Acquiring to Hire: Team Complementarity and Innovation

Zunda Winston Xu*

Stanford Graduate School of Business

August 4, 2026

Abstract

Firms sometimes acquire companies primarily for their teams, even though they can hire workers directly. Why is acquisition preferred, and what follows for innovation and knowledge diffusion? To answer these questions, I use large language models to identify 12,271 acqui-hires from transaction text and link them to worker and patent histories. I document that acqui-hires move cohesive teams into the acquirer, keep them together, and combine team knowledge with the acquirer's technology. These facts motivate a framework in which acquisition can make complementary workers' participation jointly contingent and transfer residual control over target-held technology that complements team-specific capital. I isolate team complementarity using the COVID-19 disruption to in-person work, which made coordination harder for interaction-intensive teams. Greater exposure shifts firms away from acqui-hires and erodes acquisition's relative advantage in keeping teams together and supporting joint production. Team complementarity is therefore an advantage of acquisition that also supports knowledge diffusion to other acquirer inventors. After acquisition, the team's innovation shifts toward the acquirer's technology, and closer alignment predicts less external follow-on per patent. Workers later leave, and their outside patents remain related to earlier team production. Antitrust review must therefore consider project direction and external knowledge diffusion alongside patent counts and talent concentration.

Keywords: acqui-hires, mergers and acquisitions, team complementarity, firm boundaries, human capital, corporate innovation

*Stanford Graduate School of Business. Email: zundaxu@stanford.edu. I am especially grateful to my main advisor, Amit Seru, and to Nicholas Bloom, Greg Buchak, and Vikrant Vig for their valuable guidance and feedback. I also thank Rafael Berriel, Michael Blank, Harry Cooperman, Ege Ercan, Benjamin Hébert, Naz Koont, Tianshu Lyu, James Paron, Josh Lerner, Dimitris Papanikolaou, Tim de Silva, Lulu Wang, Oliver Xie, Chun Zhao, and participants at the Stanford seminar for their helpful comments. All errors are my own.

1 Introduction

When a firm wants an innovative team, why acquire its employer rather than hire the workers directly? Recent transactions in artificial intelligence expose the choice: OpenAI acquired io and brought its team inside the firm; Microsoft instead hired members of Inflection's core team and obtained a nonexclusive license to Inflection's intellectual property without acquiring the company.[1] An acqui-hire is an acquisition undertaken primarily to obtain the target's team. Yet changing corporate control does not secure the team's participation: each worker decides whether to participate and remains free to leave. Noncompete agreements can raise the cost of direct hiring, but cannot explain what ownership adds when workers can move.[2] What does ownership add?

With complete contracts, it adds nothing to productive surplus: team-contingent employment offers and a transfer of all relevant technology rights reproduce the acquirer's productive allocation. Two frictions create an advantage. First, beyond individual human capital, a coalition can embody team-specific capital—joint productive capacity specific to its membership. This capital generates coworker complementarity: each worker is more productive when the relevant coworkers are present. Sequential bilateral agreements that cannot be conditioned on later workers' acceptance may therefore fragment the team; acquisition can instead condition a common closing on joint participation. Second, hiring the workers does not confer control over target-controlled technology. When future uses cannot all be specified ex ante, ownership assigns residual control over unspecified uses (Grossman and Hart, 1986; Hart and Moore, 1990). The relevant target technology and team-specific capital are complements: each raises the value of the other. Acquisition coordinates participation and control without transferring ownership of labor.

The empirical challenge is that transaction data omit the team motive, productive coalitions are unobserved, and the alternatives considered for each opportunity are unknown.

[1] See OpenAI (2025) and the CMA (2024).

[2] See California Business and Professions Code §16600.

I construct a documentary measure linked to worker and patent histories, and compare acqui-hires with groups of workers from a common source firm who join the same acquirer without an acquisition of that source firm. I call these observations comparison worker flows. They document realized group assembly through hiring but do not reveal whether the offers were coordinated or team-contingent. To isolate the role of coordination, I interact the timing of the 2020 disruption to in-person work with task requirements that make the disruption more binding. The disruption changes the productivity of coordination-dependent work; it is not the subject of the paper.

The evidence supports a dynamic account of the firm boundary. Acqui-hires assemble nearly all of a cohesive group's observed prior joint-mobility links, compared with roughly one-fifth for comparison flows. The group's inventions both continue its pre-transfer trajectory and move toward the acquirer's prior technology. When coordination becomes harder, acqui-hire incidence shifts away from exposed settings. Among groups transferred before the disruption, joint retention and innovative production fall more after acqui-hires than after comparison flows. The boundary later adjusts again through exit and entrepreneurial reallocation. Among workers who subsequently patent outside the acquirer, post-exit inventions are related to other group members' work inside the acquirer before exit. Patents produced inside the acquirer that lie closer to its prior technology receive less nonaffiliate follow-on per patent. Acqui-hire groups also produce more patents per incoming worker, and along observed transfer paths nonaffiliate follow-on per incoming worker is greater than after comparison flows. Acquisition can thus preserve team-specific capital and transfer control over target technology to the acquirer without fixing the eventual allocation of workers or knowledge. The evidence identifies these margins, but not innovation when acquisition is unavailable—the counterfactual relevant for merger review.

Studying this organizational choice requires identifying transactions undertaken to obtain a team—a motive standard M&A databases do not record. I assemble 109,299 qualifying PitchBook acquisitions from 2000 through 2025 and retrieve and classify public transaction

text for 76,341 of them. The documentary measure identifies 12,271 acqui-hires, requiring a source to state explicitly that obtaining the team was a motive. Subsequent employment and patenting do not enter the classification, leaving worker histories available for independent validation: classified transactions bring a greater concentration of high-skill workers into the acquirer. The identified targets are also smaller, younger, more likely to be venture-backed, and more concentrated in information technology. The main worker and patent analyses focus on public acquirers, for which firm, worker, and inventor links are consistently available.

The linked data show that acqui-hires assemble cohesive groups and keep them economically active. Acqui-hire targets contain denser prior worker links than other acquisition targets. Among groups with at least one qualifying prior link, 95.6 percent of those links arrive at the acquirer after an acqui-hire, compared with 21.1 percent after comparison worker flows of similar arriving-group size and observable composition. The differences persist after entry. At 24 months, pairwise co-retention is 4.8 percentage points higher following an acqui-hire, and the conditional expected number of promotions per worker is 57 percent higher. At 48 months, incoming groups produce 183 percent more patents and 238 percent more observed private-value-weighted patent output per incoming worker. Their post-entry patents both continue the group's pre-transfer technological trajectory and move toward the acquirer's pre-transaction technology. These conditional comparisons show what acqui-hires select, assemble, and use; they are not causal effects of acquisition.

To explain these patterns, the framework separates coalition assembly from the allocation of residual control over target technology. Under direct hiring, sequential bilateral contracts need not assemble the full coalition. Acquisition can instead condition closing on joint participation and place the target's technology under the acquirer's control. An expanded nonacquisition arrangement that assembles the same coalition, transfers the same control rights, and implements the same retention and project direction at the same real resource cost removes the buyer-side productive advantage of acquisition. Under the sufficient condition that acquisition value is more sensitive to coordination productivity than the parties' combined no-

acquisition payoffs, a decline in coordination productivity weakly reduces acquisition feasibility. For groups acquired before an unanticipated decline, lower coordination productivity weakly reduces the relationship stock the acquirer optimally preserves. The first comparative static motivates a choice-time incidence test. The second motivates the sharper post-transfer design, which fixes organizational form before coordination conditions change.

Consistent with the choice-time prediction, more exposed industries shift away from acqui-hires during the disruption. I measure predetermined industry exposure using the prevalence of work that cannot be performed entirely from home (Dingel and Neiman, 2020). During 2020–2022, the ratio of expected acqui-hire counts to its pre-2020 level is 45.7 percent lower in an industry with one standard deviation greater exposure, comparing target industries within an acquirer-year. After dividing acqui-hire counts by total measured transfers—acqui-hires plus comparison worker flows—in the acquirer–industry–year cell, the corresponding difference is 25.8 percent. Pre-period interactions are jointly indistinguishable from zero in both designs. The estimates show that observed organizational-form incidence responds to the productivity of coordination.[3]

The sharper test fixes organizational form before coordination conditions change. I study groups transferred before April 2020 and measure the share of each outcome window exposed to the subsequent disruption. Task exposure is greater for groups whose jobs combine low teleworkability with interaction-intensive activities. Within completed acqui-hires, moving both disruption exposure and task-based exposure by one sample standard deviation reduces pairwise co-retention at 24 months by 7.6 percentage points. In the cross-form design, which uses a symmetrically measured pre-transfer task index, the analogous within-sample contrast produces an additional 2.8-percentage-point decline relative to comparison flows. Expected seniority gains and promotions per worker fall along the same differential gradient. Pre-disruption placebo windows produce no corresponding effects, and measuring the within-

[3]Because the data do not reveal every organizational form considered for the same acquirer–team opportunity, the design identifies a shift in observed incidence rather than the probability of acquisition for a fixed opportunity.

acqui-hire task index only from job assignments observed before the disruption yields similar estimates. Because every transaction predates the disruption, the estimates capture a post-transfer response after the firm boundary has been chosen, rather than selection into the original acquisition.

The retention response matters because it extends to inventive production. As disruption exposure rises, the rate at which incoming inventors patent jointly falls more after acqui-hires than after comparison flows. Production combining multiple incoming inventors with pre-existing acquirer inventors also falls relative to solo incoming-inventor output, and that response is more negative after acqui-hires; citation- and private-value-weighted output moves similarly. At longer horizons, I turn to pre-existing acquirer inventors later observed co-inventing with an incoming inventor. Their independent production falls more in fields represented in the incoming group's pre-transfer expertise than in their other fields, again with a larger disruption response after acqui-hires. This pattern is consistent with internal knowledge diffusion through interaction.[4] The sequence traces the value of preserving the incoming group from its own joint production, through integration with the acquirer's inventive base, to broader use of its expertise inside the acquirer.

The same disruption also changes who remains inside the acquirer. In the acqui-hire-versus-flow design, analogous one-standard-deviation contrasts raise differential exit at 24 months by 2.0 percentage points and departure within 180 days of another original group member by 2.6 percentage points. The incidence of at least two members exiting and later entering entrepreneurially oriented destinations or roles rises by 4.8 percentage points. Thus the firm boundary continues to adjust after the acquisition through exit, temporally proximate departures, and entrepreneurial reallocation.[5] This adjustment raises a further question: whether knowledge produced inside the acquirer later crosses the boundary with departing workers.

[4]Because the set of pre-existing acquirer inventors is defined by observed co-invention, the estimate does not identify the causal effect of a particular collaboration.

[5]The measures do not require departing workers to join the same firm or jointly found a new company, and the design does not separately attribute exit to employer demand or worker supply.

Acquisition can simultaneously change innovative output, technological direction, and the routes through which knowledge reaches outsiders. Along observed transfer paths, acqui-hire groups produce 2.31 times as many inside patents per incoming worker as comparison flows, using patents filed within 48 months and eligible for a complete 36-month citation window. Direction matters for the external reach of this output. Within patent-producing acqui-hires, patents one standard deviation closer to the acquirer's pre-transaction technology receive 23 percent fewer nonaffiliate citations in a 36-month-capped window. This is a conditional within-acquisition association, not a causal effect of project direction. At the organizational-form level, nonaffiliate follow-on per inside patent is 1.33 times as high after an acqui-hire but not statistically distinguishable between acqui-hires and comparison flows. The higher patent output means that nonaffiliate follow-on per incoming worker is 3.07 times as high. Worker mobility supplies another route: among workers who exit and subsequently patent outside the acquirer's corporate family, post-exit patents following acqui-hires are more closely related to other original group members' pre-exit work inside the acquirer than are patents following comparison flows. This selected-sample comparison documents technological continuity after exit. The evidence therefore separates the quantity of inside innovation, its direction and external follow-on per patent, and diffusion through mobile workers. None of these margins alone determines the innovation consequences of acquisition.

A separate local-market analysis asks whether economic activity changes where acqui-hire teams arrive. Using synthetic difference-in-differences, I compare markets receiving their first observed acqui-hire with weighted combinations of never-treated markets in the same broad industry (Arkhangelsky et al., 2021). In the fifth complete post-transaction year, establishments are 4.9 percent higher, employment is 8.7 percent higher, and payroll is 10.0 percent higher relative to their synthetic comparisons. In the Business Dynamics Statistics, the average post-treatment number of new employer firms is 5.2 percent higher. The establishment and firm-entry results show that receiving-market expansion includes

new-firm entry as well as growth in local economic activity.[6]

For antitrust review, the central implication is that an acqui-hire changes more than the location of a fixed stock of talent. Acquisition can preserve team-specific capital and give the acquirer residual control over uses of target technology not specified ex ante. Its external consequences depend on how the team and technology are used—how much innovation is produced, which projects are pursued, and whether workers subsequently carry knowledge across the firm boundary. Neither patent counts nor a static measure of talent concentration is therefore sufficient to evaluate the innovation consequences. Evaluating a prohibition requires the no-acquisition counterfactual, under which direct hiring, team assembly, target continuation, project direction, innovative output, and worker mobility would all adjust. The reduced-form evidence does not by itself recover that counterfactual or total welfare.[7] It instead identifies the organizational and diffusion margins that an innovation counterfactual for acqui-hire policy must combine.

This paper is closest to research on acqui-hires and on human capital as an acquisition motive. Legal and strategy work describes acqui-hire contracting and studies the integration, retention, and subsequent careers of acquired workers and teams (Coyle and Polsky, 2013; Sawicki, 2015; Chatterji and Patro, 2014; Ng and Stuart, 2022; Kim, 2024; Boyacıoğlu et al., 2024). Ng and Stuart (2022) and Kim (2024) compare acquired employees with ordinary new hires and find that the former leave sooner. My comparison instead asks whether acquisition preserves a pre-existing team more effectively than assembling all or part of it through hiring. An acquired worker can leave sooner than an ordinary hire even if acquisition initially preserves more of that worker's prior team. A broader finance literature shows that labor scarcity, mobility restrictions, and the cost of developing human capital inside the firm affect acquisition decisions (Ouimet and Zarutskie, 2020; Chen et al., 2021, 2024; Beaumont et al., 2025). Related evidence shows that workers' productivity depends on their coworkers

[6]The firm-entry analysis does not link new firms directly to former members of the incoming group, and the estimates describe local incidence rather than aggregate net creation.

[7]The paper does not quantify wage-setting power, worker surplus, product-market effects, or general-equilibrium changes in startup formation, employment, and innovation.

and organizational environment, and that joint mobility and corporate control can preserve team-specific capital when a firm boundary breaks down (Groysberg et al., 2008; Jaravel et al., 2018; Baghai et al., 2024). Most closely, Beaumont et al. (2025) study building versus buying a stock of operational human capital; Bena et al. (2026) study how ownership affects relationship-specific inventor investment and distinguish acquisitions from inventor hiring without patent-ownership transfers; and Boyacıoğlu et al. (2024) study team integration conditional on an acqui-hire. I study a different organizational margin: why an acquirer acquires the target rather than attempting to assemble all or part of its complementary team through direct hiring while contracting separately for target-held technology. The framework distinguishes the coordination of workers' participation from residual control over that technology, and the data compare acqui-hires with observed group assembly by the same acquirers without acquisition of the source firm.

The organizational-choice evidence also speaks to research on firm boundaries and the organization of innovation. Technological relatedness shapes merger formation, and changes in ownership alter the scale, direction, and organization of inventive activity (Bena and Li, 2014; Seru, 2014; Bernstein, 2015; Frésard et al., 2020). Research on inventor teams and technology acquisitions shows how team design and post-transaction integration shape knowledge recombination and innovation (Aggarwal et al., 2020; Puranam et al., 2006; Puranam and Srikanth, 2007; Li and Wang, 2023). A complementary literature shows that the returns to new technologies depend on skilled human capital and complementary organizational investments, and that organizational frictions shape which firms implement them (Bresnahan et al., 2002; Brynjolfsson et al., 2021; Babina et al., 2024; Crouzet et al., 2026). Acquisition opportunities can also change ex ante R&D incentives and the direction of innovation (Phillips and Zhdanov, 2013; Dijk et al., 2024). This paper studies a boundary choice through which an acquirer obtains a pre-existing innovative team and control over target-held technology. I follow the incoming team from preservation through its own production, production with pre-existing acquirer inventors, subsequent independent production by those

inventors in related fields, and technological redirection. The coordination design identifies the team-complementarity margin behind preservation and production; the patent evidence traces how the incoming team and target technology are integrated and directed inside the acquirer.

The subsequent movement of workers back across the firm boundary connects the paper to work on inventor mobility, knowledge diffusion, and employee entrepreneurship. Mobile inventors carry ideas across firms, while prior employment, coworker networks, and corporate R&D shape the creation of new ventures (Almeida and Kogut, 1999; Singh and Agrawal, 2011; Franco and Filson, 2006; Gompers et al., 2005; Babina and Howell, 2024). Kim (2022) shows that startup acquisitions increase employee entrepreneurship, and Ma et al. (2025) show that firms can choose the scope of innovation to influence workers' mobility. I connect these margins to the initial organizational choice by following the same incoming group beyond the acquirer. The coordination disruption that reduces retention increases exit, departures clustered in time, and entrepreneurial reallocation; patent histories then document technological continuity when departing workers produce innovation outside the acquirer.

The results also inform research on acquisitions, labor-market power, and innovation policy. An acquirer may terminate a competing project, while recent evidence on digital acquisitions finds greater outside follow-on to acquired patents despite a large contraction in the target workforce (Cunningham et al., 2021; Ederer et al., 2025). Theories of acqui-hiring emphasize preemption and monopsony power (Bar-Isaac et al., 2025; Benkert et al., 2025). Related work studies talent hoarding through hiring and retention by large incumbents (Wu and Qian, 2026). Most closely on the labor-market-power margin, He and Xue (2026) show that inventors' output, earnings, and mobility fall when a merger substantially increases concentration in an already concentrated technology-specific labor market. Their mechanism turns on a loss of outside options, not on merger status alone. Acqui-hires need not combine close substitutes in the relevant inventor labor market, although the same mechanism can apply when a transaction materially consolidates employers. I isolate different margins

by separating innovation produced inside the acquirer, technological direction, nonaffiliate follow-on per patent, and diffusion carried by departing workers. These margins explain how continued production, technological redirection, and later worker-mediated diffusion can coexist, and why the innovation effect of prohibiting an acqui-hire cannot be inferred from any one of them.

The remainder of the paper follows the economic argument. Section 2 describes the setting and data, and Section 3 develops the deal and team measures. Section 4 establishes the main descriptive facts, and Section 5 presents the framework. Sections 6 and 7 test its predictions for organizational form, team preservation, and innovation. Section 8 follows workers and knowledge beyond the acquirer and studies local market outcomes. Section 9 concludes.

2 Economic Setting and Data

2.1 Acqui-Hires as an Organizational Form

In this paper, an acqui-hire is an acquisition undertaken primarily to obtain the target's team. This definition distinguishes an acqui-hire from an acquisition undertaken principally to operate the target's business or obtain its nonhuman assets, and from direct team hiring, in which the acquirer attempts to assemble all or part of the team without acquiring the target. A transaction can nevertheless serve several purposes: an acquirer may seek the team together with the target's technology, products, or contracts. Legal scholarship likewise uses a team-centered definition and documents variation in the legal form of these transactions (Coyle and Polsky, 2013; Pollman, 2023).

Acquisition changes the firm boundary without conveying ownership of labor. Workers remain voluntary contracting parties, subject to enforceable contractual restrictions, and a change of control does not by itself bind an at-will employee to remain. Purchase consideration and compensation for employment or retention are therefore distinct payments (Coyle and Polsky, 2013; Pollman, 2023). By contrast, acquisition can transfer control over assets held

by the target. In a stock acquisition, the acquirer controls the entity that continues to own those assets; an asset acquisition can transfer specified assets directly. Hiring inventors alone does not convey target-owned patents or other intellectual-property rights (Sawicki, 2015). Target-owned code, data rights, licenses, and contracts likewise require acquisition, assignment, or license.

The transaction can also link worker participation to a common closing. An acquisition agreement may condition the acquirer's obligation to close on accepted employment agreements from designated workers or specified shares of the workforce. The disclosed Sarcos–RE2 agreement, for example, imposed participation conditions for a key executive, a scheduled employee group, the broader workforce, and individual departments. [8] Each worker still chooses whether to accept, and compensation can remain individually negotiated.

Direct team hiring instead requires separate employment agreements and, when needed, a license or asset purchase. Microsoft's 2024 arrangement with Inflection illustrates this alternative: members of Inflection's core team entered employment agreements with Microsoft, Microsoft obtained a nonexclusive intellectual-property license, and Inflection remained a separate company. [9] With complete contracts, team-contingent employment offers and contracts granting all relevant technology rights can reproduce the same allocation of the team and technology without acquisition.

The framework in Section 5 studies two departures from this benchmark. A common closing can coordinate the participation of complementary workers when separate or sequential bargaining may fragment the team. Ownership can also confer residual control over target-held technology when nonacquisition contracts leave future uses unspecified (Grossman and

[8] Section 2.2(b)(viii) of the Sarcos–RE2 Agreement and Plan of Reorganization, filed as Exhibit 2.1 on March 29, 2022, required the chief executive's specified Key Employee Agreement to remain in force and executed offer documents from at least 90 percent of a scheduled employee group, 80 percent of all current employees other than the chief executive, and 50 percent of employees in each of eight departments. The conditions were waivable by the acquirer, and the agreement did not require continued employment for a specified period.

[9] See Competition and Markets Authority, *Microsoft Corporation's Hiring of Certain Former Employees of Inflection and Its Entry into Associated Arrangements with Inflection*, CMA/14/24, paras. 2, 30, 84, 94–100, and 110–114 (October 24, 2024).

Hart, 1986; Hart and Moore, 1990). These advantages concern participation and control at closing; neither gives the acquirer ownership of workers or guarantees their subsequent retention.

2.2 Transaction and Documentary Data

The transaction data come from PitchBook. The unit is a completed merger or acquisition from 2000 through 2025 in which 100 percent of the PitchBook-defined target is acquired and at least one reported acquirer's ultimate parent is headquartered in the United States. Targets may be located anywhere and may be standalone companies or carved-out businesses. I use the recorded deal date as the completion date and event time throughout the paper. [10] PitchBook supplies transaction status, ownership stakes, industries, financing histories, and firm identifiers. Acquiring subsidiaries are mapped to their ultimate parents before applying the geographic restriction.

For each qualifying transaction, I retrieve up to five reports about the transaction, each linked to its cited source. These reports provide the evidence used to determine whether obtaining the target's team was identified as a transaction motive. Section 3 develops the documentary proxy.

2.3 Worker Histories and Firm Linkage

Revelio Labs compiles publicly available professional profiles into employment histories containing employer identifiers, job titles, occupations, seniority, education, and model-based compensation measures. [11] I link PitchBook targets and acquirers to Revelio firms and map subsidiaries to ultimate parents. The linked histories identify observed target employees

[10] For example, PitchBook defines *HTC (Mobile Division)*—rather than HTC Corporation—as the target in Google–HTC. HTC reported that employees and assets would transfer to Google while Google separately obtained a nonexclusive intellectual-property license; Alphabet treated the transaction as a business combination. See HTC Corporation, *2017 Annual Report*, p. 246; and Alphabet Inc., Form 10-Q for the quarter ended September 30, 2018, Note 7.

[11] See Revelio Labs, "Methodology."

who join the acquirer's corporate family around an acquisition, groups entering the same acquirer without acquisition of their source firm, and subsequent career outcomes. Open employment spells are administratively right-censored on September 30, 2025. Coverage depends on workers maintaining public professional profiles and therefore is not a workforce census.

2.4 Patent and Inventor Data

I link PatentsView inventors to Revelio workers and patent assignees to Revelio firms and their ultimate parents. I use the Quarter 3, 2025 PatentsView release and its granted-patent data dictionaries (United States Patent and Trademark Office, 2025b,a). The consolidated records cover granted U.S. utility patents from 1976 through September 2025 and contain filing and grant dates, inventor teams, assignees, Cooperative Patent Classification codes, and citations. Filing date assigns patents to event time; recent filing cohorts are incomplete because applications enter the data only after grant.

Patent identifiers also link to PatentsView's granted-patent long-text tables for titles, abstracts, and claims. I embed this text with a fixed pretrained model (Li et al., 2023) and use equal-patent-weighted portfolio representations to measure technological proximity. Appendix B.7 describes the corpus, preprocessing, coverage, and portfolio definitions.

I also link patents to the text-based measures of Kelly et al. (2021) and the market-based private values of Kogan et al. (2017). Both have shorter, measure-specific coverage than the base patent records. Inventor links follow a worker's inventive history, whereas assignee links identify patents assigned to the acquirer's corporate family; the two sets need not coincide. Market-based patent values cover only the linked subset and measure private, not social, value.

2.5 Local Industry-Market Data

Local outcomes come from the Quarterly Census of Employment and Wages (QCEW), 2004–2023. The QCEW reports establishments, employment, and payroll for workers covered by federal or state unemployment insurance at detailed industry and geographic levels. [12] I aggregate counties to a fixed set of Core-Based Statistical Areas, harmonize industry codes over time, and leave disclosure-suppressed cells missing.

I use the Census Bureau's Business Dynamics Statistics (BDS), 2004–2022, to separate changes in operating establishments from firm entry and young-firm activity. The BDS report firms, establishments, and employment by metropolitan area, industry sector, and firm age (Stinson et al., 2021). [13] I distinguish firms aged zero, firms aged zero through five, and mature firms aged six or older. I leave suppressed observations missing and require complete outcome histories. Because joint geography-by-firm-age tabulations are available only at the sector level, BDS industry cells are coarser than the four-digit QCEW cells.

2.6 Empirical Samples

The transaction-level measurement population includes public and private acquirers. Worker and patent analyses focus on public acquirers, for which the necessary corporate-family, worker, inventor, patent, and financial links are available. Each empirical design then imposes its own linkage, timing, support, and outcome-observability requirements and reports its estimation sample. PitchBook does not provide a comparable candidate population for hiring and licensing arrangements outside acquisition records, so the paper does not estimate their prevalence.

I next turn to the documentary and worker-history measures that make the two organizational arrangements observable.

[12] See the U.S. Bureau of Labor Statistics, "Quarterly Census of Employment and Wages: Overview."
[13] See the U.S. Census Bureau, "Business Dynamics Statistics."

3 Measuring Acqui-Hires and Innovative Teams

3.1 Documentary Measurement of Acqui-Hires

Section 2 defines an acqui-hire as an acquisition undertaken primarily to obtain the target's team. Because transaction databases do not record that latent motive, I construct a documentary proxy from public transaction coverage. The indicator equals one when the supplied text explicitly presents obtaining the target's employees, founders, or team as a transaction motive. I call these transactions identified acqui-hires. The rule establishes a stated personnel motive, not that personnel were the sole or dominant motive, and never infers motive from subsequent outcomes.

Measurement proceeds in two stages. A retrieval model collects up to five reports about each transaction and preserves their source attribution; a separate classifier applies the documentary rule. [14] The prompt excludes routine employee continuation, product shutdown, undisclosed terms, and generic references to talent. Every positive classification retains the supporting passage, source identifier, and named firms and people. Neither stage receives worker histories, patent outcomes, or later econometric estimates. Appendix B provides the complete implementation. This application follows the use of language models as measurement devices in economic research (Ludwig et al., 2025).

The final population contains 109,299 qualifying acquisitions from 2000 through 2025. Of the 76,341 with usable returned text and a completed classification, 12,271 satisfy the rule, or 16.1 percent; this rate is conditional on documentary coverage. The population comprises two nonoverlapping components. Among public-acquirer transactions, 41,050 of 54,158 eligible acquisitions are classified and 8,139 are identified as acqui-hires. The additional component, which includes both public and private acquirers, contains 55,141 eligible acquisitions, 35,291 classifications, and 4,132 identified acqui-hires. Figure 1 reports the transaction population together with the classification and worker-linkage samples used below; the latter two are

[14] The retrieval model is Perplexity Sonar and the documentary classifier is Gemini 3 Flash Preview.

drawn separately from the transaction population rather than forming sequential filters.

3.2 Validation of the Deal Measure

Each positive classification retains a passage and document identifier that can be compared with the text supplied to the classifier. I also examine target characteristics that do not enter the rule. Median target employment is 56 in identified acqui-hires and 82 in other classified acquisitions; median age is 12 rather than 15 years. Identified targets are more likely to be venture-backed (31.5 versus 17.2 percent) and to operate in information technology (38.4 versus 22.0 percent). The size and age differences remain within information technology. [15]

An independent validation uses worker records that were not supplied to either model. [16] Among 13,670 transactions with a nonempty observed target-to-acquirer-family cohort and a high-skill share below one, the probability of an identified acqui-hire rises toward the upper end of the high-skill-share distribution. The denominator is the observed acquirer-family cohort, including seamless subsidiary continuations where applicable. The comparison therefore concerns cohort composition, not the fraction of a target's pre-transaction workforce that transfers. The documentary rule selects the smaller, younger, technology-oriented targets and high-skill incoming groups expected of team-motivated acquisitions. Post-transfer retention, career progression, and patenting remain outcomes rather than inputs to the measure.

3.3 Measuring Incoming Groups across Acquisition and Hiring

The documentary indicator identifies the transaction but not the affected workers or their prior relationships. After linking each PitchBook target and acquirer to Revelio, I identify workers employed by the target near completion who either move to the acquirer's corporate family or continue there after the target becomes a subsidiary. Let E_e denote this observed

[15] These target-characteristic comparisons use a 54,394-transaction public-acquirer sample. Appendix Table A.1 reconciles that sample with the final eligibility population.

[16] A worker is high skill if seniority is at least five on Revelio's seven-point scale; the occupation is managerial, professional, scientific, engineering, computing, or healthcare; the worker has an advanced degree; or model-based compensation is at or above the year-specific seventy-fifth percentile.

acquirer-family cohort, including recorded employer transitions and seamless subsidiary continuations, and let $M_e \subseteq E_e$ denote members with a recorded employer change. I call E_e the observed incoming group. Analyses requiring a recorded entry date or labor-market move use M_e; relationship statistics use this narrower mover set. Appendix B.5 gives the linkage hierarchy and employment windows.

To construct a nonacquisition comparison, I first identify raw candidates with at least three workers from a common source firm who enter the same employer entity within an observed acquirer's corporate family within 60 days of one another, without acquisition of the source firm. Let M_c denote the candidate's recorded mover set. These deal-anchored observations document realized group assembly through hiring; their contracting arrangements are unobserved. I retain candidates whose pre-existing workplace relationships are unusually dense relative to equal-size placebo groups from the same source workforce and call the retained observations comparison worker flows. Appendix B.6 gives the clustering windows, thresholds, overlap rules, and source-group reconstruction.

The principal relationship measure is prior joint mobility. A pair is linked when its members overlapped for at least 90 days at an earlier employer and subsequently entered another common employer within 365 days of one another, all before the focal event. For a mover set M of size K, prior-mobility density is

$$\text{Cohesion}(M) = \frac{2}{K(K-1)} \sum_{i<j;\, i,j \in M} L_{ij}, \tag{1}$$

where L_{ij} indicates a prior joint-mobility link. I apply the same definition to acqui-hire mover sets M_e and raw comparison mover sets M_c. For analyses requiring a broader pre-transfer group, I augment the recorded movers with observed occupational peers at the source firm. This reconstructed group is a workplace-peer proxy, not a census of the economically relevant team; Appendix B.6 gives the exact overlap, occupation, and cap rules.

3.4 Validation and Scope of the Team Measures

Applying the same prior-joint-mobility definition before screening candidates for evidence of a shared prior workplace, mean event-level density in 2010–2022 is 9.5 percent among 2,635 identified acqui-hires with at least two recorded movers and 4.0 percent among 279,410 raw comparison candidates; pair weighting gives 14.2 and 4.4 percent. These descriptive differences support a cohesive-group interpretation of acqui-hire movers. The screen retains 45,916 candidates. Pooled prior-mobility density is 6.2 percent among retained and 0.7 percent among rejected candidates, and retained event-weighted density is 10.6 percent. This second comparison describes the groups selected by the screen.

The empirical object is the observed incoming group, not the target's complete workforce. [17] In the final public-acquirer population, 25,599 transactions link to a target in the worker data. The worker-linked and classified samples are constructed independently; 21,186 transactions appear in both, including 4,807 identified acqui-hires. The pre-transaction target-cohesion fact in Section 4.1 uses an earlier public-acquirer extraction, reconciled in Appendix Table A.1. The additional public- and private-acquirer transactions broaden the measurement population; worker, patent, and causal analyses focus on public acquirers because the required corporate-family, worker, patent, and financial links are available.

4 Acqui-Hires: Selection, Assembly, and Production

The measurement makes it possible to follow the organizational choice from selection through production. Identified acqui-hires involve unusually cohesive targets, preserve more of the incoming group's prior worker relationships than comparison flows, and are followed by greater retention and innovative output inside the acquirer. The resulting inventions both continue the group's prior work and move toward the acquirer's technology. These conditional

[17] The high-skill result describes the composition of the observed acquirer-family cohort. The available pre-transaction worker stock and aggregate incoming-worker count use different populations and timing rules, so I do not form a transfer-rate ratio.

level comparisons motivate the framework in Section 5; Sections 6 and 7 then isolate the role of team complementarity.

4.1 Pre-Transaction Selection: Teams and Technology

Identified acqui-hire targets are smaller, younger, more likely to be venture-backed, and more concentrated in information technology than targets in other acquisitions (Section 3.2). They also contain denser prior worker relationships. For each target, I freeze the observed high-skill roster during the year before completion and measure prior joint mobility among all possible worker pairs. A pair is linked when its members overlapped for at least 90 days at an earlier employer, left that employer before joining the target, and entered the target within 365 days of one another. The event-level density is the share of all possible pairs with such a link, analogous to Equation (1).

Among 27,582 classified acquisitions linked to worker histories, 20,043 targets have at least two workers in the observed high-skill roster: 4,891 identified acqui-hires and 15,152 other acquisitions. Figure 2 compares equal-transaction means within five roster-size categories. The acqui-hire mean is higher in every category. Among targets with two to five observed high-skill workers, for example, mean density is 5.15 percent for acqui-hires and 1.95 percent for other acquisitions. The difference persists when the link requires only prior coworker overlap rather than a previous joint move. Identified acqui-hire targets therefore have unusually dense prior coworker and joint-mobility ties. These ties proxy for pre-existing worker relationships; they do not measure their productivity.

Worker relationships describe the target's internal organization but not the technological relation between target and acquirer. I examine that relation among 2,826 realized acquisitions—661 identified acqui-hires—in which both firms filed a linked U.S. utility patent during the five years before completion. Table 1 distinguishes similarity in the firms' own patented output from overlap in the prior art they cite. The first two measures are patent-class portfolio proximity (Bloom et al., 2013) and claims-text similarity. The third is the share of

the smaller firm's cited-prior-art set also cited by the other, adapting the common-knowledge-base measure of Bena and Li (2014). The first two compare what the firms patent; the third compares the knowledge on which those patents draw.

Similarity in the firms' own patented output is not detectably associated with acqui-hire status. By contrast, overlap in cited prior art is. Controlling for the sizes of both patent portfolios and prior-art sets, the coefficient is 0.073 (standard error 0.034); a one-standard-deviation increase corresponds to a 2.2-percentage-point difference in acqui-hire probability. The association remains when both output-similarity measures and citation flows between the firms enter jointly. The pattern is therefore one of shared technological foundations rather than similarity in observed patented output. It motivates a role for target-held technology that is complementary to the acquirer's knowledge without implying that the two firms were already pursuing the same projects. [18]

4.2 Assembly and Persistence Across the Firm Boundary

I next compare acqui-hires with the comparison worker flows defined in Section 3.3. Their members leave a common source firm and enter the same acquirer without acquisition of the source firm. Coarsened exact matching (Iacus et al., 2012) retains cells containing both event categories, defined by transfer year, two-digit pre-transfer industry, arriving-group size, occupation composition, and prior-mobility cohesion. The resulting common-support population contains 1,133 acqui-hires and 7,095 comparison flows. The comparisons in this section are descriptive; Section 6 develops the identifying variation. Common support improves balance in observed characteristics but neither makes organizational form random nor equates an arriving group with the complete pre-transfer team.

The first outcome asks how much of the pre-existing relationship network reaches the acquirer. Within the largest connected component of each pre-transfer network, relationship

[18] This comparison conditions on realized patent-active pairs. It does not identify why a particular match formed or which technology rights changed control. Appendix Tables A.2 and A.3 report alternative specifications and text constructions.

assembly is the share of prior-mobility links for which both workers are observed in the acquirer's corporate family during the first 183 days; continuous presence throughout the window is not required. Among link-eligible events, the sample contains 609 acqui-hires and 3,458 comparison flows. Mean relationship assembly is 95.6 percent for acqui-hires and 21.1 percent for comparison flows. At least three quarters of the links reach the acquirer in 93.6 percent of acqui-hires, compared with 12.5 percent of comparison flows (Figure 3). Acqui-hires therefore bring groups that preserve most of their observed prior worker relationships, whereas similarly sized comparison flows more often assemble only part of the reconstructed source group. Because arriving-group size is balanced, the distinction concerns how much of the pre-existing relationship network arrives, not simply how many workers do.

Post-transfer persistence and career advancement. Entry-window assembly need not imply persistence. For a predefined arriving group, pairwise co-retention is the share of entrant pairs whose two members both remain in the acquirer's corporate family at a given horizon; worker retention is the corresponding individual-level outcome. These outcomes use the full common-support population rather than conditioning on observed prior links. The team-level comparisons absorb acquirer and pre-transfer-industry-by-entry-year fixed effects. Worker-level regressions instead absorb acquirer, pre-transfer-industry, and entry-year fixed effects and condition on pre-entry worker characteristics. At 24 months, pairwise co-retention is 4.8 percentage points higher following acqui-hires (standard error 1.5), and worker retention is 3.2 percentage points higher (1.3). The corresponding estimates at 36 months are 4.0 and 3.0 percentage points. By 48 months, both differences are small and statistically indistinguishable from zero. The retention difference is thus concentrated in the second and third years after entry. Initial preservation is therefore not permanent; Section 8 returns to the subsequent movement of these workers and their knowledge across firm boundaries.

Workers joining through acqui-hires also advance more rapidly inside the acquirer. At 24 months, their conditional expected number of promotions is 57.0 percent higher and their

cumulative positive seniority gain is 47.7 percent higher; at 48 months, the differences remain 37.6 and 30.7 percent. Figure 4 traces these outcomes over time. Both margins differ initially, but only advancement remains elevated at four years. The advancement differences also remain when outcomes are aggregated to the arriving-group level and scaled by group size. [19]

4.3 Innovative Production After Transfer

I next ask whether incoming workers contribute to inventive production inside the acquirer. Figure 5 reports two worker-level outcomes in the same common-support population. Within 12 months, workers joining through an acqui-hire are 2.0 percentage points more likely to appear on a patent assigned to the acquirer's corporate family; the difference grows to 3.7 percentage points within 48 months. They are also 0.7 percentage points more likely to co-invent with a pre-existing acquirer inventor within 12 months and 1.7 percentage points more likely within 48 months. The comparison-group means are small—0.27 and 1.07 percent for patenting and 0.19 and 0.69 percent for co-invention at the two horizons—so these are economically large differences. The first outcome locates incoming-worker production inside the acquirer; the second shows that this production extends beyond collaboration solely among incoming workers.[20]

The corresponding event-level estimates measure patented output per incoming worker. A qualifying patent names an incoming worker, is assigned to the acquirer, and is filed within the indicated horizon. Poisson regressions retain zero-output events, use arriving-group size as an exposure and the matching weights, and compare events within the matched support, with acquirer, pre-transfer-industry, and event-year fixed effects. Patents connected to more than one focal event are allocated fractionally across those events. At 24 months, incoming groups following an acqui-hire produce 196 percent more patents per worker than comparison

[19] Advancement is observed only while a worker remains at the acquirer and can reflect performance, contractual terms, and outside options. The level comparisons do not separately identify the acquirer's investment.

[20] Among patents that combine at least two incoming inventors with a pre-existing acquirer inventor and are observed in the private-value data, allocated private value is also higher following acqui-hires than comparison flows (Appendix Table A.6).

flows; at 48 months, they produce 183 percent more. Observed private-value-weighted patent output (Kogan et al., 2017) is 259 percent higher at 24 months and 238 percent higher at 48 months (Figure 6). [21]

Greater output does not reveal its technological direction. Panel A of Table 2 compares each post-entry patent with the incoming group's and acquirer's pre-entry patent portfolios. Relative to comparison-flow patents, the share of fields continuing the incoming group's prior work is 23.3 percentage points higher at the patent-subclass level and 27.6 points higher at the main-group level. Acqui-hire patents are also 36.4 percentage points more likely to combine at least one field from the incoming group's prior portfolio with one from the acquirer's prior portfolio.

Claims text yields the same economic pattern. At 48 months, post-entry output following an acqui-hire is 0.027 cosine units more similar to the incoming group's prior work than output following a comparison flow. Its increase in proximity to the acquirer's pre-entry portfolio is also 0.016 cosine units larger. The estimates are stable to stricter text coverage, complete potential follow-up, and excluding each incoming inventor's own earlier patents from the team portfolio (Appendix Tables A.7 and A.8). Acqui-hires are therefore followed both by continuity in the incoming group's knowledge and by movement toward the acquirer's technology. These patterns are compatible: the incoming group can preserve its expertise while shifting the technological areas in which that expertise is used inside the acquirer.

The facts show that acqui-hires assemble unusually cohesive groups with most of their prior relationships intact. They also show shared technological foundations and subsequent movement toward the acquirer's technology. These patterns motivate, but do not identify, a distinct control advantage from target-owned technology. The framework next separates that control margin from the assembly of complementary workers; the coordination design then tests whether team complementarity shapes organizational form, retention, and production.

[21] Private patent values cover 80.7 percent of acqui-hire and 95.1 percent of comparison patent–event observations. Events with uncovered patents are not assigned zero value. Appendix Tables A.4–A.5 report follow-up, linkage, and influence checks.

5 Acquisition versus Direct Hiring of Innovative Teams

Why acquire a target primarily for its innovative team rather than hire the workers? The framework separates two advantages. A common closing can coordinate the joint participation of complementary workers, while ownership transfers control over target-held technology that employment contracts do not convey. Workers remain free to reject or leave under either form; acquisition never transfers ownership of labor. The model conditions on a target and a fixed set of potential acquirers, called buyers below, and does not model target selection, bidder entry, or merger policy. Appendix C.1 contains proofs and technical derivations.

5.1 Team Production and Organizational Alternatives

A target startup i has equity claimants and an innovative team $\mathcal{T}_i = \{1, \ldots, n_i\}$, where $n_i \geq 2$. The team contains the workers whose joint participation is relevant for the buyer's intended activity; it need not contain every target employee. A founder can be both a worker and an equity claimant. I keep these roles separate: purchase consideration pays for target ownership, control, or specified assets, while employment and retention packages compensate workers for joining and remaining with the buyer.

Worker k contributes standalone human capital $h_{ik} > 0$. Workers can also have productive relationships established before the transaction. For a coalition $M \subseteq \mathcal{T}_i$, define

$$\begin{aligned} H_i(M) &= \sum_{k \in M} h_{ik}, \\ L_i(M) &= \frac{2}{n_i - 1} \sum_{\substack{k<\ell \\ k,\ell \in M}} a_{ik\ell} h_{ik} h_{i\ell}, \end{aligned} \tag{2}$$

where $a_{ik\ell} = a_{i\ell k} \geq 0$ measures the productive relationship between workers k and ℓ. The denominator is fixed at original-team size. Removing a worker therefore eliminates that worker's links without mechanically raising the weight on links that remain. Motivated by evidence that prior collaboration creates team-specific capital and that corporate reorgani-

zation can destroy or preserve it (Jaravel et al., 2018; Baghai et al., 2024), $L_i(M)$ captures productive relationships distinct from workers' standalone human capital.

Lemma 1 (Relationship stock is superadditive)**.** *For disjoint coalitions $M_1, M_2 \subseteq \mathcal{T}_i$, $L_i(M_1 \cup M_2) \geq L_i(M_1) + L_i(M_2)$, strictly when a productive relationship connects the two coalitions.*

The union contains the links within each coalition and all links between them. Thus, when direct hiring omits a worker, the buyer loses both that worker's standalone contribution and the worker's productive relationships with assembled members. This is the team-production margin in the framework (Alchian and Demsetz, 1972).

Let X_i denote target-controlled nonhuman assets and enforceable technology rights relevant to team production. These can include intellectual property, code, data, and rights under licenses and contracts. They exclude workers' human capital, tacit knowledge, and relationships. Hiring workers does not reallocate rights held by the target or another rightsholder. Particular uses can be licensed or assigned, but when future uses and adaptations cannot be fully specified ex ante, ownership allocates residual decision rights over uses left unassigned (Grossman and Hart, 1986; Hart and Moore, 1990). In an equity acquisition, the buyer controls the rights-holding target; in an asset acquisition, it receives the specified assets and rights. This is the model's control object. The framework takes the productive consequence of control as a primitive rather than deriving optimal ownership from ex ante noncontractible investment. Aghion and Tirole (1994) apply the incomplete-contract approach to the organization of innovative activity.

Let $X_i^H \preceq X_i$ denote the enforceable uses available when buyer j hires workers without acquiring the target or separately contracting with the rightsholder. The buyer's operating value from coalition M, usable target technology x, and coordination environment $q > 0$ is

$$V_{ij}(M, x; q) = \phi_{ij} H_i(M) + q\theta_i g_{ij}(x) L_i(M), \tag{3}$$

where $\phi_{ij} > 0$ values standalone inputs, $\theta_i > 0$ is the productivity of the team's pre-existing

relationships, and $g_{ij}(x) > 0$ is weakly increasing in the technology available to the buyer. Technology and team-specific capital are complements: control matters here because it changes how productively the team can use its relationships. A useful specialization sets $g_{ij}(X_i^H)/g_{ij}(X_i) = \rho_L \in [0,1]$; ρ_L is a productivity ratio, not a fraction of relationships, patents, rights, or assets.

The coordination environment q, normalized to one in the baseline environment, measures how effectively existing relationships can be used. It does not determine whether the relationships exist: $L_i(M)$ is fixed before organizational choice, whereas q can change. The benchmark omits target assets whose value is separable from the team. Such assets can raise bid capacity and price but do not change the coalition-assembly comparative static. This restriction reflects the paper's focus on transactions in which obtaining the team is the primary motive. Project direction and the acquirer's other complementary assets are held fixed here and examined separately in Section 8.2.

5.2 Sequential Hiring, Coordinated Closing, and Target-Controlled Technology

At the organizational-choice date, the parties observe the team, outside options, buyer–target match, and coordination environment q_i^0, which they expect to govern the relevant valuation horizon. Agents are risk neutral with quasilinear monetary payoffs, accepted agreements are enforceable, and there are no liquidity constraints. Under acquisition, the buyer can condition closing on the transfer of target control or specified assets and acceptance by every required worker. Each worker still signs an individual employment agreement and can refuse. Conditional on closing, the required team begins employment with the buyer and the buyer controls X_i.

Baseline direct hiring uses a different contracting technology. The buyer approaches workers sequentially. An accepted offer is binding and cannot be conditioned on every worker approached later also accepting. The buyer chooses whom to approach next and can stop

after any history. A rejected offer ends the sequence with the coalition already assembled. If bargaining ends with coalition M, the buyer produces with M and X_i^H; additional technology rights require a separate agreement. This is an assembly-of-complements problem (Kominers and Weyl, 2012): a worker's value depends on who else ultimately joins, but bilateral contracts become binding before the final coalition is known. The restriction on team contingency is a maintained contracting technology, not a claim that coordinated offers are legally impossible.

A two-worker example isolates the wedge. Suppose each worker has reservation value r, produces $u < r$ alone, the complete team produces V, and a worker receives bargaining share $\beta \in (0, 1)$. A coordinated package creates positive surplus when $V > 2r$. Sequential hiring nevertheless fails to begin when

$$2r < V < 2r + \frac{\beta}{1-\beta}(r - u). \tag{4}$$

The first bargain reflects the buyer's continuation value after the second worker captures part of the complementarity surplus. Full-team surplus is therefore not sufficient for sequential bilateral hiring to assemble the team. This is a possibility result, not a claim that direct hiring always fragments teams. Appendix C.1.3 gives the general finite-team recursion and defines the coalition $M_{ij}^H(q) \subseteq \mathcal{T}_i$ delivered by direct hiring.

An expanded nonacquisition benchmark permits a package contingent on joint acceptance by a specified coalition and a separately priced license, assignment, or asset transfer that grants the enforceable rights in X_i. When reproducing X_i requires residual decision rights over unspecified future uses, the alternative must transfer control of the relevant asset or rights-holding entity; a license for specified uses alone is not equivalent. If that arrangement assembles the same workers, transfers the same control rights, implements the same retention and project direction, and uses the same real resources, acquisition has no buyer-side productive advantage. The equivalence does not equate the independent target's continuation, total innovation, surplus division, or welfare. It states precisely what contracts would have to

reproduce for the firm boundary to be irrelevant.

5.3 Acquisition Market and Organizational-Form Choice

The productive difference between acquisition and direct hiring can be written as

$$\begin{aligned} & V_{ij}(\mathcal{T}_i, X_i; q) - V_{ij}(M_{ij}^H(q), X_i^H; q) \\ = & \underbrace{V_{ij}(\mathcal{T}_i, X_i; q) - V_{ij}(M_{ij}^H(q), X_i; q)}_{\text{coalition-assembly margin}} \\ & + \underbrace{V_{ij}(M_{ij}^H(q), X_i; q) - V_{ij}(M_{ij}^H(q), X_i^H; q)}_{\text{control-of-technology margin}}. \end{aligned} \tag{5}$$

The first term changes the worker coalition while holding control of target technology fixed. The second changes available technology while holding the coalition fixed. Either can be positive when the other is zero. Because technology and relationship stock are complements, reversing the order changes how their interaction is allocated across the two terms but not their sum; Appendix C.1.4 gives that accounting.

Production is not the same as the private acquisition decision. The buyer anticipates post-close retention and integration. Let

$$\Pi_{ij}^A(q) = \max_{R \subseteq \mathcal{T}_i} \left\{ V_{ij}(R, X_i; q) - K_{ij}^A(R; q) \right\}, \tag{6}$$

where K_{ij}^A is the real integration and retention cost of implementing coalition R. It excludes monetary compensation, which is a transfer handled separately through workers' participation constraints. The notation does not give the buyer a right to prevent workers from leaving; it summarizes the real resources needed to implement a voluntary retention outcome.

Let C_{ij}^A be the real transaction cost, $O_{ij}^B(q)$ the buyer's no-acquisition payoff (including direct hiring), $\bar{r}_{ik}^A(q)$ worker k's no-acquisition employment payoff, and $O_i^S(q)$ the target side's

continuation value. Candidate buyer j's capacity to compensate the target is

$$\bar{b}_{ij}(q) = \Pi^A_{ij}(q) - C^A_{ij} - O^B_{ij}(q) - \sum_{k \in \mathcal{T}_i} \bar{r}^A_{ik}(q), \tag{7}$$

$$\Delta_i(q) = \max_{j \in \mathcal{J}^A_i} \bar{b}_{ij}(q) - O^S_i(q). \tag{8}$$

The highest-capacity buyer acquires the target if and only if $\Delta_i(q) \geq 0$. More candidate buyers affect incidence through the highest capacity. The second-highest capacity can raise target consideration, but absent financing constraints, price divides surplus rather than creates a second feasibility condition. This is a reduced-form market closure, not a strategic auction equilibrium; Appendix C.1.5 states the price rule, bounds, and technical cases.

5.4 Coordination Shocks and Empirical Predictions

Fix a reference environment and the comparison coalition $\widetilde{M}^H_{ij}$ delivered by direct hiring. Holding that coalition and available technology fixed, the derivative of the buyer's gross productive advantage is

$$\begin{aligned} \frac{\partial}{\partial q}\Big[V_{ij}(\mathcal{T}_i, X_i; q) - V_{ij}(\widetilde{M}^H_{ij}, X^H_i; q)\Big] = \theta_i\Big[&g_{ij}(X_i)\{L_i(\mathcal{T}_i) - L_i(\widetilde{M}^H_{ij})\} \\ &+ \{g_{ij}(X_i) - g_{ij}(X^H_i)\}L_i(\widetilde{M}^H_{ij})\Big] \geq 0. \end{aligned} \tag{9}$$

The response is strict if acquisition preserves additional relationship stock or makes the retained stock more productive through greater technology access. The two terms map directly to coalition assembly and control of technology.

This productive comparison does not by itself sign acquisition feasibility. Endogenous retention, real implementation costs, workers' employment payoffs, the buyer's direct-hiring opportunity, and the target's continuation value can all respond to q; the winning buyer can also change. A sufficient condition makes the required comparison explicit.

Corollary 1 (A sufficient condition for acquisition feasibility)**.** *Suppose that, for every*

candidate buyer and $q_2 > q_1$,

$$\Pi^A_{ij}(q_2) - \Pi^A_{ij}(q_1) \geq [O^B_{ij}(q_2) - O^B_{ij}(q_1)] + \sum_{k \in \mathcal{T}_i} [\bar{r}^A_{ik}(q_2) - \bar{r}^A_{ik}(q_1)] \tag{10}$$

$$+ [O^S_i(q_2) - O^S_i(q_1)]. \tag{11}$$

Then $\Delta_i(q)$ *is weakly increasing in* q*. A decline in coordination productivity can eliminate an acquisition feasible at a higher* q*, but cannot make an infeasible acquisition feasible under this condition.*

Complementarity alone does not impose the condition: lower coordination productivity can also reduce the independent target's value or alter bargains under direct hiring. The effect on acquisition price is likewise ambiguous. This qualification separates the productive force in Equation (9) from the full private choice.

The timing of the shock determines the prediction. For transactions considered during a disruption, the choice-time environment q^0_i enters Δ_i. For transactions completed earlier, a later disruption is an unanticipated change in coordination conditions. It can affect retention and production, but it cannot explain the original organizational choice.

To formalize the post-transfer margin, let $M^A_{ij} = \mathcal{T}_i$ and $M^H_{ij} = M^H_{ij}(q^0_i)$ denote the groups initially assembled under the two forms, with available technology $x^A_{ij} = X_i$ and $x^H_{ij} = X^H_i$. The employer's retention and integration policies induce

$$R^\mu_{ij}(q_{ih}) \in \arg\max_{R \subseteq M^\mu_{ij}} \left\{ V_{ij}(R, x^\mu_{ij}; q_{ih}) - K^{\mu,\text{post}}_{ij}(R) \right\}. \tag{12}$$

This is a reduced-form implementation problem for voluntary retention, not an assumption that the employer can prevent separation.

Proposition 1 (Coordination and retained relationships)**.** *Fix organizational form* μ*, buyer* j*, available technology, implementation costs, and feasible coalitions. If* $g_{ij}(x^\mu_{ij}) > 0$*, then higher* q_{ih} *weakly raises the relationship stock of every optimal retained coalition.*

The sharp implication concerns preserved relationships, not headcount. A headcount result requires additional restrictions that make optimal coalitions nested. Employment retention, career progression, and compensation can respond to the value of preserving the team, but they are not themselves measures of team productivity or latent investment.

The framework therefore yields two predictions. Under Corollary 1, a decline in choice-time coordination productivity weakly reduces acquisition feasibility relative to direct hiring. After a pre-disruption transfer, Proposition 1 predicts that a decline weakly reduces the relationship stock the employer optimally preserves. Prior joint-mobility links proxy for the relationships in L_i; relationship assembly and pairwise co-retention record how much of that stock crosses and remains inside the acquirer; and joint patents measure coordinated production. The patent-pairing and direction facts motivate the distinct target-technology object X_i, but do not identify its control advantage. Comparison flows show realized group transfers without acquisition, although their contracts are unobserved. The disruption evidence therefore disciplines the team-complementarity margin rather than separately identifying the two advantages of acquisition.

6 Team Complementarity and Organizational Form

6.1 Coordination Shock and Empirical Strategy

Section 4 shows that acqui-hires assemble and preserve more pre-existing worker relationships than comparison worker flows. The framework yields two corresponding predictions. Lower coordination productivity can reduce acquisition feasibility relative to direct hiring (Corollary 1); after an acquisition has occurred, it can reduce the relationship stock the acquirer chooses to preserve (Proposition 1). I test these margins separately.

The empirical setting is the COVID-19 disruption to in-person work. Its timing provides a common change in the coordination environment, while occupational task content governs exposure. For organizational-form incidence, I follow Dingel and Neiman (2020) and use the

industry employment share in occupations that cannot be performed entirely from home. For post-transfer outcomes, I construct a group-specific index. At the occupation level, I multiply one minus Dingel and Neiman's teleworkability measure by the average O*NET importance rating across nine communication, interpersonal, coordination, and supervisory activities; I standardize this product and average it across the transferred high-skill workers. The index is therefore high when the incoming group's work is both difficult to perform remotely and interaction intensive. Task content is an exposure measure, not a direct measure of team complementarity. Appendix D.2.1 provides the construction.

The relevant occupation measure differs across designs because the designs study different margins. Within acqui-hires, C_e^A uses workers' occupations at the acquirer and measures the coordination demands of their post-transfer tasks. Because some recorded acquirer spells begin after April 2020, I interpret it as the realized post-transfer task environment rather than a uniformly predetermined team attribute. The acquisition-versus-flow design instead uses occupations at the target for acqui-hires and at the source firm for comparison flows, C_e^S. This index is determined before transfer and defined symmetrically across the two event categories. A predetermined-only version of C_e^A and a version based on pre-transfer occupations are reported in the appendix.[22]

For the post-transfer analyses, let d_e be the transfer-event date and define the share of an h-month outcome window that overlaps the primary disruption as

$$S_{eh} = \frac{\ell([d_e, d_e + h] \cap [\text{April 1, 2020}, \text{April 1, 2022}])}{\ell([d_e, d_e + h])}. \tag{13}$$

Thus S_{eh} is zero when the outcome window ends before April 2020 and one when it falls entirely within the two-year disruption. All events used in the post-transfer analyses occur by March 31, 2020, so organizational form is chosen before measured disruption exposure begins.[23]

[22]Each index is standardized in its own pre-estimation reference sample. Their coefficients are not mechanically comparable.

[23]Results excluding first-quarter 2020 transactions address possible anticipation.

The ideal choice-time experiment would change coordination productivity for a fixed acquirer–team opportunity before the acquirer chooses between acquisition and direct hiring. Because the data do not reveal every form considered for the same opportunity, I first study acqui-hire incidence across industries and then condition on the number of observed transfers of worker groups. The post-transfer design comes closer to its ideal experiment: among groups transferred before COVID-19, it compares outcome windows with different disruption overlap and workers with different task exposure. Acquirer-by-year effects absorb shocks common across an acquirer's target industries; the main post-transfer specifications use industry-by-transfer-year effects to absorb shocks common to an industry and cohort. Matching and weighting improve overlap on predetermined group or employer characteristics. Four pre-COVID placebo disruptions probe whether similar timing-by-task gradients were already present. Together, the designs separate the choice-time and post-transfer implications of the model.[24]

6.2 Organizational Form at the Time of Transfer

I examine the choice-time prediction in two steps. The first asks whether acqui-hire activity shifted across industries after the disruption. The second asks whether the same shift appears after normalizing by the number of measured transfers of worker groups. The designs successively address changes in the allocation of acquisitions and in overall transfer activity, but neither holds fixed a particular acquirer–team opportunity.

6.2.1 Acqui-hire activity across industries

The sample is a balanced 2017–2022 panel of 481 acquirer–target-industry pairs with an acqui-hire during 2017–2019. Let N_{ajt}^{acq} denote the number of acqui-hires by acquirer a in two-digit target industry j and year t, and let C_j^I denote the industry's predetermined share

[24] A contemporaneous shock that differentially affected highly exposed acqui-hire teams along the same timing gradient would violate the identifying restriction.

of employment that cannot be performed from home. The baseline specification is

$$\mathbb{E}\Big[N_{ajt}^{\text{acq}} \mid C_j^I, \alpha_{at}, \delta_j\Big] = \exp\Big\{\beta^E \,\text{Post}_t \, C_j^I + \alpha_{at} + \delta_j\Big\}, \tag{14}$$

where Post_t equals one in 2020–2022, α_{at} are acquirer-by-year effects, and δ_j are target-industry effects. The annual indicator accommodates lags in search, negotiation, and completion. Acquirer-by-year effects absorb acquisition capacity, financing, and other acquirer-wide shocks. Hence β^E is identified by changes in the allocation of an acquirer's acqui-hires across its historically active target industries. Standard errors are clustered by acquirer.

Column (1) of Table 3 reports $\widehat{\beta}^E = -2.487$ (s.e. 0.588). Since the standard deviation of C_j^I is 0.245, the expected-count ratio during 2020–2022 relative to the pre-period is 45.7 percent lower in an industry with one standard deviation greater exposure. This estimate captures a relative reallocation across industries within acquirer-years, not the absolute change in acqui-hire activity in a high-exposure industry. The estimate remains negative under inverse-probability weighting, linear specifications retaining all 2,886 cells, and PPML with acquirer rather than acquirer-by-year effects (Appendix Table A.9).

Figure 7 shows pre-2020 exposure coefficients close to zero and negative coefficients thereafter. In a quarterly specification, the 2017Q1–2019Q3 interactions are jointly indistinguishable from zero relative to 2019Q4 ($p = 0.924$). A pure financing explanation would predict a stronger decline among constrained acquirers. Instead, the negative coefficient persists among acquirers classified as less constrained by either the pre-pandemic SA or WW index, including the least-constrained quartiles (Hadlock and Pierce, 2010; Whited and Wu, 2006); Appendix Table A.10 reports these results.

I next ask whether the composition of observed transfers of worker groups moved in the same direction.

6.2.2 Acqui-hire incidence among observed transfers

I combine acqui-hires with the comparison worker flows defined in Section 3.3 and aggregate both to acquirer–pre-transfer-industry–year cells. Let N_{ajt}^{flow} be the flow count and $N_{ajt}^{\text{all}} = N_{ajt}^{\text{acq}} + N_{ajt}^{\text{flow}}$. For cells with at least one measured transfer, I estimate

$$\mathbb{E}\Big[N_{ajt}^{\text{acq}} \mid N_{ajt}^{\text{all}}, C_j^I, \alpha_{at}, \delta_j\Big] = N_{ajt}^{\text{all}} \exp\Big\{\beta^C \,\text{Post}_t \, C_j^I + \alpha_{at} + \delta_j\Big\}. \tag{15}$$

The total measured-transfer count enters as an offset. Thus β^C measures the change in expected acqui-hires per measured transfer as industry exposure rises; the fixed effects absorb acquirer-year shocks and permanent cross-industry differences in incidence.

Column (2) of Table 3 reports $\widehat{\beta}^C = -1.078$ (s.e. 0.396) in 2,158 cell-years and 235 acquirers. A one-standard-deviation increase in industry exposure corresponds to a 25.8 percent lower expected acqui-hire count per measured transfer during 2020–2022 relative to the pre-period. The annual estimates in Figure 8 show no differential pre-period gradient ($p = 0.254$); weighting on pre-pandemic employer size and age yields -1.334 (s.e. 0.401). Acqui-hire incidence therefore shifts even after normalizing by total measured transfers. The denominator nonetheless contains realized transfers rather than every organizational arrangement considered for a fixed opportunity. These are relative-incidence results; the post-transfer tests provide sharper evidence on the coordination-sensitive value of preserving the team.

6.3 Post-transfer Team Preservation

For a group of $K_e \geq 3$ workers arriving at an acquirer, let R_{eh} be the number employed there h months later. My primary outcome is pairwise co-retention,

$$\text{PairRet}_{eh} = \frac{R_{eh}(R_{eh} - 1)}{K_e(K_e - 1)}. \tag{16}$$

This is the share of possible entrant pairs whose members both remain. Because it is convex in retained headcount, a departure matters more when it severs potential relationships with many remaining members. Pairwise co-retention is still a headcount-based proxy, not the model's weighted relationship stock. I therefore also report the fraction retained and thresholds for retaining 50, 75, or 100 percent of the group.[25]

6.3.1 Within acqui-hires

I begin within completed acqui-hires. Organizational form is fixed before the shock, making this the closest empirical counterpart to Proposition 1. I estimate

$$Y_{eh} = \beta_h^W S_{eh} C_e^A + \delta_h^S S_{eh} + \delta_h^C C_e^A + W_e' \pi_h + C_e^A Z_e' \kappa_h + \alpha_a + \lambda_{jt} + \varepsilon_{eh}, \tag{17}$$

where α_a are acquirer effects and λ_{jt} are two-digit target-industry-by-transfer-year effects. The controls W_e include group size, target size and age, and the group's high-skill and professional-managerial shares. The interactions $C_e^A Z_e$ allow the task-exposure gradient to vary with group size and target scale and age; every constituent lower-order term is included. I retain strata with overlap in task-based exposure and apply Hájek weights that balance exposure quartiles within strata defined by industry, year, and predetermined group and target characteristics. Standard errors are clustered by acquirer. Identification comes from the interaction between event timing, which determines S_{eh}, and the task exposure of the group at the acquirer. The restriction is that, absent COVID-19, retention would not have developed a differential timing gradient across groups with different task exposure, conditional on these characteristics and fixed effects.

Panel A of Table 4 reports the results for 502 events and 133 acquirers. At 24 months, $\widehat{\beta}_{24}^W = -0.267$ (s.e. 0.065) for pairwise co-retention. A representative one-standard-deviation increase in both disruption exposure and task exposure implies a 7.6-percentage-point decline

[25]The empirical mapping requires pairwise co-retention and the model's weighted relationship stock to covary positively over the relevant changes in group composition.

in the probability that a randomly selected entrant pair remains jointly employed. The corresponding estimates are −0.181 (s.e. 0.043) for the fraction retained and −0.102 (s.e. 0.045) at the worker level. Threshold and worker-level outcomes have the same sign, so the result is not driven by the convex pairwise transformation. Estimates at 36 and 48 months are smaller and imprecise. No analogous gradient appears in four pre-COVID placebo windows or cohort-slope tests, and the estimate is similar after excluding first-quarter 2020 transactions or restricting the index to occupations observed before the shock. Using pre-transfer occupations attenuates the estimate; that specification measures prior tasks rather than the task environment in which the group works at the acquirer. Appendix Tables A.11–A.14 and Figure A.1 report these tests.

The within-acqui-hire result links group preservation to the coordination demands of its post-transfer work. It does not show that the response is specific to acquisition. I next compare acqui-hires with comparison worker flows using the symmetric pre-transfer task index.

6.3.2 Acqui-hires versus Comparison Worker Flows

Let $A_e = 1$ for an acqui-hire and zero for a comparison worker flow. I estimate

$$\begin{aligned} Y_{eh} = \beta_h^D A_e S_{eh} C_e^S + \beta_h^{SC} S_{eh} C_e^S + \beta_h^{AS} A_e S_{eh} + \beta_h^S S_{eh} \\ + \beta_h^{AC} A_e C_e^S + \beta_h^A A_e + \beta_h^C C_e^S + \alpha_a + \lambda_{jt} + \varepsilon_{eh}. \end{aligned} \tag{18}$$

Here β_h^D is the differential change in the disruption-overlap-by-task-exposure retention gradient. It is negative if retention among acqui-hire teams is more sensitive to the disruption than retention among the comparison flows. This prediction requires the relationship stock preserved under acquisition, or its marginal value, to be more coordination sensitive than under the comparison flows, if those flows approximate partial assembly. It is not an average treatment effect of acquisition.

The common-support sample contains 1,133 acqui-hires and 7,095 comparison flows,

each with at least three arriving workers. Support is defined over event year, two-digit pre-transfer industry, group size, occupation composition, and prior-mobility cohesion. The design compares events within arriving-group-size cells; the comparison group can remain a subset of a larger team at its source firm. Prior cohesion reduces the likelihood that a comparison flow is coincidental co-movement.[26] The baseline gives each realized transfer equal weight and adds no covariates beyond the matched-support construction. It includes acquirer and two-digit pre-transfer-industry-by-transfer-year effects, with standard errors clustered by acquirer; equal-mass weighting within support cells is a robustness check. The identifying restriction is that, absent COVID-19, the difference between acqui-hires and comparison flows in the task-exposure gradient would not have varied systematically with event timing.

After singleton removal, the 24-month sample contains 7,926 events. Panel B of Table 4 reports $\widehat{\beta}_{24}^{D} = -0.136$ (s.e. 0.065) for pairwise co-retention. A one-standard-deviation increase in both disruption exposure and the pre-transfer task index implies a 2.8-percentage-point more negative change for acqui-hires than for comparison flows. The comparison-flow gradient is -0.009 (s.e. 0.024), while the acqui-hire gradient is -0.146 (s.e. 0.064); the differential is not generated by improving retention among comparison flows. The estimate for the fraction retained is -0.095 (s.e. 0.047). Threshold and worker-level estimates have the same sign (Appendix Tables A.15 and A.12).

The estimate is similar after excluding first-quarter 2020 transactions and using equal-mass weighting within matching cells; no analogous gradient appears in four pre-COVID placebo windows. As in the within-acqui-hire design, estimates at 36 and 48 months are smaller and imprecise. Appendix Table A.14 and Figure A.2 report these diagnostics.

6.3.3 Selection on target and source-firm characteristics

Comparison flows can originate from much larger or older firms than acqui-hire targets. I therefore construct a second sample using coarsened exact matching on pre-transfer industry

[26] Appendix D.2.4 gives the coarsening rules and discusses the underlying group objects.

and year, target- or source-firm size and age, and mover composition (Iacus et al., 2012). It contains 876 acqui-hires and 3,215 comparison flows. With the matching weights and corresponding controls and fixed effects, the 24-month differential is −0.261 (s.e. 0.098) for pairwise co-retention, −0.153 (s.e. 0.069) for the fraction retained, and −0.121 (s.e. 0.050) at the worker level. All four pre-COVID placebo estimates are indistinguishable from zero. The estimates are similar when balance is imposed on arriving-group characteristics or on target- and source-firm characteristics, making either composition margin alone an unlikely explanation (Appendix Table A.16).

6.4 Career Progression After Transfer

Retention can reflect productive team preservation, contractual lock-in, or weak outside mobility. I therefore ask whether the retention response is accompanied by career progression inside the acquirer. Within acqui-hires, an analogous specification that also allows the task-exposure gradient to vary with workforce composition yields cumulative-seniority-gain estimates of −0.065, −0.092, and −0.121 at 24, 36, and 48 months; all are statistically significant and none of the corresponding pre-COVID placebo estimates is significant (Appendix Table A.17).

For the cross-form analysis, I measure cumulative positive seniority changes and promotion counts at the acquirer, excluding the rank change upon arrival. Because both outcomes scale with group size, I estimate

$$\mathbb{E}\left[\sum_{i\in M_e} Y_{ieh} \mid Z_{eh}\right] = |M_e| \exp(Z'_{eh}\vartheta_h)\,, \tag{19}$$

where Z_{eh} contains the triple-difference terms and fixed effects in equation (18). The group size offset makes the coefficient a differential change in the expected outcome per worker.

Table 5 reports 24-month log-mean interactions of −0.765 (s.e. 0.283) for seniority gains and −0.904 (s.e. 0.267) for promotions. A representative one-standard-deviation increase in both exposure dimensions multiplies the acqui-hire-to-comparison ratio of expected per-worker

outcomes by 0.86 and 0.83, respectively. The interactions remain negative at 36 and 48 months, and worker-level estimates are similar (Appendix Table A.12).

Career progression therefore declines along the same differential gradient as retention. The parallel decline in retention and advancement is difficult to reconcile with contractual lock-in alone.[27]

6.5 Interpretation

During 2020–2022, acqui-hire incidence shifts toward less exposed industries. Among groups transferred beforehand, retention and career progression fall more sharply with task exposure for acqui-hire teams than for comparison flows. The post-transfer results provide the sharper evidence on the team-complementarity margin because organizational form is chosen before the disruption. Section 7 asks whether the same coordination-sensitive relationships support innovative production inside the acquirer.

7 Team Complementarity and Innovation

Section 6 shows that the coordination disruption reduced the preservation of acqui-hired teams. This section asks whether the same disruption reduced innovation inside the acquirer. I follow inventive output from patents produced within the incoming group, to patents combining incoming and pre-existing acquirer inventors, and finally to patents produced independently by pre-existing acquirer inventors in technological fields brought by the incoming group. The sequence separates the team's direct contribution, its integration with the acquirer's inventive base, and internal knowledge diffusion.

[27] Earlier departure mechanically reduces time available for advancement, so these outcomes do not separately identify employer investment or worker welfare.

7.1 Predictions for Innovative Production

In the framework, coordination productivity q scales the return to the team-specific capital L_i embodied in productive relationships among incoming workers. A decline in q should therefore reduce innovation that depends on several incoming inventors. The decline should be larger following an acqui-hire if acquisition preserves a greater stock of those relationships, or makes them more productive at the margin, than a comparison worker flow. The same force can reduce production that combines the incoming group with the acquirer's inventive base. Mixed patents measure that integration. Patents subsequently produced by acquirer inventors without an incoming co-inventor provide a separate test of internal knowledge diffusion. Lucas and Moll (2014) model knowledge growth through meetings. In this setting, co-invention is an observable interaction through which knowledge may spread; the reduced-form response combines changes in the frequency and productivity of those interactions.

Because these outcomes are produced after transfer, I measure coordination exposure using the incoming inventors' roles at the acquirer, C_e^A. This differs from the source-role exposure used for the cross-form retention analysis. Section 6 asks whether pre-transfer relationships remain intact; this section asks how those relationships are used inside the acquirer. The index measures the coordination requirements of post-transfer work, not team complementarity itself.[28]

7.2 Coordinated Production by Incoming Inventors

Production within the incoming group. For an acqui-hire, the incoming inventor group is defined at the transaction level; for a comparison flow, it is the corresponding incoming worker group. A group-internal patent names at least two inventors from that group and no pre-existing acquirer inventor. The volume outcome counts inventor–patent incidences involving incoming inventors, with member-years as exposure, and therefore measures the

[28] A version based on pre-transfer occupations provides a timing check. Each index is standardized within its empirical design, so coefficients should not be compared mechanically across specifications.

average member's rate of participation in patents produced with another group member. I also weight the same incidences by adjusted forward citations (Trajtenberg, 1990) and by the private patent-value measure of Kogan et al. (2017).

Let A_e indicate an acqui-hire, S_{eh} denote the disruption overlap of the h-month outcome window, and P_{eh} denote member-year exposure. For outcome m, I estimate

$$\mathbb{E}[Y_{eh}^m \mid \mathcal{X}_e] = P_{eh} \exp\Big\{\beta_h^m A_e S_{eh} C_e^A + \mathcal{L}_h^m(A_e, S_{eh}, C_e^A) + \xi_h^m \log(1 + n_e) + \lambda_{d(e),y(e)}^m\Big\}, \quad (20)$$

where $\mathcal{L}_h^m$ contains every constituent lower-order term, n_e is incoming-inventor count, and $\lambda_{d(e),y(e)}^m$ are two-digit acquirer-industry-by-transfer-year effects. Standard errors are clustered by acquirer. The coefficient β_h^m is the acqui-hire-versus-flow difference in the disruption-by-task-exposure gradient, not an average effect of acquisition. Identification requires that, absent COVID-19, this differential timing gradient would not have varied systematically with coordination requirements.

Panel A of Table 6 reports negative volume, impact, and value coefficients throughout; all but one are statistically distinguishable from zero. At 48 months, the volume coefficient is -3.166. Since the maximum disruption overlap is one-half, a one-unit difference in standardized acquirer-role exposure multiplies the acqui-hire-versus-flow incidence-rate ratio by $\exp(-3.166 \times 0.5) = 0.205$. This is a change in a ratio of patenting rates, not in either group's patent count. The corresponding coefficients are indistinguishable from zero in the four pre-COVID placebo windows and are similar under alternative disruption windows, exposure denominators, occupation timing, and industry effects (Appendix Table A.18).

Production combining incoming and acquirer inventors. A mixed patent names at least two incoming inventors and at least one inventor already employed by the acquirer; its benchmark is solo output by one incoming inventor with no acquirer inventor. The design compares mixed and solo output within an analysis unit and then differences that response across the two event categories. An acqui-hire unit is an acquirer-by-entry-year cohort, while

a comparison-flow unit remains the incoming inventor group. The broader acqui-hire unit reflects the acquirer-wide inventive base and provides support for a within-unit contrast: mixed output is observed for only 16 of 1,365 comparison groups at the individual-flow level. This aggregation is used only for mixed production; group-internal output remains transaction level.

The reported coefficient is the acqui-hire-versus-flow difference in the mixed-versus-solo response to disruption exposure and acquirer-role coordination requirements. Analysis-unit effects absorb permanent differences in quality and overall patenting. Weights equalize the two event categories within support cells defined by pre-transfer industry, inventor count, and task exposure. Panel B reports volume coefficients of -0.735, -2.238, and -3.926 at 24, 36, and 48 months. At 48 months, maximum disruption overlap implies 1.96 fewer top-coded incoming-inventor–patent incidences in the acqui-hire-versus-flow difference between mixed and solo production. All six impact- and value-weighted estimates are negative and significant at the one-percent level. These outcomes are inverse-hyperbolic-sine transformed, so their coefficients are not percentage effects. The corresponding pre-period placebo estimates show no systematic pattern. The two designs show that the disruption reduces both production within incoming groups and production that integrates them with the acquirer's inventive base.

7.3 Internal Diffusion to Acquirer Inventors

Mixed patents show that incoming and acquirer inventors produce together. To ask whether incoming expertise subsequently appears more broadly in the acquirer's inventive base, I study patents produced by pre-existing acquirer inventors that do not name any incoming inventor. I identify these inventors through observed co-invention with an incoming inventor, then compare their production in incoming-expertise fields with production in their other active fields.

The treated fields are defined without using the technological field in which post-transfer

co-invention occurs. For aggregate and new-field patent counts, I use the acquirer-level pool of expertise fields brought by incoming groups before April 2020. Aggregate production counts patents in these fields; new-field production restricts attention to fields new to the acquirer inventor. For value deepening, a high-impact incoming-expertise field becomes treated when the relevant group arrives; the outcome is patent value in fields where the acquirer inventor had already worked. Adjusted-citation-weighted deepening provides a parallel value measure.[29]

Let T^o_{veft} indicate that field f contains the incoming expertise relevant to outcome o for acquirer inventor v. At horizon h, I estimate the acquisition and comparison-flow versions of

$$Y^o_{veft}(h) = \beta^o_{ah} T^o_{veft} \, \text{Post}_{et} \, S_{eh} C^A_e + \mathcal{L}^o_{ah}(T^o_{veft}, \text{Post}_{et}, S_{eh}, C^A_e) + \mathcal{F}^o_{ah} + \varepsilon^o_{veft}, \tag{21}$$

and report $\beta^o_{Ah} - \beta^o_{Fh}$. Count specifications include separate technology-class-by-year effects for acqui-hires and comparison worker flows, as well as inventor-by-treated-field effects. Value specifications additionally include inventor-by-year effects, which absorb changes in total annual patenting. Standard errors are clustered at the acquirer-by-event-category level. Identification comes from an acquirer inventor's allocation of output between incoming-expertise and other fields, combined with event timing, task exposure, and the comparison worker flows. It requires that no other shock changes this within-inventor field allocation through the same four-way interaction.

Table 7 reports a progressively larger decline in aggregate production: -0.993, -1.662, and -2.669 at 24, 36, and 48 months. At 48 months, maximum disruption overlap and a one-unit increase in acquirer-role exposure imply a 1.33-patent decline in cumulative output in incoming-expertise fields, relative to the inventor's other fields, in the acqui-hire-versus-flow difference. New-field production also declines, with significant estimates at 36 and

[29]When an acquirer completes several acqui-hires, assigning an acquirer inventor to one incoming group would require the realized post-transfer co-invention link. Because that link can respond to the coordination disruption, I instead pool expertise at the acquirer level. Count fields are limited to expertise brought before the shock; high-impact fields used for value outcomes become active only when the relevant incoming group arrives.

48 months. Private-value deepening is negative and significant at every horizon, including −4.238 at 24 months; the citation-weighted measure yields the same conclusion. The four pre-COVID placebo windows yield no corresponding pattern across these outcomes. The delayed quantity and breadth responses resemble the adjustment pattern emphasized in the technology-adoption J-curve literature, while the value response is already visible at 24 months (Brynjolfsson et al., 2021).

Because the sample identifies acquirer inventors through observed co-invention, these estimates capture the reduced-form response of their field allocation; they are not the causal effect of a particular collaboration. In the language of Lucas and Moll (2014), the evidence does not separately identify changes in meeting frequency from changes in the gain per meeting. It does show that coordination disruption reduces independent production by acquirer inventors in fields brought by the incoming group, over and above its effect on patents that directly name incoming inventors.

7.4 From Internal Innovation to External Diffusion

Team complementarity supports innovation at three margins inside the acquirer: production within the incoming group, production combining incoming and acquirer inventors, and subsequent production by acquirer inventors in fields brought by the group. The first two margins map to the productivity and integration of preserved team-specific capital; the third is consistent with internal knowledge diffusion through interaction. These results do not determine how much knowledge subsequently reaches firms outside the acquirer. Section 8 turns to that external margin.

8 Reallocation Beyond the Acquirer

Sections 6 and 7 show that coordination-sensitive relationships support retention and innovative production inside the acquirer. Those relationships need not remain inside the firm

permanently. Acquisition can transfer control over the target and target-held assets, but not ownership of its workers. Post-acquisition retention is therefore an endogenous margin of the firm boundary, jointly determined by the acquirer's demand for the team and workers' outside opportunities.

In the framework of Section 5, a decline in the productivity of coordination reduces the relationship stock that the acquirer chooses to preserve. Worker departure is one margin through which that stock adjusts. I first use the COVID-19 disruption to study exit and entrepreneurial reallocation. I then trace two routes by which innovation reaches beyond the acquirer: patents produced after workers leave and nonaffiliate follow-on to patents produced inside. A final analysis studies the local markets that receive acqui-hire teams.

8.1 Exit and Entrepreneurial Reallocation

Empirical design. I use the common-support sample from Section 6.3.2, which contains 1,133 acqui-hires and 7,095 comparison worker flows. The comparison flows provide a benchmark for contemporaneous changes in group mobility. The design asks whether the COVID-19 disruption changed mobility more for acqui-hire teams whose pre-transfer occupations combine low teleworkability with more interaction-intensive tasks than for similarly exposed comparison groups.

For each horizon $h \in \{24, 36, 48\}$, I estimate the worker-level analogue of equation (18). The coefficient of interest is on $A_e S_{eh} C_e^S$, where A_e indicates an acqui-hire, S_{eh} is the share of the h-month outcome window that overlaps April 1, 2020 through April 1, 2022, and C_e^S is the standardized task-based exposure to in-person coordination constructed from the group's pre-transfer occupations. The specification includes every constituent lower-order term. The identifying restriction is the same as in Section 6: absent the disruption, the acqui-hire–comparison-flow difference in the relationship between task exposure and worker mobility would not have varied with transfer timing. The coefficient therefore measures a differential mobility response, not the average effect of acquisition.

The worker regressions include acquirer, two-digit pre-transfer-industry, and transfer-year fixed effects and control for predetermined seniority, modeled compensation, experience, and occupational-skill measures. These worker characteristics account for observable differences in outside options and baseline mobility. Each observation receives its event's common-support weight divided by the number of workers in the event, so larger groups do not receive more aggregate weight solely because they contain more workers. Standard errors are two-way clustered by acquirer and transfer event.

I examine three margins. Exit is the exact complement of the retention outcome in Section 6.3.2: it indicates that the worker is no longer employed within the acquirer's corporate family at horizon h. Co-departure indicates that the worker and at least one other member of the original transferred group have both left by h, with their recorded exit dates no more than 180 days apart.[30] The third outcome is defined at the transfer-event level and indicates that at least two original members leave the acquirer's corporate family by h and subsequently appear in entrepreneurially oriented destinations or roles. [31] For the event-level outcome, I estimate an unweighted linear probability model with the same triple interaction and lower-order terms, acquirer and pre-transfer-industry-by-transfer-year fixed effects, and standard errors clustered by acquirer. The interacted fixed effects absorb industry-specific changes in labor-market and entrepreneurial conditions across transfer cohorts. Common-support weighting and a stricter outcome that requires the outside destination to begin by h serve as robustness checks.

Exit and temporally proximate departure. Exit responds more strongly to the disruption among coordination-intensive acqui-hire teams than among the comparison flows. At 24 months, the triple-interaction coefficient is 0.098 with a standard error of 0.048 (Table 8). As in Section 6, I scale the triple-interaction contrast by the dispersion of both disruption

[30]Co-departure does not require a common destination and therefore captures temporally proximate departure, not necessarily coordinated movement to another firm.

[31]The measure combines founder and leadership roles at young or small firms with external founder and startup links; Table 8 gives the exact definition. It neither requires a common destination nor implies that the workers jointly founded a firm.

exposure and task-based exposure in the estimating sample. For simultaneous one-standard-deviation increases in the two variables, the triple-interaction component corresponds to a 2.0-percentage-point increase in differential exit. This is the mobility counterpart to the negative worker-retention estimate in Section 6.3.2.

The corresponding co-departure coefficient is 0.130 with a standard error of 0.040. For the analogous simultaneous one-standard-deviation contrast, the triple-interaction component is 2.6 percentage points. The differential adjustment therefore includes an increase in departures occurring within six months of another original group member's departure.

Entrepreneurial reallocation. The disruption also raises the joint incidence of exit and subsequent entrepreneurial reallocation. At 24 months, the coefficient for at least two members exiting and later entering entrepreneurially oriented destinations or roles is 0.241 with a standard error of 0.113. For simultaneous one-standard-deviation increases in disruption exposure and task-based exposure, the triple-interaction component is 4.8 percentage points. Under common-support weighting, the estimate is 0.259 with a standard error of 0.121. It is 0.190 with a standard error of 0.111 when the outside destination must begin within 24 months. Appendix Table A.19 shows similar positive estimates when the outcome requires at least one or at least three members.

Across all three outcomes, none of the four pre-2020 placebo estimates is statistically distinguishable from zero at the 10 percent level (Appendix Figures A.3–A.5). The differential adjustment is detected at 24 months and estimated imprecisely at 36 and 48 months. The later-horizon estimates are not directly comparable because maximum disruption exposure falls as the outcome window lengthens.[32]

Entrepreneurial careers. The shock design estimates a differential response to the coordination disruption. Descriptive level differences provide a separate view of the subsequent careers of workers joining through an acqui-hire. By 48 months, the regression-adjusted

[32]Maximum overlap falls from 1.00 at 24 months to 0.67 and 0.50 at 36 and 48 months.

difference between acqui-hires and comparison flows in entry into an entrepreneurially oriented destination or role is 11.1 percentage points (Appendix Table A.20). Within acqui-hires, workers with a prior founder, CEO, or CTO role are 15.3 percentage points more likely to make such a transition by 48 months, and a one-standard-deviation increase in pre-transfer seniority is associated with a 3.1-percentage-point increase. These are conditional level differences, not shock heterogeneity.

These mobility outcomes combine firms' employment choices with workers' willingness to remain and therefore describe a joint adjustment of the firm boundary. The next subsection asks whether innovative knowledge also crosses that boundary.

8.2 Innovation Direction and External Follow-on

The preceding evidence identifies the productive role of preserving the team inside the acquirer. Innovation produced after transfer also moves toward the acquirer's pre-transaction technology, raising the question of whether external follow-on varies with technological direction. Worker exit provides a second route by which knowledge associated with the incoming group appears in outside inventions. I examine both margins along the observed transfer path.

Technological continuity after exit. I define a post-exit patent as a U.S. utility patent that names an original group member, is filed after that worker's sustained exit from the acquirer's corporate family, and is assigned at filing to a firm outside that family. I compare the patent's claims text with two reference portfolios. The first contains pre-transfer patents linked to original group members and assigned to the target in an acqui-hire or to the source firm in a comparison flow. The second contains patents produced inside the acquirer by *other* original group members before the focal worker exits. Excluding the focal worker from the second portfolio separates technological relatedness to the rest of the group from persistence in the worker's own research.

I benchmark each focal similarity against up to five placebo groups matched on pre-transfer

industry, transfer cohort, and target- or source-firm patent-portfolio size, using only patents filed before the focal worker's exit. I average excess similarity within each transfer event and weight events equally. Appendix B.8 provides the full construction. The analysis is conditional on exit, subsequent outside patenting, and observable reference portfolios.

The acqui-hire–comparison-flow difference in continuity with the incoming group's pre-transfer technology is −0.0012 (standard error 0.0040) and statistically indistinguishable from zero (Table 9). By contrast, the difference is 0.0243 (0.0056) when the reference is patents produced inside the acquirer before exit by other original group members. Among selected patenting exiters, post-exit patents following an acqui-hire are therefore more closely related to the pre-exit work of other original group members. This is descriptive technological continuity, not a causal estimate of knowledge transfer. The pattern is consistent with prior evidence associating inventor mobility with knowledge flows across firm boundaries (Almeida and Kogut, 1999; Singh and Agrawal, 2011).

Technological direction and external follow-on. Patents produced inside the acquirer can also reach other firms through nonaffiliate follow-on. For patent p from acquisition e, let Z_{pe} be the standardized claims-text similarity between the patent and the acquirer's pre-transaction patent portfolio. With acquisition fixed effects, β_δ is estimated from variation across patents produced following the same transaction. I estimate

$$\mathbb{E}\Big[C_{pe}^{\text{ext}} \mid Z_{pe}, e, f(p), t(p)\Big] = \exp\Big\{\alpha_e + \lambda_{f(p),t(p)} + \beta_\delta Z_{pe}\Big\}, \tag{22}$$

where C_{pe}^{ext} counts wholly nonaffiliate citing patents. I interpret this outcome as citation-based external follow-on per patent. The specification absorbs acquisition and technology-field-by-filing-year fixed effects, and standard errors are clustered by acquirer; Table 10 gives the citation-classification and exclusion rules. Because project direction is endogenous, β_δ is a conditional within-acquisition association, not the causal effect of redirecting a project.

Within the patent-producing acquisitions in the estimating sample, a one-standard-

deviation increase in proximity to the acquirer's pre-transaction technology predicts 21 percent fewer nonaffiliate citations over the full observed window (coefficient −0.237, standard error 0.136) and 23 percent fewer during the 36-month-capped window (−0.262, 0.110) (Table 10). Estimates for receiving any nonaffiliate citation and for the number of distinct outside citing families are imprecise. The result therefore concerns citation intensity, not the extensive margin or breadth of outside use. Appendix Table A.21 reports outcome and scope sensitivity.[33]

Innovation output and total external follow-on. The negative within-acquisition direction gradient concerns external follow-on per patent, not external follow-on per original worker. Because external follow-on per worker also depends on innovation output, I construct three quantities using a common sample and citation window:

$$Q_h^{\text{in}} = \frac{P_h^{\text{in}}}{N}, \qquad \delta_h^{\text{in}} = \frac{C_h^{\text{ext}}}{P_h^{\text{in}}}, \qquad D_h^{\text{in}} = \frac{C_h^{\text{ext}}}{N} = Q_h^{\text{in}} \delta_h^{\text{in}}. \tag{23}$$

Here N is the number of original incoming-worker observations across the corresponding transfers. P_h^{in} counts acquirer-family patents filed within h months by an original group inventor and early enough to permit a complete 36-month citation window; C_h^{ext} counts wholly nonaffiliate citations during that window. Events with no eligible patent remain in N. These definitions make the identity in equation (23) exact.

At 48 months, acqui-hires produce 0.097 eligible inside patents per original incoming worker, compared with 0.042 for comparison flows, a ratio of 2.31 (95 percent confidence interval 1.55 to 3.85) (Table 11). Nonaffiliate citations per inside patent have a ratio of 1.33 (0.80 to 2.23) and are not statistically different between acqui-hires and comparison flows. Nonaffiliate citations per original incoming worker are 0.068 versus 0.022, a ratio of 3.07 (1.68 to 5.96).

[33] Patent citations measure citation-based follow-on, not the full economic value of knowledge spillovers (Jaffe et al., 1993; Hall et al., 2005).

Across observed transfer paths, acqui-hires therefore generate more inside patents and more nonaffiliate follow-on per original incoming worker, while follow-on per patent does not differ detectably. The accounting comparison combines output and citation-based follow-on; the within-acquisition estimate relates technological direction to follow-on per patent; and post-exit patenting records a separate route through mobile workers. These observed-path margins do not recover the no-acquisition outcome, under which team assembly, target continuation, project direction, and worker mobility would all adjust.

8.3 Local Market Outcomes

Worker mobility and patent follow-on trace two routes beyond the acquirer. I next ask whether the local industry-markets receiving acqui-hire teams also expand. The outcomes are establishments, employment, payroll, and firm entry; they capture local economic incidence rather than a direct measure of knowledge diffusion.

Market definition and empirical design. I use the QCEW data described in Section 2.5. A market is the target's harmonized four-digit industry before acquisition interacted with a fixed-definition Core-Based Statistical Area. The relevant location is the modal post-transfer area of the target's transferred high-skill workers. I therefore use the group's post-transfer rather than pre-transfer location. [34]

The first qualifying acqui-hire observed in a market defines its treatment year, g. I omit the transaction year, which may be only partially treated, and examine the five complete calendar years from $g+1$ through $g+5$. Subsequent acqui-hires are allowed. The estimand therefore describes a market's outcomes after its first observed acqui-hire.

I estimate synthetic difference-in-differences separately within blocks defined by three-digit industry and treatment year (Arkhangelsky et al., 2021). Within each block, treated four-digit-industry–area markets are compared with markets in the same three-digit industry

[34]The post-transfer and pre-transfer areas coincide in 71.9 percent of qualifying transactions.

that never receive an observed acqui-hire. Synthetic weights align pre-treatment outcome paths, and block-level estimates are aggregated using treated-market counts. Inference uses a delete-one-area jackknife over post-transfer areas; Appendix D.3 provides the donor-screening and inference details.

The identifying restriction is that, absent first exposure to an acqui-hire, treated markets would have continued to follow their synthetic comparisons; an unobserved local shock that both attracts an acqui-hire and subsequently raises market activity would violate it. Figure 9 shows four pre-treatment coefficients and five post-treatment coefficients for each outcome. The pre-treatment coefficients are small and jointly indistinguishable from zero: the joint-test p-values are 0.308 for establishments, 0.960 for employment, and 0.997 for payroll. These patterns support the counterfactual fit.

Local expansion. The estimated establishment gaps are positive throughout the post-treatment period (Figure 9 and Table 12). The coefficient rises from 0.0219 in the first complete post-transaction year to 0.0478 in the fifth. The latter estimate has a standard error of 0.0259 and a p-value of 0.065. It corresponds to a 4.9 percent increase relative to synthetic comparison markets. Averaging over the five post-treatment years yields 0.0327 log points (standard error 0.0176).

The estimated post-treatment gaps are larger for employment and payroll. By the fifth complete post-transaction year, employment is 0.0835 log points higher (standard error 0.0275), and total payroll is 0.0957 log points higher (standard error 0.0270). These estimates correspond to increases of 8.7 percent and 10.0 percent, respectively. The three outcomes therefore indicate growth in both the number of operating locations and the scale of local activity. Payroll is the market's total wage bill, not pay per worker, so the payroll estimate is not an individual wage effect.

Firm births and young-firm activity. QCEW establishment counts include both independent firms and additional locations of existing firms. I therefore repeat the design using

BDS sector–metropolitan-area panels, which report firms and establishments by firm age. The BDS analysis uses the same transaction-year exclusion and five-year post-treatment window, but its industry definition is necessarily coarser.

The average post-treatment number of firms aged zero to five is 0.0566 log points higher in receiving markets (standard error 0.0154), and the number of establishments operated by those firms is 0.0539 log points higher (standard error 0.0178). These coefficients correspond to increases of 5.8 and 5.5 percent. More directly, the number of age-zero firms—new employer firms in the BDS—is 0.0504 log points, or 5.2 percent, higher (standard error 0.0251). The age-zero result provides evidence of a firm-entry margin. The broader zero-to-five category also reflects the subsequent survival and aging of young firms and is not itself a count of contemporaneous entry.

The firm-entry margin is the local-market outcome most closely related to the entrepreneurial reallocation in Section 8.1. [35]

Markets receiving an acqui-hire team subsequently expand in establishments, employment, payroll, and new employer firms relative to their synthetic comparisons. This local incidence is distinct from the worker-mobility and patent-follow-on evidence above; the three designs do not require one margin to mediate another.[36]

9 Conclusion

This paper studies why firms use acquisition to obtain innovative teams and how that organizational choice shapes the subsequent location of workers and knowledge. The framework isolates two advantages of acquisition over direct team hiring. A common closing can coordinate the participation of workers whose value depends on joining together, while ownership gives the acquirer residual control over technology and other nonhuman assets held by the

[35]The two designs do not form a mediation chain: BDS does not link firm births to former members of the incoming team, and the local-market treatment begins when the team arrives rather than when a worker later leaves.

[36]The local estimates describe receiving markets rather than aggregate net creation; activity may partly be reallocated from other markets.

target. These advantages change the firm boundary without transferring ownership of labor: workers choose whether to join and whether to remain.

I construct a documentary measure of transactions in which obtaining the target's employees, founders, or team is a stated motive, and link the resulting acqui-hires to worker histories and patent records. The descriptive evidence reveals a consistent organizational pattern. Acqui-hire targets contain denser prior worker relationships than other acquisition targets. Relative to comparison worker flows, acqui-hires assemble more of those relationships at the acquirer, initially retain more of the incoming group, and are followed by greater career progression and patented output. The resulting inventions both continue the incoming group's prior work and move toward the acquirer's technology.

The COVID-19 disruption to in-person work provides the paper's central test of team complementarity. During 2020–2022, acqui-hire activity shifted away from industries more exposed to nonteleworkable work, both in the broad acquisition-count panel and after normalizing by observed worker-group transfers. Among groups transferred before the disruption, retention and career progression respond more negatively to exposure for acqui-hire teams than for comparison flows. Patented production within the incoming group, production combining incoming and pre-existing acquirer inventors, and later innovation by the acquirer's existing inventors also decline. The labor outcomes show that preservation of the incoming group inside the acquirer responds to coordination conditions; the patent outcomes show that the relationships preserved in acqui-hires support coordinated production.

The firm boundary subsequently adjusts. When the coordination disruption is more binding, workers who entered through an acqui-hire are more likely to leave, depart near another original group member, and enter entrepreneurially oriented work. Among exiters who later patent, post-exit patents following acqui-hires are more closely related than those following comparison flows to other group members' pre-exit work. Separately, within acquisitions, patents closer to the acquirer's prior technology receive less nonaffiliate follow-on per patent. Across acqui-hires and comparison flows, acqui-hires produce substantially more

inside patents; with no detectable per-patent follow-on difference, nonaffiliate follow-on per original incoming worker is also higher along the observed path. Local markets receiving an acqui-hire team subsequently expand in establishments, employment, payroll, and new employer firms relative to synthetic comparisons.

An acqui-hire is therefore a dynamic organizational choice, not a permanent allocation of talent. Acquisition can preserve a complementary team and place target technology under the acquirer's control at closing, but workers and knowledge can later cross the boundary again. The innovation consequences depend jointly on team assembly, project direction, internal production, and subsequent mobility.

This distinction matters for merger policy. Static measures of talent concentration omit the value of preserving team-specific capital, while post-transaction patent counts omit technological direction and external follow-on. Evaluating a prohibition therefore requires an explicit no-acquisition counterfactual in which team assembly, target continuation, project direction, and worker mobility can all adjust. The reduced-form evidence identifies the central margins in that counterfactual; it does not by itself determine net innovation or total welfare.

References

Aggarwal, V. A., D. H. Hsu, and A. Wu (2020). Organizing knowledge production teams within firms for innovation. *Strategy Science 5*(1), 1–16.

Aghion, P. and J. Tirole (1994). The management of innovation. *Quarterly Journal of Economics 109*(4), 1185–1209.

Alchian, A. A. and H. Demsetz (1972). Production, information costs, and economic organization. *American Economic Review 62*(5), 777–795.

Almeida, P. and B. Kogut (1999). Localization of knowledge and the mobility of engineers in regional networks. *Management Science 45*(7), 905–917.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review 111*(12), 4088–4118.

Babina, T., A. Fedyk, A. X. He, and J. Hodson (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics 151*, 103745.

Babina, T. and S. T. Howell (2024). Entrepreneurial spillovers from corporate R&D. *Journal of Labor Economics 42*(2), 469–509.

Baghai, R. P., R. C. Silva, and L. Ye (2024). Teams and bankruptcy. *Review of Financial Studies 37*(9), 2855–2902.

Bar-Isaac, H., J. P. Johnson, and V. Nocke (2025). Acquihiring for monopsony power. *Management Science 71*(4), 3485–3496.

Beaumont, P., C. Hébert, and V. Lyonnet (2025). Build or buy? Human capital and corporate diversification. *Review of Financial Studies 38*(5), 1333–1367.

Bena, J., I. Erel, D. Wang, and M. S. Weisbach (2026). Relationship-specific investments and firms' boundaries: Evidence from textual analysis of patents. Working paper.

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

Benkert, J.-M., I. Letina, and S. Liu (2025). Startup acquisitions: Acquihires and talent hoarding. *European Economic Review 178*, 105103.

Bernstein, S. (2015). Does going public affect innovation? *Journal of Finance 70*(4), 1365–1403.

Bloom, N., M. Schankerman, and J. Van Reenen (2013). Identifying technology spillovers and product market rivalry. *Econometrica 81*(4), 1347–1393.

Boyacıoğlu, B., M. N. Özdemir, and S. Karim (2024). Acqui-hires: Redeployment and retention of human capital post-acquisition. *Strategic Management Journal 45*(2), 205–237.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics 117*(1), 339–376.

Brynjolfsson, E., D. Rock, and C. Syverson (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics 13*(1), 333–372.

Chatterji, A. and A. Patro (2014). Dynamic capabilities and managing human capital. *Academy of Management Perspectives 28*(4), 395–408.

Chen, D., H. Gao, and Y. Ma (2021). Human capital-driven acquisition: Evidence from the inevitable disclosure doctrine. *Management Science 67*(8), 4643–4664.

Chen, J., S. Hshieh, and F. Zhang (2024). Hiring high-skilled labor through mergers and acquisitions. *Journal of Financial and Quantitative Analysis 59*(6), 2762–2798.

Coyle, J. F. and G. D. Polsky (2013). Acqui-hiring. *Duke Law Journal 63*(2), 281–346.

Crouzet, N., Z. He, V. Lyonnet, and Y. Ma (2026). Why don't old firms do new things? Working paper.

Cunningham, C., F. Ederer, and S. Ma (2021). Killer acquisitions. *Journal of Political Economy 129*(3), 649–702.

Dijk, E., J. L. Moraga-González, and E. Motchenkova (2024). How do start-up acquisitions affect the direction of innovation? *Journal of Industrial Economics 72*(1), 118–156.

Dingel, J. I. and B. Neiman (2020). How many jobs can be done at home? *Journal of Public Economics 189*, 104235.

Ederer, F., R. Seibel, and T. Simcoe (2025). Digital (killer?) acquisitions. Working paper.

Franco, A. M. and D. Filson (2006). Spin-outs: Knowledge diffusion through employee mobility. *RAND Journal of Economics 37*(4), 841–860.

Frésard, L., G. Hoberg, and G. M. Phillips (2020). Innovation activities and integration through vertical acquisitions. *Review of Financial Studies 33*(7), 2937–2976.

Gompers, P., J. Lerner, and D. Scharfstein (2005). Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999. *Journal of Finance 60*(2), 577–614.

Grossman, S. J. and O. D. Hart (1986). The costs and benefits of ownership: A theory of vertical and lateral integration. *Journal of Political Economy 94*(4), 691–719.

Groysberg, B., L.-E. Lee, and A. Nanda (2008). Can they take it with them? The portability of star knowledge workers' performance. *Management Science 54*(7), 1213–1230.

Hadlock, C. J. and J. R. Pierce (2010). New evidence on measuring financial constraints: Moving beyond the KZ index. *Review of Financial Studies 23*(5), 1909–1940.

Hall, B. H., A. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *RAND Journal of Economics 36*(1), 16–38.

Hart, O. and J. Moore (1990). Property rights and the nature of the firm. *Journal of Political Economy 98*(6), 1119–1158.

He, A. X. and J. S. Xue (2026). Talent-consolidating mergers. Working paper.

Iacus, S. M., G. King, and G. Porro (2012). Causal inference without balance checking: Coarsened exact matching. *Political Analysis 20*(1), 1–24.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics 108*(3), 577–598.

Jaravel, X., N. Petkova, and A. Bell (2018). Team-specific capital and innovation. *American Economic Review 108*(4–5), 1034–1073.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–320.

Kim, J. D. (2022). Startup acquisitions, relocation, and employee entrepreneurship. *Strategic Management Journal 43*(11), 2189–2216.

Kim, J. D. (2024). Startup acquisitions as a hiring strategy: Turnover differences between acquired and regular hires. *Strategy Science 9*(2), 118–134.

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

Kominers, S. D. and E. G. Weyl (2012). Holdout in the assembly of complements: A problem for market design. *American Economic Review 102*(3), 360–365.

Li, K. and J. Wang (2023). Inter-firm inventor collaboration and path-breaking innovation: Evidence from inventor teams post-merger. *Journal of Financial and Quantitative Analysis 58*(3), 1144–1171.

Li, Z., X. Zhang, Y. Zhang, D. Long, P. Xie, and M. Zhang (2023). Towards general text embeddings with multi-stage contrastive learning. arXiv:2308.03281.

Lucas, Jr., R. E. and B. Moll (2014). Knowledge growth and the allocation of time. *Journal of Political Economy 122*(1), 1–51.

Ludwig, J., S. Mullainathan, and A. Rambachan (2025, January). Large language models: An applied econometric framework. Working Paper 33344, National Bureau of Economic Research.

Ma, S., W. Wang, and Y. Wu (2025). Steering labor mobility through innovation. Working paper.

Ng, W. and T. E. Stuart (2022). Acquired employees versus hired employees: Retained or turned over? *Strategic Management Journal 43*(5), 1025–1045.

Ouimet, P. and R. Zarutskie (2020). Acquiring labor. *Quarterly Journal of Finance 10*(3), 2050011.

Phillips, G. M. and A. Zhdanov (2013). R&D and the incentives from merger and acquisition activity. *Review of Financial Studies 26*(1), 34–78.

Pollman, E. (2023). Startup failure. *Duke Law Journal 73*(2), 327–386.

Puranam, P., H. Singh, and M. Zollo (2006). Organizing for innovation: Managing the coordination–autonomy dilemma in technology acquisitions. *Academy of Management Journal 49*(2), 263–280.

Puranam, P. and K. Srikanth (2007). What they know vs. what they do: How acquirers leverage technology acquisitions. *Strategic Management Journal 28*(8), 805–825.

Sawicki, A. (2015). Buying teams. *Seattle University Law Review 38*(2), 651–684.

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

Singh, J. and A. Agrawal (2011). Recruiting for ideas: How firms exploit the prior inventions of new hires. *Management Science 57*(1), 129–150.

Stinson, M., T. K. White, T. C. Fort, C. Goetz, N. Goldschlag, M. Chow, E. R. Perlman, and J. Lawrence (2021). Redesigning the longitudinal business database. Working Paper CES-WP-21-08, U.S. Census Bureau, Center for Economic Studies.

Trajtenberg, M. (1990). A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics 21*(1), 172–187.

United States Patent and Trademark Office (2025a). PatentsView data dictionaries and long-text tables. PatentsView. Accessed August 3, 2026.

United States Patent and Trademark Office (2025b). What's new with PatentsView—December 2025. PatentsView. Quarter 3, 2025 data release, covering data through September 30, 2025.

Whited, T. M. and G. Wu (2006). Financial constraints risk. *Review of Financial Studies 19*(2), 531–559.

Wu, S. and Z. Qian (2026). Talent hoarding and upstream innovation: Labor market distortions by large incumbents. Working paper.

Figure 1. Measurement Population and Analysis Samples

Notes: The unit is a unique PitchBook transaction. The full population contains completed mergers and acquisitions from 2000 through 2025 in which 100 percent of the target is acquired and at least one reported acquirer's ultimate parent is headquartered in the United States. Reported acquiring entities are mapped to their ultimate parents before the geographic restriction is applied, and consortium deals are counted once after all reported acquirers are evaluated. The documentary indicator equals one when transaction text explicitly presents obtaining the target's employees, founders, or team as a motive. The additional measurement component contains transactions involving both public and private acquirers and does not enter the worker-level causal analyses. The classification and Revelio-linkage samples overlap; one is not obtained by sequentially filtering the other.

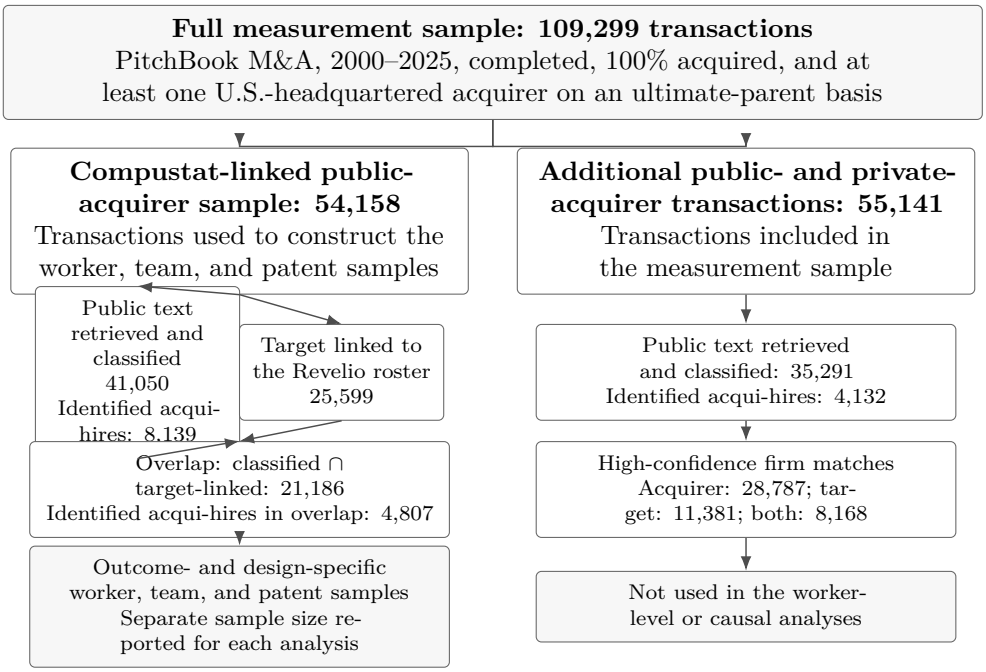


Figure 2. Prior Joint Mobility within Acquisition Targets

Notes: The unit is an acquisition target. An acqui-hire is the paper's documentary proxy: a transaction source explicitly identifies obtaining employees, founders, or the team as a motive. The source population is 27,582 classified acquisitions from 2000 through 2025 that link to worker histories; the figure retains the 20,043 targets with at least two observed high-skill workers during the year before transaction completion (4,891 acqui-hires and 15,152 other acquisitions). A pair has a prior joint-mobility link if its members overlapped for at least 90 days at the same earlier employer, both left that employer before entering the target, and joined the target within 365 days of one another. Prior joint-mobility density is the number of linked pairs divided by all possible pairs in the target's observed high-skill roster. Bars are equal-transaction means within the five displayed roster-size categories; vertical bars are 95 percent confidence intervals. Annotations report the ratio of the acqui-hire mean to the other-acquisition mean. The comparison is descriptive. Transaction data are from PitchBook, acqui-hire labels use the cited transaction sources, and worker histories are from Revelio Labs.

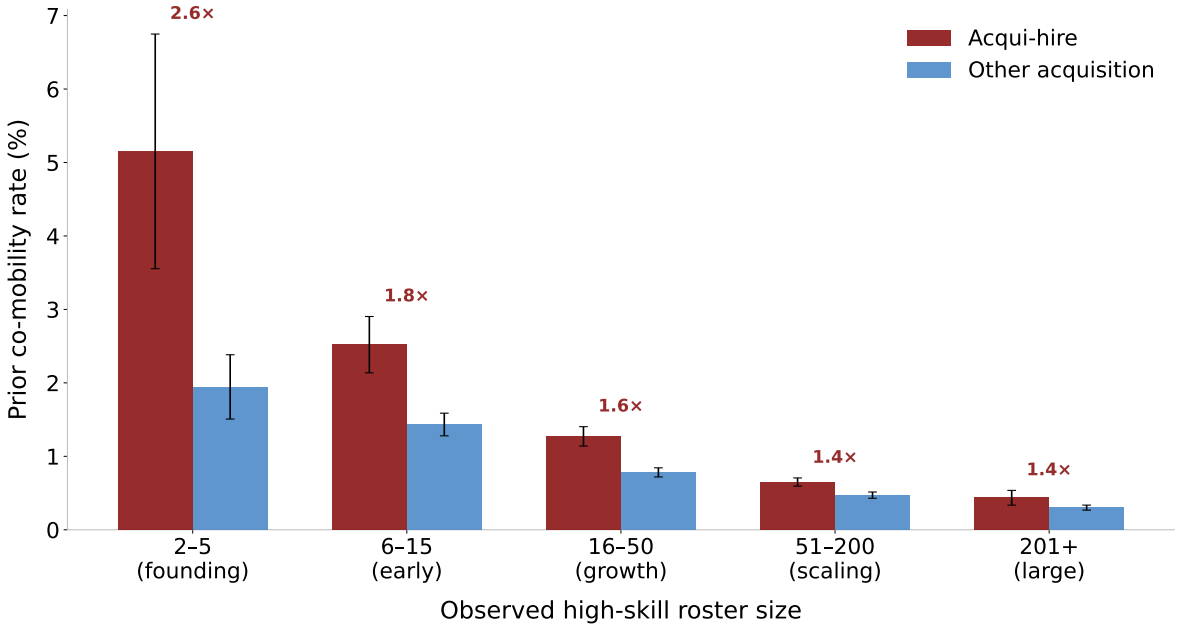


Figure 3. Assembly of Pre-existing Worker Relationships at the Acquirer

Notes: The unit is a matched worker-group event. An acqui-hire is the paper's documentary proxy: a transaction source explicitly identifies obtaining employees, founders, or the team as a motive. The common-support population contains 1,133 acqui-hires and 7,095 comparison worker flows from 2010 through March 2020. A comparison flow is a group of workers from a common source firm who enter the same employer within the acquirer's corporate family within 60 days, without acquisition of the source firm, and have prior-workplace linkage above that of equal-size placebo groups from the source firm; their contracting arrangement is not observed. The displayed sample conditions symmetrically on the pre-transfer group having at least one prior joint-mobility link, leaving 609 acqui-hires and 3,458 comparison flows. A link is an unordered pair whose members overlapped for at least 90 days at a common earlier employer and subsequently entered their pre-transfer employer within 365 days of one another. For acqui-hires, the pre-transfer group is drawn from the target's high-skill roster; for comparison flows, it is the reconstructed pre-transfer group. A link is assembled when both endpoints are observed in the acquirer's corporate group at some point from the event date through day 183; continuous presence for the full window is not required. The four bins are < 25, $[25, 50)$, $[50, 75)$, and at least 75 percent of eligible links assembled. Bars are equal-event shares within each event category and do not report regression-adjusted estimates. Transaction data are from PitchBook, acqui-hire labels use the cited transaction sources, and worker histories are from Revelio Labs.

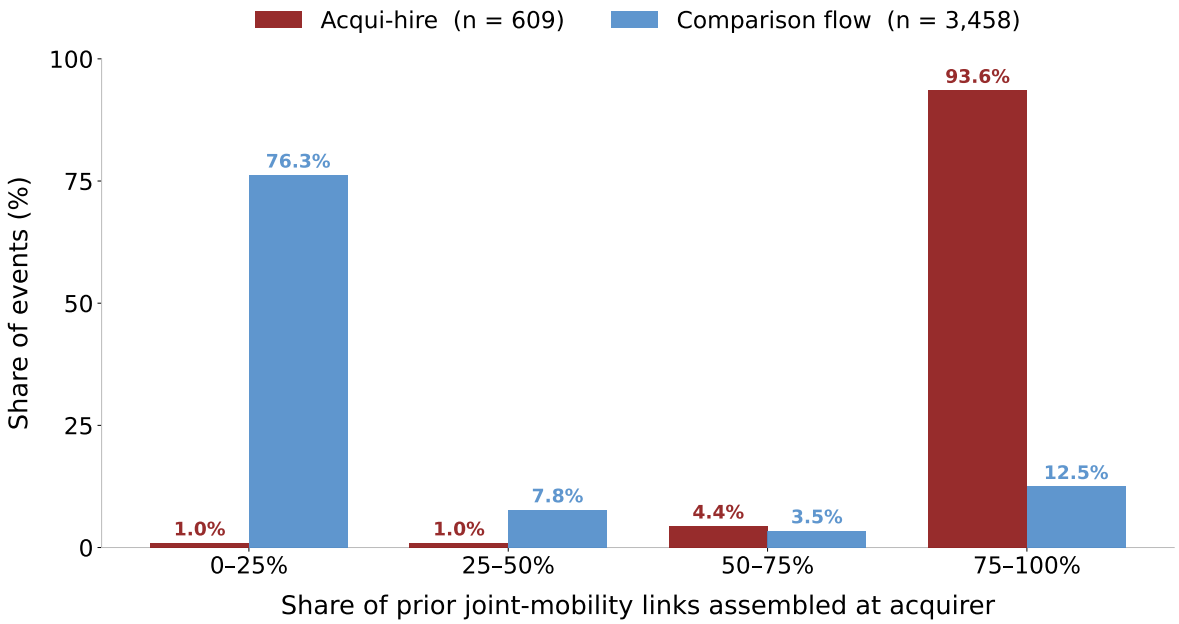
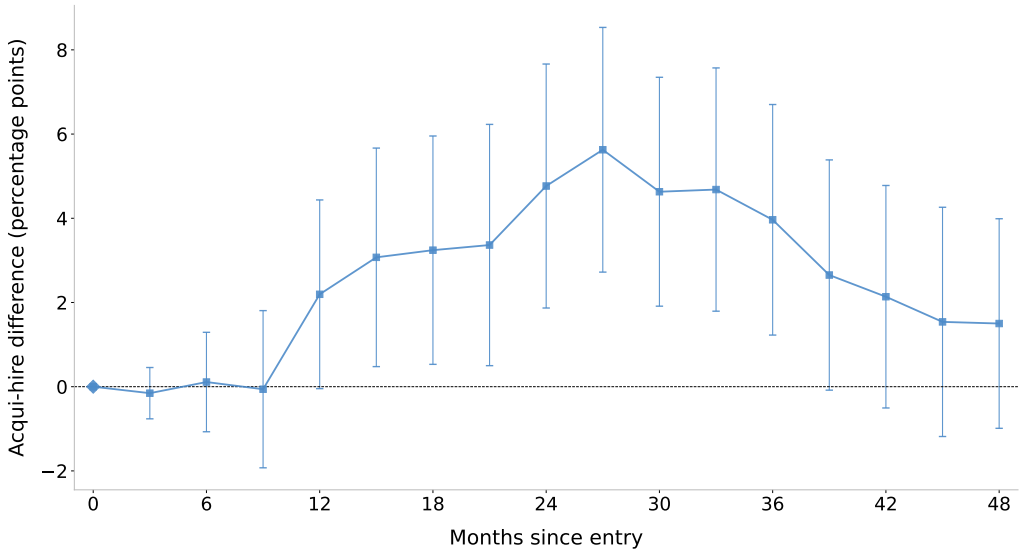
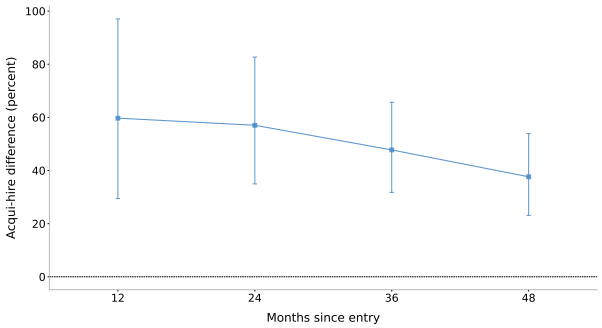


Figure 4. Post-transfer Team Persistence and Career Advancement

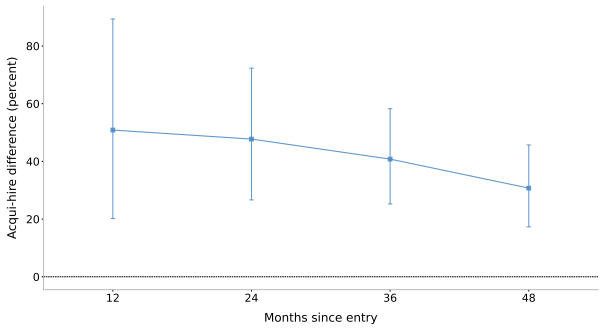
Notes: The common-support population contains 1,133 acqui-hires and 7,095 comparison worker flows from 2010 through March 2020. An acqui-hire is the paper's documentary proxy: a transaction source explicitly identifies obtaining employees, founders, or the team as a motive. A comparison flow is a group leaving a common source firm and entering a common acquirer without acquisition of that source firm; it passes the prior-workplace-link screen, but its contracting arrangement is not observed. Support is imposed on entry year, two-digit pre-transfer industry, arriving-group size, occupation composition, and prior cohesion. Panel A reports the acqui-hire coefficient from separate unweighted team-level linear regressions of pairwise co-retention at each displayed month. Pairwise co-retention is the share of entrant pairs whose two members both remain in the acquirer's corporate group. These regressions absorb acquirer and pre-transfer-industry-by-entry-year fixed effects and cluster standard errors by acquirer; fixed-effect singleton removal leaves 7,926 events. The month-zero diamond is normalized to zero. Panels B and C report worker-level Poisson pseudo-maximum-likelihood estimates for cumulative promotions and cumulative positive changes in within-firm seniority. The rank change recorded at entry is excluded. These regressions condition on pre-entry seniority, log modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status; absorb separate acquirer, pre-transfer-industry, and entry-year fixed effects; and cluster two ways by acquirer and event. They are unweighted and use between 145,524 and 152,539 worker–event observations across horizons. Panel A is in percentage points. Panels B and C report $100[\exp(\widehat{\beta}) - 1]$, the percentage difference in the conditional expected count. Vertical bars are 95 percent confidence intervals. Transaction data are from PitchBook, acqui-hire labels use the cited transaction sources, and worker histories and career outcomes are from Revelio Labs.



(a) Pairwise co-retention



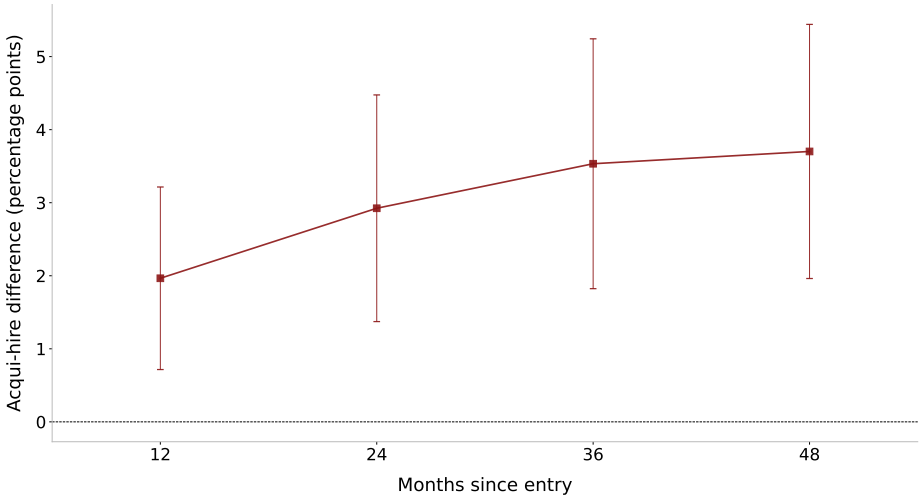
(b) Promotions



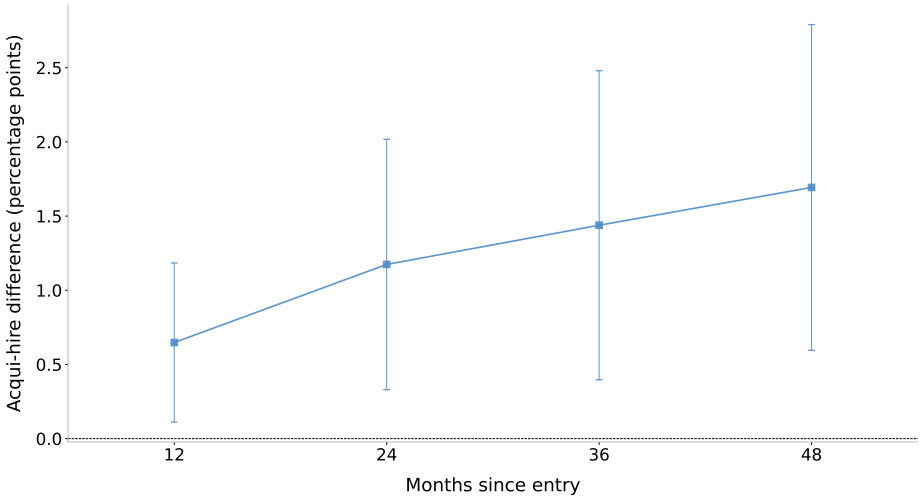
(c) Positive seniority changes

Figure 5. Worker Patenting and Co-Invention at the Acquirer

Notes: The unit is a worker–event observation. The sample is the common-support worker panel spanning 2010 through March 2020 and contains 169,559 effective observations, 142,812 distinct workers, 8,228 events, and 650 acquirers. An acqui-hire is the paper's documentary proxy: a transaction source explicitly identifies obtaining employees, founders, or the team as a motive. The comparison category consists of groups leaving a common source firm and entering a common acquirer without acquisition of that source firm; each group passes the prior-workplace-link screen, but its contracting arrangement is not observed. Support is imposed on entry year, two-digit pre-transfer industry, arriving-group size, occupation composition, and prior cohesion. Panel A's outcome equals one if the entrant is named, within the indicated number of months after joining, on a patent assigned to the acquirer or its parent. Panel B additionally requires a co-inventor outside the incoming cohort who is connected to the acquirer before the entrant arrives, as shown by either an employment spell or a patent previously assigned to that firm. Each point is the acqui-hire coefficient from a linear probability model. The regressions use matching weights normalized by event size, condition on pre-entry seniority, log modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status, absorb acquirer, two-digit pre-transfer-industry, and entry-year fixed effects, and cluster standard errors two ways by acquirer and event. Coefficients are in percentage points and vertical bars are 95 percent confidence intervals. At 12 and 48 months, comparison-group means are 0.27 and 1.07 percent in Panel A and 0.19 and 0.69 percent in Panel B. Transaction data are from PitchBook, acqui-hire labels use the cited transaction sources, worker histories are from Revelio Labs, and patent data are from PatentsView.



(a) Any patent assigned to the acquirer



(b) Co-invention with a pre-existing acquirer inventor

Figure 6. Cohort Patent Output and Observed Private Patent Value

Notes: The unit is a matched worker-group event from the common-support population spanning 2010 through March 2020: 1,133 acqui-hires and 7,095 comparison worker flows. An acqui-hire is the paper's documentary proxy: a transaction source explicitly identifies obtaining employees, founders, or the team as a motive. A comparison flow is a group leaving a common source firm and entering a common acquirer without acquisition of that source firm; it passes the prior-workplace-link screen, but its contracting arrangement is not observed. A qualifying patent names at least one incoming worker, is assigned to the acquirer or its parent, and is filed within the indicated number of months after the cohort begins to enter. A patent linked to m_p focal events contributes $1/m_p$ to each. Output is the resulting patent count per incoming worker. Observed private value is $\sum_p \xi_p/m_p$ per incoming worker, where ξ_p is the market-based patent-value allocation of Kogan et al. (2017). Events with no qualifying patent are true zeros. For the private-value series, events with patents but no covered patent are unobserved rather than coded as zero; a partially covered event contributes its observed subtotal. Each point is $100[\exp(\widehat{\beta}) - 1]$, where $\widehat{\beta}$ is the acqui-hire coefficient relative to comparison flows in a Poisson pseudo-maximum-likelihood regression with log arriving-group size as an exposure. Regressions use event matching weights, absorb acquirer, two-digit pre-transfer-industry, and event-year fixed effects, and cluster standard errors by acquirer. Output specifications use 4,591, 5,006, 5,192, and 5,243 events at 12, 24, 36, and 48 months; private-value specifications use 4,499, 4,837, 5,006, and 5,058 events. The corresponding acquirer counts are 69/102/118/122 for output and 60/85/101/105 for private value. Vertical bars are delta-method 95 percent confidence intervals. Transaction data are from PitchBook, acqui-hire labels use the cited transaction sources, worker histories are from Revelio Labs, patents are from PatentsView, and market-based private patent values are from Kogan et al. (2017).

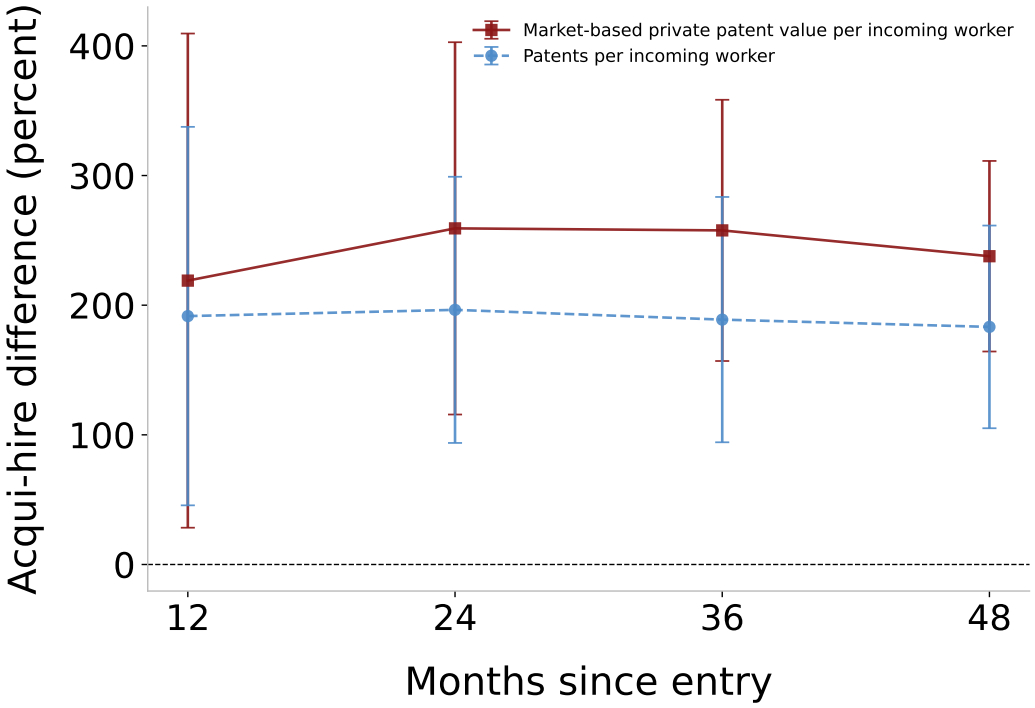


Figure 7. Industry Exposure to Nonteleworkable Work and Acqui-hire Activity

Notes: The figure reports annual coefficients from an unweighted Poisson pseudo-maximum-likelihood regression of the number of acqui-hires in an acquirer–two-digit-target-industry–year cell on calendar-year indicators interacted with the industry's predetermined employment share in occupations that cannot be performed entirely from home. The sample is a balanced 2017–2022 panel of 481 acquirer–target-industry pairs with at least one acqui-hire during 2017–2019, yielding 2,886 input cells. The regression absorbs acquirer-by-year and target-industry fixed effects and clusters standard errors by acquirer. Fixed-effect singleton and separation removal leaves 479 estimating cells and 69 acquirers. The coefficients measure annual changes, relative to 2019, in the industry-exposure gradient of expected acqui-hire counts within acquirer-years. The open circle denotes the omitted 2019 coefficient, normalized to zero. Vertical bars are 95 percent confidence intervals; shading denotes 2020–2022. Acquisition data are from PitchBook and the documentary measure; nonteleworkability follows Dingel and Neiman (2020).

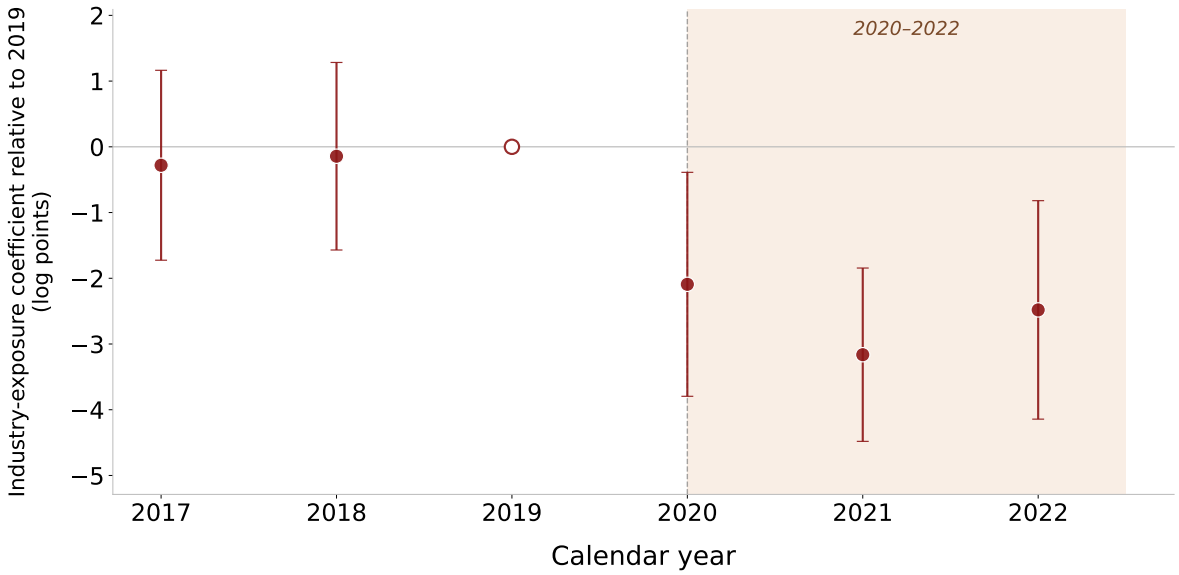


Figure 8. Industry Exposure to Nonteleworkable Work and Acqui-hire Incidence among Measured Transfers

Notes: The figure reports annual coefficients from an unweighted Poisson pseudo-maximum-likelihood regression of the acqui-hire count in a acquirer–two-digit-pre-transfer-industry–year cell on calendar-year indicators interacted with the industry's predetermined employment share in occupations that cannot be performed entirely from home. The logarithm of the number of measured transfers in the cell enters as an offset. Measured transfers comprise acqui-hires and comparison worker flows; the screen requires prior employment linkage among the movers. The regression absorbs acquirer-by-year and pre-transfer-industry fixed effects and clusters standard errors by acquirer. The input contains 3,263 positive-transfer cells; fixed-effect singleton and separation removal leaves 2,158 estimating cells and 235 acquirers. The coefficients measure annual changes, relative to 2019, in the industry-exposure gradient of expected acqui-hire counts per measured transfer. They do not estimate the probability of acquisition for a fixed acquirer–team opportunity. The 2017 and 2018 coefficients are jointly indistinguishable from zero ($p = 0.254$). The open circle denotes the omitted 2019 coefficient, normalized to zero. Vertical bars are 95 percent confidence intervals; shading denotes 2020–2022. Acquisition data are from PitchBook and the documentary measure; worker histories are from Revelio Labs; and nonteleworkability follows Dingel and Neiman (2020).

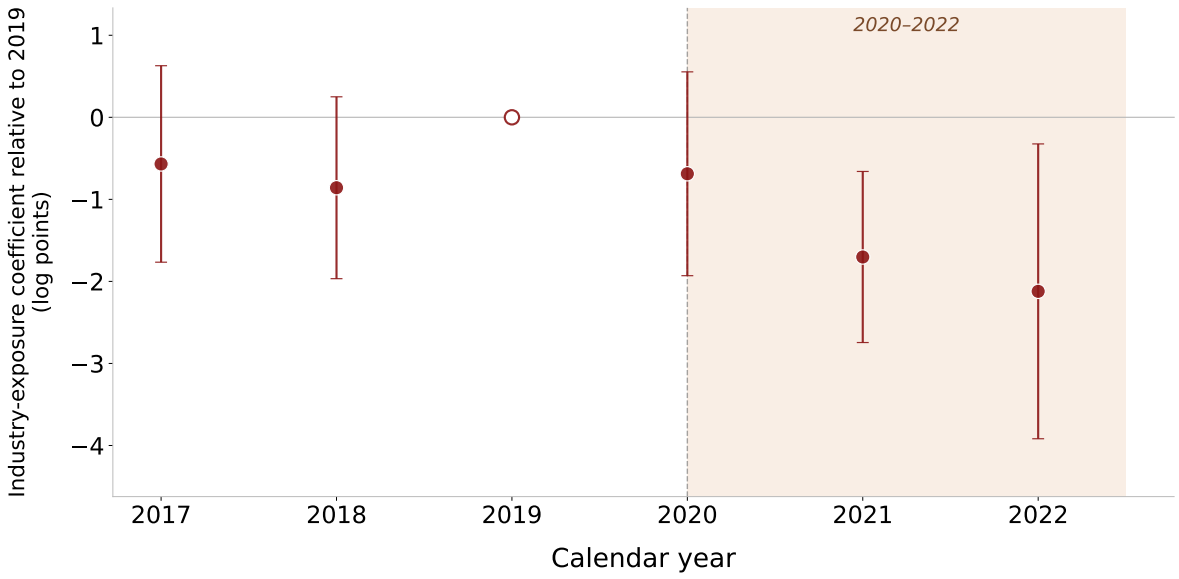


Figure 9. Local Industry-Market Outcomes after an Acqui-hire

Notes: The figure reports synthetic difference-in-differences event-study estimates for the log number of establishments, log employment, and log total annual payroll. A market is the target's harmonized pre-acquisition four-digit industry interacted with the modal destination Core-Based Statistical Area of its transferred high-skill workers. The first qualifying acqui-hire in the market defines year g. The transaction year is excluded; open circles report four pre-treatment coefficients and filled circles report the five complete post-treatment years $g+1$ through $g+5$. Estimation is conducted within three-digit-industry-by-treatment-year blocks. Treated markets are compared with never-treated markets in the same three-digit industry whose pre-treatment growth lies within the treated support; if fewer than 15 donors qualify, the screen expands to a two-standard-deviation band. Block estimates are aggregated using treated-market counts. Subsequent acqui-hires are allowed, so the estimates describe outcomes after the first observed acqui-hire rather than one isolated transaction. Vertical bars are 95 percent confidence intervals from a delete-one-area jackknife over 359 Core-Based Statistical Areas. Joint tests of the four pre-treatment coefficients yield $p = 0.308$, 0.960, and 0.997 for establishments, employment, and payroll. The three panels use 42, 42, and 40 industry–cohort blocks and 147, 143, and 139 treated markets, respectively. Outcomes are from the Quarterly Census of Employment and Wages, 2004–2023. Coefficients are in log points.

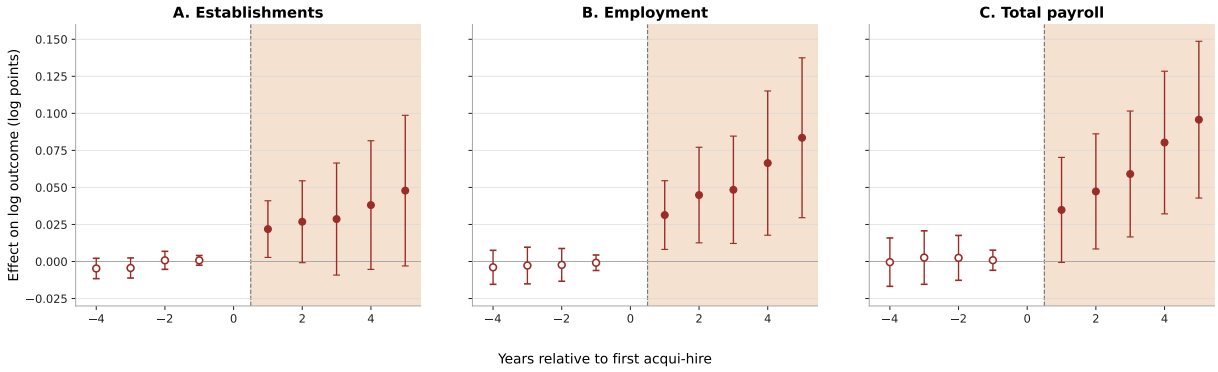


Table 1. Technological Relatedness and Acqui-hire Status

Notes: The unit is a completed acquisition from 2000 through 2025 in which the direct acquirer and target each have at least one linked U.S. utility patent filed during the five years before transaction completion. The common valid-pair population contains 2,826 transactions, including 661 identified acqui-hires; self-pairs are excluded and each transaction appears once. The dependent variable equals one when a transaction source explicitly identifies obtaining employees, founders, or the team as a motive. The regressions compare realized acquisitions and do not model the set of potential acquirers or targets. Patent-class portfolio proximity is the generalized Mahalanobis-weighted measure of Bloom et al. (2013). Claims-text similarity is the cosine similarity between equal-patent-weighted portfolio representations formed from the eventual claims content of patents filed during the five-year pre-transaction window. Historical versions of patent claims are not observed, so this measure need not equal the text publicly available at transaction completion. Common cited prior art is the number of distinct prior patents cited by both firms divided by the smaller of their two prior-art sets, adapting the common-knowledge-base idea in Bena and Li (2014). Column 4 controls for the logarithms of one plus each firm's pre-transaction patent count and distinct prior-art-set size. Column 5 includes all three relatedness measures and the logarithms of one plus citations running in each direction between the firms. All measures use the direct acquirer. Each column reports a linear probability model with target three-digit-industry-by-deal-year fixed effects. Standard errors are clustered by the acquirer's ultimate parent and reported in parentheses. Sample sizes vary with measure availability. Coefficients are in probability units. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Transaction data are from PitchBook, acqui-hire labels use the cited transaction sources, patent records are from PatentsView, and claims, titles, and abstracts are from the linked patent-text archive.

	(1) Patent-class proximity	(2) Claims-text similarity	(3) Common cited prior art	(4) Prior art and scale controls	(5) Full joint
Patent-class portfolio proximity	0.0494				0.0403
	(0.0359)				(0.0465)
Claims-text similarity		0.1732			-0.0390
		(0.2713)			(0.3915)
Common cited prior art			0.1006***	0.0725**	0.0950**
			(0.0376)	(0.0336)	(0.0442)
Patent-portfolio and prior-art-set size controls	No	No	No	Yes	No
Directional citation controls	No	No	No	No	Yes
Transactions	2,583	2,525	2,433	2,433	2,151
Ultimate-parent acquirer clusters	772	750	732	732	666
Target industry × deal-year FE	Yes	Yes	Yes	Yes	Yes

Table 2. Technological Direction of Post-entry Patenting

Notes: The underlying common-support population contains 1,133 identified acqui-hires and 7,095 comparison worker flows from 2010 through March 2020. An identified acqui-hire is an acquisition whose transaction coverage presents obtaining employees, founders, or the team as a motive. A comparison flow is a group leaving a common source firm and entering the same employer within an acquirer's corporate family within 60 days, without acquisition of the source firm, and with prior-workplace linkage above that of equal-size source-firm placebo groups; its contracting arrangement is not observed.
Panel A. The unit is a patent–event observation. The sample contains patents that name at least one incoming worker, are assigned to the acquirer or its parent, are filed within four years after at least one named entrant arrives, and have observed pre-entry patent portfolios for the incoming group and acquirer. The 4,717 observations represent 3,572 distinct patents, 398 events, and 65 acquirers. Columns 1 and 2 measure the share of the focal patent's Cooperative Patent Classification subclasses or main groups represented in source-assigned patents produced by incoming workers during the five years before entry. Column 3 equals one if the focal patent contains at least one main group represented in the incoming group's pre-entry portfolio and at least one represented in the acquirer's pre-entry portfolio. Each column reports the acqui-hire coefficient from a weighted linear regression. Matching weights are normalized across patents so that each event retains its matched weight. Regressions control for the logarithms of one plus inventor count, focal-patent CPC breadth, and the incoming group's pre-entry patent count and breadth; column 3 also controls for the acquirer's pre-entry patent count and breadth. They absorb acquirer, three-digit pre-transfer-industry, event-year, filing-year, and primary-CPC-field fixed effects.
Panel B. The unit is a transfer event with nonempty, text-supported portfolios required for the relevant outcome. The post-entry portfolio contains patents naming an incoming worker, assigned to the acquirer's corporate family, and filed within the indicated horizon after that worker enters. The incoming-group pre-entry portfolio contains source-assigned patents filed during the preceding five years; the acquirer portfolio contains patents assigned to its corporate family during the five years before the cohort's earliest entry. Each portfolio equally weights the available claims-text patent representations. Continuity is the cosine similarity between the post-entry and incoming-group pre-entry portfolios. Movement toward the acquirer is the post-entry portfolio's similarity to the acquirer's pre-entry portfolio minus the incoming group's pre-entry similarity to that portfolio. Each cell reports the acqui-hire coefficient relative to comparison flows from a weighted linear regression with acquirer, two-digit pre-transfer-industry, and entry-year fixed effects. Regressions use the common-support matching weights. In both panels, standard errors are clustered by acquirer and reported in parentheses. Panel A coefficients multiplied by 100 are percentage-point differences; Panel B coefficients are changes in cosine similarity and are not percentage-point effects. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Transaction data are from PitchBook and the cited transaction documents; worker histories are from Revelio Labs; patent records are from PatentsView; and claims are from the linked patent-text archive.

Panel A. Technological content of post-entry patents	(1) Incoming-group fields CPC subclass	(2) Incoming-group fields CPC main group	(3) Fields from incoming group and acquirer
Acqui-hire	0.233***	0.276***	0.364***
	(0.045)	(0.035)	(0.053)
Patent–event observations	4,717	4,717	4,717
Distinct patents	3,572	3,572	3,572
Events	398	398	398
Acquirer clusters	65	65	65
Panel B. Claims-text portfolio measures	24 months	36 months	48 months
Continuity with incoming group's pre-entry technology	0.0317***	0.0311***	0.0272***
	(0.0043)	(0.0036)	(0.0047)
Movement toward acquirer's pre-entry technology	0.0199***	0.0177***	0.0159***
	(0.0054)	(0.0038)	(0.0035)
Transfer events	266	314	327
Acquirer clusters	27	28	29

Table 3. Industry Exposure to Nonteleworkable Work and Acqui-hire Activity

Notes: Column 1 uses a balanced 2017–2022 panel of acquirer–target-industry pairs with at least one acqui-hire during 2017–2019. The dependent variable is the annual number of acqui-hires. Column 2 aggregates acqui-hires and comparison worker flows to acquirer–pre-transfer-industry–year cells with at least one measured transfer. Its dependent variable is the acqui-hire count, and the logarithm of the total measured-transfer count enters as an offset. The workplace screen requires prior employment linkage among the moving workers; it does not verify their contracting arrangement. Industry exposure is the predetermined employment share in occupations that cannot be performed from home. Both columns are estimated by Poisson pseudo-maximum likelihood, absorb acquirer-by-year and two-digit-industry fixed effects, are unweighted, and cluster standard errors by acquirer. The balanced input panel contains 2,886 cells; fixed-effect singleton and separation removal leaves 479 cells that identify the PPML coefficient. In column 2, the corresponding counts are 3,263 input cells and 2,158 retained estimation cells. The implied changes equal $100[\exp(\widehat{\beta}\sigma_C) - 1]$. The column 1 pre-period test uses quarterly interactions; the column 2 test uses annual interactions. The conditional-incidence estimate in column 2 conditions on a realized measured transfer and is not the probability that a fixed acquirer–team opportunity is organized through acquisition. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Acquisition data are from PitchBook and the documentary measure; worker histories are from Revelio Labs; and nonteleworkability follows Dingel and Neiman (2020).

	Acqui-hire count (1)	Acqui-hire count per measured transfer (2)
2020–2022 × industry share in nonteleworkable occupations	-2.487^{***}	-1.078^{***}
	(0.588)	(0.396)
Standard deviation of industry exposure	0.245	0.276
Implied one-standard-deviation change	−45.7%	−25.8%
Estimator	PPML	PPML
Log measured-transfer-count offset	No	Yes
Acquirer × year fixed effects	Yes	Yes
Two-digit industry fixed effects	Yes	Yes
Input cells	2,886	3,263
Estimation cells	479	2,158
Acquirer clusters	69	235
Quarterly pre-period test, p-value	0.924	—
Annual pre-period test, p-value	—	0.254

Table 4. COVID-19 Disruption Exposure and Post-transfer Team Preservation

Notes: The unit is a worker group transferred to a common acquirer; outcomes are measured 24 months after entry. Pairwise co-retention is the share of entrant pairs whose two members both remain. Disruption exposure is the share of the outcome window overlapping April 2020–April 2022. Task-based exposure averages across high-skill entrants a standardized occupation index equal to Dingel–Neiman nonteleworkability multiplied by the mean O*NET importance of nine interaction-oriented activities; the event mean is restandardized as described below. It predicts exposure to the disruption, not team complementarity itself.
Panel A. The sample contains pre-April-2020 acqui-hires with at least three entrants. Post-transfer task exposure, C_e^D, uses occupations associated with entrants' spells at the acquirer and is standardized over the 835-event matched pool before singleton removal. It measures the realized post-transfer task environment and can reflect assignments made after acquisition. Appendix Table A.14 shows that the estimates remain negative and statistically precise when only acquirer spells beginning before April 2020 contribute to the index. Weighted linear regressions include disruption exposure, task-based exposure, the logarithm of one plus arriving-group size, the logarithm of one plus target employment, target age, high-skill share, and the professional-managerial share (the fraction of entrants who are at least midlevel and work in managerial, computer, engineering, science, or health occupations). Task-based exposure is also interacted with these two logarithmic controls and target age. Regressions absorb acquirer and two-digit target-industry-by-transfer-year fixed effects. Hájek weights equalize represented exposure quartiles within matched strata defined by target industry, transfer year, and above-median indicators for arriving-group size, target employment, high-skill share, and professional-managerial share.
Panel B. The common-support sample contains 1,133 acqui-hires and 7,095 comparison worker flows; singleton removal leaves 7,926 observations. Support is restricted within transfer year, pre-transfer industry, arriving-group-size, occupation-composition, and prior-cohesion cells. Composition uses incoming workers for acqui-hires but the reconstructed pre-transfer group for comparison flows. Pre-transfer task exposure, C_e^S, uses pre-transfer occupations and is standardized over the eligible event population before trimming and support restriction. Unweighted linear regressions include the triple interaction, every lower-order term, acquirer fixed effects, and two-digit pre-transfer-industry-by-transfer-year fixed effects. Standard errors are clustered by the acquirer and reported in parentheses. Linear coefficients are in proportion units; multiplying by 100 gives percentage points. Within the observed support, simultaneous one-standard-deviation increases in disruption exposure and task-based exposure imply changes of approximately −7.6 percentage points in Panel A and −2.8 percentage points in Panel B for pairwise co-retention. Panel B coefficients are differential gradients, not average effects of acquisition. The comparison-flow and acqui-hire gradients are reported for pairwise co-retention in column 1. Threshold-retention outcomes and longer horizons are reported in the appendix. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Transaction classifications are from PitchBook and cited documents; worker histories are from Revelio Labs; and task measures combine Dingel–Neiman teleworkability with O*NET activities.

	Pairwise co-retention (1)	Fraction of workers retained (2)
Panel A. Within acqui-hires		
Disruption exposure × task-based exposure	−0.267***	−0.181***
	(0.065)	(0.043)
Observations	502	502
Acquirer clusters	133	133
Panel B. Acqui-hires versus comparison worker flows		
Acqui-hire × disruption exposure × task-based exposure	−0.136**	−0.095**
	(0.065)	(0.047)
Comparison-flow gradient	−0.009	
	(0.024)	
Acqui-hire gradient	−0.146**	
	(0.064)	
Observations	7,926	7,926
Acquirer clusters	354	354
Outcome horizon	24 months	
Acquirer fixed effects	Yes	
Two-digit pre-transfer industry × transfer-year fixed effects	Yes	

Table 5. COVID-19 Disruption Exposure and Post-transfer Career Progression

Notes: The common-support sample contains 1,133 acqui-hires and 7,095 comparison worker flows with at least three incoming workers. A comparison flow is a group from a common source firm that enters the same employer within the acquirer's corporate family within 60 days, without acquisition of the source firm, and has prior-workplace linkage above that of equal-size placebo groups from the source firm. Common support is imposed on transfer year, pre-transfer industry, incoming-group size, occupation composition, and prior-joint-mobility cohesion. The outcomes are cumulative positive within-firm seniority change and the number of promotions for the incoming group during the indicated post-transfer horizon. The seniority measure sums positive rank-change magnitudes and excludes the rank change recorded when a worker enters the acquirer. Each coefficient comes from a separate Poisson pseudo-maximum-likelihood regression with log incoming-group size as an offset, so the coefficient describes the expected outcome per worker. The reported coefficient is on acqui-hire × disruption exposure × standardized task-based exposure. Disruption exposure is the fraction of the outcome window overlapping April 2020–April 2022. Task-based exposure is the mean among high-skill movers of an occupation-level index that multiplies the complement of Dingel–Neiman teleworkability by the average O*NET importance of nine communication, interpersonal, coordination, and supervisory activities. It is constructed from pre-transfer occupations and standardized over the eligible event population before trimming and common-support restriction. The occupation-composition variables used for common support are measured over the incoming group for acqui-hires and the reconstructed pre-transfer group for comparison flows. Every lower-order term is included; regressions also absorb acquirer and two-digit pre-transfer-industry-by-transfer-year fixed effects and cluster standard errors by acquirer. At 24 months, fixed-effect singleton and separation removal leaves 6,914 of the 8,228 common-support groups and 229 acquirers; the effective sample varies with the outcome horizon as reported in the table. The exponentiated coefficient is $\exp(\widehat{\beta})$. It compares the exposure-related change in the ratio of expected outcomes following an acqui-hire relative to a comparison flow for a one-unit increase on this scale standardized before the common-support restriction; it is not a difference in career-outcome levels. Simultaneous one-standard-deviation increases in disruption exposure and task-based exposure imply 24-month multipliers of 0.86 for cumulative seniority gain and 0.83 for promotions. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Transaction classifications are from PitchBook and cited documents; worker histories are from Revelio Labs; and task measures combine Dingel–Neiman teleworkability with O*NET activities.

	Outcome horizon		
	24 months	36 months	48 months
	(1)	(2)	(3)
Cumulative positive within-firm seniority change	-0.765^{***}	-0.673^{**}	-0.576^{**}
	(0.283)	(0.268)	(0.277)
Exponentiated coefficient	0.47	0.51	0.56
Promotions	-0.904^{***}	-0.799^{***}	-0.637^{**}
	(0.267)	(0.288)	(0.288)
Exponentiated coefficient	0.40	0.45	0.53
Observations	6,914	7,017	7,049
Acquirer clusters	229	254	260
Log incoming-group-size offset	Yes	Yes	Yes
Acquirer fixed effects	Yes	Yes	Yes
Two-digit pre-transfer industry × transfer-year fixed effects	Yes	Yes	Yes

Table 6. Coordination Disruption and Innovation by Incoming Inventors

Notes: This table reports differential disruption-by-task-exposure gradients in post-transfer innovation for incoming inventor groups following identified acqui-hires and comparison worker flows. An identified acqui-hire is an acquisition whose documentary text presents obtaining the target's employees, founders, or team as a transaction motive. Comparison flows contain at least three workers moving from a common source firm to the same employer within the acquirer's corporate family within a 60-day window, with prior-workplace linkage above that of equal-size placebo groups from the source firm; their contracting arrangement is not observed. The data combine PitchBook transactions, Revelio Labs employment histories, and U.S. utility-patent records through September 2025. Transfers occur from 2010 through March 2020. All patent outcomes retain patents assigned to the acquirer or its corporate family. Disruption exposure is the share of the stated outcome window that overlaps April 1, 2020–April 1, 2022. Post-transfer task exposure is constructed from the incoming inventors' occupations at the acquirer and is high when those jobs combine low teleworkability with interaction-intensive tasks.
Panel A's unit is an incoming inventor group. A group-internal patent names at least two inventors from that group and no pre-existing acquirer inventor. The outcomes count incoming-inventor–patent incidences or weight those incidences by adjusted forward citations or allocated private patent value. Each cell reports the coefficient on acqui-hire status interacted with disruption exposure and post-transfer task exposure from a separate PPML regression. All constituent lower-order terms are included. Member-years enter as the exposure; regressions also control for the logarithm of one plus incoming-inventor count and absorb pre-transfer-industry-by-transfer-year fixed effects, and cluster standard errors by acquirer. Before outcome-specific attrition and PPML separation, the sample contains 353 acqui-hire transaction groups and 1,365 comparison groups. Coefficients are log conditional-mean gradients, not percentage-point changes or average effects of acquisition.
Panel B contrasts mixed and solo output within an analysis unit. Mixed patents name at least two incoming inventors and at least one pre-existing acquirer inventor; solo output names one incoming inventor and no pre-existing acquirer inventor. An acqui-hire analysis unit is an acquirer-by-entry-year cohort; a comparison-flow unit is its incoming inventor group. The reported coefficient is the difference between acqui-hires and comparison flows in the mixed-versus-solo disruption-by-task-exposure gradient. Regressions include every constituent lower-order term, analysis-unit fixed effects, and weights that equalize the two event categories within support cells defined by pre-transfer industry, incoming-inventor count, and post-transfer task exposure; standard errors are clustered by analysis unit. Patent volume is top-coded at the 95th percentile of positive observations. Citation- and private-value-weighted outcomes are inverse-hyperbolic-sine transformed, so those coefficients are neither patent counts nor percentage effects. Four pre-COVID placebo windows yield no corresponding pattern for the estimates used in the text. The imprecise 48-month adjusted-citation estimate in Panel A is the exception and is not used to support the claim. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Outcome horizon		
	24 months	36 months	48 months
	(1)	(2)	(3)
Panel A. Production within the incoming inventor group			
Group-internal patenting rate	−3.749***	−2.982***	−3.166**
	(0.724)	(1.098)	(1.434)
Observations	1,512	1,559	1,568
Acquirer clusters	169	175	176
Adjusted-citation-weighted rate	−6.643***	−6.924***	−5.542
	(1.283)	(1.767)	(3.839)
Observations	1,360	1,378	1,489
Acquirer clusters	150	157	167
Private-value-weighted rate	−4.837***	−5.154***	−3.931**
	(0.947)	(1.454)	(1.960)
Observations	1,479	1,486	1,532
Acquirer clusters	166	166	171
Panel B. Production combining incoming and pre-existing acquirer inventors			
Mixed-versus-solo patent volume	−0.735	−2.238*	−3.926**
	(0.792)	(1.172)	(1.772)
Adjusted-citation-weighted output	−2.084***	−2.740***	−3.362***
	(0.457)	(0.731)	(0.965)
Private-value-weighted output	−3.306***	−3.591***	−4.697***
	(0.922)	(1.344)	(1.717)
Observations	2,802	2,802	2,802
Acquirer-by-transfer-category clusters	1,401	1,401	1,401

Table 7. Coordination Disruption and Innovation by Pre-existing Acquirer Inventors

Notes: The unit is a pre-existing acquirer-inventor–technology-field–filing-year observation. An identified acqui-hire is an acquisition whose documentary text presents obtaining the target's employees, founders, or team as a transaction motive. A comparison flow contains at least three workers who move from a common source firm to the same employer within the acquirer's corporate family within 60 days and have prior-workplace linkage above that of equal-size placebo groups from the source firm; the contracting arrangement is unobserved. Transfers occur from 2010 through March 2020. A pre-existing acquirer inventor worked at the acquirer before the transfer and is subsequently observed co-inventing with an incoming inventor. Outcomes exclude patents naming any incoming inventor and retain patents assigned to the acquirer or its corporate family.

Each cell reports the acqui-hire-versus-flow difference in the coefficient on incoming-expertise field × post-transfer period × disruption exposure × acquirer-role task exposure. Disruption exposure is the share of the stated horizon overlapping April 1, 2020–April 1, 2022. Acquirer-role exposure is constructed from the incoming inventors' occupations at the acquirer and is high when those jobs combine low teleworkability with interaction-intensive tasks. All constituent lower-order terms are included. Aggregate and new-field specifications absorb transfer-category-specific technology-class-by-year and inventor-by-treated-field fixed effects. Value specifications additionally absorb inventor-by-year fixed effects. Standard errors are clustered by acquirer-by-transfer-category cells. The coefficient compares each inventor's production in incoming-expertise fields with production in other active fields and then compares that gradient across acqui-hires and comparison flows; it is not the causal effect of a particular collaboration.

For aggregate and new-field counts, incoming-expertise fields are the acquirer-level union of fields represented in pre-transfer patents of groups arriving before April 2020. Aggregate production is cumulative patenting in those fields; new-field production restricts the count to fields new to the acquirer inventor. Both counts are top-coded at the 95th percentile of positive observations. For private-value and citation-weighted deepening, high-impact incoming-expertise fields become treated when the relevant group arrives. The outcomes are the inverse-hyperbolic-sine of annual allocated private patent value and adjusted citations in fields where the acquirer inventor had previously worked. None of the four pre-COVID placebo windows yields a corresponding pattern across the twelve reported specifications. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Transaction classifications are from PitchBook and cited documents; worker histories are from Revelio Labs; patent records are from PatentsView; and task measures combine Dingel–Neiman teleworkability with O*NET activities.

	Outcome horizon		
	24 months	36 months	48 months
	(1)	(2)	(3)
Aggregate patent production	−0.993**	−1.662***	−2.669***
	(0.461)	(0.615)	(0.879)
Observations	202,027	227,303	252,578
Clusters	138	138	138
Private-value deepening	−4.238***	−2.460***	−2.709***
	(0.769)	(0.549)	(0.826)
Observations	189,883	213,641	237,396
Clusters	110	110	110
New-field patent production	−0.149	−0.748*	−1.439**
	(0.352)	(0.452)	(0.647)
Observations	202,027	227,303	252,578
Clusters	138	138	138
Citation-weighted deepening	−1.205***	−0.675***	−0.693***
	(0.432)	(0.252)	(0.231)
Observations	189,883	213,641	237,396
Clusters	110	110	110

Table 8. Coordination Disruption and Worker Reallocation

Notes: This table reports differential changes in worker reallocation following acqui-hires relative to comparison worker flows for transfers from 2010 through March 2020. An identified acqui-hire is an acquisition whose documentary text explicitly presents obtaining the target's employees, founders, or team as a transaction motive; this proxy does not establish that it was the primary motive. A comparison flow contains at least three workers who move from a common source firm to the same employer within the acquirer's corporate family within 60 days and have prior-workplace linkage above that of equal-size placebo groups from the source firm. Before outcome-specific estimation, the common-support population contains 1,133 acqui-hires and 7,095 comparison flows.
Each cell reports $\widehat{\beta}_h$ from

$$Y_{ieh} = \beta_h A_e S_{eh} C_e^S + \text{all constituent lower-order terms} + \text{fixed effects} + \varepsilon_{ieh}.$$

The index i denotes a worker in Panel A and is suppressed in the event-level regressions in Panel B. Here, A_e indicates an identified acqui-hire, S_{eh} is the share of the h-month interval following transfer event e that overlaps April 1, 2020–April 1, 2022, and C_e^S is standardized task-based exposure to in-person coordination, constructed from the incoming group's pre-transfer occupations. The coefficient is a differential disruption-by-task-exposure gradient, not an average effect of acquisition.
Panel A uses workers as observations. "Exit" equals one when the worker is no longer employed within the acquirer's corporate family at the stated horizon. "Co-departure" equals one when the worker and another original group member have both left by that horizon and their recorded exit dates are within 180 days; it does not require a common destination. Linear probability models control for pre-transfer seniority, log modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status, and absorb acquirer, two-digit pre-transfer-industry, and transfer-year fixed effects. Each worker's weight is the event's common-support weight divided by the number of workers in the event. Standard errors are two-way clustered by acquirer and transfer event.
Panel B uses transfer events as observations. The outcome equals one when at least two original members leave the acquirer's corporate family by the stated horizon and their first observed outside employment spells satisfy the paper's entrepreneurial-destination or entrepreneurial-role criteria. The broad measure uses founder or cofounder titles and founding language; CEO, CTO, or vice-president-or-higher roles at firms with at most 50 employees; employment at firms no more than five years old with at most 500 employees; Y Combinator founder matches; and PitchBook startup destinations with at most 500 employees. The outside spells can begin after the stated horizon and need not be at the same firm; Appendix Table A.19 reports a stricter timing definition. The unweighted linear probability models absorb acquirer and two-digit pre-transfer-industry-by-transfer-year fixed effects and cluster standard errors by acquirer. Fixed-effect singleton removal leaves 7,926 events—856 acqui-hires and 7,070 comparison flows—and 354 acquirers in columns 1 and 2. Column 3 additionally requires complete 48-month outcome support for every measured group member and contains 7,922 events and 354 acquirers. Coefficients are in proportion units. At 24 months, the estimator-relevant standard deviations of disruption exposure and C_e^S are 0.268 and 0.751 in Panel A and 0.280 and 0.713 in Panel B. For simultaneous one-standard-deviation increases in the two variables, the triple-interaction component therefore corresponds to 2.0 percentage points for exit, 2.6 percentage points for co-departure, and 4.8 percentage points for entrepreneurial reallocation. For comparison, under full disruption exposure and a one-standard-deviation difference in C_e^S, the corresponding components are 7.4, 9.8, and 17.2 percentage points, respectively. Maximum disruption exposure is 0.999, 0.666, and 0.500 in columns 1–3, so raw coefficients are not directly comparable across horizons. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Transaction classifications are from PitchBook and cited documents; worker histories are from Revelio Labs; and task measures combine Dingel–Neiman teleworkability with O*NET activities.

	Outcome horizon		
	24 months (1)	36 months (2)	48 months (3)
Panel A. Worker exit and co-departure			
Exit from the acquirer's corporate family	0.098**	0.054	−0.040
	(0.048)	(0.063)	(0.070)
Co-departure within 180 days	0.130***	0.078	0.038
	(0.040)	(0.057)	(0.074)
Observations	169,559	169,559	169,541
Acquirer clusters	650	650	650
Transfer-event clusters	8,228	8,228	8,228
Panel B. Entrepreneurial reallocation at the transfer-event level			
At least two members exit by horizon; later entrepreneurial destination or role	0.241**	−0.122	−0.213
	(0.113)	(0.127)	(0.140)
Transfer events	7,926	7,926	7,922
Acquirer clusters	354	354	354

Table 9. Technological Relatedness of Post-exit Patents to the Incoming Group

Notes: The unit is a transfer event. The table compares identified acqui-hires with comparison worker flows from 2010 through March 2020. Comparison flows are groups leaving a common source firm and entering a common acquirer without acquisition of that source firm; their contracting arrangement is unobserved. A post-exit patent is a U.S. utility patent that names an original group member, is filed after that worker's sustained exit from the acquirer's corporate family, and is assigned wholly outside that family at filing. Patent records are observed through September 2025. Excess similarity is the claims-text cosine similarity between a post-exit patent and the indicated focal-group portfolio minus its average similarity to up to five placebo-group portfolios. The first reference contains patents linked to original group members and assigned to their pre-transfer employer before transfer. The second contains acquirer-family patents filed before the focal worker's exit by other original group members, excluding the focal worker. Placebo groups come from the same two-digit pre-transfer industry and within two transfer years, are selected by proximity in the size of the original group's pre-transfer employer patent portfolio, and share no inventor, patent, or acquirer with the focal group. Focal and placebo portfolios use the same focal-exit cutoff. Patent-level excess similarity is averaged within each transfer event, and events receive equal weight. Columns 1 and 2 report event-level means. Column 3 reports the coefficient on the acqui-hire indicator from an otherwise uncontrolled event-level regression. Standard errors, in parentheses, are clustered by acquirer. Counts in column 3 are the pooled support for the corresponding contrast. The two reference portfolios have different availability and therefore different samples. In columns 1 and 2, stars test zero excess similarity; in column 3, they test equality across event categories. The estimates describe selected patenting exiters and do not identify knowledge transfer or the causal effect of acquisition. Data are from PitchBook, Revelio Labs, PatentsView, the patent-text archive, and the corporate-family mapping. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Reference portfolio	(1) Acqui-hire	(2) Comparison flow	(3) Acqui-hire − comparison
Original group members' pre-transfer patents			
assigned to target or source firm	0.0213***	0.0225***	-0.0012
	(0.0042)	(0.0019)	(0.0040)
Other group members' acquirer-family patents before focal exit	0.0285***	0.0042	0.0243***
	(0.0040)	(0.0035)	(0.0056)
Transfer events, target/source-firm reference	39	238	277
Acquirer clusters, target/source-firm reference	34	39	59
Transfer events, other-members reference	54	126	180
Acquirer clusters, other-members reference	38	27	50
Equal event weights	Yes	Yes	Yes
Additional controls	No	No	No

Table 10. Technological Direction and External Follow-on

Notes: The unit is a patent produced after an identified acqui-hire, naming at least one original incoming-group inventor and assigned to the acquirer or its corporate family. Transfers occur from 2010 through March 2020, and post-transfer patents are observed through September 2025; the citation archive ends July 24, 2025. The regressor is the claims-text cosine similarity between the focal patent and the acquirer's equal-patent-weighted portfolio of U.S. utility patents filed before the transaction, standardized within the acqui-hire sample. Columns 1–3 report Poisson pseudo-maximum-likelihood coefficients. Column 1 counts wholly nonaffiliate citing patents over the full observed citation history; column 2 counts such citations during the first 36 months after filing or through the citation-archive endpoint, whichever comes first; and column 3 counts distinct wholly nonaffiliate citing corporate families over the full history. Column 4 is a linear probability model in which the outcome equals one if the focal patent receives any wholly nonaffiliate citation. Nonaffiliate means that no assignee on the citing patent shares an ultimate corporate parent with the acquirer. Citations sharing an inventor with the focal patent and citations whose assignee cannot be resolved are excluded. The raw direction sample contains 2,607 patents from 289 acquisitions; fixed-effect singleton and separation removal yields the column-specific samples reported below. Each regression absorbs acquisition and technology-field-by-filing-year fixed effects and clusters standard errors by acquirer. PPML coefficients are log conditional-mean differences; the column 4 coefficient is in probability units. The specifications compare patents produced following the same acquisition and do not identify a causal effect of technological direction. Data are from PitchBook, Revelio Labs, PatentsView, the patent-text archive, and the corporate-family mapping. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	(1)	(2)	(3)	(4)
	Nonaffiliate citing patents		Distinct outside	Any outside
	Full observed	Up to 36 months	citing families	citation
Proximity to acquirer's pre-transaction technology	-0.237*	-0.262**	-0.102	-0.015
	(0.136)	(0.110)	(0.068)	(0.015)
Acquisition fixed effects	Yes	Yes	Yes	Yes
Technology-field × filing-year fixed effects	Yes	Yes	Yes	Yes
Estimator	PPML	PPML	PPML	LPM
Patents	2,304	2,153	2,304	2,525
Acquisitions	166	139	166	216
Acquirer clusters	86	73	86	113

Table 11. Innovation Output and External Follow-on along the Observed Transfer Path

Notes: This table reports descriptive accounting for the common population of 8,228 transfer events from 2010 through March 2020: 1,133 identified acqui-hires and 7,095 comparison worker flows. Comparison flows are groups leaving a common source firm and entering a common acquirer without acquisition of that source firm; their contracting arrangement is not observed. Q_{48}^{in} is the number of eligible inside-produced patents per original incoming worker in the corresponding transfer. An eligible patent names at least one original group inventor, is assigned to the acquirer's corporate family, is filed within 48 months after transfer, and has a complete 36-month forward-citation opportunity. The citation archive ends July 24, 2025, so eligible patents are filed by July 24, 2022. δ_{48}^{in} is the number of wholly nonaffiliate citations during that 36-month window per eligible patent. D_{48}^{in} is the same citations per original incoming worker. The same patent and citation universe is used throughout, so $D_{48}^{\text{in}} = Q_{48}^{\text{in}}\delta_{48}^{\text{in}}$ holds exactly. Events with no eligible patent remain in the worker denominator. The table is unweighted. Confidence intervals for the ratios are centered percentile intervals from 999 bootstrap draws that resample the 650 acquirer clusters and recompute all three objects. Stars test the null that the ratio equals one using the bootstrap standard error of the log ratio. Data are from PitchBook, Revelio Labs, PatentsView, and the corporate-family mapping. These are observed-path comparisons between acqui-hires and measured comparison flows, not causal effects of acquisition or estimates of the no-acquisition counterfactual. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	(1) Inside patents per incoming worker Q_{48}^{in}	(2) External citations per inside patent δ_{48}^{in}	(3) External citations per incoming worker D_{48}^{in}
Acqui-hire	0.097	0.697	0.068
Comparison flow	0.042	0.526	0.022
Ratio: acqui-hire / comparison flow	2.31***	1.33	3.07***
95% confidence interval for ratio	[1.55, 3.85]	[0.80, 2.23]	[1.68, 5.96]
Transfer events		8,228	
Incoming-worker observations		169,663	
Eligible inside patents		8,370	
Wholly nonaffiliate citations		4,776	
Acquirer clusters		650	

Table 12. Acqui-hires and Local Industry-Market Outcomes

Notes: Panel A reports synthetic difference-in-differences estimates for the log number of establishments, log employment, and log total annual payroll in the Quarterly Census of Employment and Wages, 2004–2023. A market is the target's harmonized pre-acquisition four-digit industry interacted with the modal destination Core-Based Statistical Area of its transferred high-skill workers. The first qualifying acqui-hire in a market defines year g. The transaction year is excluded; columns 1–5 are the five complete post-treatment years and column 6 is the average post-treatment effect. Estimation is conducted within three-digit-industry-by-treatment-year blocks. Treated markets are compared with never-treated markets in the same three-digit industry whose pre-treatment growth lies within the treated support; if fewer than 15 donors qualify, the screen expands to a two-standard-deviation band. Block estimates are aggregated using treated-market counts. Subsequent acqui-hires are allowed. Standard errors in parentheses use a delete-one-area jackknife over 359 Core-Based Statistical Areas. Joint tests of the four pre-treatment coefficients yield p-values of 0.308, 0.960, and 0.997 for establishments, employment, and payroll. The three outcomes use 42, 42, and 40 blocks and 147, 143, and 139 treated markets, respectively.
Panel B reports average post-treatment effects from the Business Dynamics Statistics, 2004–2022. Markets are industry-sector-by-metropolitan-area cells, and estimation otherwise follows the same treatment timing, donor rule, and treated-market aggregation. Standard errors use a delete-one-area jackknife over 344 metropolitan areas. The three rows use, respectively, 33, 32, and 30 sector-by-treatment-year blocks. Firms aged zero are new employer firms; the stock of firms aged zero to five and their establishments also reflects survival and aging. Suppressed BDS cells are treated as missing, and each row requires a balanced outcome history. The BDS industry cells are coarser than the QCEW cells, so coefficients should not be differenced across panels. All coefficients are in log points. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Panel A. Local industry-market outcomes

	(1) g+1	(2) g+2	(3) g+3	(4) g+4	(5) g+5	(6) Average post
Establishments	0.0219**	0.0268*	0.0286	0.0381*	0.0478*	0.0327*
	(0.0097)	(0.0141)	(0.0193)	(0.0222)	(0.0259)	(0.0176)
Employment	0.0314***	0.0448***	0.0484***	0.0664***	0.0835***	0.0549***
	(0.0118)	(0.0165)	(0.0185)	(0.0248)	(0.0275)	(0.0181)
Total payroll	0.0348*	0.0473**	0.0591***	0.0803***	0.0957***	0.0634***
	(0.0181)	(0.0198)	(0.0217)	(0.0246)	(0.0270)	(0.0195)

Panel B. Firm births and young-firm activity

	Average post	Treated markets
Firms aged 0–5	0.0566***	97
	(0.0154)	
Establishments of firms aged 0–5	0.0539***	95
	(0.0178)	
Firms aged 0 (births)	0.0504**	92
	(0.0251)	

A Figures and Tables

A.1 Figures

Figure A.1. Cohort-Specific Post-Transfer Task-Exposure Gradients in Pairwise Co-retention

Notes: The outcome is pairwise co-retention 24 months after entry: the share of distinct pairs of workers joining through the acqui-hire whose two members both remain employed by the acquirer. Each marker is the coefficient on standardized post-transfer task exposure interacted with the indicated transfer-cohort indicator. Thus, each marker is the difference between the indicated cohort's post-transfer task-exposure slope and the slope for the omitted 2017 cohort. Coefficients are in proportion units; multiplying by 100 gives percentage-point differences. The first marker pools transactions completed from 2010 through 2014; the 2020 cohort contains transactions completed through March 31. The shaded region identifies cohorts whose 24-month outcome windows can overlap the April 2020–April 2022 disruption.

Post-transfer task exposure multiplies the complement of Dingel–Neiman teleworkability by the average O*NET importance of nine communication, interpersonal, coordination, and supervisory activities. It uses the occupations associated with workers' spells at the acquirer, is averaged among high-skill workers, and is standardized over the 835-event matched pool before fixed-effect singleton removal. The weighted linear regression controls for the logarithm of one plus arriving-group size, the logarithm of one plus target employment, target age, high-skill share, and professional-managerial share. It also interacts post-transfer task exposure with the two logarithmic controls and target age, absorbs acquirer and two-digit target-industry-by-transfer-year fixed effects, and uses Hájek weights that equalize represented exposure quartiles within strata defined by target industry, transfer year, and above-median indicators for arriving-group size, target size, high-skill share, and professional-managerial share. The estimating sample contains 502 transactions and 133 acquirers. Vertical bars are 95 percent confidence intervals based on standard errors clustered by acquirer. The differences in the 2014-or-earlier, 2015, and 2016 post-transfer task-exposure slopes relative to 2017 are jointly zero ($p = 0.963$).

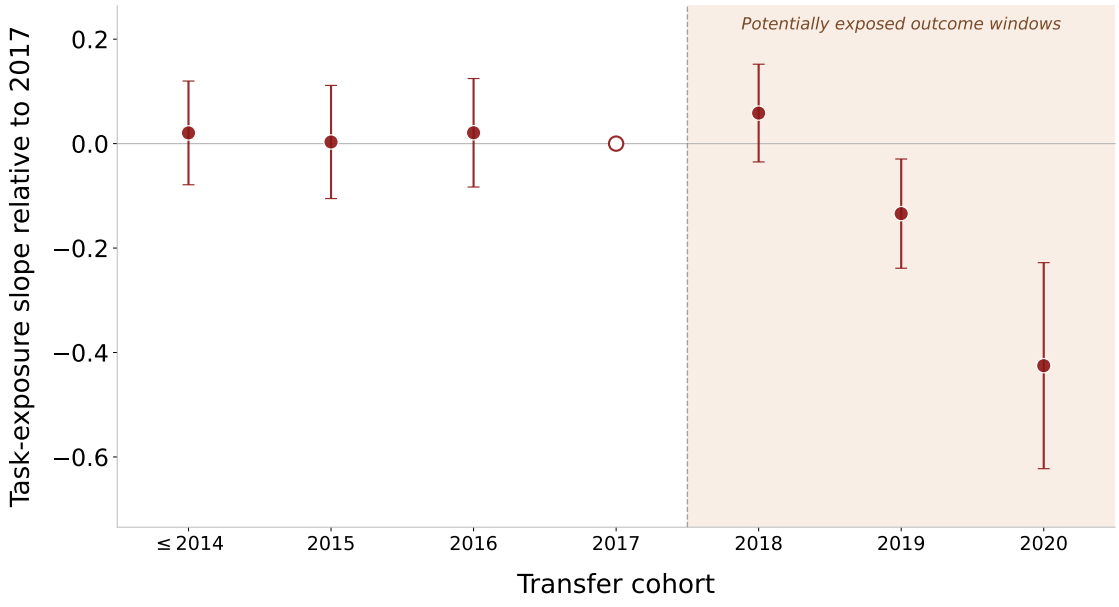


Figure A.2. Actual and Placebo Disruptions for Acqui-hires and Comparison Worker Flows

Notes: This figure compares acqui-hires with comparison flows: groups of at least three workers from the same source firm who enter the same employer within the acquirer's corporate family within 60 days and have prior-workplace linkage above that of equal-size placebo groups from the source firm. Appendix B.6 gives the construction. The screen does not identify the workers' contracting arrangement or establish that the flow was direct hiring. Each marker is the coefficient on acqui-hire status interacted with disruption exposure and standardized pre-transfer task exposure. Filled markers use the April 2020–April 2022 disruption; open markers replace it with the indicated two-year pre-pandemic window and exclude observations whose outcome window overlaps the actual disruption. The coefficients are differential disruption-by-pre-transfer-task-exposure gradients, not average effects of acquisition.

The first four panels report 24-month retention outcomes estimated by unweighted linear regressions. Their coefficients are in proportion units; multiplying by 100 gives percentage points. The last two panels report 36-month cumulative positive within-acquirer seniority change and promotion counts for the arriving group. They are estimated by unweighted PPML with log arriving-group size as an offset. Exponentiating a coefficient gives the multiplicative change, for a one-unit increase in pre-transfer task exposure, in the acqui-hire-to-comparison expected-outcome rate ratio when disruption exposure rises from zero to one. The actual-disruption retention regressions contain 7,926 transfer events and 354 acquirers; their placebo counterparts contain 5,015 events. The actual-disruption career regressions contain 7,017 transfer events and 254 acquirers; their placebo counterparts contain 3,159 events.

Before outcome-specific attrition, the common-support sample contains 1,133 acqui-hires and 7,095 comparison flows. Support is restricted within cells defined by transfer year, the two-digit industry of the target for an acqui-hire or the source firm for a comparison flow, arriving-group size, occupation composition, and prior-joint-mobility cohesion. Pre-transfer task exposure is constructed from movers' pre-transfer occupations and standardized over the eligible event universe before trimming and support restriction. Every regression includes the complete set of lower-order interactions, absorbs acquirer fixed effects and two-digit-industry-by-transfer-year fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, and clusters standard errors by acquirer. Vertical bars are 95 percent confidence intervals.

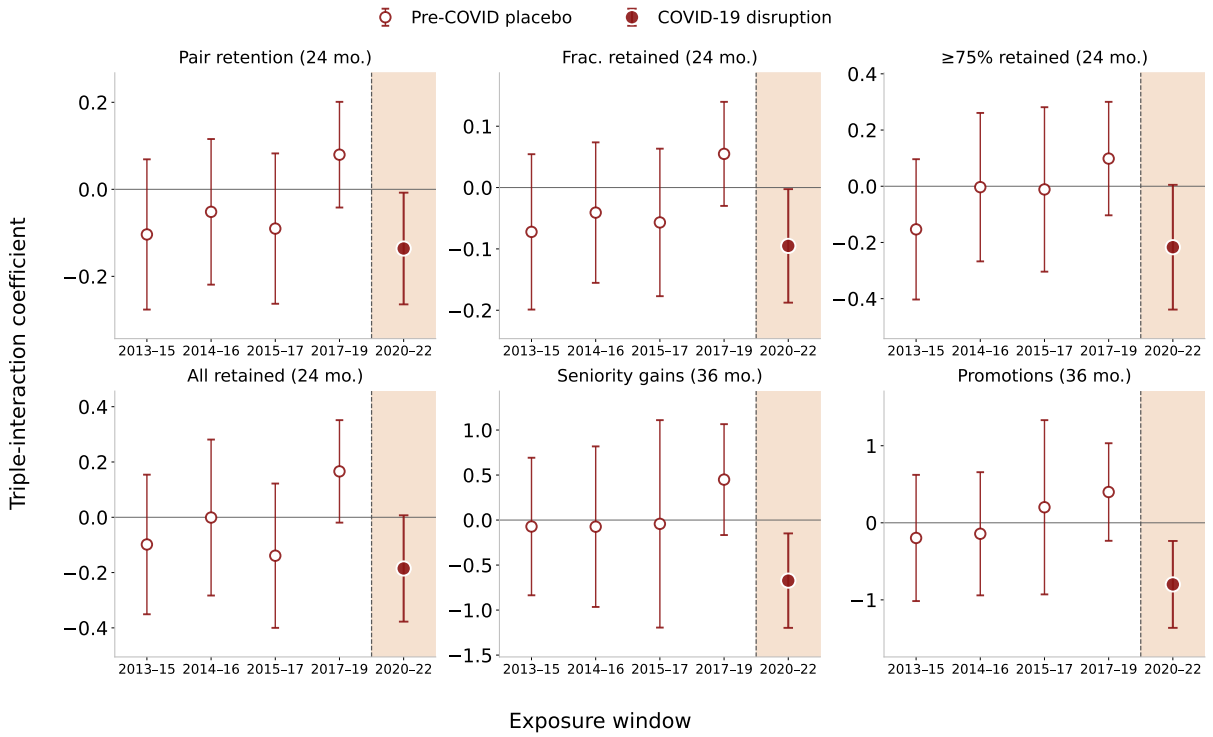


Figure A.3. COVID-19 and Placebo Estimates: Exit from the Acquirer

Notes: This figure reports differential disruption-by-pre-transfer-task-exposure gradients for worker exit 24 months after a transfer. Exit equals one when a worker is no longer employed within the acquirer's corporate family. Each marker is the coefficient on acqui-hire status interacted with disruption exposure and standardized pre-transfer task exposure. The filled marker uses exposure to the April 1, 2020–April 1, 2022 disruption. Open markers replace the actual disruption with the indicated two-year pre-pandemic window and restrict the sample to events whose 24-month outcome windows do not overlap the actual disruption.

Identified acqui-hires are acquisitions for which documentary text presents obtaining the target's employees, founders, or team as a transaction motive; the proxy does not establish that it was the primary motive. A comparison worker flow is a reconstructed group of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, these observations are called comparison flows and are not interpreted as verified direct hiring. The common-support sample compares workers joining through an acqui-hire with workers in comparison flows for transfers from 2010 through March 2020. Every weighted linear probability model includes all lower-order terms; controls for pre-transfer seniority, log modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status; and absorbs acquirer and transfer-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. Each worker's weight is the event's common-support weight divided by event size. The actual-disruption regression contains 169,559 workers, 8,228 transfer-event clusters, and 650 acquirer clusters. Each placebo regression contains 105,903 workers, 5,256 transfer-event clusters, and 498 acquirer clusters. Vertical bars are 95 percent confidence intervals based on standard errors two-way clustered by acquirer and transfer event. The regression coefficients and intervals are multiplied by 100 and plotted in percentage points; they are not average effects of acquisition.

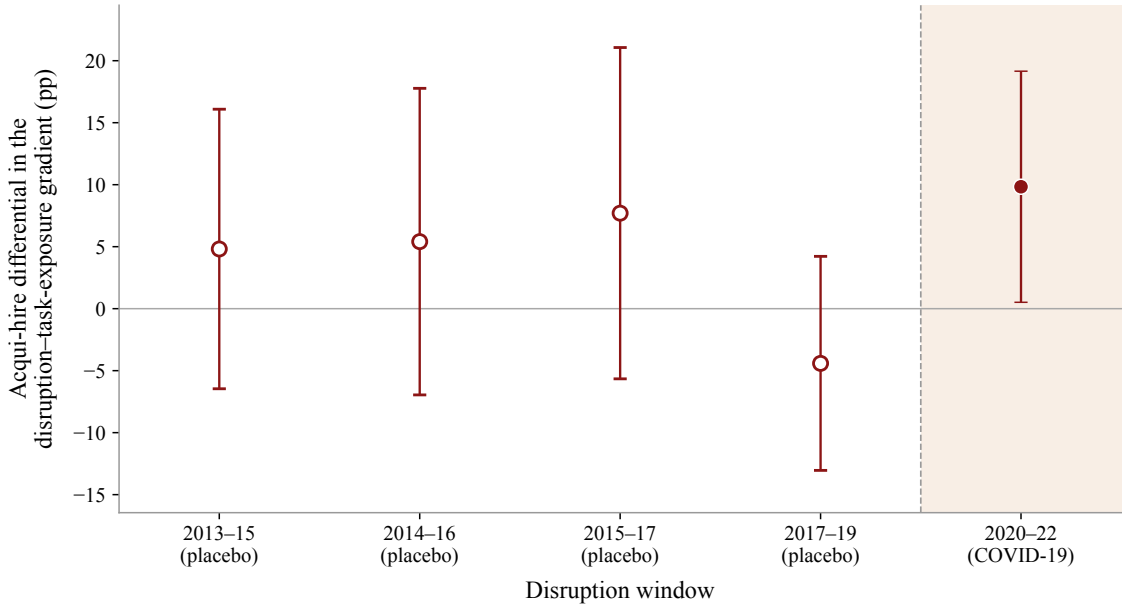


Figure A.4. COVID-19 and Placebo Estimates: Co-departure

Notes: This figure reports differential disruption-by-pre-transfer-task-exposure gradients for co-departure 24 months after a transfer. Co-departure equals one when a worker and another member of the original transferred group have both left the acquirer's corporate family by 24 months and their recorded exit dates are within 180 days. It does not require a common outside employer. Each marker is the coefficient on acqui-hire status interacted with disruption exposure and standardized pre-transfer task exposure. The filled marker uses exposure to the April 1, 2020–April 1, 2022 disruption. Open markers replace the actual disruption with the indicated two-year pre-pandemic window and restrict the sample to events whose outcome windows do not overlap the actual disruption.

Identified acqui-hires are acquisitions for which documentary text presents obtaining the target's employees, founders, or team as a transaction motive; the proxy does not establish that it was the primary motive. A comparison worker flow is a reconstructed group of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, these observations are called comparison flows and are not interpreted as verified direct hiring. The common-support sample compares workers joining through an acqui-hire with workers in comparison flows for transfers from 2010 through March 2020. Every weighted linear probability model includes all lower-order terms; controls for pre-transfer seniority, log modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status; and absorbs acquirer and transfer-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. Each worker's weight is the event's common-support weight divided by event size. The actual-disruption regression contains 169,559 workers, 8,228 transfer-event clusters, and 650 acquirer clusters. Each placebo regression contains 105,903 workers, 5,256 transfer-event clusters, and 498 acquirer clusters. Vertical bars are 95 percent confidence intervals based on standard errors two-way clustered by acquirer and transfer event. The regression coefficients and intervals are multiplied by 100 and plotted in percentage points; they are not average effects of acquisition.

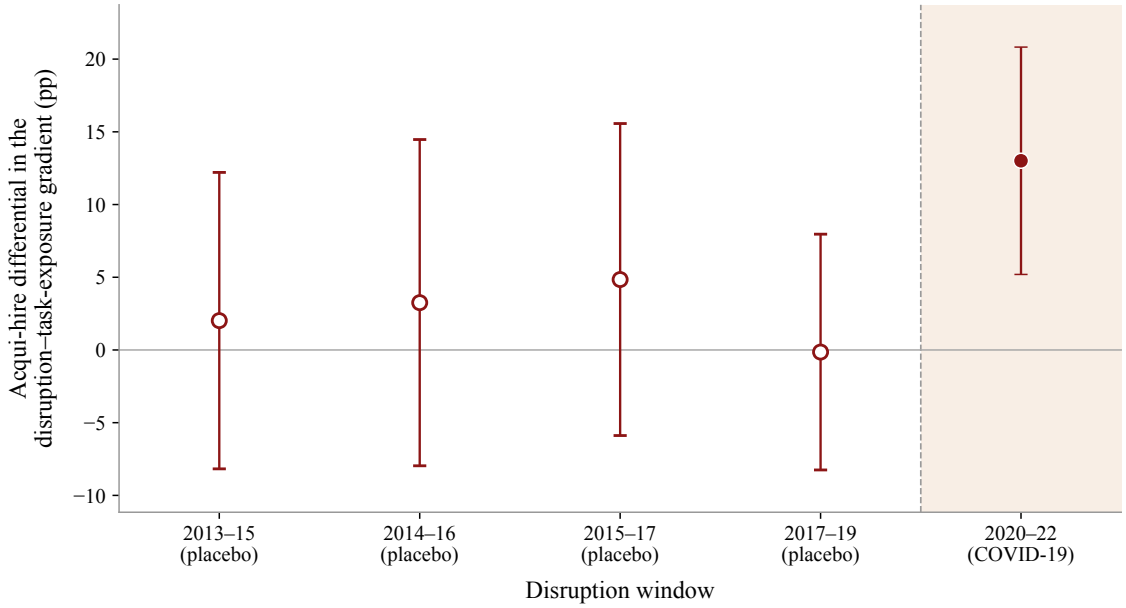


Figure A.5. COVID-19 and Placebo Estimates: Entrepreneurial Reallocation

Notes: This figure reports differential disruption-by-pre-transfer-task-exposure gradients for entrepreneurial reallocation. The event-level outcome equals one when at least two original group members leave the acquirer's corporate family within 24 months and their first observed outside employment spells satisfy the paper's entrepreneurial-employer or entrepreneurial-role criteria. Those spells can begin after the 24-month horizon and need not be at the same firm. The broad measure uses founder or cofounder titles and founding language; CEO, CTO, or vice-president-or-higher roles at firms with at most 50 employees; employment at firms no more than five years old with at most 500 employees; Y Combinator founder matches; and PitchBook startup firms with at most 500 employees. Each marker is the coefficient on acqui-hire status interacted with disruption exposure and standardized pre-transfer task exposure. The filled marker uses exposure to the April 1, 2020–April 1, 2022 disruption. Open markers replace the actual disruption with the indicated two-year pre-pandemic window and restrict the sample to events whose outcome windows do not overlap the actual disruption.
Identified acqui-hires are acquisitions for which documentary text presents obtaining the target's employees, founders, or team as a transaction motive; the proxy does not establish that it was the primary motive. A comparison worker flow is a reconstructed group of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, these observations are called comparison flows and are not interpreted as verified direct hiring. The common-support sample compares acqui-hires with comparison flows for transfers from 2010 through March 2020. Every unweighted linear probability model includes all lower-order terms, absorbs acquirer fixed effects and two-digit-industry-by-transfer-year fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. The actual-disruption regression contains 7,926 transfer events and 354 acquirers; each placebo regression contains 5,015 transfer events and 263 acquirers. Vertical bars are 95 percent confidence intervals based on standard errors clustered by acquirer. The regression coefficients and intervals are multiplied by 100 and plotted in percentage points; they are not average effects of acquisition.

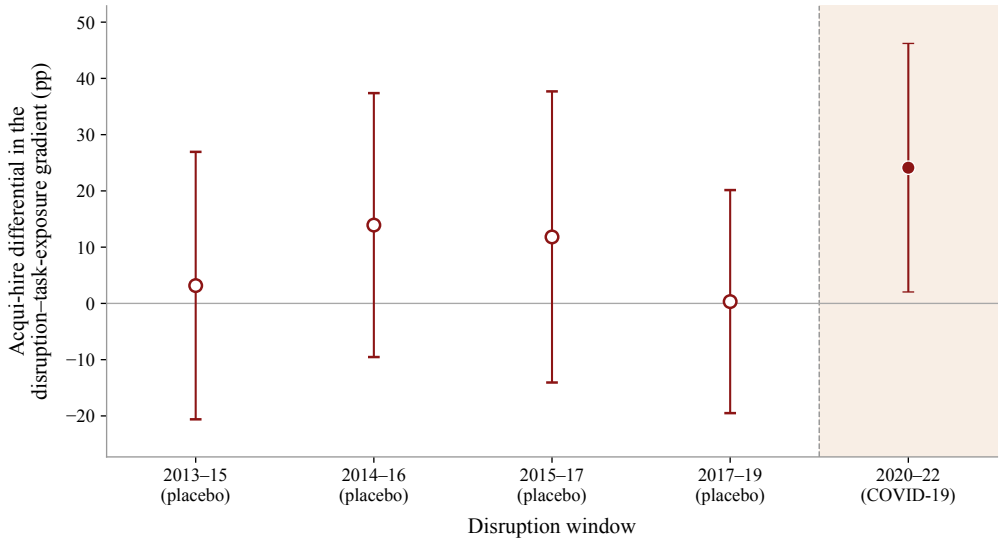
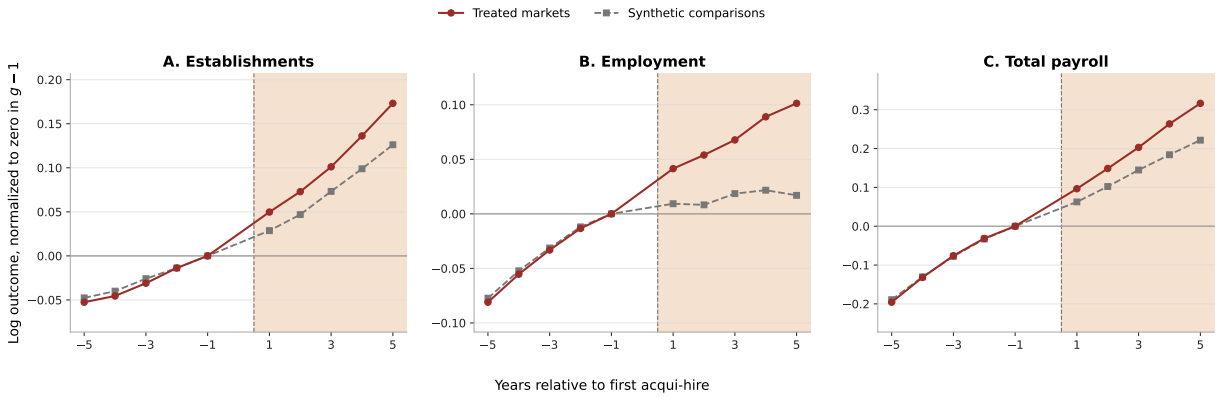


Figure A.6. Treated and Synthetic Paths in Receiving Local Industry-Markets

Notes: The figure compares treated and synthetic outcome paths for the log number of establishments, log employment, and log total annual payroll. Each path is normalized to zero in the last complete pre-transaction year, $g-1$. A market is the target's harmonized pre-acquisition four-digit industry interacted with the modal destination Core-Based Statistical Area of its transferred high-skill workers. The first qualifying acqui-hire defines year g, which is excluded. Synthetic controls are estimated within three-digit-industry-by-treatment-year blocks from never-treated markets screened on pre-treatment growth. Paths are then aggregated using treated-market counts. Subsequent acqui-hires are allowed. The establishment and employment panels use 42 blocks and 147 and 143 treated markets; the payroll panel uses 40 blocks and 139 treated markets. Outcomes are from the Quarterly Census of Employment and Wages, 2004–2023. The figure is a pre-treatment-fit diagnostic; statistical inference is reported in Figure 9 and Table 12.



A.2 Tables

Table A.1. Transaction Classification and Worker-Data Sample Lineage

Notes: The unit is a unique PitchBook transaction in the public-acquirer sample used to construct the worker-linked analysis files. "Earlier sample" reports the counts before the final eligibility rule was applied. "Final eligibility" restricts the sample to completed mergers and acquisitions whose PitchBook deal date falls between 2000 and 2025 and in which 100 percent of the target is acquired and at least one reported acquirer's ultimate parent is headquartered in the United States. Acquiring subsidiaries are mapped to their ultimate parents; target location is unrestricted. The classification and Revelio-linkage samples overlap; neither is obtained by sequentially filtering the other. Their intersection contains transactions present in both files. "Explicit personnel-motive evidence" identifies transactions in that intersection for which the documentary classifier finds an explicit statement that obtaining employees, founders, or a team motivated the acquisition. The indicator establishes a stated personnel motive, not that personnel acquisition was the transaction's primary motive. These counts describe data lineage; each empirical design applies additional outcome- and support-specific restrictions.

	Earlier sample	Final eligibility
Classified transactions	54,394	41,050
Target linked to Revelio roster	33,815	25,599
Classified ∩ target-linked	27,582	21,186
Explicit personnel-motive evidence in overlap	5,992	4,807

Table A.2. Robustness of Pre-Transaction Patent Pairing

Notes: The unit is a realized acquisition in which the direct acquirer and target each have at least one linked U.S. utility patent filed during the five years before transaction completion. The dependent variable equals one for an identified acqui-hire. Portfolio proximity measures similarity between the acquirer's and target's own patent portfolios; common cited prior art measures overlap in the prior art cited by the two firms, scaled by the smaller citation set. The separate specifications include one measure at a time. The joint specification includes both measures and controls for the logarithm of one plus citations from the acquirer to the target and from the target to the acquirer. The first two rows absorb target three-digit-NAICS-by-transaction-year fixed effects; the last two instead absorb broad PitchBook-target-industry-by-year fixed effects and therefore retain a somewhat different sample. Standard errors are clustered either by the acquirer's ultimate parent or by the direct acquirer, as indicated. Coefficients are from linear probability models and are in probability units. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Separate specifications		Joint specification		Transactions		
Specification	Portfolio proximity	Common cited prior art	Portfolio proximity	Common cited prior art	Proximity	Prior art	Joint
Target NAICS-3 × year; ultimate-parent clustering	0.0494	0.1006***	0.0286	0.1086**	2,583	2,433	2,416
	(0.0359)	(0.0376)	(0.0368)	(0.0430)			
Target NAICS-3 × year; direct-acquirer clustering	0.0494	0.1006***	0.0286	0.1086**	2,583	2,433	2,416
	(0.0349)	(0.0372)	(0.0366)	(0.0431)			
Broad target industry × year; ultimate-parent clustering	0.0490	0.0942***	0.0322	0.1029**	2,668	2,505	2,488
	(0.0324)	(0.0345)	(0.0338)	(0.0415)			
Broad target industry × year; direct-acquirer clustering	0.0490	0.0942***	0.0322	0.1029**	2,668	2,505	2,488
	(0.0321)	(0.0354)	(0.0340)	(0.0433)			

Table A.3. Robustness of Claims-Based Acquirer–Target Similarity

Notes: The unit is a completed acquisition in which the direct acquirer and target each have at least one linked U.S. utility patent filed during the five years before transaction completion. The dependent variable equals one for an identified acqui-hire. Each row is a separate linear probability model of acqui-hire status on the indicated measure of similarity between the acquirer's and target's patented output. Portfolio-centroid measures compute the cosine between equal-patent-weighted portfolio representations. "Any text support" requires at least one embedded claims patent on each side; the higher-coverage row requires at least two-thirds of each side's eligible patents to have claims text. The title-and-abstract row replaces claims with titles and abstracts on the same higher-coverage subset. Mean pairwise cosine averages similarity over all acquirer–target patent pairs rather than normalizing portfolio centroids. The final row restricts both portfolios to patents granted by transaction completion, while still using their eventual claims content because historical claims versions are not observed. Every regression absorbs target three-digit-industry-by-deal-year fixed effects and clusters standard errors by the acquirer's ultimate parent. Coefficients are in probability units. The table tests sensitivity to text representation, coverage, and timing; statistically insignificant estimates are not interpreted as proof of equivalence. Appendix B.7 gives the patent-text construction.

Text representation or sample	Coefficient	Standard error	Transactions	Acquirer clusters
Portfolio centroid; any text support	0.1732	0.2713	2,525	750
Portfolio centroid; higher text coverage	0.1840	0.2825	2,313	706
Title-and-abstract portfolio centroid	0.0619	0.2178	2,313	706
Mean pairwise claims cosine	0.4903	0.4295	2,525	750
Portfolio centroid; patents granted by completion	-0.0107	0.3411	1,740	558

Table A.4. Robustness of Worker Patenting and Co-invention

Notes: The underlying common-support population contains 1,133 identified acqui-hires and 7,095 comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. Support is imposed on transfer year, the two-digit industry of the target for an acqui-hire or the source firm for a comparison flow, arriving-group size, occupation composition, and prior-mobility cohesion. Panels A and B use worker-level linear probability models. Panel A indicates whether a worker appears on a patent assigned to the acquirer's corporate family by the stated horizon. Panel B additionally requires a pre-existing acquirer inventor as a co-inventor. The reported coefficient is the acqui-hire indicator. Baseline regressions use the event's common-support weight divided by the number of workers in the event; multiplicity-adjusted weights additionally account for workers linked to more than one comparison event. The complete-follow-up rows retain only cohorts whose stated horizon lies within the observed patent-coverage window. Regressions control for pre-transfer seniority, modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status; absorb acquirer and transfer-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow; and cluster standard errors by acquirer and transfer event. Coefficients are in probability units. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Months since entry			
	12	24	36	48
Panel A. Worker appears on a patent assigned to the acquirer				
Baseline	0.0197***	0.0292***	0.0353***	0.0370***
	(0.0064)	(0.0079)	(0.0087)	(0.0089)
Complete follow-up	0.0208***	0.0327***	0.0432***	0.0495***
	(0.0068)	(0.0097)	(0.0109)	(0.0123)
Multiplicity-adjusted weights	0.0196***	0.0289***	0.0349***	0.0365***
	(0.0064)	(0.0079)	(0.0087)	(0.0089)
Observations: baseline and adjusted weights	169,559	169,559	169,559	169,559
Observations: complete follow-up	163,650	124,263	95,332	71,807
Panel B. Worker co-invents with a pre-existing acquirer inventor				
Baseline	0.0065**	0.0117***	0.0144***	0.0169***
	(0.0027)	(0.0043)	(0.0053)	(0.0056)
Complete follow-up	0.0070**	0.0130**	0.0168***	0.0223***
	(0.0029)	(0.0051)	(0.0062)	(0.0069)
Multiplicity-adjusted weights	0.0066**	0.0117***	0.0142***	0.0166***
	(0.0028)	(0.0043)	(0.0053)	(0.0056)
Observations: baseline and adjusted weights	169,559	169,559	169,559	169,559
Observations: complete follow-up	163,650	124,263	95,332	71,807

Table A.5. Robustness of Private-Value-Weighted Patent Output

Notes: The underlying common-support population contains 1,133 identified acqui-hires and 7,095 comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. The outcome sums the market-based private-value allocation over patents produced by workers in the arriving group and assigned to the acquirer, allocating patents linked to several focal events across those events. Each cell reports the acqui-hire coefficient from a Poisson pseudo-maximum-likelihood regression of observed private-value-weighted patent output. Regressions use the event matching weight, include the logarithm of arriving-group size as an offset, absorb acquirer and transfer-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, and cluster standard errors by acquirer. The complete-follow-up row retains only cohorts whose stated horizon lies within the observed patent-coverage window. The two influence checks omit, respectively, the acqui-hire event or acquirer contributing the most observed value at each horizon. Coefficients are log conditional-mean differences, not percentage-point effects. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Months since entry			
	12	24	36	48
Baseline	1.1599***	1.2788***	1.2745***	1.2172***
	(0.3049)	(0.2039)	(0.1437)	(0.1110)
Complete follow-up	1.1599***	1.2788***	1.2752***	1.3307***
	(0.3049)	(0.2039)	(0.1438)	(0.1343)
Exclude the highest-value event	1.1599***	1.2434***	1.2181***	1.1511***
	(0.3049)	(0.1977)	(0.1379)	(0.0970)
Exclude the highest-value acquirer	1.1599***	1.2531***	1.2256***	1.1570***
	(0.3049)	(0.2002)	(0.1386)	(0.0973)
Events: baseline	4,499	4,837	5,006	5,058
Events: complete follow-up	4,499	4,837	4,943	4,002
Events: excluding highest-value event	4,499	4,836	5,005	5,057
Events: excluding highest-value acquirer	4,499	4,824	4,993	5,045

Table A.6. Allocated Private Value of Mixed-Inventor Patents

Notes: The unit is a patent–event observation in the common-support population described in Section 4.2. This population compares acqui-hires with comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. The sample is restricted to patents assigned to the acquirer's corporate family, filed within 48 months after a worker in the arriving group enters, observed in the private-value data, and naming at least two workers in the arriving group and at least one pre-existing acquirer inventor. A pre-existing acquirer inventor is outside the arriving group and was connected to the acquirer's corporate family before entry through an employment spell or an earlier patent assigned to that family. The dependent variable is $\log(1+\xi_p)$, where ξ_p is the market-based private-value amount allocated to patent p by Kogan et al. (2017). Because that measure allocates an issuer's grant-day stock-market reaction across patents granted on the same day, the outcome is an allocated private-value measure rather than a separately identified value for each patent. The regression weights each patent–event observation by the event's matching weight divided by the number of qualifying patents linked to that event and by the number of focal events to which the patent is linked; it absorbs acquirer, event-year, filing-year, and primary CPC-field fixed effects and three-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow; and clusters standard errors by acquirer. The omitted category is patents following comparison flows. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	$\log(1+\text{allocated private patent value})$
Acqui-hire	0.662^{**}
	(0.291)
Patent–event observations	466
Distinct patents	454
Transfer events	50
Acquirer clusters	20

Table A.7. Decomposition of Claims-Based Movement Toward Acquirer Technology

Notes: The underlying common-support population contains 1,133 identified acqui-hires and 7,095 comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. The table uses the subset with nonempty claims-text portfolios for the arriving group's five-year pre-entry work, the acquirer's five-year pre-entry technology, and patents produced by workers in the arriving group inside the acquirer by the stated horizon. Post-entry proximity is the cosine similarity between the post-entry and acquirer pre-entry portfolios. Pre-entry group proximity is the cosine similarity between the arriving-group and acquirer pre-entry portfolios. Movement toward the acquirer is the first measure minus the second, observation by observation.

Panel A reports the acqui-hire coefficient relative to comparison flows for each component on the same horizon-specific sample. Regressions use the event matching weights, absorb acquirer and entry-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, and cluster standard errors by acquirer. Panel B estimates movement separately within each event category on that category's usable sample, using the same weights, industry and entry-year fixed effects, and acquirer-clustered standard errors, but no acquirer fixed effects. Panel B reports adjusted movement estimates by event category; its category-specific samples need not sum to the common samples in Panel A. Coefficients are changes in cosine similarity, not percentage-point effects. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Appendix B.7 defines the portfolios and text representation.

	24 months	36 months	48 months
Panel A. Acqui-hire–comparison-flow differences on a common sample			
Post-entry proximity to acquirer technology	0.0065	0.0059*	0.0039
	(0.0044)	(0.0033)	(0.0031)
Pre-entry team proximity to acquirer technology	-0.0134**	-0.0118***	-0.0120***
	(0.0053)	(0.0035)	(0.0033)
Movement toward acquirer technology	0.0199***	0.0177***	0.0159***
	(0.0054)	(0.0038)	(0.0035)
Transfer events	266	314	327
Acquirer clusters	27	28	29
Panel B. Movement within each event category			
Comparison flows	-0.0102***	-0.0069***	-0.0038
	(0.0028)	(0.0025)	(0.0028)
Acqui-hires	0.0052*	0.0093***	0.0106***
	(0.0028)	(0.0027)	(0.0026)
Comparison-flow events	235	284	294
Comparison-flow acquirer clusters	33	39	39
Acqui-hire events	70	75	79
Acqui-hire acquirer clusters	53	55	57

Table A.8. Robustness of Claims-Based Technological Direction

Notes: The unit is a transfer event and the outcome horizon is 48 months. A comparison worker flow is a reconstructed group of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, it is called a comparison flow. Each reported coefficient is the acqui-hire difference relative to comparison flows from a matching-weighted linear regression with acquirer and entry-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, with acquirer-clustered standard errors. Continuity is the cosine similarity between post-entry output and the arriving group's pre-entry portfolio. Movement toward the acquirer is the post-entry portfolio's similarity to the acquirer's pre-entry portfolio minus the arriving group's pre-entry similarity to that portfolio. Two-sided proximity is the minimum of continuity and post-entry proximity to the acquirer and is defined only when both components are observed.
The first row requires at least one embedded claims patent in every required portfolio. The next three rows require the indicated minimum share of eligible patents in each portfolio to have claims text. Complete follow-up restricts events to those for which the full 48-month patent-filing window can be observed. In the focal-inventor exclusion, each post-entry patent is compared with a group portfolio that omits earlier patents written by the focal patent's own arriving inventor or inventors while retaining earlier patents by other members of the arriving group. The final row reports the acqui-hire–comparison difference in own-minus-placebo continuity, where the placebo is average continuity with all other text-supported arriving-group portfolios. A dash denotes an outcome that was not estimated, not a zero. Coefficients are changes in cosine similarity. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels. Appendix B.7 gives the complete construction.

Sample or measure	Continuity with incoming-group technology	Movement toward acquirer technology	Two-sided proximity	Events	Clusters
Any text support	0.0272***	0.0159***	0.0210***	327	29
	(0.0047)	(0.0035)	(0.0046)		
At least 50 percent text coverage	0.0270***	0.0145***	0.0207***	305	28
	(0.0048)	(0.0033)	(0.0046)		
At least two-thirds text coverage	0.0282***	0.0168***	0.0218***	269	26
	(0.0062)	(0.0039)	(0.0061)		
At least 75 percent text coverage	0.0299***	0.0159***	0.0228***	248	26
	(0.0063)	(0.0035)	(0.0062)		
Complete 48-month follow-up	0.0243***	—	0.0186***	284	27
	(0.0049)		(0.0049)		
Remove focal inventors' own prior patents	0.0209***	—	—	274	24
	(0.0054)				
Own-minus-placebo continuity	0.0241***	—	—	327	29
	(0.0039)				

Table A.9. Robustness of Acqui-hire Activity Across Industries

Notes: The unit is an acquirer–two-digit target-industry–year cell in a balanced 2017–2022 panel of acquirer–industry relationships with at least one acqui-hire during 2017–2019. The dependent variable is the annual acqui-hire count, except in column 5, where it is the logarithm of one plus that count. The reported regressor interacts an indicator for 2020–2022 with the industry's predetermined employment share in occupations that cannot be performed from home. Columns 1–3 use Poisson pseudo-maximum likelihood (PPML). Columns 1 and 2 absorb acquirer-by-year and target-industry fixed effects; separation and singleton removal leave 479 of the 2,886 balanced panel cells. Column 3 instead includes additive acquirer, year, and target-industry fixed effects and retains all 2,886 cells. Columns 4 and 5 are linear regressions with acquirer-by-year and target-industry fixed effects on the count and log-count outcomes, respectively.
The inverse-probability weights define high exposure as an above-median industry share and estimate its propensity using the relationship's pre-2020 mean log target employment and mean target age. High-exposure observations receive weight one; low-exposure observations receive the corresponding odds weight, capped at the 99th percentile of the low-exposure weight distribution. In the PPML columns, the coefficient is the change in the log conditional mean associated with moving the industry share from zero to one during 2020–2022. Column 4 is in acqui-hire-count units and column 5 is in log-count units. Because the estimators and dependent-variable scales differ, coefficients should not be compared across columns. Standard errors are clustered by acquirer and reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	PPML			Linear	
	Acqui-hire count (1)	Acqui-hire count (2)	Acqui-hire count (3)	Acqui-hire count (4)	log(1+ acqui-hire count) (5)
2020–2022 × industry share in nonteleworkable occupations	−2.487***	−2.321***	−2.217***	−0.668*	−0.227**
	(0.588)	(0.543)	(0.443)	(0.400)	(0.106)
Inverse-probability weights	No	Yes	Yes	Yes	Yes
Acquirer × year fixed effects	Yes	Yes	No	Yes	Yes
Additive acquirer and year fixed effects	No	No	Yes	No	No
Two-digit target-industry fixed effects	Yes	Yes	Yes	Yes	Yes
Estimation cells	479	479	2,886	2,886	2,886
Acquirer clusters	69	69	395	395	395

Table A.10. Acqui-hire Activity by Acquirer Financial Constraints

Notes: The table repeats the industry-level organizational-form analysis separately by acquirers' pre-pandemic financial-constraint measures. The underlying population is a balanced 2017–2022 panel of acquirer–two-digit target-industry relationships with at least one acqui-hire during 2017–2019; the table restricts that population to Compustat-linked acquirers with the required fiscal-year 2019 information. The outcome is the annual acqui-hire count in an acquirer–industry cell. The reported coefficient interacts 2020–2022 with the industry's predetermined employment share in occupations that cannot be performed from home. The size–age index follows Hadlock and Pierce; the Whited–Wu index follows Whited and Wu. Both are measured using fiscal-year 2019 information, and larger values indicate greater financial constraints. Panel A divides acquirers at the sample median. Panel B retains the least-constrained and most-constrained quartiles.
Each column is estimated separately by PPML with additive acquirer, year, and two-digit target-industry fixed effects. The coefficient is the PPML log-conditional-mean change associated with moving the industry exposure share from zero to one during 2020–2022. Unlike the baseline specification in Table 3, these regressions do not absorb acquirer-by-year shocks. They are therefore diagnostics for the possibility that the industry pattern is confined to financially constrained acquirers, not preferred estimates of heterogeneous effects. The table reports no formal test of equality across subgroup coefficients; differences in statistical significance across columns do not establish coefficient differences. The pre-period test in Panel A is a joint test of the 2017 and 2018 exposure interactions relative to 2019. Standard errors are clustered by acquirer and reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Size–age index		Whited–Wu index	
	Less constrained	More constrained	Less constrained	More constrained
	(1)	(2)	(3)	(4)
Panel A. Median split				
2020–2022 × industry share in nonteleworkable occupations	−1.621***	−1.832**	−1.445***	−2.195***
	(0.495)	(0.768)	(0.474)	(0.797)
Pre-period test, *p*-value	0.649	0.406	0.625	0.510
Observations	978	996	990	924
Acquirer clusters	133	134	129	129
Panel B. Extreme quartiles				
2020–2022 × industry share in nonteleworkable occupations	−1.644***	−1.694*	−1.371*	−2.536*
	(0.620)	(0.881)	(0.746)	(1.431)
Observations	522	474	510	462
Acquirer clusters	67	67	65	65
Acquirer fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Two-digit target-industry fixed effects	Yes	Yes	Yes	Yes

Table A.11. Within-Acqui-hire Team Preservation Across Outcome Horizons

Notes: The unit is an acqui-hire through which at least three workers join the acquirer. Each cell is the coefficient on disruption exposure interacted with standardized post-transfer task exposure from a separate weighted linear regression. Disruption exposure is the fraction of the indicated post-entry outcome window overlapping April 2020–April 2022. Post-transfer task exposure is constructed from high-skill workers' post-transfer occupations and equals Dingel–Neiman nonteleworkability multiplied by the mean O*NET importance of nine interaction-oriented activities. It is standardized over the 835-event matched pool before fixed-effect singleton removal. Pairwise co-retention is the share of distinct worker pairs whose two members both remain at the acquirer. The remaining outcomes are the fraction of workers retained and indicators for retaining at least 50 percent, at least 75 percent, or every worker.

Regressions control for disruption exposure, post-transfer task exposure, the logarithm of one plus arriving-group size, the logarithm of one plus target employment, target age, high-skill share, and professional-managerial share. They also interact post-transfer task exposure with the two logarithmic controls and target age, absorb acquirer and two-digit target-industry-by-transfer-year fixed effects, and use Hájek weights that equalize represented exposure quartiles within strata defined by target industry, transfer year, and above-median indicators for arriving-group size, target size, high-skill share, and professional-managerial share. Standard errors are clustered by acquirer. Each coefficient is the change in the outcome's post-transfer task-exposure gradient when disruption exposure rises from zero to one. Coefficients are in proportion units, so multiplying by 100 gives the percentage-point change in that gradient for a one-unit increase in standardized post-transfer task exposure. The 36- and 48-month estimates are reported to show the horizon profile; the causal retention claim rests on the 24-month results and their pre-period diagnostics. The dagger marks complete retention, which has a significant 2015–2017 placebo and a rejected cohort-slope test and is not used to support the retention claim. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Outcome horizon	Pairwise co-retention (1)	Fraction of workers retained (2)	At least 50% retained (3)	At least 75% retained (4)	All workers retained† (5)
24 months	−0.267***	−0.181***	−0.203***	−0.428***	−0.151*
	(0.065)	(0.043)	(0.071)	(0.119)	(0.088)
36 months	−0.075	−0.078	−0.225*	−0.162	0.026
	(0.065)	(0.053)	(0.124)	(0.133)	(0.052)
48 months	−0.009	−0.017	−0.098	0.066	0.046
	(0.067)	(0.063)	(0.142)	(0.121)	(0.048)
Observations per regression	502	502	502	502	502
Acquirer clusters	133	133	133	133	133
Acquirer fixed effects	Yes	Yes	Yes	Yes	Yes
Two-digit target industry × transfer-year fixed effects	Yes	Yes	Yes	Yes	Yes
Hájek weights	Yes	Yes	Yes	Yes	Yes

Table A.12. Worker-Level Retention and Career-Progression Diagnostics

Notes: Each cell in Panel A comes from a separate worker-level regression. Column 1 uses the within-acqui-hire sample and reports the coefficient on disruption exposure interacted with standardized post-transfer task exposure. Columns 2–4 compare acqui-hires with comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. These columns report the coefficient on acqui-hire status interacted with disruption exposure and standardized pre-transfer task exposure. The placebo rows replace the April 2020–April 2022 disruption with the indicated two-year pre-pandemic window and exclude workers whose outcome window overlaps the actual disruption. "Retained" indicates that the worker remains at the acquirer 24 months after entry. "Seniority gain" is cumulative positive movement in Revelio's ordered seniority measure within the acquirer, excluding the entry-rank change.

Column 1 is a weighted linear probability model. It controls for disruption exposure, standardized post-transfer task exposure, log arriving-group size, log target employment, target age, high-skill share, and professional-managerial share; interacts post-transfer task exposure with the two logarithmic controls and target age; and includes a linear post-transfer-task-exposure-by-transfer-year trend. It also includes worker controls for pre-transfer seniority, log modeled compensation, experience, estimated age, advanced and doctoral degrees, employer prestige, high-skill status, and professional-managerial status; and acquirer, target-industry-by-transfer-year, post-transfer-occupation-by-year, and pre-transfer-occupation-by-year fixed effects. Each worker receives the event's Hájek weight without division by arriving-group size, and standard errors are clustered by acquirer. Column 2 is a weighted linear probability model with every lower-order term of the triple interaction, worker controls for pre-transfer seniority, log modeled compensation, experience, high-skill status, and professional-managerial status, and additive acquirer and transfer-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. Its weight is the matching-cell event weight divided by arriving-group size, so the total weight of an event does not rise mechanically with its number of workers. Columns 3 and 4 are unweighted PPML regressions with the same acqui-hire-versus-comparison-flow controls and fixed effects. Their standard errors are two-way clustered by acquirer and transfer event.

Linear retention coefficients are in proportion units; multiplying by 100 gives percentage points. PPML coefficients are changes in the log conditional mean of cumulative seniority gain. Exponentiating the triple-interaction coefficient gives the multiplicative change, for a one-unit increase in the specified task exposure, in the acqui-hire-to-comparison expected-outcome rate ratio when disruption exposure rises from zero to one. It is not a change in seniority-rank levels or promotion probabilities. Differences in observations across outcomes reflect outcome observability and PPML separation. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Within acqui-hires Retained, 24 months (1)	Acqui-hire versus flow Retained, 24 months (2)	Acqui-hire versus flow Seniority gain, 24 months (3)	Acqui-hire versus flow Seniority gain, 36 months (4)
Panel A. Actual and placebo disruption estimates				
Actual disruption	−0.102**	−0.098**	−0.777***	−0.681**
	(0.045)	(0.048)	(0.273)	(0.265)
2013–2015 placebo window	−0.018	−0.055	0.440	0.167
	(0.060)	(0.060)	(0.294)	(0.346)
2014–2016 placebo window	−0.039	−0.048	0.337	0.236
	(0.074)	(0.058)	(0.488)	(0.407)
2015–2017 placebo window	−0.033	−0.064	−0.386	−0.087
	(0.050)	(0.062)	(0.471)	(0.563)
2017–2019 placebo window	0.071	0.049	0.292	0.161
	(0.106)	(0.044)	(0.306)	(0.286)
Panel B. Sample and estimation				
Observations, actual disruption	24,929	169,559	151,255	152,274
Observations per placebo regression	15,994	105,903	92,888	67,241
Estimator	Linear	Linear	PPML	PPML
Weights	Hájek	Matching-cell event weight ÷ group size	None	None

Table A.13. Within-Acqui-hire Pre-pandemic Retention Diagnostics

Notes: Each placebo row replaces the April 2020–April 2022 disruption with the indicated two-year pre-pandemic window. Placebo exposure is the fraction of the 24-month post-entry outcome window that overlaps that interval. Transactions whose actual 24-month outcome window overlaps April 2020–April 2022 are excluded. Each weighted linear regression controls for placebo exposure, standardized post-transfer task exposure, the logarithm of one plus arriving-group size, the logarithm of one plus target employment, target age, high-skill share, and professional-managerial share. It interacts post-transfer task exposure with the two logarithmic controls and target age; absorbs acquirer and two-digit target-industry-by-transfer-year fixed effects; uses Hájek weights that equalize represented task-exposure quartiles within strata defined by target industry, transfer year, and above-median indicators for arriving-group size, target size, high-skill share, and professional-managerial share; and clusters standard errors by acquirer. The coefficient is the interaction of placebo exposure and standardized post-transfer task exposure and is reported in proportion units. Pairwise co-retention is the share of entrant pairs whose members both remain at the acquirer; fraction retained is the share of entrants who remain; and the threshold outcomes indicate whether at least 50, 75, or 100 percent of the group remains.
The cohort-slope row reports the p-value from a joint test that the post-transfer task-exposure slopes for the pooled 2010–2014 cohort, 2015 cohort, and 2016 cohort equal the slope in the omitted 2017 cohort. Appendix Figure A.1 plots the pairwise co-retention coefficients. The dagger marks complete retention: its 2015–2017 placebo and cohort-slope test reject at the five-percent level, so the paper reports but does not rely on that outcome. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Pairwise co-retention (1)	Fraction of workers retained (2)	At least 50% retained (3)	At least 75% retained (4)	All workers retained† (5)
2013–2015 placebo window	−0.022	−0.019	−0.052	−0.080	−0.029
	(0.057)	(0.040)	(0.073)	(0.093)	(0.095)
2014–2016 placebo window	−0.111	−0.083*	−0.106	−0.171	−0.034
	(0.072)	(0.050)	(0.086)	(0.134)	(0.122)
2015–2017 placebo window	−0.067	−0.032	0.022	0.064	−0.164**
	(0.056)	(0.038)	(0.067)	(0.136)	(0.082)
2017–2019 placebo window	0.005	−0.006	−0.021	−0.066	0.066
	(0.072)	(0.047)	(0.066)	(0.118)	(0.116)
Pre-2017 cohort-slope test, p-value	0.963	0.795	0.598	0.111	0.024†
Observations per placebo regression	288	288	288	288	288
Acquirer clusters	79	79	79	79	79
Acquirer fixed effects	Yes	Yes	Yes	Yes	Yes
Two-digit target industry × transfer-year fixed effects	Yes	Yes	Yes	Yes	Yes
Hájek weights	Yes	Yes	Yes	Yes	Yes

Table A.14. Sensitivity of Team-Preservation Gradients

Notes: Panel A uses the within-acqui-hire design. The baseline post-transfer task-exposure index is constructed from occupations associated with workers' post-transfer spells at the acquirer. The second row restricts the index to high-skill workers whose acquirer spell begins before April 1, 2020; the third removes transactions completed during the first quarter of 2020; and the fourth constructs the index from workers' pre-transfer occupations. Each alternative index is standardized over its applicable matched pool. Each weighted linear regression controls for disruption exposure, the applicable task exposure, the logarithm of one plus arriving-group size, the logarithm of one plus target employment, target age, high-skill share, and professional-managerial share. It interacts the applicable task exposure with the two logarithmic controls and target age; absorbs acquirer and two-digit target-industry-by-transfer-year fixed effects; uses Hájek weights that equalize represented task-exposure quartiles within strata defined by target industry, transfer year, and above-median indicators for arriving-group size, target size, high-skill share, and professional-managerial share; and clusters standard errors by acquirer. The pre-transfer task-exposure index measures the workers' prior task environment rather than the tasks in which the acquirer deploys them, so the final row is an alternative measurement object rather than a replication of the post-transfer task gradient.

Panel B compares 1,133 acqui-hires with 7,095 comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. These counts precede outcome-specific fixed-effect attrition. Common support is imposed within transfer year, the two-digit industry of the target for an acqui-hire or the source firm for a comparison flow, six arriving-group-size bins, eight occupation-composition categories, and two prior-mobility-cohesion bins. The composition categories cross four bins of the professional-managerial share with an indicator that all movers are high-skill. The baseline is unweighted. The second row excludes first-quarter 2020 transfers. The final row gives the two event categories equal total weight within each common-support cell; it does not balance target size or age against source-firm size or age. Pre-transfer task exposure is constructed symmetrically from movers' pre-transfer occupations and standardized over the eligible event universe before trimming and support restriction. Every Panel B regression includes acqui-hire status, disruption exposure, pre-transfer task exposure, and all lower-order terms of their triple interaction, and absorbs acquirer fixed effects and two-digit-industry-by-transfer-year fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. Standard errors are clustered by acquirer. Coefficients are in proportion units; multiplying by 100 gives percentage points. Panel A reports the within-acqui-hire disruption-by-post-transfer-task-exposure gradient. Panel B reports how much that gradient differs between acqui-hires and comparison flows; it is not an average effect of acquisition. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Pairwise co-retention (1)	Fraction of workers retained (2)	Observations (3)	Acquirer clusters (4)
Panel A. Within acqui-hires				
Baseline: post-transfer occupations	−0.267***	−0.181***	502	133
	(0.065)	(0.043)		
Use only acquirer spells beginning before April 2020	−0.229***	−0.160***	500	133
	(0.071)	(0.046)		
Exclude transactions completed in 2020Q1	−0.244***	−0.171***	488	130
	(0.070)	(0.046)		
Use pre-transfer occupations	−0.060	−0.053	502	133
	(0.059)	(0.040)		
Panel B. Acqui-hires versus comparison worker flows				
Baseline: unweighted	−0.136**	−0.095**	7,926	354
	(0.065)	(0.047)		
Exclude transactions completed in 2020Q1	−0.154**	−0.101*	7,738	344
	(0.075)	(0.054)		
Reweight categories within common-support cells	−0.120*	−0.083*	7,926	354
	(0.072)	(0.047)		

Table A.15. Differential Team-Preservation Gradients: Acqui-hires versus Comparison Worker Flows

Notes: The unit is a group of at least three workers transferred to a common acquirer. The common-support sample contains 1,133 acqui-hires and 7,095 comparison worker flows before fixed-effect singleton removal. A comparison worker flow contains at least three workers who move from the same source firm to the same employer within the acquirer's corporate family within 60 days and has prior-workplace linkage above that of equal-size source-firm placebo groups; Appendix B.6 gives the construction, which does not verify the workers' contracting arrangement. Below, it is called a comparison flow. Common support is imposed on transfer year, the two-digit industry of the target for an acqui-hire or the source firm for a comparison flow, arriving-group size, occupation composition, and prior-mobility cohesion. Pre-transfer task exposure is constructed from movers' pre-transfer occupations and standardized over the eligible event universe before trimming and support restriction. Each cell reports the coefficient on acqui-hire status interacted with disruption exposure and pre-transfer task exposure from a separate unweighted linear regression. Every regression includes all constituent lower-order terms, absorbs acquirer fixed effects and two-digit-industry-by-transfer-year fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, and clusters standard errors by acquirer. Disruption exposure is the fraction of the indicated post-entry outcome window overlapping April 2020–April 2022. Coefficients are in proportion units and measure how much more the disruption-by-pre-transfer-task-exposure retention gradient changes for acqui-hires than for comparison flows. They are not average acquisition effects. The longer-horizon estimates are shown to characterize the timing of the response; the retention claim rests on the 24-month estimates and their pre-pandemic placebo tests. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Outcome horizon		
	24 months	36 months	48 months
	(1)	(2)	(3)
Pairwise co-retention	−0.136**	−0.049	0.071
	(0.065)	(0.071)	(0.073)
Fraction of workers retained	−0.095**	−0.059	0.038
	(0.047)	(0.060)	(0.067)
At least 75% retained	−0.217*	−0.053	0.038
	(0.113)	(0.119)	(0.116)
All workers retained	−0.185*	−0.019	0.041
	(0.098)	(0.076)	(0.068)
Observations	7,926	7,926	7,926
Acquirer clusters	354	354	354
Acquirer fixed effects	Yes	Yes	Yes
Two-digit pre-transfer industry × transfer-year fixed effects	Yes	Yes	Yes

Table A.16. Actual and Placebo Gradients in the Target–Source-Firm-Matched Sample

Notes: The target–source-firm-matched sample begins with 876 acqui-hires and 3,215 comparison worker flows, reconstructed groups from a common source firm entering a common acquirer; below, they are called comparison flows. Matching uses two-digit industry, transfer year, employment and age bins, and mover occupation composition, with industry, employment, and age defined for the target in an acqui-hire and the source firm in a comparison flow. The comparison flows are not interpreted as verified direct-hiring transactions. The sample contains 3,893 group-level observations after data-quality exclusions and fixed-effect singleton removal. Group outcomes are pairwise co-retention, the fraction retained, and indicators for retaining at least 50 percent, at least 75 percent, or every arriving worker 24 months after entry. "Worker retained" is an indicator that an individual remains at the acquirer 24 months after entry.
Column 1 reports the coefficient on acqui-hire status interacted with exposure to the April 2020–April 2022 disruption and standardized pre-transfer task exposure. Columns 2–5 replace the actual disruption with the indicated two-year pre-pandemic window, exclude observations whose outcome windows overlap the actual disruption, and otherwise preserve the corresponding specification. Pre-transfer task exposure uses movers' pre-transfer occupations and is standardized once over the original 4,104-event matched population. Team regressions include all constituent lower-order terms of the triple interaction; controls for log arriving-group size; log target employment for an acqui-hire or source-firm employment for a comparison flow; the corresponding target or source-firm age; and the high-skill and professional-managerial shares; and interactions of pre-transfer task exposure with those five controls. They use the matching-cell Hájek weights, absorb acquirer fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, plus transfer-year effects interacted with five broad sectors, and cluster standard errors by acquirer. The worker regression adds pre-transfer seniority, log modeled compensation, years of experience, advanced-degree and doctoral-degree indicators, employer prestige, and high-skill and professional-managerial indicators. Its weight equals the inherited event-level Hájek weight divided by arriving-group size, so each event's total worker weight equals its team-level weight. Team and worker samples differ because worker-level covariates and outcomes impose additional observability requirements. All coefficients are differential disruption-by-pre-transfer-task-exposure gradients in proportion units; multiplying by 100 gives percentage points. Standard errors are clustered by acquirer and reported in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Actual disruption 2020–2022 (1)	Placebo 2013–2015 (2)	Placebo 2014–2016 (3)	Placebo 2015–2017 (4)	Placebo 2017–2019 (5)
Pairwise co-retention	−0.261***	0.062	−0.074	−0.104	0.075
	(0.098)	(0.095)	(0.104)	(0.098)	(0.080)
Fraction retained	−0.153**	0.067	−0.023	−0.074	0.034
	(0.069)	(0.073)	(0.076)	(0.071)	(0.061)
At least 50 percent retained	0.040	0.136	0.038	−0.134	−0.032
	(0.092)	(0.108)	(0.112)	(0.096)	(0.077)
At least 75 percent retained	−0.465***	−0.048	−0.145	−0.057	0.146
	(0.125)	(0.143)	(0.145)	(0.177)	(0.122)
Complete retention	−0.422***	0.017	−0.037	0.007	0.148
	(0.114)	(0.110)	(0.132)	(0.106)	(0.096)
Worker retained	−0.121**	0.039	−0.016	−0.027	0.044
	(0.050)	(0.061)	(0.065)	(0.056)	(0.051)
Team-level observations	3,893	2,290	2,290	2,290	2,290
Worker-level observations	41,011	25,765	25,765	25,765	25,765

Table A.17. Within-Acqui-hire Career Progression and Pre-pandemic Placebos

Notes: The unit is an acqui-hire through which at least three workers join the acquirer. The outcome is the average across these workers of cumulative positive changes in Revelio's ordered seniority measure during the indicated post-entry horizon. The change recorded when a worker first enters the acquirer is excluded. The outcome therefore measures within-acquirer advancement in seniority-rank units; it is not a training expenditure, a direct measure of human-capital accumulation, or a worker-welfare measure.

Panel A reports the coefficient on exposure to the April 2020–April 2022 disruption interacted with standardized post-transfer task exposure. Panel B replaces the actual disruption with the indicated pre-pandemic two-year window and excludes observations whose outcome window overlaps the actual disruption. Regressions are weighted linear models. They control for disruption exposure, post-transfer task exposure, the logarithm of one plus arriving-group size, the logarithm of one plus target employment, target age, high-skill share, and professional-managerial share. For the career outcome, post-transfer task exposure is interacted with all five group controls. Regressions absorb acquirer and two-digit target-industry-by-transfer-year fixed effects, use the Hájek matching weights, and cluster standard errors by acquirer. Each coefficient is the change in the post-transfer task-exposure gradient in average cumulative positive seniority gain when disruption exposure rises from zero to one. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Outcome horizon		
	24 months	36 months	48 months
	(1)	(2)	(3)
Panel A. COVID-19 disruption			
Disruption exposure × task-based exposure	−0.065*	−0.092**	−0.121***
	(0.033)	(0.038)	(0.044)
Observations	502	502	502
Acquirer clusters	133	133	133
Panel B. Pre-COVID placebo disruptions			
2013–2015 placebo window	0.010	0.023	−0.043
	(0.028)	(0.074)	(0.119)
2014–2016 placebo window	0.039	−0.024	−0.011
	(0.037)	(0.085)	(0.119)
2015–2017 placebo window	0.044	−0.140	−0.017
	(0.051)	(0.086)	(0.120)
2017–2019 placebo window	−0.042	−0.003	0.058
	(0.034)	(0.083)	(0.113)
Observations per placebo regression	288	214	153
Acquirer clusters	79	60	46
Acquirer fixed effects	Yes	Yes	Yes
Two-digit target industry × transfer-year fixed effects	Yes	Yes	Yes
Hájek weights	Yes	Yes	Yes

Table A.18. Robustness of Coordination-Sensitive Innovative Production

Notes: This table reports robustness and pre-disruption diagnostics for innovation following identified acqui-hires relative to comparison worker flows, reconstructed groups of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, they are called comparison flows. Transfers occur from 2010 through March 2020. In the upper panel, the unit is an arriving inventor group: the inventors joining through a particular acqui-hire and the reconstructed arriving group for a comparison flow. Each cell reports the coefficient on acqui-hire status interacted with disruption exposure and the applicable task exposure from a separate PPML regression. The disruption measure is the share of the stated post-transfer horizon that overlaps April 1, 2020–April 1, 2022; the second row of the upper panel instead ends the disruption window on April 1, 2021. The baseline post-transfer task exposure is constructed from occupations at the acquirer; the final row instead uses pre-transfer task exposure. Member-years enter as exposure except in the third row, which uses the number of members present.

Regressions include every constituent lower-order term, control for the logarithm of one plus the number of incoming inventors, absorb pre-transfer-industry-by-transfer-year fixed effects except in the fourth row, and cluster standard errors by acquirer. The outcomes count arriving-inventor–patent incidences on patents produced by at least two members of the group and no pre-existing acquirer inventor, or weight those incidences by adjusted forward citations or allocated private patent value. The coefficients are log conditional-mean gradients, not percentage-point changes or average effects of acquisition.

Four pre-COVID pseudo-shock regressions accompany each reported estimate. † marks a cell for which at least one corresponding pseudo-shock estimate is negative with a p-value below 0.10. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Group-internal patenting rate			Adjusted-citation-weighted rate			Private-value-weighted rate		
	24 mo. (1)	36 mo. (2)	48 mo. (3)	24 mo. (4)	36 mo. (5)	48 mo. (6)	24 mo. (7)	36 mo. (8)	48 mo. (9)
Baseline	-3.749^{***}	-2.982^{***}	-3.166^{**}	-6.643^{***}	-6.924^{***}	$-5.542^{\dagger}$	-4.837^{***}	-5.154^{***}	-3.931^{**}
	(0.724)	(1.098)	(1.434)	(1.283)	(1.767)	(3.839)	(0.947)	(1.454)	(1.960)
Disruption through April 2021	-4.389^{***}	-5.321^{**}	-6.408^{**}	-7.853^{***}	-14.534^{***}	$-9.577^{\dagger}$	-6.443^{***}	-10.214^{***}	-7.892^{**}
	(1.110)	(2.570)	(2.660)	(1.575)	(5.295)	(5.882)	(1.321)	(3.550)	(3.298)
Members present as exposure	-4.225^{***}	-3.407^{**}	-2.132^{*}	-12.236^{***}	-7.871^{***}	-2.998	-5.128^{***}	-5.450^{***}	-4.267^{*}
	(0.893)	(1.473)	(1.190)	(3.748)	(2.322)	(2.556)	(1.131)	(1.487)	(2.341)
Pre-transfer-industry × transfer-year fixed effects	-3.471^{***}	-2.321^{**}	-2.641^{**}	-8.166^{***}	-6.862^{***}	-4.373^{*}	-4.183^{**}	-3.262^{**}	-3.136^{**}
	(1.322)	(1.017)	(1.203)	(1.415)	(1.173)	(2.250)	(1.652)	(1.408)	(1.579)
Pre-transfer task exposure	-1.147^{*}	-1.182^{*}	-1.685	$-4.633^{\dagger}$	$-2.426^{\dagger}$	$0.286^{\dagger}$	0.140	$-1.512^{**\dagger}$	-3.596^{*}
	(0.593)	(0.609)	(1.345)	(3.008)	(2.875)	(2.177)	(0.770)	(0.680)	(1.879)

Table A.19. Threshold and Timing Sensitivity for Entrepreneurial Reallocation

Notes: This table reports 24-month estimates for alternative definitions of entrepreneurial reallocation. Identified acqui-hires are acquisitions for which documentary text presents obtaining the target's employees, founders, or team as a transaction motive; the proxy does not establish that it was the primary motive. A comparison worker flow is a reconstructed group of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, it is called a comparison flow. The common-support population compares these acquisitions with comparison flows for transfers from 2010 through March 2020. Fixed-effect singleton removal leaves 7,926 transfer events—856 acqui-hires and 7,070 comparison flows—and 354 acquirers. Each cell is the coefficient on acqui-hire status interacted with disruption exposure and standardized pre-transfer task exposure. Disruption exposure is the share of the 24-month interval following the transfer that overlaps April 1, 2020–April 1, 2022. Pre-transfer task exposure is constructed from the arriving group's pre-transfer occupations. Every constituent lower-order term is included.

The first row equals one when the stated number of original group members leave the acquirer's corporate family within 24 months and their first observed outside employment spells satisfy the paper's entrepreneurial-employer or entrepreneurial-role criteria; those spells can begin after the 24-month horizon. The second row requires the qualifying outside spell itself to begin within 24 months. The broad entrepreneurial measure combines founder titles or founding language, CEO or CTO positions at firms with at most 50 employees, vice-president-or-higher positions at firms with at most 50 employees, employment at firms no more than five years old with at most 500 employees, and matches to Y Combinator founders or PitchBook startups with at most 500 employees. Multiple members need not enter the same outside firm.

Each cell in Panel A comes from a separate unweighted linear probability model with acquirer fixed effects and two-digit-industry-by-transfer-year fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. Panel B reports the column 2 specification using the event's common-support weight. Standard errors are clustered by acquirer and reported in parentheses. Coefficients are in proportion units. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Required number of original group members		
	At least one (1)	At least two (2)	At least three (3)
Panel A. Threshold and timing definitions			
Exit by 24 months; later entrepreneurial reallocation	0.204*	0.241**	0.277***
	(0.122)	(0.113)	(0.096)
Entrepreneurial reallocation begins by 24 months	0.268**	0.190*	0.305***
	(0.121)	(0.111)	(0.095)
Panel B. Weighting sensitivity			
Exit by 24 months; later entrepreneurial reallocation		0.259**	
		(0.121)	

Table A.20. Descriptive Associations with Entrepreneurial Reallocation

Notes: This table reports descriptive conditional associations with worker-level entrepreneurial reallocation for transfers from 2010 through March 2020. Identified acqui-hires are acquisitions for which documentary text presents obtaining the target's employees, founders, or team as a transaction motive; the proxy does not establish that it was the primary motive. The outcome equals one when a worker leaves the acquirer's corporate family by the stated horizon and the worker's first observed outside employment spell satisfies the paper's entrepreneurial-employer or entrepreneurial-role criteria. The entrepreneurial measure combines founder or cofounder titles and founding language; CEO, CTO, or vice-president-or-higher roles at firms with at most 50 employees; employment at firms no more than five years old with at most 500 employees; Y Combinator founder matches; and PitchBook startup firms with at most 500 employees. The qualifying outside spell can begin after the stated horizon.

A comparison worker flow is a reconstructed group of at least three workers moving from the same source firm to the same employer within the acquirer's corporate family within 60 days and whose prior-workplace linkage exceeds that of equal-size source-firm placebo groups; below, it is called a comparison flow. Panel A compares workers joining through acqui-hires with workers in comparison flows. The displayed estimate is the average adjusted acqui-hire–comparison-flow difference over the weighted estimating-sample distribution of standardized pre-transfer task exposure. Linear probability models include acqui-hire status, pre-transfer task exposure, their interaction, and controls for pre-transfer seniority, log modeled compensation, experience, and indicators for broad high-skill status and a narrower senior technical or managerial status. They absorb acquirer and transfer-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow. Each worker's weight is the event's common-support weight divided by event size, and standard errors are two-way clustered by acquirer and transfer event.

Panel B restricts the sample to workers joining through an acqui-hire. "Prior founder" identifies workers with a founder, CEO, or CTO role at the target before acquisition. Pre-transfer seniority is standardized within the common-support worker panel. The unweighted regressions include the remaining worker controls and the same additive fixed effects; standard errors are clustered by acquirer. The estimates in both panels are conditional associations, not differential responses to the COVID-19 disruption. Coefficients are in proportion units, standard errors are in parentheses, and ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	Exit by horizon; later entrepreneurial destination or role		
	24 months	36 months	48 months
	(1)	(2)	(3)
Panel A. Acqui-hire versus comparison worker flow			
Average adjusted difference	0.040***	0.073***	0.111***
	(0.005)	(0.007)	(0.009)
Observations	169,559	169,559	169,541
Acquirer clusters	650	650	650
Transfer-event clusters	8,228	8,228	8,228
Panel B. Associations within acqui-hires			
Prior founder, CEO, or CTO role	0.026*	0.088***	0.153***
	(0.014)	(0.015)	(0.016)
Pre-transfer seniority (standardized)	0.014***	0.024***	0.031***
	(0.003)	(0.003)	(0.003)
Observations	22,358	22,358	22,340
Acquirer clusters	591	591	591

Table A.21. Technological Direction and External Follow-on: Scope and Outcome Sensitivity

Notes: The unit is a patent produced after an identified acqui-hire, naming at least one original inventor who joined through the acqui-hire and assigned to the acquirer's corporate family. The regressor is the standardized claims-text proximity of the focal patent to the acquirer's pre-transaction patent portfolio. The citation archive ends July 24, 2025; the 36-month column follows a patent for at most 36 months and is shorter for recent filing cohorts. "Observed target patent portfolio" means that at least one target-assigned patent linked to an original group inventor is observed before acquisition; absence of such a portfolio does not imply that the target had no technology. Each column first reports the pooled all-acqui-hire slope. The next three rows come from one pooled model that interacts proximity with an indicator for no observed target portfolio; the final row is the formal difference between the two slopes. Columns 1–3 use Poisson pseudo-maximum likelihood. The raw direction sample contains 2,607 patents from 289 acquisitions; fixed-effect singleton and separation removal yields the column-specific samples reported below. Column 4 also uses PPML for the binary any-citation outcome, so its coefficients are log conditional-mean differences and should not be read as probability changes. Nonaffiliate citing patents share no ultimate corporate parent with the acquirer; citations that share a focal-patent inventor and citations with unresolved assignees are excluded. All specifications absorb acquisition and technology-field-by-filing-year fixed effects and cluster standard errors by acquirer. The estimates describe conditional relationships within completed acquisitions, not causal effects of technological direction. Data are from PitchBook, Revelio Labs, PatentsView, the patent-text archive, and the corporate-family mapping. Standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

	(1) Nonaffiliate citations, full	(2) Nonaffiliate citations, up to 36m	(3) Distinct outside citing families	(4) Any outside citation
All acqui-hires	-0.237*	-0.262**	-0.102	-0.038
	(0.136)	(0.110)	(0.068)	(0.035)
With observed target patent portfolio	0.006	-0.119	-0.036	-0.040
	(0.138)	(0.139)	(0.064)	(0.034)
Without observed target patent portfolio	-0.507***	-0.464**	-0.213*	-0.035
	(0.190)	(0.232)	(0.116)	(0.076)
Without minus with observed portfolio	-0.514**	-0.345	-0.177	0.005
	(0.256)	(0.316)	(0.126)	(0.082)
Acquisition fixed effects	Yes	Yes	Yes	Yes
Technology-field × filing-year fixed effects	Yes	Yes	Yes	Yes
Patents	2,304	2,153	2,304	2,304
Acquisitions	166	139	166	166
Acquirer clusters	86	73	86	86

Table A.22. Post-exit Patenting and Citation Coverage

Notes: The table starts from the common population of 8,228 transfer events from 2010 through March 2020: 1,133 identified acqui-hires and 7,095 comparison worker flows. A comparison worker flow is a reconstructed group leaving a common source firm and entering a common acquirer without acquisition of the source firm; below, it is called a comparison flow. Its contracting arrangement is unobserved. Distinct workers count people once within each event category. Worker-by-transfer-event exposure counts a worker once for every qualifying transfer event to which the worker is linked and is the denominator in the observed-path accounting. An outside-assigned post-exit patent names an original group member, is filed after the worker's sustained exit from the acquirer's corporate family, and is assigned at filing wholly outside that family. A complete 36-month citation opportunity requires filing by July 24, 2022, because the citation archive ends July 24, 2025. The two indented relatedness rows report the main relatedness table's final analytic samples; percentages use the preceding row as the denominator. Citation rows classify citing patents into four mutually exclusive categories by assignee relationship to the producing firm: affiliate, wholly nonaffiliate, jointly assigned, and unresolved. The table displays the first, second, and fourth categories; jointly assigned citations are kept separately in the underlying data. Citations sharing an inventor with the focal patent are excluded from the displayed citation counts. Counts are descriptive and do not identify the effect of acquisition on exit, patenting, or follow-on use. Data are from PitchBook, Revelio Labs, PatentsView, the patent-text archive, and the corporate-family mapping.

	(1) Acqui-hires	(2) Comparison flows
Transfer events	1,133	7,095
Distinct original workers	22,339	120,570
Worker-by-transfer-event exposures	22,364	147,299
Workers with an outside-assigned patent after exit	252	669
Events with at least one such worker	141	603
In source-team relatedness sample	39 (27.7%)	238 (39.5%)
In other-members relatedness sample	54 (38.3%)	126 (20.9%)
Outside-assigned post-exit patents	1,044	2,759
With complete 36-month citation opportunity	891	2,492
Wholly nonaffiliate citations to those patents	815	1,269
Affiliate citations to those patents	132	426
Citations with unresolved assignees	294	453

B Data Construction and Measurement

B.1 Measurement Population and Sample Lineage

The transaction universe follows Section 2.2. The unit is a unique PitchBook transaction, and acquiring subsidiaries are mapped to their ultimate parents before applying the U.S.-headquarters restriction. The target's location is unrestricted.

The resulting population contains 109,299 transactions. It combines two nonoverlapping components: 54,158 transactions involving Compustat-linked public acquirers and 55,141 additional transactions involving public and private acquirers. Of this population, 76,341 transactions have nonempty returned-text bundles and completed classifications, and 12,271 satisfy the paper's documentary rule. The corresponding 16.1 percent share is therefore conditional on returned-text coverage and completed classification.

The worker and patent datasets were constructed from earlier public-acquirer files. Table A.1 reconciles those files with the final eligibility rule. Before that rule is applied, the classification file contains 54,394 transactions, the Revelio target-linkage file contains 33,815, and their intersection contains 27,582. Applying the final rule produces 41,050 classified transactions, 25,599 target-linked transactions, and an intersection of 21,186. The documentary indicator equals one for 5,992 transactions in the earlier intersection and 4,807 in the final intersection. Classification and target linkage are overlapping samples rather than sequential filters. Each empirical design begins from the relevant intersection and imposes its own outcome-observability and support restrictions.

B.2 Document Retrieval and Input Construction

For each transaction, the retrieval query supplies the target name, acquirer name, and transaction date. It asks Perplexity's `sonar` model for up to five relevant items of public transaction coverage and requests a structured record containing the title, URL, publication, and returned text for each item. The implementation uses temperature zero and top-p equal to one. The stored record contains the query, API response, reported citations, normalized document objects, and run timestamp. The classifier receives the text returned by Sonar, which can be an article body, summary, or snippet.

The cleaning procedure first attempts to recover an article array from the stored raw response when the initial JSON parse produces no articles. It then maps alternative field names into a common schema, removes HTML and Markdown formatting and trailing source lists, and collapses whitespace. It does not truncate the returned text by characters or tokens. Within a transaction, records with nonempty URLs are deduplicated by a normalized URL and the bundle is capped at five; records without URLs are not deduplicated. The normalizer assigns a stable within-transaction document identifier and constructs one deal-level bundle containing the transaction context and up to five cleaned records. Records without usable returned text are not submitted to the classifier.

The settings in this subsection describe retrieval for the additional public- and private-acquirer transactions. The public-acquirer sample was assembled in an earlier retrieval run. Both components use the same economic decision rule described below.

B.3 Classification, Parsing, and Deal-Level Aggregation

Transactions are classified independently. The classifier receives the same transaction context and the complete cleaned bundle. The operative instruction asks Gemini to extract verbatim statements indicating that the acquisition was motivated, in whole or in part, by obtaining the target's employees, founders, or team, and directs the model not to infer motives beyond the supplied text. It instructs the model not to treat routine employee continuation, deal size, undisclosed terms, product shutdown, or generic references to talent as evidence. The complete prompt and output schema are retained with the classification code. The call uses neither a separate system instruction nor few-shot demonstrations.

Only passages satisfying the explicit documentary rule in Section 3.1 enter the empirical measure. The prompt treats an explicit acqui-hire description or an unambiguous statement that named employees, founders, or a team will join or be hired by the acquirer as qualifying evidence. Two weaker diagnostic fields and the code's broader diagnostic flag are unused in every reported analysis. As the main text explains, this rule establishes a stated personnel motive, not its primacy.

For each extracted passage, the prompt requires the evidence category, verbatim quotation, document identifier, people described, relationship to the acquirer, and named target and acquirer. A schema validator checks that these fields are present and that the category is admissible. The parser removes formatting fences, isolates the JSON object, and validates it against the schema; malformed outputs receive one schema-only reformatting call. The classifier is `gemini-3-flash-preview`, run at temperature zero with top-p equal to 0.95 and top-k equal to 64. Changes in output-token allowance and truncated-JSON recovery do not alter the prompt or documentary rule.

For a successfully parsed record, deal-level aggregation is deterministic. The documentary indicator equals one if and only if at least one validated item carries the explicit evidence code. The aggregation routine does not reclassify the passage. The full evidence record retains the raw response, parsed evidence, document identifiers, success status, model name, and timestamp. The analysis file retains the indicator, evidence counts, model identifier, and up to three supporting quotations. Together, these records permit reconstruction of the deal-level aggregation step.

Calls with no usable input, an API failure, or an unrepaired parse error are marked unsuccessful rather than treated as documentary negatives. Both consolidated label files used here contain only successful records, so the completed-classification denominator excludes unsuccessful and skipped cases.

B.4 Validation Scope

Section 3.2 evaluates the documentary indicator using transaction characteristics and worker histories that do not enter either language-model prompt. These exercises assess construct validity: they ask whether the measure selects the smaller, younger, technology-oriented targets and high-skill worker cohorts expected of team-motivated acquisitions. The classifier also retains the supporting passage and document identifier for each positive item, permitting direct comparison between the deal-level label and the text supplied to the classifier.

These checks do not recover the transaction's latent primary motive. The paper therefore interprets the indicator as a documentary measure of an explicitly stated personnel motive, not as a direct measure of the relative value of people and technology in the transaction.

B.5 PitchBook–Revelio Linkage and Acqui-hire Cohorts

Worker measurement uses the Compustat-linked public-acquirer sample. The target matcher first searches Revelio firms already recorded within the acquirer's corporate family and then searches stand-alone Revelio firms. Within each search set, it applies the same ordered match sequence: LinkedIn company URL, normalized exact company name, and LinkTransformer. Model-based matches must have a similarity score of at least 0.95. The stand-alone model match also uses country and state information when available; the within-family search does not impose that geographic filter. Transactions without an accepted target link do not enter the worker-linked samples.

The cleaned employment histories are right-censored on September 30, 2025. The transition code restricts both target and candidate next-employer spells to records for which the cleaned-history stale-profile flag equals zero.

For an identified acqui-hire e closing at date t_e, the construction first selects target employment spells overlapping $[t_e - 12 \text{ months}, t_e + 6 \text{ months}]$. For each worker, it finds the earliest qualifying spell at a different employer that begins no more than 30 days before the target spell ends and no later than 24 months after closing. A recorded move enters the acquirer-family cohort when the new employer is the acquirer or one of its controlled subsidiaries and the gap from the target spell is no more than 90 days. The 30-day allowance accommodates imprecision in reported spell boundaries and permits limited overlap between the target and next-employer spells.

Revelio sometimes already represents the target as a subsidiary of the acquirer, so a legal ownership change need not produce a new employer identifier. In that case, the cohort also includes a worker with no qualifying different-employer spell through 24 months after closing whose target spell continues for at least 24 months after closing. Let E_e denote the union of these recorded moves and seamless subsidiary continuations, and let $M_e \subseteq E_e$ contain only workers with a recorded employer change. Measures requiring a move date or prior joint mobility use M_e. Neither set is a census of the target workforce. In particular, the current data do not provide a consistently measured pre-transaction roster denominator and therefore cannot estimate the fraction of that roster that transfers.

B.6 Comparison Worker Flows and Prior-Coworker Reconstruction

To construct a comparison worker flow, I begin with workers entering the same employer within an observed acquirer's corporate family, within six months on either side of the associated transaction date. I refer below to a retained comparison worker flow as a comparison flow. The procedure removes workers employed by the focal target, assigns each remaining worker to the most recent prior employer different from that acquirer-family employer whose spell ends no later than 30 days after entry, and calls that prior employer the source firm. It removes workers whose source firm is already in the acquirer's corporate family. Ties in the source-firm assignment are resolved deterministically, and transactions with unresolved target links are removed before clustering. Each row retains the employer-entity identifier and is deduplicated at the acquirer-family-employer–source-firm–worker level. A raw candidate contains at least three workers from the same source firm who enter the same employer within the acquirer's corporate family within a maximum span of 60 days. Within an acquirer-family-employer–source-firm pair, the algorithm retains every maximal set whose earliest and latest arrivals are no more than 60 days apart; overlapping sets can share workers. Clustering does not reset at calendar-year boundaries, and the event date is the median arrival date among the movers. These candidates are anchored to observed acquirers, their employer entities, and the dates of observed acqui-hires, not an economy-wide sample of direct hiring.

To exclude coincident arrivals that do not represent a pre-existing group, I measure whether candidate members worked together before the transfer. Within a raw candidate, a pair is connected if the workers overlap at the source firm for at least 180 days in the same six-digit O*NET occupation or if they satisfy the prior-joint-mobility definition in equation (1). The candidate's connected-pair density is compared with 500 equal-size placebo groups sampled without replacement from nonmoving source-firm workers whose employment intersects the 24 months before the event date. The placebo universe excludes the focal movers and is capped deterministically at the larger of 2,000 workers or ten times the candidate's mover count. If S_c is the observed combined density and S_{cb}^P is the source-firm workplace density in placebo draw b, the screening statistic is

$$p_c = \frac{1 + \sum_{b=1}^{500} \mathbf{1}\{S_{cb}^P \geq S_c\}}{501}. \tag{24}$$

The random-number generator is seeded by the candidate identifier. Candidates with $p_c < 0.05$ are retained. The placebo routine does not recompute prior joint mobility for placebo pairs: it compares the observed combined density with the source-firm workplace placebo distribution.

It is therefore a screening rule, not an exact permutation test of the complete combined network.

Let M_c denote the movers in retained candidate c. For each mover, the prior-coworker reconstruction first selects the mover's source-firm spell that ends within 365 days on either side of arrival and is closest to that date. It then considers coworkers at the same source firm whose employment overlaps the selected spell by at least 180 days and whose six-digit O*NET occupation is the same. If more than 200 coworkers qualify, the procedure retains the 200 with the longest overlap, breaking ties by worker identifier, and then takes the union across movers. With $\mathcal{P}_i$ denoting the retained peers of mover i, the observed prior-coworker proxy is

$$T_c = M_c \cup \bigcup_{i \in M_c} \mathcal{P}_i. \tag{25}$$

Thus $M_c \subseteq T_c$ by construction. The cap is applied separately to each mover, so the union may contain more than 200 nonmovers. This set is an observed occupational-peer proxy, not a census of the productive team.

I retain candidates that pass the screen as comparison flows. They document a group from a common source firm joining a common acquirer without acquisition of that source firm. The workplace screen supports a pre-existing-group interpretation, but the worker histories do not reveal whether offers were coordinated or team-contingent. The paper therefore uses these observations as a comparison for realized group assembly through hiring, not as evidence about the underlying contracts.

B.7 Patent-Text Measures of Technological Proximity

The patent-text archive draws claims, titles, and abstracts from PatentsView's granted-patent long-text tables and keys them to the identifiers used elsewhere in the paper (United States Patent and Trademark Office, 2025a). The baseline text representation uses the available claims. I encode each document with the pretrained GTE-large text-embedding model of Li et al. (2023), using the fixed `thenlper/gte-large` model revision `4bef63f39fcc5e2d6b0aae83089f307af4970164`. The resulting patent vector has 1,024 dimensions. Claims longer than the model's context window are divided into 448-token segments with 32-token overlap. Segment vectors are normalized and combined using effective token counts as weights; the resulting document vector is normalized again. Titles and abstracts, encoded with the same model, provide an alternative text representation.

Let z_p denote the normalized vector for patent p, and let $\mathcal{P}^{\text{text}} \subseteq \mathcal{P}$ be the patents in an eligible portfolio with an observed vector. For any nonempty $\mathcal{P}^{\text{text}}$, I form the equal-patent-weight representation

$$\bar{z}(\mathcal{P}) = \frac{1}{|\mathcal{P}^{\text{text}}|} \sum_{p \in \mathcal{P}^{\text{text}}} z_p. \tag{26}$$

Technological proximity between two portfolios is the cosine of their representations. Thus, patents with more claims or longer claims do not receive greater portfolio weight. A portfolio with $\mathcal{P}^{\text{text}} = \varnothing$ remains missing; it is never coded as zero similarity. The resulting claims-text archive contains vectors for 620,479 patents.

The baseline construction assigns patents to event windows by filing date and uses their eventual claims content. Historical versions of claims are not available, so the measure should not be interpreted as the exact text observed by the parties on the event date. For the pre-transaction pairing exercise, a timing check restricts the portfolios to patents granted by transaction completion. Complete-follow-up restrictions address the separate truncation created by the grant requirement for recent filing cohorts.

The claims representation distinguishes technological relationships defined without reference to acqui-hire status. Mean cosine similarity is 0.872 for patent pairs in the same CPC subclass, compared with 0.838 for same-filing-year pairs in different subclasses. It is 0.922 for pairs sharing an inventor, compared with 0.849 for CPC-section- and filing-year-matched pairs without a shared inventor, and 0.909 for citation-linked pairs, compared with 0.850 for CPC-section- and filing-year-matched unlinked pairs. These comparisons validate the ordering of the similarity measure; they do not impose an economic threshold for whether two portfolios are related.

Pre-transaction acquirer–target similarity. For acquisition d, let $\mathcal{P}^{-}_{ad}$ and $\mathcal{P}^{-}_{td}$ contain distinct U.S. utility patents assigned to the direct acquirer and target, respectively, and filed during the five years before completion. The claims-text measure in Table 1 is

$$\text{ClaimsSimilarity}_d = \cos\Big(\bar{z}(\mathcal{P}^{-}_{ad}), \bar{z}(\mathcal{P}^{-}_{td})\Big). \tag{27}$$

The baseline support rule requires at least one embedded claims patent on each side. Of the 2,826 valid patent-active pairs, 2,745 satisfy this rule before fixed-effect estimation; the corresponding regression retains 2,525 transactions and 750 ultimate-parent acquirer clusters. Appendix Table A.3 varies the text representation, coverage requirement, and timing restriction. The scale controls in column 4 of Table 1 are the logarithms of one plus the number of patents and the number of distinct cited prior-art patents on each side of the transaction.

Post-entry continuity and movement toward the acquirer. For transfer event e, let $\mathcal{P}^{T}_{e}$ contain patents that name an incoming worker, are assigned before transfer to the target for an acqui-hire or the source firm for a comparison flow, and are filed during the five years before the worker enters the acquirer. Let $\mathcal{P}^{A}_{e}$ contain patents assigned to the acquirer or its parent during the five years before the cohort's earliest entry. Finally, let $\mathcal{P}^{P}_{eh}$ contain patents that name an incoming worker, are assigned to the acquirer's corporate family, and are filed within h months after the named worker enters. Patent links are deduplicated within an event; a patent that is legitimately linked to several transfer events can enter each event's portfolio.

I define continuity with the incoming group's prior work, pre-entry alignment with the acquirer, post-entry proximity to the acquirer, and movement toward the acquirer as

$$\mathcal{C}_{eh} = \cos\Big(\bar{z}(\mathcal{P}^{P}_{eh}), \bar{z}(\mathcal{P}^{T}_{e})\Big), \tag{28}$$

$$\mathcal{B}_{e} = \cos\Big(\bar{z}(\mathcal{P}^{T}_{e}), \bar{z}(\mathcal{P}^{A}_{e})\Big), \tag{29}$$

$$\mathcal{A}_{eh} = \cos\Big(\bar{z}(\mathcal{P}^{P}_{eh}), \bar{z}(\mathcal{P}^{A}_{e})\Big), \qquad \mathcal{M}_{eh} = \mathcal{A}_{eh} - \mathcal{B}_{e}. \tag{30}$$

The secondary two-sided measure is $\min\{\mathcal{C}_{eh}, \mathcal{A}_{eh}\}$, which requires both components to be observed. The main regressions compare acqui-hires with comparison flows in the Section 4.2 common-support population. They use the event matching weights, absorb acquirer and entry-year fixed effects and two-digit industry fixed effects defined using the target for an acqui-hire and the source firm for a comparison flow, and cluster standard errors by acquirer.

The estimation samples contain only events with every text-supported portfolio required for the corresponding outcome. Appendix Tables A.7 and A.8 report the decomposition and the principal coverage, follow-up, and portfolio-construction checks. These measures describe technological direction among patent-observed events; they do not identify the causal effect of acquisition on that direction.

B.8 External Follow-on and Post-exit Technological Relatedness

The external-follow-on analysis begins from the patent lineage described in Section 8.2. An inside-produced patent names at least one original incoming inventor and is assigned at filing to the acquirer's corporate family. A post-exit patent names an original incoming worker, is filed after the worker's sustained exit from that family, and is assigned at filing wholly outside it. Corporate families are defined using ultimate-parent links observed in the company mapping.

For every citing patent, I classify its assignee relationship to the focal firm as affiliate only, wholly nonaffiliate, joint, or unresolved. The main external-follow-on outcome counts wholly nonaffiliate citing patents and excludes citations for which a focal-patent inventor also appears on the citing patent. Jointly assigned and unresolved citations are kept as separate categories and do not enter the baseline external count. This construction uses patent citations as a measure of observed follow-on use; it does not assign a welfare value to a citation.

The direction measure in equation (22) compares each inside-produced patent with the acquirer's pre-transaction claims-text portfolio. Subtracting the incoming group's pre-transaction similarity to that same acquirer portfolio would change the regressor only by an acquisition-level constant, which is absorbed by the acquisition fixed effect. The specification therefore retains acquisitions for which the target has no observed pre-transaction patent portfolio. Appendix Table A.21 reports the corresponding split without treating the absence of a target portfolio as evidence that the target had no technology.

For post-exit relatedness, each claims-text portfolio is the normalized, equal-patent-weighted mean defined in equation (26). The pre-transfer prior-work reference contains patents linked to original group members and assigned before transfer to the target for an acqui-hire or the source firm for a comparison flow. The other-members reference contains patents filed inside the acquirer by other original group members strictly before the focal worker's exit. Placebo groups are selected from the same two-digit industry of the target for an acqui-hire or the source firm for a comparison flow and within two transfer years, ordered by the distance in the size of the original group's pre-transfer target or source-firm patent portfolio. They cannot share an inventor, patent, or acquirer with the focal group. Both focal and placebo reference portfolios are cut off at the same focal-worker exit date. Appendix Table A.22 reports the resulting sample and citation coverage.

Finally, the accounting in equation (23) uses one eligible inside-patent universe for Q_h^{in}, δ_h^{in}, and D_h^{in}. The common 36-month citation opportunity removes differences in forward-citation exposure across filing cohorts. The identity is an accounting identity among realized outcomes. It does not imply that the organizational form is as-if randomly assigned or that the same

patents would exist under direct hiring or target continuation.

C Model Derivations and Proofs

C.1 Proofs and Derivations for the Benchmark Framework

C.1.1 Superadditivity of the relationship stock

Proof of Lemma 1. Because M_1 and M_2 are disjoint, every unordered pair in $M_1 \cup M_2$ is either internal to M_1, internal to M_2, or a cross-pair with one endpoint in each coalition. Therefore,

$$L_i(M_1 \cup M_2) = \frac{2}{n_i - 1} \sum_{\substack{k<\ell \\ k,\ell \in M_1}} a_{ik\ell} h_{ik} h_{i\ell} + \frac{2}{n_i - 1} \sum_{\substack{k<\ell \\ k,\ell \in M_2}} a_{ik\ell} h_{ik} h_{i\ell} \tag{31}$$

$$+ \frac{2}{n_i - 1} \sum_{k \in M_1} \sum_{\ell \in M_2} a_{ik\ell} h_{ik} h_{i\ell}. \tag{32}$$

The first two terms are $L_i(M_1)$ and $L_i(M_2)$. Subtracting them leaves the cross-pair term in (32). Since $n_i \geq 2$, $a_{ik\ell} \geq 0$, and $h_{ik} > 0$, every term in the remaining sum is nonnegative. The inequality is strict if and only if at least one cross-pair has $a_{ik\ell} > 0$. □

C.1.2 The two-worker sequential-hiring example

Let $P(M)$ denote the acquirer's gross stopping value after the workers in M have signed, excluding packages already promised. In the example in Section 5.2,

$$P(\varnothing) = 0, \qquad P(\{1\}) = P(\{2\}) = u, \qquad P(\{1,2\}) = V, \tag{33}$$

and each worker's reservation value is $r > u$. Suppose one worker has signed. Adding the second worker creates bilateral surplus $V - u - r$. Generalized Nash bargaining with worker share β gives the second worker $r + \beta(V - u - r)$ and leaves the acquirer with

$$W_1 = u + (1 - \beta)(V - u - r). \tag{34}$$

When bargaining with the first worker, the bilateral surplus is therefore $W_1 - r$. The first agreement is feasible if and only if

$$u + (1 - \beta)(V - u - r) \geq r, \quad \text{or equivalently} \quad V \geq 2r + \frac{\beta}{1 - \beta}(r - u). \tag{35}$$

The workers are symmetric, so allowing the acquirer to choose which worker to approach first does not change this threshold. A team-contingent package is feasible if and only if $V \geq 2r$,

and it produces strictly positive joint surplus when $V > 2r$. Hence

$$2r < V < 2r + \frac{\beta}{1-\beta}(r-u)$$

is precisely the region in which coordinated hiring creates positive surplus but sequential hiring does not begin. At $V = 2r$, coordinated full-team surplus is zero. At $V = 2r+\beta(r-u)/(1-\beta)$, the first bilateral bargain has zero surplus even though full-team surplus is positive. Behavior at either boundary therefore depends on the tie-breaking convention; the strict inequalities in (4) avoid those knife-edges.

C.1.3 Recursive bilateral bargaining for a general team

Fix acquirer j, coordination environment q, and a state $M \subseteq \mathcal{T}_i$ containing the workers already under binding contracts. Past payments are sunk. If the acquirer stops at M, its gross continuation value, excluding those sunk payments, is

$$P_{ij}^H(M;q) = V_{ij}(M, X_i^H; q). \tag{36}$$

Let $r_{ik}^j(M;q)$ be worker k's monetary reservation value if bargaining at state M does not produce an agreement and the acquirer stops at M. All stopping and reservation values are finite. Let $W_{ij}^H(M;q)$ be the acquirer's value immediately before it stops or selects the next worker. The terminal condition is

$$W_{ij}^H(\mathcal{T}_i;q) = P_{ij}^H(\mathcal{T}_i;q). \tag{37}$$

For $k \notin M$, define the bilateral surplus from adding worker k as

$$\Delta_{ijk}^H(M;q) = W_{ij}^H(M \cup \{k\};q) - P_{ij}^H(M;q) - r_{ik}^j(M;q). \tag{38}$$

If this surplus is nonnegative, generalized Nash bargaining with worker share β gives the worker $r_{ik}^j(M;q) + \beta\Delta_{ijk}^H(M;q)$ and gives the acquirer $P_{ij}^H(M;q) + (1-\beta)\Delta_{ijk}^H(M;q)$. The continuation recursion is therefore

$$W_{ij}^H(M;q) = \max\Bigg\{ P_{ij}^H(M;q), \max_{\substack{k\notin M \\ \Delta_{ijk}^H(M;q)\geq 0}} \Big[P_{ij}^H(M;q) + (1-\beta)\Delta_{ijk}^H(M;q)\Big] \Bigg\}, \tag{39}$$

where the inner maximum is omitted when no worker yields nonnegative bilateral surplus. A deterministic tie-breaking rule selects one action whenever values coincide.

Proposition 2 (Recursive bilateral-bargaining solution)**.** *Given the finite team, stopping values, reservation values, bargaining weight, and deterministic tie-breaking rule, Equation* (39) *determines a unique continuation value and selected action at every state. Starting from* $M = \varnothing$*, the selected path terminates at a well-defined coalition* $M_{ij}^H(q) \subseteq \mathcal{T}_i$*.*

Proof. Proceed by backward induction on the number of workers not yet contracted. The value at $M = \mathcal{T}_i$ is fixed by (37). Suppose the continuation value is known at every state with fewer workers remaining than at M. Equation (38) then determines the bilateral surplus from each possible next worker. Nash bargaining determines the acquirer's value for each nonnegative-surplus negotiation, and the acquirer compares those values with stopping at $P_{ij}^H(M; q)$. Finiteness guarantees that a maximum exists, and the tie-breaking rule selects a unique action. Induction determines the value and action at every state. Each agreement adds one worker, so the selected path reaches a terminal coalition in finitely many steps. □

C.1.4 Assembly and control of target-controlled technology

Let $M = M_{ij}^H(q)$. Adding and subtracting $V_{ij}(M, X_i; q)$ yields (5). Substituting (3), the coalition-assembly margin is

$$\mathcal{A}_{ij}(M; q) = \phi_{ij} \left[H_i(\mathcal{T}_i) - H_i(M) \right] \tag{40}$$

$$+ q\theta_i g_{ij}(X_i) \left[L_i(\mathcal{T}_i) - L_i(M) \right], \tag{41}$$

and the control-of-technology margin is

$$\mathcal{X}_{ij}(M; q) = q\theta_i \left[g_{ij}(X_i) - g_{ij}(X_i^H) \right] L_i(M). \tag{42}$$

For $M \subseteq \mathcal{T}_i$, both bracketed differences in (41) are nonnegative. If $M \subsetneq \mathcal{T}_i$, then

$$H_i(\mathcal{T}_i) - H_i(M) = \sum_{k \in \mathcal{T}_i \setminus M} h_{ik} > 0,$$

so $\mathcal{A}_{ij}(M; q) > 0$. Weak monotonicity of g_{ij} implies $\mathcal{X}_{ij}(M; q) \geq 0$; the inequality is strict exactly when $L_i(M) > 0$ and $g_{ij}(X_i) > g_{ij}(X_i^H)$, because $q > 0$ and $\theta_i > 0$.

The decomposition in (5) first changes the coalition while holding technology at X_i, and then changes technology while holding the coalition at M. Reversing the order gives

$$V_{ij}(\mathcal{T}_i, X_i; q) - V_{ij}(M, X_i^H; q)$$

$$= \underbrace{V_{ij}(\mathcal{T}_i, X_i; q) - V_{ij}(\mathcal{T}_i, X_i^H; q)}_{\text{control margin at the complete team}} + \underbrace{V_{ij}(\mathcal{T}_i, X_i^H; q) - V_{ij}(M, X_i^H; q)}_{\text{assembly margin at direct-hiring technology}} . \tag{43}$$

The amount assigned differently across the two orderings is

$$\begin{aligned} & \mathcal{A}_{ij}(M; q) - \left[V_{ij}(\mathcal{T}_i, X_i^H; q) - V_{ij}(M, X_i^H; q)\right] \\ = & \left[V_{ij}(\mathcal{T}_i, X_i; q) - V_{ij}(\mathcal{T}_i, X_i^H; q)\right] - \mathcal{X}_{ij}(M; q) \\ = & q\theta_i \left[g_{ij}(X_i) - g_{ij}(X_i^H)\right] \left[L_i(\mathcal{T}_i) - L_i(M)\right] \geq 0. \end{aligned} \tag{44}$$

Thus the total productive advantage and the condition under which that total is zero are invariant to the ordering. The magnitudes, and potentially the zero-versus-positive status, of the individual components are not.

If an expanded arrangement assembles $\mathcal{T}_i$, places the same enforceable decision rights represented by X_i under the acquirer's control, induces the same retention and project direction, and incurs the same real resource costs, then its acquirer-side productive value equals the value under acquisition by direct substitution into (3). This equality concerns the acquirer-side production possibilities. It does not equate the independent target's continuation, total innovation, transaction price, surplus division, or welfare.

C.1.5 Properties of the acquisition-market closure

Write $\bar{b}_{i(1)}$ and $\bar{b}_{i(2)}$ for the highest and second-highest bid capacities. When there is only one candidate acquirer, the second-highest term below is omitted. Conditional on a feasible acquisition, define the payment floor and reduced-form price by

$$F_i = \max\{O_i^S, \bar{b}_{i(2)}\}, \qquad P_i^* = \alpha_P \bar{b}_{i(1)} + (1 - \alpha_P) F_i, \qquad \alpha_P \in [0, 1]. \tag{45}$$

The target side requires consideration of at least O_i^S, while candidate acquirer j is willing to pay at most $\bar{b}_{ij}$. If $\bar{b}_{i(1)} < O_i^S$, no target consideration can satisfy both sides. If $\bar{b}_{i(1)} \geq O_i^S$, the highest-capacity acquirer can make a mutually acceptable payment. Moreover,

$$\bar{b}_{i(1)} \geq O_i^S \quad \text{and} \quad \bar{b}_{i(1)} \geq \bar{b}_{i(2)}$$

imply

$$F_i \leq \bar{b}_{i(1)}. \tag{46}$$

Equation (45) then gives

$$F_i \leq P_i^* \leq \bar{b}_{i(1)}. \tag{47}$$

On a local region in which $O_i^S < \bar{b}_{i(2)} < \bar{b}_{i(1)}$, candidate-acquirer ranks are fixed, and all other objects are fixed,

$$\frac{\partial P_i^*}{\partial \bar{b}_{i(2)}} = 1 - \alpha_P, \qquad \frac{\partial \left(\bar{b}_{i(1)} - P_i^*\right)}{\partial \bar{b}_{i(2)}} = -(1 - \alpha_P). \tag{48}$$

Thus a stronger second capacity affects the division of surplus but not acquisition feasibility when the highest capacity and target outside value are held fixed. At $\bar{b}_{i(2)} = O_i^S$ or at an order-statistic tie, the analogous statements use one-sided derivatives because the identity of the binding floor or ranked acquirer can change.

C.1.6 Coordination productivity and acquisition feasibility

For the fixed comparison coalition $\widetilde{M}_{ij}^H$, direct differentiation of (3) gives (9). To establish the properties of the acquisition continuation value under q-invariant implementation costs, define

$$\eta_{ij}^A(R) = \phi_{ij} H_i(R) - K_{ij}^A(R), \qquad s_{ij}(R) = \theta_i g_{ij}(X_i) L_i(R) \geq 0. \tag{49}$$

Then

$$\Pi_{ij}^A(q) = \max_{R \subseteq \mathcal{T}_i} \{\eta_{ij}^A(R) + q s_{ij}(R)\}. \tag{50}$$

Each term in braces is affine and weakly increasing in q. Their finite pointwise maximum is therefore nondecreasing, convex, and piecewise linear. If the optimizer is unique at q, the slope of the active affine function is $s_{ij}(R_{ij}^{A,*}(q))$, so

$$\frac{\mathrm{d}\Pi_{ij}^A(q)}{\mathrm{d}q} = \theta_i g_{ij}(X_i) L_i\left(R_{ij}^{A,*}(q)\right). \tag{51}$$

If several coalitions are optimal at an interior q, the right derivative is the largest and the left derivative is the smallest value of $s_{ij}(R)$ among those coalitions. At a boundary of the maintained interval, only the relevant one-sided derivative is defined.

In the fixed-coalition benchmark described in the text, acquisition retains $\mathcal{T}_i$, O_i^S is q-invariant, and the aggregate no-acquisition payoff of the acquirer and workers changes with direct-hiring production according to

$$\frac{\mathrm{d}}{\mathrm{d}q}\left[O_{ij}^B(q) + \sum_{k \in \mathcal{T}_i} \bar{r}_{ik}^A(q)\right] = \theta_i g_{ij}(X_i^H) L_i(\widetilde{M}_{ij}^H). \tag{52}$$

Differentiating (7) therefore gives

$$\frac{\mathrm{d}}{\mathrm{d}q}\left[\bar{b}_{ij}(q) - O_i^S\right] = \theta_i\left[g_{ij}(X_i)L_i(\mathcal{T}_i) - g_{ij}(X_i^H)L_i(\widetilde{M}_{ij}^H)\right] \geq 0, \tag{53}$$

where K_{ij}^A and C_{ij}^A are also q-invariant. Bargaining transfers between the acquirer and workers do not enter because their payoffs are aggregated. The inequality is strict if $L_i(\mathcal{T}_i) > L_i(\widetilde{M}_{ij}^H)$, or if $g_{ij}(X_i) > g_{ij}(X_i^H)$ and $L_i(\widetilde{M}_{ij}^H) > 0$.

Proof of Corollary 1. Because $O_i^S(q)$ does not vary across candidate acquirers,

$$\Delta_i(q) = \max_{j \in \mathcal{J}_i^A}\left\{\bar{b}_{ij}(q) - O_i^S(q)\right\}. \tag{54}$$

For any $q_2 > q_1$, Equation (7) gives

$$\begin{aligned}
&\left[\bar{b}_{ij}(q_2) - O_i^S(q_2)\right] - \left[\bar{b}_{ij}(q_1) - O_i^S(q_1)\right] \\
&= \Pi_{ij}^A(q_2) - \Pi_{ij}^A(q_1) - \left[O_{ij}^B(q_2) - O_{ij}^B(q_1)\right] - \sum_{k \in \mathcal{T}_i}\left[\bar{r}_{ik}^A(q_2) - \bar{r}_{ik}^A(q_1)\right] - \left[O_i^S(q_2) - O_i^S(q_1)\right] \geq 0,
\end{aligned}$$

where the transaction cost C_{ij}^A cancels and the inequality follows from (11). Thus every function inside the maximum in (54) is weakly increasing. The candidate set is fixed and finite, so their pointwise maximum is weakly increasing. The acquisition-incidence statement follows from the rule $\Delta_i(q) \geq 0$. □

C.1.7 Coordination productivity and retained relationships

Proof of Proposition 1. For fixed i, j, μ, write

$$\eta^\mu(R) = \phi_{ij}H_i(R) - K_{ij}^{\mu,\mathrm{post}}(R), \qquad s^\mu(R) = \theta_i g_{ij}(x_{ij}^\mu)L_i(R).$$

Optimality of R_1 at q_1 and R_2 at q_2 implies

$$\begin{aligned}
\eta^\mu(R_1) + q_1 s^\mu(R_1) &\geq \eta^\mu(R_2) + q_1 s^\mu(R_2), \\
\eta^\mu(R_2) + q_2 s^\mu(R_2) &\geq \eta^\mu(R_1) + q_2 s^\mu(R_1).
\end{aligned}$$

Adding these inequalities and rearranging gives

$$(q_2 - q_1)\theta_i g_{ij}(x_{ij}^\mu)\left[L_i(R_2) - L_i(R_1)\right] \geq 0. \tag{55}$$

Because $q_2 > q_1$, $\theta_i > 0$, and $g_{ij}(x_{ij}^{\mu}) > 0$, it follows that $L_i(R_2) \geq L_i(R_1)$. This argument orders the relationship stock, not cardinality. For nonnested coalitions, it is possible that $L_i(R_2) > L_i(R_1)$ while $|R_2| < |R_1|$. A headcount result therefore requires additional restrictions. □

C.1.8 Stochastic retention policies

The deterministic implementation problem in (12) is a benchmark. Suppose instead that the employer chooses a policy e from a q-independent set $\mathcal{E}$, and the policy induces a distribution over retained coalitions. If the employer's expected payoff can be written as

$$\eta^{\mu}(e) + q\theta_i g_{ij}(x_{ij}^{\mu})\mathbb{E}[L_i(R) \mid e]\,, \tag{56}$$

where $\eta^{\mu}(e)$ is independent of q, the revealed-preference argument above applies with $\mathbb{E}[L_i(R) \mid e]$ in place of $L_i(R)$. Expected retained relationship stock under an optimal policy is then weakly increasing in q. In a more general stochastic problem, increasing differences between q and an ordered retention policy are sufficient only when higher policies also generate weakly higher expected relationship stock.

D Supplementary Empirical Analysis

D.1 Additional Organizational-Form Evidence

Appendix Table A.9 asks whether the industry-level result depends on a particular estimator, weighting scheme, or fixed-effect structure. The specifications do not all estimate the same object: the PPML columns are log conditional-mean gradients, whereas the linear columns are expressed in counts or log counts. I therefore use these columns to assess whether the negative association is stable across specifications, not to compare coefficient magnitudes. Across these alternatives, acqui-hire activity shifts away from industries with greater predetermined exposure to nonteleworkable work during 2020–2022.

Appendix Table A.10 examines whether this pattern is confined to acquirers that appear financially constrained before the pandemic. The split-sample estimates remain negative among acquirers classified as less constrained by either the size–age or Whited–Wu index. Because these regressions use additive acquirer and year effects rather than the baseline acquirer-by-year effects, and because the table does not test equality across subgroups, I treat it as a diagnostic for a purely financing-based interpretation rather than as evidence of heterogeneous effects.

D.2 Additional Evidence from the Coordination Disruption

D.2.1 Task Exposure to the Coordination Disruption

Let T_o denote Dingel and Neiman's teleworkability measure for occupation o. Their original classification uses 2010 SOC codes. I map it to 2018 SOC codes and, when several 2010 codes map to the same 2018 code, take the unweighted mean of their classifications. Consequently, T_o can be fractional for consolidated occupations. Let I_o be the average O*NET importance rating across the following nine Generalized Work Activities: interpreting information for others; communicating with supervisors, peers, or subordinates; communicating with people outside the organization; establishing and maintaining interpersonal relationships; coordinating the work and activities of others; developing and building teams; training and teaching others; guiding, directing, and motivating subordinates; and coaching and developing others. I form the occupation-level product

$$R_o = (1 - T_o)I_o \tag{57}$$

and standardize R_o across the 768 six-digit 2018 SOC occupations for which both inputs are observed.

For each transfer, I average this standardized occupation-level product across the high-skill movers. I then standardize the event-level average on the analysis population specified for each design. The within-acqui-hire post-transfer task-exposure index is standardized over the 835-event matched analysis pool before fixed-effect singleton removal. Because the baseline within-acqui-hire index uses the occupation attached to each worker's spell at the acquirer after joining through the acqui-hire, some contributing spells begin after April 1, 2020. In a timing check, I reconstruct the index using only high-skill workers whose acquirer spell begins before that date and standardize the resulting event means over the 830 matched events with at least one eligible worker. The principal acqui-hire–comparison-flow pre-transfer task-exposure index is standardized over the eligible event universe before the group-size floor, trimming, and common-support restriction. The product is high only when an occupation is both difficult to perform remotely and intensive in the selected interpersonal activities. It is therefore a predicted exposure to disruption of interaction-dependent work, not a direct measure of realized face-to-face interaction, pre-existing team relationships, or team complementarity.

D.2.2 Timing, Measurement, and Weighting

Appendix Table A.14 collects the sensitivity analyses that change one element of the exposure measure or sample at a time. Within acqui-hires, restricting post-transfer occupations to

spells that begin before April 2020 and excluding first-quarter 2020 transactions produce estimates close to the baseline. The pre-transfer occupation index produces smaller, imprecise estimates, but it measures workers' prior tasks rather than the tasks they perform at the acquirer. In the acqui-hire-versus-comparison-flow design, the table also excludes first-quarter 2020 transfers and equalizes the two event categories' total weight within each support cell. This reweighting does not balance target size or age against source-firm size or age; the target–source-firm-matched analysis addresses that separate selection margin.

D.2.3 Retention Horizons and Pre-period Diagnostics

Appendix Table A.11 reports the 24-, 36-, and 48-month estimates, while Appendix Table A.13 and Appendix Figure A.1 report pre-pandemic placebo and cohort-slope tests. The graded retention outcomes are strongest at 24 months, with no corresponding pre-period pattern. Complete retention yields one significant placebo estimate and a nonzero cohort slope, so I report it as a sensitivity outcome rather than as support for the retention claim.

D.2.4 Comparison Worker Flows

Appendix Table A.15 and Appendix Figure A.2 report the horizon and placebo results for the acqui-hire-versus-comparison-flow comparison. The coefficients are differential disruption-by-pre-transfer-task-exposure gradients, not average acquisition effects. The principal comparison restricts support using characteristics of the group reaching the acquirer. Appendix Table A.16 instead imposes common support on target size, age, and industry for acqui-hires and source-firm size, age, and industry for comparison flows, along with transfer year and mover composition. The samples address different observable selection margins; neither makes the acquisition-versus-flow category as-if random.

D.2.5 Career Progression and Worker-Level Outcomes

Appendix Table A.17 reports the within-acqui-hire career-progression gradient and its pre-pandemic placebos. Appendix Table A.12 reports worker-level retention and seniority estimates for both empirical designs. Retention is estimated by linear probability models and seniority by PPML. Because the controls and weights differ from the team-level specifications, these results provide corroborating evidence rather than an algebraic restatement of the team estimates.

D.2.6 Innovative Production

Appendix Table A.18 varies the disruption window, the exposure denominator, the industry-by-cohort fixed effects, and the occupations used to construct task exposure for production within the arriving inventor group. It also summarizes the pre-disruption diagnostics for that design and for the within-group mixed-versus-solo comparison. The exercise asks whether the production response is specific to one construction of exposure or appears during analogous pre-pandemic windows; it does not turn the comparison flows into randomly assigned alternatives to acquisition.

D.2.7 Worker Reallocation

Appendix Figures A.3–A.5 compare the 2020–2022 disruption estimates with four pre-pandemic pseudo-shock windows for exit, co-departure, and entrepreneurial reallocation. Appendix Table A.19 varies the number of group members required to enter entrepreneurially oriented work and the timing rule for their outside employment spells. Appendix Table A.20 separately reports conditional level differences in entrepreneurial reallocation; those estimates are not differential responses to the disruption.

D.2.8 External Follow-on and Post-exit Patenting

Appendix Table A.21 reports the formal comparison between acquisitions with and without an observed target-assigned pre-transaction patent portfolio. The full-history nonaffiliate-citation slope is more negative in the latter group, but that difference is imprecise under a 36-month citation cap and for the number of distinct outside citing families. The extensive-margin outcome is also imprecise. These patterns are why the main text describes a relationship in external citation intensity rather than a change in the probability or breadth of outside use.

Appendix Table A.22 reports the support for the post-exit analysis. Patenting after exit is uncommon, and the main relatedness table conditions further on an observed focal and placebo reference portfolio. The table therefore documents a technological link among selected patenting exiters; it is not a population estimate of the share of workers carrying knowledge outside. Citations from joint or unresolved assignees remain separate throughout and are not reclassified as nonaffiliate follow-on.

D.3 Local Industry-Market Outcomes

The local-market analysis aggregates annual QCEW county–industry records to a fixed set of Core-Based Statistical Areas and maps historical NAICS codes to a harmonized four-digit industry. A treated market combines the target's pre-acquisition industry with the modal post-transfer area of its transferred high-skill workers. The first qualifying acqui-hire defines the treatment year; the transaction year is omitted, and subsequent transactions do not reset event time. Disclosure-suppressed cells remain missing. The estimation window is balanced within each industry–cohort block.

Within a three-digit-industry-by-treatment-year block, candidate donors are markets with no observed acqui-hire during the sample. Donors are admitted when their own pre-treatment outcome growth lies within the range spanned by treated markets. If that rule produces fewer than 15 donors, I use a band of two standard deviations around the treated mean. Synthetic difference-in-differences weights are estimated within each block, and the resulting effects are aggregated using treated-market counts. Appendix Figure A.6 compares the weighted treated and synthetic paths. The close pre-treatment fit complements the event-study lead tests reported with Figure 9.

The preferred variance estimator deletes one geographic area at a time and recomputes the treated-market-weighted aggregate. This keeps the set of industry–cohort blocks fixed in every replication. A synchronized-area pairs bootstrap that re-estimates the synthetic weights within the fixed donor sets produces the same statistical-significance conclusions, but some bootstrap draws remove all treated markets from a block and therefore change the set of effects entering the aggregate. I use the fixed-estimand jackknife in the main text and treat the bootstrap as a sensitivity analysis. Neither procedure repeats the initial donor-admission screen.

For the BDS analysis, markets are metropolitan areas interacted with the published industry sectors. I retain a market only when every firm-age category needed to construct an outcome is reported in every year of its estimation window; disclosure and data-quality suppressions are never coded as zeros. The same transaction-year exclusion, donor rule, and treated-market aggregation are then applied within sector-by-treatment-year blocks. Because the published support can differ across age-specific outcomes, Table 12 interprets each BDS estimate relative to its own synthetic comparison and does not form differences between the young- and mature-firm coefficients. Firms aged zero measure employer-firm births. Firms aged zero to five and the establishments they operate are stocks that also incorporate survival and aging.