007	0.196	0.845
Q	1.134	0.963	-0.172	0.566	0.576
Leverage	0.226	0.141	-0.086	-1.087	0.282
Cash/AT	0.216	0.195	-0.021	0.117	0.907
R&D/AT	0.108	0.080	-0.027	-0.099	0.922
CAPX/AT	0.038	0.035	-0.004	1.939	0.058
MKVALT ($M)	4,156	3,660	-496	1.048	0.305
$\xi^{\text{real}}_{\text{mean}}$	6.408	5.610	-0.798	-0.341	0.739
$\text{Citations}_{\text{mean}}$	18.584	15.642	-2.942	0.059	0.954
CPC unique count	16.000	33.676	17.676	-0.873	0.402

Notes. This table compares realized **targets** to their matched feasible candidates within a five-year lookback window. Matching restricts to firms in the same two-digit SIC and with nonzero patenting activity. Candidates are chosen using standardized Euclidean distance over firm characteristics $\{\ln AT, ROA, Q, LEV, CASH/AT, XRD/AT, CAPX/AT, MKVALT, \xi^{\text{real}}_{\text{mean}}, \text{Citations}_{\text{mean}}, \text{CPC unique count}\}$. Differences represent candidate minus target means across deals; the absence of significant differences indicates strong technological comparability. ξ represents the KPSS market response to each patent.

4.3 Baseline Target–Choice Results (OLS)

Using the acquirer–year risk sets, I estimate the probability that candidate k becomes the chosen target in deal-year t as a function of its technological threat to the acquirer. The baseline specification is:

$$\boldsymbol{is\_target_{ikt} = \alpha_i + \gamma_t + \beta\, T5\_euc, ikt + X'_{ikt}\delta + \varepsilon_{ikt}.}$$

where α_i denotes acquirer fixed effects, γ_t denotes deal-year fixed effects, $T5_{euc}$ is the five-year technological threat index, and X_{ikt} includes firm-level covariates such as changes in assets, ROA, q, leverage, cash holdings, R&D intensity, capital investment, market value, and innovation activity.

Table 2: Baseline OLS: Threat and Target Selection

	Coefficient	**Std. Error**	**t-stat**
$T5\_euc$	-0.174	0.040	-4.38
Δ ln(Assets)	0.0013	0.0010	1.38
ΔROA	0.0001	0.0013	0.12
Δq	0.00009	0.00010	0.96
ΔLeverage	0.0018	0.0027	0.68
ΔCash/AT	0.0016	0.0021	0.74
ΔR&D/AT	-0.0010	0.0030	-0.34
ΔCAPX/AT	-0.0011	0.0080	-0.14
ΔMktVal	2.8×10^{-9}	4.1×10^{-9}	0.68
ΔPatents	1.7×10^{-6}	1.3×10^{-6}	1.23
ΔCitations	-3.9×10^{-8}	3.5×10^{-8}	-1.13
Acquirer FE (α_i)	Yes		
Deal-year FE (γ_t)	Yes		
Clusters	255		
Observations	18,513		

Notes: This table reports OLS estimates from the target-choice regression using acquirer–year risk sets. The dependent variable equals one if candidate k is the realized target in deal-year t. $T5\_euc$ is the five-year technological threat index. All specifications include acquirer fixed effects and deal-year fixed effects. Standard errors are clustered at the acquirer level.

The OLS estimate shows a negative and statistically significant association between threat and target choice: a one–standard-deviation increase in $T5_{euc}$ (SD = 0.018) lowers the probability of being selected as the target by roughly 0.32 percentage points.

Economically, firms are *less likely* to acquire candidates that are rapidly converging toward their technological core. This pattern suggests defensive avoidance rather than aggressive consolidation. Since $T5_{euc}$ is measured with error due to patent-grant randomness and imperfect distance measurement, attenuation bias likely pushes the OLS coefficient toward zero; the true effect is expected to be stronger.

The measured threat index $T5_{euc}$ depends on observed patent grants, which are not exogenous. Firms anticipating a merger may shift their patenting toward technologies that appear attractive to a potential acquirer, creating reverse causality. Common shocks like industry-level technology cycles or funding conditions can simultaneously affect patenting intensity and the likelihood of M&A activity, generating omitted-variable bias. In addition, random noise in patent embeddings introduces measurement error that attenuates the OLS coefficient toward zero. Together, these issues can bias the estimated relationship between technological threat and target choice, potentially masking the true negative effect. To recover exogenous variation in perceived innovation intensity, I use examiner leniency exposure as an instrument, which affects patent grant outcomes but is orthogonal to firms' acquisition intent.

4.4 Identification Strategy–Instrumental Variable

Examiner-specific heterogeneity in granting behavior provides plausibly exogenous variation in patent approvals, following Frakes and Wasserman (2014, 2020), Galasso and Schankerman (2018), and Sampat and Williams (2019). I construct examiner leniency using the USPTO Patent Examination Research Dataset (PatEx, 2022 release), which covers more than 13 million publicly visible provisional and non-provisional U.S. applications and over one million PCT filings. PatEx is compiled from the Patent Examination Data System (PEDS) as of June 2023 and includes examiner identifiers, art units, filing dates, disposition outcomes, and status codes.

I apply the following filters before constructing the instrument. First, I restrict to applications with (i) a named examiner, (ii) a valid art unit, (iii) a non-missing filing date, and (iv) an issued U.S. patent number. I then match issued applications to the Kogan et al. (2017b) patent panel by patent number in order to attach a CRSP firm identifier (`permco`). Second, I collapse examiners to an "*examiner key,*" defined as the examiner's full name interacted with the art unit. I drop examiner keys with fewer than five observed applications in a given year to avoid high-variance examiners. Third, I drop art units with fewer than twenty observed applications and remove art units whose grant rate is essentially degenerate (grant ratio $\leq 1\%$ or $\geq 99\%$). Figure 3 shows the resulting distribution of art-unit size and grant rates.

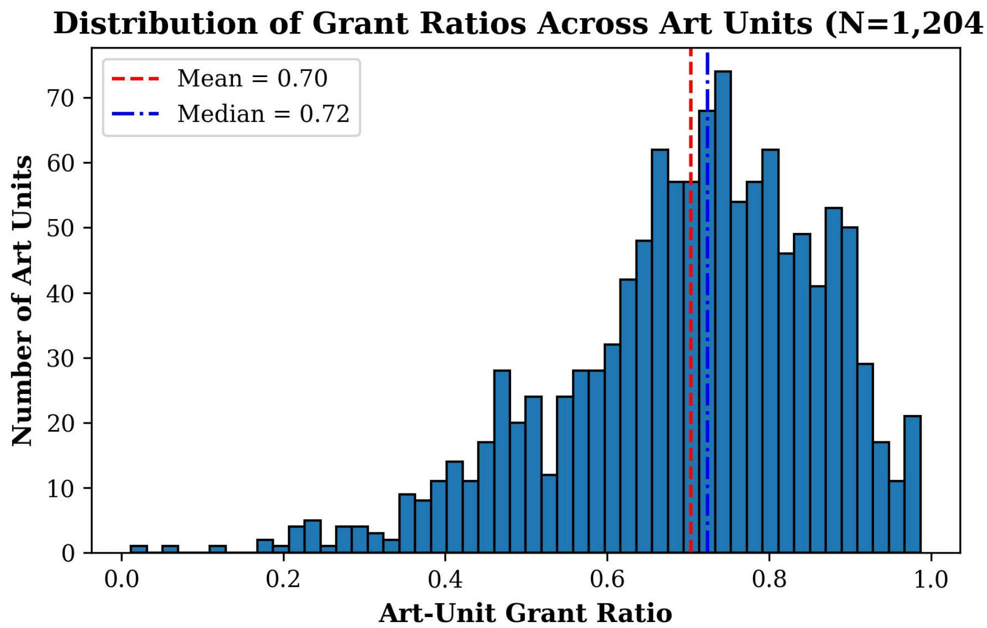

Figure 3: Art-unit level application grant rates after filtering on examiner coverage and art-unit activity.

For each application a examined by examiner e in year t, I compute (i) the examiner's leave-one-out grant rate and (ii) its residual within art-unit–year:

$$Z_{aet} = \frac{G_{e,t} - Grant_a}{N_{e,t} - 1}, \qquad Z^{resid}_{aet} = Z_{aet} - \overline{Z}_{aut},$$

where $N_{e,t}$ and $G_{e,t}$ are the total examined and granted applications for examiner e in year t, $Grant_a$ is the grant outcome for application a, and $\overline{Z}_{aut}$ is the mean leniency in art unit au and year t. The first term removes mechanical dependence on the focal application; the second removes art-unit–year shifts in examination policy.

Finally, I aggregate these examiner-level shocks to the firm–year level. Let $\mathcal{A}_{f,t}$ be the set of granted applications filed in year t by firm f (matched via `permno`). The firm's exposure to examiner leniency is:

$$Z^{\text{firm}}_{f,t} = \frac{1}{|\mathcal{A}_{f,t}|} \sum_{a \in \mathcal{A}_{f,t}} Z_{aet}, \qquad Z^{\text{resid}}_{f,t} = \frac{1}{|\mathcal{A}_{f,t}|} \sum_{a \in \mathcal{A}_{f,t}} Z^{resid}_{aet}.$$

This produces a panel of 8,191 CRSP-linked firms over 1970–2022 (53 years), with 2,268,139 granted applications contributing to the instrument. Figure 4 shows the cross-sectional distribution of $Z^{\text{resid}}_{f,t}$ across firm–years. These measures provide firm-year shocks to the probability of patent grant, which I later use to construct instruments for patenting success and firm dynamics in the M&A analysis.

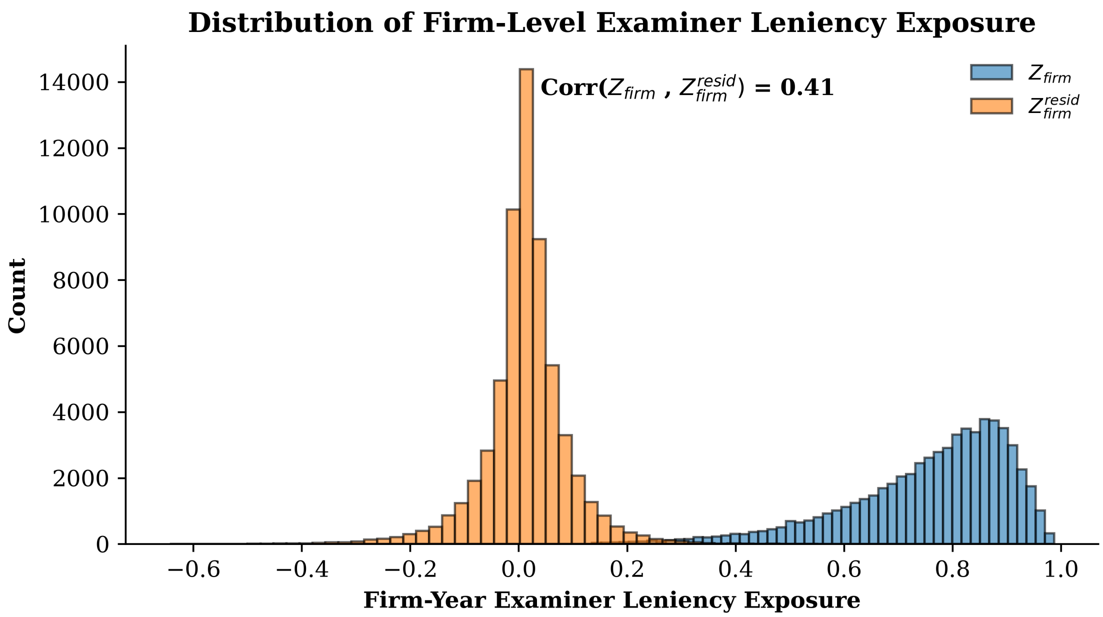

Figure 4: Distribution of firm-year exposure to examiner leniency, $Z_{f,t}$, and its residualized counterpart, $Z_{f,t}^{\text{resid}}$. Both measures reflect variation in examiner grant propensity across firms and years, with $Z_{f,t}^{\text{resid}}$ purging art-unit–year effects.

4.5 Instrumental Variable Results

This section reports the two–stage least squares estimates using examiner–leniency exposure as an instrument for the technological threat index $T5_{euc}$. As described earlier, examiner assignment within art units is quasi-random, and examiners differ persistently in grant propensity. Firms exposed to more lenient examiners appear more technologically active not because they strategically accelerate innovation, but because a greater share of their applications are approved. This quasi-random variation generates exogenous shifts in measured convergence and threat that are orthogonal to acquisition intent.

4.5.1 First Stage: Leniency and Perceived Threat

The first-stage regression relates the threat index $T5_{euc}$ for acquirer–candidate pair (i, k) in year t to the five-year residualized examiner-leniency exposure:

$$\boldsymbol{T5_{euc,ikt} = \pi_1 Z5_{ikt}^{resid} + X_{ikt}'\rho + \mu_i + \tau_t + \nu_{ikt}.}$$

where μ_i are acquirer fixed effects, τ_t are deal-year fixed effects, and X_{ikt} contains changes in firm-level characteristics.

Table 3 reports the results. The coefficient on $Z5^{resid}$ is positive and statistically significant: more lenient examiner exposure increases the measured technological threat between an acquirer and a candidate. This reflects the fact that when a candidate's patents are more likely to be granted for reasons unrelated to strategy, its recent patent trajectory shifts inward toward the acquirer in embedding space. The Kleibergen–Paap statistic of 11.93 indicates a relevant first stage.

Table 3: First Stage: Examiner Leniency and Perceived Threat

	Coefficient	Std. Error	t-stat
$Z5^{resid}$	0.0100	0.0029	3.45
Δ ln(Assets)	−0.00135	0.00033	−4.09
ΔROA	−0.00015	0.00081	−0.19
Δq	−0.00013	0.00008	−1.55
ΔLeverage	0.00035	0.00135	0.26
ΔCash/AT	0.00000	0.00093	0.00
ΔR&D/AT	−0.00294	0.00170	−1.73
ΔCAPX/AT	0.00190	0.00401	0.47
ΔMktVal	3.4×10^{-9}	1.3×10^{-9}	2.56
ΔPatents	1.26×10^{-6}	2.94×10^{-7}	4.27
ΔCitations	1.2×10^{-9}	1.8×10^{-8}	0.07
Acquirer FE	Yes		
Deal-Year FE	Yes		
Kleibergen–Paap F	11.93		
Clusters	239		
Observations	15,634		

Notes: The dependent variable is the five-year technological threat index $T5_{euc}$. The instrument is the five-year residualized examiner-leniency exposure $Z5^{resid}$. Standard errors cluster at the acquirer level.

4.5.2 Second Stage: Threat and Target Choice

The second stage estimates the causal effect of technological threat on the probability that a candidate is chosen as the target:

$$\boldsymbol{is\_target_{ikt} = \alpha_i + \gamma_t + \beta\,\widehat{T5_{euc,ikt}} + X'_{ikt}\delta + \varepsilon_{ikt}.}$$

Table 4 reports the results. The IV estimate of β is negative and statistically significant: a one–standard-deviation increase in the threat index (SD = 0.018) decreases the probability of being selected by approximately 3.6 percentage points. This effect is substantially larger in magnitude than the corresponding OLS estimate, consistent with attenuation from measurement error in observed patent signals.

Table 4: Second Stage: Threat and the Probability of Being Acquired

	Coefficient	Std. Error	t-stat
$\widehat{T5_{euc}}$	-1.953	0.854	-2.29
$\Delta\ln(\text{Assets})$	-0.00040	0.00179	-0.22
ΔROA	-0.00030	0.00213	-0.14
Δq	-0.00008	0.00026	-0.30
ΔLeverage	0.00361	0.00417	0.87
ΔCash/AT	0.00127	0.00329	0.38
ΔR&D/AT	-0.00604	0.00568	-1.06
ΔCAPX/AT	0.00201	0.01344	0.15
ΔMktVal	7.0×10^{-9}	5.4×10^{-9}	1.29
ΔPatents	4.05×10^{-6}	1.91×10^{-6}	2.12
ΔCitations	-6.96×10^{-8}	6.31×10^{-8}	-1.10
Acquirer FE	Yes		
Deal-Year FE	Yes		
Clusters	239		
Observations	15,634		

Notes: The dependent variable equals one if candidate k is the realized target for acquirer i in year t. The threat index is instrumented using examiner-leniency exposure. Standard errors cluster at the acquirer level.

4.6 Empirical Patterns in Technological Motion

The threat index $T5_{euc}$ can be positive or negative because it reflects the ratio of convergence speed to current technological distance. A firm may be close but drifting away (small d and negative S), or distant but moving sharply toward the incumbent (large d but strongly positive S). The two panels in Figure 5 illustrate this geometry. When the technological gap is large, even a moderate directional shift creates a "big jump" in measured convergence. Conversely, for firms already close to the incumbent's core, small portfolio adjustments generate only a "small jump." This variation produces both positive and negative values of $T5_{euc}$ and provides the identifying variation needed for the empirical tests.

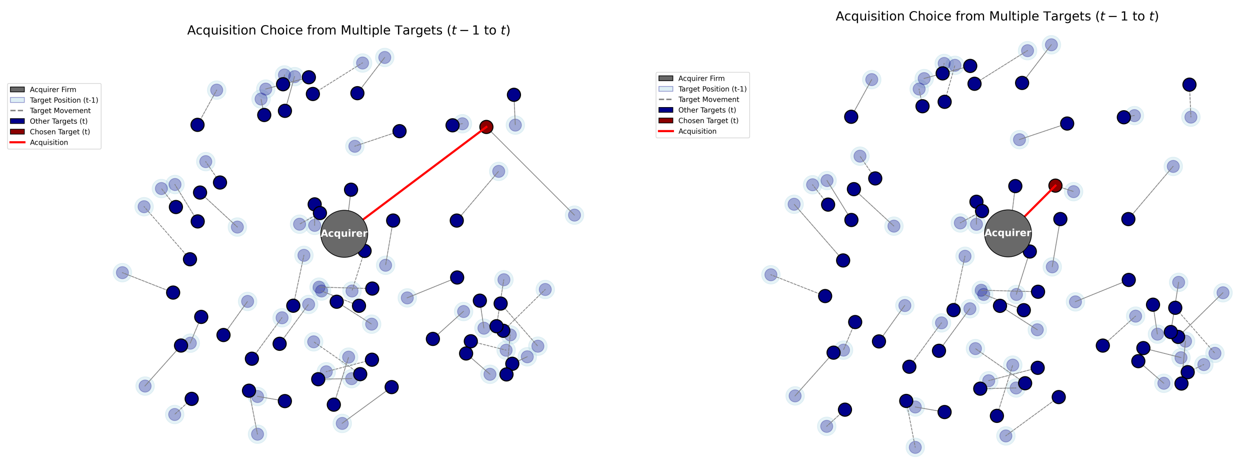

Big Gap, Big Jump

Small Gap, Small Jump

Figure 5: Big Gap, Big Jump (left) and Small Gap, Small Jump (right). The figure is to visually illustrate the difference, and it is not from the actual deals.

The IV estimates in Table 5 show that these geometric patterns have economically meaningful consequences. First, the positive coefficient on technological distance confirms that acquirers favor candidates in the absorbable band: firms that are neither too close nor too distant. This is consistent with the model's hump-shaped surplus in $S(d)$, in which intermediate distances yield the strongest complementarities. Second, the strongly negative coefficient on convergence speed indicates that firms moving rapidly toward the incumbent's core are systematically avoided. This aligns with the collision-zone logic of the framework: high s raises the target's bargaining power and expected regulatory burden, lowering $V_I(d, s)$ and making the acquisition unattractive.

Table 5: Instrumental Variable Results: Complementarity vs. Speed

	First Stage		Second Stage (2SLS)	
	KP F	p-val	IV Coef. (β)	Effect of $+1\sigma$ in regressor
(1) Complementarity: d_{euc}	12.08	0.001	$+0.192^{**}$	$1\sigma \uparrow \Rightarrow$ $+2.1$ pp $\uparrow$
(2) Speed: $S5_{\text{euc}}$	11.32	0.001	-1.363^{**}	$1\sigma \uparrow \Rightarrow$ -3.4 pp $\downarrow$
Obs.	26,765		26,765 / 16,424	
Clusters (acq_permco)	244		244 / 239	

Notes: This table reports instrumental-variable estimates of technological complementarity (d_{euc}) and convergence speed ($S5_{\text{euc}}$) on target choice. The first-stage columns show Kleibergen–Paap F-statistics and p-values for the examiner-leniency instruments. The second-stage columns report IV coefficients and the implied marginal effects of a one–standard-deviation increase in the regressor on the probability of being acquired. Standard errors are clustered at the acquirer level.

Table 6 provides a complementary test using within–risk-set differences. Conditional on being feasible candidates in the same acquirer–year, realized targets do not differ meaningfully in static distance or overall threat level, but they exhibit significantly lower convergence speed. This pattern confirms that avoidance of fast-approaching rivals holds even after conditioning on observables that define the feasible set. Although acquirers and targets tend to converge technologically before the deal, the chosen targets are those whose trajectories remain relatively slower and more absorbable.

Table 6: Targets vs. Feasible Candidates: Technological Characteristics

Variable	Mean	Std. Dev.
$\Delta T5_{\text{euc}}$ (Threat Index)	-0.0013	0.0220
$\Delta S5_{\text{euc}}$ (Convergence Speed)	-0.0026^{***}	0.0077
Δd_{euc} (Technological Distance)	0.0055	0.1114

Notes: Differences are computed within each acquirer–year risk set as (Target − mean of candidate pool). t-tests are for H_0: mean difference = 0. $^{***}p < 0.01$.

Taken together, the evidence indicates that negative values of the threat index correspond to dynamics favorable for acquisition, firms drifting away or converging only gradually, while positive and high values capture the rapid motion associated with the collision zone. The empirical patterns mirror the theoretical predictions: incumbents acquire complementary, slow-moving targets within the absorbable band and systematically avoid candidates whose innovation paths are bending too quickly toward their technological core.

4.7 Interpretations and Connection to the Theoretical Framework

The IV results align closely with the predictions of the simple framework developed in the theory section. In the model, merger surplus is hump-shaped in technological distance and decreasing in convergence speed once the candidate enters the collision zone: for a given distance in the absorbable band (d_L, d_H), higher s raises the target's bargaining power and expected regulatory costs, reducing the incumbent's net gain $V_I(d, s)$. The empirical threat index $T5_{euc}$ is a direct analogue of this geometry—capturing convergence speed scaled by distance—and the examiner-leniency instrument generates exogenous shifts in the observable component of this motion.

Integration frictions and absorptive capacity. The positive IV effect for technological complementarity is consistent with mergers occurring in the absorbable region of the model. At intermediate distances, firms' knowledge bases differ enough to generate synergies but not so much that integration becomes infeasible. Candidates that are too distant contribute little new capability, while candidates that are too close offer little incremental synergy. The empirical fact that realized targets cluster in this middle region matches the hump-shaped synergy term $S(d)$.

Bargaining power and dynamic rivalry. The strong negative causal effect of convergence speed reflects the mechanism by which high s enhances the target's bargaining weight $\beta(d, s)$. A fast-approaching candidate represents a future competitive threat, increasing its outside option and reducing the incumbent's share of surplus. The IV estimates show that when a quasi-random shock makes a candidate *appear* to be converging more rapidly, the incumbent becomes significantly less likely to acquire it. This is precisely the comparative static implied by the collision-zone threshold $\bar{s}(d)$.

Regulatory scrutiny in the collision zone. Rapid convergence also amplifies expected regulatory costs. When portfolios grow too similar, consolidation concentrates innovation capacity, raising the likelihood of antitrust review. In the model, this is captured by $R(d)$ increasing as distance shrinks. Empirically, the fact that convergence speed has a larger causal effect than distance suggests that acquirers respond not only to static similarity but to the rate at which a transaction would move the pair into a regulatory-sensitive region.

Information frictions and perceived trajectory risk. The IV results highlight a perception channel. Examiner-leniency shocks inflate observed innovation motion by increasing the likelihood that marginal patents are granted. This makes the candidate *look* more convergent even if its underlying R&D direction is unchanged. Because incumbents base acquisition choices on observed (d, s) rather than latent fundamentals, measurement error in $T5_{euc}$ attenuates OLS, whereas the IV estimates recover the causal avoidance response predicted by the model.

Unified mechanism. Taken together, the evidence supports a mechanism in which incumbents partition potential targets into three geometric regions. Firms that are too distant generate little synergy; fast-approaching or overly similar firms fall into the collision zone because of bargaining power and regulatory frictions; and acquisitions occur primarily among candidates in the absorbable band—those at intermediate distances with moderate or slow convergence. The strengthened IV coefficient provides a causal validation of the model's core prediction: dynamic innovation motion, not static technological overlap, determines which firms incumbents choose to acquire and which they prefer to leave as future competitors.

5 Conclusion

This paper shows that the dynamics of firms' technological portfolios play an important role in selection into mergers. By representing firms as evolving trajectories in a patent based technology space, the analysis separates two primitives of innovation motion: the technological distance between firms at a given moment and the speed at which a candidate's research direction moves toward or away from an incumbent. A simple framework illustrates why incumbents value candidates at intermediate distances but avoid those whose portfolios converge too quickly. Fast convergence strengthens the target's outside option and increases expected regulatory scrutiny, creating a collision region where acquisition becomes unattractive even when firms are technologically relatexd.

Empirically, transformer based patent embeddings make it possible to trace these trajectories at the firm year level. Acquirer target pairs lie in an interior band of technological distances and exhibit gradual pre deal convergence. Conditional on distance, the likelihood that a candidate is acquired declines sharply with convergence speed. This pattern becomes stronger when I instrument for the evolution of granted patent portfolios using examiner driven variation in grant propensity. The evidence indicates that dynamic innovation motion, not only static technological similarity, is a first order determinant of which firms are absorbed and which remain independent competitors.

Overall, the results suggest that merger analysis should account for both the location of firms in technology space and the speed at which they move through it. Dynamic technological motion contains economically meaningful information about future rivalry, bargaining, and regulatory exposure, and these forces help shape the structure of acquisition flows. Future research can build on this perspective by studying how technological trajectories interact with acquirer heterogeneity, post merger integration, and industry level innovation cycles.

References

Ahuja, G. and Katila, R. (2001). Technological acquisitions and the innovation performance of acquiring firms: A longitudinal study. *Strategic Management Journal*, 22(3):197–220.

Althammer, S., Buckley, M., Hofstätter, S., and Hanbury, A. (2021). Linguistically informed masking for representation learning in the patent domain. *CoRR*, abs/2106.05768. Presented at the PatentSemTech Workshop, SIGIR 2021.

Andrade, G., Mitchell, M., and Stafford, E. (2001). New evidence and perspectives on mergers. *Journal of Economic Perspectives*, 15(2):103–120.

Barber, B. M., Jiang, W., and Morse, A. (2022). Can firms avoid tough patent examiners through examiner shopping? *Strategic Management Journal*. Forthcoming / early view. Add volume/issue once final.

Bekamiri, H., Hain, D. S., and Jurowetzki, R. (2024). Patentsberta: A deep NLP based hybrid model for patent distance and classification using augmented SBERT. *Technological Forecasting and Social Change*, 206:123536.

Bena, J. and Li, K. (2014). Corporate innovations and mergers and acquisitions. *Journal of Finance*, 69(5):1923–1960.

Cloodt, M., Hagedoorn, J., and van Kranenburg, H. (2006). Mergers and acquisitions: Their effect on the innovative performance of companies in high-tech industries. *Research Policy*, 35(5):642–654.

Dyer, T. A., Lang, M., and Shroff, N. (2024). The effect of patent disclosure quality on innovation. *Research Policy*. In press. Add vol/issue/pages when available.

Farré-Mensa, J., Hegde, D., and Ljungqvist, A. (2019). What is a patent worth? evidence from the u.s. patent "lottery". *American Economic Review*, 109(9):3038–3073.

Frakes, M. D. and Wasserman, M. F. (2014). Is the time allocated to review patent applications inducing examiners to grant invalid patents? NBER Working Paper 20337, National Bureau of Economic Research.

Frakes, M. D. and Wasserman, M. F. (2020). Procrastination at the patent office? *Journal of Public Economics*, 185:104–122. Final details: fill in exact pages if you have them.

Galasso, A. and Schankerman, M. (2018). Patent rights, innovation, and firm exit. *Journal of Law and Economics*, 61(3):545–573.

Gaulé, P. (2014). Patents and the success of venture-capital-backed companies. Working Paper 546, CERGE-EI.

Healy, P. M., Palepu, K. G., and Ruback, R. (1992). Does corporate performance improve after mergers? *Journal of Financial Economics*, 31(2):135–175.

Hoberg, G. and Phillips, G. (2010). Product market synergies and competition in mergers and acquisitions: A text-based analysis. *Review of Financial Studies*, 23(10):3773–3811.

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.

Hoberg, G. and Phillips, G. (2025). Conglomerate firms, product scope, and the organization of production. *Journal of Finance*. Forthcoming.

Hoberg, G., Phillips, G., and Prabhala, N. (2014). Product market threats, payouts, and financial flexibility. *Journal of Finance*, 69(1):293–324.

Klette, T. and Kortum, S. (2004). Innovating firms and aggregate innovation. *Journal of Political Economy*, 112(5):986–1018.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017a). Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics*, 132(2):665–712.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017b). Technological innovation, resource allocation, and growth (extended data). GitHub repository. Available at https://github.com/KPSS2017/Technological-Innovation-Resource-Allocation-and-Growth-Extended-Data.

Lee, J.-S. and Hsiang, J. (2020). Patentbert: Patent classification with fine-tuning a pre-trained bert model. *World Patent Information*, 62:101965.

Loughran, T. and Vijh, A. M. (1997). Do long-term shareholders benefit from corporate acquisitions? *Journal of Finance*, 52(5):1765–1790.

Sampat, B. N. and Williams, H. (2019). How do patents affect follow-on innovation? evidence from the human genome. *American Economic Review*, 109(1):203–236.

Seru, A. (2014). Firm boundaries matter: Evidence from conglomerates and r&d activity. *Journal of Financial Economics*, 111(2):381–405.

Tur, E. M., D'Este, P., and Cascón-Katchadourian, J. (2025). Repeated examiner–attorney interaction and patent approval. *Scientometrics*. Advance online publication.

Appendix

Table 7 summarizes the variables I have used for the regressions:

Table 7: Summary of Variables in the Firm–Year Panel

Variable	**Definition and Construction**
Identifiers and Timing	
`permco`, `gvkey`	Firm identifiers (CRSP and Compustat linkage)
`filing_year`	Patent application filing year; used as panel year index
`cal_year`	Calendar-aligned fiscal year (based on `fyr`)
Patent and Innovation Measures (KPSS)	
`n_patents_window`	Number of granted patents in a 10-year rolling window
`centroid_1024`, `radius_cos`	Patent embedding centroid and portfolio dispersion (technological breadth)
`cites10_sum`, `cites10_mean`	Forward citations within 10 years of grant (sum and mean)
`xi10_sum`, `xi10_mean`	Innovation impact index (sum and mean)
`cpc10_unique`	Count of unique CPC subclasses in 10-year window
Examiner-Based Instrument	
`Z`	Examiner leniency (leave-one-out grant rate) averaged at firm–year level
`Z_resid`	Residualized leniency (within art-unit $\times$ year)
`n_grants`	Number of granted applications used in instrument computation
Financial Variables (Compustat)	
`at`, `sale`, `ib`	Total assets, net sales, and income before extraordinary items
`roa`	Return on assets: ROA $= \frac{IB}{AT}$
`lev`	Leverage ratio: LEV $= \frac{DLTT+DLC}{AT}$
`cash_at`	Cash holdings ratio: $\frac{CHE}{AT}$
`xrd_at`	R&D intensity: $\frac{XRD}{AT}$
`capx_at`	Investment intensity: $\frac{CAPX}{AT}$
`wc_at`	Working capital ratio: $\frac{WCAP}{AT}$
`q`	Tobin's $Q = \frac{MKVALT+PSTKRV-CEQ+DT}{AT}$
Dynamic and Log-Transformed Measures	
`at_lag`, `sale_lag`	One-year lagged values of total assets and sales
`asset_g`, `sales_g`	Growth rates: $\frac{X_t-X_{t-1}}{X_{t-1}}$ for assets and sales
`ln_at`, `ln_sale`, `ln_emp`	Logarithms of firm size, sales, and employment
Deal Indicators (CRSP/SDC)	
`deal_acq`, `deal_tgt`	Binary indicators for acquirer and target firm–years
`linkdt`, `linkenddt`, `link_span_days`	Link coverage period between CRSP and Compustat