Local Spillovers of Startup Innovation

June 6, 2026

Abstract

When startups innovate, does the surrounding area prosper? Exploiting quasi-random assignment of patent examiners with varying leniency, we show that patent grants to startups increase nearby annualized real estate returns by 0.4 to 0.6 percentage points for residential properties and 0.5 to 0.9 percentage points for commercial properties and vacant land. Effects are largest for the first successful startup in an area and decline with additional successes. We also find that patent-granted startups spur nearby patenting, establishments, and jobs. Yet these real effects emerge only over the following decade, while real estate return premia appear immediately.

Keywords: innovation, real estate, local effects, agglomeration.

1 Introduction

Cities and states spend roughly $100 billion a year competing to attract large firms (Kline and Moretti, 2014), betting that such a firm will raise local incomes and household wealth. A large literature finds that when one opens a plant, relocates its headquarters, or shuts down, nearby employment, wages, and property values all move with it (Greenstone, Hornbeck and Moretti, 2010; Bernstein, Colonnelli, Giroud and Iverson, 2019; Butler, Fauver and Spyridopoulos, 2019; Hu, Tsang and Wan, 2025; Pope and Pope, 2015). In every one of these settings, however, the actor is a large, established firm, and the event a single, salient episode.

Young startups are different. A single one employs only a handful of people, so its direct footprint might be faint. But its success can still move the surrounding economy (Haltiwanger, Jarmin and Miranda, 2013; Akcigit and Kerr, 2018; Walsh, 2025), and when it does, it does so less through the firm's own scale than through what that success signals, seeding the entry of other firms and the agglomeration that follows. The local consequences of a startup's success are, in this sense, a claim on the future. That forward-looking quality is also what makes the effect hard to identify. The promising future a startup's success seems to point to is often what drew the startup to the area to begin with, so its success is entangled with the local conditions that preceded it.

We measure startup success by patenting outcomes and ask how that success affects the local economy. Innovation is a principal channel through which firms shape the economy around them, both directly (Moretti, 2021) and by building a local cluster (Chatterji, Glaeser and Kerr, 2014), so a startup's success in innovating is a natural place to look for local spillovers. The patent-examination process also offers a credible way around the selection problem. We compare local outcomes around startups whose first patent application is granted with those around startups whose application is rejected, instrumenting approval with the leniency of the quasi-randomly assigned examiner (Sampat and Williams, 2019; Farre-Mensa, Hegde and Ljungqvist, 2020).

We find that startup innovation generates real local spillovers. A single patent-granted startup raises annualized real estate returns within 10 miles by 0.4 to 0.9 percentage points across property types. It also spurs nearby patenting and increases establishments and employment in the startup's own industry by roughly 9 to 18 percent. The real estate premium emerges immediately after the examiner's first action on the application, typically years before the patent itself is issued, whereas the local innovation and business activity it anticipates build only gradually over the following decade, consistent with the market capitalizing the signal well before its real effects materialize.

Our analysis draws on transaction-level and establishment-level micro-data assembled nationwide around every startup, with each record carrying an exact location and date. This granularity lets us compare properties and establishments in the same county that lie closer to or farther from a startup, and observe them just before and after its patent decision. We can therefore measure how local real estate and business activity respond to innovation in fine spatial and temporal detail.

Our first set of results shows spillovers in nearby real estate markets. A patent grant raises annualized returns within 10 miles by about 0.4 percentage points for single-family homes, 0.6 for multi-family residential, 0.5 for commercial, and 0.9 for vacant land. These premia are driven by both faster price growth and shorter holding periods in treated areas. Controlling for the holding period roughly halves the premia but leaves a positive, statistically significant residual that reflects price growth. The premia decay with distance, with most of the decline inside 5 miles. They are absent in placebo samples of repeat sales completed entirely before the patent decision, so pre-existing local trends do not drive them. They are also absent among repeat sales completed entirely afterward, indicating that they accrue to investors who owned property before that decision rather than to later buyers.

Do these premia reflect expectations of future neighborhood change, or mere speculation? We find that the activity they anticipate does materialize. A patent grant raises nearby patenting over the following 3 years and, over the next decade, draws in business activity concentrated in the startup's own industry. By the tenth year, the area around the average successful startup gains roughly 310 establishments and 2,570 jobs, with little response in other industries. The real activity thus arrives, but only gradually.

The expectations channel also appears in how the premia accumulate across innovators. Exploiting variation in the number of nearby innovators, we find that the spillovers are concave. The premium rises with the number of nearby grants, but each additional one contributes less. For single-family homes, for example, the per-startup premium falls from about 0.75 percentage points where only a few grants are nearby to about 0.17 where many are, so the first successful startup in an area generates the largest spillover. We interpret it as the first innovator both signaling most strongly that a neighborhood is on an upward trajectory and seeding the most new activity, with later ones mostly confirming and building on what the first set in motion.[1]

This paper contributes to several strands of literature. A large body of work shows that big firms reshape their surroundings. New manufacturing plants and government-led industrial expansions raise local productivity, wages, and employment, and corporate relocations and big-box entry move nearby housing and business outcomes, although the effects can be heterogeneous or negative.[2] A parallel literature places new firms and local knowledge spillovers at the center of employment growth and urban dynamism, and builds macro and spatial models for how shocks to startup activity shape output and regional trajectories.[3] Both literatures, however, study either large established firms or the formation of new startups. We shift the focus to the success of small,

[1]The return concavity is what a concave agglomeration technology can produce, with each additional nearby success adding less real activity as finite local capacity fills. Our conceptual framework in Appendix D formalizes it in Proposition 1. This technology and the expectations channel are observationally close here and not mutually exclusive, and we do not attempt to separate them.

[2]On manufacturing plants and industrial expansions, see Greenstone et al. (2010); Garin and Rothbaum (2025). On corporate relocations and big-box retail entry, see Hu et al. (2025); Pope and Pope (2015). Qian and Tan (2021) find that high-skilled firm entry benefits high-skilled homeowners but harms nearby low-skilled renters, and Bernstein et al. (2019) find that liquidation, relative to reorganization, produces highly localized negative spillovers, particularly in services.

[3]On knowledge spillovers, young firms, and local employment growth, see Glaeser, Kallal, Scheinkman and Shleifer (1992); Haltiwanger et al. (2013); Adelino, Ma and Robinson (2017), and Glaeser, Kerr and Kerr (2015) for reduced-form evidence linking entrepreneurship to urban growth. For macro and spatial models of how firm and startup shocks propagate, see Monte, Redding and Rossi-Hansberg (2018); Gourio, Messer and Siemer (2016); Carlino and Drautzburg (2020); Walsh (2025).

early-stage innovators, and provide causal evidence that a single startup's innovation generates localized spillovers in both real estate returns and business activity.

It also contributes to the economics of innovation, which asks what a granted patent does for the firm that earns it and for other firms. Studies that exploit examiner leniency show that approval raises the patenting firm's own growth and later innovation (Farre-Mensa et al., 2020), and a related strand traces how its patenting shapes the innovation and investment of others.[4] A spatial strand instead studies how labor and capital markets sort innovative activity across places and how knowledge spills over with proximity.[5] We add micro-level causal evidence that connects the two, showing that a startup's success not only raises nearby patenting but draws in real business activity, mostly within its own industry, linking the innovation it produces to the agglomeration that follows.

Finally, the paper contributes to the literature on how local economic activity is capitalized into housing markets. Prior studies show that house prices reflect expectations about regional growth and a range of local shocks, and that prices correlate with metro-level innovation.[6] We provide causal evidence that startup patent approvals generate return premia across all major property types.

The rest of the paper proceeds as follows. Section 2 describes the patent examination process and our empirical setting. Section 3 introduces the three datasets and the construction of our key variables. Section 4 presents descriptive evidence, develops the instrumental-variable strategy,

[4] Hegde, Ljungqvist and Raj (2022) show that patent grant delays slow a startup and harm its product-market rivals, and Matray (2021) documents local knowledge spillovers, in which innovation by one firm fosters innovation by neighboring firms. Grieser and Liu (2019) examine how competitor financial constraints shape firm investment, and Kogan, Papanikolaou, Seru and Stoffman (2017) tie the market value of patented innovation to medium-run growth. Ma (2025) documents that technological obsolescence is followed by declines in growth, productivity, and capital reallocation.

[5] Derrien, Kecskés and Nguyen (2023) find that firms in younger labor markets are more innovative, and Hombert and Matray (2017) trace a spatial sorting of young inventors out of areas hit by credit-supply shocks. Cornaggia, Mao, Tian and Wolfe (2015) link banking competition to small-firm innovation, and Jaffe, Trajtenberg and Henderson (1993) show that patent citations are geographically localized.

[6] See Coulson, Liu and Villupuram (2013) on how expectations about metropolitan growth enter house prices. Butler et al. (2019) document local spillovers from stock market listings, and Hartman-Glaser, Thibodeau and Yoshida (2023) link IPO wealth to higher house prices near firm headquarters. Beracha, Kim, Wintoki and Xi (2025) and Aiello, Kotter and Schubert (2022) study high-skilled migration and movers carrying housing-equity gains, and Beracha, He, Wintoki and Xi (2022) relate metro-level innovation to house prices.

and reports the causal estimates of startup patent approvals on nearby real estate returns and local business activity. Section 5 concludes.

2 Institutional Background and Empirical Setting

Our empirical strategy rests on two features of the U.S. patent examination process. Applications are assigned to examiners in a way that is as good as random, and examiners differ widely in their leniency. This section describes the assignment mechanism, explains why the resulting variation in patent approval is plausibly exogenous, and lays out our empirical setting.

2.1 The Patent Application and Examiner Assignment Process

When an application is filed with the U.S. Patent and Trademark Office (USPTO), it undergoes a preliminary review for completeness and is classified by technology area. It is then forwarded to the appropriate Technology Center and assigned to an *art unit* specializing in that field. Within the art unit, applications enter an examiner's queue ordered by filing date, and assignment to a particular examiner is as good as random conditional on the art unit and filing time (Farre-Mensa et al., 2020). This assignment is therefore orthogonal to the quality of the invention and the characteristics of the applicant. Because examiners differ substantially in leniency, the random assignment produces plausibly exogenous variation in whether an application is approved. This variation is the foundation of our empirical strategy, which we formalize and validate in Section 4.2.

After taking up an application, the examiner assesses whether it satisfies the legal requirements for patentability, namely that the claims are directed to patent-eligible subject matter, that the specification adequately describes and enables the invention, and that the claims distinctly define it. The examiner conducts a prior-art search to assess novelty and non-obviousness and issues an initial decision known as the *first action on the merits*. Most first actions are non-final rejections that outline the examiner's objections. Only about 10 to 15 percent of applications receive a first-action allowance with no rejections (Carley, Hegde and Marco, 2015). The USPTO notifies the applicant by official letter, at which point the applicant also learns the examiner's identity (Kogan et al., 2017).

The applicant typically responds by amending claims or presenting arguments to address the examiner's concerns. Prosecution may involve multiple rounds. After each response, the examiner can either allow the application or issue a final rejection if issues remain. A final rejection does not end the process. The applicant may abandon the application, file a Request for Continued Examination to reopen prosecution, or appeal to the Patent Trial and Appeal Board. The process concludes only when the application is either granted or permanently abandoned. The USPTO never conclusively "rejects" an application. Any rejection can in principle be overcome or appealed, so a "rejected" application is one the applicant elects to abandon after an adverse decision (Kogan et al., 2017).[7]

Patent examination is lengthy. In our sample, an application waits about eight months on average to be taken up by an examiner and roughly another year for the first-action decision. The final disposition (a grant or an abandonment) typically follows several more rounds, about three years after filing on average. The first action therefore provides an early but informative signal about an application's prospects, long before the ultimate outcome is decided. Following the prior literature (Kogan et al., 2017; Kline, Petkova, Williams and Zidar, 2019; Farre-Mensa et al., 2020), we anchor the timing of treatment to the first-action decision, and we define the treatment indicator by the application's ultimate disposition, equal to one if the application is granted and zero if it is rejected.

[7] We refer to such abandoned applications as rejections throughout the paper.

2.2 Empirical Setting

Our empirical setting consists of U.S. startups that file their first patent application between 2001 and 2009 and receive a first-action decision by 2009, with a final disposition by 2013. Several features make this setting well suited to identifying the local spillovers of startup innovation.

First, as discussed above, the random assignment of applications to examiners within art units, combined with persistent differences in examiner leniency, generates plausibly exogenous variation in whether a given startup obtains a patent. This variation lets us compare the areas around startups that are assigned more lenient examiners, and therefore obtain patents at higher rates, with the areas around otherwise similar startups assigned stricter examiners, isolating the causal effect of successful innovation on nearby economic outcomes.

Second, the American Inventors Protection Act of 1999 required publication of patent applications filed on or after November 2000, so for applications filed in 2001 and later the applicant's identity is observable regardless of the eventual outcome (Carley et al., 2015; Farre-Mensa et al., 2020). This disclosure rule lets us observe both granted and rejected applicants, which is essential for a design that compares treated and control startups defined by the patent decision itself.

Third, startups are observable in the patent record well before they appear in conventional firm-level datasets. Most firms in our sample are young at filing. The median startup is two years old and employs eight people, and many have no presence in standard administrative data. Patent filings therefore provide a window onto the early-stage innovative activity that drives much of the local economic growth documented in the agglomeration literature (Glaeser et al., 1992; Carlino and Kerr, 2015).

Fourth, restricting treatment to first-action decisions made by 2009 leaves a long subsequent window in which to observe local economic outcomes. This window lets us measure how successful innovation affects nearby real estate and economic activity over a long horizon, rather than only in the years immediately following the patent decision.

Together, these features make the patent decision a credible source of variation in the success of startup innovation, whose local effects we can trace over a long horizon. We detail the construction of the startup sample in Section 3.1.

3 Data

Our empirical analysis draws on three datasets. The USPTO Patent Examination Research Dataset identifies startups and the examiners assigned to their applications. ATTOM property transaction records measure local real estate outcomes. The Data Axle business registry measures local economic activity. This section describes each in turn, along with the key variables we construct from them.

3.1 Patent Data and Startups

Our patent data come from the USPTO's Patent Examination Research Dataset (PatEx), which contains detailed records of patent applications and their prosecution histories. PatEx is derived from the USPTO's Patent Application Information Retrieval (PAIR) system and provides research-ready information on application filings, examiner actions, and outcomes, covering both unpublished and abandoned applications.

Following Farre-Mensa et al. (2020), we restrict the sample to first-time patent applicants that are private, incorporated, U.S.-based businesses. We exclude not-for-profit organizations such as universities and government labs, along with any applicant that is publicly traded or a subsidiary of an existing firm. We further require "small entity" status under the U.S. Small Business Act, that is, below the employee-size threshold for reduced filing fees. To capture genuinely new ventures, we require that each firm's first patent application be filed on or after January 1, 2001, and that the firm have no patenting activity in the prior 25 years. The resulting sample contains 29,487 startups across 8,504 ZIP codes, spanning all 50 states.

For each startup, we identify the firm's address at the time of the patent application from the assignee address listed in the USPTO assignment records. We geocode each address with the Bing Maps API and retain only addresses that correspond to the firm's headquarters or primary place of business rather than the inventor's personal address.[8] Appendix Figure A.1 maps the resulting geographic distribution of startups, with the largest counts in well-known innovation hubs such as Santa Clara, Los Angeles, San Diego, and Orange County. We use original (non-continuation) utility patent applications that these startups filed between 2001 and 2009.

One limitation of PatEx is that it does not record industry classifications for patent applications. Because we estimate within-industry spillovers, we would ideally observe the industry of each startup and each nearby establishment. In practice, however, we observe only technology classifications. Each application is assigned to an art unit and one or more classes in the U.S. Patent Classification (USPC) system. To bridge this gap, we use the probabilistic concordance developed by Goldschlag, Lybbert and Zolas (2020), which links USPC classes to 6-digit NAICS industries by matching patent abstracts to industry descriptions, and we assign each startup to the sector with the highest implied match probability.

Examiner leniency instrument.

A key variable for our empirical strategy is the leniency of the patent examiner assigned to each startup's first application. We define the leniency of examiner j in art unit a at time t as the share

[8] A USPTO patent application has three key addresses on file: the firm's address (from the USPTO Patent Assignment database), the inventor's residence (from the Application Data Sheet), and the correspondence address (also from the Application Data Sheet, typically the patent attorney's or agent's office). We use the firm's address from the assignment event closest in time to the patent's first-action date.

of applications that examiner j has approved among all applications the examiner has examined in that art unit before startup i's application:

$$\text{Examiner Leniency}_{j,a,t} = \frac{n_{\text{granted}, j,a,t}}{n_{\text{reviewed}, j,a,t}}, \tag{1}$$

where $n_{\text{granted}, j,a,t}$ and $n_{\text{reviewed}, j,a,t}$ are the numbers of applications examiner j has granted and reviewed, respectively, in art unit a up to time t (the filing date of startup i's application). We compute this rate from the examiner's prior examination history in the PatEx data, drawing on the full universe of patent filings rather than startup applications alone, and we exclude observations in which the assigned examiner has reviewed fewer than 10 applications before the focal one. Because the calculation excludes the focal application itself, the instrument is free of any mechanical correlation with the startup's outcome. Sections 4.2 and 4.3 use this instrument in the two-stage least squares design.

Innovation spillovers measure.

We also use the patent data to construct a measure of innovation spillovers to nearby startups. For each focal startup application, we count the number of applications filed by other startups within one-mile distance bands out to 10 miles, over the 3-year window following the focal application's first-action date. Because our sample consists of startups with first-action dates between 2001 and 2009, we restrict the spillover analysis to focal applications with first-action dates through 2006, so that a full 3-year post-decision window is observed for every application. We measure spillovers as the change in nearby counts relative to the year before the focal application's first-action date, so the spillover sample begins with first-action dates in 2002.

Appendix Table A.2 reports summary statistics for this sample. Panel A summarizes the change in the number of nearby applications over the 3-year post-decision window, relative to the year before the first-action date. Pooling across all art units, the number of nearby applications rises by 35 on average (standard deviation 54); within the focal patent's own industry sector, it rises by 11 (standard deviation 22). Panel B reports the analogous statistics for successful applications (approvals). Nearby approvals rise by 20 on average across all art units (standard deviation 33) and by 6 within the same sector (standard deviation 13). The large standard deviations relative to the means indicate substantial variation in spillovers across focal patents.

3.2 ATTOM Property Tax and Deed Data

We use property transaction data from ATTOM's deed records and tax assessor files to track real estate market outcomes around each startup. The data provide property-level information on transaction dates and prices, buyer and seller names, and property characteristics. For each transaction, we observe the exact property address and detailed use codes, which we map into four broad property categories: single-family residential, multi-family residential, commercial, and vacant land.[9] From the tax records, we extract property characteristics (lot size for all property types; building square footage and year built additionally for commercial, single-family, and multi-family properties; and number of bedrooms and bathrooms additionally for single-family and multi-family properties), which we use as sample filters and as controls in the property-level analysis in Section 4.2.

We compile a repeat-sales dataset of properties that transact at least twice. Each repeat-sale observation consists of a purchase (buy) and a subsequent sale (sell) of the same property. For every such pair, we compute the annualized unlevered return between the two transactions,

$$R = \left(\frac{P_{\text{sell}}}{P_{\text{buy}}}\right)^{1/T} - 1, \tag{2}$$

where P_{buy} and P_{sell} are the purchase and sale prices of the property and T is the holding period in years. All transaction prices are inflation-adjusted to 2010 dollars using the Consumer Price Index.

[9] See Appendix Table A.3 for the mapping of detailed property use codes into these four categories.

This annualized unlevered return R is our key outcome variable, reflecting both price appreciation and the holding period over which the return accrues. We report results in Section 4 from specifications both with and without controls for the number of holding years interacted with county fixed effects. The specifications without these controls capture the combined effect of price growth and the length of the holding period, whereas those with the controls isolate price growth.

We link each repeat sale to nearby patent-applying startups using the geocoded coordinates of both the property and the startup's address described in Section 3.1. For each repeat-sale transaction, we identify all startups in our sample located within 10 miles of the property and compute the distance between each property-startup pair. Our baseline definition of a "local" transaction is one within 10 miles of an innovating startup, which provides a broad local real estate market around each firm and roughly corresponds to typical pre-pandemic commuting distances in the U.S.[10] In Section 4.2, we also report results by one-mile distance bin and discuss results extending up to 15 miles.

Estimation Samples.

We apply two filters to the repeat-sales data before aggregation. First, we drop observations with outlier values in buy price, sell price, or annualized return, defined using the interquartile range rule calculated at the county-by-property-sub-category level. Second, we link each repeat sale to the property's tax assessment history to detect major property improvements. For non-vacant properties, we compare the property's effective year built recorded after the second transaction to the timing of the first transaction, and we exclude any repeat sale in which the first purchase date precedes a recorded major renovation or rebuild.

We then construct two aggregations of the repeat-sales data, each tailored to a different empirical exercise and each chosen to keep estimation computationally tractable given the large number

[10] According to the Census Bureau's American Community Survey, the average one-way commute in 2006 was about 25 minutes. Proprietary data from Gusto suggest that the mean commute distance had risen to approximately 10 miles by 2019.

of startups and nearby transactions. The *firm-level sample* treats each startup as the unit of analysis and asks how its patent decision moves nearby returns, whereas the *property-level sample* treats each property as the unit and asks how its returns respond to the full set of innovation events around it.

The *firm-level sample* groups transactions into cells indexed by startup, property sub-category, county, buy year, sell year, and distance band (in miles from the startup). Within each cell, we compute the median unlevered return R, purchase price, and sale price, and in regressions we weight observations by the number of underlying transactions. We further partition the firm-level sample into three subsamples by the timing of transactions relative to the associated startup's first-action date. The *baseline* sample contains transactions whose buy date precedes and whose sell date follows the first-action date. The *post* sample contains transactions whose buy and sell dates both fall after it, and the *placebo* sample contains transactions whose buy and sell dates both fall before it. Appendix Table A.4 reports summary statistics for the firm-level sample.

The *property-level sample* uses an alternative aggregation. Instead of collapsing transactions around each startup, it aggregates patent outcomes around each property. We retain individual repeat-sale transactions, and for each one we count the patent-applying startups located within 10 miles whose first-action decisions fall between the property's buy and sell dates, classifying these decisions by outcome to record separately the number of approvals and the number of rejections. This construction lets us evaluate how the return premium at a property responds to the cumulative arrival of nearby innovation events, distinguishing properties exposed to many successful patent grants from those exposed to many unsuccessful applications or to none at all. The property-level sample inherits the same baseline, post, and placebo timing distinctions. Appendix Table A.5 reports summary statistics for the property-level sample.

3.3 Data Axle Business Registry

We measure local economic activity around each startup using Data Axle (formerly infoUSA), an establishment-level business registry available annually from 1997 onward that compiles information on nearly all U.S. businesses from telephone directories, business filings, surveys, and web mining. For each establishment, the data reports its name, address, industry (NAICS and SIC), and employment. Relative to the Census Bureau's ZIP Code Business Patterns, Data Axle's establishment counts closely track the official totals while offering exact addresses and annual frequency. Appendix C reports this comparison.

The raw registry, however, sometimes overcounts establishments by listing a single business as several records. For example, each attorney in a law firm can appear as a separate "business." We collapse such duplicate listings into a single establishment, treating records as the same business when they share either Data Axle's business identifier or the same company name and 5-digit ZIP code, and we set the establishment's employment to the median across its constituent records.[11] We drop records that list only a P.O. Box address and exclude establishments that Data Axle leaves unclassified by industry.

For each startup, we then build an annual panel of nearby businesses. The panel includes every establishment within 10 miles of the startup's geocoded address, grouped by 2-digit NAICS sector and by year. The panel gives two measures of nearby activity in each sector and year, the *number of establishments* and their *total employment* (the combined number of employees across those establishments). Appendix Tables A.6 and A.7 report summary statistics for the number of nearby establishments and for total nearby employment, respectively. For each, we report the level within 10 miles in the year before the first-action decision and its change over the following decade.

[11] We link records through connected components. We treat each business identifier and each distinct company-name-by-ZIP-code combination as a node, place an edge between a business identifier and a name-by-ZIP combination whenever the two appear together on a record, and assign all records in a connected component to a single establishment. Records that match on neither key retain their own identifier.

4 Empirical Analysis

This section presents our main empirical results. We first document, descriptively, how local real estate markets and business activity evolve around each startup's patent application's first-action (FA) decision date. We then develop our instrumental variable strategy, instrumenting a startup's patent approval with the leniency of its assigned examiner, and apply it first to local real estate returns and then to local innovation and business activity.

4.1 Descriptive Evidence

Patent grants and local real estate outcomes.

We begin by describing how local real estate markets evolve around the patent decision, using the firm-level repeat-sales sample from Section 3.2. Figure 1a plots the time series of annualized unlevered returns for nearby repeat-sales transactions, separately for startups whose first application was eventually granted (*treated*) and rejected (*control*), relative to the associated firm's FA date. These returns are adjusted for calendar-year effects.[12] This characterizes the path of real estate returns in treated versus control neighborhoods, adjusting for the macro real estate market trend, from five years before to ten years after the decision.

A common pattern emerges around the FA decision across property types. Returns near patent-granted startups are flat to modestly declining before treatment and turn up sharply at the treatment

[12] We construct the calendar-year adjustment as follows. For each property type and each calendar year in which sales occur, we compute the transaction-weighted average annualized return across all repeat sales of that type sold in that year, pooling treated and control observations. We subtract this calendar-year average from each cell's return and add back the property type's overall mean, so that the adjusted return measures how a sale's return compares with the prevailing market return for the same property type in the same year of sale, while remaining on the original return scale. We then average these adjusted returns by years since the FA date, separately for treated and control startups. This adjustment helps characterize the trend in returns after accounting for cyclical patterns in real estate markets, especially because part of our post-treatment period coincides with the financial crisis. Netting out the calendar-year average removes this macro movement and isolates the gap between treated and control areas with transactions in the same year. The unadjusted return series is reported in Appendix Figure A.3a.

date, rising above or closer to returns near rejected startups over the first several years afterward. This turning point is clearest for single- and multi-family properties, present but weaker for commercial properties, and noisier for vacant land. The differences are small in absolute terms, however, on the order of a few tenths of a percentage point. Similar patterns hold when excluding the financial crisis years, 2008 and 2009, or when using a smaller bandwidth, 5 miles instead of 10 miles, as presented in Appendix Figures A.3b and A.3c.

Returns near patent-rejected startups follow a different path from those near patent-granted startups both before and after the patent decision. Before the decision, control returns tend to move in the opposite direction from treated returns, and afterward they remain largely flat, with vacant control properties showing a noisier path throughout. Control returns are also higher in level, by roughly 0.1 to 0.2 percentage points, outside the few years immediately after the decision when the treated series rises above them for residential properties.

Taken together, the descriptive return evidence points in two directions. Treated areas show a spike in returns immediately after the decision that, although small in magnitude, suggests a potential positive spillover from innovative startups. At the same time, the level and trend differences that appear before the FA date suggest that treated and control startups locate in systematically different neighborhoods. We sharpen this raw comparison with the weighted least squares regressions that follow. Specifically, we estimate

$$y_{i,k,c,t_0,d} = \beta_0 + \beta_1\,\mathbf{1}(\text{Approved})_i + \lambda_{gau_i \times appl.\ year_i} + \lambda_k + \lambda_{c \times t_0} + \lambda_d + \varepsilon_{i,k,c,t_0,d}, \tag{3}$$

where the unit of observation is startup i, disaggregated by property type k, county c, initial purchase year t_0, and distance bin d (in miles from the startup). The indicator $\mathbf{1}(\text{Approved})_i$ equals one if the firm received a patent grant, our treatment of interest, and zero for firms whose first application was rejected, which serve as the comparison group. The outcome y is the median unlevered return, and we also examine the log median buy price, the log median sell price, and the holding term, the three components that determine the unlevered return. We include fixed effects for art unit by application year ($gau_i \times appl.\ year_i$), property type, county by purchase year, and distance bin, and weight the regressions by the number of repeat-sale transactions in each i-k-c-t_0-d cell. The coefficient β_1 therefore compares granted and rejected startups that are otherwise alike along these dimensions.

Figure 1b reports the estimated β_1 and its 95% confidence interval for each outcome, sample, and property use category. In the baseline sample, which collects repeat sales that straddle the patent decision, the estimated β_1 for the unlevered return is positive for all four property types, with confidence intervals above zero. Areas near granted startups therefore show higher annualized returns than areas near rejected startups. The magnitudes are modest, with point estimates of about 0.12 to 0.16 percentage points for residential and commercial properties and about 0.4 percentage points for vacant land. The placebo coefficients, estimated on transactions completed before the decision, are near zero, either statistically indistinguishable from zero at the 95% level or slightly negative, so before the decision returns are comparable across the two groups of areas, or marginally lower near patent-granted startups. The post coefficients are smaller than in the baseline and, for most property types, indistinguishable from zero. That is, the return differential is concentrated among repeat sales that span the decision rather than those transacted entirely beforehand or afterward.

The remaining panels in Figure 1b decompose this return differential into its three components, the holding term, the sell price, and the buy price. In the baseline sample, median sell prices are higher near granted startups for every property type and the holding term is shorter, the two components that account for the higher annualized return. Median buy prices are also somewhat higher, but by less than sell prices and significantly so only for single-family homes. Together the components describe a wider gap between sale and purchase prices realized over a shorter holding period near patent-granted startups.

In the placebo sample, sell prices show no differential and buy prices are only slightly higher, leaving the return near zero, whereas the holding term is already shorter, though by less than in the baseline. The shorter holding is thus present even before the decision, and it is the sell-price gap, which appears only around it, that sets the baseline apart. In the post sample, both price differentials are close to zero and the holding term remains shorter, by a smaller margin than in the baseline, so turnover near patent-granted startups stays modestly faster among transactions completed entirely afterward but, absent a price gap, no longer produces a higher annualized return.

These estimates are descriptive rather than causal, and they align with the immediate spike in returns near patent-granted startups after the patent decision in Figure 1a. The level and trend differences in the raw pre-decision series are largely absorbed once the fixed effects are added, leaving the placebo return coefficients near zero. Differences in prices and holding terms, however, persist even after this conditioning, so granted and rejected startups still differ in ways the fixed effects do not capture. We address them with the quasi-experimental design developed in Section 4.2.

Patent grants and local business activity.

Local business activity tracks the same pattern as local real estate. Areas near granted startups begin with fewer establishments and lower employment than areas near rejected startups, yet grow faster once the patent decision is revealed. We describe this pattern using the establishment- and employment-level Data Axle sample from Section 3.3. Figure 2a plots the average number of nearby establishments and the level of nearby employment by year relative to the FA date, separately for startups whose first application was eventually granted (treated) and rejected (control), in levels (top row) and as the change from the pre-FA baseline (bottom row). A common pattern emerges. The level panels show a persistent gap, with granted-startup areas sitting below rejected-startup areas throughout, consistent with granted startups locating in somewhat less dense areas. The change panels show similar pre-trends, and after the FA date the employment series near granted startups pulls above the control, whereas the two establishment series stay closer.

We sharpen this raw comparison with the regressions that follow. Specifically, for the Data Axle sample we estimate

$$y_{i,k,c,d,t} = \beta_0 + \beta_1 \mathbf{1}(\text{Approved})_i + \lambda_{gau_i \times appl.\ year_i} + \lambda_k + \lambda_{c \times appl.\ year_i} + \lambda_d + \varepsilon_{i,k,c,d,t}, \tag{4}$$

where $y_{i,k,c,d,t}$ is the change, or the percent change, in nearby establishments or jobs for startup i in sector k, county c, and distance ring d, measured t years from its FA date, and $\mathbf{1}(\text{Approved})_i$ equals one for granted startups. We include fixed effects for art unit by application year, NAICS sector, county by application year, and distance ring, and cluster standard errors at the art-unit level. As in the real-estate specification, β_1 compares granted and rejected startups that are otherwise alike along these dimensions, and the estimates are descriptive rather than causal.

Figure 2b reports the estimated β_1 and its 95 percent confidence interval for the change and the percent change in establishments and in employment, each at a placebo horizon (3 years before the decision) and at 10 years. Every placebo coefficient is small and statistically indistinguishable from zero, so the two groups evolve in parallel beforehand. At 10 years, the coefficients are positive and significant throughout. The change near granted startups is about 0.5 more establishments and 6.5 more jobs. The effects are modest in absolute terms but amount to proportional increases of about 1 percent in establishments and 3 percent in employment.

Taken together, the descriptive evidence points in two directions, as with real estate. The post-FA divergence, with employment near granted startups rising above the control, suggests a positive spillover from successful innovation. At the same time, the persistent level gap indicates that granted and rejected startups locate in systematically different areas. These comparisons remain descriptive, and we turn next to a causal interpretation.

4.2 Effects on Local Real Estate Returns

Instrumental variable strategy.

The descriptive estimates presented earlierf are not causal. Startups choose where to locate, and these choices may correlate with local real estate and economic outcomes. For example, unobservably stronger startups may locate in lower-cost, lower-appreciation areas to economize on rent and labor, or in high-growth areas to attract talent. Either pattern would bias the OLS comparison of granted and rejected startups, since the patent outcome would then be correlated with unobserved determinants of local returns. We address this by instrumenting a startup's patent approval with the leniency of the examiner assigned to its first application, defined in Equation 1. As discussed in Section 2, applications are quasi-randomly assigned to examiners within art units conditional on application timing, so leniency is plausibly orthogonal to the local economic conditions.

We estimate the model by two-stage least squares. In the first stage, we project the approval indicator onto examiner leniency, conditional on the same fixed effects as the second stage. The second-stage equation is

$$y_{i,k,c,t_0,d} = \beta_0 + \beta_1 \mathbf{1}(\widehat{\text{Approved}})_i + \lambda_{gau_i \times appl.\ year_i} + \lambda_k + \lambda_{c \times t_0} + \lambda_d + \varepsilon_{i,k,c,t_0,d}, \tag{5}$$

where $\mathbf{1}(\widehat{\text{Approved}})_i$ is the fitted value from the first stage and the primary outcome $y_{i,k,c,t_0,d}$ is the median annualized unlevered return.[13] Regressions are weighted by the number of repeat-sale transactions in each i-k-c-t_0-d cell, and standard errors are clustered at the art-unit level. The coefficient of interest, β_1, captures the local average treatment effect of a patent grant on nearby real estate returns, identified from variation in approval probabilities induced by examiner assignment. These estimates therefore describe the marginal startups whose approval hinges on examiner as-

[13] We also examine the log median buy price, the log median sell price, and the holding term as alternative outcomes. As a robustness check, we re-estimate the baseline specification using the mean, rather than median, annualized unlevered return and obtain results that are quantitatively similar and remain statistically significant.

signment, not the full population of applicants. A first patent grant also bundles several channels. Prior work shows that approval shapes a startup's own growth, financing, and survival (Farre-Mensa et al., 2020), so the nearby effects we estimate need not stem from innovation alone.

The instrument exhibits substantial variation and a strong first stage. Appendix Figure A.2 plots the distribution of examiner leniency after residualizing on art-unit-by-application-year fixed effects; within a cell, the gap in residual leniency between the 25th- and 75th-percentile examiners is 17.5 percentage points. First-stage coefficients range from 0.65 to 0.68, with standard errors of 0.024 to 0.029, across specifications (Panel A of Appendix Table A.8). Together these imply that a single application, had it been assigned to a 75th- rather than a 25th-percentile examiner within its art unit and application year, would face an approval probability roughly 11.4 to 11.9 percentage points higher.

Innovation and local return premia.

We first examine the baseline sample, which consists of repeat-sale transactions in which the property was purchased before and sold after the FA date of the associated patent application. Panel A of Table 2 reports IV estimates from Equation 5. All specifications include the full set of fixed effects, with standard errors clustered at the art-unit level. As mentioned above, the corresponding first-stage results are reported in Panel A of Appendix Table A.8, with coefficient estimates ranging from 0.65 to 0.67 and cluster-robust Kleibergen–Paap rk first-stage F-statistics between roughly 425 and 774 across subsamples, far above conventional weak-instrument thresholds, such as the Stock–Yogo 10% maximal-size critical value of 16.38.

Columns (1), (3), (5), and (7) of Panel A.1 in Table 2 show that patent-granted startups raise local real estate returns by 0.38 percentage points for single-family homes, 0.64 percentage points for multi-family properties, 0.50 percentage points for commercial real estate, and 0.91 percentage points for vacant land, all significant at the 1 percent level. Return premia are larger for multi-family and commercial properties than for single-family homes, consistent with innovation effects being most pronounced in denser, typically downtown, neighborhoods. Vacant land exhibits the largest premium, consistent with expectations of future development being capitalized strongly into land values, though the estimates are noisy because of sparse vacant land transactions.

Compared to the OLS estimates discussed in Section 4.1, the IV coefficients are larger across all four property categories. The OLS baseline estimates imply return premia of roughly 0.1 to 0.2 percentage points for single-family, multi-family, and commercial properties and about 0.4 percentage points for vacant land, whereas the IV estimates reported here are two to four times larger. This pattern is consistent with endogenous location choice attenuating the OLS estimates. If unobservably stronger startups systematically locate in lower-appreciation areas (for instance, to economize on rent or labor costs, or because they can attract high-quality workers and investment even to less desirable locations), the OLS estimator understates the impact of successful innovation on nearby real estate returns.

We assess robustness to four alternative specifications, reported in Appendix Figure A.4: an unweighted 2SLS estimator (rather than weighting by transaction counts); a sample restricted to properties exposed to only one patent-applying firm within 10 miles between the buy and sell dates ("Excl. overlap"); a more granular ZIP-by-buy-year (and, where applicable, ZIP-by-holding-year) fixed-effects structure ("ZIP level"); and a sample that excludes the 2008–2009 financial crisis. Coefficient estimates remain statistically indistinguishable from the baseline at the 10 percent level across all specifications, with the exception of the commercial and vacant land premia in the ZIP-level analysis, where the granular ZIP-by-buy-year and ZIP-by-holding-year fixed effects absorb most of the variation.

We next decompose these return premia into their three underlying components, log buy price, log sell price, and holding term. We re-estimate Equation 5 with each as the outcome variable. Panels A.2 through A.4 of Table 2 show that the higher returns are driven primarily by lower buy prices and shorter holding periods in treated areas, rather than by higher sell prices. IV estimates for log median sell prices (Panel A.3) are statistically insignificant at the 10 percent level, indicating similar exit prices across treated and control areas. In contrast, the estimates for log median buy prices (Panel A.2) are negative and generally larger in absolute value, indicating that properties in treated areas are purchased at lower prices, which likely reflects differences in quality or size. Estimates for holding term (Panel A.4) are negative and statistically significant, showing that properties in treated areas re-sell more quickly and thereby amplify annualized returns.

To assess the role of the holding-term channel more directly, we augment Equation 5 and its first-stage counterpart with county-by-holding-term (in years) fixed effects and report the resulting IV estimates in Columns (2), (4), (6), and (8) of Panels A.1 through A.3. The return premia in Panel A.1 remain positive and statistically significant, specifically, 0.20 percentage points for single-family homes (1 percent level), 0.39 for multi-family (1 percent), 0.23 for commercial (1 percent), and 0.58 for vacant land (10 percent). The reductions relative to the baseline are broadly comparable across property types, ranging from about 54 percent for commercial properties (from 0.50 to 0.23) to about 36 percent for vacant land (from 0.91 to 0.58). These estimates indicate that the shorter holding term in treated areas accounts for a substantive share of the observed premia but does not fully explain them; the residual reflects economically and statistically significant stronger underlying price growth in treated areas. The price components in Panels A.2 and A.3 are consistent with this. Although neither the buy-price nor the sell-price estimate is individually significant at the 10 percent level, the sell-price differential exceeds the buy-price differential, and this gap underlies the residual return premia.

We next turn to the placebo and post-treatment samples. Panel B.1 of Table 2 reports results for placebo transactions in which both the purchase and sale occurred before the FA date of any associated startup patent application. The IV coefficients are statistically insignificant at the 10 percent level across all four property types. This null result supports the validity of our identification strategy. The baseline return premia are not driven by pre-existing local trends that happen to coincide with later innovation activity, as the instrument is orthogonal to price movements that precede innovation.

Panel B.2 of Table 2 reports results for post-treatment transactions, in which both the purchase and sale occurred after the innovation signal was revealed. The estimates provide little evidence that later buyers capture the return premium. Coefficients are statistically insignificant at the 10 percent level across all property types and remain so when holding-term fixed effects are added in Columns (2), (4), (6), and (8). The sole exception is commercial real estate in Column (6), where the estimate of 0.07 percentage points is significant at the 10 percent level but far smaller than the corresponding baseline estimate of 0.23. Taken together, these patterns imply that most of the return premia accrue to investors who purchased before the innovation signal.

How far do return premia extend?

We examine how the return premia vary with distance from the innovating startup. We modify Equation 5 by interacting $\mathbf{1}(\text{Approved})_i$ with indicators for each one-mile distance bin $d \in \{1, \ldots, 10\}$, and instrument each interaction with the corresponding examiner-leniency-by-distance-bin interaction. Formally, let $d(i, p) \in \{1, \ldots, 10\}$ denote the one-mile distance bin between startup i and property p, and define the bin-specific treatment $A_{i,d'} \equiv \mathbf{1}(\text{Approved})_i \cdot \mathbf{1}(d(i, p) = d')$. We estimate the following system by two-stage least squares. The first-stage equations are

$$A_{i,d'} = \pi_{0,d'} + \sum_{d=1}^{10} \pi_{1,d,d'} Z_{i,d} + \tilde{\lambda}_{gau_i \times appl.\ year_i} + \tilde{\lambda}_k + \tilde{\lambda}_{c \times t_0} + \nu_{i,k,c,t_0,d'}, \tag{6}$$

estimated jointly for each $d' \in \{1, \ldots, 10\}$. The corresponding second-stage equation is

$$y_{i,k,c,t_0,d} = \beta_0 + \sum_{d'=1}^{10} \beta_{1,d'} \widehat{A}_{i,d'} + \lambda_{gau_i \times appl.\ year_i} + \lambda_k + \lambda_{c \times t_0} + \varepsilon_{i,k,c,t_0,d}, \tag{7}$$

where $\widehat{A}_{i,d'}$ denotes the fitted value from Equation 6. The coefficient of interest, $\beta_{1,d'}$, captures the return premium from a patent grant for properties located in distance bin d' from the innovating startup. This yields ten separate treatment-effect coefficients from a single regression, each identifying the return premium at a given distance from the focal startup. As in the baseline specification, we weight by the number of repeat-sale transactions in each cell and cluster standard errors at the art-unit level.

Figure 3a plots the resulting IV estimates and 95 percent confidence intervals, with the baseline specification in the left panel and county-by-holding-term fixed effects added in the right panel. Two patterns stand out. First, return premia decline monotonically with distance for all four property types, with most of the decay occurring within the first five miles. In the baseline specification, the estimated premium within one mile of the innovating startup is 0.70 percentage points for single-family homes, 0.87 for multi-family properties, 1.02 for commercial real estate, and 1.96 for vacant land, and all are significant at the 5 percent level. By 10 miles, the corresponding estimates decline to 0.27, 0.51, 0.33, and 0.74 percentage points, respectively. The ordering across property types is largely preserved across the full distance range, with vacant land exhibiting the largest premia at every distance, consistent with development expectations being capitalized most sharply near the innovating firm.

Second, adding county-by-holding-term fixed effects (right panel) uniformly attenuates the estimates and accelerates the rate at which premia decay with distance. At one mile, the attenuation is modest. Roughly 25 to 27 percent for single-family, multi-family, and commercial properties, and only about 5 percent for vacant land. At 10 miles, by contrast, the estimates are attenuated by 46 to 76 percent, with the single-family and commercial premia falling to roughly 0.09 and 0.08 percentage points. This faster decay at greater distances is consistent with part of the baseline premium reflecting shorter holding periods in treated areas, a channel that operates relatively uniformly across the 10-mile window, so that once it is absorbed, the residual premium tracks more tightly the localized capitalization of innovation into property values. Estimates remain statistically significant at the 5 percent level out to 7 miles for commercial real estate and out to 6 miles for vacant land, whereas single-family and multi-family premia remain significant across the full 10-mile range. These patterns confirm that the real estate return premia, particularly the component attributable to price growth itself, are localized but span several miles.

Appendix Figure A.5 extends the analysis to 15 miles from the innovating startup. The pattern is consistent with the baseline results. Return premia decline steadily up to roughly six miles and then flatten out.

How do return premia evolve over time?

We now turn to the evolution of return premia over time. The specification is analogous to the distance-gradient design in Equation 7, replacing the distance bins with event-time bins. We interact $\mathbf{1}(\text{Approved})_i$ with indicators for each year of sale relative to the patent's FA date and instrument each interaction with the corresponding examiner-leniency-by-event-time interaction, yielding a vector of estimates $\{\beta_{1,\tau}\}$, where τ denotes the event-time year. We retain the art-unit-by-application-year, county-by-buy-year, and distance-bin fixed effects from that specification, and we additionally include county-by-holding-term fixed effects so that the event-time profile is not confounded with holding-term composition.[14] For $\tau > 0$ we pool the baseline and post samples, so that $\beta_{1,\tau}$ draws on every repeat sale sold τ years after the FA date, regardless of when the property was purchased. This keeps the profile comparable across τ.[15]

Figure 3b plots the IV estimates by τ. The pre-period estimates ($\tau = -5, \ldots, -1$) are small and close to zero across all property types, supporting the identifying assumption of no differential pre-trends between treated and control areas. Following the innovation ($\tau > 0$), all four property types exhibit a sharp return premium that peaks immediately at $\tau = 1$. The ordering at the peak mirrors the baseline and distance-gradient results. Vacant land commands the largest premium, whereas multi-family, single-family, and commercial properties show similar, smaller premia. The

[14] This control is especially important here. Because τ is the time from the FA date to the *sale*, longer-τ transactions are mechanically held longer, so absent the holding-term fixed effects the event-time estimates would partly trace the relationship between holding period and annualized returns rather than the timing of capitalization.

[15] Only using baseline sample, in which every transaction is bought before and sold after the FA date, raises a truncation problem. There, event time and holding period are mechanically linked. A sale observed at τ must have been held at least τ years, so the baseline subsample is increasingly truncated toward long-held properties as τ grows, and $\beta_{1,\tau}$ at long horizons would reflect a selected subpopulation rather than the premium at that horizon. Including the post subsample, whose transactions are purchased after the FA date, repopulates the short- and medium-holding cells at every τ and breaks this link.

premia then decline steadily, converging toward zero by roughly $\tau = 4$, and remain near zero at longer horizons. This front-loaded profile echoes the shorter holding terms documented in the descriptive evidence of Section 4.1 and the larger premia in the baseline than in the post sample reported earlier. Returns are realized disproportionately by properties that transact soon after the innovation signal.

Several mechanisms potentially explain this pattern. First, the market may price in the innovation signal quickly upon patent FA decision, capitalizing the expectations about neighborhood changes it entails. Second, supply responses may gradually erode the demand-driven price premium as new construction absorbs the innovation shock. Third, transaction selection may also play a role. Properties that benefit most from local spillovers, such as those closest to the startup, are disproportionately likely to be sold shortly after the innovation signal, as owners capitalize on the premium. As this high-return pool is exhausted, the composition of sellers shifts toward properties with smaller treatment exposure, mechanically attenuating the estimated premium at longer horizons. Overall, the results show front-loaded capitalization of the innovation signal, with the bulk of the return premium realized within the first three years.

How do return premia accumulate with exposure to multiple innovators?

The firm-level estimates are an average across treated properties, some near a single successful startup and others near many. To characterize how the premium scales with the number of nearby innovators, we examine the property-level sample, described in Section 3.2. Each repeat-sale transaction is now associated with N_i nearby patent-applying startups within 10 miles and $N_{1,i}$ of those whose FA decisions occurred between the property's buy and sell dates. The estimating equation is

$$r_{i,c,k,t_0,T} = \beta_0 + \beta_1 \mathbf{1}(\widehat{N_{1,i} > 0}) + \beta_2 N_i + \mathbf{X}_i'\gamma + \lambda_k + \lambda_{c\times t_0} + \lambda_{c\times T} + \varepsilon_{i,c,k,t_0,T}, \tag{8}$$

where r denotes the annualized unlevered return and $\mathbf{X}_i$ is a vector of property characteristics whose content depends on the property type: lot size for all types, additionally building square footage and year built for residential and commercial properties, and additionally the number of bedrooms and bathrooms for single-family and multi-family properties (so that vacant land enters with lot size alone). We instrument the binary indicator $\mathbf{1}(N_{1,i} > 0)$ with the average leniency of nearby examiners. Standard errors are clustered at the county level. Panel B of Appendix Table A.8 reports the first-stage estimates, which remain strong, with coefficients ranging from 0.39 to 0.67, and the cluster-robust Kleibergen–Paap rk first-stage F-statistics range from roughly 69 to 251 across subsamples.

Table 3 reports the IV estimates of β_1 from Equation 8. Panel A shows the overall property-level return premia. Columns (1), (3), (5), and (7) omit the property-characteristics controls, while Columns (2), (4), (6), and (8) include them. The estimates with controls indicate that exposure to at least one nearby patent-granted startup raises annualized returns by 5.2 percentage points for single-family homes, 16.5 percentage points for multi-family properties, 6.5 percentage points for commercial real estate, and 10.0 percentage points for vacant land, all significant at the 1% level. Treated properties are typically exposed to multiple successful startups. The average number of nearby grants conditional on exposure is 17.9 for single-family, 29.9 for multi-family, 23.4 for commercial, and 10.5 for vacant land. The implied per-startup effects are 0.29, 0.55, 0.28, and 0.95 percentage points, respectively, closely aligning with the firm-level estimates in Section 4.2.

Panel B of Table 3 examines how the marginal effect of each additional innovator varies with exposure. The estimates reveal pronounced diminishing returns. Across all four property types, the per-startup effect is largest when exposure is low and declines sharply as the cluster grows. For single-family properties, the return premium rises from 1.3 percentage points for transactions exposed to a single nearby patent grant, to 2.1 percentage points in the bottom quartile of exposure (mean of 2.8 nearby grants), and to 10.8 percentage points in the top quartile (mean of 62.7). The implied per-startup effect declines from 0.75 percentage points in the bottom quartile to 0.17 in the top quartile, which is a more than four-fold attenuation. The pattern is similar for the other property types. Per-startup effects decline from 1.25 to 0.32 percentage points for multi-family, from 0.90 to 0.16 for commercial, and from 2.44 to 0.68 for vacant land.

This concavity is consistent with real estate markets pricing the arrival of nearby innovation more than its scale. The first patent-granted startup in an area provides the strongest signal of future neighborhood economic activity, whereas subsequent innovators confirm but do not multiply the implied capitalization. Innovation events thus generate concave cumulative spillovers, with the largest marginal effects realized at the first innovation.

Panels C and D of Table 3 examine the placebo and post-treatment samples, as in the firm-level analysis. Panel C, restricted to repeat-sale transactions in which both buy and sell dates precede any nearby FA decision, shows return premia that are statistically insignificant and economically small across all property types, providing additional support for the validity of our instrumental variable strategy. Panel D, restricted to transactions in which both dates occur after the FA decision, similarly yields estimates that are insignificant or very small, consistent with the firm-level result that later buyers do not capture the return premium.

4.3 Effects on Local Innovation and Business Activity

Section 4.2 shows that successful patent grants are immediately and locally capitalized into nearby property values. For this capitalization to be rational rather than speculative, the patent grant must signal real economic activity that market participants reasonably expect to materialize. We now examine the two channels through which a successful innovating startup plausibly reshapes its surrounding neighborhood. We first show that a patent grant spurs further patenting among nearby firms, indicating that innovation diffuses locally. We then turn to broader agglomeration effects, documenting that successful startups draw additional business activity and employment into the area in the years following the grant. Together, these results validate the forward-looking expectations embedded in the real estate return premia documented above.

Knowledge spillovers.

We first explore the effect of a patent grant on patenting activity in the area. To do so, we estimate the following regression

$$y_{i,k,c} = \beta_0 + \beta_1 \mathbf{1}(\text{Approved})_i + \lambda_{gau_i \times appl.\ year_i} + \lambda_{c \times appl.\ year_i} + \epsilon_{ikc}, \tag{9}$$

where $y_{i,k,c}$ denotes the change in the number of patent applications (or approvals) by other startups within a 10-mile distance of startup i and within a 3-year window after the FA date of startup i. The change is relative to the year before a startup's first action date. The variable $\mathbf{1}(\text{Approved})_i$ is binary and indicates whether the startup received a patent grant. The variable is instrumented by examiner leniency, as before. The specification includes art unit-by-application year fixed effects ($\lambda_{gau_i \times appl.\ year_i}$) and county-by-application year fixed effects ($\lambda_{c \times appl.\ year_i}$). Standard errors are clustered at the art unit level. Since the sample comprises startups with a first action date from 2001 to 2009, we focus on the period from 2002 to 2006, allowing a year of prior activity and a 3-year window within which spillovers may occur.

Figure 4 presents results for the change in the number of patent applications within 10 miles of the focal granted patent and within a 3-year window after its FA date. The specifications include art unit-by-application year and county-by-application year fixed effects, and we plot the coefficient on the binary treatment $\mathbf{1}(\text{Approved})$, which indicates whether the focal application was approved. Panel (a) plots two coefficients, one for the change in nearby applications irrespective of the focal patent's sector and one for the change in applications by startups in the focal patent's own sector. The coefficient irrespective of sector is positive and statistically significant, whereas the own-sector coefficient is not. The pattern implies that a patent approval raises nearby patent applications in general rather than specifically within a sector.

Panel (b) of Figure 4 repeats the exercise for the change in the number of approved applications, again irrespective of sector and within the focal patent's own sector. The pattern is similar. Patent approval raises nearby approvals in general but not specifically within a sector.

Appendix Figure A.6 presents results for different distance bands. The band coefficients are not statistically distinguishable from one another, although the point estimates rise linearly out to 3 or 4 miles and then remain similar or decline slightly. Because each band farther from the focal startup covers a larger area, these coefficients imply a lower effect per unit area, consistent with a spillover that diminishes with distance.

Overall, a patent grant raises nearby patenting activity, but the spillover is broad rather than narrowly sectoral. The increase shows up in nearby applications and approvals overall but is not statistically distinguishable within the focal startup's own sector. This pattern suggests that successful innovation diffuses across technological boundaries within the local innovation cluster rather than only among direct same-sector competitors.

Establishment counts.

We document that a patent grant produces an increase in the number of nearby establishments in the applicant's own industry, with no detectable spillover to other industries over a 10-year horizon. To do so, we count establishments within 10 miles of each startup, disaggregated by 2-digit NAICS sector k and by 2-mile distance ring d, in each year from 3 years before to 10 years after the FA date. Let $N_{i,k,c,d,\tau}$ denote this count for startup i in county $c = c(i)$ at event time τ, and let $\Delta N_{i,k,c,d,\tau} \equiv N_{i,k,c,d,\tau} - N_{i,k,c,d,-1}$ be its change from the pre-treatment baseline $N_{i,k,c,d,-1}$, the count 1 year before the FA date. We construct ΔN for each $\tau \in \{-3, -2, 0, 1, \ldots, 10\}$ and winsorize it at the 0.1 and 99.9 percentiles.

To isolate within-industry from cross-industry effects, we estimate by two-stage least squares, for each event time τ,

$$\begin{aligned}\Delta N_{i,k,c,d,\tau} &= \beta_{1,\tau}\,\mathbf{1}(\widehat{\text{Approved}})_i \;+\; \beta_{2,\tau}\,\widehat{X}_{i,k} \;+\; \psi_\tau\,\mathbf{1}(i \in k) \\ &\quad + \; \gamma_\tau\,N_{i,k,c,d,-1} \;+\; \lambda_{gau_i \times appl.\ year_i} \;+\; \lambda_{c \times appl.\ year_i} \;+\; \lambda_{(appl.\ year_i+\tau)\times k} \;+\; \varepsilon_{i,k,c,d,\tau},\end{aligned} \tag{10}$$

where $\mathbf{1}(i \in k)$ equals one when the nearby sector k matches the applicant's own 2-digit NAICS sector and zero otherwise, $X_{i,k} \equiv \mathbf{1}(\text{Approved})_i \cdot \mathbf{1}(i \in k)$, and the lagged count $\tau N_{i,k,c,d,-1}$ controls for the baseline density.[16] The fixed effects absorb time-varying shocks at the art-unit and county levels and industry-specific calendar shocks in the year being measured. We instrument both endogenous regressors, $\mathbf{1}(\text{Approved})_i$ and $X_{i,k}$, with examiner leniency and its interaction with $\mathbf{1}(i \in k)$, exactly as in Section 4.2. The system is exactly identified, with one instrument per endogenous regressor, and we cluster standard errors at the art-unit level.

Two key parameters are of interest. $\beta_{1,\tau}$ is the effect of a patent grant that applies across all industries, and $\beta_{2,\tau}$ is the additional effect in the applicant's own industry, so that $\beta_{1,\tau} + \beta_{2,\tau}$ is the total own-industry effect.

Figure 5a plots $\beta_{2,\tau}$ and $\beta_{1,\tau}$ coefficient estimates and their 95 percent confidence intervals. The same-sector response is large and builds steadily. By year 10, the same-sector differential reaches about 6.51 additional establishments per startup-by-NAICS-by-ring cell, significant at the 1 percent level, so the total own-industry effect is about 7.19 establishments per cell. The different-sector effect is essentially zero, at about 0.68 establishments per cell and statistically insignificant at the 10 percent level, with a confidence interval that rules out effects above roughly 1.5 establishments per cell, less than a quarter of the same-sector estimate.

[16] Although $N_{i,k,c,d,-1}$ also enters the dependent variable, including it as a regressor is innocuous. The specification is algebraically equivalent to a levels regression of $N_{i,k,c,d,\tau}$ on the treatment with an unrestricted coefficient on the baseline, an ANCOVA adjustment, so the treatment coefficients are unchanged. The baseline term relaxes the unit-coefficient restriction implicit in a pure change-score and absorbs mean reversion in local density. We do not interpret γ_τ, which is mechanically attenuated by regression to the mean.

Two patterns reinforce this interpretation. First, the pre-period estimates at $\tau = -3$ and -2 are small and statistically indistinguishable from zero for both series, as the identifying assumption requires. Second, the same-sector effect turns positive and significant at $\tau = 1$ and rises almost monotonically through $\tau = 10$, whereas the different-sector effect stays near zero throughout, with its confidence interval covering zero at every horizon.

The timing and concentration of the response point to knowledge spillovers and supplier-network thickening among firms in the same technology area rather than to a generic local demand shock, which would lift all industries (Marshall, 1890; Glaeser et al., 1992; Ellison, Glaeser and Kerr, 2010). The magnitudes are sizable. Against a same-sector baseline of 40.8 establishments per cell, the year-10 effect of 7.19 is a 17.6 percent increase, whereas the cross-industry effect, set against a baseline of 64.8, is economically and statistically zero. A typical applicant has same-sector establishments in roughly 43 cells within 10 miles, so the implied firm-level effect by year 10 is about 310 additional establishments (about 260 at the median count of 36 cells).[17,18]

Local employment.

Patent approval raises nearby employment as well, and again the response falls overwhelmingly within the applicant's own industry. We re-estimate Equation 10 with the change in nearby employment, $\Delta E_{i,k,c,d,\tau} \equiv E_{i,k,c,d,\tau} - E_{i,k,c,d,-1}$, in place of the establishment count. Sample, instruments, fixed effects, and clustering are unchanged.

Figure 5b shows the same shape as for establishments. By year 10, the same-sector differential is about 53.5 jobs per cell, significant at the 1 percent level, for a total own-industry effect of about 59.7 jobs per cell. The different-sector estimate is about 6.22 jobs per cell and statistically

[17] The first stage is strong throughout. Because the instruments and the fixed-effect partition are time-invariant, the first stage is identical across event times. The first-stage F-statistic exceeds 281,000 for each instrumented regressor.

[18] Appendix B traces how the per-cell response varies with distance from the applicant, separately by 2-mile ring (Appendix Figures A.8a and A.9a).

insignificant at the 10 percent level, with a confidence interval that rules out effects above roughly 13.8 jobs per cell, again about a quarter of the same-sector estimate.

The pre-period estimates are small next to the post-treatment response. At $\tau = -2$ the different-sector estimate is statistically indistinguishable from zero and the same-sector estimate is marginally negative but an order of magnitude below the year-10 response. After the FA date the same-sector effect turns positive and significant within a year and climbs through $\tau = 10$, whereas the different-sector effect stays near zero early and drifts up only modestly at long horizons, never statistically distinguishable from zero.

Together the two outcomes pin down the size of the establishments that approval induces. Each new same-sector establishment carries about 8 employees (59.7 jobs spread over 7.19 establishments), roughly half the size of the average incumbent same-sector establishment (about 16 employees). Because the ratio divides one per-cell effect by another, it does not depend on the number of cells and is unaffected by how we aggregate to the firm level. In percentage terms, the year-10 same-sector response is about 9.2 percent of its baseline, against a cross-industry change of at most 2.1 percent. Summed across the roughly 43 same-sector cells around the average applicant, the implied firm-level effect is about 2,570 additional jobs by year 10 (about 2,150 at the median count of 36 cells). The establishment and employment responses together show that successful innovation generates real local economic activity, concentrated in the applicant's own industry, with at most a small cross-industry response that emerges only at longer horizons.[19]

Taken together, the real-estate premium appears immediately upon the patent decision, whereas the real-side effects that at least partly justify it, including nearby patenting, additional establishments, and employment, build only gradually over the following decade. A simple asset-pricing logic makes sense of the gap. A property is a claim on local rents, rents rise with local economic activity, and forward-looking buyers capitalize a startup's anticipated success as soon as it becomes public, ahead of the activity itself. The same logic organizes the cross-sectional patterns we doc-

[19]Appendix B documents the corresponding spatial gradient for nearby employment (Appendix Figures A.8b and A.9b).

ument. The premium decays with distance from the startup and diminishes with each additional nearby success. Appendix D sets out the framework formally and derives the predictions that the estimates bear out.

5 Conclusion

Real estate markets price a startup's success almost as soon as it happens, long before the growth that success brings arrives. Instrumenting patent approval with the quasi-random leniency of the assigned examiner, we find that a grant on a startup's first application raises annualized real estate returns within ten miles by between four-tenths and nine-tenths of a percentage point across property types. The premium is sharply local, fading within a few miles of the firm, and concave: the first successful startup in an area drives most of it, and each later one adds less. Prices respond almost immediately to the patent decision, whereas the real activity, nearby patenting, new establishments, and jobs concentrated in the startup's own industry, arrives only over the following decade.

The gains from a startup's innovation are not confined to the firm. A meaningful share is capitalized into the land and buildings around it. For example, a single nearby successful startup raises the value of a median single-family home by roughly \$1,000 a year, on the order of \$7,000 to \$9,500 over a typical holding period. Two implications follow. First, the windfall accrues substantially to incumbent property owners, so the social value of early-stage innovation is larger, and more geographically concentrated, than firm-level measures capture. Second, to the extent that place-based programs succeed in seeding the kind of local innovation we study, part of the public money spent to attract or support innovative firms capitalizes into surrounding property values, a transfer to nearby landowners that cost-benefit assessments of such programs should weigh.

Two questions stand out for future work. First, although we show that activity does materialize, we do not explore whether it fully accounts for the premium, so a direct test would ask whether the premium is larger where future activity turns out larger. Second, we find that the gains accrue to owners who held property across the patent decision rather than to later buyers, but how they are shared more broadly, among owners, renters, and the buyers drawn to these areas, remains unexplored. Even so, the findings suggest that a single early-stage innovator can shift the economic trajectory of its immediate surroundings, and that local real estate is an early and visible register of that change.

6 Tables and Figures

Table 1 — Summary Statistics of Startups

	By First Patent Outcome	
	Approved	Rejected
Panel A. # Observations		
No. of firms	18,923	10,564
% of firms	64.2	35.8
Panel B. Firm Characteristics (At Application Filing)		
Age (years), Median	2	2
Normalized HPI at first patent filing, Median	100	100
Employees at Filing Date		
Mean	29.5	29.4
Median	8	7
SD	61.7	62.1
Sales at Filing Date ($ million)		
Mean	4.3	4.3
Median	0.8	0.7
SD	9.7	10.0
Pre-Patent-Filing Employment Growth (%)		
Mean	15.2	14.7
SD	64.0	62.8
Pre-Patent-Filing Sales Growth (%)		
Mean	19.3	17.4
SD	84.2	80.4
Pre-Patent-Filing Establishment Growth (%)		
Mean	2.0	1.9
SD	5.1	4.9
Pre-Patent-Filing HPI Growth (%)		
Mean	7.3	7.7
SD	7.0	7.6
Total Establishments in ZIP Code at Filing Date		
Mean	976.3	1,070.4
Median	845	946
SD	683.5	714.6

Note. This table reports summary statistics for the 29,487 startups in our analysis sample, split by the outcome (granted vs. rejected) of their first patent application. Panel A reports the number and share of firms in each group. Panel B reports firm characteristics measured in the year the firm filed its first patent application. Employees, sales, and the count of nearby establishments are winsorized at the 2% level; pre-application growth rates compare the application year to the year immediately prior and are winsorized at the 1% level. Sales are deflated to 2001 USD using BLS CPI factors. The Normalized House Price Index sets the value at the application year equal to 100 within each ZIP code; the median therefore equals 100 by construction.

Figure 1 — Descriptive Evidence on Local Real Estate Outcomes

(a) Annualized returns by year relative to the patent decision

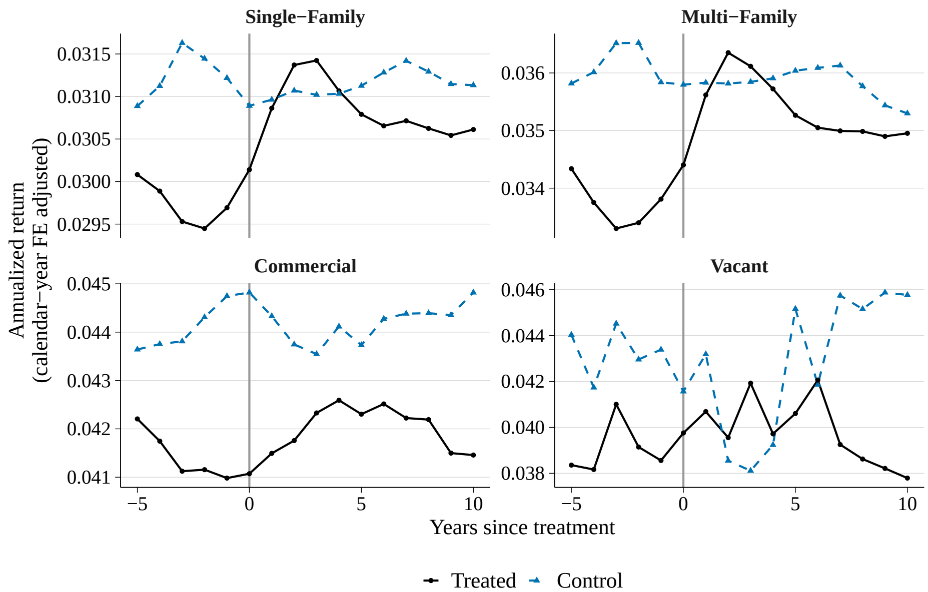

(b) Descriptive regression estimates

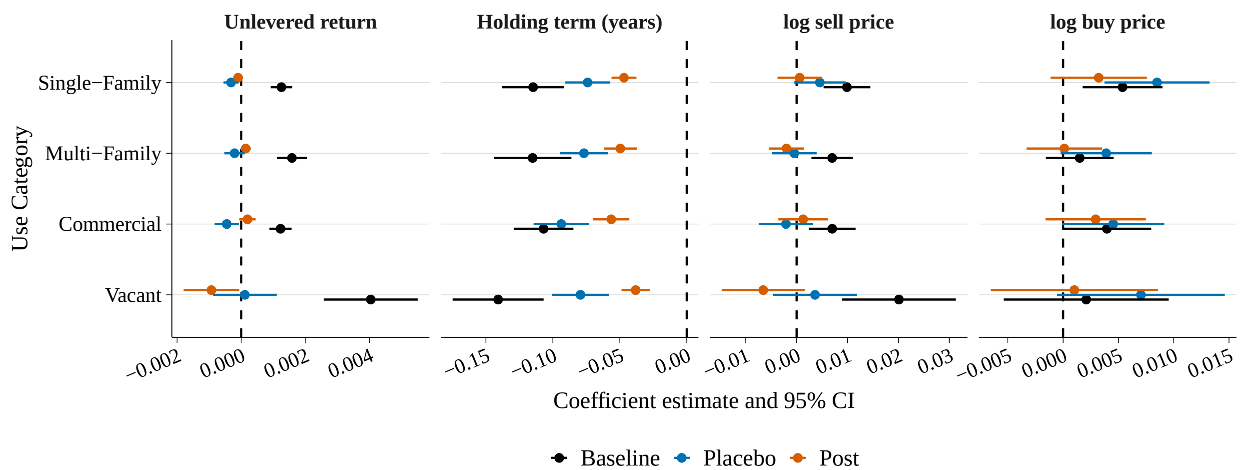

Notes: Panel (a) plots transaction-weighted average annualized returns (y-axis) from the firm-level repeat-sales sample against the year of sale relative to the firm's first-action (FA) date, separately for startups whose first application was eventually granted (treated, solid line) and rejected (control, dashed line). Returns are adjusted for property-type-by-calendar-year mean returns and recentered to the property-type mean, as described in Section 4.1. Observations pool the placebo, baseline, and post samples and are restricted to repeat sales within 10 miles of the startup. Facets correspond to property use categories, and the vertical line marks the associated patent application's first-action year. Panel (b) shows coefficient estimates of $\mathbf{1}$(Approved) and their 95% confidence intervals from Equation 3, across four outcome variables in each facet: unlevered return, log sell price, log buy price, and holding term. All regressions in Panel (b) include art-unit-by-application-year fixed effects, property type fixed effects, county-by-buy-year fixed effects, and distance-bin fixed effects, and report results for three samples (baseline, placebo, and post) across four property use categories (single-family, multi-family, commercial, and vacant). Standard errors are clustered at the art-unit level.

Figure 2 — Descriptive Evidence on Local Business Activity

(a) Establishment Counts and employment by year relative to the patent decision

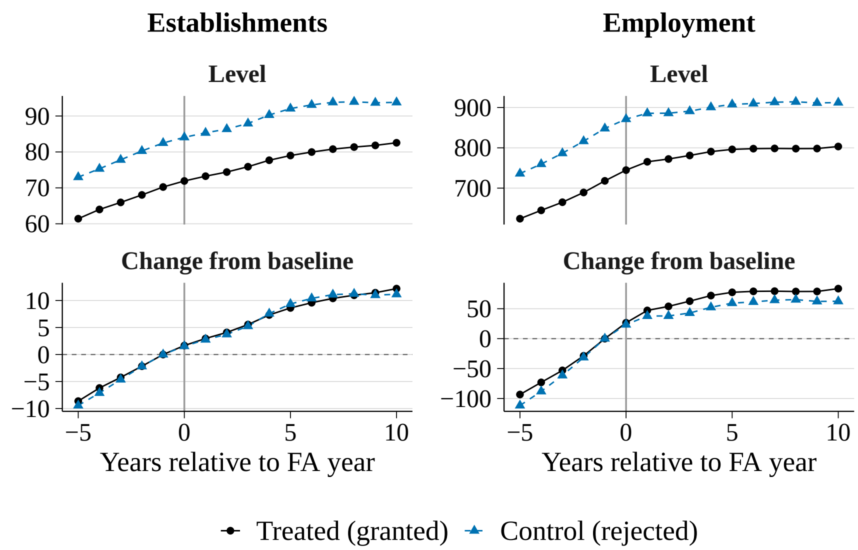

(b) Descriptive regression estimates

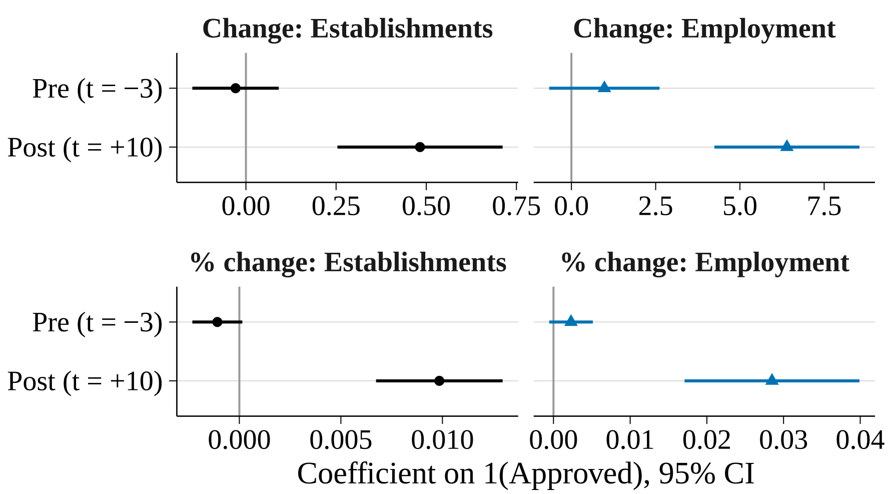

Note. Panel (a) plots the average number of nearby establishments and the level of nearby employment by year relative to the firm's first-action (FA) date, separately for patent-granted (treated) and rejected (control) startups; the top row shows levels and the bottom row the change from the pre-FA baseline, each outcome on its own vertical scale. Panel (b) displays the coefficient β_1 on $\mathbf{1}$(Approved) and its 95% confidence interval from Equation 4, for the change and the percent change in establishments and in employment, at a pre-period placebo horizon (three years before the decision) and at the ten-year horizon. All regressions in Panel (b) include art-unit-by-application-year, NAICS-sector, county-by-application-year, and distance-ring fixed effects, and standard errors are clustered at the art unit.

Table 2 — Effects on Local Real Estate

	Single-Family		Multi-Family		Commercial		Vacant	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Baseline sample — properties bought before and sold after treatment								
A.1. Outcome — median annualized unlevered return								
1(Approved)	0.0038***	0.0020***	0.0064***	0.0039***	0.0050***	0.0023***	0.0091***	0.0058***
	(0.0007)	(0.0005)	(0.0011)	(0.0007)	(0.0008)	(0.0005)	(0.0027)	(0.0021)
A.2. Outcome — log median buy price								
1(Approved)	-0.0060	0.0011	-0.0135**	-0.0067	-0.0152*	-0.0083	-0.0228	-0.0194
	(0.0067)	(0.0070)	(0.0057)	(0.0056)	(0.0084)	(0.0094)	(0.0158)	(0.0148)
A.3. Outcome — log median sell price								
1(Approved)	0.0043	0.0097	0.0046	0.0097	-0.0073	0.0009	0.0022	-0.0026
	(0.0084)	(0.0086)	(0.0076)	(0.0075)	(0.0090)	(0.0102)	(0.0230)	(0.0214)
A.4. Outcome — holding term (in years)								
1(Approved)	-0.4591***		-0.4995***		-0.5186***		-0.5827***	
	(0.0610)		(0.0737)		(0.0634)		(0.0784)	
Observations	38,101,658	38,101,658	35,986,905	35,986,905	5,237,993	5,237,993	1,595,151	1,595,151
Holding term FE		✓		✓		✓		✓

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel B. Placebo and post-treatment samples — outcome: median annualized unlevered return								
B.1. Placebo — repeat sales fully before treatment								
1(Approved)	5.78×10^{-5}	0.0001	0.0003	0.0001	0.0002	0.0002	-0.0034	-0.0008
	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0007)	(0.0007)	(0.0036)	(0.0035)
Observations	10,331,945	10,331,945	12,662,816	12,662,816	1,781,337	1,781,337	377,622	377,622
B.2. Post-treatment — repeat sales fully after treatment								
1(Approved)	6.48×10^{-5}	6.49×10^{-5}	0.0002	0.0002	0.0007	0.0007*	0.0002	0.0005
	(0.0003)	(0.0003)	(0.0002)	(0.0002)	(0.0004)	(0.0004)	(0.0021)	(0.0019)
Observations	31,409,014	31,409,014	28,623,627	28,623,627	4,172,772	4,172,772	2,201,112	2,201,112
Holding term FE		✓		✓		✓		✓

Standard errors clustered at the art unit level in parentheses.
Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

This table reports coefficient estimates of β_1 from the IV regressions in Equation 5. Panel A uses the baseline sample, consisting of transactions for which the property was purchased before and sold after the first-action date of the associated patent application. The outcome variables are the median annualized unlevered return, the log median buy price, the log median sell price, and the holding term. Panel B uses the placebo sample (transactions for which both the purchase and sale occurred before the first-action date) and the post-treatment sample (transactions for which both the purchase and sale occurred after the first-action date), with the median annualized unlevered return as the outcome variable. Each subpanel is estimated separately for four property use categories, namely single-family residential, multi-family residential, commercial, and vacant real estate. Columns (1), (3), (5), and (7) report the baseline estimates, including art-unit-by-application-year, property type, county-by-buy-year, and distance-bin fixed effects. Columns (2), (4), (6), and (8) additionally include county-by-holding-term (in years) fixed effects. Because the holding term is mechanically absorbed by these fixed effects in the corresponding subpanel of Panel A, those columns are left blank.

Figure 3 — Real Estate Return Premia by Distance and Time

(a) By Distance from the Innovating Startup

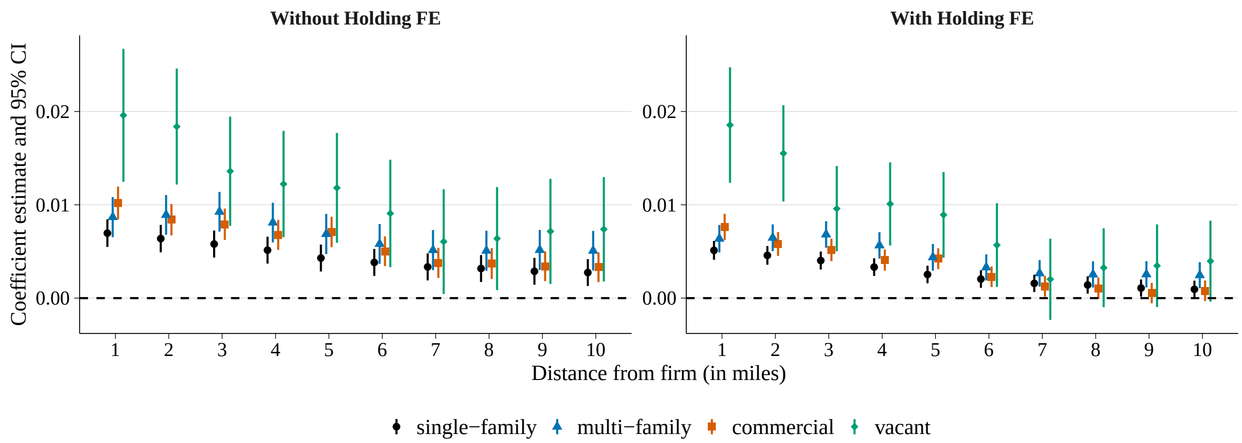

(b) By Years Since Treatment

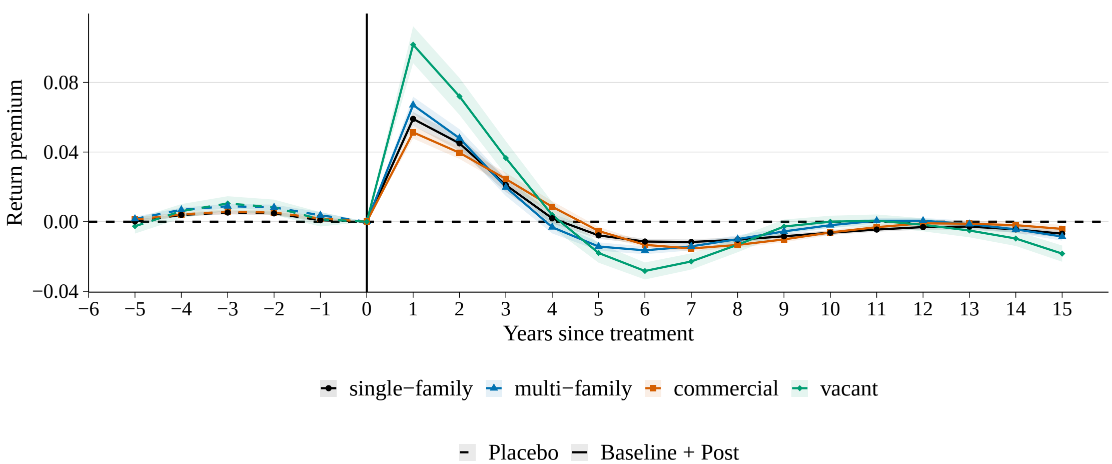

Note. Panel (a) plots IV estimates of the distance-specific return premia $\beta_{1,d'}$ from Equation 7, with 95 percent confidence intervals. Each point corresponds to the estimated return premium at distance bin $d' \in \{1, \ldots, 10\}$ miles from the innovating startup, estimated jointly within a single regression that instruments ten treatment-by-distance interactions with the corresponding examiner-leniency-by-distance interactions. The left panel reports the baseline specification with art-unit-by-application-year, property type, and county-by-buy-year fixed effects; the right panel additionally includes county-by-holding-term fixed effects. Panel (b) plots IV estimates of the event-time return premia $\beta_{1,\tau}$, with 95 percent confidence intervals. For $\tau > 0$, estimates use the baseline and post samples; for $\tau < 0$, estimates use the placebo sample. Solid lines correspond to the post-treatment period and dashed lines to the placebo period. The specification includes county-by-holding-term fixed effects in addition to the baseline controls. In both panels, estimates are shown separately by property use type: single-family residential (purple), multi-family residential (blue), commercial (green), and vacant land (yellow). Regressions are weighted by the number of repeat-sale transactions in each cell, and standard errors are clustered at the art-unit level.

Table 3 — Cumulative Effects of Nearby Innovative Startups

	Single-Family		Multi-Family		Commercial		Vacant	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Baseline repeat sales — exposure to at least one patent grant								
$\mathbf{1}(\widehat{N_1 > 0})$	0.0560***	0.0522***	0.1706***	0.1651***	0.0757***	0.0648***	0.1010***	0.0999***
	(0.0064)	(0.0061)	(0.0258)	(0.0254)	(0.0110)	(0.0100)	(0.0231)	(0.0231)
Mean N_1 cond. $N_1 > 0$	17.87		29.89		23.42		10.48	
Per-startup (pp)	0.31	0.29	0.57	0.55	0.32	0.28	0.96	0.95
Observations	11,696,097	11,696,097	4,142,706	4,142,706	222,234	222,234	163,678	163,678
Panel B. Heterogeneity by number of nearby patent grants (N_1)								
$\mathbf{1}(\widehat{N_1 = 1})$	0.0149***	0.0133***	0.0219***	0.0213***	0.0201***	0.0183***	0.0382**	0.0379**
	(0.0026)	(0.0024)	(0.0059)	(0.0056)	(0.0073)	(0.0070)	(0.0182)	(0.0182)
$\mathbf{1}(\widehat{N_1 \in Q_{N_1,1}})$	0.0231***	0.0211***	0.0482***	0.0460***	0.0314***	0.0291***	0.0589***	0.0582***
	(0.0026)	(0.0026)	(0.0069)	(0.0065)	(0.0062)	(0.0060)	(0.0187)	(0.0186)
$\mathbf{1}(\widehat{N_1 \in Q_{N_1,2}})$	0.0431***	0.0391***	0.0978***	0.0935***	0.0634***	0.0563***	0.1003***	0.1002***
	(0.0046)	(0.0042)	(0.0134)	(0.0132)	(0.0100)	(0.0096)	(0.0285)	(0.0283)
$\mathbf{1}(\widehat{N_1 \in Q_{N_1,3}})$	0.0668***	0.0626***	0.1633***	0.1587***	0.0946***	0.0828***	0.2142***	0.2119***
	(0.0080)	(0.0074)	(0.0224)	(0.0226)	(0.0137)	(0.0132)	(0.0456)	(0.0456)
$\mathbf{1}(\widehat{N_1 \in Q_{N_1,4}})$	0.1113***	0.1082***	0.3090***	0.3090***	0.1435***	0.1264***	0.2900***	0.2843***
	(0.0183)	(0.0182)	(0.0747)	(0.0741)	(0.0279)	(0.0259)	(0.0768)	(0.0742)
$Q_{N_1,1}$ range [mean]	2–4 [2.81]		2–6 [3.68]		2–5 [3.22]		2–3 [2.39]	
$Q_{N_1,2}$ range [mean]	5–9 [6.72]		7–14 [10.06]		6–11 [8.19]		4–6 [4.84]	
$Q_{N_1,3}$ range [mean]	10–23 [15.15]		15–37 [23.68]		12–28 [18.46]		7–15 [10.09]	
$Q_{N_1,4}$ range [mean]	24–842 [62.71]		38–844 [97.21]		29–834 [81.04]		16–791 [41.81]	
Panel C. Placebo — repeat sales fully before treatment								
$\mathbf{1}(\widehat{N_1 > 0})$	-0.0021	0.0047	0.0264	0.0301	-0.0046	-0.0018	-0.0038	-0.0052
	(0.0065)	(0.0062)	(0.0199)	(0.0209)	(0.0114)	(0.0115)	(0.0183)	(0.0182)
Observations	9,152,691	9,152,691	3,092,237	3,092,237	139,884	139,884	111,828	111,828
Panel D. Post-treatment — repeat sales fully after treatment								
$\mathbf{1}(\widehat{N_1 > 0})$	0.0042	0.0094*	0.0483	0.0590	-0.0333	-0.0298	-0.0026	-0.0005
	(0.0066)	(0.0056)	(0.0475)	(0.0500)	(0.0319)	(0.0304)	(0.0387)	(0.0389)
Observations	5,464,306	5,464,306	1,517,271	1,517,271	85,421	85,421	159,973	159,973
Property Characteristics		✓		✓		✓		✓

Standard errors clustered at the county level in parentheses.
Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Note. This table presents the property-level analysis results from Equation 8, where the outcome is the annualized unlevered return and the treatment indicator $\mathbf{1}(N_{1,i} > 0)$ equals one if any nearby patent-applying startup received a first-action grant between the property's buy and sell dates, and is instrumented by the average leniency of nearby examiners. Specifications include property type, county-by-buy-year, and county-by-holding-term fixed effects, as well as the total number of nearby patent-applying startups N_i as a control. In Panel B, the treatment is partitioned by the number of nearby grants: $N_1 = 1$ for transactions exposed to a single grant, and $Q_{N_1,1}$ through $Q_{N_1,4}$ for transactions in the quartiles of the remaining distribution.

Figure 4 — Effects on Knowledge Spillovers

(a) Patent Applications **(b) Patent Approvals**

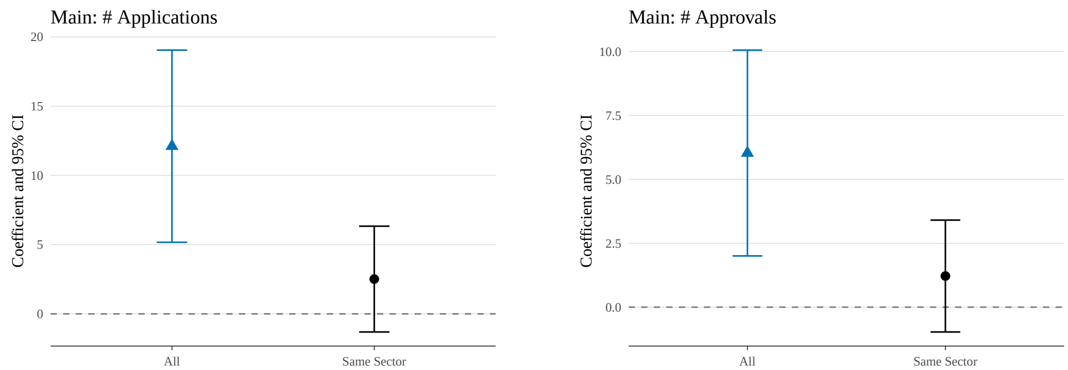

Note. The figure plots the coefficient on the binary treatment **1**(Approved), which indicates whether the focal application was approved, from the patent-spillover regression in Equation 4.3. The outcome is the change in the number of nearby patent applications (Panel a) or approved applications (Panel b) by other startups within 10 miles of the focal startup and within a 3-year window after its first-action date, measured relative to the year before that date. Within each panel, the first coefficient uses the change irrespective of the focal patent's sector and the second uses only startups in the focal patent's own sector. The specifications include art unit-by-application year and county-by-application year fixed effects. Bars denote 95 percent confidence intervals, with standard errors clustered at the art-unit level.

Figure 5 — Effects on Local Business Activity

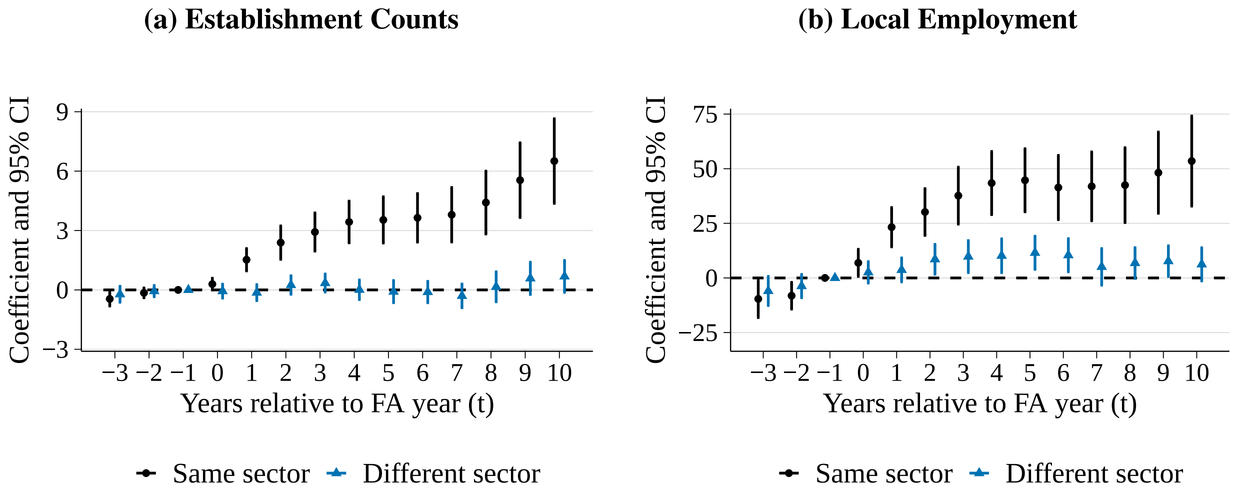

Note. Event-study IV estimates of $\beta_{1,\tau}$ (different sector) and $\beta_{2,\tau}$ (same sector) from Equation 10, plotted against event time τ (years from the first-action date). The same-sector series is the differential $\beta_{2,\tau}$, so the total same-sector effect is the sum of the two series. Panel (a) uses $\Delta N_{i,k,c,d,\tau}$, the change in nearby establishments from the pre-treatment baseline ($\tau = -1$); Panel (b) uses $\Delta E_{i,k,c,d,\tau}$, the change in nearby employment from the same baseline. Each point is a separate IV regression at event time τ, and the pre-period estimates ($\tau \leq -2$) serve as a within-figure placebo. Sample, instruments, and fixed effects are identical across the two panels. Bars denote 95 percent confidence intervals, with standard errors clustered at the art-unit level.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work, the authors used ChatGPT and Claude to review code and proofread the draft. After using these tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

References

Adelino, M., Ma, S., Robinson, D. 2017. Firm age, investment opportunities, and job creation. The Journal of Finance, 72, 999–1038.

Aiello, D., Kotter, J. D., Schubert, G. 2022. When money moves in: The consequences of housing wealth. Available at SSRN 4280776.

Akcigit, U., Kerr, W. R. 2018. Growth through heterogeneous innovations. Journal of Political Economy, 126, 1374–1443.

Beracha, E., He, Z., Wintoki, M. B., Xi, Y. 2022. On the relation between innovation and housing prices–a metro level analysis of the us market. The Journal of Real Estate Finance and Economics, 1–27.

Beracha, E., Kim, E., Wintoki, M. B., Xi, Y. 2025. Highly skilled immigrants and local housing prices: Evidence from the h-1b visa lottery. Journal of Real Estate Research.

Bernstein, S., Colonnelli, E., Giroud, X., Iverson, B. 2019. Bankruptcy spillovers. Journal of Financial Economics, 133, 608–633.

Butler, A. W., Fauver, L., Spyridopoulos, I. 2019. Local economic spillover effects of stock market listings. Journal of Financial and Quantitative Analysis, 54, 1025–1050.

Carley, M., Hegde, D., Marco, A. 2015. What is the probability of receiving a us patent. Yale Journal of Law and Technology, 17, p. 203.

Carlino, G., Drautzburg, T. 2020. The role of startups for local labor markets. Journal of Applied Econometrics, 35, 751–775.

Carlino, G., Kerr, W. R. 2015. Agglomeration and innovation. Handbook of regional and urban economics, 5, 349–404.

Chatterji, A., Glaeser, E., Kerr, W. 2014. Clusters of entrepreneurship and innovation. Innovation policy and the economy, 14, 129–166.

Cornaggia, J., Mao, Y., Tian, X., Wolfe, B. 2015. Does banking competition affect innovation? Journal of financial economics, 115, 189–209.

Coulson, N. E., Liu, C. H., Villupuram, S. V. 2013. Urban economic base as a catalyst for movements in real estate prices. Regional Science and Urban Economics, 43, 1023–1040.

Derrien, F., Kecskés, A., Nguyen, P.-A. 2023. Labor force demographics and corporate innovation. The Review of Financial Studies, 36, 2797–2838.

Ellison, G., Glaeser, E. L., Kerr, W. R. 2010. What causes industry agglomeration? evidence from coagglomeration patterns. American Economic Review, 100, 1195–1213.

Farre-Mensa, J., Hegde, D., Ljungqvist, A. 2020. What is a patent worth? evidence from the us patent "lottery". The Journal of Finance, 75, 639–682.

Garin, A., Rothbaum, J. 2025. The long-run impacts of public industrial investment on local development and economic mobility: Evidence from world war ii. The Quarterly Journal of Economics, 140, 459–520.

Glaeser, E. L., Kallal, H. D., Scheinkman, J. A., Shleifer, A. 1992. Growth in cities. Journal of political economy, 100, 1126–1152.

Glaeser, E. L., Kerr, S. P., Kerr, W. R. 2015. Entrepreneurship and urban growth: An empirical assessment with historical mines. Review of Economics and Statistics, 97, 498–520.

Goldschlag, N., Lybbert, T. J., Zolas, N. J. 2020. Tracking the technological composition of industries with algorithmic patent concordances. Economics of Innovation and New Technology, 29, 582–602.

Gourio, F., Messer, T., Siemer, M. 2016. Firm entry and macroeconomic dynamics: A state-level analysis. American Economic Review, 106, 214–218.

Greenstone, M., Hornbeck, R., Moretti, E. 2010. Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings. Journal of political economy, 118, 536–598.

Grieser, W., Liu, Z. 2019. Corporate investment and innovation in the presence of competitor constraints. The Review of Financial Studies, 32, 4271–4303.

Haltiwanger, J., Jarmin, R. S., Miranda, J. 2013. Who creates jobs? small versus large versus young. Review of Economics and Statistics, 95, 347–361.

Hartman-Glaser, B., Thibodeau, M., Yoshida, J. 2023. Cash to spend: Ipo wealth and house prices. Real Estate Economics, 51, 68–102.

Hegde, D., Ljungqvist, A., Raj, M. 2022. Quick or broad patents? evidence from us startups. The Review of Financial Studies, 35, 2705–2742.

Hombert, J., Matray, A. 2017. The real effects of lending relationships on innovative firms and inventor mobility. The Review of Financial Studies, 30, 2413–2445.

Hu, M., Tsang, D., Wan, W. X. 2025. Corporate relocation and housing market spillovers. Management Science, 71, 4344–4376.

Jaffe, A. B., Trajtenberg, M., Henderson, R. 1993. Geographic localization of knowledge spillovers as evidenced by patent citations. the Quarterly journal of Economics, 108, 577–598.

Kline, P., Moretti, E. 2014. People, places, and public policy: Some simple welfare economics of local economic development programs. Annu. Rev. Econ. 6, 629–662.

Kline, P., Petkova, N., Williams, H., Zidar, O. 2019. Who profits from patents? rent-sharing at innovative firms. The quarterly journal of economics, 134, 1343–1404.

Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N. 2017. Technological innovation, resource allocation, and growth. The Quarterly Journal of Economics, 132, 665–712.

Ma, S. 2025. Technological obsolescence. The Review of Financial Studies, p. hhaf059.

Marshall, A. 1890. Principles of Economics. London, Macmillan.

Matray, A. 2021. The local innovation spillovers of listed firms. Journal of Financial Economics, 141, 395–412.

Monte, F., Redding, S. J., Rossi-Hansberg, E. 2018. Commuting, migration, and local employment elasticities. American Economic Review, 108, 3855–3890.

Moretti, E. 2021. The effect of high-tech clusters on the productivity of top inventors. American Economic Review, 111, 3328–3375.

Pope, D. G., Pope, J. C. 2015. When walmart comes to town: Always low housing prices? always? Journal of Urban Economics, 87, 1–13.

Qian, F., Tan, R. 2021. The effects of high-skilled firm entry on incumbent residents. Stanford Institute for Economic Policy Research (SIEPR) Working Paper, 21–039.

Sampat, B., Williams, H. L. 2019. How do patents affect follow-on innovation? evidence from the human genome. American Economic Review, 109, 203–236.

Walsh, C. 2025. The entry multiplier. Available at SSRN 5468409.

Appendix

A Supplementary Tables and Figures

Figure A.1 — Spatial Distribution of Startups in the Analysis Sample

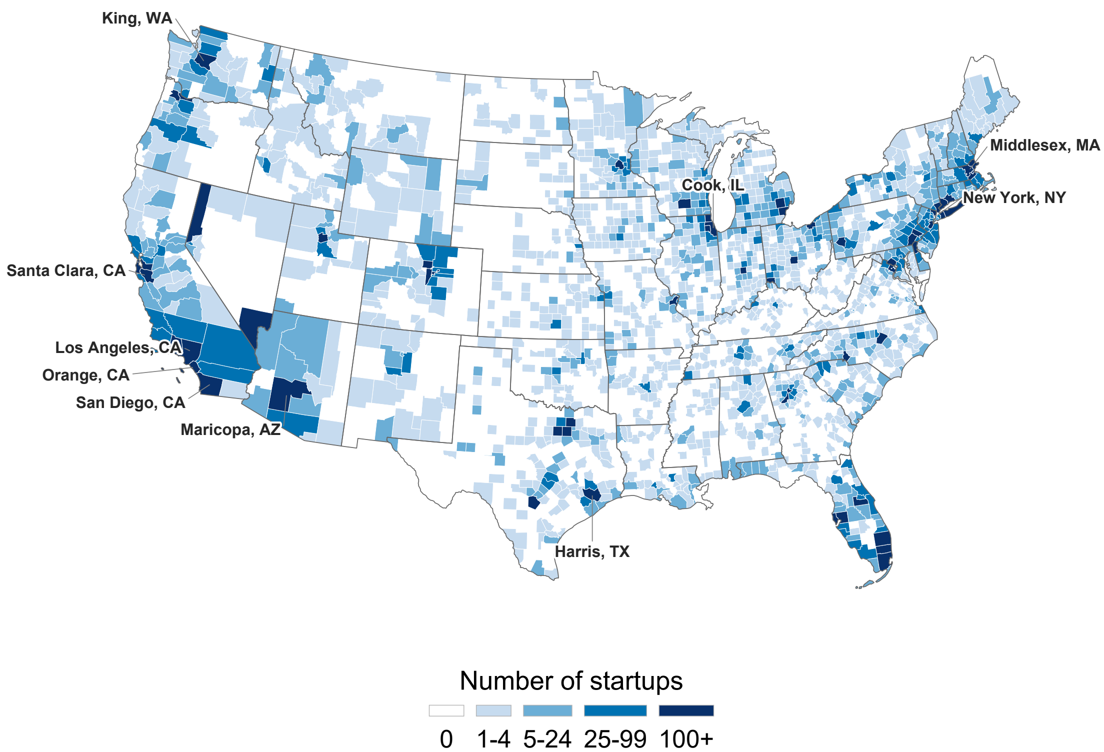

Note. This figure displays the spatial distribution of startups in our analysis sample, showing the number of startups located in each county across the continental United States. An additional 29 startups in the sample are located in Alaska and Hawaii.

Table A.1 — Summary Statistics of Startups

	By First Patent Outcome	
	Approved	Rejected
Panel A. # Observations		
No. of firms	18,923	10,564
% of firms	64.2	35.8
Panel B. Firm Characteristics (at first-action year)		
Age at first-action (years), Median	4	4
Normalized HPI at first-action, Median	100	100
Employees at First-Action Date		
Mean	29.6	29.2
Median	7	7
SD	62.1	62.9
Sales at First-Action Date ($ million)		
Mean	4.2	4.1
Median	0.8	0.6
SD	9.8	9.9
Pre-First-Action Employment Growth (%)		
Mean	11.8	10.0
SD	53.9	51.7
Pre-First-Action Sales Growth (%)		
Mean	14.7	11.7
SD	69.0	65.5
Pre-First-Action Establishment Growth (%)		
Mean	1.4	0.8
SD	4.8	4.5
Pre-First-Action HPI Growth (%)		
Mean	4.4	2.2
SD	9.2	10.0
Total Establishments in ZIP Code at First-Action Date		
Mean	992.6	1,088.7
Median	863	965
SD	685.9	716.0

Note. This table reports summary statistics for the 29,487 startups in our analysis sample, split by the final outcome (granted vs. rejected) of their first patent application. Panel A reports the number and share of firms in each group. Panel B reports firm characteristics measured in the year of the first-action decision on the first patent application. Employees, sales, and the count of nearby establishments are winsorized at the 2% level; pre-first-action growth rates compare the first-action year to the year immediately prior and are winsorized at the 1% level. Sales are deflated to 2001 USD using BLS CPI factors. The Normalized House Price Index sets the value at the first-action year equal to 100 within each ZIP code; the median therefore equals 100 by construction.

Figure A.2 — Distribution of Examiner Approval Rates (Residualized)

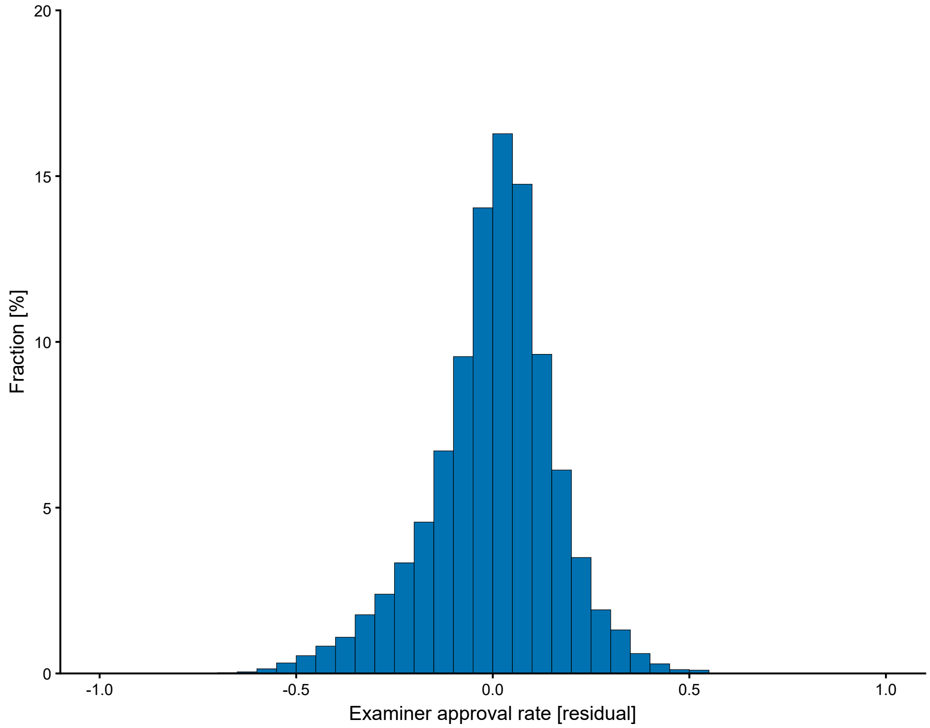

Note. Histogram of residuals from a regression of the examiner leniency instrument (Section 3.1) on art-unit-by-application-year fixed effects, restricted to the first-stage estimation sample of 29,487 startups. Each observation is one application. The residualization absorbs all variation that is common to a cohort of cases in the same technology art unit in the same application year; what remains is the within-cell dispersion in examiner leniency. The gap in residual leniency between the 25th and 75th percentile examiners is 17.5 percentage points. Combined with the first-stage coefficient $\widehat{\theta} = 0.671$, this implies an 11.7-percentage-point gap in approval probability between an otherwise identical application drawn by a 25th-percentile versus a 75th-percentile examiner. Bin width is 0.05; the y-axis reports the fraction of applications in each bin (in percent).

Table A.2 — Summary Statistics on Patent Application Spillover

	N	Mean	P1	P50	P99	SD
Panel A.						
# Applications	17,222	35.00	-4.00	14.00	274.00	54.49
# Applications (Same Sector)	17,222	11.38	-4.00	4.00	115.00	21.93
Panel B.						
# Approvals	17,222	19.60	-4.00	8.00	183.00	33.20
# Approvals (Same Sector)	17,222	6.23	-3.00	2.00	73.00	13.12

Note. This table reports summary statistics on measures of patenting spillovers. For each patent application of a startup, we count the number of nearby applications by other startups within a 10 mile distance. The count period pertains to a 3-year window after the focal patent application's first action date. The table presents statistics on the counts post the focal patent application's first-action date relative to the counts in the year prior to the focal patent application's first action date.

Table A.3 — Property Types by Use Category

Use Category	Property Type (as originally coded in ATTOM)	# Repeat Sales in Sample
Single-family	Single family residence	14,046,388
	Townhouse	880,924
	Single family residential	102,984
	Row house	101,188
	Zero lot line (residential)	59,843
	Residential (general/single)	33,895
	Bungalow (residential)	9,312
Multi-family	Condominium	3,552,058
	Duplex (2 units, any combination)	398,742
	Apartments (generic)	124,486
	Triplex (3 units, any combination)	94,263
	Multi-family dwellings (generic, any combination)	82,236
	Quadplex (4 units, any combination)	61,174
	Apartment house (5+ units)	56,352
	Cluster home	52,649
	Residential income (general/multi-family)	40,623
	Garden apt, court apt (5+ units)	8,750
Commercial	Retail stores (personal services, photography, travel)	64,612
	Commercial (general)	51,621
	Office building	49,149
	Commercial office/residential (mixed use)	34,680
	Restaurant	19,706
	Condominium offices	19,410
	Commercial office (general)	12,883
Vacant	Residential - vacant land	94,999
	Vacant land	23,393
	Vacant commercial	23,262

Table A.4 — Summary Statistics of Firm-level Repeat Sales Sample

Sample	1(Approved)	Annualized Unlevered Return (%)					Buy Price ($1000s)					Sell Price ($1000s)				
		Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD
Panel A. Single-Family																
Baseline	0	-2.11	0.18	2.81	0.43	5.30	192.79	283.42	432.41	350.42	391.36	189.31	303.85	491.22	392.96	535.82
Baseline	1	-1.73	0.63	3.49	1.07	5.42	183.74	269.94	413.12	332.25	344.84	186.83	300.07	491.41	385.72	469.11
Placebo	0	-2.51	0.80	3.91	1.00	5.47	188.67	283.57	397.62	312.06	173.13	206.71	302.84	440.97	353.30	218.75
Placebo	1	-2.60	0.68	3.84	0.95	5.57	186.29	281.05	394.59	308.53	170.29	203.35	299.38	434.44	347.07	213.13
Post	0	-1.72	1.14	4.70	1.53	6.17	171.12	275.07	430.47	343.70	446.51	183.77	292.40	471.79	390.03	601.77
Post	1	-2.31	0.29	3.79	0.76	6.22	171.99	271.55	425.48	335.95	385.38	173.22	273.08	444.10	364.81	513.65
Panel B. Multi-Family																
Baseline	0	-2.76	0.38	3.97	0.80	8.02	174.29	287.10	460.99	408.98	1120.56	170.66	317.45	544.96	484.21	1488.83
Baseline	1	-2.33	0.96	4.91	1.63	8.22	162.81	267.25	436.59	382.11	986.04	167.30	313.44	536.92	467.43	1315.37
Placebo	0	-4.32	-0.29	4.13	0.17	9.78	189.49	268.66	394.60	341.75	333.60	178.09	267.59	424.59	362.93	425.99
Placebo	1	-4.47	-0.47	3.94	0.02	9.76	190.15	267.71	392.82	340.78	328.51	175.14	263.19	417.97	356.75	408.76
Post	0	-2.36	1.42	6.41	2.18	9.72	160.52	291.34	477.57	442.04	1465.19	164.39	307.06	544.87	525.08	1885.11
Post	1	-3.17	0.38	5.25	1.16	9.65	160.29	284.28	470.96	429.28	1362.87	151.16	282.27	513.61	486.35	1734.61
Panel C. Commercial																
Baseline	0	-2.60	1.53	6.86	2.47	10.03	223.99	433.00	820.95	909.00	3364.53	260.53	537.27	1053.11	1147.98	4028.36
Baseline	1	-2.31	2.10	7.81	3.16	10.24	209.74	408.57	795.60	815.53	2700.67	256.28	529.04	1051.91	1057.62	3276.55
Placebo	0	-5.94	-0.35	6.66	0.54	12.61	212.28	388.92	767.28	680.50	1051.21	208.73	391.36	798.43	716.08	1299.12
Placebo	1	-6.17	-0.59	6.44	0.35	12.64	214.29	394.57	780.59	686.57	1042.91	207.01	391.40	802.12	709.10	1243.33
Post	0	-3.14	1.47	8.62	2.84	12.04	227.57	478.28	901.64	1126.08	4688.36	235.67	524.93	1086.11	1360.74	5590.35
Post	1	-3.70	0.66	7.29	2.04	11.77	226.90	471.17	897.90	1003.14	3793.54	222.18	493.89	1029.30	1172.32	4522.81
Panel D. Vacant																
Baseline	0	-11.19	-1.00	7.39	-1.68	23.29	53.24	162.69	401.29	375.27	1072.36	27.60	145.46	444.76	559.17	2669.18
Baseline	1	-9.53	-0.03	8.96	-0.05	23.39	48.83	144.56	360.88	330.97	910.36	29.01	143.24	429.30	493.90	2246.12
Placebo	0	-8.28	-0.33	7.77	0.08	25.33	91.00	198.30	396.23	304.87	334.30	78.13	193.86	397.99	323.69	426.29
Placebo	1	-8.01	-0.42	7.72	0.23	25.33	90.90	201.66	396.00	302.83	329.49	77.76	190.66	390.04	313.94	407.35
Post	0	-13.10	0.67	14.20	0.80	32.14	28.55	121.27	355.04	358.27	1282.14	23.60	117.20	376.10	506.91	2754.51
Post	1	-14.37	-0.61	12.07	-0.54	31.51	29.03	115.96	340.26	326.80	1137.57	20.77	101.26	337.64	427.00	2320.37

Note. This table reports the summary statistics of the firm-level repeat-sale sample described in Section 3.2. Each panel presents a subsample by property use category: single-family in Panel A, multi-family in Panel B, commercial in Panel C, and vacant in Panel D. Within each panel, three sample periods are shown: the baseline period, which includes transactions bought before and sold after the first-action date of the associated patent application; the placebo period, which includes transactions bought and sold before; and the post period, which includes transactions bought and sold after. For each sample period, the table reports summary statistics for annualized unlevered returns, buy prices (inflation-adjusted to a 2010 baseline), and sell prices (inflation-adjusted) by first-grant status (**1**(Approved)).

Table A.5 — Summary Statistics of Property-level Repeat Sales Sample

		# Repeat Sales				Annualized Unlevered Return (%)					Buy Price ($1000s)					Sell Price ($1000s)				
Sample	**Holding**	N_1=0	N_1>0	**Avg** N	**Avg** N_1	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD
Panel A. Single-Family																				
Baseline	All	747,433	10,948,664	26.97	17.87	-2.25	0.43	4.63	1.47	7.33	146.85	214.75	322.03	255.61	175.74	146.83	224.85	345.12	274.02	206.81
	≤ 5	257,676	3,500,177	15.06	10.67	-1.79	3.67	10.46	4.13	9.42	147.05	211.42	316.79	253.64	174.78	157.52	235.22	358.66	283.82	201.83
	(5, 10]	288,069	3,940,533	28.49	18.73	-3.34	-0.56	3.33	0.05	6.65	147.51	217.26	326.50	258.81	179.62	138.44	214.70	332.27	264.42	206.79
	(10, 15]	201,688	3,507,954	37.29	24.08	-1.58	0.07	1.94	0.40	4.27	145.61	215.74	322.25	253.94	172.14	145.67	226.35	345.68	275.05	211.28
Placebo	All	372,852	8,779,839	57.29	36.27	-0.68	3.20	8.38	4.05	7.75	140.19	202.89	299.67	241.54	159.30	159.20	233.36	355.48	284.62	196.11
	≤ 5	237,559	5,705,236	55.39	35.11	-0.56	4.05	10.12	4.81	8.47	141.60	202.01	299.18	242.10	159.70	157.19	227.89	343.87	276.04	186.43
	(5, 10]	103,649	2,417,296	61.41	38.90	-1.07	2.27	6.12	2.73	6.24	137.36	202.52	297.83	239.37	158.79	161.30	238.66	366.55	293.23	205.67
	(10, 15]	31,644	657,307	58.62	36.70	-0.14	2.09	4.33	2.35	4.73	137.15	213.08	310.44	244.61	157.57	173.29	272.01	420.75	327.06	231.54
Post	All	191,061	5,273,245	40.07	25.82	0.88	4.10	8.19	4.81	7.04	113.38	174.91	267.60	209.35	162.20	143.62	220.32	328.65	260.83	201.07
	≤ 5	107,017	3,017,977	40.84	26.29	0.40	4.13	9.08	4.91	7.91	117.98	180.00	273.02	214.09	166.39	137.27	209.78	312.46	247.69	197.13
	(5, 10]	76,710	2,078,453	39.25	25.30	1.36	4.06	7.33	4.68	5.72	108.60	169.66	262.17	204.76	156.51	153.91	234.67	347.99	277.49	204.26
	(10, 15]	7,334	176,815	36.61	23.79	1.68	4.18	7.21	4.72	5.35	87.12	150.14	236.00	182.55	151.20	154.03	248.67	367.89	288.71	212.03
Panel B. Multi-Family																				
Baseline	All	162,778	3,979,928	48.94	29.89	-3.20	0.28	6.40	1.80	9.73	125.93	200.90	321.74	272.36	738.65	121.06	205.78	345.08	292.48	875.12
	≤ 5	61,721	1,443,462	30.40	19.53	-2.14	5.32	13.75	5.46	12.32	124.88	200.14	323.65	279.09	828.61	140.04	228.50	368.68	309.58	814.65
	(5, 10]	58,914	1,386,627	55.39	33.28	-5.03	-1.21	3.85	-0.51	8.50	127.86	205.00	331.23	282.09	843.74	110.13	193.59	333.77	289.22	1080.09
	(10, 15]	42,143	1,149,839	64.51	38.79	-2.42	-0.44	1.96	-0.03	4.82	124.85	197.10	308.24	252.07	406.58	114.13	191.94	324.39	274.82	638.86
Placebo	All	80,601	3,011,636	96.82	58.45	-1.31	4.42	11.69	5.31	10.81	116.32	185.26	283.35	239.60	530.46	135.76	214.49	344.83	287.15	585.54
	≤ 5	55,638	1,998,776	93.10	56.17	-0.56	6.28	13.98	6.62	11.70	115.92	183.56	286.39	243.51	616.13	135.47	213.39	341.10	284.58	659.10
	(5, 10]	19,334	764,703	105.20	63.74	-2.35	2.39	8.16	3.02	8.76	113.52	183.79	275.06	227.20	290.11	133.46	211.72	344.27	284.64	381.33
	(10, 15]	5,629	248,157	101.04	60.48	-1.64	1.50	4.45	1.73	5.96	131.42	202.56	286.65	246.27	308.05	145.41	234.65	376.15	315.76	459.03
Post	All	28,253	1,489,018	95.14	57.38	1.33	5.39	10.52	6.24	8.92	80.21	150.52	267.28	246.50	1028.03	114.50	195.65	334.15	324.76	1282.10
	≤ 5	15,559	864,949	96.76	58.39	1.00	6.06	12.22	6.76	10.28	82.44	153.67	272.11	263.87	1122.99	105.02	183.57	321.59	325.51	1355.56
	(5, 10]	11,565	569,934	92.99	56.01	1.63	4.82	8.69	5.55	6.59	78.44	147.96	261.93	224.55	909.29	127.26	209.02	347.87	319.99	1135.55
	(10, 15]	1,129	54,135	91.86	55.62	1.71	4.57	8.15	5.25	5.70	71.44	139.01	247.23	200.63	459.35	143.10	225.77	367.63	363.04	1510.94

Table A.5 — Summary Statistics of Property-level Repeat Sales Sample (continued)

		# Repeat Sales				Annualized Unlevered Return (%)					Buy Price ($1000s)					Sell Price ($1000s)				
Sample	**Holding**	N_1=0	N_1>0	**Avg** N	**Avg** N_1	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD
Panel C. Commercial																				
Baseline	All	13,048	209,186	37.41	23.42	-3.69	0.69	7.53	2.21	11.55	152.84	299.29	599.67	625.00	2492.88	157.60	325.16	674.47	701.44	2804.38
	≤ 5	4,299	63,805	22.63	14.64	-3.82	3.67	14.33	4.74	14.99	155.21	300.07	595.43	618.27	2469.64	167.00	335.47	675.26	688.10	2675.91
	(5, 10]	4,716	76,575	39.45	24.47	-4.79	0.02	6.58	1.18	10.79	153.45	304.07	626.31	672.28	2679.98	153.13	322.58	680.07	734.50	2953.88
	(10, 15]	4,033	68,806	48.94	30.38	-2.77	0.19	4.19	0.99	7.64	149.07	292.42	575.30	578.53	2288.57	153.66	318.54	667.41	677.01	2749.61
Placebo	All	6,918	132,966	67.11	40.96	-2.59	4.45	13.22	5.38	13.61	138.65	263.29	520.31	493.28	1181.94	169.03	328.39	657.72	618.22	1495.92
	≤ 5	4,416	79,824	64.43	39.53	-2.86	5.32	15.81	6.08	15.10	144.35	270.66	533.15	509.31	1349.41	163.25	314.87	628.38	594.40	1612.29
	(5, 10]	1,787	38,918	71.26	43.26	-2.19	4.19	10.91	4.65	11.60	132.29	249.73	492.91	472.20	902.22	176.34	340.93	688.10	642.28	1257.33
	(10, 15]	715	14,224	70.93	42.69	-1.76	2.74	7.51	3.40	8.60	128.50	259.58	511.40	460.40	768.95	188.93	368.00	741.83	686.95	1406.65
Post	All	2,921	82,500	58.64	36.60	-2.03	3.62	11.70	4.48	13.33	105.00	233.42	529.99	718.73	3672.72	126.51	283.26	656.79	917.31	4763.22
	≤ 5	1,587	46,559	62.04	38.57	-3.50	3.64	14.28	4.59	15.29	109.11	248.08	580.47	845.48	4176.29	118.16	275.31	668.20	1076.27	5741.62
	(5, 10]	1,222	32,157	54.63	34.26	-1.12	3.74	9.58	4.45	10.42	99.90	215.96	480.46	556.41	2820.42	134.53	291.66	642.37	709.15	3021.56
	(10, 15]	112	3,784	51.07	32.10	-1.12	2.59	7.13	3.35	8.68	109.05	226.89	467.31	543.18	3407.24	151.80	304.23	649.71	736.35	3336.10
Panel D. Vacant																				
Baseline	All	17,952	145,726	15.74	10.48	-14.85	-1.65	10.89	-0.52	32.05	22.46	68.67	181.15	171.08	439.79	12.91	55.06	179.45	186.59	814.75
	≤ 5	9,899	70,198	9.90	6.87	-24.48	-0.89	22.82	1.07	41.89	21.32	63.90	173.16	165.24	395.08	14.03	54.70	170.23	166.07	422.28
	(5, 10]	4,655	41,679	19.18	12.51	-13.19	-2.30	6.86	-2.22	20.92	23.35	73.16	190.03	184.90	532.60	11.36	55.59	188.20	207.49	842.21
	(10, 15]	3,398	33,849	24.02	15.45	-8.12	-1.54	3.72	-1.83	13.60	24.21	73.41	188.61	166.47	401.37	12.35	55.45	187.86	204.73	1284.35
Placebo	All	11,221	100,607	27.49	17.74	-11.00	3.93	24.77	6.85	39.53	26.39	68.67	164.88	146.95	266.79	23.50	73.43	195.03	179.90	387.82
	≤ 5	9,743	79,265	25.54	16.65	-15.13	4.90	31.01	7.86	43.22	24.70	65.27	159.74	144.80	273.51	22.19	66.69	176.85	165.63	365.72
	(5, 10]	1,060	16,545	36.13	22.36	-4.94	2.81	12.06	3.49	20.07	32.65	82.50	179.58	154.50	241.23	30.88	107.17	259.96	235.04	479.42
	(10, 15]	418	4,797	31.45	19.78	-3.41	1.95	6.96	1.12	13.79	34.77	86.10	188.71	158.15	229.37	30.60	112.46	290.74	237.35	386.62
Post	All	10,490	149,483	24.82	16.52	-14.01	2.35	18.14	2.13	35.90	9.86	39.69	119.85	129.14	534.71	9.82	40.77	132.61	145.41	667.33
	≤ 5	7,332	103,835	25.20	16.78	-22.71	0.79	20.66	0.42	40.48	10.79	42.13	123.88	129.61	378.13	9.20	36.94	121.82	133.70	607.83
	(5, 10]	3,011	43,166	23.97	15.95	-3.62	4.30	15.11	6.15	21.98	8.78	34.88	111.48	129.56	802.22	12.31	49.74	155.96	173.24	802.88
	(10, 15]	147	2,482	23.75	15.63	-4.15	2.94	11.02	4.09	15.11	6.77	31.88	100.00	102.00	217.11	9.72	47.89	157.07	151.46	365.82

Note. This table reports summary statistics for the property-level repeat-sale sample described in Section 3.2, where each observation is a repeat-sale transaction located within 10 miles of at least one patent-applying startup. Each panel presents a property use category: single-family (Panel A), multi-family (Panel B), commercial (Panel C), and vacant (Panel D). Within each panel, rows are grouped by sample period—*Baseline* (bought before and sold after the first-action date of an associated patent application), *Placebo* (bought and sold before), and *Post* (bought and sold after)—and further binned by holding term. *# Repeat Sales* reports the number of transactions with no nearby patent-granted startup (N_1=0) and with at least one (N_1>0). *Avg* N is the mean number of nearby patent-applying startups (granted or not); *Avg* N_1 is the mean number of nearby patent-granted startups, conditional on $N_1 > 0$. Returns are annualized and unlevered; buy and sell prices are inflation-adjusted to a 2010 baseline and reported in thousands of dollars. The sample is restricted to holding terms of at most 15 years.

Table A.6 — Summary Statistics of Local Business Sample

2-digit NAICS	1(Approved)= 0					1(Approved) = 1				
	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD
Panel A. # Establishments within 10 miles (Pre-innovation)										
11	25	40	62	47.09	31.21	22	36	58	42.93	29.94
21	5	12	23	36.47	94.28	4	10	19	32.19	87.59
22	9	18	29	25.09	27.58	7	14	25	20.67	23.86
23	914	1,912	3,218	2,434.14	2,358.58	687	1,574	2,734	2,012.15	2,023.99
31	47	111	245	336.52	809.28	35	88	198	256.26	673.36
32	116	278	518	427.87	544.54	83	226	459	356.90	468.53
33	246	571	1,111	875.62	986.59	185	482	997	751.70	881.37
42	489	1,232	2,402	2,031.02	2,857.03	364	1,010	2,075	1,676.01	2,435.62
44	1,070	2,413	4,710	3,993.74	5,576.52	818	1,965	4,012	3,220.33	4,629.29
45	468	979	1,887	1,597.60	2,216.86	358	821	1,618	1,294.04	1,836.47
48	151	369	708	691.81	1,207.34	111	286	578	531.90	966.16
49	25	48	85	74.02	93.58	20	40	73	60.83	78.21
51	200	501	1,128	986.10	1,642.16	137	384	914	770.17	1,348.58
52	525	1,333	2,524	1,921.24	2,145.05	374	1,028	2,089	1,534.97	1,814.31
53	509	1,194	2,424	1,849.15	2,283.80	369	959	2,008	1,485.36	1,906.01
54	1,129	2,937	6,252	4,510.01	4,951.76	750	2,262	5,062	3,636.28	4,259.71
55	4	10	25	25.62	52.33	3	8	17	18.51	42.10
56	435	1,028	2,023	1,653.93	2,295.84	310	823	1,703	1,310.58	1,881.58
61	245	555	1,061	854.32	1,094.37	179	456	888	690.18	914.59
62	949	2,293	4,708	3,763.55	4,945.62	667	1,800	3,836	2,991.18	4,134.47
71	180	373	707	580.79	760.09	130	309	580	463.04	627.62
72	619	1,392	2,816	2,390.70	3,506.40	460	1,141	2,347	1,906.26	2,885.36
81	1,317	2,825	5,375	4,602.94	6,225.08	988	2,390	4,707	3,739.87	5,180.32
92	240	448	731	604.27	606.43	195	385	645	506.68	517.91
All sectors	10,226	23,241	45,678	36,304.07	46,425.17	7,525	18,773	38,144	29,301.47	38,831.23
Panel B. Δ Establishments within 10 miles (Pre-innovation to 10 Years After)										
11	6	15	29	21.04	25.19	6	15	28	19.60	21.45
21	-1	2	7	0.08	21.81	-1	2	7	1.53	18.22
22	1	6	15	9.89	14.47	2	7	15	10.19	13.16
23	-47	103	517	276.93	719.72	0	180	606	363.94	669.72
31	0	11	29	-36.06	223.94	1	12	29	-17.20	183.05
32	-112	-49	-13	-104.63	187.17	-87	-31	-5	-74.82	150.48
33	-136	-41	-2	-121.82	244.87	-91	-18	6	-81.88	204.17
42	-535	-227	-51	-470.53	764.69	-387	-134	-14	-329.06	617.30
44	-300	-24	192	-92.61	866.58	-111	28	327	82.49	758.70
45	-139	-22	37	-88.40	375.67	-72	0	69	-21.89	304.20
48	0	40	143	53.60	398.21	7	53	165	93.73	324.49
49	-1	3	10	3.94	16.68	0	4	11	5.72	15.08
51	-4	26	112	42.91	207.25	1	34	118	63.43	184.38
52	-1	82	354	276.34	686.51	14	127	404	312.21	610.41
53	12	136	420	392.45	922.76	22	154	435	375.97	806.35
54	-2	201	731	483.48	1,019.23	24	277	838	566.89	955.99
55	6	18	41	40.84	79.31	4	14	32	30.77	63.28
56	-16	68	275	106.31	672.47	7	116	352	204.53	599.86
61	20	92	203	155.55	251.94	20	90	203	146.26	220.70
62	910	2,340	4,528	3,306.07	3,648.03	600	1,792	3,643	2,613.06	3,053.63
71	30	101	227	170.74	230.83	25	91	205	150.33	199.10
72	64	243	534	437.35	749.97	58	224	493	392.44	652.19
81	-91	59	343	53.69	781.88	-23	109	398	168.19	663.99
92	-33	3	36	-9.08	130.86	-10	13	49	13.37	109.09
All sectors	601	2,983	6,913	4,906.60	7,304.80	724	3,214	7,199	5,088.66	6,849.44

Note. This table reports the summary statistics of the firm-level establishment sample described in Section 3.3. Panel A shows the number of establishments within 10 miles of the startup one year prior to the startup's first-action date, by industry sector defined as the first two digits of the NAICS code. Panel B shows the change in the number of establishments from one year prior to ten years after the first-action date. The summary statistics are provided by the startup's first-grant status (**1**(Approved)). The final *All sectors* row of each panel pools all NAICS sectors within ten miles of the startup. NAICS sector 99 (Unclassified) is excluded to match the estimation sample.

Table A.7 — Summary Statistics of Local Employment Sample

2-digit NAICS	1(Approved)= 0					1(Approved) = 1				
	Q1	Q2	Q3	Avg	SD	Q1	Q2	Q3	Avg	SD
Panel A. Employment within 10 miles (Pre-innovation)										
11	115	226	440	340.55	374.08	92	192	406	310.70	372.91
21	30	82	221	459.05	1,744.43	20	63	177	415.15	1,635.96
22	134	427	1,122	986.48	1,442.64	88	310	838	775.65	1,245.52
23	6,395	15,134	25,589	18,915.76	17,372.25	4,516	12,035	22,317	15,646.44	15,307.15
31	791	2,177	5,168	5,605.38	10,685.88	572	1,651	4,120	4,291.13	8,878.07
32	2,132	5,098	9,565	7,863.08	9,435.90	1,416	4,146	8,105	6,434.60	8,145.44
33	5,225	13,471	29,390	20,453.52	20,641.89	3,671	10,875	24,478	17,951.06	20,152.59
42	4,900	13,320	24,674	19,868.51	24,657.37	3,326	10,369	21,181	16,256.91	21,264.49
44	10,268	24,004	45,429	33,999.03	37,946.65	7,204	18,682	37,097	27,261.73	31,977.86
45	4,474	10,434	18,523	13,996.83	15,145.12	3,112	8,115	15,510	11,233.09	12,716.96
48	1,787	5,168	11,201	9,533.58	14,470.56	1,233	3,897	8,702	7,451.95	11,943.28
49	406	1,038	2,305	1,920.77	2,782.30	308	877	1,915	1,597.75	2,400.10
51	2,452	7,419	16,944	14,602.51	21,832.72	1,562	5,430	13,759	11,423.15	18,076.06
52	4,127	13,203	26,561	21,203.07	27,208.43	2,553	9,877	20,653	16,475.43	22,925.02
53	2,867	7,628	15,579	12,451.25	17,013.79	1,939	5,851	12,724	9,706.69	14,056.59
54	6,507	18,932	40,614	30,657.39	35,567.84	3,868	13,970	32,376	24,122.71	30,196.37
55	39	162	583	722.04	1,597.70	19	100	360	528.36	1,379.97
56	3,275	9,354	18,668	13,902.09	16,259.13	2,168	7,128	15,750	11,224.15	13,875.58
61	7,723	17,835	33,449	28,435.11	38,385.84	5,648	14,402	28,312	22,729.41	31,897.37
62	8,944	20,901	39,290	32,045.55	41,071.38	6,378	17,100	33,221	25,677.04	34,210.94
71	1,706	4,308	8,634	7,130.42	9,585.42	1,173	3,235	6,888	5,527.58	7,889.17
72	9,884	23,187	45,614	34,544.89	38,239.37	7,371	18,514	38,089	27,915.44	32,288.57
81	5,994	14,228	27,334	22,910.62	30,310.50	4,318	11,571	23,356	18,210.59	25,102.15
92	5,554	13,331	25,118	20,543.70	24,399.23	4,088	10,506	20,383	16,285.11	20,354.71
All sectors	103,927	253,870	488,965	372,967.97	435,127.07	73,760	200,438	407,047	299,343.78	366,462.85
Panel B. Δ Employment within 10 miles (Pre-innovation to 10 Years After)										
11	-16	45	146	71.48	257.67	-9	44	137	74.63	266.91
21	-13	14	75	140.73	1,004.35	-7	16	76	145.97	954.36
22	-22	101	537	359.74	1,132.63	-2	117	566	376.25	1,029.83
23	-1,501	-26	1,889	679.54	4,926.34	-674	401	2,769	1,300.93	4,574.99
31	-430	65	596	-510.62	3,449.12	-195	136	727	-88.37	2,866.51
32	-1,685	-477	135	-1,086.35	3,246.72	-1,179	-161	418	-510.23	2,816.15
33	-3,970	-520	1,043	-1,413.43	5,951.21	-2,540	-36	1,763	-391.66	5,960.31
42	-3,981	-916	221	-2,960.88	7,057.83	-2,584	-191	671	-1,687.05	5,863.37
44	-1,559	365	3,854	1,739.61	8,583.72	-196	1,278	5,607	3,213.26	8,138.84
45	-36	877	2,979	1,762.73	3,806.99	109	1,214	3,500	2,195.25	3,683.15
48	-245	210	1,368	729.06	2,834.44	-58	314	1,490	980.10	2,791.05
49	-412	-80	32	-433.20	1,671.80	-312	-52	37	-286.53	1,448.93
51	-549	219	2,118	1,130.00	4,640.49	-290	235	1,870	1,088.83	4,256.42
52	-2,809	-206	777	-451.02	6,273.28	-1,542	26	1,255	194.98	5,294.48
53	-110	458	2,288	2,112.00	6,161.00	1	605	2,390	1,992.83	5,164.75
54	292	2,520	7,147	4,962.62	7,515.03	339	2,551	7,076	4,816.07	6,849.97
55	-1	66	306	445.81	1,999.34	0	57	260	345.74	1,641.09
56	-764	164	1,713	815.34	4,500.14	-270	386	2,026	1,142.89	4,149.99
61	-840	666	3,393	1,742.84	7,187.11	-103	1,105	4,024	2,386.31	6,685.52
62	2,019	5,866	11,568	8,625.10	10,326.17	1,412	4,947	10,096	7,355.38	9,001.58
71	444	1,794	4,197	3,105.02	4,640.83	329	1,507	3,700	2,779.83	4,488.44
72	228	1,978	5,455	3,991.25	7,649.51	388	2,245	5,894	4,251.95	7,154.09
81	56	1,090	3,494	1,996.55	4,088.82	148	1,222	3,486	2,171.00	3,832.17
92	-2,087	57	2,271	-94.14	6,397.92	-762	455	3,014	1,000.79	5,737.98
All sectors	-910	11,391	39,908	27,447.26	59,483.42	1,496	17,645	50,467	34,832.76	58,164.33

Note. This table reports the summary statistics of the firm-level employment sample described in Section 3.3. Panel A shows total employment within 10 miles of the startup one year prior to the startup's first-action date, by industry sector defined as the first two digits of the NAICS code. Panel B shows the change in employment from one year prior to ten years after the first-action date. The summary statistics are provided by the startup's first-grant status (**1**(Approved)). The final *All sectors* row of each panel pools all NAICS sectors within ten miles of the startup. NAICS sector 99 (Unclassified) is excluded to match the estimation sample.

Figure A.3 — Annualized returns: unadjusted series and robustness checks

(a) Unadjusted annualized returns

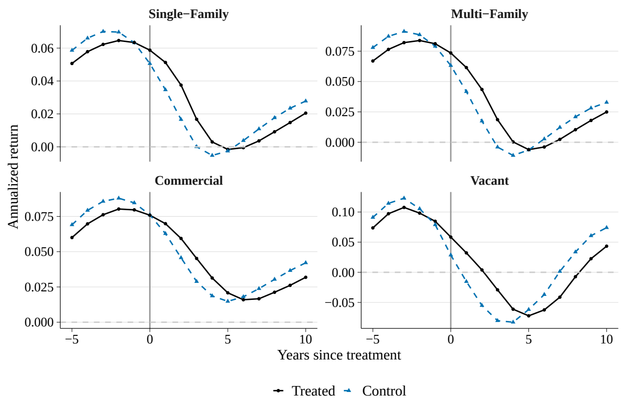

(b) Calendar-adjusted returns, excluding 2008 and 2009

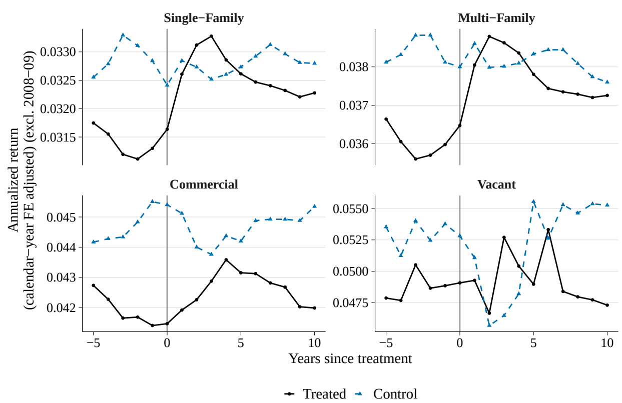

(c) Calendar-adjusted returns, 5-mile radius

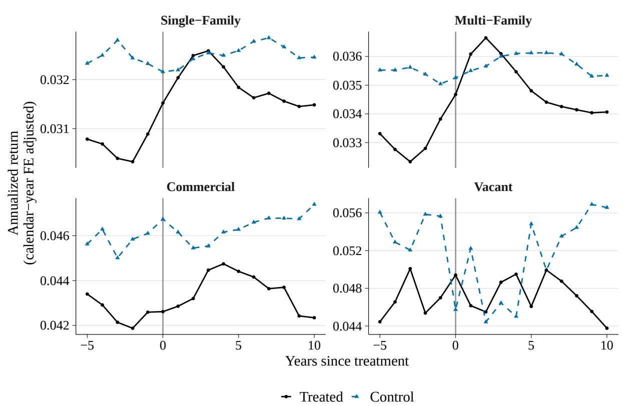

Notes: Each panel plots transaction-weighted average annualized returns against the year of sale relative to the firm's first-action (FA) date, separately for startups whose first application was eventually granted (treated, solid line) and rejected (control, dashed line), with facets by property type. Panel (a) shows the unadjusted series, without the calendar-year adjustment used in Figure 1a. Panel (b) applies the calendar-year adjustment after excluding sales in 2008 and 2009. Panel (c) applies the calendar-year adjustment to repeat sales within 5 miles of the startup rather than 10.

Table A.8 — IV First-Stage Regression Results

	Single-Family		Multi-Family		Commercial		Vacant	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A. Firm-level — dependent variable: $\mathbf{1}$*(Approved)*								
Examiner leniency	0.6576***	0.6570***	0.6726***	0.6722***	0.6584***	0.6579***	0.6534***	0.6525***
	(0.0236)	(0.0237)	(0.0326)	(0.0326)	(0.0292)	(0.0293)	(0.0273)	(0.0274)
Holding-year FE		✓		✓		✓		✓
Observations	38,101,658	38,101,658	35,986,905	35,986,905	5,237,993	5,237,993	1,595,151	1,595,151
R^2	0.336	0.336	0.380	0.380	0.364	0.365	0.378	0.381
Within R^2	0.061	0.061	0.067	0.067	0.063	0.063	0.062	0.061
Kleibergen–Paap rk F	774.5	770.8	426.7	425.0	508.3	505.2	573.5	565.6
Panel B. Property-level — dependent variable: $\mathbf{1}(N_1 > 0)$								
Avg. leniency	0.5335***	0.5375***	0.3904***	0.3919***	0.5027***	0.5048***	0.6722***	0.6748***
	(0.0341)	(0.0339)	(0.0471)	(0.0471)	(0.0400)	(0.0400)	(0.0807)	(0.0805)
N_i	0.0003***	0.0003***	0.0001***	0.0001***	0.0002***	0.0002***	0.0007***	0.0007***
	(8.19×10^{-5})	(7.78×10^{-5})	(3.26×10^{-5})	(3.22×10^{-5})	(4.87×10^{-5})	(4.86×10^{-5})	(1.38×10^{-4})	(1.35×10^{-4})
Property char's		✓		✓		✓		✓
Observations	11,696,097	11,696,097	4,142,706	4,142,706	222,234	222,234	163,678	163,678
R^2	0.367	0.368	0.389	0.389	0.498	0.498	0.510	0.511
Within R^2	0.069	0.071	0.043	0.043	0.061	0.061	0.104	0.106
Kleibergen–Paap rk F	245.2	250.9	68.8	69.1	157.9	159.3	69.4	70.3

Note. This table reports first-stage estimates from the IV specifications in Equation 5 (Panel A) and Equation 8 (Panel B). The dependent variable is the endogenous treatment instrumented in each design: $\mathbf{1}$(Approved) in Panel A and $\mathbf{1}(N_1 > 0)$ in Panel B. Columns (1), (3), (5), and (7) report the baseline specifications; Columns (2), (4), (6), and (8) add county-by-holding-term fixed effects in Panel A and property characteristics in Panel B. Each design is just-identified with a single excluded instrument, so the reported Kleibergen–Paap F-statistic equals the squared cluster-robust t-statistic on the instrument and coincides with the Montiel-Olea–Pflueger effective F. Standard errors (in parentheses) are clustered at the art-unit level in Panel A and at the county level in Panel B. Significance levels: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Figure A.4 — Robustness Checks

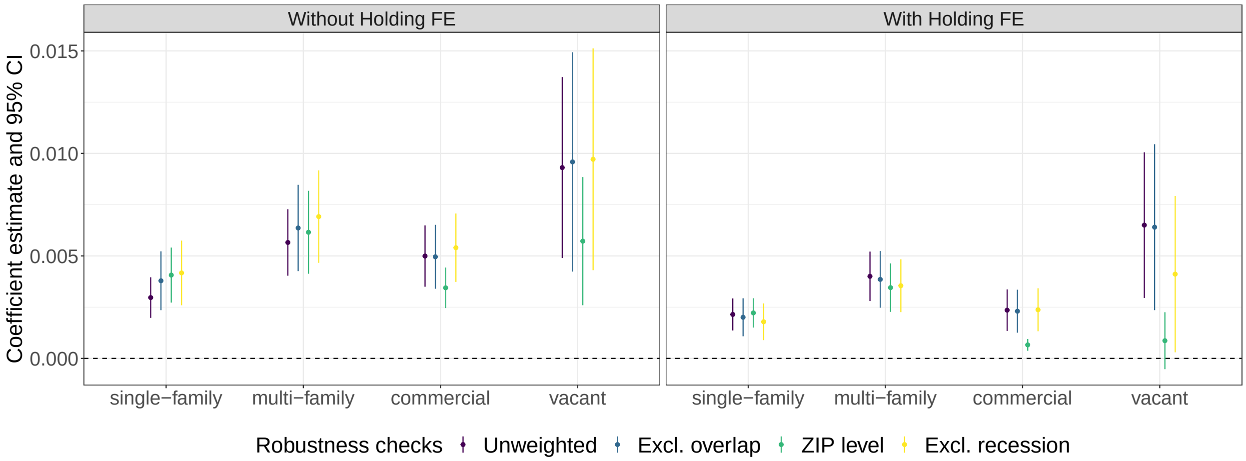

Note. The figure reports coefficient estimates of β_1 from IV regressions in Equation 5 under alternative specifications described in Section 4.2. Estimates are shown separately by property use type. The left subpanel reports the baseline IV specification, and the right subpanel adds holding-term fixed effects. Standard errors are clustered at the art-unit level throughout.

Figure A.5 — Return Premia by Distance (Up to 15 Miles)

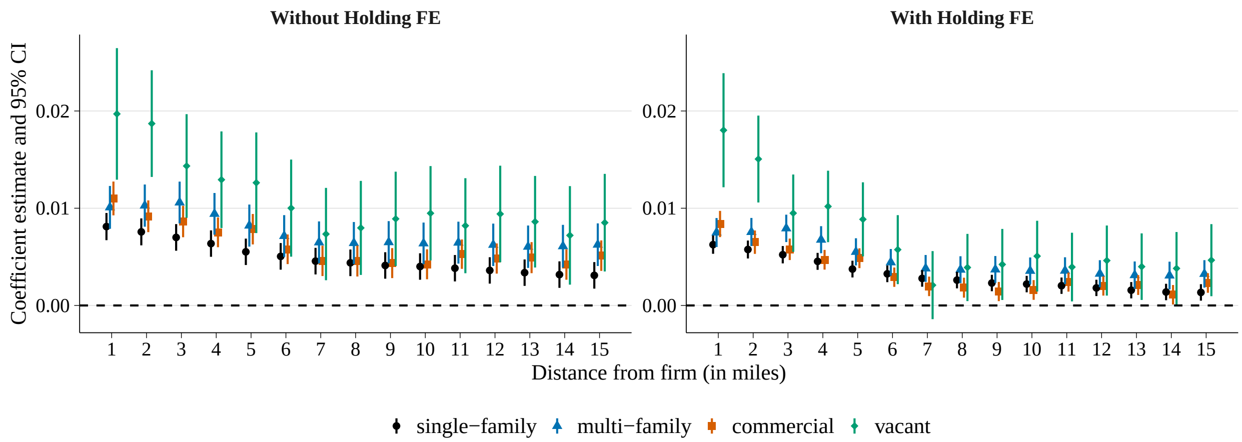

Note. This figure plots IV estimates of the distance-specific return premia $\beta_{1,d'}$ from Equation 7, with 95 percent confidence intervals. Each point corresponds to the estimated return premium at distance bin $d' \in \{1, \ldots, 15\}$ miles from the innovating startup, estimated jointly within a single regression that instruments ten treatment-by-distance interactions with the corresponding examiner-leniency-by-distance interactions. The left panel reports the baseline specification with art-unit-by-application-year, property type, and county-by-buy-year fixed effects; the right panel additionally includes county-by-holding-term fixed effects. Estimates are shown separately by property use type: single-family residential (purple), multi-family residential (blue), commercial (green), and vacant land (yellow). Regressions are weighted by the number of repeat-sale transactions in each cell, and standard errors are clustered at the art-unit level.

Figure A.6 — Knowledge Spillovers - By Distance

(a) Patent Applications (b) Patent Approvals

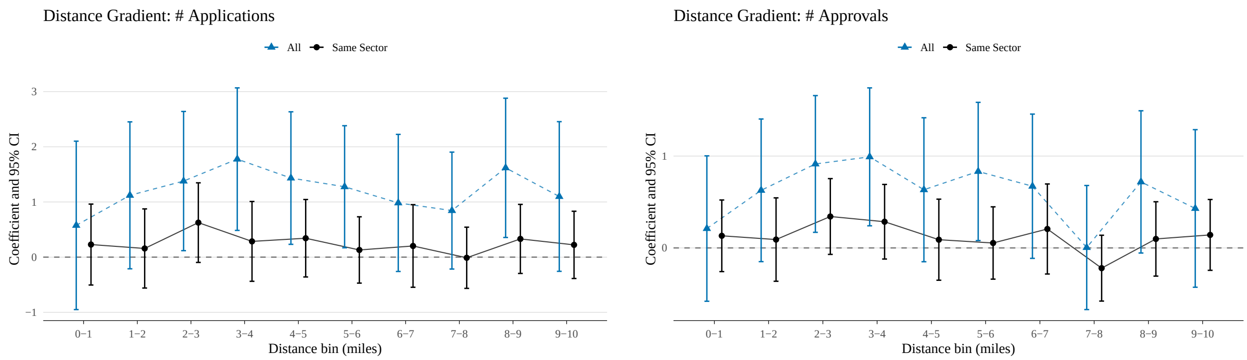

Note. This figure presents the regression results that involves the change in the number of patent applications within a 10-mile distance of the focal granted patent and within a 3-year window after the first-action date of the granted patent. The specifications include art unit-by-application year and county-by-application year fixed effects. The coefficient of the explanatory variable of interest, **1**(Approved), that is binary and depicts if the focal patent application was approved is plotted. Panel (a) depicts the regression results that is based on the change in the number of patent applications irrespective of sector of the focal patent, and the regression results that is based on the change in the number of patent applications by startups that have the same sector as the focal patent. Panel (b) depicts the regression results that is based on the change in the number of approved patent applications irrespective of sector of the focal patent, and the regression results that is based on the change in the number of approved patent applications by startups that have the same sector as the focal patent. Bars denote 95 percent confidence intervals, with standard errors clustered at the art-unit level.

B Spatial Gradient of the Local-Activity Response

The main-text estimates of $\beta_{1,\tau}$ and $\beta_{2,\tau}$ in Equation 10 pool across all cells within ten miles of the applicant. To trace how the per-cell response decays with distance, we re-estimate the same equation separately on each of five 2-mile rings, $d \in \{0–2, 2–4, 4–6, 6–8, 8–10\}$ miles. We also estimate an all-sectors variant of Equation 10 that drops the interaction terms $\widehat{X}_{i,k}$ and $\mathbf{1}(i \in k)$, summarizing the per-cell response without the same- versus different-sector decomposition. In both cases, the estimating equation and instrumental variable are unchanged; only the sample varies.

We begin with the pooled all-sectors estimate. By the tenth year the per-cell effect is about 2.58 additional establishments (Figure A.7a) and about 24.5 additional jobs (Figure A.7b), each significant at the 1 percent level throughout the post-treatment window. Because the all-sectors specification averages the small minority of same-sector cells, where effects are large, against the far more numerous different-sector cells, where they are near zero, its magnitude is dominated by the different-sector cells; the decomposition below makes the underlying asymmetry visible.

Breaking the pooled estimate out by ring, the all-sectors response is roughly flat across distance (Figures A.8a and A.8b): the ten-year establishment estimates cluster between about 2.1 and 3.0 per cell with no sharp gradient and are significant at the 1 percent level in every ring, and the employment estimates range from about 12.5 jobs per cell in the 2–4 mile ring to about 28.8 in the 4–6 mile ring, with four of the five rings significant.

Once decomposed, the same-sector response declines with distance while the different-sector response stays near zero throughout (Figures A.9a and A.9b). For establishments, the same-sector estimate at ten years falls from about 9.1 per cell in the innermost ring (0–2 miles) to about 4.8 in the outermost (8–10 miles), all significant at conventional levels, while the different-sector estimate is indistinguishable from zero in three of five rings and only marginally positive, at roughly 1.3 establishments per cell, in the other two. For employment, the same-sector estimate is positive and significant in every ring and larger in the inner rings than the outer ones, while the different-sector estimate is near zero in four of five rings. The within-industry response is therefore localized but spans several miles, mirroring the distance decay documented for real estate returns in Section 4.2.

Figure A.7 — Pooled All-Sectors Estimates: Local Business Activity

(a) Establishment Counts **(b) Local Employment**

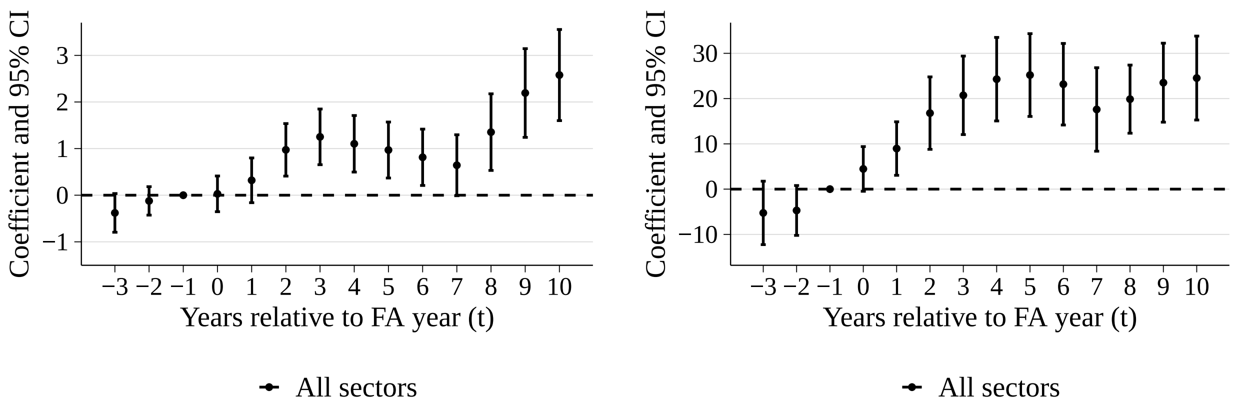

Note. Event-study IV estimates of β_τ from the all-sectors variant of Equation 10 (dropping the within-sector interaction terms), with 95 percent confidence intervals from cluster-robust standard errors at the art-unit level. In Panel (a) the outcome is the change in nearby establishment count from the pre-treatment baseline; in Panel (b) the outcome is redefined as the change in nearby employment.

Figure A.8 — Distance Gradient (All Sectors): Local Business Activity

(a) Establishment Counts

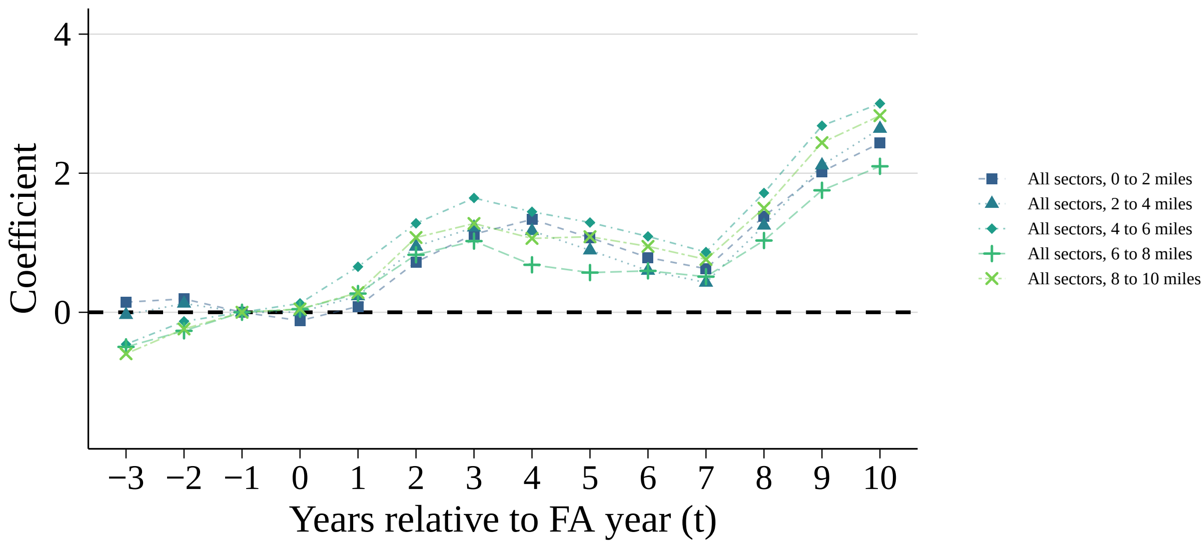

(b) Local Employment

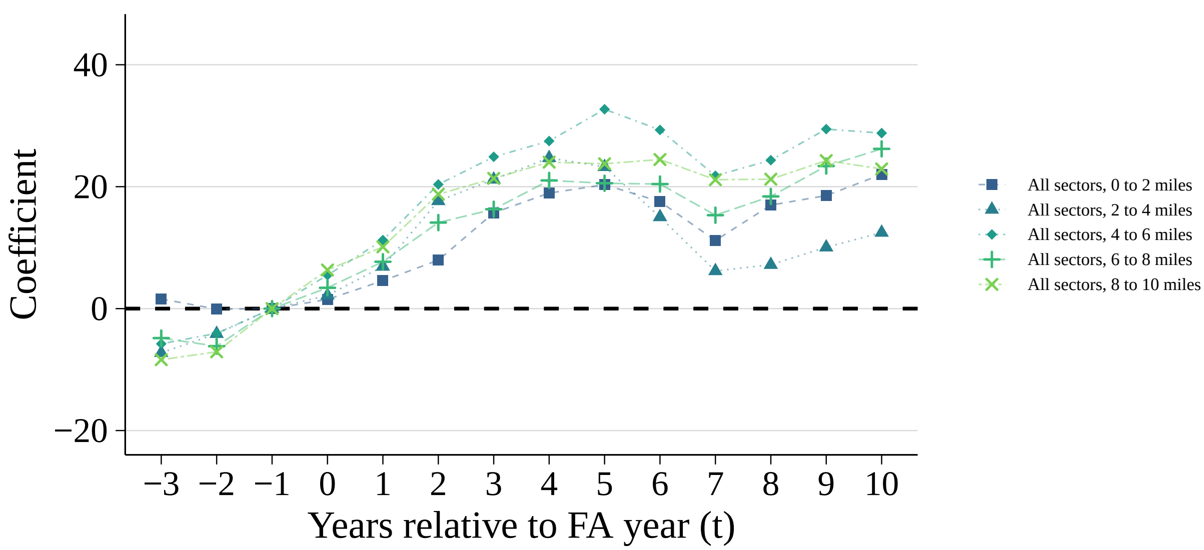

Note. Ring-specific IV estimates of $\beta_\tau^{(d)}$ from the all-sectors variant of Equation 10, estimated separately on each of five 2-mile distance rings $d \in \{0\text{–}2, 2\text{–}4, 4\text{–}6, 6\text{–}8, 8\text{–}10\}$ miles. Standard errors are clustered at the art-unit level. In Panel (a) the outcome is the change in nearby establishment count; in Panel (b) the outcome is redefined as the change in nearby employment.

Figure A.9 — Distance Gradient (Same vs. Different Sector): Local Business Activity

(a) Establishment Counts

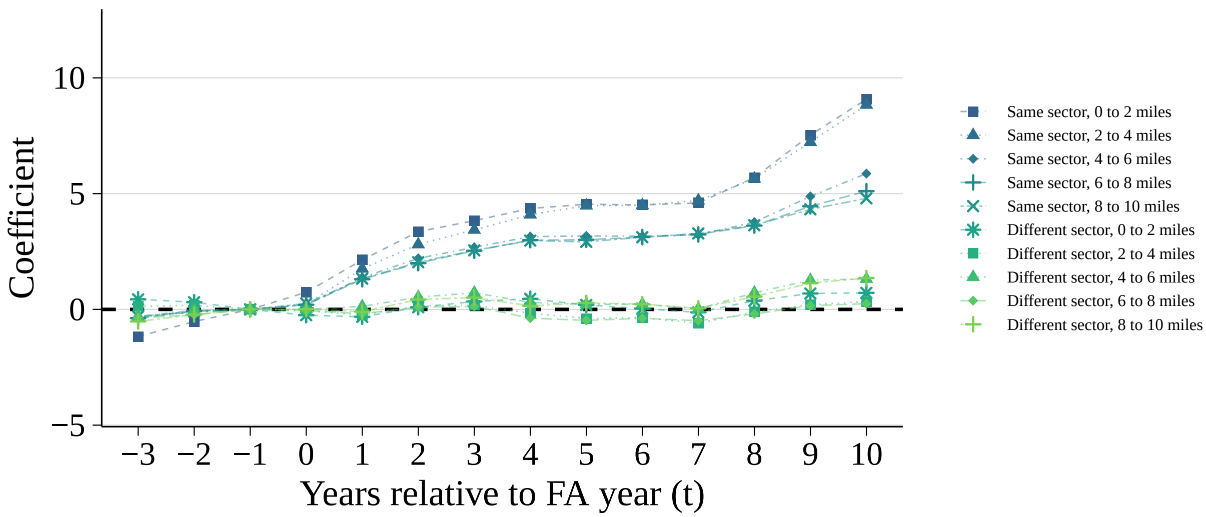

(b) Local Employment

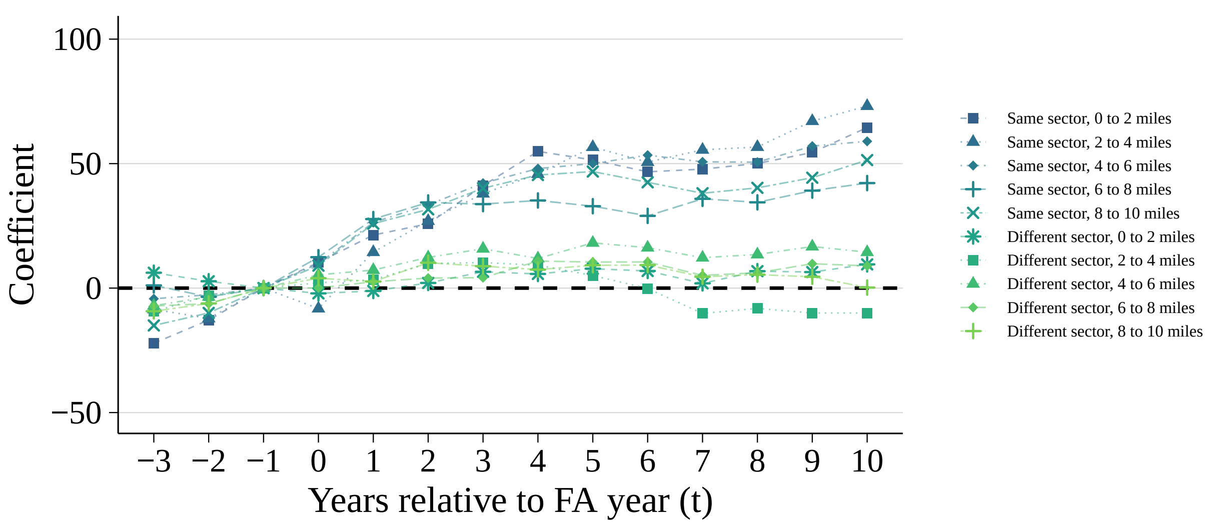

Note. Ring-specific IV estimates of $\beta_{1,\tau}^{(d)}$ (different sector) and $\beta_{2,\tau}^{(d)}$ (same sector) from Equation 10, estimated separately on each of five 2-mile distance rings $d \in \{0\text{–}2, 2\text{–}4, 4\text{–}6, 6\text{–}8, 8\text{–}10\}$ miles. The figures overlay ten series, color-coded so that the within-industry response visually separates from the cross-industry response. Standard errors are clustered at the art-unit level. In Panel (a) the outcome is the change in nearby establishment count; in Panel (b) the outcome is redefined as the change in nearby employment.

C Comparing Data Axle and Zip Code Business Patterns

We find Data Axle provides a reliable measure of local business activity and aligns closely with statistics from the Census Bureau's Zip Code Business Patterns (ZBP) data. At the 3-digit ZIP code level, the correlation between the number of establishments reported by Data Axle and ZBP is remarkably high (0.966), as shown in Panel (a) of Figure A.10. Although Data Axle generally reports a higher number of establishments, this is partly due to its broader coverage. Unlike ZBP, which includes only payroll-based establishments, Data Axle captures a wider range of entities such as farms, railroads, postal services, notaries, and government offices without payroll employees—segments excluded from ZBP. Moreover, we notice that Data Axle assigns business identifiers at the individual professional level (e.g., each lawyer within a firm) that further explains the discrepancy.

The differences between the two datasets become more pronounced among larger establishments, as illustrated in Figures A.10 and A.11, presumably due to the disaggregation mentioned above. Even so, correlations remain strong across size categories, with values of 0.81 for firms with 250–499 employees, 0.68 for those with 500–999 employees, and 0.65 for those with 1,000 or more employees.

A notable limitation of Data Axle lies in its industry coding. As shown in Figure A.12, some records are missing NAICS codes, and the correlation between industry-level establishment counts in Data Axle and ZBP varies considerably, as detailed in Table A.9. These discrepancies are relatively minor in service-oriented sectors but are more pronounced in industries such as Federal Government and Agriculture, Forestry, Fishing, and Hunting. Although caution is warranted when conducting industry-specific analyses, we find Data Axle to be a consistent and comprehensive measure of local business activity, particularly useful for high-resolution spatial and longitudinal analyses.

Figure A.10 — Data Axle vs. Zip Code Business Patterns

(a) All establishments

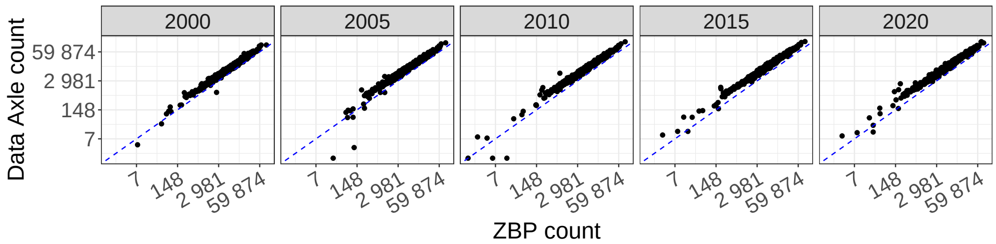

(b) Establishments with 1–4 employees

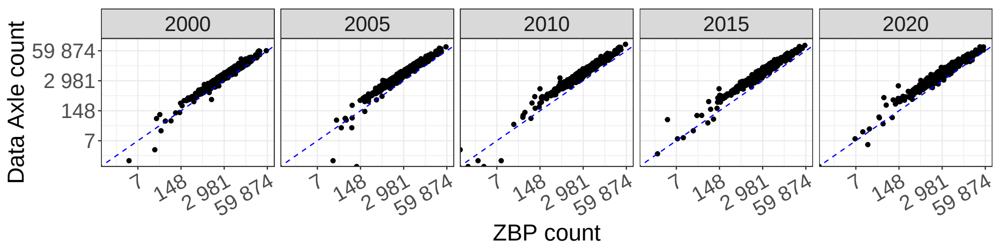

(c) Establishments with 5–9 employees

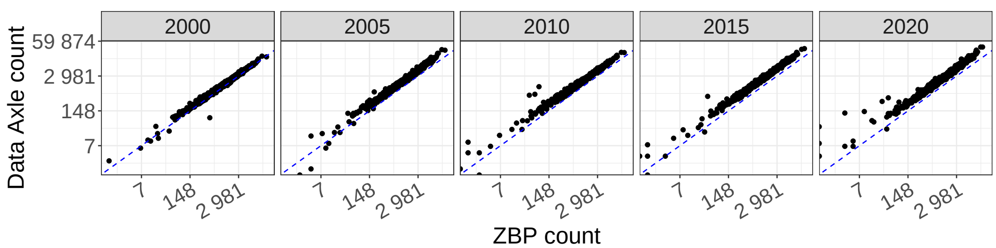

(d) Establishments with 10–19 employees

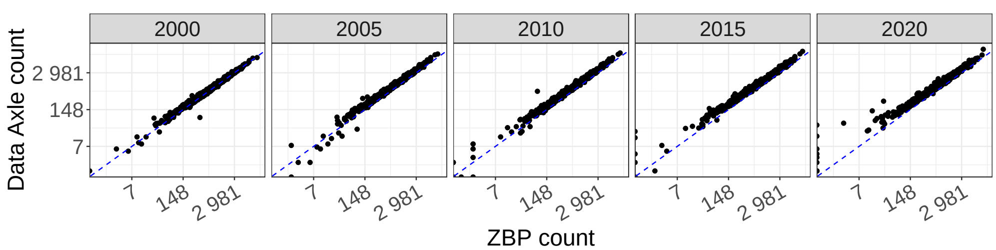

(e) Establishments with 20–49 employees

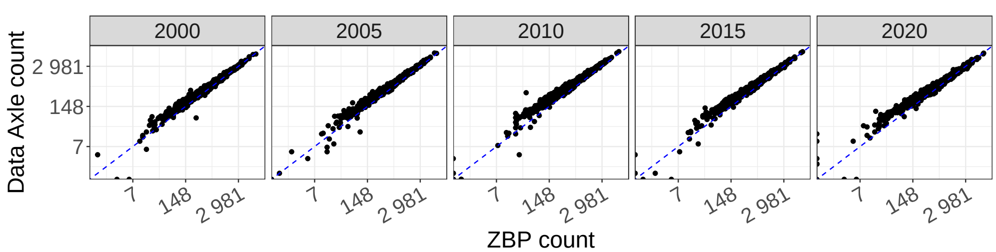

Figure A.11 — Data Axle vs. Zip Code Business Patterns (Cont'd)

(a) Establishments with 50–99 employees

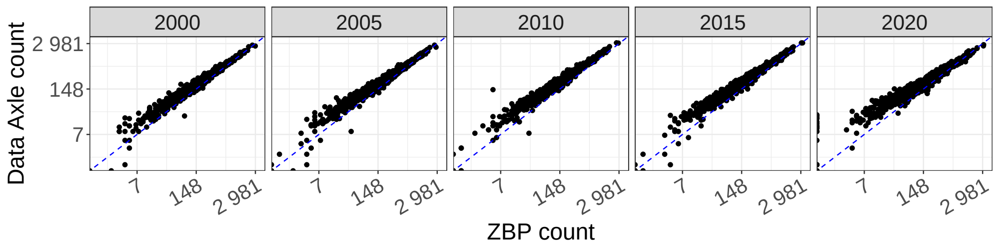

(b) Establishments with 100–249 employees

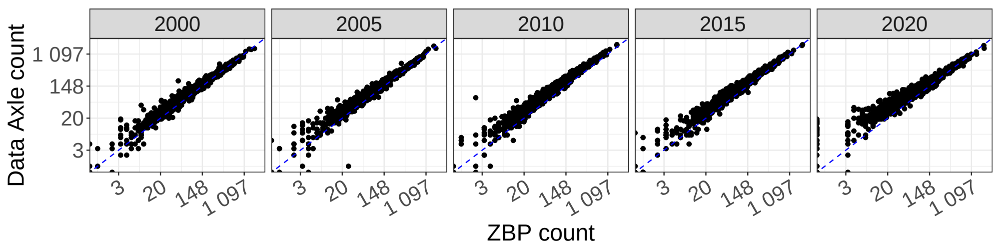

(c) Establishments with 250–499 employees

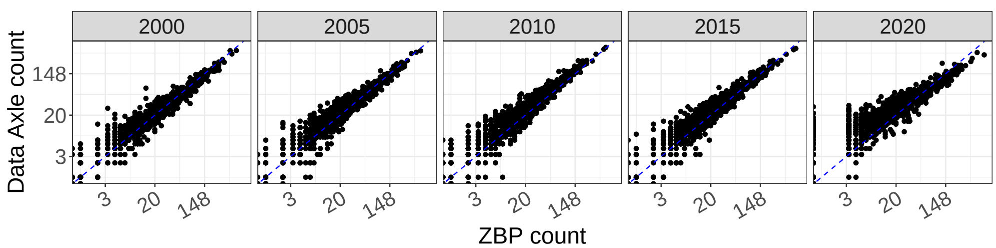

(d) Establishments with 500–999 employees

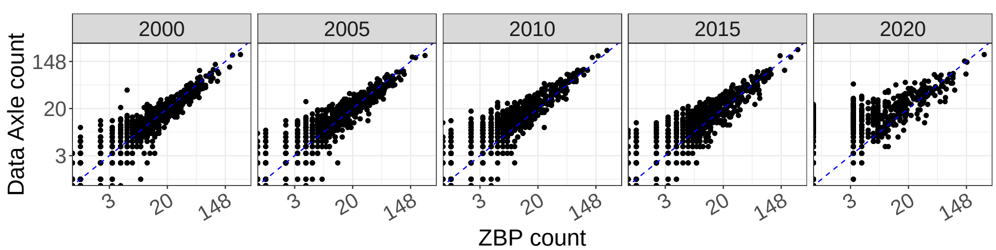

(e) Establishments with ≥ 1000 employees

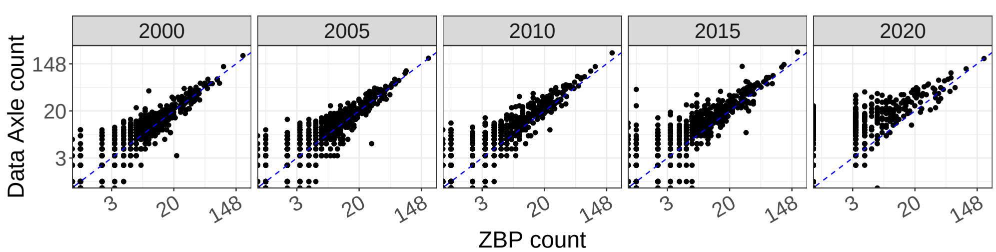

Figure A.12 — Data Axle vs. Zip Code Business Patterns by NAICS filled

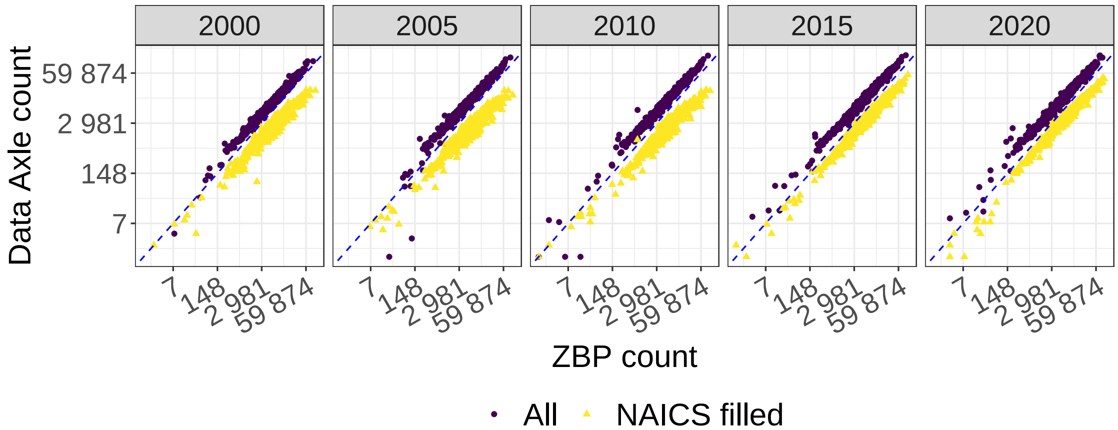

Table A.9 — Zip Code-Year Level Correlation between ZBP and Data Axle by Sector

2-Digit NAICS	Industry Description	Correlation
72—-	Accommodation and Food Services	0.8738
81—-	Other Services (except Public Admin.)	0.8626
42—-	Wholesale Trade	0.8468
56—-	Administrative and Support Services	0.7985
44—-	Retail Trade	0.7981
54—-	Professional, Scientific, and Technical Services	0.7959
62—-	Health Care and Social Assistance	0.7933
52—-	Finance and Insurance	0.7736
23—-	Construction	0.7703
61—-	Educational Services	0.7138
21—-	Mining, Quarrying, and Oil and Gas Extraction	0.7109
51—-	Information	0.7063
53—-	Real Estate and Rental and Leasing	0.6966
31—-	Manufacturing	0.6777
48—-	Transportation and Warehousing	0.5459
71—-	Arts, Entertainment, and Recreation	0.4455
22—-	Utilities	0.3605
55—-	Management of Companies and Enterprises	0.3241
11—-	Agriculture, Forestry, Fishing and Hunting	0.1759
99—-	Unclassified Establishments	0.0692
95—-	Federal Government	0.0055

D Conceptual Framework and Derivations

A patent's first action is a quasi-random shock to a startup's prospects: a favorable action sends the firm onto a higher growth path, and its success becomes visible locally well before the wider activity it seeds. Forward-looking property buyers capitalize that anticipated activity early, while the activity itself arrives only gradually. This appendix states a simple asset-pricing framework and derives every prediction from its primitives, flagging which assumption each prediction requires.

D.1 Setup

Consider the neighborhood around a single startup, and let property buyers and owners value real estate there as a claim on the local rents it will earn. The startup has filed its first patent application, and its prospects hinge on the examiner's review. We index time s in years since the *first-action date*, the date on which the examiner issues the first action on the merits. This action is the informative event: the firm learns the identity of its assigned examiner, and the action itself, much like a revise-and-resubmit decision, resolves most of the remaining uncertainty about whether a patent will eventually issue. The formal grant, when it comes, arrives years later and carries little additional news. We therefore set $s = 0$ at the first action. The property market does not observe the action itself; it responds to the startup's success that follows, which becomes apparent locally soon after and well before the wider activity it seeds. The forward-looking buyers who price this anticipated change are, near a successful startup, disproportionately developers and institutional investors who follow local conditions, and they pay the expected present value of future rents.

For a property p that sits d miles from startup i, this value is

$$V_{p,t} = \mathbb{E}_t\left[\sum_{k\geq 1} \delta^k R_{p,\,t+k}\right], \qquad R_{p,s} = \bar{r} + \phi\, A_{i,s}(d), \tag{11}$$

where $\delta \in (0, 1)$ is the discount factor and $R_{p,s}$ is the rent at time s. Rent has two parts. The first, $\bar{r}$, is a constant baseline that every property earns whatever the startup's outcome. Because $\bar{r}$ is common to startups with and without a favorable first action, it cancels from the price response we study and plays no role below. The second part ties rent to local economic activity $A_{i,s}(d)$ through the capitalization coefficient $\phi > 0$, which measures how much one unit of nearby activity adds to rent.

Local activity responds to the startup's success. Let $G_i \in \{0, 1\}$ record a favorable first action, the quasi-random shock to the startup's prospects. A favorable first action raises nearby activity by $g > 0$ in the long run, working through a much higher probability of an eventual grant and the firm growth that follows; because the first action resolves most of the uncertainty, g is the reduced-form effect of the signal, already net of the chance that the grant does not ultimately issue. The startup's effect reaches a property only in proportion to how close it sits, through a spatial kernel $\rho(d) \in (0, 1]$ that declines with distance ($\rho(0) = 1$ and $\rho'(d) < 0$). Activity at the property is then

$$A_{i,s}(d) = \theta_i + g\, G_i\, m_s\, \rho(d) + \varepsilon_{i,s}, \tag{12}$$

where θ_i is a location baseline that does not depend on the startup's outcome, m_s is the materialization profile, and $\varepsilon_{i,s}$ is mean-zero noise. The profile m_s sets the timing: it is the share of g that has materialized by year s, equal to zero through the first action ($m_s = 0$ for $s \leq 0$), rising over time, and approaching one in the long run. This profile is what separates the timing of prices from the timing of real activity. The key restriction, stated formally in Assumption 5, is that the first action G_i is as good as randomly assigned within an art unit and application year, through the quasi-random assignment of examiners who differ in leniency, so its only route to activity, rents, and prices runs through g.

We study the price response to a favorable first action. Define

$$\Delta P(d) \;\equiv\; \mathbb{E}\Big[V_{p,0} \mid G_i = 1\Big] - \mathbb{E}\Big[V_{p,0} \mid G_i = 0\Big], \tag{13}$$

the gap at the first-action date between the value of a property near a startup with a favorable action and an otherwise identical property near one without. Substituting (11) and (12), the baseline rent $\bar{r}$, the location baseline θ_i, and the noise all cancel, since none of them depends on G_i. Only the success effect remains, which decays with distance and is discounted over the years in which it arrives:

$$\Delta P(d) = \phi\, g\, \rho(d)\, M > 0, \qquad M \equiv \sum_{k \geq 1} \delta^k m_k = \sum_{k \geq 0} \delta^k m_k, \tag{14}$$

where the two sums agree because $m_0 = 0$. The response lands at $s = 0$, when the first action resolves the uncertainty, rather than at the later formal grant: the market capitalizes the expected activity path well before any of it appears.

The constant M collects the timing of the gain. It discounts each future year and weights it by the share of activity that has arrived by then, converting one unit of permanent activity into present value. Slower materialization lowers M but keeps it positive, so $0 < M \leq \delta/(1 - \delta)$, and the upper bound is reached only if the whole gain were present at once ($m_k \equiv 1$ for all $k \geq 1$). For the geometric example $m_k = 1 - \lambda^k$, with $\lambda \in (0, 1)$ setting the speed of arrival (λ near zero is fast, λ near one is slow),

$$M = \sum_{k \geq 1} \delta^k (1 - \lambda^k) = \frac{\delta}{1 - \delta} - \frac{\delta\lambda}{1 - \delta\lambda} = \frac{\delta(1 - \lambda)}{(1 - \delta)(1 - \delta\lambda)}. \tag{15}$$

D.2 Primitives

Six assumptions underlie the framework. The first three are standard or simplifying and only fix functional form; the fourth fixes form as well, but its signs do the work. The last two, a clean signal and rational pricing, are essential, since they are what let the price response identify the grant's effect and capitalize it at the first action.

Assumption 1 (Discounting) *Future rents are discounted at a common factor $\delta \in (0, 1)$. This assumption is simplifying: a constant risk premium folded into δ changes nothing below.*

Assumption 2 (Asset value) *Property value equals the expected present value of future rents, as in (11). This equation is the standard present-value identity and the foundation for everything that follows.*

Assumption 3 (Linear rents) *Rent is linear in local activity, $R_{p,s} = \bar{r} + \phi A_{i,s}(d)$ with $\phi > 0$. Linearity makes the price response exactly proportional to the activity gain (Prediction P2); a concave rent schedule would weaken that proportionality but reverse no prediction.*

Assumption 4 (Activity process) *Local activity follows (12): its effect decays with distance ($\rho'(d) < 0$), a favorable first action raises it ($g > 0$), and the gain builds over time from zero ($m_s = 0$ for $s \leq 0$, non-decreasing, $m_s \uparrow 1$). The functional form is simplifying, but the three signs are essential, since the distance decay produces P3 and the delayed build-up produces the timing in P1 and P6.*

Assumption 5 (Clean signal) *The first action G_i is independent of the location's potential θ_i and the noise ε within an art unit and application year, so its only route to activity, rents, and prices is through g. This is the essential identifying assumption, delivered by the quasi-random assignment of examiners who differ in leniency. Without it the price gap could reflect where successful startups choose to locate rather than their success itself.*

Assumption 6 (Rational expectations) *Buyers price property with correct conditional expectations and update as the startup's success becomes apparent, around the first action ($s = 0$). This assumption is essential: it makes the market capitalize the expected activity at $s = 0$ rather than wait for it to appear, and it is what the matching test in P2 takes to the data.*

D.3 Predictions

The framework yields six predictions. Each follows from the pricing result in (14) and the lemmas below, and each has a counterpart in the paper's estimates. For each we give the statement, the reason it holds, and the evidence.

P1. Front-loaded capitalization.

The price level jumps once, at the first action, so the return premium is front-loaded: it is largest among properties that transact soon after the signal and fades with the time to sale, and it is absent in windows lying entirely before (placebo) or after (post) it. Lemma 1 gives the one-time level gain and Lemma 2 shows that, once the gain is captured, returns are normal; a one-time gain realized as an annualized return is therefore mechanically larger for quick resales and smaller the longer a property is held.[20] In the data the event-time premium peaks one year after the decision, declines over the next few years, and is null in the placebo and post-treatment samples.

P2. Expectations match fundamentals.

The size of the jump is proportional to the expected gain in local activity. This is immediate from $\Delta P(d) = \phi g \rho(d) M$, which is linear in g: a startup whose success will draw more nearby firms and jobs earns a larger premium. P2 is the prediction that separates rational capitalization from a mechanical or sentiment-driven jump, because it ties the cross-section of price responses to the cross-section of *realized* future activity. The matching test regresses each location's price response on its later change in nearby establishments, jobs, and patenting, and the model predicts a positive slope.

[20] The framework delivers the front-loading and the eventual return to zero, not the full year-by-year shape. The decline over the first few years and the slight dip below zero at intermediate horizons in the event-time panel of Figure 3 also reflect transaction selection, which properties sell when, and the supply response of Prediction P6, neither of which the stylized model formalizes. The model's cleanest counterpart is a price-level event study, in which the one-time jump appears directly.

P3. Distance decay.

The premium strictly falls with distance from the startup. A grant's activity reaches nearby property more than distant property through the spatial kernel $\rho(d)$, and differentiating the pricing result gives $\partial \Delta P / \partial d = \phi g M \rho'(d) < 0$. The estimated distance gradient matches this, with most of the decline inside about five miles.

P4. Incidence.

The gain goes to owners who held property across the first action, not to later buyers. Because the repricing happens once, at $s = 0$, the owner of record then captures the jump, while a buyer who arrives afterward pays the already-higher price and earns only the normal return thereafter (Lemma 2). This one-time repricing is why the post-treatment sample, in which a property is both bought and sold after the decision, shows no premium.

P5. Concavity.

The marginal price response falls as the number of nearby successful startups N_1 rises. A neighborhood has finite room to absorb new activity, so the first nearby success does the most to seed a local cluster and each later one adds less. Proposition 1 makes this precise: the activity gain $\Phi(N_1)$ is concave, so the premium $\Delta P = \phi \rho(d) M \, \Phi(N_1)$ rises with N_1 but with a strictly diminishing marginal contribution. The concavity comes from how real activity aggregates, not from any selection in which areas receive grants, so it is consistent with the as-good-as-random first action of Assumption 5. Empirically, the implied per-startup premium falls from about three-quarters of a percentage point at low exposure to roughly one-sixth of a point at high exposure.

P6. Partial reversal.

After the jump the property earns the ordinary return $1/\delta - 1$, so there is no further abnormal return, and the premium can give part of itself back. Capitalization is a one-time level shift (Lemma 2), which on its own leaves later returns normal; if housing supply responds, new construction absorbs part of the added demand and erodes some of the price gain (Proposition 2). The event-time premium behaves this way, fading toward zero within a few years and dipping slightly below it. The sharper test is that premia should be larger and more persistent where local supply is inelastic.

D.4 Pricing and immediate capitalization

The framework rests on two results. The first sizes the price jump at the first action: the market pays today for activity that arrives only over many years, so the jump equals the present value of the expected activity gain, faded by distance. The second shows that the jump is a one-time event, after which the property is again an ordinary asset that earns the discount rate and no abnormal return. Together they place the whole of the grant's price effect at $s = 0$.

Lemma 1 (Capitalization) *Under Assumptions 1–6, the price response defined in* (13) *equals* $\Delta P(d) = \phi\, g\, \rho(d)\, M$*, with* M *as in* (14).

The price gap reads directly off its parts. A favorable first action raises long-run local activity by g; each unit of activity adds ϕ to rent; distance scales the effect by $\rho(d)$; and M discounts the gain over the years in which it arrives. Their product $\phi g \rho(d) M$ is what a forward-looking buyer pays today for that future stream. The proof differences the present-value price across the two values of the signal and uses the clean-signal assumption to clear away every term that does not depend on the grant.

Proof. Take the price at the first action from (11) and difference it across the two values of the signal. The price is a discounted sum of expected rents, and the baseline rent $\bar{r}$ is identical in the two cases, so it drops out:

$$\Delta P(d) = \mathbb{E}[V_{p,0} \mid G_i{=}1] - \mathbb{E}[V_{p,0} \mid G_i{=}0] = \phi \sum_{k\geq 1} \delta^k \Big(\mathbb{E}[A_{i,k}(d) \mid G_i{=}1] - \mathbb{E}[A_{i,k}(d) \mid G_i{=}0]\Big).$$

From the activity process (12), $A_{i,k}(d) = \theta_i + g\, G_i\, m_k\, \rho(d) + \varepsilon_{i,k}$. Assumption 5 makes θ_i and ε mean-independent of G_i, so their conditional means are the same under $G_i = 1$ and $G_i = 0$ and cancel in the difference. Only the grant term survives, with G_i moving from 0 to 1:

$$\mathbb{E}[A_{i,k}(d) \mid G_i{=}1] - \mathbb{E}[A_{i,k}(d) \mid G_i{=}0] = g\, m_k\, \rho(d).$$

Substituting this into the sum and factoring out $\rho(d)$,

$$\Delta P(d) = \phi \sum_{k\geq 1} \delta^k\, g\, m_k\, \rho(d) = \phi\, g\, \rho(d) \sum_{k\geq 1} \delta^k m_k = \phi\, g\, \rho(d)\, M,$$

where $M = \sum_{k\geq 1} \delta^k m_k = \sum_{k\geq 0} \delta^k m_k$ because $m_0 = 0$. Every factor is positive, so $\Delta P(d) > 0$; the closed form for the example $m_k = 1 - \lambda^k$ is (15). □

Lemma 2 (Normal returns after the jump) *Let $P_t = \sum_{k\geq 1} \delta^k \mathbb{E}_t R_{t+k}$ be the price at any date t. Then $\mathbb{E}_t[(R_{t+1} + P_{t+1})/P_t] = 1/\delta$, a constant, so the expected net return is $1/\delta - 1$ in every period.*

The ratio $(R_{t+1} + P_{t+1})/P_t$ is the one-year holding return on the property: a buyer pays P_t, collects the rent R_{t+1}, and ends the year owning an asset worth P_{t+1}, so income plus resale value over purchase price is the return per dollar invested. The lemma says this return equals the discount rate $1/\delta - 1$ in expectation, at every date and whether or not the startup succeeded, because a price set equal to the discounted value of expected rents earns exactly its discount rate going forward. The grant shifts the level of rents, and hence the level of the price, but leaves the expected return unchanged. This constancy is why the premium appears once, as the jump in Lemma 1, rather than as an ongoing abnormal return, and it gives the normal-return benchmark behind the placebo, post-period, and incidence predictions.

Proof. Separate the first rent from the rest of the price and apply the law of iterated expectations to the continuation value:

$$\delta\, \mathbb{E}_t[R_{t+1} + P_{t+1}] = \delta\, \mathbb{E}_t R_{t+1} + \sum_{k\geq 1} \delta^{k+1} \mathbb{E}_t R_{t+1+k} = \sum_{j\geq 1} \delta^j \mathbb{E}_t R_{t+j} = P_t,$$

where the middle equality uses $\mathbb{E}_t \mathbb{E}_{t+1}[\cdot] = \mathbb{E}_t[\cdot]$ and the last re-indexes $j = k + 1$ and re-collects the first rent. Dividing by δP_t gives the gross return $1/\delta$. This value is the same whether or not the property sits near a granted startup, so beyond the one-time revision at $s = 0$ there is no abnormal return. □

D.5 Proofs of Predictions P1–P4

We take the four predictions in turn; each follows from the two lemmas above.

P1 (front-loaded capitalization). Proof. Lemma 1 shows that the price near a favorable first action is higher by $\Delta P(d) > 0$, and that this gap opens at $s = 0$, the date the signal arrives. A repeat sale bought before and sold after $s = 0$ brackets the jump, so it realizes $\Delta P(d)$ as part of its return and shows a positive abnormal return. A repeat sale lying entirely before $s = 0$ (the placebo window) or entirely after it (the post window) does not bracket the jump; by Lemma 2 it earns only the ordinary return $1/\delta - 1$, which is the same near granted and rejected startups, so its abnormal return is zero. The premium therefore appears only in windows that span $s = 0$. □

P2 (expectations match fundamentals). Proof. Lemma 1 gives $\Delta P(d) = \phi\rho(d)M \cdot g$, which is linear in g, the size of the activity gain the grant will bring. A location whose grant produces a larger gain therefore has a proportionally larger price response. Under rational expectations (Assumption 6) the expected gain equals the realized gain on average, so the cross-section of price responses lines up with the cross-section of realized activity. This alignment is exactly what the matching test checks. □

P3 (distance decay). Proof. Differentiate the capitalization result with respect to distance: $\partial \Delta P(d)/\partial d = \phi g M \rho'(d)$. Since ϕ, g, and M are positive and $\rho'(d) < 0$ by Assumption 4, the derivative is negative, so $\Delta P(d)$ strictly declines with distance. □

P4 (incidence). Proof. The premium is a one-time level shift at $s = 0$ (Lemma 1). The owner of record at $s = 0$ holds the property through the repricing and captures the gain. A buyer at any $t > 0$ instead pays the price P_t, which already embeds the gain, and from then on earns only the ordinary return $1/\delta - 1$ (Lemma 2). The gain therefore goes to pre-signal owners, not to later buyers. □

D.6 Concavity (P5), the conditional prediction

The earlier sections follow a single startup. P5 asks what happens when several nearby startups succeed, so the question is how their effects combine. A neighborhood has only so much room to absorb new activity: finite land, a finite local labor pool, and infrastructure that fills up. The first nearby success does the most, because it seeds a local cluster where none was forming; once a cluster is under way, each further success adds less. The activity gain is therefore concave in the number of nearby successes, and because the price tracks expected activity, the premium per startup falls as successes accumulate. This concavity is a property of how real activity aggregates, not of any tendency for grants to land in better areas, so it sits comfortably with the as-good-as-random first action of Assumption 5.

We extend the activity process to a count of nearby successes. Let N_1 be the number of successful startups within the relevant distance, and let the long-run activity gain they generate be $\Phi(N_1)$, with $\Phi(0) = 0$, Φ increasing ($\Phi' > 0$), and Φ concave ($\Phi'' < 0$). This function replaces the single-success term: the gain g of the earlier sections is the first increment, $g = \Phi(1)$. The pricing argument of Lemma 1 is linear in the activity gain, so it carries over with g replaced by $\Phi(N_1)$.

Proposition 1 (Concavity) *Under Assumptions 1–6 and concave Φ, the premium with N_1 nearby successes is*

$$\Delta P(N_1, d) = \phi\, \rho(d)\, M\, \Phi(N_1),$$

increasing and strictly concave in N_1*. The marginal premium of the nth success,* $\phi\rho(d)M\,[\Phi(n) - \Phi(n-1)]$*, strictly declines in n, and the per-startup premium* $\Delta P(N_1, d)/N_1$ *declines in* N_1.

The result reads directly off the shape of Φ. Concavity means each extra success raises activity by less than the one before, so the increment $\Phi(n) - \Phi(n-1)$ shrinks as n grows. The first success therefore carries the largest premium and later ones carry less, which is the diminishing per-startup effect in the data. A convenient example is $\Phi(N_1) = \bar{A}\,(1 - \lambda^{N_1})$ with $\lambda \in (0, 1)$, in which activity approaches a ceiling $\bar{A}$ and the marginal contribution of the nth success is $\bar{A}\,(1 - \lambda)\lambda^{n-1}$, falling geometrically; with $\lambda = \frac{1}{2}$ each success adds half as much as the one before.

Proof. Replace the single-success gain g in Lemma 1 by $\Phi(N_1)$. Because the price is linear in the activity gain, $\Delta P(N_1, d) = \phi\rho(d)M\,\Phi(N_1)$, and the positive factor $\phi\rho(d)M$ makes ΔP inherit the monotonicity and concavity of Φ. The marginal premium is $\phi\rho(d)M\,[\Phi(n) - \Phi(n-1)]$, and concavity gives $\Phi(n) - \Phi(n-1) < \Phi(n-1) - \Phi(n-2)$, so it declines in n. Concavity together with $\Phi(0) = 0$ implies that $\Phi(N_1)/N_1$ decreases in N_1, which is the per-startup premium. □

Remark 1 *P5 is conditional on the one added assumption, concavity of* Φ*. With constant returns (*Φ *linear) the premium would rise in proportion to* N_1 *and the per-startup effect would be flat; with increasing returns it would rise. The estimates show the per-startup premium falling from about* 0.75 *to* 0.17 *percentage points as nearby successes accumulate, which identifies* Φ *as strictly concave. Diminishing returns of this kind follow from finite local capacity, and they require nothing about which startups succeed, so the prediction leaves the identification undisturbed.*

D.7 Partial reversal (P6), the structure-dependent prediction

Lemma 2 settles what happens to returns once the price has jumped, but it holds the housing stock fixed. P6 asks whether the jump itself lasts, and the answer turns on supply. If the local stock of housing cannot grow, the higher rents the grant brings are capitalized once and stay capitalized, so the jump is permanent. If builders can respond, the high prices that follow the grant make new construction pay, the added units compete for the same tenants, and part of the initial gain is competed away. The premium then settles below its first-day peak, a partial reversal rather than a full one.

To see this, let the local housing stock be $H_{i,s}$ and let it weigh on rents, since more units spread the same local demand more thinly. The rent equation extends to $R_{p,s} = \bar{r} + \phi A_{i,s} - \psi H_{i,s}$, with $\psi > 0$ the amount by which one extra nearby unit lowers rent. New units appear only when they pay: builders add to the stock while the price sits above the cost c of constructing a unit, and they do so with a lag, so the stock climbs gradually in the years after the grant. Let $\eta \geq 0$ summarize how strongly supply responds, with $\eta = 0$ a fixed stock and larger η a more elastic market. In the long run the extra stock settles at $\Delta H = \eta\,\Delta$, increasing in the lasting price gain Δ.

Proposition 2 (Partial reversal) *Let $J = \phi g \rho(d) M > 0$ be the fixed-supply price gain of Lemma 1, and let $\beta = \kappa\psi > 0$ be the amount by which one extra unit of local housing lowers the capitalized price, where $\kappa = \delta/(1-\delta)$ is the price-to-rent multiple implied by Lemma 2. With the long-run supply response $\Delta H = \eta\,\Delta$, the gain that survives in the long run is*

$$\Delta = \frac{J}{1+\beta\eta} \in (0, J].$$

It equals J when supply is fixed ($\eta = 0$) and falls toward zero as supply becomes perfectly elastic ($\eta \to \infty$). The attenuation factor $1/(1+\beta\eta)$ strictly decreases in the supply response η and in the rent sensitivity ψ.

The argument is supply and demand in the housing market. The grant shifts demand out and would raise the price by the full J if no new units appeared. The higher price draws in construction, $\Delta H = \eta\Delta$, and each new unit shaves β off the price through the stock term, so the net gain solves $\Delta = J - \beta\,\Delta H$. The more readily builders respond (larger η) or the more an extra unit depresses rents (larger ψ), the more of the jump is handed back, yet as long as supply is not infinitely elastic some premium always survives.

Proof. In the long run the extra stock built in response to a price gain Δ is $\Delta H = \eta\,\Delta$, and each unit of extra stock lowers the capitalized price by β, so the surviving gain is $\Delta = J - \beta\,\Delta H = J - \beta\eta\,\Delta$. Collecting terms gives $\Delta(1+\beta\eta) = J$, hence $\Delta = J/(1+\beta\eta)$. With $\beta, \eta \geq 0$ the factor $1/(1+\beta\eta)$ lies in $(0, 1]$, equals 1 at $\eta = 0$, and decreases in both η and ψ through $\beta = \kappa\psi$. □

Remark 2 *The direction of P6 is not in doubt: under any elastic supply the long-run premium stays positive but smaller than the first-day jump. Two things are not pinned down, which is why P6 is structure-dependent rather than a clean universal result. First, the size of the reversal depends on the local parameters (η, ψ, c), so it differs across markets. Second, the path to the long run depends on construction lags. Prices jump at $s = 0$ because buyers look forward, but building takes time; if the stock overshoots its long-run level before settling, the premium can dip briefly below where it eventually rests, even a little below zero, before recovering. This transitory dip is the small negative stretch in the event-time premium at intermediate horizons. Because its size and timing hinge on the lag structure, the model delivers it only qualitatively. The clean, testable* *content is cross-sectional: premia should be larger and more persistent where local supply is inelastic, that is, where η is small.*

D.8 Assumption ledger

This ledger records, for each prediction, what it rests on and how secure it is. Every prediction shares a common *core*: Assumptions A1, A2, A5, and A6, together with the two sign restrictions $\phi > 0$ (part of A3) and $g > 0$ (part of A4). The core is what delivers the causal one-time price jump $\Delta P(d) = \phi g \rho(d) M > 0$ at the first action. Beyond the core, each prediction adds at most one further ingredient, shown in the middle column.

The status grades how safe each prediction is. *Unconditional* means it follows from the core, using at most a sign or timing restriction already contained in A3 or A4, so it holds unless the asset-pricing setup itself is wrong. *Conditional* means it needs one substantive extra assumption that the data must bear out. *Structure-dependent* means it needs added structure whose size depends on local conditions, so only its direction is guaranteed.

Prediction	Beyond the core, also needs	Status
P1 front-loaded capitalization	gradual build-up, $m_s = 0$ for $s \leq 0$ (A4)	Unconditional
P2 expectations $\propto$ fundamentals	linear rent (A3)	Unconditional
P3 distance decay	$\rho' < 0$ in the kernel (A4)	Unconditional
P4 incidence	nothing beyond the core	Unconditional
P5 concavity	concave Φ	Conditional
P6 partial reversal	elastic, lagged supply	Structure-dependent

Read from top to bottom, the predictions weaken in one direction only. The first four stand on the core: P1 adds the build-up timing already in A4, namely that the activity effect is zero until the first action and then arrives gradually; P2 adds the linearity already in A3; and P3 adds the downward slope of the distance kernel in A4, while P4 needs nothing more. P5 then adds a single assumption the estimates confirm, a concave Φ. P6 adds a supply block whose direction is firm but whose size depends on local conditions. The four core predictions are the ones to stake the paper on; the last two are where a skeptic has room to push, and this table marks exactly where.

D.9 Mapping to the evidence

This section connects the model to the paper one feature at a time. The model is a single mechanism: a quasi-random success signal that the market capitalizes at once into a claim on local activity that then arrives only gradually. Each row of the table takes one implication of that mechanism, states it in the middle column, and points to the specific estimate that bears it out in the right column. The first and last rows, magnitude and timing, are overall features of the mechanism; the six in between are the numbered predictions. To read a row, take P1: the model says the price jumps once at the signal and then earns ordinary returns, so in event time the return premium should appear right after the decision and be absent both before and after. That is what the estimates show, a peak one year out with null placebo and post samples.

	What the model implies	Where the paper shows it
Magnitude	A positive one-time capitalization, $\Delta P(d) = \phi\, g\, \rho(d)\, M > 0$	A grant raises nearby annualized returns by 0.4 to 0.9 percentage points across property types (Table 2, Panel A.1).
P1	Front-loaded: the premium appears at the signal, not before or after	The event-time premium peaks at $\tau = 1$ and declines over the next few years (Figure 3, panel b); the placebo and post samples are null (Table 2, Panel B).
P2	Fundamentals: the response is proportional to the activity gain	The anticipated activity does appear (Figure 5); the cross-sectional proportionality test is the natural next step.
P3	Distance decay: the premium falls with distance	The premium falls monotonically with distance, with most of the decay inside five miles (Figure 3, panel a).
P4	Incidence: the gain goes to owners holding across the signal	The post-treatment sample, bought and sold after the decision, shows no premium (Table 2, Panel B.2).
P5	Concavity: the marginal effect falls with the number of nearby successes	The per-startup premium falls about fourfold across exposure quartiles; for single-family homes, from 0.75 to 0.17 percentage points (Table 3, Panel B).
P6	Reversal: normal returns after the jump, partial reversal under elastic supply	The premium fades to near zero by about $\tau = 4$ and dips slightly below before recovering (Figure 3, panel b).
Timing	Prices jump at the signal while activity builds as $m_s \to 1$	Returns react at $\tau = 1$, while nearby establishments and jobs build over the following decade (Figure 5).

Three things keep this an honest map rather than a perfect one. First, the model's clean object is the price level, while the paper's outcome is an annualized return that also moves through shorter holding periods and lower purchase prices; the capitalization in Lemma 1 is the price-growth part of that return. Second, P2 is the one prediction without a direct test yet: the paper shows that the anticipated activity appears, but the sharper check, that larger price responses occur where more activity later appears, remains to be run. Third, the reversal in P6 is real but modest: the event-time premium dips only slightly below zero at intermediate horizons, and the dip is clearest for vacant land, so we read it as suggestive of the supply channel rather than decisive.

D.10 Falsification

Each prediction is refutable, and the design makes the relevant test concrete. P1 fails if a pre-signal (placebo) abnormal return appears, since the premium should open only at the first action. P2 fails if the cross-section of price responses is uncorrelated with realized future activity; this is the central test that separates rational capitalization from sentiment. P3 fails if the premium does not fall with distance. P4 fails if the post-treatment sample, bought and sold after the decision, carries a premium. P5 fails if the per-startup effect rises rather than falls with exposure. P6 fails if the premium reverses fully and permanently, which would point to over-reaction rather than capitalization with a supply response. The timing claim that prices lead activity fails if nearby activity rises before or together with the price jump rather than after it.