Talent Hoarding and Upstream Innovation: Labor Market Distortions by Large Incumbents

Shaolong Wu [*], Zefan Qian [†‡]

Latest Version: July 2026

Abstract

Why might large incumbents hire more frontier talent than they effectively deploy? How do they shape the allocation of scarce frontier talent before innovation is realized? We study this question through a model of talent hoarding upstream of product markets, in which a large incumbent and a specialized rival compete for scarce frontier researchers whose allocation shapes discovery probabilities. The incumbent trades off the gains from innovation against the risk of cannibalizing legacy rents, while also facing external displacement risk from the rival. Because hiring a researcher both raises the incumbent's probability of success and reduces the rival's, external displacement risk creates a retention motive that can lead incumbents to keep key researchers inside the firm boundary even when they are not fully deployed. The model predicts that when complementary inputs are costly, incumbents may retain more frontier researchers than they can efficiently deploy in innovation activity, leading some retained workers to remain on the "bench" without proportional capital investment. We test this prediction using *Alice v. CLS Bank* (2014), which weakened appropriability for software-related claims. In a firm-month panel linking labor-market flows to firm outcomes, firms with pre-shock software patent exposure increase net hiring after *Alice*, primarily through reduced separations, with effects concentrated among workers with advanced degrees and negligible capital expansion. Using worker career histories data, we show that occupations more plausibly tied to incumbent-specific knowledge (e.g., Science and Research) experience slower upward career progression following *Alice*, whereas more administrative occupations exhibit little evidence of comparable effects.

JEL classification: J24, O31, D4, D61, L1, L26

Keywords: Talent hoarding, Labor misallocation, Large incumbents, Upstream innovation, Innovation strategy, Strategic hiring, Market equilibrium

[*]Harvard Business School, Harvard University: lorrywu@g.harvard.edu

[†]Department of Economics, Georgetown University: zq49@georgetown.edu

[‡]We thank Axel Anderson, David Autor, Chris Avery, Cauê Dobbin, Chiara Farronato, Bob Gibbons, Shane Greenstein, Oliver Hart, Namrata Kala, Quan Le, Josh Lerner, Todd Lensman, Arik Levinson, Max Miller, Kyle Myers, Maria Roche, John Rust, Kathy Spier, Chris Stanton, Dennis Yao, and Feng Zhu, as well as seminar participants at Harvard and Georgetown, for their helpful suggestions. All errors are our own.

1 Introduction

Large incumbents sometimes pay scarce research and engineering workers to remain inside the firm even when they are not fully deployed in frontier R&D. Industry observers often describe such positions as a "bench": skilled workers remain inside the firm boundary, but are not matched to projects, capital, or organizational roles that fully use their capabilities.[1] For ordinary production labor, underutilization may reflect adjustment costs, demand uncertainty, or slack capacity. For research labor central to innovation, however, deliberate underdeployment raises a sharper puzzle: if these workers are scarce, expensive, and important for innovation, why would a firm pay to keep them inside while leaving them only partially deployed? This paper proposes a strategic motivation for why firms may underutilize or sideline scarce research workers.

The stakes are large because modern innovation combines rising R&D investment with concentrated control over scarce technical labor. U.S. businesses performed $692 billion of R&D in 2022, with firms employing 250 or more domestic workers accounting for roughly 91% of that spending.[2] While sustaining technological progress requires expanding overall research effort (Bloom et al., 2020), the current sorting of talent toward incumbents reduces direct innovative productivity while crowding out young entrants. Specifically, when inventors join incumbents rather than young firms, earnings rise while innovative output falls (Akcigit and Goldschlag, 2023); when large firms account for more local hiring demand, startups pay more for critical roles and grow more slowly (Bessen et al., 2023). The allocation of scarce research labor across incumbents and entrants can therefore shape not only the price of talent, but also the direction and speed of innovation.

We study two linked margins in innovation labor markets. The first is talent hoarding: an incumbent hires or retains scarce R&D workers primarily because doing so limits rivals' access to the knowledge and discovery capacity those workers carry. The second is benching: the incumbent keeps the worker inside the firm but does not deploy her at full innovative capacity.

We formalize these mechanisms in a model in which an incumbent with legacy rents competes with a specialized rival whose innovation can displace those rents. The incumbent chooses both how much scarce research labor to control and how fully to deploy that labor inside the firm. This choice reflects two pressures. External displacement makes rival innovation costly because a successful rival can erode the incumbent's legacy rents. Internal cannibalization makes the incumbent's own innovation less valuable when new products

[1] Recent reporting describes highly paid AI scientists placed on extended paid "garden leave" under strict noncompete terms, fully compensated but temporarily barred from joining rival labs; other accounts from the post-2020 technology hiring boom describe elite technical workers hired into roles but then spending long stretches without a clear project, team, or deployable research agenda. See, e.g., https://www.wsj.com/articles/these-tech-workers-say-they-were-hired-to-do-nothing-762ff158; https://www.businessinsider.com/google-deepmind-ai-talent-war-aggressive-noncompetes-2025-4; https://techcrunch.com/2025/04/07/google-is-allegedly-paying-some-ai-staff-to-do-nothing-for-a-year-rather-than-join-rivals/. Industry commentary has used similar language for technology labor markets; see Bain Insights, https://www.bain.com/insights/tech-talent-war-tech-report-2021, and *The Economist*, https://www.economist.com/business/2016/11/05/tech-firms-shell-out-to-hire-and-hoard-talent.

[2] National Center for Science and Engineering Statistics, *Business R&D Performance in the United States Nears $700 Billion in 2022*, NSF 24-334.

erode profits from existing lines.

The model separates how much scarce research talent the incumbent keeps inside the firm from how much of that talent it actually deploys in innovation. Scarcity makes control strategically valuable: hiring one researcher raises the incumbent's probability of innovation success, but it also lowers the rival's probability of innovation success by reallocating a key input away from the outside market. When outside innovation threatens the incumbent's legacy business, the incumbent therefore values the worker not only for her internal contribution, but also for the reduction in rival discovery risk. This gives rise to talent hoarding. Deployment is a separate margin. After the worker is retained, full deployment may require complementary inputs such as compute, specialized equipment, proprietary data, managerial attention, and organizational bandwidth. When these inputs are costly or slow to scale, or when the worker's internal contribution is limited by cannibalization concerns, the incumbent may rationally keep the worker inside while investing less aggressively in the complementary inputs needed to make her fully productive. External displacement risk raises the value of retention; deployment costs determine how much retained talent becomes effective internal innovation. Instead of existing products or patents that can be acquired and shelved ex post, we focus on scarce research labor that shapes who can innovate, who can implement, and which technological trajectories remain feasible.

This logic delivers a sharp empirical prediction: when external displacement risk rises, exposed incumbents should retain more scarce researchers inside firm boundaries, even without proportional capital expansion. We test this prediction using *Alice Corp. v. CLS Bank*, a 2014 Supreme Court decision that weakened appropriability for abstract-idea-like software and business-method patent claims. In June 2014, the Court held that implementing an abstract idea on a generic computer does not make it patent-eligible; in *Alice*-affected technologies, the likelihood that a first office action contained a Section 101 rejection rose by 31% in the following 18 months. We use firms' predetermined software-related patent portfolios to proxy for exposure to this shock and compare more- and less-exposed firms before and after the decision. In the model, the shock reduces incumbents' ability to rely on broad IP-based exclusion and raises the perceived threat from outside innovators, increasing the private value of retaining scarce frontier workers. Consistent with this interpretation, recent entrant-side evidence shows that venture-capital funding rose substantially in *Alice*-exposed industries after the decision (Feng and Williams, 2026).

Using a firm-month panel that links labor-market flows to firm outcomes, we find that firms with pre-shock software patent exposure increase monthly net hiring by about 0.34 percentage points after *Alice* relative to less exposed firms in our preferred specification. This corresponds to roughly 3.4 additional net hires per 1,000 employees in a month and amounts to about 63% of the mean net-hiring rate in the software sample. Decomposing the result shows that the effect is driven primarily by reduced separations rather than increased gross hiring, consistent with retention-based hoarding. The response is concentrated among workers with advanced degrees and is not accompanied by proportional capital expansion. Thus, exposed incumbents retain more high-skill labor after the shock, but the response does not look like a standard

scale-up in which labor, hiring, and complementary capital expand together.

We then use a worker-year panel of career histories to examine the occupational trajectories of workers attached to exposed incumbents around the shock. Among workers whose 2013 occupations are more plausibly tied to incumbent-specific knowledge, such as Science and Research, those working at exposed incumbents in 2014 subsequently appear in occupations with lower O*NET Job Zones relative to comparable workers at less-exposed firms. O*NET Job Zone is an occupation-level measure of the preparation, training, and experience typically required for an occupation, so we interpret this result as a flatter occupational trajectory rather than as a literal promotion measure. The relative Job-Zone gap grows over the post period and reaches roughly 0.06–0.07 Job-Zone points by the fifth year after *Alice*. The worker-level evidence therefore complements the firm-level labor-flow and capital facts: exposed incumbents retain more frontier labor, while affected research workers move into occupations that, on average, involve less upward movement in required preparation and task complexity.

Our contribution is to shift the analysis of incumbent defense from downstream products and patents to the upstream allocation of frontier talent. Existing work often studies strategic control after an asset is already visible: a rival product can be acquired, a patent can be asserted or slept on, or a downstream input can be tied or withheld. We study control before the asset exists, when the strategic lever is scarce research capacity itself. The model separates two incumbent choices that are often bundled together in innovation-investment models: whether to keep scarce frontier workers inside the firm and whether to deploy those workers at full innovative capacity. This distinction is important because wages can reflect the incumbent's private value from attracting and retaining talent, including preemptive and defensive value, rather than the broader social value of implemented discovery. As a result, incumbents can attract frontier workers even when alternative employers would use them more intensively. The welfare concern is thus both monopoly pricing or ex post shelving and slower discovery or delayed diffusion caused by the upstream misallocation of frontier talent.

The remainder of the paper is organized as follows. Section 2 positions the paper within related work on research-labor allocation, strategic hiring, inventor labor-market power, and incumbent defense. Section 3 develops a model in which an incumbent chooses how much frontier talent to retain and deploy when outside innovation threatens legacy rents and internal innovation may cannibalize them. Section 4 describes the empirical setting, data construction, and the *Alice* shock. Section 5 tests the model's comparative statics using firm-level labor-flow evidence and worker-level career histories. Section 6 discusses limitations and extensions, and Section 7 concludes.

2 Literature Review

This paper studies how incumbent firms use scarce research labor when outside innovation threatens legacy rents. We situate the paper in three core literatures. Work on R&D productivity and inventor allocation

shows that where researchers work affects innovative output. Work on strategic hiring and inventor labor-market power shows that R&D labor markets are not well described by perfect competition. Work on incumbent defense shows that firms may protect existing rents by controlling assets that threaten them. We connect these facts through a labor-input channel: an incumbent can retain researchers who would raise rivals' discovery capacity while failing to expand the complementary inputs needed to make those workers productive inside the firm. The review proceeds from research-labor allocation to strategic hiring, incumbent defense, and worker careers.

Evidence on R&D and talent allocation gives reason to treat research labor as allocation-sensitive. Aggregate research effort has risen while measured research productivity has fallen, so the organization of research inputs is crucial for realized technological progress (Bloom et al., 2020). Large-firm hiring demand raises startup labor costs and slows startup growth (Bessen et al., 2023). Inventors have increasingly moved toward large incumbents, and moves into incumbents are associated with higher earnings but lower innovative output (Akcigit and Goldschlag, 2023). These papers establish the premise that the same researcher can have different innovative value across employers. We study one source of that allocation: incumbents' retention of scarce research workers after patent-based exclusion weakens. Exposed incumbents retain more high-skill workers, mainly through lower separations, without a matching expansion in observable capital or worker advancement.

The closest theoretical work studies strategic hiring, acqui-hiring, and team absorption in thin labor markets. Acqui-hiring and related team purchases move specialized workers across firm boundaries and can substitute for buying a product-market asset (Benkert et al., 2024; Friedmann et al., 2025). Bar-Isaac et al. (2025) shows that an acquihire can buy labor-market power by removing a startup as a source of counteroffers, lowering workers' outside options, and making a privately attractive deal socially inefficient. Fernández-Villaverde et al. (2025) study defensive hiring of researchers by incumbents with market power and inelastic research labor. We study a different way incumbents can use scarce research labor strategically: talent hoarding combines defensive retention with underdeployment. In our setting, no acquisition is required, no rival product needs to exist, and the key decision is not only whether the researcher remains inside the incumbent or becomes available to outside innovators, but also whether retained talent is deployed at full innovative capacity. By retaining research and engineering workers, the incumbent can reduce the probability that the same worker helps a rival discover and commercialize a replacement technology. This effect can arise before a rival patent, product, or startup becomes visible. Recent work on inventor monopsony studies a related allocation problem. Lehr (2026) shows that large firms face less elastic inventor labor supply and that monopsony can reduce R&D employment and misallocate inventors across firms. We add evidence on defensive retention after patent-based exclusion weakens. These mechanisms point in different employment directions for different reasons: wage-setting power can reduce inventor employment by making additional hiring costly, while displacement risk can raise retention by making outside access to researchers costly to the incumbent.

We are also related to the industrial-organization literature on incumbent defense through control of products, patents, or market access, dating back to the Arrow replacement effect, where incumbents may invest less in innovations that cannibalize existing rents (Arrow, 1962). Other papers study foreclosure (Whinston, 1990), strategic patenting and sleeping patents (Gilbert and Newbery, 1982; Hall and Ziedonis, 2001), non-practicing entities (Cohen et al., 2019), and killer acquisitions (Cunningham et al., 2021). Across these settings, the relevant asset already exists and may later be underused: an input can be withheld, a patent can be asserted or left unused, a project can be acquired, or a startup can be bought. To our knowledge, this is the first paper to examine an earlier labor-side channel of incumbent defense: retaining research workers before a rival product, patent, or startup becomes visible. This generates a distinct testable empirical pattern: a firm may retain more frontier labor after external displacement risk rises, yet fail to expand capital or worker advancement in a way consistent with full deployment.

Search and personnel economics help interpret the worker-level evidence. Search and matching models explain why frictions generate inefficient matches and wage dispersion (McCall, 1972; Postel-Vinay and Robin, 2002). Personnel economics links job design, internal assignment, and promotion incentives to worker allocation inside firms (Kahn and Huberman, 1988; Lazear and Shaw, 2007; Waldman and Zax, 2016). Recent work on within-firm talent hoarding shows that managers may keep strong workers inside low-fit units rather than release them to better assignments (He and Waldman, 2024; Haegele, 2024). These papers differ from our main mechanism, but they help interpret the worker-level career patterns. If retained researchers are valuable partly because they reduce rivals' access to frontier know-how, their occupational paths need not resemble those of researchers fully deployed on frontier projects.

3 Model

This section develops a theoretical framework of strategic talent hoarding. Rather than modeling labor as a frictionless factor of production, the mechanism here posits that an incumbent firm retains workers with knowledge of its legacy business not because all of those workers are actively deployed in internal innovation, but to strategically reduce a competitor's chance of discovering a replacement technology. To formalize this friction, the model introduces a structural distinction between labor deployment – workers actively utilized in ongoing innovation activities – and an internal labor bench, where excess talent is hoarded primarily to foreclose the external pool of ideas and prevent entry.

The market structure consists of two asymmetric players: an incumbent firm that initially monopolizes a legacy technology, and a potential entrant seeking to develop a disruptive, replacement technology. The critical bottleneck to innovation is a fixed mass $m > 0$ of workers endowed with incumbent-specific knowledge. This specialized talent pool possesses proprietary expertise regarding the incumbent's underlying technology, client architecture, and operational routines – assets that are highly complementary to, and uniquely valuable for, the development of the entrant's competing product. Initially, this entire mass m of skilled labor resides

within the incumbent firm.

Crucially, this framework is not intended to model aggregate employment dynamics or routine personnel flows at either firm. Both entities may actively hire or separate workers across a variety of orthogonal, non-core business segments, and those baseline labor flows may be empirically large. Instead, the endogenous labor margin in our model is purposefully narrow: it isolates the structural allocation of the initial stock m of specific talent between the incumbent's internal organization and the entry threat.

3.1 Setup

3.1.1 Talent, deployment, and benching

The incumbent chooses a total mass of initial talent to retain, denoted by $\ell_I \in [0, m]$. The remaining mass $m - \ell_I$ separates from the firm and enters the external labor market, becoming available to the entrant. We formally operationalize the distinction between talent retention and talent deployment by modeling the production of replacement R&D. Crucially, a firm requires innovation capital to actively deploy its specific talent. For simplicity, we assume these inputs are perfect complements combined via a symmetric, one-to-one Leontief production technology. If a firm allocates L units of talent and k units of capital, its effective R&D input is given by: $X = \min\{L, k\}$.

Because innovation capital is supplied elastically at an exogenous rental rate $r > 0$ and yields a zero marginal product unless matched with talent, cost-minimizing firms will never choose to rent capital in excess of deployed labor. Consequently, active deployment is equivalent to the volume of matched talent-capital pairs. Among its retained workers, the incumbent chooses how many to deploy to its own replacement project: $x \in [0, \ell_I]$, which implicitly requires renting exactly x units of innovation capital at cost $r \cdot x$. The residual retained talent is structurally benched: $z \equiv \ell_I - x$. While benched workers do not enter the replacement R&D success technology, they are successfully foreclosed from the external labor pool.

If the potential entrant attempts the replacement project, its available hiring pool is bounded by the unretained talent supply. Let ℓ_E denote the entrant's employment of this specific talent, subject to the feasibility constraint: $0 \leq \ell_E \leq m - \ell_I$. By the same cost-minimizing Leontief logic, the entrant optimally pairs this talent with an identical quantity of innovation capital. Thus, a single variable ℓ_E completely characterizes the entrant's effective innovation input, capturing both the specific labor hired and the capital rented at a unit cost of $r \cdot \ell_E$.

3.1.2 Replacement success

Both firms can attempt to develop a replacement product, denoted by b. Let $p(u)$ denote the probability of technical success when a firm uses u effective units of replacement R&D input.

Assumption 1 (Replacement technology) *The success probability function $p : [0, m] \to [0, 1]$ is continuous on $[0, m]$, twice continuously differentiable on $(0, m)$, strictly increasing, and strictly concave, satisfying*

$p(0) = 0$, $p'(u) > 0$, *and* $p''(u) < 0$ *for all* $u \in (0, m)$. *Where* $p'(0)$ *appears in subsequent derivations, it denotes the right-sided derivative evaluated at zero.*

3.1.3 Product-market payoffs and commercialization

The incumbent currently earns an optimized, reduced-form product-market profit $\pi_0 > 0$ from its legacy business. A replacement product generates a gross product-market profit $\pi_b > 0$ only upon successful market implementation. Crucially, both π_0 and π_b represent realized payoffs at the final product-market stage; they are structurally decoupled from, and independent of, upstream R&D input expenditures, which are subtracted separately.

For the incumbent, implementation requires only internal technical success. For the potential entrant, however, market implementation is a joint outcome requiring both technical R&D success and successful commercialization. To capture commercialization barriers, we assume that even conditional on technical success, the entrant implements the replacement product with probability $Q \in [0, 1]$. A higher Q reflects weaker entry frictions, such as lower customer lock-in, reduced distribution barriers, or minimal regulatory uncertainty.

The potential entrant also maintains an outside innovation trajectory, project a, which yields an exogenous payoff $\pi_a \geq 0$. The entrant only attempts the replacement project if its expected return from joint implementation dominates this alternative baseline resource allocation.

The incumbent's product-market operating value conditional on no successful external entry is

$$\Omega_I(x) = \pi_0 + p(x)(\pi_b - \pi_0)\mathbf{1}\{\pi_b > \pi_0\}. \tag{3.1}$$

The incumbent always secures the legacy baseline π_0. If $\pi_b > \pi_0$, internal innovation allows the incumbent to upgrade its product line and capture the incremental rents $\pi_b - \pi_0$. If $\pi_b \leq \pi_0$, successful internal R&D yields no private product-market benefit, as the incumbent will not voluntarily cannibalize its superior legacy line.

If the entrant attempts the replacement project by hiring ℓ_E units of specific talent against the incumbent's internal deployment x, the entrant's gross match surplus before wages is

$$S_E^b(\ell_E, x, Q) = Q\pi_b\big(1 - p(x)\big)p(\ell_E) - r\ell_E. \tag{3.2}$$

The entrant captures π_b only if it achieves technical success, commercializes, and the incumbent's internal replacement effort fails. The final term $r\ell_E$ represents the elastically supplied innovation-capital cost pinned down by Leontief complementarity.

The term $1-p(x)$ highlights the strategic duality of incumbent deployment: internal R&D not only drives upgrading but directly suppresses the entrant's expected returns. Ultimately, our equilibrium mechanics rest on two robust forces: external commercialization erodes incumbent rents, and internal incumbent deployment

strategically depresses external replacement surplus.

3.1.4 Timing and equilibrium

The game consists of four stages:

1. **Retention and Deployment:** The incumbent chooses total labor retention $\ell_I \in [0, m]$ and deployment $x \in [0, \ell_I]$. The remaining benched labor is $z = \ell_I - x$.
2. **Project Choice:** The entrant observes (ℓ_I, x) and chooses a project $j \in \{a, b\}$, where a denotes the outside project and b denotes the replacement project.
3. **Employment Choice:** If the entrant chooses project b, it hires a quantity of incumbent-specific talent $\ell_E \leq m - \ell_I$ and rents the matching amount of innovation capital.
4. **Payoffs:** Innovation and commercialization outcomes are realized, and payoffs are distributed.

The entrant breaks ties in favor of project a. Thus, it attempts the replacement project only when doing so yields strictly more surplus than the outside project. The incumbent breaks payoff-equivalent ties in favor of the smaller total retention level ℓ_I. This tie-breaking rule rules out retaining redundant workers after deterrence has already been achieved. The solution concept is subgame-perfect equilibrium. We solve by backward induction.

3.2 Entrant choice, thresholds, and retention premia

The entrant's decision is the first object to characterize. Given the incumbent's choices (ℓ_I, x), the entrant compares the outside project a with the replacement project b. The key question is how much incumbent-specific talent must be retained by the incumbent to make the entrant abandon the replacement trajectory.

3.2.1 Entrant employment conditional on replacement entry

The entrant's marginal surplus from the ℓ_E-th unit of incumbent-specific talent on project b is

$$w_E^{\max}(\ell_E; x, Q) = Q\pi_b\big(1 - p(x)\big)p'(\ell_E) - r. \tag{3.3}$$

This object is not itself a wage offer. It is the entrant's maximum willingness to pay for an additional unit of talent before the total match surplus is divided between the entrant and workers. The mechanism of interest is talent scarcity: the entrant would like to use more incumbent-specific talent than it can access. We thereby focus on the following parameter region.

Assumption 2 (Labor-hungry replacement region) *On the parameter region of interest,*

$$Q\pi_b\big(1 - p(x)\big)p'(\ell_E) > r \tag{3.4}$$

for every relevant (x, ℓ_E) before the entrant is deterred. Thus, conditional on choosing project b, the entrant's match surplus is increasing in accessible incumbent-specific talent.

Assumption 2 defines the parameter space where specialized labor is a scarce and binding resource for the entrant. This assumption focuses the analysis on the economically interesting region where labor market withholding operates as a strategic instrument of defense.

Lemma 1 (Entrant employment on the replacement project) *Suppose Assumptions 1 and 2 hold. If the entrant chooses project b, then it hires all accessible incumbent-specific talent:*

$$\ell_E^*(\ell_I) = m - \ell_I. \tag{3.5}$$

The proof is Proof A.1. Lemma 1 lets us summarize the entrant's replacement surplus as a function of the incumbent's total retention ℓ_I and deployment x. Conditional on choosing b, the entrant uses all talent not retained by the incumbent.

3.2.2 The total-retention threshold

Using Lemma 1, the entrant's surplus from the replacement project is

$$S_E^b(m - \ell_I, x, Q) = Q\pi_b\big(1 - p(x)\big)p(m - \ell_I) - r(m - \ell_I). \tag{3.6}$$

The entrant chooses project b if and only if

$$S_E^b(m - \ell_I, x, Q) > \pi_a. \tag{3.7}$$

At equality, the entrant chooses project a. Define

$$G(\ell_I; x, Q) = Q\pi_b\big(1 - p(x)\big)p(m - \ell_I) - r(m - \ell_I) - \pi_a. \tag{3.8}$$

The function G is the entrant's replacement surplus relative to the outside project. The entrant chooses b when $G > 0$ and chooses a otherwise. For each (x, Q), define the total-retention threshold

$$\ell_I^a(x, Q) = \inf\{\ell \in [0, m] : G(\ell; x, Q) \leq 0\}. \tag{3.9}$$

This is the smallest amount of incumbent retention that makes the entrant prefer project a.

Proposition 1 (Entrant best response and total-retention threshold) *Suppose Assumptions 1 and 2 hold. Then:*

(i) $G(\ell_I; x, Q)$ is strictly decreasing in ℓ_I.

(ii) The threshold $\ell_I^a(x,Q)$ is well defined and unique.

(iii) The entrant chooses b if $\ell_I < \ell_I^a(x,Q)$ and chooses a if $\ell_I \geq \ell_I^a(x,Q)$.

(iv) Whenever the threshold is interior, it is characterized by

$$Q\pi_b\big(1-p(x)\big)p\big(m-\ell_I^a(x,Q)\big) - r\big(m-\ell_I^a(x,Q)\big) = \pi_a. \tag{3.10}$$

(v) Whenever the threshold is interior,

$$\frac{\partial \ell_I^a(x,Q)}{\partial Q} > 0, \qquad \frac{\partial \ell_I^a(x,Q)}{\partial x} < 0. \tag{3.11}$$

The proof is Proof A.2.

3.2.3 Retention premia

The incumbent does not pay the entrant's willingness to pay for every worker it retains. The relevant cost is the premium required to prevent workers from being drawn into the entrant's replacement project. Once the incumbent has retained enough talent that the entrant switches to project a, additional retained workers no longer command a replacement-project poaching premium. Let

$$\kappa(\ell_I, x, Q) = \min\{\ell_I, \ell_I^a(x,Q)\} \tag{3.12}$$

be the amount of incumbent retention that is relevant for removing surplus from the entrant's replacement option. We normalize ordinary compensation and legacy-task productivity into π_0. The object C_R is the incremental retention premium required to keep these workers from being reallocated to the entrant's replacement project. Equivalently, workers or their outside options appropriate the marginal replacement surplus that the entrant would generate from access to them.

The retention premium is

$$\begin{aligned} C_R(\ell_I, x, Q) &= \int_{m-\kappa(\ell_I,x,Q)}^{m} \big[Q\pi_b\big(1-p(x)\big)p'(s) - r\big]\, s \\ &= Q\pi_b\big(1-p(x)\big)\,\big[p(m) - p\big(m-\kappa(\ell_I,x,Q)\big)\big] - r\kappa(\ell_I,x,Q). \end{aligned} \tag{3.13}$$

When $\ell_I < \ell_I^a(x,Q)$, the entrant still chooses project b, and the premium rises with additional retention. When $\ell_I \geq \ell_I^a(x,Q)$, the entrant is already deterred, and the premium is capped. If the deterrence threshold is interior and $\ell_I \geq \ell_I^a(x,Q)$, then the threshold equality (3.10) implies

$$C_R(\ell_I, x, Q) = Q\pi_b\big(1-p(x)\big)p(m) - rm - \pi_a. \tag{3.14}$$

Expression (3.14) shows that, in the deterrence region, deployment lowers the retention premium by reducing the entrant's expected replacement surplus.

3.3 The incumbent problem and deterrence regimes

The incumbent chooses total retention and deployment. The key tradeoff is between leaving the entrant on the replacement project and retaining enough incumbent-specific talent to make the entrant abandon that project.

3.3.1 Incumbent payoff

Given (ℓ_I, x), the incumbent's continuation value is

$$V_I(\ell_I, x, Q) = \begin{cases} \big(1 - Qp(m - \ell_I)\big)\Omega_I(x) - rx - C_R(\ell_I, x, Q), & \text{if } \ell_I < \ell_I^a(x, Q), \\ \Omega_I(x) - rx - C_R(\ell_I, x, Q), & \text{if } \ell_I \geq \ell_I^a(x, Q). \end{cases} \tag{3.15}$$

If $\ell_I < \ell_I^a(x, Q)$, the entrant remains on the replacement project. If $\ell_I \geq \ell_I^a(x, Q)$, the entrant switches to the outside project and the incumbent avoids the replacement-entry threat. Let

$$\mathcal{L} = \{(\ell_I, x) : 0 \leq x \leq \ell_I \leq m\}$$

denote the feasible set. The incumbent solves

$$\max_{(\ell_I, x) \in \mathcal{L}} V_I(\ell_I, x, Q). \tag{3.16}$$

A useful benchmark is no retention:

$$(\ell_I, x, z) = (0, 0, 0).$$

When the entrant chooses b absent incumbent retention, this benchmark gives the incumbent value

$$J^A(Q) = (1 - Qp(m))\pi_0. \tag{3.17}$$

The entrant is an active replacement threat absent incumbent retention when

$$G(0; 0, Q) = Q\pi_b p(m) - rm - \pi_a > 0. \tag{3.18}$$

Equivalently, since $p(m) > 0$, the replacement threat is active when

$$Q > Q^T \equiv \frac{rm + \pi_a}{\pi_b p(m)}. \tag{3.19}$$

If (3.18) fails, the entrant chooses project a even without incumbent retention. The incumbent then chooses $\ell_I = x = z = 0$ by the tie-breaking rule.

3.3.2 No self-replacement incentives

The main case of interest is

$$\pi_b \leq \pi_0. \tag{3.20}$$

We call this the no self-replacement case. The replacement product may be valuable to the entrant because it can attack the incumbent's legacy position, but it is not privately more valuable to the incumbent than the legacy line it already owns. In this case, internal replacement R&D has no direct product-market benefit for the incumbent:

$$\Omega_I(x) = \pi_0, \forall x.$$

Deployment may still be useful, but only because it reduces the entrant's replacement surplus and hence reduces the retention premium needed for deterrence. The following result is the key simplification.

Proposition 2 (No positive-retention competition under no self-replacement) *Suppose $\pi_b \leq \pi_0$. Then no equilibrium with positive incumbent retention of incumbent-specific talent has the entrant remaining on project b. Any equilibrium with $\ell_I^* > 0$ deters the entrant to project a.*

The proof is Proof A.3. Proposition 2 rules out positive retention that merely softens competition while leaving the entrant on b. The next lemma shows that, when the replacement threat is active, the incumbent does not choose the no-retention benchmark either.

Lemma 2 (Active threats are deterred under no self-replacement) *Suppose $\pi_b \leq \pi_0$ and $G(0; 0, Q) > 0$. Then every equilibrium deters the entrant. In particular, equilibrium retention satisfies*

$$\ell_I^*(Q) \geq \ell_I^a(x^*(Q), Q).$$

The proof is Proof A.4.

3.3.3 Reduced deterrence problem and regimes

Lemma 2 lets us analyze the active-threat case as a deterrence problem. For a fixed deployment level x, the smallest total retention level that deters the entrant is

$$\ell_I^D(x, Q) = \max\{x, \ell_I^a(x, Q)\}. \tag{3.21}$$

The corresponding amount of benching is

$$z^D(x, Q) = \ell_I^D(x, Q) - x = \max\{\ell_I^a(x, Q) - x, 0\}. \tag{3.22}$$

The incumbent's reduced deterrence value is

$$\widetilde{V}^D(x,Q) = V_I\big(\ell_I^D(x,Q),x,Q\big). \tag{3.23}$$

Let

$$x^D(Q) \in_{x\in[0,m]} \widetilde{V}^D(x,Q) \tag{3.24}$$

denote the incumbent's deployment choice in the reduced deterrence problem. If there are multiple maximizers, we select the smallest one, consistent with the incumbent's tie-breaking rule.

On the region where the deterrence threshold is interior and positive, (3.14) gives

$$\widetilde{V}^D(x,Q) = \Omega_I(x) - rx - \big[Q\pi_b\big(1-p(x)\big)p(m) - rm - \pi_a\big]. \tag{3.25}$$

Differentiating with respect to deployment gives

$$\frac{\partial \widetilde{V}^D(x,Q)}{\partial x} = p'(x)(\pi_b-\pi_0)\mathbf{1}\{\pi_b>\pi_0\} - r + Q\pi_b p'(x)p(m). \tag{3.26}$$

Under no self-replacement, this becomes

$$\frac{\partial \widetilde{V}^D(x,Q)}{\partial x} = Q\pi_b p'(x)p(m) - r. \tag{3.27}$$

Thus deployment can arise even when replacement is not privately more profitable for the incumbent. In this case, deployment is valuable because it makes deterrence cheaper.

To distinguish partial benching from no-bench deterrence, define

$$H(x,Q) = G(x;x,Q) = Q\pi_b\big(1-p(x)\big)p(m-x) - r(m-x) - \pi_a. \tag{3.28}$$

This is the entrant's replacement surplus relative to project a when the incumbent retains exactly the workers it deploys, so $\ell_I = x$ and $z = 0$. Let

$$x^{pd}(Q) = \inf\{x\in[0,m] : H(x,Q)\le 0\}. \tag{3.29}$$

If $x \ge x^{pd}(Q)$, deployment together with the retention mechanically required for that deployment is enough to deter the entrant without benching. If $x < x^{pd}(Q)$, deterrence requires positive benching.

Theorem 1 (Deterrence regimes under no self-replacement) *Suppose Assumptions 1 and 2 hold. Suppose also that $\pi_b \le \pi_0$ and $G(0;0,Q) > 0$. Then the entrant is deterred in equilibrium. The equilibrium takes one of the following forms:*

(i) ***Pure benching deterrence.*** *If* $x^D(Q) = 0$*, then*

$$x^*(Q) = 0, \qquad \ell_I^*(Q) = z^*(Q) = \ell_I^a(0, Q).$$

A sufficient condition in the interior active-deterrence region is

$$Q\pi_b p(m)p'(0) \leq r. \tag{3.30}$$

(ii) ***Partial benching deterrence.*** *If*

$$0 < x^D(Q) < x^{pd}(Q),$$

then

$$x^*(Q) = x^D(Q), \qquad \ell_I^*(Q) = \ell_I^a(x^D(Q), Q), \qquad z^*(Q) = \ell_I^a(x^D(Q), Q) - x^D(Q) > 0.$$

Deployment and benching jointly deter the entrant.

(iii) ***Pure deployment, or no-bench, deterrence.*** *If*

$$x^D(Q) \geq x^{pd}(Q),$$

then

$$\ell_I^*(Q) = x^*(Q) = x^D(Q), \qquad z^*(Q) = 0.$$

Deployment alone, together with the retention required to deploy, deters the entrant.

The proof is Proof A.5.

3.4 Commercialization probability and equilibrium labor allocation

We now ask how equilibrium labor allocation responds to an increase in Q, the entrant's commercialization probability conditional on technical success. Throughout this subsection, suppose $\pi_b \leq \pi_0$ and focus on the active-threat region $Q > Q^T$. By Lemma 2, the incumbent deters entry in equilibrium. The key object is the distinction between total retention and deployment:

$$\ell_I^*(Q) \quad \text{is total retention,} \qquad x^*(Q) \quad \text{is deployment,} \qquad z^*(Q) = \ell_I^*(Q) - x^*(Q) \quad \text{is benching.}$$

The following proposition gives the prediction in the partial-benching and no-bench regimes.

Proposition 3 (Commercialization shocks and labor allocation) *Suppose Assumptions 1 and 2 hold, suppose* $\pi_b \leq \pi_0$*, and consider an interval of* Q *on which the equilibrium regime is unchanged and the relevant thresholds are interior.*

(i) ***Partial benching deterrence.*** *If the equilibrium remains in an interior partial-benching regime, then*

$$\ell_I^*(Q) = \ell_I^a(x^*(Q), Q), \qquad z^*(Q) = \ell_I^a(x^*(Q), Q) - x^*(Q) > 0.$$

Define

$$A(Q) = \frac{\partial \ell_I^a(x^*(Q), Q)}{\partial Q} > 0, \qquad B(Q) = -\frac{\partial \ell_I^a(x^*(Q), Q)}{\partial x} > 0. \tag{3.31}$$

The deployment response is

$$\frac{dx^*(Q)}{dQ} = -\frac{p'(x^*(Q))}{Qp''(x^*(Q))} > 0. \tag{3.32}$$

Total retention and benching satisfy

$$\frac{d\ell_I^*(Q)}{dQ} = A(Q) - B(Q)\frac{dx^*(Q)}{dQ}, \tag{3.33}$$

and

$$\frac{dz^*(Q)}{dQ} = A(Q) - \big(1 + B(Q)\big)\frac{dx^*(Q)}{dQ}. \tag{3.34}$$

Hence the partial-benching regime has two empirically relevant responses:

$$\frac{dx^*(Q)}{dQ} < \frac{A(Q)}{1 + B(Q)} \quad \Longrightarrow \quad \frac{d\ell_I^*(Q)}{dQ} > 0 \text{ and } \frac{dz^*(Q)}{dQ} > 0, \tag{3.35}$$

$$\frac{A(Q)}{1 + B(Q)} < \frac{dx^*(Q)}{dQ} < \frac{A(Q)}{B(Q)} \quad \Longrightarrow \quad \frac{d\ell_I^*(Q)}{dQ} > 0 \text{ and } \frac{dz^*(Q)}{dQ} < 0. \tag{3.36}$$

In the first case, the incumbent retains more talent and some of the additional retained talent is benched. In the second case, total retention still rises, but the incumbent uses the stronger deployment response to draw down the existing bench.

(ii) ***Pure deployment, or no-bench, deterrence.*** *If the equilibrium remains in pure deployment deterrence, then*

$$\ell_I^*(Q) = x^*(Q), \qquad z^*(Q) = 0.$$

In the interior cases in which either the deployment first-order condition or the zero-benching boundary governs the choice of $x^*(Q)$,

$$\frac{d\ell_I^*(Q)}{dQ} = \frac{dx^*(Q)}{dQ} > 0, \qquad \frac{dz^*(Q)}{dQ} = 0. \tag{3.37}$$

The proof is Proof A.6. Proposition 3 provides the bridge to the empirical analysis. An exogenous increase in commercialization probability (Q) raises the entrant's replacement value and therefore increases the incumbent's incentive to retain incumbent-specific talent. The model allows this additional retention to be allocated in two ways.

In the pure-deployment region, additional retained workers are matched with innovation capital and immediately deployed, so total retention and deployment rise one-for-one. Empirically, this deployment margin should be accompanied by a proportional expansion in the inputs required to use those workers productively, such as R&D capital, capital equipment, or other complementary innovation resources. Equivalently, if we observe that additional retained workers are not matched by proportional increases in these deployment inputs and instead appear as benched labor, then the firm was not operating on the pure-deployment margin before the exogenous increase.

In the partial-benching region, it can retain additional workers without immediately deploying all of them, or it can expand deployment by matching retained workers with innovation capital. In the short run, increased deployment may be less effective as a deterrence device because deployment is not immediately observed. We therefore anticipate that incumbents first respond by increasing retention, with some of the additional retained workers not immediately deployed. Over longer horizons, firms may also expand active deployment. This is analogous to defensive patenting strategies: firms may invest aggressively to secure control over a technology space, but not all protected ideas are immediately commercialized or actively used.

The empirical prediction is therefore twofold. First, firms exposed to an exogenous increase in commercialization probability should increase retention of incumbent-specific workers relative to otherwise similar firms. Second, the allocation of this additional retention distinguishes the deployment regime in the short run. If the increase in retained workers is accompanied by a proportional expansion in R&D capital or other deployment inputs, the response is consistent with pure deployment both before and after the increase in commercialization probability. If retention rises without a proportional expansion in these inputs, then the firm is operating on the partial-benching margin both before and after the shock: some incumbent-specific workers are retained inside the firm but are not immediately matched with the capital required for active deployment.

4 Setting and Data

We proxy the shift in the commercialization probability using a plausibly exogenous legal shock, the Supreme Court's 2014 decision in *Alice Corp. v. CLS Bank International*.

4.1 Institutional Background: Patent-Eligibility Doctrine and the Alice Shock

Prior to *Alice*, U.S. patent-eligibility doctrine had evolved in a way that made patents on computerized methods commercially salient in software-intensive and financial-technology settings. In practice, this doctrinal environment supported the accumulation and enforcement of large portfolios containing software and business-method claims, including claims written at a relatively high level of abstraction (e.g., method steps implemented "on a computer").

The case arose when CLS Bank challenged patents held by Alice Corporation covering a computerized

scheme for mitigating settlement risk in financial transactions. CLS argued that the patents were directed to an abstract idea and therefore were not eligible for patent protection under 35 U.S.C. §101. In a unanimous decision issued on June 19, 2014, the Supreme Court held that merely implementing an abstract idea on a generic computer does not render the invention patent-eligible.

The Court applied a two-step framework: it first asks whether the claims are directed to a patent-ineligible concept (including abstract ideas), and if so, whether additional claim elements supply an "inventive concept" sufficient to amount to "significantly more" than the ineligible concept itself. *Alice* tightened the patent-eligibility screen for claims that recite an abstract idea plus only conventional or generic computer implementation. Because *Alice* changes a legal standard, the practical impact on portfolio expectations and prosecution depends on implementation and standardization in examination practice. In the empirical design, we use June 2014 as the main post start as a conservative implementation cutoff.

By making abstract-idea-like software claims harder to obtain and enforce when they relied only on generic computer implementation, *Alice* weakened a specific appropriability margin: the ability to use broad software claims to block or deter close substitutes.[3] This maps directly to the comparative static in the model. For incumbents with greater pre-shock exposure to such claims, weaker patent-based exclusion raises the perceived threat from outside innovators and makes retaining frontier workers inside the firm a more valuable defensive instrument. The empirical design therefore uses predetermined software-patent exposure to compare firms differentially exposed to this exclusion-margin shock, rather than treating *Alice* as a broad shock to software profitability, technological opportunities, or production technology.

Tables 1 and 2 provide the institutional background for this mapping. Table 1 clarifies that *Alice* was followed quickly by USPTO implementation guidance, while Table 2 illustrates the specific appropriability margin most affected by the decision: abstract software claims implemented on generic computers became harder to enforce, increasing exposed incumbents' perceived risk of outside displacement.

Table 1: Timeline of the Alice shock and administrative implementation

Date	Event and interpretation for the empirical design
1998	*State Street* (Fed. Cir.) rejects a categorical business-method exclusion; articulates a permissive framing ("useful, concrete and tangible result").
June 19, 2014	Supreme Court decision in *Alice v. CLS Bank*: abstract idea + generic computer implementation is not patent-eligible under §101.
June 25, 2014	USPTO issues preliminary examination instructions following *Alice*.
Dec. 2014	USPTO issues "2014 Interim Guidance on Patent Subject Matter Eligibility" to standardize examination practice under *Alice*.

[3] A related public-firm study similarly operationalizes *Alice* as a loss of patent protection for firms whose pre-shock patent portfolios were exposed to affected business-method, software, and related claims (Acikalin et al., 2026).

Table 2: One concrete illustration of the *Alice* shock: computerized intermediated settlement (stylized)

	Pre-*Alice* (stylized)	Post-*Alice* (stylized)
Claiming style most affected	Business logic implemented "on a computer" (e.g., settlement-risk mitigation, account reconciliation, transaction processing)	"Abstract idea + generic computer" claims face heightened §101 eligibility risk; surviving claims typically require more concrete technical content ("significantly more")
Role of patents in competition	Portfolios can support the exclusion threats, licensing leverage, and defensive deterrence in software-enabled services	Patent-based protection weakens for abstract-idea-like claims; uncertainty increases about enforceability and scope
Interpretation for incumbent incentives (link to model)	Stronger expected exclusion can partially limit close imitation in the affected domain	Weaker appropriability increases the perceived exposure to close substitutes and imitation; integrated incumbents have stronger incentives to keep scarce frontier talent inside the firm boundary (e.g., via retention) to manage capability development and the timing of diffusion

Notes: This table is a stylized illustration meant to fix intuition for the empirical design. It summarizes the type of claim implicated in *Alice* (intermediated settlement implemented on a generic computer) and the corresponding shift in expected enforceability under §101.

4.2 Data and Sample Construction

Our empirical designs combine three data systems that use different firm identifiers and have different coverage. ORBIS is organized at the legal-entity level (`bvdid`) and includes both public and private companies worldwide. Patent exposure is constructed using DISCERN, which is linked to Compustat and indexed by `gvkey`. Because `gvkey` is a Compustat identifier, all analyses that use `gvkey`-based patent exposures are, by construction, conducted in the public-firm universe (i.e., public firms that can be linked to Compustat). Finally, Lightcast provides employer identifiers used to enrich outcomes (e.g., hiring and separations) and never to define treatment status.

To implement patent-based exposure to the shock and incumbency definitions for ORBIS firms, we construct a crosswalk from ORBIS legal entities (`bvdid`) to Compustat-linked parent firms (`gvkey`) using a patent-side name universe assembled from DISCERN. We standardize names on both sides and assign deterministic matches when a cleaned ORBIS name maps uniquely to a single `gvkey`. For the remaining cases, we use a conservative LLM-assisted adjudication restricted to a small candidate set generated by string-similarity heuristics; all model outputs are cached for reproducibility. We then conduct a structured human audit of all matches flagged as high-risk by transparent rules (e.g., short/generic names, low similarity, or unusually large `gvkey` clusters). Human decisions override automated outputs.

For the *Alice* shock, we construct a firm-level panel from ORBIS covering the 2012-2016 event window. The firm sample consists of active firms whose primary NACE Rev. 2 classification corresponds to software and related digital services, specifically: 58.2 (Software publishing), 62.* (Computer programming, consultancy, and related activities), and 63.1 (Data processing, hosting, and related activities; web portals). These categories capture firms operating in software development, IT consulting, and online platform services, which were directly affected by the change in software patent eligibility following *Alice*.

For the firm-level *Alice* panel, we retain unconsolidated operating companies with at least one non-missing observation of total employment in the 2012-2016 event window. For each firm, we extract core identifiers, including `bvdid`, company name, and financial variables used in the firm-level analysis, such as operating revenue, total assets, intangible assets, and reported R&D expenses where available. These ORBIS variables provide scale measures, firm covariates, and denominators for the Lightcast labor-flow rates used in the empirical specifications. ORBIS coverage is primarily annual for the early part of our sample, and we treat ORBIS outcomes as firm-year measures.

To connect the model to the data in the context of the *Alice* shock, we measure each firm's pre-shock exposure using patent information observed through 2014. Our baseline exposure measure is the firm's cumulative pre-*Alice* patent stock in the DISCERN data. In the empirical analysis, pre-period patent stock serves as the exposure variable for the appropriability shock affecting software-intensive incumbents.

Let P_{it}^{soft} denote the number of software-related patents granted to parent firm i (identified by `gvkey`)

in year t. We construct pre-shock exposure as the cumulative stock of such patents through 2014:

$$P_i^{\text{soft}} = \sum_{t=1980}^{2014} P_{it}^{\text{soft}}.$$

In the main specifications, we operationalize exposure as a binary indicator, $\text{HasPatent}_i = 1\left\{P_i^{\text{soft}} > 0\right\}$, equal to one if the firm held at least one software-related patent prior to the *Alice* decision and zero otherwise. This classification captures whether a firm had pre-shock intellectual property exposure in the affected domain and was therefore more likely to be affected by the post-*Alice* weakening of software patent protection.

To measure high-frequency labor adjustments, we use Lightcast employment-history data to construct firm-month labor-flow measures. For each employer entity i and calendar month t (year-month), we define hires as the number of worker-employer spells with a start date in month t, and separations as the number of worker spells that end in month t. These flows are constructed directly from employment spell start and end dates observed in the Lightcast data. We further construct education-tier variants, such as BA and MA+ hires and separations, by restricting worker profiles to those that satisfy the relevant education criteria.

Accordingly, we use one main ORBIS-Lightcast estimation sample and construct patent-based treatment measures within that sample. The estimation sample consists of firms for which total employment is available in ORBIS and labor-flow data, including hires and separations, are available in Lightcast over the relevant event window. This core sample contains 145 public firms observed from 2012 to 2016, yielding 6,528 firm-month observations in the main specifications. All entities in this subset can be matched to a unique Compustat parent identifier (`gvkey`) through our audited ORBIS-Compustat crosswalk. Because `gvkey` identifiers are assigned only to publicly listed firms in Compustat, all firms in the estimation sample are public companies.

We then use DISCERN to construct the patent-based treatment variables for these public firms. Firms with positive software patenting in DISCERN prior to *Alice* are treated as more exposed to the shock, with exposure measured using pre-2014 software patent stock or software patent share. Public firms in the estimation sample that do not appear as patent-active in DISCERN are assigned zero observed patenting in the relevant pre-shock window. Thus, DISCERN does not define a separate estimation sample; it defines the pre-shock patent exposure used to classify treatment intensity within the ORBIS-Lightcast public-firm sample.

For the worker-level career-progression analysis, we use a separate Lightcast full-worker-history extract. The extract begins from the matched *Alice* employer universe and retrieves the full observed job histories of workers who appear in those employers during 2009-2019, including jobs held outside the focal employers. The raw records are job spells with employer names, start and end dates, job titles, locations, and standardized occupation fields. We convert these spells into a worker-year panel. When a worker has multiple observed spells in a year, we retain the spell with the highest O*NET Job Zone; ties are broken using a fixed random seed. O*NET Job Zones are ordered occupation-level measures of the education, experience, and training

typically required for a job, so they provide a coarse measure of occupational preparation rather than compensation or direct output. We define each worker's baseline occupation using the worker's 2013 career area and Job Zone.

4.3 Descriptive Evidence and the Empirical Paradox

Before presenting formal difference-in-differences estimates, we establish the baseline characteristics of our sample and examine the raw, non-parametric trajectories of real-side economic outcomes surrounding the *Alice* decision. Table 3 reports the baseline descriptive statistics and group balancing tests for our analytical sample during the pre-shock window (2011–2013). The balancing diagnostics show no statistically detectable difference in baseline employment levels between treated incumbents and comparison firms, with a raw difference of 902 workers (t-stat = 0.18). This ensures that any post-shock divergence in hiring patterns cannot be driven by an inherent size or scale imbalance in labor trajectories. Furthermore, the two groups are structurally well-balanced across all four core NACE software and IT industry categories, with no statistically significant differences in sub-sector distribution shares.

Table 3: Baseline Summary Statistics and Group Balance (Pre-*Alice* Window)

	Exposed Incumbents (Treatment, $N = 72$)		Comparison Firms (Control, $N = 55$)		Difference (T − C)	
Variable	Mean	SD	Mean	SD	Diff.	t-stat
Firm Real-Side Outcomes						
Employment (Count)	7,957.2	19,873.7	7,055.0	26,655.8	902.2	(0.18)
Total Revenue ($M)	3,378.2	11,319.9	718.7	2,283.5	2,659.6*	(1.86)
Total Assets ($M)	6,113.3	21,013.1	753.0	2,072.5	5,360.3**	(2.08)
R&D Capital Investment ($M)	468.3	1,583.2	18.9	23.8	449.4**	(2.17)
Firm Age (Years)	15.6	14.2	11.4	11.1	4.20*	(1.74)
Industry Distribution (Shares)						
NACE 5829: Software Publishing	0.440	0.496	0.431	0.495	0.009	(0.10)
NACE 6201: Computer Programming	0.365	0.481	0.310	0.463	0.055	(0.61)
NACE 6209: Other IT Services	0.030	0.171	0.095	0.293	-0.065	(-1.29)
NACE 6311: Data Processing & Hosting	0.165	0.371	0.164	0.370	0.001	(0.02)
Unique Firms	72		55		127	
Firm-Year Observations	2,400		1,392		3,792	

Notes: This table reports baseline descriptive statistics and balancing tests calculated during the pre-shock window (2011–2013). "Exposed Incumbents" are defined as software firms with pre-shock patent in 2013. "Comparison Firms" are defined as unexposed software firms with no pre-shock patent in 2013. Financial metrics are scaled to millions of USD ($M). Industry shares represent the proportion of the treatment and control samples falling within each NACE Rev. 2 core 4-digit code. Columns (5) and (6) report the raw difference in means and the corresponding t-statistic from a two-tailed test of equal means. Standard errors are clustered at the firm level (`bvdid`). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A notable feature of the baseline data in Table 3 is the extreme cross-sectional dispersion in innovative scale. For instance, the standard deviation of R&D capital investment among exposed incumbents ($1,583.2M) is nearly triple the sample mean, and a similarly stark skewness characterizes the comparison group. This scale gap is partly explained by structural lifecycle differences: comparison firms are, on average,

younger by roughly 4 years (t-stat = 1.74) and thus heavily concentrated in earlier, less asset-intensive operational phases. While such distributions are standard in innovation settings, they present a major empirical challenge: estimations run in raw levels would be heavily distorted by scale outliers and severe heteroskedasticity. To ensure structural comparability across heterogeneous firms, we apply log transformations to total assets and R&D expenditure.

This operational choice is directly validated by the stable, empirical log-linear relationship between firm inputs documented in Table 4. Consistent with our model's core complementarity assumptions, we estimate the structural mapping of technological inputs using the following specification:

$$\ln(\text{R\&D Investment})_{it} = \alpha_i + \delta_t + \theta_1 \ln(\text{Employment})_{it} + \theta_2 \left[\ln(\text{Employment})_{it} \times \text{Post}_t\right] + \varepsilon_{it} \tag{4.1}$$

where α_i and δ_t represent firm and year fixed effects, respectively. Under this layout, θ_1 isolates the baseline elasticity between input vectors, while θ_2 explicitly captures any post-treatment structural break.

As shown in Columns 1–3 of Table 4, the baseline elasticity θ_1 is tightly pinned down between 0.65 and 0.68 across all firm types. Crucially, the model demonstrates that this relationship is structurally invariant to the regulatory environment: the interaction parameter θ_2 in Columns 4 and 5 is statistically indistinguishable from zero. By demonstrating that the log-linear mapping remains stable before and after the shock, we confirm that our logarithmic transformation successfully linearizes a persistent technological feature of these firms while neutralizing the massive scale variance visible in the summary statistics.

Table 4: Employment and R&D Capital Investment

	Pre-*Alice* Relationship			Post-*Alice* Stability	
	(1) Full Sample	(2) Incumbents	(3) Entrants	(4) Incumbents	(5) Entrants
Log Employment	0.678***	0.679***	0.646***	0.708***	0.865***
	(0.112)	(0.143)	(0.149)	(0.162)	(0.109)
Log Employment × Post				0.019	-0.014
				(0.022)	(0.042)
Firm FE	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y
Observations	614	372	242	769	494
R Squared	0.975	0.974	0.976	0.973	0.978

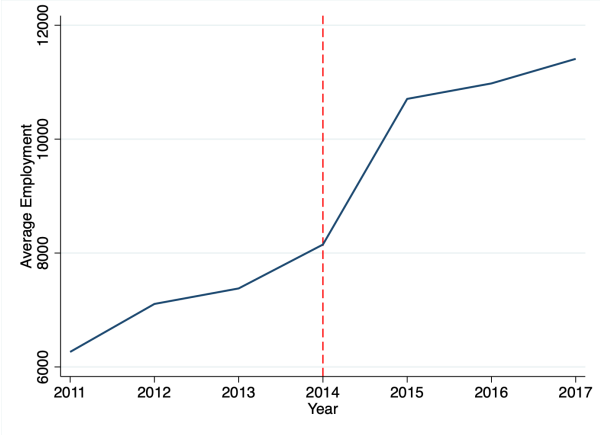
Notes: The dependent variable is log R&D capital investment. All specifications estimate variations of the structural log-linear mapping formalized in Equation (4.1). Columns (1)–(3) restrict the sample to the pre-*Alice* period to isolate the baseline input elasticity (θ_1). Columns (4)–(5) leverage the full longitudinal panel to evaluate post-*Alice* technological stability via the interaction parameter (θ_2). All models include firm and calendar-year fixed effects. Standard errors are robustly clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Standard neoclassical models of the firm suggest that a weakening of property rights acts as a negative wealth or value shock. Under ordinary conditions, such a shock should lead to a contraction in firm scale. However, the raw data present a striking departure from this standard intuition, revealing an empirical

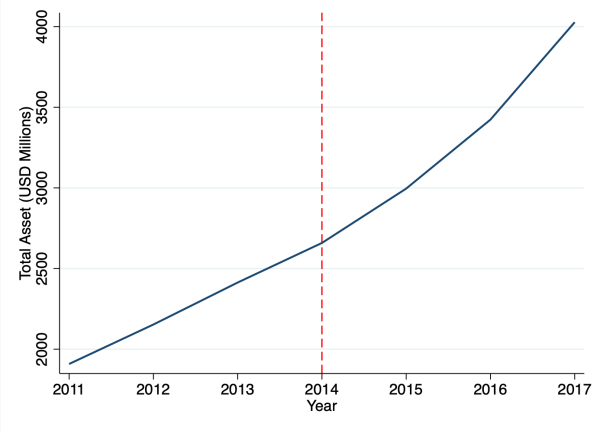
paradox that maps directly to our model's specific comparative statics.

Figure 1 plots the unconditional industry-average levels of employment, total assets, and R&D expenditures across our sample window. Following the 2014 decision, exposed software incumbents exhibit a sharp, visual acceleration in employment. Crucially, this surge is completely isolated: total assets and R&D expenditures follow smooth, uninterrupted trajectories without a comparable discrete shift. This stark divergence between labor and capital accumulation rules out a generic, balanced scale-up explanation and points toward an input-specific strategic reallocation. Instead, this combination provides raw, visual validation of our model's talent-retention mechanism: when patent-based exclusion margins weaken, the strategic asset that prevents displacement shifts from legal barriers to human capital. Incumbents respond not by liquidating or downscaling, but by aggressively expanding or retaining their frontier talent pools inside the firm boundary to manage capability development.

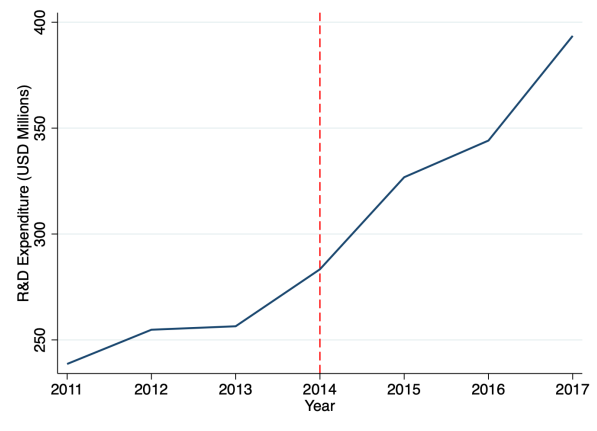
Figure 1: *Alice* Case Descriptive Series: Industry-Average Real-Side Dynamics



(a) Employment (Count)



(b) Total Assets ($M)



(c) R&D Expenditure ($M)

Notes: This figure plots the unconditional, industry-average trajectories of firm employment, total assets, and R&D expenditures across our sample window. Vertical dashed lines mark the 2014 Supreme Court ruling in *Alice Corp. v. CLS Bank International*. Financial metrics are scaled to millions of USD ($M).

As an illustration of the exposure margin, the linked sample includes patent-heavy software incumbents such as Adobe, Microsoft, and Oracle. Adobe is a useful example because its document and workflow businesses, including Acrobat, Reader, LiveCycle, and later Document Cloud services, operated in software domains where firms had historically relied on patents over computerized document handling, forms, signatures, routing, and workflow automation. In our exposure measure, Adobe held 2,044 pre-2014 software-related patents. Its average net-hiring ratio rose from 1.06% before *Alice* to 1.64% afterward, with the change driven mainly by lower separations rather than a surge in new hiring. We view this example as illustrative rather than causal on its own: it helps make concrete the type of patent-heavy incumbent variation that enters the pooled difference-in-differences estimates below.

While these raw descriptive patterns provide clear visual and illustrative support for our theoretical mechanism, they represent unconditional averages that may mask compositional shifts, time-varying industry trends, or firm-specific macroeconomic exposures. To isolate the causal impact of the *Alice* shock and to rigorously evaluate the statistical significance of these divergent margins, we turn to a formal conditional difference-in-differences framework in the following section.

5 Empirical Analyses

5.1 Mapping Theory to the Empirical Design

The stylized model predicts that holding fixed the responsiveness of technological success to labor inputs and the profitability of the new technology, an exogenous increase in the commercialization probability raises the incumbent's incentive to hoard frontier talent, so that net frontier hiring increases one-for-one with hoarding, with all marginal hires reflecting defensive retention rather than expansionary investment.

The empirical setting for testing these predictions is one in which perceived commercialization probability shifts plausibly exogenously while underlying technological success probabilities and profitability remain unchanged. As discussed in Section 4, we study *Alice* as the quasi-experimental shock to the appropriability environment for software and business-method innovations.

5.1.1 Assumptions and empirical diagnostics

The maintained assumptions are plausible in this setting. Because *Alice* is a legal decision about patent eligibility, it should not mechanically change the technological feasibility of software innovation. We therefore interpret the decision as weakening a specific appropriability margin: the ability to use broad, abstract-idea-like software claims to block or deter close substitutes. For firms with larger pre-*Alice* stocks of affected software patents, this weakening increases external displacement risk because rivals and entrants can more easily pursue similar software-based innovations without being constrained by broad patent claims. In the model, a higher risk of displacement raises the strategic value of retaining frontier workers inside the firm boundary. The empirical design therefore compares firms with greater predetermined exposure to affected

software patents to less exposed software firms and asks whether exposed firms exhibit differential post-*Alice* labor responses. This quasi-experiment speaks most directly to the retention/bench channel: it tests whether an increase in external displacement risk leads exposed incumbents to retain frontier workers, and whether this labor response occurs without parallel evidence of productive scale-up in complementary capital.

A related concern is that software markets were also experiencing declining startup experimentation costs around this period, partly due to cloud computing and related digital infrastructure (Ewens et al., 2018; Kalyani et al., 2025). This contemporaneous force is unlikely to mechanically generate our positive incumbent-exposure result. If lower experimentation costs differentially improved the growth and hiring prospects of smaller or less-exposed software firms, they would attenuate the post-*Alice* differential between more exposed incumbents and comparison firms. Therefore, this background trend only makes our estimates on the talent consequences documented below more conservative.

A remaining concern is that any post-*Alice* increase in employment could reflect a coincident rise in expected profitability in software/IT rather than a talent-retention response. Appendix B.1 addresses this concern using decision-date profitability placebos around the June 2014 Supreme Court ruling: (i) industry-portfolio return tests based on FF49 monthly and daily returns and (ii) firm-quarter Compustat tests implemented both in the full Shock 1 public-firm treated/control universe and in the exact main-regression firm sample. Across these exercises, treated firms do not exhibit a positive post-*Alice* re-rating in industry returns, Tobin's Q, market-to-book, or ROA. If anything, the full-sample Compustat estimates are negative, while the exact main-regression sample estimates are statistically indistinguishable from zero.

5.2 Identification Strategy and Empirical Specification

A key design principle is that treatment intensity (exposure) is defined strictly using pre-shock information and is therefore predetermined with respect to post-shock employment dynamics. We construct exposure measures using patent data observed prior to 2014 and then estimate difference-in-differences specifications using employment outcomes before and after the *Alice* decision. This timing structure ensures that exposure classification is not mechanically influenced by post-shock labor adjustments.

We implement one core empirical component: a Lightcast-ORBIS labor-flow difference-in-differences design. We combine Lightcast employment-history data with ORBIS firm-level controls, such as total employment and operating revenue, to construct firm-month measures of hiring, separations, and net headcount flows. Flow variables from Lightcast are scaled by total employment to obtain comparable rate measures across firms. To distinguish a generic employment response from one concentrated among workers more closely tied to frontier innovative capacity, we decompose net employment adjustments into hiring and separation margins and further separate hiring flows by education tier in the enriched Lightcast extract, distinguishing BA from MA+ workers.

In the model, an increase in the perceived probability that new products displace existing technologies raises incumbents' incentives to hoard frontier talent, as skilled workers become more valuable for defending

legacy product lines. In a benchmark setting with two firms, changes in hiring map directly into changes in talent hoarding. However, in a richer multi-firm environment, firms may also hire more in order to attack competitors' weakened legacy technologies, which is a distinct response from hoarding. When firms operate in the same technological space, they face similar opportunities to attack rivals' legacy positions. We therefore assume that, conditional on the shock, the attacking incentive is symmetric across firms operating in the same technology space, while heterogeneous exposure affects only the defensive hoarding aspect. Formally, let the change in firm i's employment following an increase in perceived external displacement risk be decomposed as

$$\Delta\text{Employment}_i = \Delta\text{Employment}_i^{\text{hoard}} + \Delta\text{Employment}^{\text{attack}},$$

where $\Delta\text{Employment}^{\text{attack}}$ captures the common incentive to hire talent in order to exploit competitors' weakened legacy technologies. Under the maintained assumption that $\Delta\text{Employment}^{\text{attack}}$ does not depend on firm-specific exposure, differencing firms with different pre-shock exposure nets out the attacking component:

$$\Delta\text{Employment}_i - \Delta\text{Employment}_j = \Delta\text{Employment}_i^{\text{hoard}} - \Delta\text{Employment}_j^{\text{hoard}}.$$

Thus, cross-sectional variation in pre-shock treatment intensity identifies the differential hoarding response predicted by the model. In the empirical implementation, $\Delta\text{Employment}_i$ is measured as the firm's scaled net hiring flow, defined as new hires minus separations and normalized by total employment. The treatment-intensity measure is constructed strictly from pre-shock patent information: firms with larger software patent stocks, or larger software patent shares, prior to 2014 are treated as more exposed to the *Alice* shock.

Specifically, we interpret heterogeneity in treatment intensity through pre-shock exposure to the appropriability regime most affected by *Alice*. When all firms face the same aggregate shock (i.e., *Alice*), cross-sectional differences in employment responses reflect differences in whether firms hold legacy positions in the affected technological domain. We therefore assume that firms with pre-shock patents in the software domain are exposed to the weakening of appropriability in that domain, while firms without such patents are not directly affected through the changes in the external displacement risk. This identifying assumption does not require any claim about the valuation of individual patents or monotonicity in patent counts; it relies only on the presence of pre-shock patent holdings in the affected domain as an indicator of exposure.

Our baseline specification interacts post-shock timing with a pre-shock exposure measure defined strictly using pre-period information:

$$Y_{it} = \beta\big(\text{Post}_t \times \text{HasPatent}_i\big) + \alpha_i + \lambda_{\text{Year}(t)} + \delta_{\text{Month}(t)} + \varepsilon_{it}. \tag{5.1}$$

where Y_{it} denotes firm i's monthly labor-flow outcome at time t. Our main variable of interest is the net

hiring rate, defined as monthly net hires divided by the firm's corresponding annual total employment:

$$\text{NetHireRate}_{it} = \frac{\text{NewHires}_{it} - \text{Separations}_{it}}{\text{Employment}_{i,\text{Year}(t)}}.$$

To examine the granular margins underlying net employment adjustments, we also estimate the specification separately for the hiring rate and the separation rate:

$$\text{HireRate}_{it} = \frac{\text{NewHires}_{it}}{\text{Employment}_{i,\text{Year}(t)}}, \qquad \text{SeparationRate}_{it} = \frac{\text{Separations}_{it}}{\text{Employment}_{i,\text{Year}(t)}}.$$

The indicator HasPatent_i equals one if firm i held pre-shock patents in the affected technological domain, and Post_t equals one beginning in January 2015, following the USPTO's December 2014 issuance of revised examination guidelines implementing the *Alice* decision. Firm fixed effects α_i absorb time-invariant firm characteristics. Calendar-year fixed effects $\lambda_{\text{Year}}(t)$ control for common macroeconomic shocks, and month-of-year fixed effects $\delta_{\text{Month}(t)}$ account for seasonality in labor flows.

Because larger firms may mechanically exhibit different growth patterns over time, we augment the baseline by allowing flexible year-specific trends by initial firm size. Specifically, we interact revenue measured in the first year of the estimation window with calendar-year indicators:

$$Y_{it} = \beta\big(\text{Post}_t \times \text{HasPatent}_i\big) + \sum_y \gamma_y\big(\text{Revenue}_{i,t_0} \times \mathbb{1}\{\text{Year}(t) = y\}\big) + \alpha_i + \lambda_{\text{Year}(t)} + \delta_{\text{Month}(t)} + \varepsilon_{it}. \quad (5.2)$$

The interaction between initial revenue and calendar-year indicators permits firms of different initial sizes to follow flexible year-specific growth trajectories, ensuring that the estimated treatment effect is not driven by differential expansion patterns among larger firms. Unless otherwise noted, standard errors are clustered at the firm (parent `gvkey`) level.

The model's comparative static applies most directly to net employment adjustment. An increase in perceived external displacement risk strengthens the incumbent's incentive to hoard frontier talent, which in equilibrium implies a higher net hiring rate among firms exposed to the threatened technological domain. Accordingly, the core empirical prediction is that firms with greater pre-shock exposure should experience a larger post-shock increase in net hiring. We therefore focus the baseline tests on the net hiring response. The hiring and separation rates are reported to decompose that net response, not because the model delivers independent comparative statics for each gross-flow component. Likewise, the design does not directly observe worker-by-worker under-utilization; the benching interpretation comes from combining the retention-based flow decomposition with the absence of parallel observable scale-up in capital proxies.

5.3 Fact 1: Talent Hoarding

Table 5 reports difference-in-differences estimates using the net hiring ratio as the dependent variable, presenting the baseline framework from Equation (5.1) in Columns (1)–(3) and the controlled specification from

Equation (5.2) in Columns (4)–(6). The coefficient of interest, β, captures the differential margin of labor adjustment for patent-exposed firms relative to comparison firms following the *Alice* legal shock. Because our monthly panel spans a symmetric, high-frequency window covering approximately 24 months before and 24 months after the shock event, plotting a full set of individual monthly leads and lags within the main text would be visually redundant and statistically noisy. Instead, we present the consolidated difference-in-differences parameters as our primary estimates and relegate the comprehensive verification of the parallel trends assumption to Appendix B.2. There, we report formal joint F-tests evaluating whether all monthly pre-treatment interaction coefficients are simultaneously equal to zero, providing rigorous statistical support for our baseline identifying assumptions.

Table 5: The *Alice* Legal Shock and Net Hiring Ratios

Dependent Variable:	*Net Hiring Ratio*					
	Baseline Specification			+ Revenue × Year Controls		
	(1)	(2)	(3)	(4)	(5)	(6)
Sample Split:	Full	BA	MA & Above	Full	BA	MA & Above
Post-*Alice* × Exposure	0.0036***	0.0001	0.0023*	0.0034**	-0.0000	0.0023*
	(0.0013)	(0.0014)	(0.0013)	(0.0013)	(0.0014)	(0.0012)
Firm Fixed Effects	Y	Y	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y	Y	Y
Year Fixed Effects	Y	Y	Y	N	N	N
Revenue × Year Controls	N	N	N	Y	Y	Y
Observations	6,528	1,851	1,851	6,528	1,851	1,851
R-squared	0.0955	0.0662	0.0484	0.0956	0.0662	0.0484

Notes: This table reports difference-in-differences estimates corresponding to the empirical specification outlined in Equation (5.1) and Equation (5.2). The dependent variable is the net hiring ratio constructed from ORBIS employment and Lightcast labor flows. Worker subgroups are partitioned by highest educational attainment (BA or MA and above) at the time of hire. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Across the full-sample specifications, the estimated effect is positive and statistically significant. In the specification with flexible year-specific trends by baseline firm size (Column 4), exposed firms increase monthly net hiring by about 0.34 percentage points relative to unexposed firms after the *Alice* implementation window. Evaluated at the average pre-period employment of exposed firms in the core sample, this implies retention of a few dozen workers per month, cumulating to several hundred additional retained workers over a two-year post period. Relative to the 0.54 percentage-point mean net hiring ratio in the software sample, the estimate is about 63% of the sample mean, suggesting an economically meaningful effect. Its stability after introducing baseline revenue-year interactions further indicates that the result is not driven by differential expansion among larger firms, but instead reflects exposure-related adjustments consistent with the model.

This estimate captures the increase in hoarding induced by higher perceived external displacement risk. Because exposure is predetermined and the shock is common across firms, the interaction term identifies the cross-sectional response: exposed firms expand net employment relative to unexposed firms. The positive sign matches the model's comparative statics, in which greater external displacement risk leads incumbents

to retain and accumulate frontier talent to protect legacy lines.

Columns (2)-(3) and (5)-(6) decompose this response by education tier. The effect is concentrated among workers with MA degrees and above, with no statistically meaningful change among BA workers. The MA+ coefficient remains positive and significant across specifications. This pattern is consistent with the proposed mechanism: hoarding should be strongest in labor-market segments where talent is both scarce and strategically important for innovation competition.

Supplementary composition evidence from the enriched Lightcast extract points in the same direction. When we classify hires using Lightcast's technical and research occupation buckets, the post-*Alice* differential is concentrated in frontier-oriented categories rather than in broad nontechnical hiring. We view these occupation-based estimates as supportive rather than stand-alone because they are estimated on a smaller and more selected subsample, but they sharpen the interpretation of the MA+ result: the response is not simply more hiring of educated workers in general, but a tilt toward technical and research labor categories most closely tied to frontier innovative capacity.

To understand the source of the net hiring increase, Tables 6 and 7 report results using the new hiring ratio and the separation ratio as dependent variables. This decomposition is important for interpretation: the economically large net-hiring effect in Table 5 does not reflect a surge in gross recruitment, but instead arises almost entirely from lower worker outflows among exposed firms. The new hiring ratio exhibits no statistically significant post-shock change for exposed firms. In contrast, the separation ratio declines significantly. The magnitude of the decline closely mirrors the increase in net hiring, indicating that the net employment effect is driven primarily by reduced separations rather than increased gross hiring. Specifically, Lightcast separations measure the end of a worker-firm pair spell, and the decline in separations thus means a reduced outflow of incumbent workers. The lower outflow aligns well with the retention/hoarding mechanism in our framework: exposed firms may (i) reduce voluntary exits via retention pay, internal career options, and mobility frictions, and/or (ii) avoid involuntary separations by redeploying workers internally, including into short-run bench positions when complementary inputs or deployable frontier projects adjust slowly. Because the identifying variation compares exposed and unexposed firms while absorbing common time shocks, the central finding is that exposed firms keep more frontier workers inside the firm boundary after *Alice*. The interpretation does not require separately classifying the decline in separations into fewer quits versus fewer layoffs.

This decomposition suggests that talent retention, rather than aggressive new recruitment, is the main channel of adjustment. A natural interpretation is that when firms foresee their current products more likely to be displaced by external innovations, the strategic cost of losing employees who possess firm-specific knowledge and business secrets increases, as their departure may facilitate imitation or accelerate rival innovation of similar technologies. In this sense, incumbent workers are not merely units of labor; they are carriers of proprietary knowledge that can weaken the firm's legacy position if reallocated to competitors. Consequently, reducing separations becomes a particularly effective way to mitigate external displacement

Table 6: The *Alice* Legal Shock and New Hiring Ratios

Dependent Variable:	*New Hiring Ratio*					
	Baseline Specification			+ Revenue × Year Controls		
	(1)	(2)	(3)	(4)	(5)	(6)
Sample Split:	Full	BA	MA & Above	Full	BA	MA & Above
Post-*Alice* × Exposure	0.0003	-0.0001	-0.0003	0.0003	-0.0001	-0.0003
	(0.0010)	(0.0008)	(0.0014)	(0.0010)	(0.0008)	(0.0014)
Firm Fixed Effects	Y	Y	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y	Y	Y
Year Fixed Effects	Y	Y	Y	N	N	N
Revenue × Year Controls	N	N	N	Y	Y	Y
Observations	6,528	1,851	1,851	6,528	1,851	1,851
R-squared	0.2728	0.2751	0.0477	0.2729	0.2752	0.0478

Notes: This table reports difference-in-differences estimates corresponding to the empirical specification outlined in Equation (5.1) and Equation (5.2). The dependent variable is the new hiring ratio constructed from ORBIS employment data and Lightcast labor flows. Worker subgroups are partitioned by highest educational attainment (BA or MA and above) at the time of hire. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 7: The *Alice* Legal Shock and Separation Ratios

Dependent Variable:	*Separation Ratio*					
	Baseline Specification			+ Revenue × Year Controls		
	(1)	(2)	(3)	(4)	(5)	(6)
Sample Split:	Full	BA	MA & Above	Full	BA	MA & Above
Post-*Alice* × Exposure	-0.0031***	-0.0001	-0.0025**	-0.0031**	-0.0001	-0.0025**
	(0.0012)	(0.0008)	(0.0012)	(0.0012)	(0.0009)	(0.0012)
Firm Fixed Effects	Y	Y	Y	Y	Y	Y
Month Fixed Effects	Y	Y	Y	Y	Y	Y
Year Fixed Effects	Y	Y	Y	N	N	N
Revenue × Year Controls	N	N	N	Y	Y	Y
Observations	6,528	1,851	1,851	6,528	1,851	1,851
R-squared	0.3384	0.0291	0.0025	0.3384	0.0269	0.0002

Notes: This table reports difference-in-differences estimates corresponding to the empirical specification outlined in Equation (5.1) and Equation (5.2). The dependent variable is the separation ratio constructed from ORBIS employment data and Lightcast labor flows. Separations reflect the structural conclusion of a worker-firm spell and do not distinguish between voluntary quits and involuntary layoffs. Worker subgroups are partitioned by highest educational attainment (BA or MA and above) at the time of hire. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

risk. Retention protects both innovative capacity and the informational boundaries of the legacy technology, whereas marginal new hiring does not provide the same defensive benefit.

This retention pattern is also consistent with a firm-specific knowledge channel. When formal exclusion weakens, incumbents may place greater weight on protecting tacit implementation knowledge, deployment know-how, internal applications, and complementary trade secrets embodied in incumbent workers. In that case, the relevant defensive response is not broad-based recruiting, but reducing separations among workers who already sit at the interface of legacy systems and frontier development. This interpretation helps explain why the response appears primarily on the reduced separations rather than through a large increase in gross hiring.

The evidence here presents a coherent picture: following an increase in perceived displacement risk, exposed firms increase net employment primarily by reducing separations rather than by expanding gross hiring, and the response is concentrated among highly educated workers. We interpret this as a retention-based hoarding response. The *Alice* design supports the claim that incumbents can respond to weaker appropriability by keeping scarce frontier labor inside the firm boundary, even without commensurate evidence of productive scale-up. Appendix B.3 shows that the main results are robust to alternative definitions of patent exposure based on pre-*Alice* patent-stock thresholds.

5.4 Fact 2: Bench Position? Retention Without Proportional Scale-Up

The benching implication of the model relies on complementarity between frontier labor and complementary capital. Empirically, however, the capital result is not intended to prove the literal idleness of workers. Rather, it serves as a scale-up diagnostic: the capital stock did not deepen proportionally to support the additional employees. If the post-*Alice* labor response primarily reflected productive expansion of frontier activity, one would expect some parallel expansion in complementary scale proxies, such as assets, R&D-supporting expenditure, or other observable investment measures. In the software industry, much of the relevant capital adjustment would plausibly occur through purchases of licenses, computing resources such as GPUs, and related equipment, which can typically be scaled within a few months. Instead, we observe a labor response without corresponding evidence of differential capital scale-up. That contrast is especially informative here because the labor response is not only statistically significant but also economically large relative to the baseline net hiring rate, whereas the capital estimates are small and statistically indistinguishable from zero. This pattern is consistent with a retention or benching response, while still allowing that some retained workers may be redeployed internally rather than fully idle.

This distinction yields a clear empirical implication. If the post-*Alice* increase in employment among exposed firms reflected broad expansionary investment, one would expect parallel increases in capital or scale measures. Conversely, if employment rises without corresponding capital growth, the evidence is more consistent with firms retaining frontier workers faster than they expand the complementary inputs needed to deploy them productively.

Table 8 reports estimates of Equation (5.1) and Equation (5.2) using total assets (as well as other capital measures such as R&D capital, see the Appendix) as the dependent variable. Columns (1)–(2) focus on the 2012–2016 event window, while Columns (3)–(4) extend the sample to 2009–2019. Across specifications, the interaction coefficient is economically small and statistically indistinguishable from zero once flexible size trends are included. In particular, when allowing for baseline employment × year interactions, the estimated effect on total assets is close to zero in both sample windows.

Table 8: The *Alice* Legal Shock and Total Expenditures

Dependent Variable:	*Total Asset ($M)*			
	2012–2016 Window		2009–2019 Window	
	(1)	(2)	(3)	(4)
Post-*Alice* × Exposure	1,600.00	-79.40	3,870.00	-196.00
	(969.00)	(214.00)	(2,500.00)	(485.00)
Firm Fixed Effects	Y	Y	Y	Y
Year Fixed Effects	Y	N	Y	N
Revenue × Year Controls	N	Y	N	Y
Observations	708	707	1,523	1,520
R-squared	0.970	0.998	0.843	0.987

Notes: This table reports difference-in-differences estimates corresponding to the empirical specification outlined in Equation (5.1) and Equation (5.2). The dependent variable is firm total asset extracted from ORBIS and scaled to millions of USD ($M). Columns (1)–(2) limit the panel sample to a tight 2012–2016 window, while Columns (3)–(4) expand the horizon to a long 2009–2019 window. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Table 9 provides ratio-based scale-up diagnostics. Though our findings are identical using either the native R&D expense ratio or reconstructed log expenditures, since there exists a subset of observations contains valid ratios but missing revenue components, we report the raw ratio to maximize sample coverage. Across these specifications, R&D intensity rises imprecisely, while log total assets show no statistically meaningful increase across windows. This pattern reinforces our primary capital result: the post-*Alice* labor response is not accompanied by a robust, observable scale-up in a firm's asset base or innovation capacity.

Taken together with the labor-flow results, the capital evidence shows that the observed employment response among exposed firms is not accompanied by parallel observable scale-up. The pattern is consistent with a setting in which firms retain or absorb frontier labor faster than they expand the complementary inputs needed for full deployment, because the lack of differential capital support appears in both the shorter 2012-2016 window and the longer 2009-2019 window, in an industry where relevant complementary inputs such as cloud compute, API-based services, software licenses, and related equipment can often be scaled within just days and weeks.

The worker-history evidence below moves the test from firm scale to individual career progression. If the retained research workers were being absorbed into expanding frontier projects, their occupational advancement should have remained comparable to that of similar workers attached to less exposed firms. Instead, exposed science-and-research workers exhibit slower upward career progression after *Alice*. The worker-level

Table 9: The *Alice* Legal Shock, R&D Expense Ratios, and log(Total Assets)

Dependent Variable:	*R&D Expense Ratio*			*log(Total Assets)*		
	(1)	(2)	(3)	(4)	(5)	(6)
Time Window	2012–2016	2012–2019	2009–2019	2012–2016	2012–2019	2009–2019
Post-*Alice* × Exposure	2.1160	2.7920*	3.2320	0.0495	0.0286	0.3100
	(1.3320)	(1.6070)	(2.4050)	(0.1510)	(0.1850)	(0.2860)
Dependent Variable Mean	18.75	19.39	18.62	—	—	—
Firm Fixed Effects	Y	Y	Y	Y	Y	Y
Employment × Year Controls	Y	Y	Y	Y	Y	Y
Observations	377	577	731	491	753	992
R-squared	0.885	0.864	0.767	0.981	0.971	0.960

Notes: This table reports difference-in-differences estimates corresponding to the empirical specification outlined in Equation (5.1) and Equation (5.2). Columns (1)–(3) use the R&D expense ratio (measured in percentage points) as the dependent variable, alongside its sample mean. Columns (4)–(6) use the natural logarithm of total assets from ORBIS. Standard errors are clustered at the firm level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

evidence therefore sharpens Fact 2: the retained labor pattern is accompanied by slower upward progression among research workers, consistent with limited deployment into higher-responsibility innovation roles.

5.5 Fact 3: The Effect of Hoarding on Workers' Future Career

The firm-level evidence shows that exposed firms retain more workers following *Alice* without proportionate increases in observable capital investment. We next examine whether this increased retention is associated with changes in workers' occupational trajectories. This distinction is important because a productive-expansion interpretation predicts that retained workers should continue advancing into more senior occupations as firms scale frontier projects. By contrast, a talent-hoarding interpretation allows firms to retain labor even when workers experience weaker occupational advancement.

To study worker career dynamics, we construct a worker-year panel from the full Lightcast employment histories. When a worker holds multiple jobs within the same year, we retain the occupation with the highest occupational rank measure; ties are broken randomly. The sample contains approximately 9.9 million worker-year observations spanning 1.4 million workers employed at least once by one of the 173 sample firms between 2009 and 2019. Our primary outcome is O*NET Job Zone, an occupation-level measure of required preparation, skill intensity, and seniority ranging from 2 to 5. Within a given occupational track, higher Job Zones correspond to more senior occupations. We focus on comparisons within narrowly defined 2013 career areas to avoid comparing occupational ladders across fundamentally different fields, since Job Zones are not directly comparable across broad occupational categories.

We begin by estimating event-study specifications to trace the dynamic evolution of workers' occupational trajectories around *Alice*:

$$Y_{ijt} = \sum_{k \neq -1} \beta_k \left(1\{t = k\} \times Treated_i\right) + \alpha_i + \gamma_j + \varepsilon_{ijt}, \tag{5.3}$$

where Y_{ijt} denotes the Job Zone of worker i employed at firm j in year t. $Treated_i$ equals one if worker i was employed at an incumbent patent-holding firm at the time of *Alice*. The specification includes worker fixed effects α_i and firm fixed effects γ_j, so identification comes from within-worker changes in occupational trajectories relative to comparable workers employed at non-incumbent firms within the same career area.

Table 10 summarizes the heterogeneous occupational patterns implied by the event-study estimates. For brevity, we report only a qualitative classification of the estimated responses across occupations, while the full set of occupation-specific event-study figures is reported in Appendix B.4.

Table 10: Summary of Event Study Trajectories Across Occupations

Observed Pattern	**Occupation / Career Area**
Strongest Negative Response	Science and Research; Education and Training; Law, Compliance, and Public Safety
Moderate Negative Response	Manufacturing and Production; Information Technology and Computer Science
Negligible Response	Customer and Client Support; Human Resources
Not Causally Interpretable	Design, Media, and Writing; Healthcare; Finance; Business Management and Operations

Notes: This table synthesizes the heterogeneous post-treatment patterns observed across occupation-specific event study estimations. Groupings are categorized by the presence of a sharp, post-2014 break from baseline trends, the statistical significance of post-treatment coefficients, and the validity of parallel pre-trends prior to the *Alice* legal shock.

These patterns are consistent with a strong correlation between the types of workers most likely to be strategically retained after *Alice* and those experiencing weaker long-run occupational advancement. In particular, the largest negative effects appear in Science and Research; Education and Training; and Law, Compliance, and Public Safety occupations. These are precisely the categories in which workers are either closely tied to innovative activity, responsible for preserving and transmitting firm-specific knowledge, or directly exposed to the institutional value of intellectual property. The results, therefore, suggest that the workers most likely to be hoarded are also the ones whose subsequent career progression slows the most.

The negative effects in Science and Research occupations align naturally with the talent-hoarding interpretation. Earlier results showed that exposed firms retained more workers primarily through lower separations rather than increased hiring, without proportional increases in observable capital investment. This pattern is consistent with firms preserving existing innovative capacity and accumulated expertise even when opportunities for frontier expansion weakened. Under this mechanism, research-oriented workers remain employed but experience slower occupational advancement over time.

The even stronger effects observed in Education and Training occupations are consistent with the idea that incumbent firms placed greater value on workers who possessed and transmitted accumulated firm-specific expertise after *Alice*. Because the employment response operated mainly through retaining incumbent

workers rather than expanding external hiring, firms may have relied more heavily on employees responsible for onboarding, mentoring, technical training, and the diffusion of specialized internal knowledge. In software-intensive industries, these workers often possess substantial tacit knowledge themselves, including familiarity with proprietary systems, workflows, coding practices, and accumulated organizational know-how. The strong negative effects for these occupations therefore further support the idea that workers most valuable for preserving incumbent-specific capabilities were also among those experiencing weaker long-run occupational advancement.

For Law, Compliance, and Public Safety occupations, the negative effects likely reflect a related but distinct channel. By weakening the scope of software patent protection, *Alice* plausibly reduced the long-run demand for certain forms of patent-related legal expertise and compliance activity. The weaker occupational progression observed among legal workers is therefore consistent with a decline in the economic value of patent-intensive legal functions following the decision.

The moderate negative effects observed in Manufacturing and Production, Finance, and Information Technology and Computer Science are likewise consistent with a broader correlation between the extent to which workers appear likely to have been strategically retained after *Alice* and the magnitude of their subsequent decline in occupational progression. Relative to Science and Research, Education and Training, or Law occupations, workers in these categories may have been somewhat less central to the preservation of incumbent-specific capabilities or less intensively targeted by retention incentives. Correspondingly, the estimated declines in Job Zone progression are smaller, consistent with a gradient in which occupations more strongly associated with strategic labor retention exhibit larger long-run career effects.

By contrast, we find little evidence of meaningful effects in Customer and Client Support; Hospitality, Food, and Tourism; and Human Resources. These occupations are less plausibly associated with the types of workers firms would strategically retain in response to weaker appropriability incentives or concerns about preserving specialized incumbent knowledge. The absence of substantial declines in occupational progression in these categories therefore further supports the broader correlation between the likelihood of being hoarded and the magnitude of subsequent career stagnation.

Finally, several occupational groups exhibit imprecise or difficult-to-interpret estimates, including Design, Media, and Writing; Healthcare; Maintenance, Repair, and Installation; and Business Management and Operations occupations. Because these occupations may reflect heterogeneous forces unrelated to strategic labor retention or may have ambiguous relationships to incumbent-specific knowledge, we avoid drawing strong conclusions from these estimates and report them primarily for completeness.

We also estimate the magnitude of these effects more parsimoniously with the following difference-in-differences specification:

$$Y_{ijt} = \beta \left(Post_t \times Treated_i \right) + \alpha_i + \gamma_j + \varepsilon_{ijt}, \tag{5.4}$$

where $Post_t$ is an indicator for the post-*Alice* period. The coefficient of interest, β, captures the average post-shock change in occupational rank among exposed workers relative to comparable control workers within

the same career area.

Table 11: Occupations Exhibiting Strong Negative Effects

Sample Split:	Science & Research	Education & Training	Law / Public Safety
	(1)	(2)	(3)
Post-*Alice* × Exposure	-0.018**	-0.068***	-0.041***
	(0.008)	(0.012)	(0.010)
Worker Fixed Effects	Y	Y	Y
Firm Fixed Effects	Y	Y	Y
Observations	147,489	131,365	109,655
R-squared	0.693	0.775	0.849

Notes: This table reports difference-in-differences estimates corresponding to the empirical specification outlined in Equation (5.4), restricted to individual occupational sub-samples. The dependent variable O*NET Job Zones capture structural shifts in job seniority. Standard errors are clustered at the worker level. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

This specification complements the event-study estimates in two ways. First, it collapses the dynamic post-*Alice* responses into a single average treatment effect, providing a parsimonious summary of the magnitude of the career-progression effects documented above. Second, because the outcome is measured in Job Zone units, the coefficient has a direct interpretation as the average change in workers' occupational rank within a career area. A negative coefficient therefore indicates that exposed workers experienced slower movement into occupations requiring greater preparation, skill intensity, or seniority relative to comparable workers at non-incumbent firms.

Table 11 reports the estimates for the occupational groups with the clearest negative responses in the event-study analysis. Across these occupations, the estimates are negative and statistically significant, indicating that exposed workers experienced weaker upward occupational progression after *Alice*. Because Job Zone is measured on a discrete 2-5 scale, the coefficients can be read as average changes in occupational seniority measured in Job Zone units. For Science and Research workers, the estimate of -0.018 implies that exposed workers' post-*Alice* occupational rank was lower by 0.018 Job Zone points relative to comparable control workers. This is equivalent to about 1.8 percent of a full one-zone movement.

Put differently, these estimates mean that exposure to *Alice* reduced average occupational progression by a small but economically interpretable fraction of a move up the Job Zone ladder. Since workers do not move fractionally across occupations, these coefficients should be understood as changes in the average probability or intensity of moving into higher-Job-Zone occupations. For example, the Education and Training estimate is consistent with about 6.8 fewer workers per 100 moving up by one Job Zone, or an equivalent combination of smaller changes across the distribution of occupational transitions.

Appendix B.5 reports the analogous difference-in-differences estimates for the remaining occupational groups. These estimates corroborate the event-study evidence: occupations more plausibly associated with strategic retention exhibit negative average post-*Alice* effects, whereas occupations less closely connected to this mechanism show smaller or less systematic changes in Job Zone progression.

5.6 Ruling out Alternative Explanations from the Results

In the *Alice* design we have found that net hiring rises, gross hiring does not, separations fall, observable capital proxies do not differentially expand, and the workers in occupations most plausibly tied to innovation, firm-specific knowledge transmission, and patent-related institutional expertise subsequently experience weaker occupational progression. That combination is not the pattern one would naturally expect from screening or pipeline stories, from broad productive scale-up, or from simple managerial empire building. The labor-flow decomposition in Tables 6 and 7, a capital response too small to equip the retained workforce in Table 8, and the worker-history evidence in Tables 10 and 11 jointly help distinguish these channels.

Screening/up-or-out. In screening stories, firms expand the pool of workers to learn match quality and then reallocate via higher turnover. This predicts changes in gross flows—either higher separations (as low matches are culled) and/or higher hiring (to replenish the pool)—rather than a net headcount increase driven almost one-for-one by reduced separations. But we do not see the evidence of intensified screening in our data- the new-hire ratio is economically and statistically close to zero, while separations fall by a magnitude that mirrors the net-hiring increase, especially for MA+ workers.

Productive substitution toward labor. One alternative is that weaker patent protection raises the value of active engineering effort, so exposed firms productively substitute toward labor in order to innovate faster or stay ahead of imitation. A pure productive-substitution account would typically predict either broader scale-up in complementary inputs or a meaningful increase in gross hiring to expand active project capacity. Our evidence instead shows essentially no gross-hiring response, a separation-driven increase in net hiring, and no parallel evidence of capital expansion. That pattern is more consistent with defensive retention than with a broad expansion of productive frontier activity.

Hiring talent as part of the development pipeline. An alternative view is that exposed firms retain or hire high-skill workers to develop them for future frontier projects. Such a pipeline mechanism would naturally predict continued movement into higher-preparation roles as workers are matched to expanding projects, and often greater churn as firms rotate workers through assignments and select a final lineup. However, the data suggests that the net employment increase comes from reduced separations rather than gross hiring, capital does not expand proportionally, and the full worker histories show that exposed science-and-research workers become less likely to move above their 2013 Job Zone after *Alice*. This pattern aligns more closely with retention and limited advancement than with an expanding training-and-promotion pipeline.

Managerial empire building. Another theory is that managers expand headcount to increase span or internal influence. Our identification relies on an external legal shock (Alice) and heterogeneity in predetermined exposure, and we do not observe a positive treated-versus-control profitability or valuation re-rating

around Alice in Compustat placebos. Combined with the absence of capital deepening, this makes a pure "empire building" interpretation Haegele (2024) less consistent with the timing and margins of adjustment than a defensive response tied to external displacement risk and the value of keeping frontier talent inside the boundary.

6 Limitations and Future Research

Our baseline deliberately keeps talent, research lines, and complementary inputs simple. Researchers differ in ability, fit, learning trajectories, outside options, risk preferences, and redeployability across projects. Research programs also differ in uncertainty, modularity, and fit with existing assets. The one-for-one complementarity between frontier labor and innovation capital is meant to make the retention-deployment margin clear. Relaxing these assumptions would move the magnitudes and the contract design –workers differ in ability, fit, outside options, and how readily they redeploy across projects; teams produce jointly; retained know-how and within-firm spillovers can make even partial retention productive. None of this removes the wedge on its own; it survives unless retained workers also receive the project authority, capital, data, compute, and managerial attention that active discovery requires.

The empirical design speaks most clearly to public, patent-linked software firms around a particular legal shock. *Alice* weakened a specific appropriability margin for abstract-idea-like software claims, and the estimates use predetermined patent-based exposure to the shock to compare firms more and less affected by that change. This setting is useful because the shock is external to individual staffing plans and maps naturally to external displacement risk, but it should not be read as a claim that all large-firm hiring in frontier markets reflects hoarding. The mechanism should be most relevant where frontier workers are scarce, worker knowledge is firm- or technology-specific, and complementary inputs or project slots adjust more slowly than retention decisions. In addition, We observe labor flows, education composition, observable capital, and later occupational trajectories. We do not observe project rosters, worker-time allocation, retention contracts, internal assignment records, or the terms limiting worker mobility. Matched employer–worker–project data would allow a cleaner separation between capability preservation and defensive retention. It would also let us test which contracts keep researchers connected to active discovery, which merely keep them inside the firm, and when retained teams are redeployed, spun out, or left without frontier work.

The model also abstracts from richer payoff linkages between the legacy product and the replacement technology. In the main comparative statics, *Alice* is interpreted as shifting the commercialization probability rather than directly shifting the incumbent's payoff from self-replacement. A more general model could replace the reduced-form legacy and replacement payoffs with a joint-profit function that allows the new line to substitute for, complement, or be independent of the legacy line. Future work could also add bargaining in idea markets, where a rival's successful discovery gives it leverage through licensing, acquisition, litigation, or settlement. These channels would raise the value of keeping scarce frontier workers inside the firm, but they

are not required for the main prediction that higher external displacement risk raises retention incentives.

7 Conclusion

This paper studies how large incumbents shape upstream innovation through the allocation of scarce frontier research talent. The central mechanism is a retention-deployment wedge. Inside the incumbent, the worker can raise the firm's own chance of innovation. Outside it, the same worker may help a rival develop a replacement technology. When outside discovery threatens legacy rents, the incumbent's private value from retaining the worker can exceed that worker's immediate internal R&D value. If the complements – capital, data, compute, project slots, and managerial attention – cannot scale at the same pace, retention can rise faster than the firm's ability to deploy the talent it holds.

The model separates retention from deployment and clarifies why incumbent labor-market responses matter for innovation even before products, patents, or acquisitions appear. The relevant strategic choice is control over a scarce discovery input. Retention reduces rival discovery probabilities, while deployment determines how much retained talent becomes effective innovation input inside the incumbent. This distinction generates the possibility of benching and upstream misallocation: frontier workers can remain employed by the incumbent while contributing less to active discovery and diffusion than they would in alternative specialized matches.

The empirical analysis examines this mechanism around *Alice v. CLS Bank*. The decision weakened the expected exclusion value of broad abstract-idea-like software claims, raising external displacement risk for firms with greater pre-shock exposure to affected software patents. Exposed firms increase net hiring after *Alice* relative to less exposed software firms. The effect is economically meaningful, about 0.34 percentage points per month, or roughly 3.4 additional net hires per 1,000 employees. The decomposition is central: the increase is driven primarily by reduced separations rather than gross hiring, is concentrated among workers with advanced degrees, and occurs without proportional capital expansion. Worker career histories provide a second diagnostic. Science-and-research workers attached to exposed incumbents experience slower upward occupational progression after *Alice*, with the gap reaching roughly 0.06–0.07 Job-Zone points by the fifth post-shock year.

Taken together, the evidence supports a retention-based view of incumbent response to external displacement risk. When formal exclusion weakens, exposed incumbents appear to place greater value on keeping frontier workers and firm-specific knowledge inside the boundary. The labor response does not look like a broad productive scale-up: it is concentrated in separations, concentrated among high-skill workers, accompanied by weak capital adjustment, and followed by flatter career trajectories for research workers. These patterns are consistent with the model's prediction that control over scarce labor can become a substitute for more direct forms of exclusion.

While there's no proposed legal rule, we provide significant measurement implications for policy. Inno-

vation and competition policy usually look for visible assets – patents, products, acquisitions, exclusionary contracts, or product-market conduct. This paper points to an earlier margin: scarce researchers can be kept from rival discovery even before a rival technology exists. Deal review therefore does not exhaust an incumbent's influence over frontier research capacity.

The relevant records are project rosters, internal assignments, worker-time allocation, retention packages, mobility restrictions, and links between retained researchers and active frontier projects. Universities and public funders could also track whether publicly trained researchers move into roles with frontier project responsibility, publication opportunities, and mobility. Firms face the same distinction internally: broad hiring builds options when projects, inputs, and decision rights grow with headcount; it creates a bench when retention rises without capital, responsibility, or advancement. Paying a researcher, assigning that researcher to frontier work, and helping that researcher advance are separate outcomes.

Appendix

Contents

A Model Proofs **2**

B Additional Robustness Checks for the Comparative Statics Test **6**

B.1 Ex Ante Profitability Placebo Tests 6

B.1.1 Market-based placebo: industry portfolio excess returns (FF49) 6

B.1.2 Firm-sample placebo: valuation ratios and profitability in the Shock 1 GVKEY universe 8

B.2 Robustness of the Parallel Trends Assumption: High-Frequency Joint Tests 12

B.3 Robustness to Alternative Patent-Stock Thresholds 14

B.4 Occupation-Specific Event-Study Estimates 14

B.5 Occupation-Specific Difference-in-Differences Estimates 17

A Model Proofs

Proof A.1 (of Lemma 1) *Conditional on project b, the entrant chooses $\ell_E \in [0, m - \ell_I]$ to maximize*

$$S_E^b(\ell_E, x, Q) = Q\pi_b\big(1 - p(x)\big)p(\ell_E) - r\ell_E.$$

The derivative with respect to ℓ_E is

$$Q\pi_b\big(1 - p(x)\big)p'(\ell_E) - r.$$

By Assumption 2, this derivative is strictly positive over the relevant feasible range. Hence the entrant's surplus is strictly increasing in ℓ_E, and the optimal choice is $\ell_E^*(\ell_I) = m - \ell_I$. □

Proof A.2 (of Proposition 1) *Differentiate $G(\ell_I; x, Q)$ with respect to ℓ_I:*

$$\frac{\partial G(\ell_I; x, Q)}{\partial \ell_I} = -Q\pi_b\big(1 - p(x)\big)p'(m - \ell_I) + r.$$

Assumption 2 implies this derivative is strictly negative. Therefore G is strictly decreasing in ℓ_I.

Since G is continuous and strictly decreasing, the set

$$\{\ell \in [0, m] : G(\ell; x, Q) \le 0\}$$

is a closed interval whenever it is nonempty. It is nonempty because

$$G(m; x, Q) = -\pi_a \le 0.$$

Hence $\ell_I^a(x, Q)$ is well defined and unique. The entrant chooses b exactly when $G(\ell_I; x, Q) > 0$, and chooses a when $G(\ell_I; x, Q) \le 0$, using the tie-breaking rule.

If the threshold is interior, it satisfies

$$G(\ell_I^a(x, Q); x, Q) = 0,$$

which is (3.10)*. Let*

$$D(x, Q) = Q\pi_b\big(1 - p(x)\big)p'(m - \ell_I^a(x, Q)) - r.$$

Assumption 2 implies $D(x, Q) > 0$. The implicit-function theorem gives

$$\frac{\partial \ell_I^a}{\partial Q} = \frac{\pi_b\big(1 - p(x)\big)p(m - \ell_I^a)}{D(x, Q)} > 0,$$

and

$$\frac{\partial \ell_I^a}{\partial x} = -\frac{Q\pi_b p'(x)p(m - \ell_I^a)}{D(x, Q)} < 0.$$

This proves the proposition. □

Proof A.3 (of Proposition 2) *Fix x and suppose the entrant remains on b, so $\ell_I < \ell_I^a(x, Q)$. On this region,*

$$C_R(\ell_I, x, Q) = Q\pi_b\big(1 - p(x)\big)\big[p(m) - p(m - \ell_I)\big] - r\ell_I.$$

The incumbent's value is

$$V_I(\ell_I, x, Q) = \big(1 - Qp(m - \ell_I)\big)\Omega_I(x) - rx - C_R(\ell_I, x, Q).$$

Differentiating with respect to ℓ_I gives

$$\frac{\partial V_I}{\partial \ell_I} = r + Qp'(m - \ell_I)\left[\Omega_I(x) - \pi_b\big(1 - p(x)\big)\right].$$

If $\pi_b \le \pi_0$, then $\Omega_I(x) = \pi_0$ and

$$\Omega_I(x) - \pi_b\big(1 - p(x)\big) = \pi_0 - \pi_b\big(1 - p(x)\big) \ge \pi_0 - \pi_b \ge 0.$$

Therefore

$$\frac{\partial V_I}{\partial \ell_I} > 0$$

throughout the competing-b region. Any positive-retention policy with the entrant still on b can be improved by increasing ℓ_I until the deterrence threshold is reached. At the threshold, the entrant switches to project a, which weakly increases the incumbent's payoff further. Hence no positive-retention competing policy can be optimal. □

Proof A.4 (of Lemma 2) *By Proposition 2, no positive-retention equilibrium can leave the entrant on b. It remains to rule out no retention when $G(0; 0, Q) > 0$.*

Consider the path with $x = 0$ and ℓ_I increasing from 0 to $\ell_I^a(0, Q)$. Along this path, the entrant remains on b until the threshold is reached. From the derivative in Proof A.3, evaluated at $x = 0$,

$$\frac{\partial V_I}{\partial \ell_I} = r + Qp'(m - \ell_I)(\pi_0 - \pi_b).$$

Under $\pi_b \le \pi_0$, this derivative is strictly positive. Therefore the value at the deterrence threshold $\ell_I^a(0, Q)$ is strictly larger than the no-retention value $J^A(Q)$. Hence no retention is dominated whenever the replacement threat is active. Since positive retention with continued entry is also ruled out, every equilibrium deters the entrant. □

Proof A.5 (of Theorem 1) *By Lemma 2, every equilibrium deters the entrant. For a fixed deployment level x, deterrence requires*

$$\ell_I \ge \ell_I^a(x, Q),$$

and feasibility requires

$$\ell_I \geq x.$$

By the incumbent's tie-breaking rule, the incumbent chooses the smallest total retention level satisfying both inequalities:

$$\ell_I^D(x, Q) = \max\{x, \ell_I^a(x, Q)\}.$$

Thus every equilibrium is represented by the reduced deterrence problem

$$\max_{x \in [0,m]} \widetilde{V}^D(x, Q).$$

If $x^D(Q) = 0$, then $\ell_I^D(0, Q) = \ell_I^a(0, Q)$ and all retained talent is benched. This gives pure benching deterrence. In the no self-replacement case, the positive-premium part of the reduced deterrence value is concave in x, and its derivative at zero is

$$Q\pi_b p(m)p'(0) - r.$$

Hence (3.30) *is sufficient for $x^D(Q) = 0$ in the interior active-deterrence region.*

Next, suppose $0 < x^D(Q) < x^{pd}(Q)$. Since $x^D(Q) < x^{pd}(Q)$, deployment alone is not enough to deter the entrant. Thus $\ell_I^a(x^D(Q), Q) > x^D(Q)$, and the incumbent must retain

$$\ell_I^*(Q) = \ell_I^a(x^D(Q), Q)$$

with positive benching

$$z^*(Q) = \ell_I^a(x^D(Q), Q) - x^D(Q) > 0.$$

This is partial benching deterrence.

Finally, suppose $x^D(Q) \geq x^{pd}(Q)$. By definition of $x^{pd}(Q)$, choosing $\ell_I = x^D(Q)$ is sufficient to deter the entrant. Since $\ell_I = x$, no talent is benched:

$$z^*(Q) = 0.$$

This is pure deployment, or no-bench, deterrence. □

Proof A.6 (of Proposition 3) *In pure benching deterrence,*

$$x^*(Q) = 0, \qquad \ell_I^*(Q) = z^*(Q) = \ell_I^a(0, Q).$$

Proposition 1 implies

$$\frac{\partial \ell_I^a(0, Q)}{\partial Q} > 0$$

whenever the threshold is interior. Hence

$$\frac{d\ell_I^*(Q)}{dQ} > 0, \qquad \frac{dx^*(Q)}{dQ} = 0, \qquad \frac{dz^*(Q)}{dQ} > 0.$$

In partial benching deterrence,

$$\ell_I^*(Q) = \ell_I^a(x^*(Q), Q), \qquad z^*(Q) = \ell_I^a(x^*(Q), Q) - x^*(Q).$$

Differentiating the first identity gives

$$\frac{d\ell_I^*(Q)}{dQ} = \frac{\partial \ell_I^a(x^*(Q), Q)}{\partial Q} + \frac{\partial \ell_I^a(x^*(Q), Q)}{\partial x}\frac{dx^*(Q)}{dQ}.$$

Using the definitions of $A(Q)$ and $B(Q)$, this becomes

$$\frac{d\ell_I^*(Q)}{dQ} = A(Q) - B(Q)\frac{dx^*(Q)}{dQ}.$$

Differentiating $z^*(Q) = \ell_I^*(Q) - x^*(Q)$ gives

$$\frac{dz^*(Q)}{dQ} = A(Q) - \big(1 + B(Q)\big)\frac{dx^*(Q)}{dQ}.$$

Under no self-replacement, an interior deployment choice satisfies

$$Q\pi_b p(m)p'(x^*(Q)) = r.$$

Differentiating this first-order condition gives

$$\pi_b p(m)p'(x^*(Q)) + Q\pi_b p(m)p''(x^*(Q))\frac{dx^*(Q)}{dQ} = 0,$$

so

$$\frac{dx^*(Q)}{dQ} = -\frac{p'(x^*(Q))}{Qp''(x^*(Q))} > 0,$$

where the inequality follows from strict concavity of p. Conditions (3.35) *and* (3.36) *follow immediately from the formulas for $d\ell_I^*/dQ$ and dz^*/dQ, using*

$$\frac{A(Q)}{1 + B(Q)} < \frac{A(Q)}{B(Q)}.$$

In pure deployment deterrence,

$$\ell_I^*(Q) = x^*(Q), \qquad z^*(Q) = 0.$$

If the deployment first-order condition governs $x^*(Q)$, the argument above gives $dx^*(Q)/dQ > 0$. If the

zero-benching boundary governs $x^*(Q)$, then $x^*(Q) = x^{pd}(Q)$, where $H(x^{pd}(Q), Q) = 0$. Since

$$\frac{\partial H(x,Q)}{\partial Q} = \pi_b\big(1 - p(x)\big)p(m-x) > 0$$

and

$$\frac{\partial H(x,Q)}{\partial x} = r - Q\pi_b p'(x)p(m-x) - Q\pi_b\big(1-p(x)\big)p'(m-x) < 0$$

by Assumption 2, the implicit-function theorem gives

$$\frac{dx^{pd}(Q)}{dQ} = -\frac{\partial H/\partial Q}{\partial H/\partial x} > 0.$$

Thus, in the interior no-bench cases,

$$\frac{d\ell_I^*(Q)}{dQ} = \frac{dx^*(Q)}{dQ} > 0, \qquad \frac{dz^*(Q)}{dQ} = 0.$$

This proves the proposition. □

B Additional Robustness Checks for the Comparative Statics Test

B.1 Ex Ante Profitability Placebo Tests

We address the potential concern that the post-period employment increase reflects a coincident rise in expected profitability for software/IT rather than a hoarding response. If the *Alice* decision merely revalued software/IT, we should observe a positive treated-versus-control re-rating in (i) public-market industry portfolio returns around the decision and/or (ii) firm-level valuation ratios and profitability in the Shock 1 public-firm universe. This appendix implements these placebo tests around the June 2014 Supreme Court decision, using May 2014 as the last pre-decision month and June 19, 2014 as the high-frequency daily event date. For quarterly Compustat outcomes, 2014Q2 is the decision quarter; we omit this mixed quarter in the difference-in-differences specification and treat 2014Q3 as the first clean post-decision quarter. The tests find no evidence of a positive treated-versus-control re-rating at the cutoff.

B.1.1 Market-based placebo: industry portfolio excess returns (FF49)

We use Ken French's "49 Industry Portfolios" value-weighted returns. We map treated software/IT industries (NACE 58.2, 62.*, 63.1) to the `Softw` portfolio and, as a robustness mapping, to `Softw+BusSv`. We map control industries (NACE 20, 21, 26.5) to `Chems`, `Drugs`, and `LabEq`, and in robustness additionally include `MedEq`. For each treated-control pairing, we estimate:

$$\big(R^T - R^C\big)_t \;=\; \alpha \;+\beta\,\mathbb{1}\{t \geq \text{2014-06}\} \;+\varepsilon_t,$$

where R^T and R^C are monthly returns. Since the dependent variable is the treated-minus-control return differential, subtracting the risk-free rate from both sides leaves the differential unchanged. The window is 2012-04 to 2016-06, aligned with the main Shock 1 event window; the post indicator begins in 2014-06, the month of the Supreme Court decision. Standard errors are Newey–West HAC with maxlags=6.

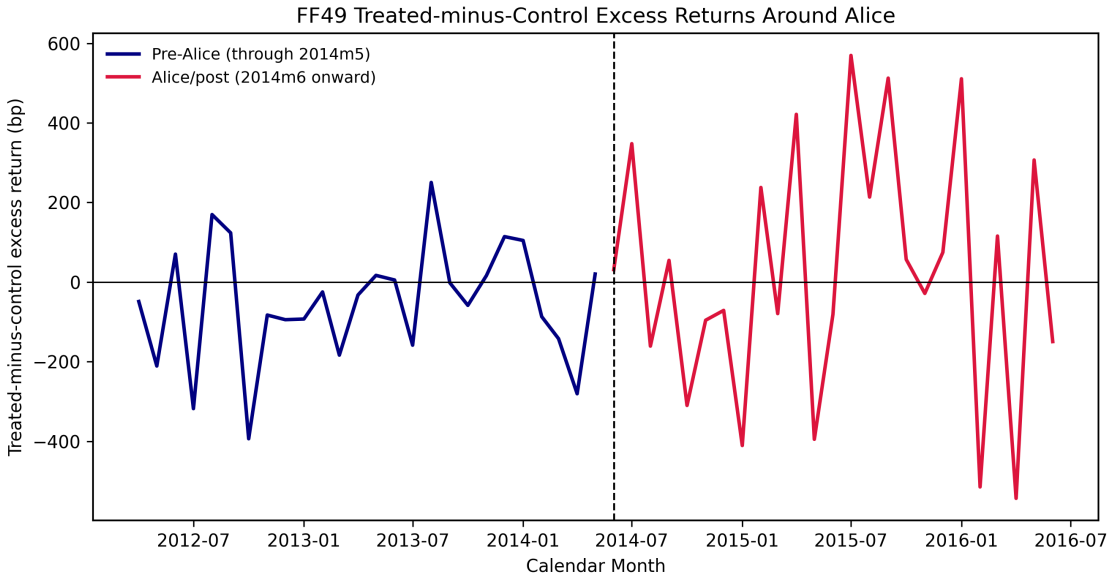
To probe pre-trends, we also estimate a linear slope in the pre-period for $(R^T - R^C)_t$. Table 12 reports the post coefficient and the pre-trend slope. Across all reasonable treated/control definitions, the post coefficients are small and statistically indistinguishable from zero, and pre-trend slopes are not significant. Figure 2 plots the treated-minus-control excess return series.

Table 12: Market-Based Placebo: Industry Excess Returns Around the *Alice* Decision (FF49 Portfolios)

Treated Portfolio	Control Portfolios	Post Coef. (β)	Pre-Trend Slope
`Softw`	`Chems+Drugs+LabEq`	0.008 (0.006)	0.0003 (0.0003)
`Softw`	`Chems+Drugs+LabEq+MedEq`	0.004 (0.006)	0.0003 (0.0003)
`Softw+BusSv`	`Chems+Drugs+LabEq`	0.005 (0.004)	-0.0000 (0.0002)
`Softw+BusSv`	`Chems+Drugs+LabEq+MedEq`	0.002 (0.004)	0.0000 (0.0002)

Notes: This table reports structural-break and baseline pre-trend estimates using Ken French's 49 Industry Portfolios. The dependent variable is the treated-minus-control monthly return differential. The post indicator equals one beginning in 2014-06, the month of the Supreme Court's *Alice* decision; 2014-05 is the final pre-period month. Newey–West HAC standard errors with maxlags=6 are reported in parentheses. Window: 2012-04 to 2016-06. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Figure 2: Treated-Minus-Control Excess Returns Around the *Alice* Shock



Notes: This figure plots monthly value-weighted returns of the treated software portfolio (`Softw`) minus the control portfolio average (`Chems`, `Drugs`, `LabEq`) from Ken French's 49 Industry Portfolios. The vertical dashed line indicates June 2014, the event month; May 2014 is the final pre-shock month.

We also examine a shorter high-frequency window around the Supreme Court decision date. Table 13

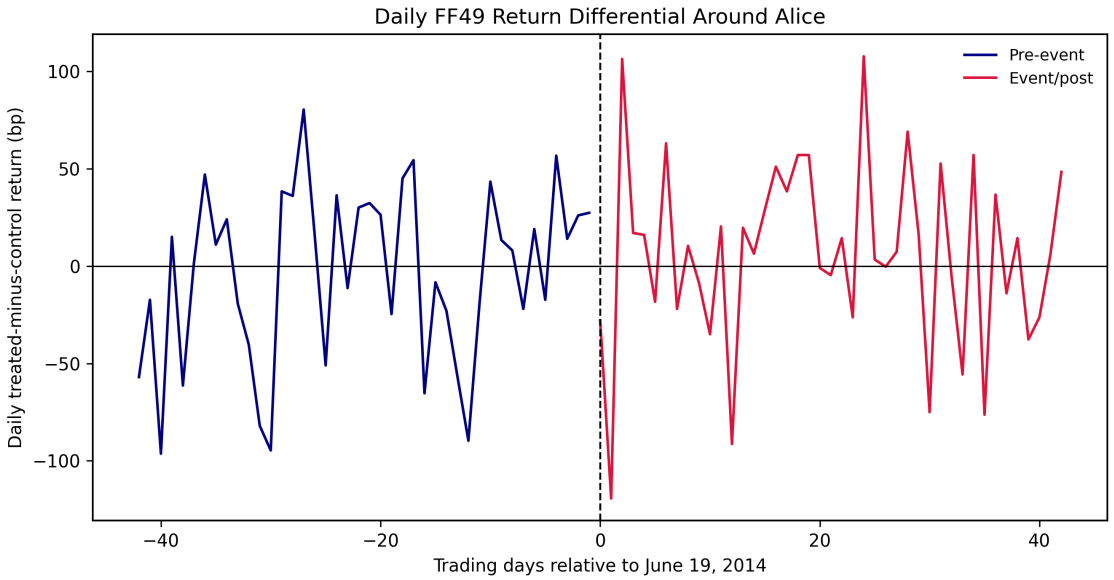
reports cumulative treated-minus-control return differentials in daily windows around June 19, 2014, and Figure 3 plots the raw daily return differential. The daily evidence does not show a statistically significant discrete re-rating at the decision date.

Table 13: High-Frequency Market-Based Placebo Around the *Alice* Decision (FF49 Daily Portfolios)

Treated Portfolio	Control Portfolios	Pre $[-20,-1]$	Post $[0,+20]$	Pre $[-42,-1]$	Post $[0,+42]$
`Softw`	`Chems+Drugs+LabEq`	0.06	1.68	-1.57	2.79
`Softw`	`Chems+Drugs+LabEq+MedEq`	0.17	1.85	-1.61	2.87
`Softw+BusSv`	`Chems+Drugs+LabEq`	-0.31	1.79	-1.29	2.46
`Softw+BusSv`	`Chems+Drugs+LabEq+MedEq`	-0.20	1.96	-1.33	2.54

Notes: Entries are cumulative treated-minus-control return differentials in percentage points using daily value-weighted FF49 industry portfolios. The event date is June 19, 2014 ($t = 0$). The primary mapping compares `Softw` to the average of `Chems`, `Drugs`, and `LabEq`. Since the statistic is a treated-minus-control return differential, subtracting the daily risk-free rate from both sides would cancel out.

Figure 3: Daily FF49 Return Differential Around the *Alice* Decision



Notes: This figure plots raw daily treated-minus-control returns for `Softw` minus the average of `Chems`, `Drugs`, and `LabEq`. The vertical dashed line indicates June 19, 2014, the date of the Supreme Court's *Alice* decision.

B.1.2 Firm-sample placebo: valuation ratios and profitability in the Shock 1 GVKEY universe

We could further implement a stricter placebo that uses firm-level quarterly accounting and market outcomes. We draw quarterly accounting and market variables from Compustat North America Quarterly Fundamentals (`comp.fundq`) via WRDS. We report two samples. The first is the full Shock 1 public-firm treated/control universe, classified using the ORBIS primary NACE in the audited mapping (treated: 58.2, 62.*, 63.1;

controls: 20, 21, 26.5), dropping ambiguous mappings that span both treated and control codes. The second is the exact 145-GVKEY firm sample used in the main Lightcast regressions, where we split firms by incumbent/pre-shock-patent status from `lightcast_incumbent.dta`, merged to GVKEY through the ORBIS–Lightcast matching table.

Variables. For each GVKEY-quarter (`gvkey`, `datadate`, `fyearq`, `fqtr`), we pull: (i) balance-sheet items: total assets (`atq`), total liabilities (`ltq`), and equity (`ceqq`, with `seqq` as a fallback book-equity measure); (ii) income-statement items: sales (`saleq`), cost of goods sold (`cogsq`), operating income before depreciation (`oibdpq`), and income before extraordinary items (`ibq`); and (iii) market inputs: quarter-end price (`prccq`) and shares outstanding (`cshoq`).

For each firm-quarter (i, t), we construct market value of equity as $\text{MVE}_{it} = \texttt{PRCCQ}_{it} \times \texttt{CSHOQ}_{it}$ and compute:

$$Q_{it} = \frac{\text{MVE}_{it} + \texttt{LTQ}_{it}}{\texttt{ATQ}_{it}}, \qquad \text{M/B}_{it} = \frac{\text{MVE}_{it}}{\text{BookEquity}_{it}}, \qquad ROA_{it} = 4 \times \frac{\texttt{IBQ}_{it}}{\texttt{ATQ}_{it}}, \tag{B.1}$$

where BookEquity_{it} is `CEQQ` when available and otherwise `SEQQ`. We winsorize each outcome at the 1st/99th percentiles and use $\ln(Q)$ and $\ln(\text{M/B})$ in the valuation specifications.

Timing and specification. We convert the Compustat period-end date (`datadate`) to a calendar-quarter index and align the placebo timing to the *Alice* decision by treating 2014Q2 as the event quarter. Because 2014Q2 contains both pre- and post-decision days, we omit it from the difference-in-differences specification and define $\text{Post}_t = 1$ beginning in 2014Q3. We estimate:

$$y_{it} \;=\; \alpha_i \;+ \delta_t \;+ \beta \left(G_i \times \text{Post}_t\right) \;+ \varepsilon_{it}, \tag{B.2}$$

with firm fixed effects α_i and calendar-quarter fixed effects δ_t; standard errors are clustered by firm (GVKEY). In the full Shock 1 universe, G_i identifies treated software/IT firms relative to the control industries. In the main-regression sample, G_i identifies incumbent/pre-shock-patent firms relative to non-incumbent firms. Table 14 reports estimates in the 2012Q2–2016Q2 window. Across valuation ratios and ROA, we find no evidence of a positive treated-versus-control or incumbent-versus-non-incumbent re-rating after *Alice*; if anything, the full-sample point estimates are negative.

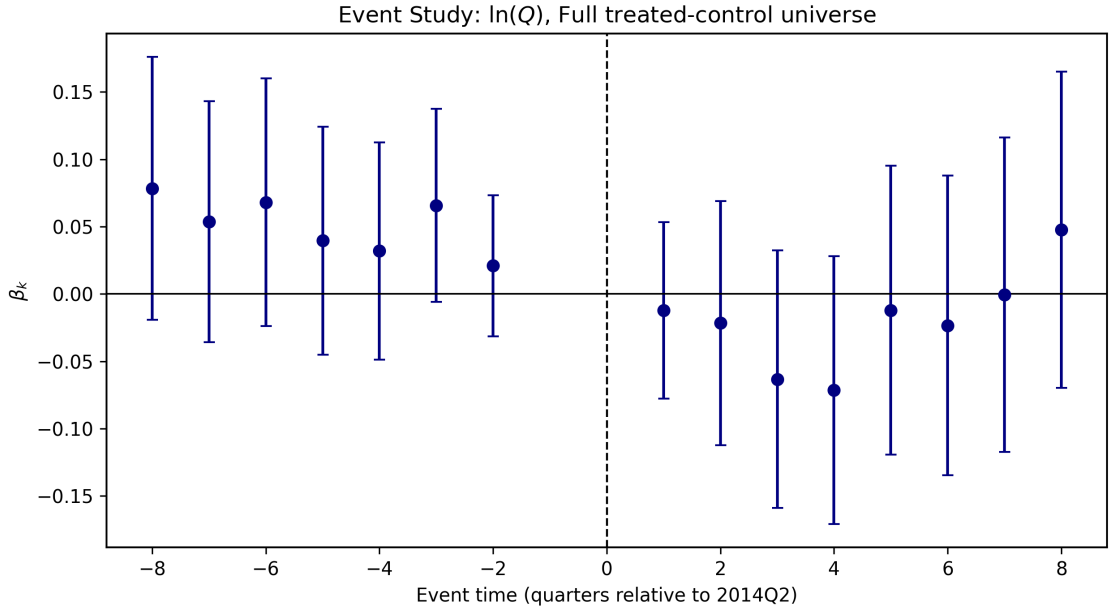
Finally, we estimate an event-study version over $k \in [-8, +8]$ quarters relative to 2014Q2, omitting $k = 0$ (decision quarter) and normalizing to $k = -1$ (2014Q1). Figure 4 and Figure 5 report the full Shock 1 public-firm universe; Figure 6 and Figure 7 report the exact main-regression firm sample. The event-study plots show no positive valuation break at the event date.

Table 14: Ex-Ante Profitability Placebo Using Compustat: June 2014 *Alice* Cutoff

Dependent Variable (y_{it})	Differential Post Coef. (β)	Observations	Unique Firms
Panel A. Full Shock 1 public-firm treated/control universe			
ln(Q)	-0.064 (0.043)	8,182	599
ln(M/B)	-0.111** (0.051)	7,111	552
ROA (Annualized)	-0.282*** (0.107)	8,475	620
Panel B. Main-regression firm sample; incumbent/pre-shock-patent versus non-incumbent			
ln(Q)	0.050 (0.135)	1,734	129
ln(M/B)	-0.057 (0.129)	1,546	120
ROA (Annualized)	0.132 (0.380)	1,827	136

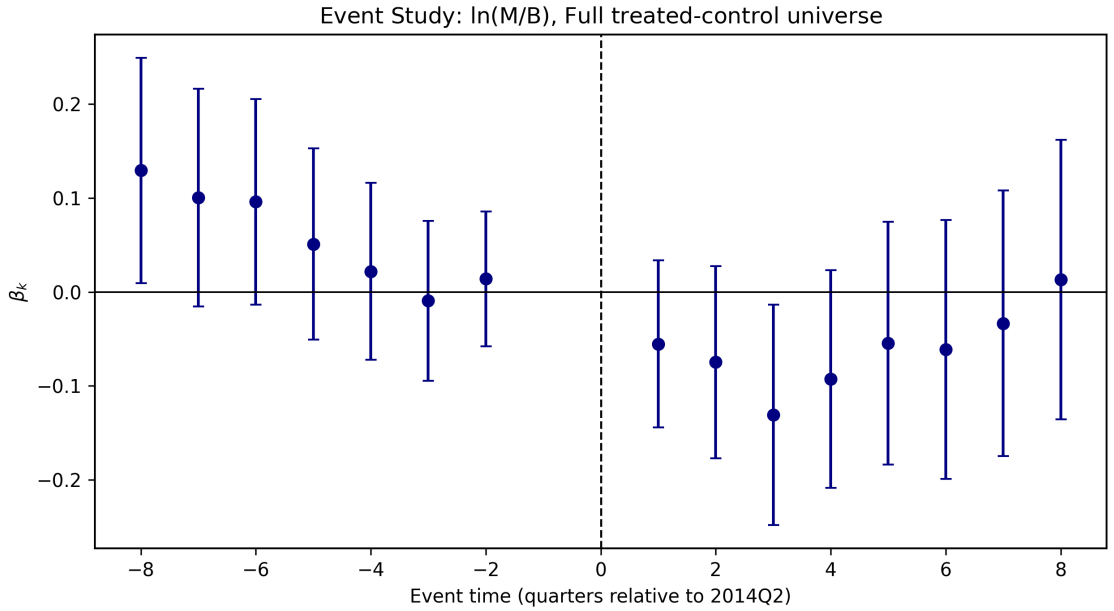
Notes: This table reports difference-in-differences estimates corresponding to Equation (B.2). In Panel A, G_i indicates treated software/IT firms (NACE 58.2, 62.*, 63.1) relative to control industries (NACE 20, 21, 26.5). In Panel B, the sample is restricted to the 145-GVKEY Lightcast/ORBIS main-regression firm sample and G_i indicates firms with incumbent/pre-shock-patent status from `lightcast_incumbent.dta`, merged to GVKEY through the ORBIS–Lightcast matching table. The event quarter is 2014Q2; 2014Q2 is omitted from the DiD and Post_t begins in 2014Q3. Window: 2012Q2–2016Q2. Outcomes are winsorized at the 1st/99th percentiles. Firm and calendar-quarter fixed effects are included; standard errors clustered by GVKEY are reported in parentheses. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

Figure 4: Event Study: ln(Q), Full Shock 1 Public-Firm Universe



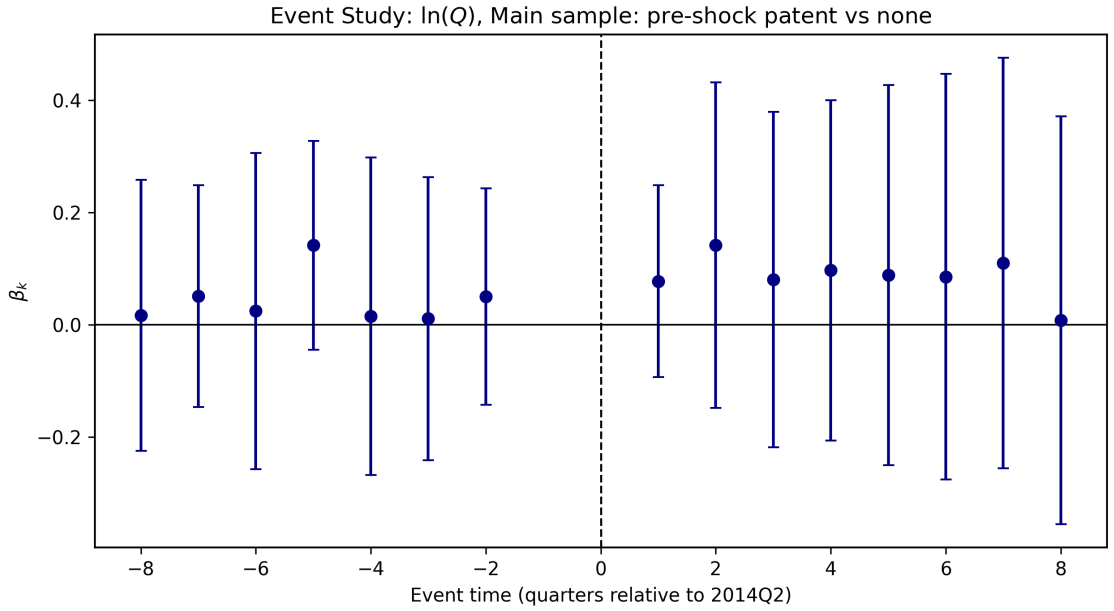
Notes: This figure plots quarterly event-study coefficient estimates for ln(Q) in the full Shock 1 public-firm universe. The omitted event quarter is $k = 0$ (2014Q2), and coefficients are normalized to $k = -1$ (2014Q1). Vertical bars report 95 percent confidence intervals using standard errors clustered by GVKEY.

Figure 5: Event Study: ln(M/B), Full Shock 1 Public-Firm Universe



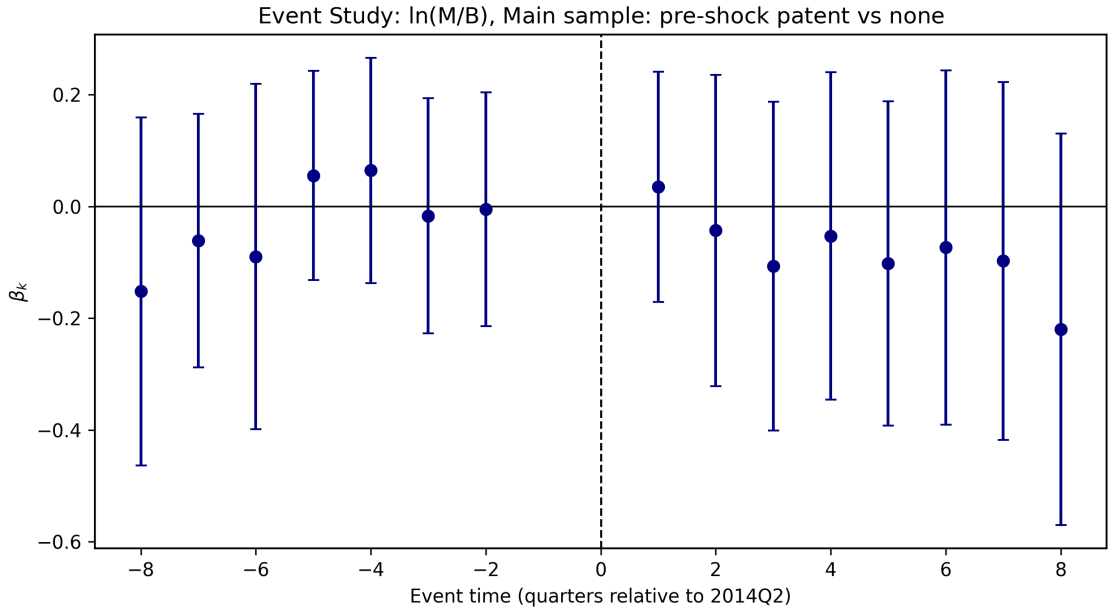
Notes: This figure plots quarterly event-study coefficient estimates for ln(M/B) in the full Shock 1 public-firm universe. The omitted event quarter is $k = 0$ (2014Q2), and coefficients are normalized to $k = -1$ (2014Q1). Vertical bars report 95 percent confidence intervals using standard errors clustered by GVKEY.

Figure 6: Event Study: ln(Q), Main-Regression Firm Sample



Notes: This figure plots quarterly event-study coefficient estimates for ln(Q) in the exact 145-GVKEY Lightcast/ORBIS main-regression firm sample. The omitted event quarter is $k = 0$ (2014Q2), and coefficients are normalized to $k = -1$ (2014Q1). Vertical bars report 95 percent confidence intervals using standard errors clustered by GVKEY.

Figure 7: Event Study: ln(M/B), Main-Regression Firm Sample



Notes: This figure plots quarterly event-study coefficient estimates for ln(M/B) in the exact 145-GVKEY Lightcast/ORBIS main-regression firm sample. The omitted event quarter is $k = 0$ (2014Q2), and coefficients are normalized to $k = -1$ (2014Q1). Vertical bars report 95 percent confidence intervals using standard errors clustered by GVKEY.

B.2 Robustness of the Parallel Trends Assumption: High-Frequency Joint Tests

To confirm that the difference-in-differences parameters reported in the main text are not driven by pre-existing divergent paths between exposed and unexposed firms, this subsection outlines a formal high-frequency test of the parallel trends assumption. Given that our underlying tracking data is recorded at the monthly level, generating traditional event-study coefficient plots across our entire estimation timeline would create a highly fractionated and statistically noisy visual matrix. Thus, we instead estimate a comprehensive monthly event-study specification and report formal joint Wald F-tests to verify that the vector of pre-treatment leads is statistically indistinguishable from zero.

For a symmetric window spanning June 2012 through June 2016 ($t \in [-24, +24]$ relative to the June 2014 *Alice* shock), we map monthly event time to a shifted relative calendar index $k = \text{ym}_t - \text{tm}(2014\text{m}6) + 25$. We estimate the following dynamic linear panel regression:

$$y_{it} = \alpha_i + \delta_t + \sum_{\substack{k=1 \\ k \neq 24}}^{49} \beta_k \left(\mathbb{1}\{\text{Relative Month} = k\} \times \text{Exposure}_i \right) + \Gamma \left(\text{Revenue}_{2012,i} \times \text{Year}_t \right) + \varepsilon_{it} \qquad \text{(B.3)}$$

where y_{it} alternates between the *Net Hiring Ratio*, *New Hiring Ratio*, and the *Separation Ratio* for firm i in month t. The model absorbs firm fixed effects (α_i) via our audited matching crosswalk and billing-month

fixed effects (δ_t). To control for time-varying scale shocks, we interact baseline firm size ($\text{Revenue}_{2012,i}$) with calendar-year indicators. Standard errors are conservatively clustered at the firm level (`matched_gvkey_final_str`).

We omit the interaction parameter immediately preceding the shock ($k = 24$, corresponding to relative month -1, or May 2014) to serve as our constant reference baseline.

Rather than examining individual, isolated point estimates across 23 distinct pre-treatment periods, we evaluate the structural validity of our parallel-trend restriction globally. We construct a joint Wald test assessing whether the entire vector of pre-shock parameters is simultaneously equal to zero:

$$H_0 : \beta_1 = \beta_2 = \cdots = \beta_{22} = \beta_{23} = 0 \tag{B.4}$$

Table 15 reports the resulting F-statistics and corresponding p-values extracted from this multi-degree-of-freedom test configuration.

Across all primary margins of worker adjustment, we fail to reject the null hypothesis of parallel pre-trends at any standard significance thresholds. For the net hiring ratio ($p = 0.7508$) and separation ratio ($p = 0.4376$), the joint pre-event trajectories are highly orthogonal, while the new hiring ratio ($p = 0.0994$) remains statistically marginal. This pre-trend justifies the consolidated difference-in-differences parameters presented in the main text.

Table 15: Joint F-Tests for Parallel Pre-Trends in Monthly Labor Flows

Dependent Variable (y_{it}):	Net Hiring Ratio (1)	New Hiring Ratio (2)	Separation Ratio (3)
Joint Pre-Trend Test ($H_0 : \beta_{\mathbf{pre}} = 0$)			
Wald F-Statistic	0.7798	1.4516	1.0267
Prob $> F$ (p-value)	0.7508	0.0994	0.4376
Reference Period Omitted	2014M05 ($k = 24$)	2014M05 ($k = 24$)	2014M05 ($k = 24$)
Pre-Shock Restrictions Tested	23 Months	23 Months	23 Months
Firm Fixed Effects	Y	Y	Y
Month Fixed Effects	Y	Y	Y
Revenue $\times$ Year Controls	Y	Y	Y
Observations	5,154	5,154	5,154
Unique Firms (Clusters)	131	131	131

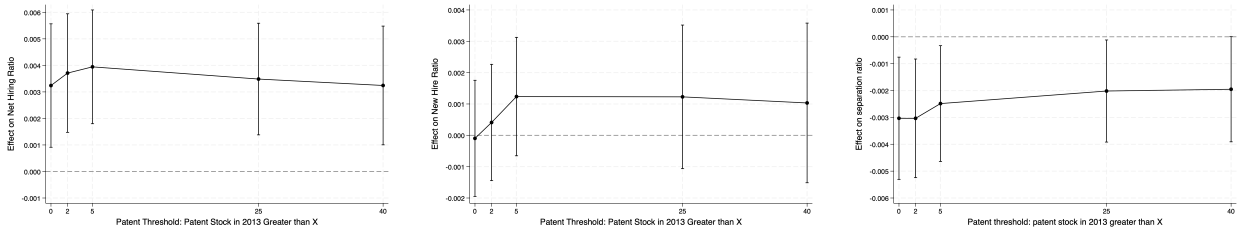
Notes: This table reports formal joint Wald F-tests verifying the parallel trends assumption prior to the June 2014 *Alice* shock, corresponding to the empirical framework detailed in Equation (B.3). The estimation sample spans a symmetric 24-month pre- and post-shock window from June 2012 through June 2016. The reported test diagnostics check the simultaneous linear equality to zero of all pre-treatment interaction terms spanning relative months -24 through -2 ($k = 1$ to $k = 23$). All specifications include firm fixed effects and calendar-month fixed effects. Standard errors are robustly clustered at the firm level.

B.3 Robustness to Alternative Patent-Stock Thresholds

This subsection examines whether the estimated post-*Alice* labor-market response is sensitive to the threshold used to define pre-shock patent exposure. In the main specification, patent exposure is defined using an indicator for whether a firm had positive patent stock prior to the shock. As a robustness exercise, we re-estimate the same difference-in-differences specification using alternative thresholds for 2013 patent stock. Specifically, we define exposed firms as those with 2013 patent stock greater than 0, 2, 5, 25, or 40 patents. This exercise keeps the estimating sample and regression specification fixed, and varies only the patent-stock threshold used to define exposure.

Figure 8 plots the corresponding point estimates and confidence intervals for net hiring, new hires, and separations. The results are robust to these alternative patent-stock thresholds.

Figure 8: Robustness of Labor Flow Responses to Alternative Patent-Stock Thresholds



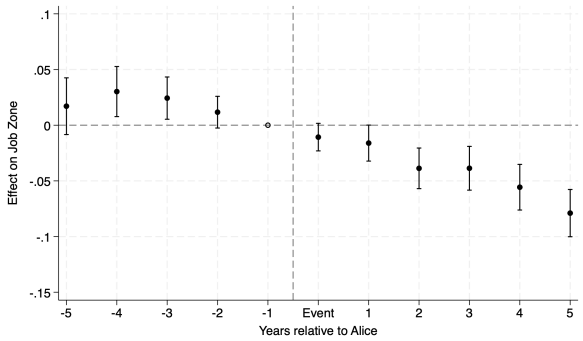
(a) Net Hire Ratio (b) New Hire Ratio (c) Separation Ratio

Notes: This figure plots difference-in-differences point estimates alongside their corresponding 95 percent confidence intervals across alternative 2013 pre-shock patent stock thresholds (greater than 0, 2, 5, 25, or 40 patents) used to define firm exposure. Panel (a) reports results for the net hiring ratio, Panel (b) for the new hire ratio, and Panel (c) for the separation ratio. Vertical bars denote confidence intervals calculated using standard errors clustered at the firm level.

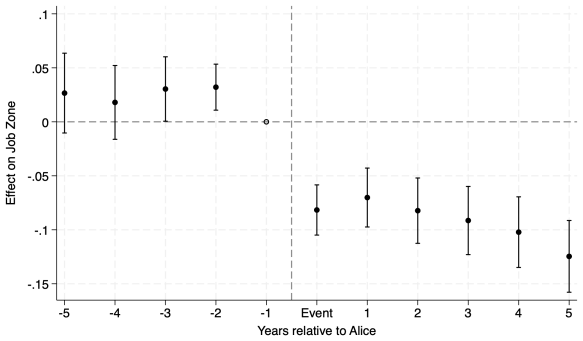
B.4 Occupation-Specific Event-Study Estimates

This appendix reports the full set of occupation-specific event-study estimates corresponding to the qualitative patterns summarized in Table 10. Figures are grouped according to the magnitude and interpretability of the estimated occupational progression effects following *Alice*. All specifications follow equation (5.3) in the main text and include worker and firm fixed effects. Coefficients are normalized relative to the year immediately preceding *Alice*. Error bars denote 95% confidence intervals.

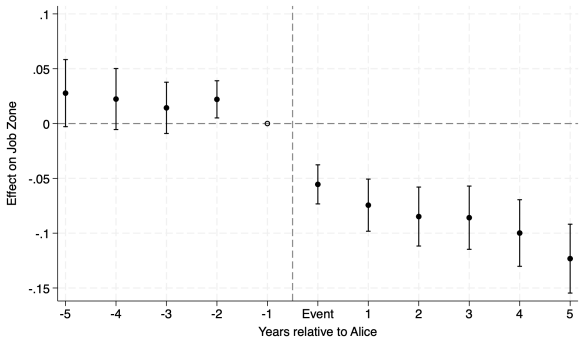
Figure 9: Occupations with Strong Negative Effects



(a) Science and Research



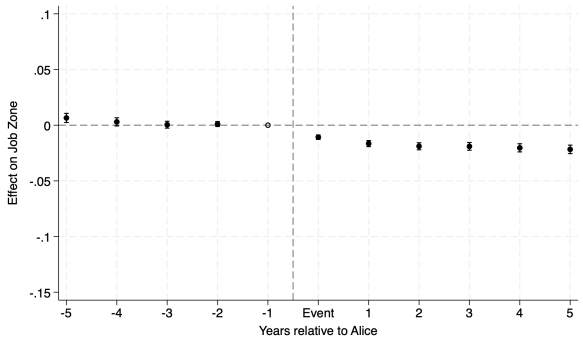
(b) Education and Training



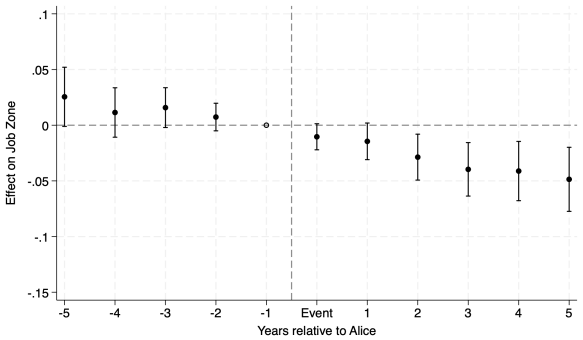
(c) Law, Compliance, & Public Safety

Notes: This figure plots event-study estimates across specific occupational sectors showing distinct structural declines. Panels (a), (b), and (c) display independent category models. Standard errors are clustered at the firm level.

Figure 10: Occupations with Moderate Negative Effects



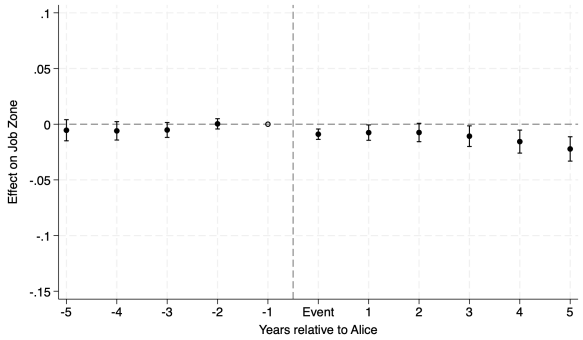
(a) Information Technology and Computer Science



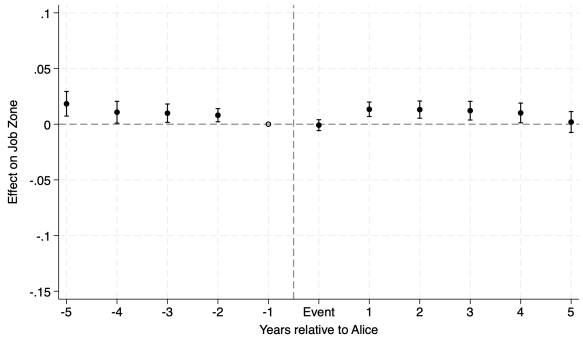
(b) Finance

Notes: This figure plots event-study estimates for occupational sectors showing moderate structural post-shock impacts. Panels (a) and (b) display independent category models. Standard errors are clustered at the firm level.

Figure 11: Occupations with Limited or Negligible Effects



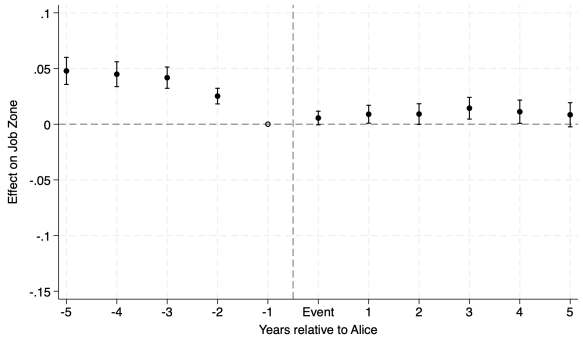
(a) Customer and Client Support



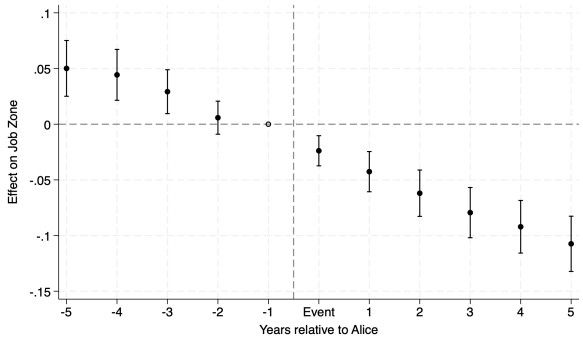
(b) Human Resources

Notes: This figure plots event-study estimates for occupational sectors showing no statistically or economically significant deviations post-intervention. Standard errors are clustered at the firm level.

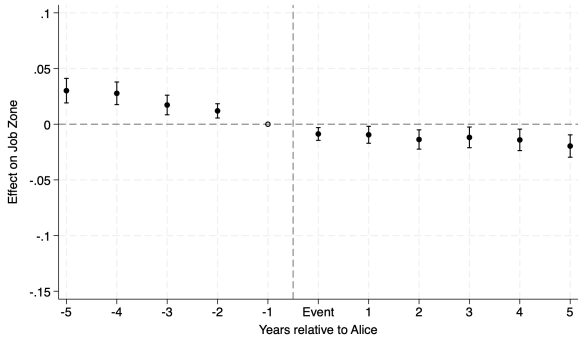
Figure 12: Occupations with Ambiguous or Difficult-to-Interpret Effects



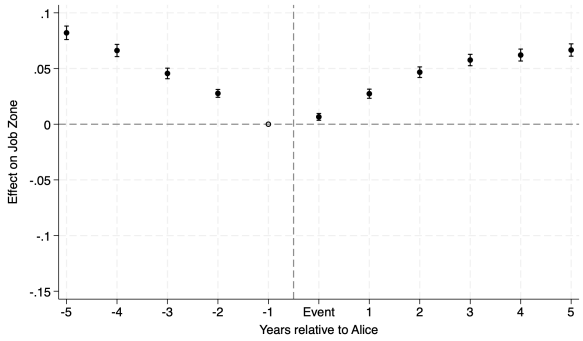
(a) Design, Media, and Writing



(b) Healthcare



(c) Finance



(d) Business Management and Operations

Notes: This figure plots event-study estimates for categories exhibiting highly volatile or statistically ambiguous trajectories. Panels (a) through (d) display independent category models. Standard errors are clustered at the firm level.

B.5 Occupation-Specific Difference-in-Differences Estimates

This appendix presents the remaining occupation-specific difference-in-differences estimates underlying the heterogeneity patterns summarized in Table 10. The estimating equation is the same as in the main text, with Job Zone as the outcome and with worker and firm fixed effects included in all specifications. The coefficient on $Post_t \times Treated_i$ measures the average post-*Alice* change in occupational rank for exposed workers relative to comparison workers in the same occupational category.

Tables 16-17 report estimates for the remaining occupational groups. Table 16 reports occupations with moderate negative responses; Table 17 reports occupations with limited or negligible responses. We report these estimates for completeness and to show that the main negative effects are not uniformly present across all occupations. Instead, the pattern is concentrated in occupations that are more plausibly connected to the strategic retention mechanism discussed in the main text.

Table 16: Occupations with Moderate Negative Effects

	Manufacturing & Production	IT & Computer Science
Post × Treated	-0.013*	-0.009***
	(0.008)	(0.001)
Worker Fixed Effects	Yes	Yes
Firm Fixed Effects	Yes	Yes
Observations	113,801	3,408,917
R^2	0.940	0.757

Notes: This table reports difference-in-differences estimates of the effect of exposure to *Alice* on occupational progression for occupations exhibiting moderate estimated declines.

Table 17: Occupations with Limited or Negligible Effects

	Customer / Client Support	Human Resources
Post × Treated	0.002	0.004
	(0.003)	(0.004)
Worker Fixed Effects	Yes	Yes
Firm Fixed Effects	Yes	Yes
Observations	224,679	535,872
R^2	0.942	0.499

Notes: This table reports difference-in-differences estimates of the effect of exposure to *Alice* on occupational progression for occupations exhibiting limited estimated effects.

References

Acikalin, U. U., Caskurlu, T., Hoberg, G., and Phillips, G. M. (2026). Intellectual property protection lost and competition: An examination using large language models. *Journal of financial economics*, 182.

Akcigit, U. and Goldschlag, N. (2023). Where Have All the "Creative Talents" Gone? Employment Dynamics of US Inventors. *NBER Working Paper Series*. Place: Cambridge.

Arrow, K. J. (1962). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pages 609–626. Princeton University Press.

Bar-Isaac, H., Johnson, J. P., and Nocke, V. (2025). Acquihiring for Monopsony Power. *Management science*, 71(4):3485–3496.

Benkert, J.-M., Letina, I., and Liu, S. (2024). Startup Acquisitions: Acquihires and Talent Hoarding. *arXiv.org*. Place: Ithaca.

Bessen, J. E., Pöge, F., and Röttger, R. (2023). Competing for Talent: Large Firms and Startup Growth. *SSRN Electronic Journal*.

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4):1104–1144. Place: NASHVILLE.

Cohen, L., Gurun, U. G., and Kominers, S. D. (2019). Patent Trolls: Evidence from Targeted Firms. *Management science*, 65(12):5461–5486. Place: Linthicum.

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

Ewens, M., Nanda, R., and Rhodes-Kropf, M. (2018). Cost of experimentation and the evolution of venture capital. *Journal of financial economics*, 128(3):422–442. Place: Amsterdam.

Feng, J. and Williams, P. (2026). Abstract patents and startup funding: Evidence from Alice v. CLS Bank. *Research policy*, 55(5):105462–.

Fernández-Villaverde, J., Yu, Y., and Zanetti, F. (2025). Defensive Hiring and Creative Destruction. Working Paper 33588, National Bureau of Economic Research. Series: Working Paper Series.

Friedmann, J., Glennon, B., and Hernandez, E. (2025). Substituting Talent with Transactions: Acquisitions as Responses to Immigration Restrictions. *NBER Working Paper Series*. Place: Cambridge.

Gilbert, R. J. and Newbery, D. M. G. (1982). Preemptive Patenting and the Persistence of Monopoly. *The American economic review*, 72(3):514–526. Place: Menasha, Wis.

Haegele, I. (2024). Talent Hoarding in Organization. *SSRN Electronic Journal*.

Hall, B. H. and Ziedonis, R. H. (2001). The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995. *The Rand journal of economics*, 32(1):101–128.

He, Q. and Waldman, M. (2024). Intra-firm Employer Learning, Talent Hoarding and Worker Visibility. *SSRN Electronic Journal*.

Kahn, C. and Huberman, G. (1988). Two-Sided Uncertainty and "Up-or-Out" Contracts. *Journal of Labor Economics*, 6(4):423–444. Place: Chicago, Ill.

Kalyani, A., Bloom, N., Carvalho, M., Hassan, T., Lerner, J., and Tahoun, A. (2025). The Diffusion of New Technologies. *The Quarterly journal of economics*, 140(2):1299–1365.

Lazear, E. P. and Shaw, K. L. (2007). Personnel Economics: The Economist's View of Human Resources. *Journal of Economic Perspectives*, 21(4):91–114. Place: Nashville.

Lehr, N. H. (2026). Does Monopsony Matter for Innovation? *IMF Working Paper*.

McCall, J. J. (1972). Economics of Information and Job Search: Reply. *Quarterly Journal of Economics*, 86(1):132–134.

Postel-Vinay, F. and Robin, J.-M. (2002). Equilibrium Wage Dispersion with Worker and Employer Heterogeneity. *Econometrica*, 70(6):2295–2350. Place: Malden, MA.

Waldman, M. and Zax, O. (2016). An Exploration of the Promotion Signaling Distortion. *Journal of Law, Economics, and Organization*, 32(1):119–149. Place: Oxford.

Whinston, M. D. (1990). Tying, Foreclosure, and Exclusion. *The American economic review*, 80(4):837–859.